

**EMC Enterprise Storage**

**EMC Navisphere  
Server Software for Solaris**

**ADMINISTRATOR'S GUIDE**

**P/N 069001068-01**

**EMC Corporation**

*Corporate Headquarters: (508) 435-1000, (800) 424-EMC2*

171 South Street, Hopkinton, MA 01748-9103

*Fax: (508) 435-5374 Service: (800) SVC-4EMC*

**Copyright © EMC Corporation 2000, 2001. All rights reserved.**

Printed February 2001

No part of this publication may be reproduced or distributed in any form or by any means, or stored in a database or retrieval system, without the prior written consent of EMC Corporation.

The information contained in this document is subject to change without notice. EMC Corporation assumes no responsibility for any errors that may appear.

All computer software programs, including but not limited to microcode, described in this document are furnished under a license, and may be used or copied only in accordance with the terms of such license. EMC either owns or has the right to license the computer software programs described in this document. EMC Corporation retains all rights, title and interest in the computer software programs.

EMC Corporation makes no warranties, expressed or implied, by operation of law or otherwise, relating to this document, the products or the computer software programs described herein. EMC CORPORATION DISCLAIMS ALL IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE. In no event shall EMC Corporation be liable for (a) incidental, indirect, special, or consequential damages or (b) any damages whatsoever resulting from the loss of use, data or profits, arising out of this document, even if advised of the possibility of such damages.

### **Trademark Information**

EMC<sup>2</sup>, EMC, MOSAIC:2000, Symmetrix, CLARiiON and Navisphere are registered trademarks and, EMC Enterprise Storage, The Enterprise Storage Company, The EMC Effect, Connectrix, EDM, SDMS, SRDF, Timefinder, PowerPath, InfoMover, FarPoint, EMC Enterprise Storage Network, EMC Enterprise Storage Specialist, EMC Storage Logic, Universal Data Tone, E-Infostructure, Celerra, and Access Logix are trademarks of EMC Corporation.

Solaris is a registered trademark of Sun Microsystems.  
Windows and Windows NT are registered trademarks of Microsoft Corporation.  
UNIX is a U.S. registered trademark of The Open Group.

All other trademarks used herein are the property of their respective owners.

<b>Preface</b> .....	xi
<b>Chapter 1 About EMC Navisphere Host Agent, CLI, and ATF</b>	
Terminology .....	1-2
Storage Components .....	1-3
Storage-System Configurations with ATF.....	1-6
Unshared Direct Configuration.....	1-6
Shared or Clustered Direct Configuration.....	1-7
Shared Switched Configuration .....	1-8
How ATF Handles Hardware Failures .....	1-9
<b>Chapter 2 Installing ATF, Host Agent and CLI</b>	
Hardware and Software Requirements .....	2-2
Power-Cycling the Storage System .....	2-2
Installing ATF Software .....	2-3
Pre-Install Checks .....	2-5
Installing ATF with the Boot Disk Outside a Storage System (First-Time Install) .....	2-6
Installing ATF with the Boot Disk in a Storage System (First-Time Install) .....	2-9
Installing a Newer Revision of ATF on a System That Already Uses ATF.....	2-12
Changing the Maximum Number of LUNs Supported by ATF.....	2-17
Installing the Host Agent and CLI Software.....	2-18
Modifying User Login Scripts.....	2-21
Configuring the Host Agent for non-FC4700 Storage Systems.....	2-21

	Adding a User to the Agent Configuration File.....	2-22
	Editing the Host Agent Configuration File .....	2-23
	Verifying Device Entries in the Host Agent Configuration File .....	2-23
	Using the Event Monitor Configuration File.....	2-24
	Configuring Storage Systems.....	2-24
	Persistent Bindings .....	2-25
	atf_configure Utility .....	2-26
	Adding and Modifying Persistent Bindings.....	2-28
	Editing Existing Persistent Bindings .....	2-30
	Verifying the HBA Driver and ATF Connections to SPs.....	2-32
	Verifying the Server Can See Its LUNs.....	2-33
	Starting and Stopping the Host Agent .....	2-34
	Logging Host Agent Status on the Host.....	2-35
	Removing (CDE) or ATF Software .....	2-36
	Removing the Host Agent or CLI.....	2-37
<b>Chapter 3</b>	<b>Making LUNs Available to Solaris</b>	
	Specifying Solaris Disk Names for LUNs .....	3-2
	Labeling and Partitioning LUNs .....	3-2
	Making File Systems on Partitions with newfs .....	3-3
	Mounting File Systems .....	3-4
<b>Chapter 4</b>	<b>Bootng From a Storage-System Device</b>	
	Using a Storage-System LUN as a Boot Device (SBUS) .....	4-2
	Installing Solaris on a Storage-System LUN (drive).....	4-2
	Using a Storage-System LUN as a Boot Device (PCI) .....	4-7
	Prerequisites .....	4-7
	Boot Disk Configuration.....	4-8
	Boot Setup .....	4-10
	Troubleshooting Suggestions.....	4-13
	Helpful ok Prompt Commands .....	4-13
<b>Chapter 5</b>	<b>Using ATF</b>	
	Testing ATF After Installation.....	5-2
	Understanding Device Names.....	5-3
	Adding Devices to a System That Has ATF Installed .....	5-4
	Failover and Other Messages.....	5-6
	ATF Automatic Failover .....	5-6
	Detecting Multiple Paths to LUN.....	5-6

	ATF Log Messages .....	5-7
	Restoring the Original Path .....	5-7
<b>Chapter 6</b>	<b>Modifying VERITAS FirstWatch Scripts</b>	
	Modifying Firstwatch Scripts in a Split-Bus Asymmetrical Configuration.....	6-2
	Modifying FirstWatch Scripts on the Primary Host .....	6-2
	Modifying the S05Clariion Startup Script .....	6-3
	Modifying the Surrender Script.....	6-4
	Modifying the FirstWatch Scripts on the Secondary Host..	6-5
	Modifying FirstWatch Scripts in a Split-Bus Symmetrical Configuration.....	6-8
	Polling Cycle Messages .....	6-8
	Modifying the S05Clariion Startup Script .....	6-9
<b>Chapter 7</b>	<b>Reassigning LUN Ownership</b>	
	Methods for Reassigning LUN Ownership.....	7-2
	Failover Using Application Transparent Failover (ATF)....	7-3
	Failover Using trespass_array and rescan_array .....	7-3
	Failover Using FirstWatch or VERITAS Cluster Server (VCS) Software.....	7-5
	Updating Disk Names After Reassigning LUNs.....	7-6
	Updating Disk Names if an SP Fails .....	7-6
	Updating Disk Names if You Manually Change LUN Ownership.....	7-6
	Examples — Manually Reassigning LUNs .....	7-7
	Unshared Direct Single Server Configuration.....	7-7
	Reassigning LUN Ownership Using Navisphere Manager .....	7-9
	Failing Over a LUN Using Trespass_array .....	7-10
	Sample Unshared Direct Dual-Server Configuration.....	7-11
	Reassigning LUN Ownership Using Navisphere Manager .....	7-13
	Failing Over a LUN Using trespass_array .....	7-15
<b>Appendix A</b>	<b>Troubleshooting</b>	
	Only LUN 0 Is Available to the Sun Host, or the Sun Host Cannot See New LUNs .....	A-2
	Cannot Find LUNs .....	A-2

**Appendix B    Configuring the Host Agent**

Configuring the Host Agent Using a Text Editor ..... B-2  
     Host Agent Configuration File Pathname ..... B-2  
     Sample Agent Configuration File (agent.config) ..... B-3  
     Specifying the Agent Description and Contact ..... B-4  
     Specifying the Communication Channel ..... B-5  
     To Specify a Communication Channel ..... B-5  
     Communicating When a Storage System Has No Bound  
     LUNs ..... B-9  
     Specifying Who Can Configure a Managed  
     Storage System ..... B-11  
     Specifying a Polling Frequency for the Agent ..... B-12  
     Specifying the Number of SP Events to Read ..... B-12

**Glossary** .....g-1

**Index** ..... 1-1

1-1	Sample Navisphere Shared Storage Configuration .....	1-4
1-2	Sample Navisphere Unshared Storage Configuration .....	1-5
1-3	Sample Unshared Direct Configuration .....	1-6
1-4	Sample Shared Direct Configuration (FC4700 only) .....	1-7
1-5	Sample Clustered Direct Configuration (non-FC4700) .....	1-7
1-6	Sample Shared Switched Configuration .....	1-8
7-1	Unshared Direct Single-Server Configuration Before Assigning LUN 0 .....	7-8
7-2	Unshared Direct Single-Server Configuration After Assigning LUN 0 .....	7-9
7-3	Unshared Direct Dual-Server Configuration Before Assigning LUN 0 .....	7-12
7-4	Unshared Direct Dual-Server Configuration After Assigning LUN 0 .....	7-13



1-1	How ATF Handles Hardware Failures .....	1-9
7-1	Unshared Direct Single-Server Configuration Before Assigning LUN 0 .....	7-7
7-2	Unshared Direct Single-Server Configuration After Assigning LUN 0 .....	7-8
7-3	Unshared Direct Dual-Server Configuration Before Assigning LUN 0 .....	7-12
7-4	Unshared Direct Dual-Server Configuration After Assigning LUN 0 .....	7-12



---

This manual describes how to set up a Sun server to use Fibre Channel disk-array storage systems with Fibre Channel interfaces to the host. These storage systems may or may not have storage processors (SPs). Storage systems without SPs consist of Disk-Array Enclosures (DAE) in a JBOD (just a bunch of disks) configuration.

This manual assumes that you are familiar with the server hardware and the Solaris<sup>®</sup> operating system, especially with how the operating system handles physical disks, disk partitioning, and file systems.

---

See the interface kit's release notice for information on supported servers.

Here is an overview of where information is located in this manual.

Chapter 1, *About EMC Navisphere Host Agent, CLI, and ATF*, provides an overview of the Navisphere<sup>®</sup> Host Agent, CLI, and ATF.

Chapter 2, *Installing ATF, Host Agent and CLI*, provides the steps for installing and configuring the Navisphere Host Agent, CLI, and ATF.

Chapter 3, *Making LUNs Available to Solaris*, describes how to make any LUNs available to the Solaris operating system.

Chapter 4, *Booting From a Storage-System Device*, describes how to boot the operating system from a storage-system device.

Chapter 5, *Using ATF*, describes how to use, test and remove ATF.

Chapter 6, *Modifying VERITAS FirstWatch Scripts*, describes how to modify the Veritas FirstWatch scripts so that FirstWatch<sup>®</sup> works properly in the split-bus asymmetrical configuration and the split-bus symmetrical configuration.

Chapter 7, *Reassigning LUN Ownership*, describes various methods for reassigning LUN ownership from one SP to the other.

Appendix A, *Troubleshooting*, describes what to do when LUNs are not available.

Appendix B, *Configuring the Host Agent*, describes how to edit the Agent configuration file without using Manager.

### Related Documentation

- *PCI Host Bus Adapter and Driver for Solaris Installation Guide*, P/N 014003032
- *SBus Host Bus Adapter and Driver for Solaris Installation Guide*, P/N 014003031
- *EMC Navisphere Command Line Interface (CLI) Version 5.X Reference*, P/N 069001038
- *EMC Navisphere Manager 5.X User Guide*, P/N 069000882
- *CDE Driver Extensions Software Administrator's Guide*, P/N 069001072
- *FC-Series and C-Series Storage System and Navisphere Event Codes Version 5.X Reference*, P/N 069001061

### Conventions Used in This Manual

EMC uses the following conventions for notes, cautions, warnings, and danger notices.

---

A note presents information that is important, but not hazard-related.

---



#### **CAUTION**

---

**A caution contains information essential to avoid damage to the system or equipment. The caution may apply to hardware or software.**

---



#### **WARNING**

---

***A warning contains information essential to avoid a hazard that can cause severe personal injury, death, or substantial property damage if you ignore the warning.***

---

**DANGER**

*A danger notice contains information essential to avoid a hazard that will cause severe personal injury, death, or substantial property damage if you ignore the warning.*

EMC uses the following type style conventions in this guide:

<b>Boldface</b>	<ul style="list-style-type: none"> <li>• Specific filenames or complete paths.</li> <li>• Dialog box names and menu items in text.</li> <li>• Selections you can make from the user interface, including buttons, icons, options, and field names.</li> <li>• Emphasis in cautions and warnings.</li> </ul>
<i>Italic</i>	<ul style="list-style-type: none"> <li>• New terms or unique word usage in text.</li> <li>• Command line arguments when used in text.</li> </ul>
Fixed space	Examples of specific command entries that you would type, displayed text, or program listings. For example:  <pre>QUERY [CUU=cuu VOLSER=volser]</pre>
<i>Fixed italic</i>	Arguments used in examples of command line syntax.

**Where to Get Help**

Obtain technical support by calling your local sales office.

If you are located outside the USA, call the nearest EMC office for technical assistance.

For service, call:

**United States:** (800) 782-4362 (SVC-4EMC)

**Canada:** (800) 543-4782 (543-4SVC)

**Worldwide:** (508) 497-7901

and ask for Customer Service.



## About EMC Navisphere Host Agent, CLI, and ATF

This chapter describes the EMC Navisphere® storage-system management configurations and architecture. Major topics are

- Terminology ..... 1-2
- Storage Components ..... 1-3
- Storage-System Configurations with ATF ..... 1-6

## Terminology

Term	Meaning
Host Agent	EMC Navisphere Agent that runs on a storage-system server
SP Agent	EMC Navisphere Agent that runs on the SPs in an FC4700 storage system
managed Agent	Host Agent or SP Agent managed by EMC Navisphere management software
managed storage system	Storage system managed by EMC Navisphere management software
RAID Group storage system	Storage system whose storage processors (SPs) are running Core or Base Software that has RAID Group functionality
non-RAID Group storage system	Storage system whose SPs are running Core or Base Software that does not have RAID Group functionality
shared storage system	Storage system with the EMC Access Logix™ option, which provides data access control (Storage Groups) and configuration access control. A shared storage system is always a RAID Group storage system
unshared storage system	Storage system without the EMC Access Logix option
FC-series storage system	FC4300/4500, FC4700, FC5600/5700, FC5400/5500, FC5200/5300, or FC5000 series storage system
C-series storage system	C1000, C1900, C2x00, or C3000 series storage system
JBOD storage system	Storage system <i>without</i> storage processors (SPs); that is, one that contains only DAEs and <i>no</i> DPEs
ATF	EMC Navisphere Application Transparent Failover
CDE	CLARiiON Driver Extensions, which provide basic failover features when ATF is not present
CLI	EMC Navisphere Command Line Interface
Event Monitor	EMC Navisphere Event Monitor
Manager	EMC Navisphere Manager

## Storage Components

The basic components of a storage-system configuration are

- One or more storage systems
- One or more servers connected to the storage systems directly or through hubs or switches. A server can run Microsoft® Windows NT®, Microsoft Windows® 2000, or one of several UNIX® operating systems, such as the Sun Solaris® or IBM AIX® operating system.
- A host computer (called a *management station*) that is running Navisphere storage-management software and is connected to any storage-system servers and any storage processors (SPs) in an FC4700 storage system over a local area network (LAN). A management station can also be a server if it is connected to a storage system. The management station must run the Microsoft Windows NT or the Windows 2000 operating system

You must install the Navisphere Host Agent on all servers. The SP Agent is installed at the factory on the SPs in FC4700 storage systems. The Host Agents and SP Agents communicate with the management station's storage-management software and with the Core or Base Software in the storage system.

Navisphere Manager, the storage-management software, provides a graphical user interface that lets you configure the storage systems connected to the servers on the LAN.

The CLI is a command line interface to the Host Agent and SP Agent that can run on a storage-system server or a management station. It provides an alternative to Manager. With the CLI, you can configure, control, and retrieve status from any managed storage system. You can also use the CLI to automate disk-storage management functions by writing shell scripts or batch files. Installing the CLI is optional; the Host Agent and SP Agent can function normally without the CLI.

ATF (Application Transparent Failover) is an optional software package designed for high-availability installations. ATF is required for a host that has two HBAs connected to the same storage system, and is optional for other hosts. ATF works with storage systems to let applications continue running after the failure of an SP, host-bus adapter (HBA), switch, or storage-system cable. Without human intervention, ATF can route I/O through a secondary path to the LUNs the applications need.

Figure 1-1 shows a sample Navisphere shared-storage configuration. The Windows NT host is a management station for shared storage systems connected to their servers. Note that the management station is connected by a LAN to the storage-system servers and to the SPs in the FC4700 storage system. The management station shown is not a storage-system server, but it could be.

The CLI is running on multiple hosts, but only one instance is needed to manage the storage systems connected to any server.

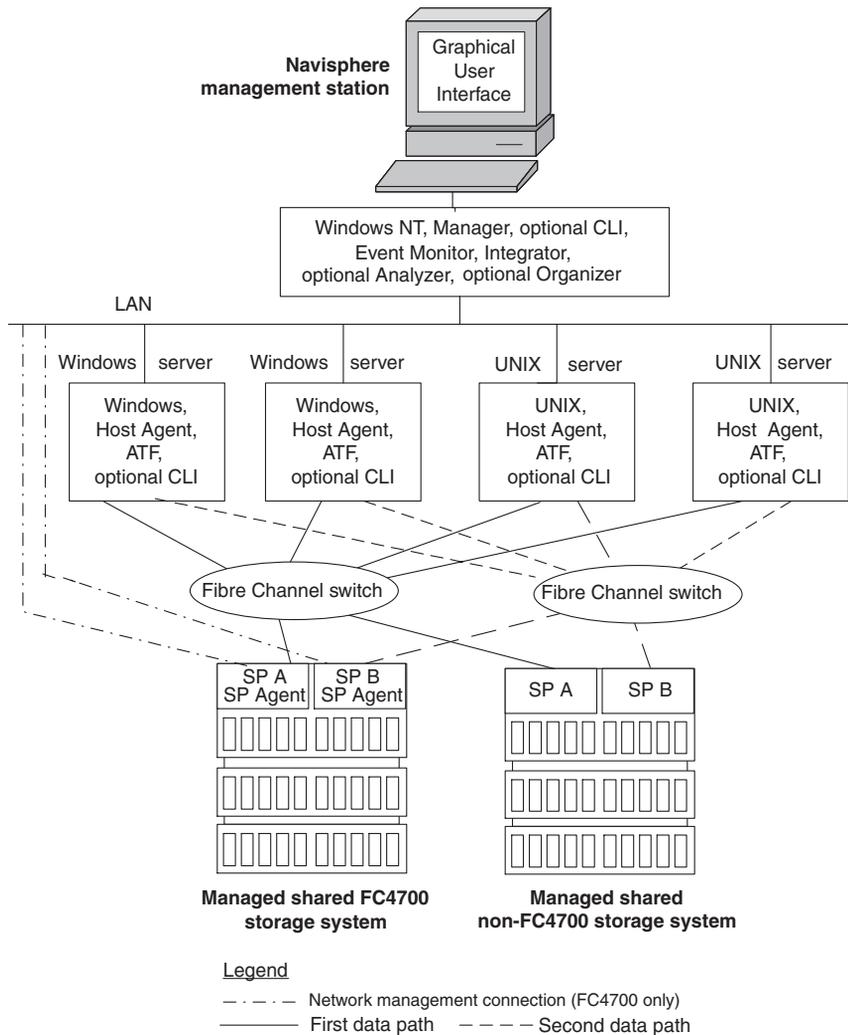


Figure 1-1 Sample Navisphere Shared Storage Configuration

Figure 1-2 shows a sample Navisphere unshared storage configuration in which a Windows NT host is a management station for unshared storage systems connected directly to their servers. The management station shown is not a storage-system server, but it could be. The CLI is running on multiple hosts, but only one instance is needed to the storage systems connected to any server.

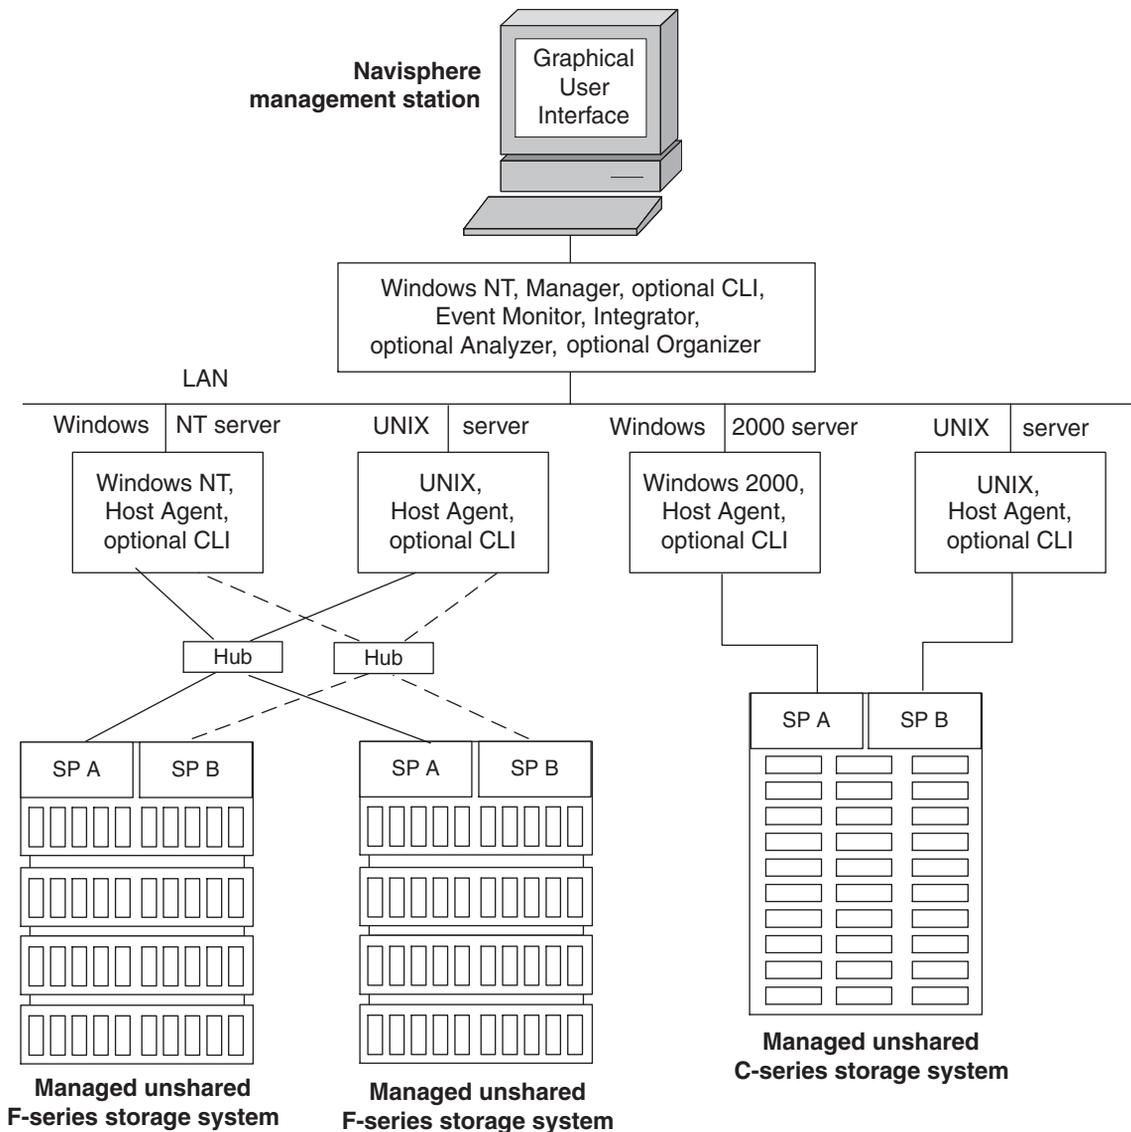


Figure 1-2 Sample Navisphere Unshared Storage Configuration

## Storage-System Configurations with ATF

The storage-system configurations that benefit from ATF are

- Unshared direct
- Shared-or-clustered direct
- Shared switched

Each configuration requires two HBAs in the server connected to each storage system and ATF running on the server. For best SP performance, you can bind some LUNs on one SP and the other LUNs on the other SP. The SP that binds a LUN is the default owner of that LUN and determines the primary I/O path to that LUN. The route through the other SP to the LUN is the secondary path, and is available if a component in the primary route fails.

### Unshared Direct Configuration

The unshared direct configuration (Figure 1-3) has one host connected to a storage system without the EMC Access Logix option.

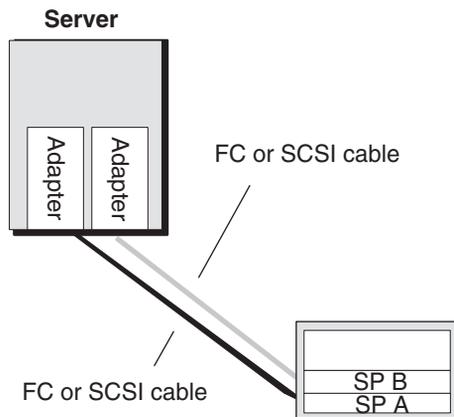


Figure 1-3 Sample Unshared Direct Configuration

## Shared or Clustered Direct Configuration

The shared direct configuration (Figure 1-4) uses the Access Logix option in the storage systems to control LUN access. It is supported for FC4700 storage systems only. A clustered direct configuration (Figure 1-5) uses clustering software. It is supported for non-FC storage systems without the Access Logix option and by Solaris servers when the servers are running clustering software. This configuration provides the highest availability by protecting against host failure.

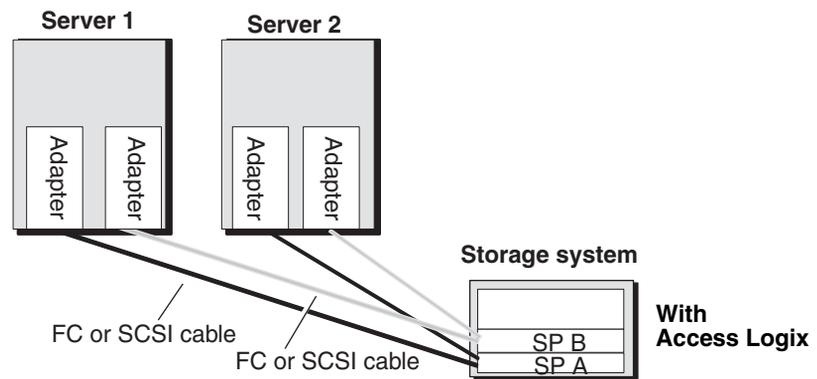


Figure 1-4 Sample Shared Direct Configuration (FC4700 only)

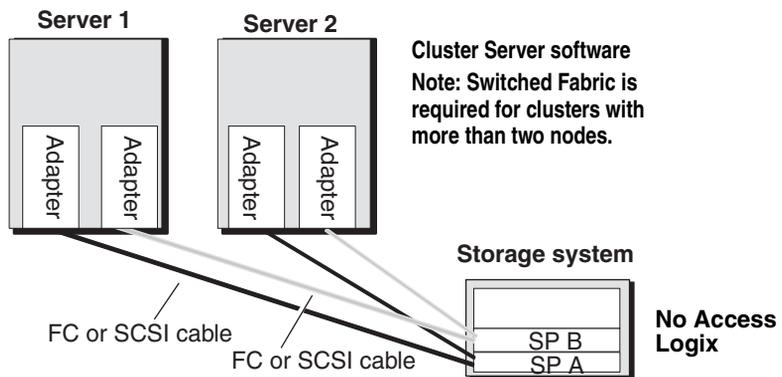


Figure 1-5 Sample Clustered Direct Configuration (non-FC4700)

## Shared Switched Configuration

Switches add flexibility and availability to a Fibre Channel site. Essentially, a switch lets you expand any of the previous configurations to include multiple servers and storage systems. A shared switch configuration uses the Access Logix option in the storage systems to control LUN access and one or two switch fabrics, depending on the switch hardware model. This configuration provides the highest availability when the servers run cluster software.

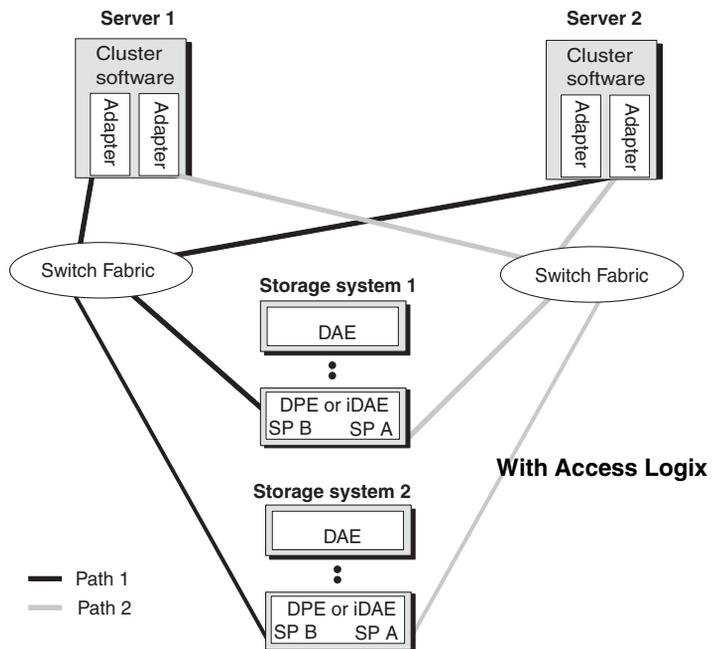


Figure 1-6 Sample Shared Switched Configuration

## How ATF Handles Hardware Failures

Table 1-1 describes ATF features for recovery from hardware failure. Note that for hardware failures of a disk, fan module, or power supply, redundant hardware will allow operation to continue regardless of ATF software.

**Table 1-1 How ATF Handles Hardware Failures**

Hardware	Failing Component	What Happens with ATF
<u>Unshared Direct Configuration</u> One host, Two HBAs, Two SPs	HBA or HBA cable	ATF transfers the LUNs in the failed path to the working path.
	SP or SP cable	
	LCC or SP-to-LCC cable	
<u>Shared Direct-or-Clustered Configuration</u> Two hosts each with two HBAs, Two SPs	Host	With FirstWatch, VCS or Sun Clusters, ATF transfers the LUNs in the failed path to the working path.
	HBA or HBA cable	ATF transfers the LUNs in the failed path to the working path.
	SP or SP cable	
	LCC or SP-to-LCC cable	
<u>Shared Switched Configuration</u> Multiple hosts each with two HBAs, Two switches or hubs, Two SPs	Host	With FirstWatch, VCS or Sun Clusters, ATF transfers the LUNs in the failed path to the working path.
	HBA or HBA cable	ATF transfers the LUNs in the failed path to the working path.
	FC switch or hub	
	SP or SP cable	
	LCC or SP-to-LCC cable	



This chapter describes the following topics:

- Hardware and Software Requirements .....2-2
- Power-Cycling the Storage System .....2-2
- Installing ATF Software.....2-3
- Changing the Maximum Number of LUNs Supported by ATF2-17
- Installing the Host Agent and CLI Software.....2-18
- Modifying User Login Scripts .....2-21
- Configuring the Host Agent for non-FC4700 Storage Systems.2-21
- Using the Event Monitor Configuration File .....2-24
- Configuring Storage Systems .....2-24
- Persistent Bindings .....2-25
- Verifying the HBA Driver and ATF Connections to SPs .....2-32
- Verifying the Server Can See Its LUNs .....2-33
- Starting and Stopping the Host Agent.....2-34
- Logging Host Agent Status on the Host.....2-35
- Removing (CDE) or ATF Software .....2-36
- Removing the Host Agent or CLI.....2-37

## Hardware and Software Requirements

To run the Host Agent, CLI, and ATF, the storage-system server must meet the following requirements:

- Be a Sun SPARCserver™, SPARCstation®, or Enterprise® host running a supported revision of a Solaris® operating system. For the latest requirements see the software product Release Notes.

---

ATF interoperates with the Sun Solaris **rdnexus** and **rdriver** drivers used with Sun Storage Devices such as the A1000.

---

- Have the HBA hardware and driver installed.
- Have the server's **sd.conf** file edited to include all existing LUNs in connected storage systems as explained in either the SBus or PCI HBA and Driver Installation Guide shipped with the HBA. For additional information on the SUN operating system, refer to the manpage on **driver.conf**.
- Be connected to both SPs in each storage system, either directly or through a switch or hub. This assumes that the switch is configured (zoned) properly.

## Power-Cycling the Storage System

UNIX does not generally support the powering down of a storage device while the system is running. There are, however, occasions when we recommend power-cycling a storage system as part of service. When power cycling a storage system, it is important to observe the following guidelines:

- If the system is booted from the storage system, the system must be shut down before the storage system is powered down.
- If the system does not boot off the storage system, it is still important to unmount any file systems on the storage system before power-cycling the storage system.
- It is recommended that all I/O, including raw I/O, be stopped before power-cycling the storage system. However, the system will not panic if raw I/O occurs when the storage system is powered down. If there is any attempt at I/O while the storage system is off or while it is going through its boot sequence, the Navisphere ATF driver will attempt to failover. This means that

the **atf\_restore** command must be executed before I/O to the storage system can continue. The correct **atf\_restore** command line is included in the Navisphere ATF system console messages.

## Installing ATF Software



### **CAUTION**

**If the server has two HBAs connected to the same storage system, you must install ATF.**

For all installation procedures, when we use the command `reboot -- -r`, we assume that the bootable disk on which you installed ATF is the default boot device. If the disk on which you installed ATF is not the default boot device, you must type a `reboot -- -r` command that specifies the full boot device name. For example, `reboot -- devicename -r`.

The following sections describe the installation prerequisites, the pre-install checks that are in place for Navisphere ATF, and three methods for installing ATF listed below:

- For the first time with the boot disk outside a storage system.
- For the first time with the boot disk in a storage system.
- Installing a newer revision of ATF on a system that already uses ATF.



### **CAUTION**

**Before installing ATF software, verify or do the following:**

- On non-fibre channel storage systems, check that the storage system is set to return **BUSY** instead of **QUEUE-FULL**.
- Do not mix `sd` and `atf` devices on the same external bus. Doing this will cause problems.
- Verify that both storage-system SPs are properly connected. ATF will not install and configure properly if one of the SPs is not properly connected.

- If a version of the `clsp###.pkg` driver package is already installed, you must remove it before installing ATF. *Do not* reinstall any version of the `clsp` driver. CLSP and ATF must not be installed on the same server. ATF provides the CLSP functionality.
- If the boot disk is a LUN within a Storage Group, make sure the sharing state of the Storage Group is set to Dedicated (the Storage Group is able to connect to only one server). Refer to the Manager manual for setting Storage Group properties.
- A boot device cannot be encapsulated by (be under the control of) Solstice Disk Suite or VERITAS.
- If you have VERITAS software installed and the version is earlier than 3.0.4, disable VXDMP (Refer to the VERITAS manual for instructions).
- Do not proceed with the ATF installation unless you know that all users have logged off and the system is ready for shutdown.
- If you are installing ATF for the first time, and if any of the Solstice Disk Suite database replicas or metadevices contain CLARiiON devices, then the devices must be removed from the database replicas and the metadevices must be decomposed. Since these devices are currently under `sd` control, the device will no longer be accessible through `sd` after ATF is installed. This makes the metadvice unusable with Solstice DiskSuite. You do not need to decompose any metadevices that do not contain CLARiiON devices.
- As part of installation, ATF replaces the program `/usr/sbin/disks`. If another install package does the same thing and they are removed in a different order than they are installed, problems may occur.

## Pre-Install Checks

The Navisphere ATF installation checks that the host system is running a supported release of the operating system and prerequisite software conditions are met. Warnings are issued if these checks fail and the installation will terminate with the following messages:

```
Warnings cause the installation to fail. To override
the warnings and allow the installation to complete,
remove the failed installation and set and export the
environment variable:
```

```
NONABI_SCRIPTS=TRUE
export NONABI_SCRIPTS
```

```
Reissue the pkgadd command and following the warnings,
answer 'yes' when asked if you want to continue the
installation.
```

```
=====
```

```
pkgadd: ERROR: preinstall script did not complete
successfully
```

```
Installation of <Clatf> failed.
```



### **CAUTION**

**Overriding installation warnings creates an unsupported Navisphere ATF system configuration.**

## What Next?

Do one of the following:

- Install ATF for the first time with the boot disk outside a storage system (on page 2-6).
- Install ATF for the first time with the boot disk in a storage system (on page 2-9).
- Install a newer revision of ATF on a system that already uses ATF (on page 2-12).

## Installing ATF with the Boot Disk Outside a Storage System (First-Time Install)

When you are ready to install ATF for the first time on a system with the boot disk outside a storage system, follow these steps.

Before installing ATF, be sure to read the **Caution** on page 2-3.

1. Insert the CD-ROM labeled “EMC Navisphere Application Transparent Failover vX.X...” in the CD-ROM drive.

The CD will be mounted automatically on `/cdrom/unnamed_cdrom`, which is linked to `/cdrom/cdrom0`. In the following example, we use `/cdrom/cdrom0`.

2. Log in to the server as root.
3. Load from the CD by typing

```
# /usr/sbin/pkgadd -d /cdrom/cdrom0/CLatf.pkg
```

The following messages appear:

```
The following packages are available:
```

```
1      CLatf Application Transparent Failover Driver
      for EMC...
      (sparc) X.xx
```

```
Select package(s)...to process (or "all" to process all
packages). (default:all)
```

4. Press **Enter** to install the ATF package.

The system console displays a series of copyright messages. Then it prompts

```
The selected base directory </opt/CLatf> must exist
before installation is attempted.
```

```
Do you want this directory created now [y,n,?] y
```

5. If this directory question is asked, enter **y** to create the directory.

The installation continues with the warning

```
This package contains scripts which will be executed
with super-user permission during the process of
installing this package.
```

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

6. Enter **y** to continue installation.

The installation script executes and on completion displays the message

```
=====
You should now verify the installation of ATF using
the command:
    pkginfo -l CLatf

Then you must reboot the system using the command:
    reboot -- -r

Installation of <CLatf> was successful.
```

7. Use the **pkginfo** command to verify the installation. Enter **# /usr/sbin/pkginfo -l CLatf**

```
PKGINST:CLatf
NAME      Application Transparent Failover Driver
          for EMC. . .
CATEGORY: system
.
PSTAMP:   CLatf: 2000:04:28:13:03:01
INSTDATE: April28 2000 13:03
STATUS:   Completely installed
FILES:    xxxxxxxx
          xxxxxxxx xxxxxxxx
```

8. Reboot the server by entering the following command:

```
# reboot -- -r
```

If the disk on which you installed ATF is not the default boot device, you must type a **reboot** command that specifies the full boot device name. For example, **reboot -- devicename -r**.

You have finished the first-time installation of ATF on a server where the system disk will not be in the storage system. The ATF software is in directory **/opt/CLatf**.

## For Servers in a Fabric Environment

At this point, there are no persistent bindings in either the `/kernel/drv/lpfc.conf` or `/kernel/drv/fcaw.conf` files, since there are no valid nodes to communicate to the storage system's LUN 0.

All `fcaw.conf` variables are set except for the following:

- `fcw-bind-WWPN`

All `lpfc.conf` variables are set except for the following:

- `fcw-bind-WWPN`
- `zone-rscn = 1`
- `rscn-ns-delay = 3`

---

Add persistent bindings after you finish configuring the storage system. (page 2-25).

---

## What Next?

- To increase or decrease the maximum number of LUNs that ATF supports, continue to the section, *Changing the Maximum Number of LUNs Supported by ATF* on page 2-17.
- To install the Host Agent and the optional Navisphere CLI software, continue to the section, *Installing the Host Agent and CLI Software* on page 2-18.

## Installing ATF with the Boot Disk in a Storage System (First-Time Install)

When you are ready to install ATF for the first time on a system with the boot disk (directory `/root`, `/usr`, or `/tmp`) in a storage system, follow these steps.

Before installing ATF, be sure to read the **Caution** on page 2-3.

When you install Navisphere ATF onto a host that is booted from a storage system, the storage system automatically reboots twice. The process of configuring the host, and passing control of the boot device to the Navisphere ATF driver, necessitates these reboots. An update installation of Navisphere ATF onto a host that is booted from a storage system includes no automatic reboots.

Until the host is rebooted, the initial Navisphere ATF driver is still in system memory and will be in control of all Navisphere ATF devices

1. Configure the system using the default `sd` driver. If the LUNs are already bound, each LUN must be assigned to its original owner SP. (This will always be true unless you have manually transferred control of a LUN to the other SP.)
2. Insert the CD labeled "EMC Navisphere Application Transparent Failover vX.X..." in the CD-ROM drive.

The CD will be mounted automatically on `/cdrom/unnamed_cdrom`, which is linked to `/cdrom/cdrom0`. In the following example, we use `/cdrom/cdrom0`.

3. Log in to the server as root.
4. Load from the CD by typing

```
# /usr/sbin/pkgadd -d /cdrom/cdrom0/CLatf.pkg
```

The following messages appear:

```
The following packages are available:
```

```
1      CLatf Application Transparent Failover Driver
      for EMC....
      (sparc) X.xx
```

```
Select package(s)...to process (or "all" to process all
packages). (default:all)
```

5. Press **Enter** to install the ATF package.

The system console displays copyright messages. Then it prompts

```
The selected base directory </opt/CLatf> must exist
before installation is attempted.
```

```
Do you want this directory created now [y,n,?] y
```

6. If this directory question is asked, enter **y** to create the directory.

The installation continues with the warning

```
This package contains scripts which will be executed
with super-user permission during the process of
installing this package.
```

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

7. Enter **y** to continue installation.

The installation script executes and on completion, displays the messages

```
=====
Installing ATF on a . . . used as a boot device is a
multi-phase operation that requires multiple
reboots. The first phase of the ATF installation is
complete. The system will now be rebooted
automatically, the second phase of the ATF install
will be executed, and the system will be again
rebooted automatically.
```

```
Initiating the atf_reboot_command script
Sending output to nohup.out
```

```
Installation of <CLatf> was successful.
```

After the second reboot, the installation of ATF is complete.

8. Use the **pkginfo** command to verify the installation. Enter **# /usr/sbin/pkginfo -l CLatf ↵**

```

PKGINST:CLatf
  NAME    Application Transparent Failover Driver
          for EMC . . .
  CATEGORY:system
  .
  PSTAMP:CLatf: 2000:04:28:13:03:01
  INSTDATE:April28 2000 13:03
  STATUS:Completely installed
  FILES: xxxxxxxx
          xxxxxxxx xxxxxxxx

```

You have finished the first-time installation of ATF on a server where the system disk will be in the storage system. The ATF software is in directory `/opt/CLatf`.

### For Servers in a Fabric Environment

At this point, there are no persistent bindings in either the `/kernel/drv/lpfc.conf` or `/kernel/drv/fcaw.conf` files, since there are no valid nodes to communicate to the storage system's LUN 0.

All `fcaw.conf` variables are set except for the following:

- `fcw-bind-WWPN`

All `lpfc.conf` variables are set except for the following:

- `fcw-bind-WWPN`
- `zone-rscn = 1`
- `rscn-ns-delay = 3`

---

Add persistent bindings after you finish configuring the storage system (page 2-25).

---

### What Next?

- To increase or decrease the maximum number of LUNs that ATF supports, continue to the section, *Changing the Maximum Number of LUNs Supported by ATF* on page 2-17.
- To install the Host Agent and the optional Navisphere CLI software, continue to the section, *Installing the Host Agent and CLI Software* on page 2-18.

## Installing a Newer Revision of ATF on a System That Already Uses ATF

This section explains installing ATF in an installation that already uses ATF in a configuration you want.

If the ATF version you have installed is 1.2.3 or later, you can add an ATF package before removing the previous revision. (Two versions of ATF can exist on the same server.) If you add an ATF package when ATF is already installed, the script names the new installation **CLatf.2**; then when you remove the old revision, **CLatf.2** becomes the only revision. If you then add an ATF package while **CLatf.2** is the only revision, the script names the newer one **CLatf**. The installation software alternates the names **CLatf** and **CLatf.2**. We recommend adding an ATF package before removing the previous revision if you can because it simplifies the installation. Continue and add the package as explained next.

If you already have ATF installed, and the version you have is lower than 1.2.3, then you must remove the previous revision before you can add an ATF package of 1.2.3 or later. Removing a revision is explained starting on page 2-36.

When you are ready to install ATF on a system that has ATF installed, follow these steps.

1. Insert the CD labeled "EMC Navisphere Application Transparent Failover vX.X..." in the CD-ROM drive.
2. If the Host Agent is running, stop the Agent.

The CD will be mounted automatically on `/cdrom/unnamed_cdrom`, which is linked to `/cdrom/cdrom0`. In the following example, we use `/cdrom/cdrom0`.

3. Log in to the server as root.
4. Load from the CD by typing

```
# /usr/sbin/pkgadd -d /cdrom/cdrom0/CLatf.pkg
```

The following messages appear:

```
The following packages are available:
```

```
1      CLatf Application Transparent Failover Driver
      for EMC...
      (sparc) X.xx
```

```
Select package(s)...to process (or "all" to process all
packages). (default:all)
```

5. Press **Enter** to install the ATF package.

The system console displays a series of copyright messages. Then it prompts

```
The following files are already installed on the system
and are being used by another package:
```

```
/.atf_load.sh
/.atf_load_driver
...
```

```
Do you want to install these conflicting files
[y,n,?,q] y
```

6. Enter **y** to install the files.

The installation continues with the warning

```
This package contains scripts which will be executed
with super-user permission during the process of
installing this package.
```

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

7. Enter **y** to continue installation.

The installation script continues. You will see many messages, including this one:

NOTICE:

```
Please read carefully any update information in the
Product Release Notes. You may especially want to
check for any changes in the default time-out and
retry values given in /opt/CLatf/install/atf.conf
and /opt/CLatf/install/atfi.conf. Compare these
values to those you are already using. In general,
your time-out and retry values should be AT LEAST as
large as the defaults in the above files.
```

```
=====
```

```
Your next step should be to remove the old instance
of CLatf -- just type:
pkgrm CLatf
```

8. If the boot disk is in a storage system, you will see the following message about multiple reboots:

```
=====
Installing ATF on a . . . used as a boot device is a
multi-phase operation that requires multiple
reboots. The first phase of the ATF installation is
complete. The system will now be rebooted
automatically, the second phase of the ATF install
will be executed, and the system will be again
rebooted automatically.
```

```
Initiating the atf_reboot_command script
Sending output to nohup.out
```

```
Installation of <CLatf> was successful.
```

After the second reboot, the update of ATF is complete. Skip to step 10.

9. If the boot disk is *not* in a storage system, you will see the following message:

```
=====
You should now verify the installation of ATF using
the command:
  pkginfo -l CLatf*
```

```
Then you must reboot the system using the command:
reboot -- -r
```

```
Installation of <CLatf> was successful.
```

10. Use the **pkginfo** command to verify the installation. Enter

```
# /usr/sbin/pkginfo -l CLatf
```

```
PKGINST:CLatf
  NAME:Application Transparent Failover Driver
        for EMC
  CATEGORY:system
.
PSTAMP:CLatf: 1998:08:28:13:03:01
INSTDATE:Sept 21 1998 17:0
STATUS:Completely installed
...
PKGINST:          CLatf.2
NAMEApplication Transparent Failover Driver..
CATEGORY:system
.
PSTAMP:CLatf: 1998:08:28:13:03:01
INSTDATE:Sept 21 1998 17:02
STATUS:Completely installed
PKGINST:CLatf
  NAME:Application Transparent Failover Driver
        for EMC . . .
  CATEGORY:system
.
PSTAMP:CLatf: 2000:04:28:13:03:01
INSTDATE:April28 2000 13:03
STATUS:Completely installed
...

```

The INSTDATE entries in the pkginfo display above show which CLatf (**CLatf** or **CLatf.2**) is more recent. The display above shows that **CLatf.2** is more recent.

11. Remove the previous version of ATF that was installed before you installed this version. To continue the example above, you would type

```
# pkgrm CLatf
```

```
The following package is currently installed:
CLatf Application Transparent Failover...
```

```
Do you want to remove this package?
```

12. Confirm the answers to each script question by entering **y**.

13. Use the **pkginfo** command to verify the removal. The display should show just one CLatf. Enter

```
# /usr/sbin/pkginfo -l Clatf ↵
```

```
PKGINST:CLatf.2
  NAMEApplication Transparent Failover Driver..
  CATEGORY:system
  ...
```

14. Perform a reconfigure reboot:

```
# reboot -- -r ↵
```

You have finished installing the newer release of ATF.

---

The Host Agent must be compatible with ATF. For information about the compatibility between Agent and ATF revisions, consult the product release notice.

---

### What Next?

- To increase or decrease the maximum number of LUNs that ATF supports, continue to the next section, *Changing the Maximum Number of LUNs Supported by ATF*.
- To install the Host Agent and the optional Navisphere CLI software, continue to the section, *Installing the Host Agent and CLI Software* on page 2-18.

## Changing the Maximum Number of LUNs Supported by ATF

The Navisphere ATF driver supports up to 223 LUNs per storage system. The LUN IDs range from 0 to 222. The default maximum number of LUNs that ATF supports is 32 per storage system. The LUN IDs range from 0 to 31. If there are more than 32 LUNs bound on a storage system, they will not be visible and usable with the Navisphere ATF driver unless the maximum LUN ID is increased.

---

When you increase the number of LUNs supported by ATF, you proportionally increase the time needed to reboot the system.

---

This section outlines the recommended procedure to increase or decrease the maximum LUN ID.

1. Edit the `/kernel/drv/atf.conf` and the `/kernel/drv/atfi.conf` files by changing the value of `atf_scsi_max_lun` to the new value.

The maximum number of LUNs that Navisphere ATF supports is 223.

2. Reboot the system.

---

A reconfiguration reboot is not necessary here.

---

3. Enter the following command:

```
atf_configure -y
```

4. Reboot the system using the following command:

```
reboot -- -r
```

### What Next?

Install the Host Agent and the optional Navisphere CLI software as explained in *Installing the Host Agent and CLI Software* on page 2-18.

## Installing the Host Agent and CLI Software

The Navisphere Installation CD-ROM contains two Navisphere software packages: Host Agent and Navisphere Command Line Interface (CLI).

1. If you are installing Host Agent, make sure that each storage system is connected to the Solaris server where you are installing Agent.
2. At the Solaris server, log in as superuser (root).
3. If any previous revision of the Host Agent or CLI is installed on the server, remove it before continuing (see page 2-37); otherwise continue to the next step.
4. Insert the Agent/CLI CD-ROM into the host's CD-ROM drive.
5. In a shell, enter the following commands:

```
cd /cdrom/cdrom0
```

```
pkgadd -d ./naviagnt.pkg
```

The following appears:

```
1 NAVIAGENT  Navisphere Disk Array Management Tool
  (AGENT)
                (sparc) 4.x.x
2 NAVICLI    Navisphere Disk Array Management Tool
  (CLI)
                (sparc) 4.x.x
Select package(s) you wish to process (or 'all' to
process
all packages). (default: all) [?,??,q]:
```

6. Specify the Agent and/or CLI.

Enter **1** (to install the Host Agent)

or

Enter **2** (to install Navisphere CLI)

or

Enter **3** (to install both the Host Agent and CLI)

The following appears:

```
Processing package instance <NAVIAGENT> from...
.
Do you want to install these conflicting files [y,n]
```

7. Enter **y**.

```
This package contains scripts which will be executed
with the super-user permission during the process of
installing this package
```

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

8. Enter **y**.

Now the installation program looks for any Agent configuration files you may already have on your system. If the program does not find any agent configuration files, then you have finished installing the Agent; go to step 10.

If the program does finds any existing configuration files, it displays a message like the following.

```
At Least 1 saved config file exists for Navisphere
Agent.
```

```
Please select 1 of the following:
```

```
[1]
Restore/etc/Navisphere/.Naviagent-config.000120:105
9
[2]
Restore/etc/Navisphere/.Naviagent-config.000121:140
8
[3] Quit
```

```
Select number 1 - 3.
```

9. Select the existing file you want to serve as the agent configuration file. The software will retain that file and rename it to the required Agent filename, **agent.config**.

Generally, you will want to use the most recent file, as shown by the numeric date suffix. To use the default configuration file, specify the number for the Quit option.

10. When the installation of the Host Agent is complete, the following message appears:

```
Installation of <NAVIAGENT> was successful.
```

If you are also installing the CLI, the program prompts

```
:Processing package instance <NAVICLI> from...
.
.
```

11. Answer **y** to the CLI installation prompts as you did for the Agent.
12. When installation is complete, exit the **/cdrom** directory (for example, execute **cd /**), and then remove the CD from the host's CD-ROM drive.

---

Any user who can access the management station can change or delete the Navisphere files you just installed. You may want to change permissions on these files to restrict access to them.

---

**What Next?** Before you can use the Navisphere storage management software and the CLI, you must modify the user login scripts. Continue to the next section.

## Modifying User Login Scripts

Use a text editor to modify login scripts as described below:

- If you are running Common Desktop Environment, uncomment the last line in `$HOME/.dtprofile`. The line should read:

`DTSOURCEPROFILE=true`

- Make the following additions to the specified paths in `$HOME/.profile` or `$HOME/.cshrc`, and export each path:
  - - add `/opt/Navisphere/bin` to `PATH`
  - - add `/opt/Navisphere/man` to `MANPATH`
  - - add `/opt/Navisphere/lib` to `LD_LIBRARY_PATH`

You must **export** each path after you have modified it.

### What Next?

**Non-FC4700 storage systems** - go to the next section, *Configuring the Host Agent for non-FC4700 Storage Systems*.

**Only FC4700 storage systems** - What you do next depends on whether you will use Navisphere Manager to configure the storage systems connected to the server.

- **Using Manager** - If you are setting up a new server, go to *Using the Event Monitor Configuration File* on page 2-24. If you are just upgrading the Host Agent or CLI on an existing system, you are finished with the upgrade.
- **Not using Manager** - Continue to the section, *Using the Event Monitor Configuration File* on page 2-24.

## Configuring the Host Agent for non-FC4700 Storage Systems

You need to configure the Host Agent for non-FC4700 storage systems only. To configure the Host Agent, you edit the Host Agent configuration file. The easiest way to edit this file is with the remote agent configuration feature of Navisphere Manager. Before someone can use this feature to configure the Host Agent, you must manually add a user entry for that person to the Host Agent's configuration file.

---

Any user who can log in to a host that is a Navisphere management station can monitor the status of any of the managed storage systems.

---

The path-name to the Agent configuration file is `/etc/Navisphere/.naviagent.config.xxxx` <link> , where `.xxxx` represents the time-date-stamp.

---

## Adding a User to the Agent Configuration File

1. Add a name entry to the `/etc/Navisphere/agent.config` file:

**user** *name*

where

*name* is the person's username. The format of this name differs depending on whether the person will be using Manager on a local or remote host.

**For a local host** - The format is *user*

where

*user* is the person's user account name.

**For a remote host** - The format is *user@hostname*

where

*user* is the person's user account name.

*hostname* is the name of the remote host.

For example, if you want to allow user **anne** to edit the Host Agent configuration file and configure a host's storage system using the Navisphere Manager running on either remote host **img01** or remote host **img02**, you must add the following entries to the server's Agent configuration file:

**user** **anne@img01**

**user** **anne@img02**

---

To let a user configure a server's storage system(s) from the Navisphere Manager on different hosts, you must add a user entry for each host to the Agent configuration file.

---

2. Save the `/etc/Navisphere/agent.config` file.

---

## Editing the Host Agent Configuration File

Edit the host Agent configuration file so the Agent can communicate with the storage system, and then save the **agent.config** file. Use one of the following methods to edit the **agent.config** file.

**Navisphere Manager** - Refer to the *Navisphere Manager Administrator's Guide*.

**A text editor on the Solaris host** (vi, emacs or textedit) - Refer to *Configuring the Host Agent*, Appendix B.

**What Next?** To verify the device enters in the Agent configuration file, continue to the next section.

---

## Verifying Device Entries in the Host Agent Configuration File

Once you have edited the Host Agent configuration file, verify the device entries.

1. Open the Host Agent configuration file.
2. Verify that there is a clsp entry for each SP.

There will be no entries for FC4700 storage systems.

3. Stop and then start the Agent (page 2-34).

**What Next?** What you do next depends on whether you will use Navisphere Manager to manage the storage systems connected to the server.

**Using Manager** - If you are setting up a new server, go to *Using the Event Monitor Configuration File* on page 2-24. If you are just upgrading the Host Agent or CLI on an existing system, you are finished with the upgrade.

**Not using Manager** - Continue to the next section, *Using the Event Monitor Configuration File* on page 2-24.

## Using the Event Monitor Configuration File

The Navisphere Agents can monitor storage-system events and take such action as sending email or paging you if specified events occur. Navisphere Manager provides an interactive way to define these events and actions. If you do not have Navisphere Manager, you can still define such events and actions by editing the Event Monitor configuration file:

**/etc/Navisphere/Navimon.cfg**

The file is self-documenting; that is, text in it describes how to define events and the actions you want taken if the events occur. You can test the file after editing it using the Navisphere CLI command **responsetest**, explained in the Navisphere 5.X CLI reference.

## Configuring Storage Systems

Configure the storage systems connected to the server as described in the Navisphere Manager User manual.

Configure the storage system only once. That is, bind all LUNs, create all RAID Groups and, for shared storage systems, create all Storage Groups that you want on the storage system.

### What Next?

Do one of the following:

**For a switched fabric environment** - Continue with the next section to add persistent bindings to the HBA driver configuration file.

**For an FC-AL environment** - Go to Chapter 3, *Making LUNs Available to Solaris*.

## Persistent Bindings

ATF adds the persistent binding information to the HBA driver configuration file by specifying the HBA instance and the World Wide Port Name (WWPN) for the target SP. The binding sets a connection path from the HBA port through the switch to the target SP in the storage system. The persistent binding allows you to replace an SP without rebooting the server.



### **CAUTION**

**This section contains notes and warnings relevant to persistent bindings and ATF.**

- **In a SAN configuration, the persistent bindings target IDs assigned to SPs on one host may not be identical to the target IDs assigned to SPs on another host. Currently, ATF does not match the WWPN target ID combination on one host to that found on other hosts on the fabric. These target IDs can be modified by changing the ATF persistent binding entries in the `/kernel/drv/hba` config file (i.e `lpfc.conf/fcaw.conf/fca.conf`).**
- **If the host is booted with a path down, `atf_configure` will not reflect the bindings properly.**
- **After each successful `atf_configure`, a backup copy of the HBA configuration file will be kept as `xxxx.conf.atf`. This file can be used should an error occur during the latest `atf_configure`. The user may want to back this file up separately to ensure a good HBA config file at all times.**

## **atf\_configure Utility**

This utility creates the `/kernel/drv/atf.conf` file and if in a SAN environment, will perform the persistent bindings. The following options are supported by the **atf\_configure** utility:

- y** eliminates the **atf\_configure** queries and indicates a **yes** response to them. The queries pertain to retaining old versions of **atf.conf** and **sd.conf** as well as retaining parameter settings in **atf.conf** if the value differs from the default.
- h** displays a help message listing the various options supported by **atf\_configure**.
- a #** adds  $2 * \#$  (where  $\#$  is the number of arrays to add to the **atf.conf** file). Targets 0 through 7 are automatically added to the file during the basic execution of **atf\_configure**. With this option,  $7 + 2 * \#$  will be the number of targets residing in **atf.conf**. A variable called `atf_max_id` contains the max number of targets located in the **atf.conf** file.
- J** adds the `clfca` and `clfcaw` entries to the **atfi.conf** file so that during a subsequent reconfiguration reboot, the `atfi` driver can detect any devices under the `clfca` or `clfcaw` driver. After the reconfiguration reboot, an execution of **atf\_configure** will process the `clfca/clfcaw` prototypes in the `\kernel\drv\atf.conf` file. The `atf` driver will detect these devices on a subsequent reconfiguration reboot.
- t** calculates new throttle settings for the arrays located on the system.
- C** creates the `/kernel/drv/atf_conf.tble` file with `clsp` information. The output file contains the WWPN (if fabric), target ID, and the associated `/dev/clsp` entry for each CLARiON SP detected on the system.

The following options pertain only to a SAN configuration and persistent bindings:

- b uses the actual SP ID as the target ID for performing the persistent bindings. Without this option, **atf\_configure** will calculate the target ID beginning with 0 for SPAs and 1 for SPBs. Any existing bindings in the HBA configuration files will be ignored. During **atf\_configure**, the user will be prompted with a table displaying the SP (whether SPA or SPB), the WWPN for this array, and the target ID. The utility asks if you want to use this information or exit **atf\_configure**. If you exit, no changes will be made to **atf.conf** or the HBA configuration file. If you use these target IDs, then the persistent bindings will be performed using the displayed information.
- d deletes a target ID from the persistent bindings. If the storage system's WWPN is still detected by the **atf\_configure** utility, a binding will exist for this array in the HBA configuration file. The option is executed as follows:  
**atf\_configure -d #** (where # is the target ID to be deleted)
- c cleans up the persistent bindings. Once all the bindings have been preserved, **atf\_configure** will detect which bindings no longer have nodes on the system and will display one at a time - the target and WWPN. The user can retain the binding or delete it.
- B does not preserve existing persistent bindings that exist in the HBA configuration file in **/kernel/drv**. Persistent bindings will be performed, but the target IDs associated with each WWPN may not be the same as before since this option tells **atf\_configure** to ignore previous bindings.

### What Next?

If you need to edit existing DID or WWNN bindings, continue to the section, *Editing Existing Persistent Bindings* on page 2-30. If you do not need to edit existing bindings, go to the next section.

---

## Adding and Modifying Persistent Bindings

This section describes how to add the persistent bindings to the driver configuration files (`/kernel/drv/fca.conf`, `/kernel/drv/fcaw.conf`, or `/kernel/drv/lpfc.conf`).

1. After configuring the storage system, add persistent bindings to the driver configuration file by entering the following command:

### **atf\_configure**

This command adds persistent bindings to the driver configuration file and sets the remaining driver configuration file variables. At this time, ATF has set the following `fcaw/fca.conf/lpfc.conf` variables:

#### **Fcaw/fca variables set:**

```
fcw-bind-WWPN
automap = 1;
linkdown-timeout = 10;
linkdown-fastfail = 2;
target-down-timeout = 10;
target-down-fastfail = 2;
```

#### **Lpfc variables set:**

```
fcw-bind-WWPN
automap = 1;
linkdown-tmo = 10;
nodev-tmo = 20;
dqfull-throttle = 0;
zone-rscn = 1;
rscn-ns-delay = 3;
```

2. Verify the target IDs of the persistent binding entries by viewing the contents of the driver configuration file.

3. If you need to modify the target IDs, do the following:
  - a. Open the driver configuration file (**fcaw.conf** or **lpfc.conf**), and scroll through until you find the following text and the persistent binding entries.

```
#EMC-CLARiiON ATF BINDINGS (today's date)
#used ONLY if supported by present version of the HBA
driver.
```

```
fcplib-bind-WWPN = "50060160fe12510b:lpfc0t0",
                  "50060168fe12510b:lpfc1t1";
```

where

50060160fe12510b is the World Wide Port Name (WWPN) for the target SP, t0.

lpfc0 is the instance of the HBA, lpfc0, connected to t0.

t0 is the SP target ID

In this persistent binding example, a dual SP (t0, t1) storage system is connected through the switch to two PCI bus adapters installed in the server.

- b. Modify the *tn* values in the persistent binding entries to match the switch settings on the back of each SP.
- c. Save any changes you make to the driver configuration file, and enter the following command:  
**atf\_configure**
- d. Reboot the server using the following command:  
**reboot -- -r.**

## Editing Existing Persistent Bindings

This section provides steps for editing existing persistent binding information.

1. After installing ATF, the ATF configuration file, **atf.conf**, should reflect the targets that were assigned during persistent reservations.

Neither the JNIfca nor the lpfc driver supports simultaneous binding by DID, WWPN or WWNN.

2. If DID bindings were previously used and you want to continue to use them, do the following:
  - a. Comment out the WWPN bindings in the HBA configuration file.
  - a. Check and verify that the **/kernel/drv/atf.conf** file contains prototype entries for LUNs 0-31 for each target that is in the DID bindings. If the entries are not present, then execute **atf\_configure -a #**. This command adds two prototype target IDs for each additional storage system. See *atf\_configure Utility* on page 2-26.

You must, once again, comment out the WWPN bindings in the HBA configuration file.

3. If DID or WWNN bindings were previously used and you want to use the ATF WWPN bindings, do the following:
  - a. Transfer the target IDs to the WWPN bindings to ensure continuity on the system. Refer to step 3, on page 2-29.
  - b. Log in to the fabric switch and use the command **switchshow** (for Brocade switches) to show which HBAs and storage systems are logged in.

The WWPN of the storage system and adapters will be displayed along with the corresponding port on the storage system to which they are logged in. The DID information should provide the port to which a specific target is connected.

- c. Find the WWPN at that port and edit the corresponding WWPN for the binding to reflect the old DID target ID.

- d. Perform **atf\_configure** again to ensure that the **atf.conf** file reflects the target prototypes for the persistent bindings.

Below is an example of what the binding will look like:

```
# EMC-CLARiION ATF BINDINGS Tue Apr 18 16:01:46 2000
# Used ONLY if supported by present version of the
HBA driver.

fcp-bind-WWPN= "50060168fe89a2ce:fcaw0t1",
               "50060160fe89a2ce:fcaw1t0";
```

The above example shows a single storage system connected to the switch from the two JNIfca adapters. The bindings will be time-stamped and a backup version of the last HBA config file will be kept in **/kernel/drv/xxxx.conf.atf** (where xxx is: lpfc, fca or fcaw).

The template for bindings is:

```
WWPN:driver_name#t%
```

where:

WWPN is the World Wide Port Name for the SP.

driver\_name is the driver's name - lpfc, fcaw, or fca.

# is the instance number of the driver.

t% is the target ID assigned to that SP.

**Note:** If the user wants to modify the target ID, edit only the % value in the applicable HBA driver config file. (i.e **lpfc.conf** or **fcaw.conf/fca.conf**).

At this time, ATF has set the following **fcaw/fca.conf/lpfc.conf** variables:

**Fcaw/fca variables set:**

```
fcp-bind-WWPN
automap = 1;
linkdown-timeout = 10;
linkdown-fastfail = 2;
target-down-timeout = 10;
target-down-fastfail = 2;
```

**Lpfc variables set:**

```
fcplib-bind-WWPN
automap = 1;
linkdown-tmo = 10;
nodev-tmo = 20;
dqfull-throttle = 0;
zone-rscn = 1;
rscn-ns-delay = 3;
```

---

If the target IDs are not present in the **atf.conf** file, and you reboot the system, these targets will not be configured by the system.

---

4. When you have finished editing the bindings, reboot the system using the following command:

```
reboot -- -r
```

---

## Verifying the HBA Driver and ATF Connections to SPs

---

For Fabric environments only. After adding the persistent bindings to the HBA configuration file, verify the driver and ATF connections to the SP.

---

1. List the contents of the **/dev** directory using the following commands:

```
cd /dev
ls
```

2. In the listing, look for the following entries for each SP:

```
atf -sp:cDtS where
```

*D* is the HBA number and *S* is the SP ID.

An `atf -sp:cDtS` entry means ATF sees the SP.

3. If any of these entries are not in **/dev**, make sure you have the correct persistent binding information in the HBA driver file.

## Verifying the Server Can See Its LUNs

Use Navisphere Manager to verify that the server can see its LUNs.

1. For each storage system connected to the server, use Navisphere Manager to verify the server can see its LUNs.
  - a. Select the storage system for management.
  - b. Click the **Hosts** tab.
  - c. Double-click the icon for the Solaris server.
  - d. Double-click the LUNs icon.
  - e. Look for an icon for each LUN that the server should have.
2. If an icon for each LUN exists, the persistent bindings are set correctly, go to Chapter 3, *Making LUNs Available to Solaris*.
3. If the icon for any LUN is missing, check that the LUN belongs to a Storage Group connected to the server as follows:
  - a. Double-click the Storage Groups icon.
  - b. Right click the icon for the Storage Group that should contain the missing LUN, and click **Properties**.
  - c. In the **Storage Group Properties** dialog box, look for the missing LUN in the **LUNs in Storage Group** list.
  - d. If the LUN is not listed, click the **Select LUNs** button.
  - e. In the **Modify Storage Group** dialog box, look for the missing LUN in the **Select LUNs for Storage Group** list.
  - f. If the LUN is not listed, click **Show LUNs in Other Storage Groups**.

If the LUN is listed, then move the LUN from its current Storage Group into the Storage Group for the server by clicking the right arrow and then **OK**.

If the LUN is still not listed, then the persistent bindings value in the **fcaw.conf** or **lpfc.conf** file is incorrect or the LUN is missing from the **sd.conf** file.

## Starting and Stopping the Host Agent

On a Solaris host, the Host Agent starts automatically when you bring the host up to init level 3. When you first start the Host Agent, look at the system log for the host's operating system to make sure the Agent started and no device errors occurred.

When a storage system experiences heavy input/output traffic (that is, the applications are using the storage system), information may not be reported to the Host Agent in a timely manner, resulting in the Agent taking several minutes to execute a storage system management task. This behavior is most evident when one Agent is managing multiple storage systems. Also, if the SP event log is large and the Agent configuration file is set up to read all the events, it may take a few minutes for the Host Agent to start.

### To Start the Host Agent

Log in as root and enter this command:

```
/etc/init.d/agent start ↵
```

### To Stop the Host Agent

Log in as root and enter this command:

```
/etc/init.d/agent stop ↵
```

### What Next?

To learn about logging Host Agent status and removing the Host Agent and CLI, continue to the next section.

---

## Logging Host Agent Status on the Host

The operating system logs Core Software error messages about Host Agent events and serious Agent errors in its system log (`/var/adm/messages`). It logs messages about such Agent events as startup and shutdown, and errors such as the Agent's inability to access a device in its configuration file. If you need information about the Agent's status, look for system log messages about the Agent.

---

Informational messages about Agent events are recorded in `/var/adm/log/agent.log`.

## Removing (CDE) or ATF Software

If you need to remove CDE (CLARiiON Driver Extensions) or ATF software from a server, follow these steps:

In the examples given, CDE is being removed. If ATF is being removed, the package names will be **CLatf**, instead of **CLcde**.

CDE uses some commands and support software also used by ATF. As a result, many messages, filenames, and device names use the name *atf* instead of *cde*. These occurrences of the name *atf* are accurate and are not errors in the software or documentation.

1. If you are removing the only existing copy of CDE or ATF, and if any logical volumes and metadb copies that use CDE or ATF were created under the Solstice Disk Suite, you must remove them. If you do not remove these devices before removing CDE or ATF, your system will not be bootable.
2. If the Host Agent is running, stop the Agent.
3. Copy the **/kernel/drv/sd.conf.atf** file saved by the CDE or ATF installation process to **/kernel/drv/sd.conf**.

The destination **sd.conf** file must contain all controller, ID, and LUN entries for all disk-array storage-system devices that will be removed from CDE or ATF control.

4. Reboot the system using the following command:

```
# reboot -- -r
```

5. Remove CDE or ATF by entering the following command:

```
# pkgrm CLxxx where xxx is either cde or atf.
```

The following package is currently installed:

```
CLcde CLARiiON Driver Extensions
      Driver for EMC . . .
      (sparc) X.X.X
Do you want to remove this package?
```

6. If the package listed is the correct one, enter **y** to remove the package.

```
## Removing installed package instance <CLcde>
. . .
Do you want to continue with the removal of this
package [y,n,?,q] y
```

7. Enter **y** to continue. Several messages will appear, including

```
Once the de-install is complete, you
  must reboot the system using the command:
  # reboot -- -r
  ...
Removal of <CLCde> was successful.
  ...
```

8. Reboot the system using the reconfiguration form of the reboot command:

```
# reboot -- -r
```

You have removed either the CDE or ATF software from this server.

9. Verify which failover package is running by typing the following command.

```
atf_control
```

---

If you have removed both CDE and ATF, **atf\_control** will not be available.

---

## Removing the Host Agent or CLI

---

To avoid losing site-specific information, copy the Host Agent configuration file (**agent.config**) to a safe location before removing any Navisphere software.

---

The following steps describe how to uninstall the Host Agent (and the CLI if required).

1. Log in to the host as superuser (root) and start the X Window System™ if it is not started.
2. At the command line in an X Window terminal emulator, enter the following commands:

```
/etc/init.d/agent stop ↵
cd ↵
pkgrm package ↵
```

where

*package* is **NAVIAGENT** or **NAVICLI** (these names must be capitalized)

The removal program begins.

The following package is currently installed:

```
....  
Do you want to remove this package?
```

3. Enter **y** ↵

```
This package contains scripts which will be executed  
with super-user permission during the process of  
removing this package.
```

```
Do you want to continue with the removal of this  
package [y,n,?,q]  
Enter y ↵  
#Removing installed package instance ....  
.  
Removal of .....package..... was successful.
```

You have completed removing the Host Agent or CLI.

---

The removal program removes only the files that were installed from the installation CD. Files that the CLI creates (for example, log files) are not removed.

---

## Making LUNs Available to Solaris

This chapter describes how to specify Solaris disk names for LUNs and describes the following tasks that you must perform to make LUNs in the server's Storage Group available to Solaris.

- Specifying Solaris Disk Names for LUNs.....3-2
- Labeling and Partitioning LUNs.....3-2
- Making File Systems on Partitions with newfs .....3-3
- Mounting File Systems.....3-4

## Specifying Solaris Disk Names for LUNs

This section describes the format for specifying Solaris disk names for LUNs.

For Solaris, *diskname* has the format

**cDtSdLsP**

where

- |          |   |
|----------|---|
| <b>D</b> | is the number of the HBA in the server. Solaris assigns these numbers. For example, the number for an HBA in slot 1 is <b>0</b> ; for an HBA in slot 2, it is <b>1</b> . Under some conditions, Solaris may assign other numbers to HBAs. |
| <b>S</b> | is the target ID of the SP that is connected to HBA <b>cX</b> . Y can range from 0 through 125.   |
| <b>Z</b> | is the LUN number. Z can range from 0 through 222. For example, for LUN 2 in the storage system, Z is <b>2</b> .  |
| <b>P</b> | is the partition number on the target.  |

For example, if the HBA is 1, the target ID of the SP is 0, and the LUN number is 2, you would format the disk by entering the following command:

```
format c1t0d2 ↵
```

**What Next?** Continue to the next section to label and partition the LUNs.

## Labeling and Partitioning LUNs

This section describes how to use the **format** command to label and partition LUNs. The version of the **format** command that ships with Solaris has an Auto configure option that configures LUNs (disks) for you.

1. At the root prompt, enter the following command:

```
format diskname ↵
```

where *diskname* is in the form explained in the previous section, *Specifying Solaris Disk Names for LUNs*. If you use the **format** command with no arguments, it displays all disks.

The **Format** menu opens.

2. From the **Format** menu, select **Available Drive Types**.
3. At the `Specify disk type` prompt, enter **0** to select the Auto configure option.
4. If you want to partition the disk, use the **format** command again.  
See the Solaris documentation for information on the **format** command.
5. If you have not labeled the disk, label it using the **label** command.  
See the Solaris documentation for information on the **label** command.

**What Next?** Continue to the next section to make file systems on the LUN partitions.

---

## Making File Systems on Partitions with `newfs`

Use the `newfs` command to make file systems on all newly created partitions. You must specify a character device name to indicate the partition where you want to create the file system.

For example, to make a file system on partition 0 on the first disk (LUN), enter the following command:

```
newfs /dev/rdisk/c1t0d0s0 ↵
```

**What Next?** Continue to the next section to mount the file systems.

## Mounting File Systems

This procedure is the same as for any LUN used with a Sun operating system. Use the **mount** command to mount the file systems that you have created on your storage system. You must specify the block device name of the partition where you created the file system, and the mount point directory.

For example, to mount the file system partition **0** of the first disk configured, at the mount point **/temp**, enter the following command:

```
mount /dev/dsk/c1t0d0s0 /temp ↵
```

### What Next?

If you have more than one server connected to the storage system and you will run the FirstWatch failover software, or the VERITAS Cluster Server (VCS) software, continue with Chapter 6, *Modifying VERITAS FirstWatch Scripts*. Otherwise, the server is ready to use the file systems.

---

This chapter describes how to boot from a storage-system device using either an SBus or PCI Fibre Channel adapter.

- Using a Storage-System LUN as a Boot Device (SBus) .....4-2
- Using a Storage-System LUN as a Boot Device (PCI) .....4-7

## Using a Storage-System LUN as a Boot Device (SBus)

Booting from a LUN on a storage system has been verified for Ultra 2 and Enterprise 3x00, 4x00, and 6x00 systems. Since the Sun installation CD does not support the JNI device, installation directly to a storage-system LUN is not possible using the Sun installation CD. The following is a modified procedure supplied by JNI. It entails building an installation on a non-storage-system disk, saving the installation with **ufsdump**, restoring the installation using **ufsrestore** to a properly prepared the storage-system LUN, and modification of the restored installation to permit booting from the storage-system LUN.

### Installing Solaris on a Storage-System LUN (drive)

The following is a procedure for installing the Solaris operating system on a storage-system LUN (drive). When the installation procedure is complete, the host can boot and run from the LUN (drive).

- Partition**
1. Partition the storage-system drive to accommodate the Sun OS partitions.
    - a. Use the **format** command to examine the partitions of the system drives where the current OS resides.

**format**

(choose disk 0)

format>**verify**

sample output:

part	tag	cyl	size
0	root	0_335	124.69 MB
1	swap	336_792	169.59 MB
2	backup	0_5846	2.12 GB
	.		
	unassigned		
	.		
6	usr	739_2305	561.4 MB
7	home	2306_5846	1.28 GB

- b. Quit partition, and choose the disk on which you want to install the OS:

**quit**

```
format> disk
```

```
AVAILABLE SELECTIONS ____
```

```
specify disk:
```

```
(Enter number of disk)
```

```
1
```

```
format> partition
```

```
partition> print
```

- c. Repartition the drive (if necessary):

```
partition> label
```

```
exit format
```

## Filesystems

2. Create file systems on the required partitions. In this example the OS is being copied to target 3 of the storage-system drives on controller 1.

To use a different target, change t3 to t(*target id*).

To use a different controller, change c1 to c(*controller number*):

```
newfs /dev/rdisk/c1t3d0s0
```

```
newfs /dev/rdisk/c1t3d0s6
```

```
newfs /dev/rdisk/c1t3d0s7
```

## Bootblk

3. Install a **bootblk** on the root partition of the drive:

- a. To determine the directory in **/usr/platform** where the **bootblk** will be found, enter the following command:

```
uname -m
```

Solaris displays the machine hardware class for your system, for example:

```
sun4u
```

- b. Insert this machine hardware class in the pathname for the **bootblk**, as shown in the following example.

```
/usr/sbin/installboot /usr/platform/sun4u/lib/fs/ufs/bootblk  
/dev/rdisk/c1t3d0s0
```

where

*sun4u* is the machine hardware class for your system.

*/dev/rdisk/c1t3d0s0* is the root device for your system.

- Files**
- 4. Create the required files and directories on the Fibre Channel root directory as follows:

- a. Mount the Fibre Channel root directory to **/mnt** so that files can be copied to it.

```
mount /dev/dsk/c1t3d0s0 /mnt  
cd /mnt
```

In this example the drive is controller 1, SCSI target 3, and LUN 0.

- b. Dump and restore the root filesystem to this drive:

```
ufsdump 0f - /dev/dsk/c0t0d0s0 | ufsrestore rf -
```

- /etc/vfstab**
- 5. In the **/mnt/etc** directory update the **vfstab** file to indicate the drives to be mounted during boot:

```
cd /mnt/etc  
vi vfstab
```

Before modification:

```
#device device mount FS fsck mount mount  
#to mount to fsck point type pass at boot options  
#  
  
/dev/dsk/c0t0d0s1 _ _ swap _ no _  
  
/dev/dsk/c0t0d0s0 /dev/rdsk/c0t0d0s0 / ufs 1 no _  
swap _ /tmp tmpfs _ yes _
```

After modification:

```
#device device mount FS fsck mount mount
#to mount to fsck point type pass at boot options

#

/dev/dsk/c1t3d0s1 _ _ swap _ no _
/dev/dsk/c1t3d0s0 /dev/rdisk/c1t3d0s0 / ufs 1 no _
swap _ /tmp tmpfs _ yes _
```

- /etc/system** 6. If ATF is installed, add a **rootdev** entry to **/etc/system**, in the following format:

**rootdev:pathname**

where

*pathname* is the root partition you just installed on the storage system. For example, enter:

**rootdev:/sbus@1,f800000/fcaw@2,0/atf@3,0**

- /usr** 7. Copy the current **/usr** directory structure to partition 6 of the Fibre Channel drive with the **ufsdump** and **ufsrestore** commands:

In the current root directory:

```
cd /
umount /mnt
mount c1t3d0s6 /mnt
cd /mnt
ufsdump 0f - /dev/dsk/c0t3d0s6 | ufsrestore rf -
```

- OpenBoot Alias** 8. Create an OpenBoot alias to boot from the Fibre Channel drive.:

- a. Power cycle the host to the OBP (Open Boot Prompt) **ok>** prompt.

Do not allow the system to reboot.

- b. Issue the **printenv** command to see the default boot device in the environment variables:

**boot-device disk disk**

- c. Issue the **devalias** command to see the device associated with the alias for the default boot device:

**devalias**

```
disk /sbus/SUNW,fas@e,8800000/sd@0,0
```

- d. Issue the `show-devs` command to see the device name for the drive:

#### **show-devs**

```
/sbus@1f,0/SUNW,fas@e,8800000/sd
/sbus@1f,0/fcaw@1,0/sd@0,0
```

- e. Create a similar alias for the drive and save the changes to the nvram:

---

Even though the device name may include **atf**, use **sd** in the device name at the OBP level.

---

```
nvalias fcadisk /sbus@1f,0/fcaw@1,0/sd@0,0
nvstore
```

- f. When booting in a SAN environment, you must set the World Wide Port Name (WWPN) by entering the commands and text indicated with **bold** text.

```
ok "/sbus@e/fcaw@1,0" select-dev
nport init complete
ok set-bootn0-wwn
Boot Node 0 WWN = 50 6 1 60 20 2 e0 b9
Boot Node 1 WWN = ff ff ff ff ff ff ff ff
Sure you want to change it? (y/n) y
Set Boot Node 0 WWN Addr
Byte 1 of 8 , enter 2 hex digits: 50
Byte 2 of 8 , enter 2 hex digits: 06
Byte 3 of 8 , enter 2 hex digits: 01
Byte 4 of 8 , enter 2 hex digits: 68
Byte 5 of 8 , enter 2 hex digits: 20
Byte 6 of 8 , enter 2 hex digits: 02
Byte 7 of 8 , enter 2 hex digits: e0
Byte 8 of 8 , enter 2 hex digits: b9
Boot Node 0 WWN = 50 6 1 68 20 2 e0 b9
Boot Node 1 WWN = ff ff ff ff ff ff ff ff
```

where **"/sbus@e/fcaw@1,0"** is the storage systems's boot path, and the WWPN entered is the SP's WWPN that the HBA uses at boot.

- g. Power cycle the host so that the alias is saved.  
h. Boot from the drive:

```
boot fcadisk
```

The normal boot process begins.

- Default Boot Device**
9. If you want the drive to be the permanent boot drive, make it the default boot device:
    - a. At the `ok` prompt type the following to see the environmental variables:  
**printenv**  
One of the variables is:  
`boot-device disk disk`
    - b. Change the **boot-device** value with the following command:  
**setenv boot-device fcdisk**

---

## Using a Storage-System LUN as a Boot Device (PCI)

This version of Open Boot supports FC-AL, Private loop, Public loop, and Fabric Point-to-Point from a single host. It also supports multi-initiators and multi-LUNs (0 through 255).

The following is a procedure for installing the Solaris operating system on a storage-system LUN (drive). When the installation procedure is complete, the host can boot and run from the LUN (drive).

---

Be sure to back up any critical files on the local boot disk.

---

---

### Prerequisites

Before starting this procedure make sure that:

- The host has a "local" boot disk (internal hard drive).
- The capacity of the intended (target) Fibre Channel disk is large enough to hold all the file systems and swap slice of the local current boot disk.
- You use persistent bindings (assigning a specific SCSI ID to a specific WWPN or DID) if the Fibre Channel adapter is connected to more than one logical disk drive.
- You correlate the intended boot disk displayed by the **probe-scsi-all** command to the devices displayed by the operating system.

## Boot Disk Configuration

You can use this procedure to make a bootable system disk from any size logical disk. This procedure also requires that a local boot disk be used to build the Fibre Channel boot disk.

1. At the `ok` prompt, boot the server's internal disk to setup the 'new' disk by entering the following command:

```
boot
```

2. If you do not see any Fibre Channel disks after booting, you may need to do one or both of the following
  - Enter the following, **touch /reconfigure**, and then reboot.
  - Edit the `/kernel/drv/lpfc.conf` file on the boot disk to change the topology from 4 to 2 or the reverse.
3. When Solaris is up, use the **format** command to do the following:

---

The following example assumes that you have selected logical drive designated as **c1t0d0** (controller 1, target 0, lun 0).

---

- Select the logical drive to become boot disk (for example **c1t0d0**).
  - Select the partition option and partition disk as desired.
  - Select the label option and write volume label to target disk.
4. Install boot on partition 0 of target disk with the following command:

```
installboot /usr/platform'uname-i'/lib/fs/ufs/bootblk  
/dev/rdisk/c1t0d0s0
```

5. To create a filesystem for each partition that will contain a mounted filesystem, enter the following commands:

```
newfs -v /dev/rdisk/c1t0d0s0 (will become root)
```

```
newfs -v /dev/rdisk/c1t0d0s6 (will become usr)
```

```
newfs -v /dev/rdisk/c1t0d0s7 (will become export/home)
```

6. To create temporary mount points for the new partitions so they can be mounted for copying, enter the following commands:

```
mkdir /root2
```

```
mkdir /usr2
```

```
mkdir /export2
```

7. Mount the new partitions to the appropriate directories.
8. Copy the applicable filesystems. An example command sequence for the root filesystem is

```
cd /root2
```

```
ufsdump 0f - /dev/dsk/c0t0d0s0 | ufsrestore rf -
```

9. Edit `/root2/etc/vfstab` to point to the new boot disk as the physical location of mount points for `root`, `swap`, `usr`, etc. Change the current entries as follows:

- `/dev/dsk/c0t0d0s0` and `/dev/rdisk/c0t0d0s0` to show `c1t0d0s0` (root)
- `/dev/dsk/c0t0d0s6` and `/dev/rdisk/c0t0d0s6` to show `c1t0d0s6` (usr)
- `/dev/dsk/c0t0d0s7` and `/dev/rdisk/c0t0s0s7` to show `c1t0d0s7` (export)
- Change entry for swap to show `c1t0d0s3`

10. If multiple targets exist on a Fibre Channel adapter, you must edit `/root2/kernel/drv/lpfc.conf` to set up a persistent binding that will ensure that the desired logical disk is always assigned the correct target ID. See instructions in the `lpfc.conf` file to set persistent binding. The disk is now ready to use as a boot disk.

11. Shut down the system.
12. At the `ok` prompt enter the following command:

```
reset -all
```

## Boot Setup

1. To display the open boot driver version, enter the following command sequence at the `ok` prompt:

```
select /lpfc
```

```
.version
```

```
unselect-dev
```

2. Change the alias for the current local boot disk from **disk** to another name (for example, **localdisk**) as follows:

- a. Enter the following command:

```
devalias
```

Record the long <path name> after the first mention of **disk**, and then use that pathname in the following command:

```
nvalias localdisk <path name>
```

This assigns the alias **localdisk** to the local boot disk.

3. Determine the device WWPN, DID or ALPA by entering the following command:

```
probe-scsi-all
```

4. Record the actual WWPN, DID or ALPA and the corresponding boot entry.
5. Select appropriate topology (point-to-point or FC-AL) for intended (target) Fibre Channel boot disk.
  - a. If point-to-point, enter the following commands:

```
select /xxxx where
```

*xxxx* is the full pathname to the boot device that you recorded from the **probe -scsi -all** command. For example:

```
/pci@2,2000/lpfc@2
```

```
/set-ptp
```

```
unselect-dev
```

```
reset -all
```

- b. If you are using a public or private loop, enter the following commands:

**select /xxxx** where

*xxxx* is the full pathname to the boot device that you recorded from the **probe -scsi -all** command. For example:

**/pci@2,2000/lpfc@2**

**set-fc-al**

**unselect-dev**

**reset -all**

6. Set boot device ID.

- a. To use WWPN, enter the following commands:

**select /xxxx** where

*xxxx* is the full pathname to the boot device that you recorded from the **probe -scsi -all** command. For example:

**/pci@2,2000/lpfc@2**

**wwpn <wwpn> <lun> <target id> set-boot-id**

**unselect-dev**

**reset -all**

where

*<wwpn>* : must be in hex

*<lun>* : default is in hex; to change to decimal type **d# <lun>**

*<target id>* : default is in hex, to change to decimal type **d# <target id>**

For example, a valid boot device ID using the WWPN is:

**22000020.37997f70 d# 140 d# 10 set-boot-id**

- b. To use DID \ (public loop), enter the following commands:

**select /xxxx** where

*xxxx* is the full pathname to the boot device that you recorded from the **probe -scsi -all** command. For example:

**/pci@2,2000/lpfc@2**

**did <did> <lun> <target id> set-boot-id**

**unselect-dev**

**reset -all**

where

*<did>* : must be in hex

*<lun>* : default is in hex; to change to decimal, type **d# <lun>**

*<target id>* : default is in hex; to change to decimal, type **d# <target id>**

For example, a valid boot device ID using the DID is:

**2009e1 d# 100 d# 10 set-boot-id**

- c. To use alpa \ (private loop), enter the following commands:

**select /xxxx** where

*xxxx* is the full pathname to the boot device that you recorded from the **probe -scsi -all** command. For example:

**/pci@2,2000/lpfc@2**

**did <alpa> <lun> <target id> set-boot-id**

**unselect-dev**

**reset -all**

where

*<alpa>* : must be in hex

*<lun>* : default is in hex; to change to decima, type **d# <lun>**

*<target id>* : default is in hex; to change to decimal, type **d# <target id>**

For example, a valid boot device ID using ALPA is:

**e1 d# 100 d# 10 set-boot-id**

7. To remove the boot ID you just assigned, enter the following commands:

**select /xxxx** where

xxxx is the full pathname to the boot device that you recorded from the **probe -scsi -all** command. For example

**/pci@2,2000/lpfc@2**

**remove-boot-id** (to clear boot ID settings.)

**unselect-dev**

**reset -all**

## Troubleshooting Suggestions

The following table lists symptoms and possible causes for problems you may experience when using a storage-system LUN as a boot device.

Symptom	Possible Cause
During system initialization, the boot sequence cannot mount or filecheck <b>/usr vfstab</b> on boot disk, or boot disk is not updated correctly.	Persistent binding not implemented or not implemented correctly.
After entering boot disk, a message displays stating that the file that was loaded is not executable.	No boot on target disk because copy from local disk to target disk was unsuccessful.
System hangs for a long time after displaying driver information, or System reboots right after displaying driver information.	The topology may be incorrect. Check the <b>/kernel/drv/lpfc.conf</b> file on the target disk.

## Helpful ok Prompt Commands

<b>select /lpfc</b>	Selects the lpfc device.
<b>.boot-id</b>	Prints current boot device ID.
<b>.version</b>	Print current version of open boot.
<b>.topology</b>	Prints current topology.
<b>.nvram</b>	Prints out current flags for open boot.
<b>unselect-dev</b>	Unselects the lpfc device.



---

This chapter describes how to use the Application Transparent Failover (ATF) software. The topics include:

- Testing ATF After Installation ..... 5-2
- Understanding Device Names ..... 5-3
- Adding Devices to a System That Has ATF Installed..... 5-4
- Failover and Other Messages ..... 5-6
- ATF Log Messages ..... 5-7
- Restoring the Original Path ..... 5-7

## Testing ATF After Installation

After installing ATF, you may want to test it. At least one LUN must be bound for the test to work.

1. Start I/O to one or more LUNs assigned to an SP (or both SPs).
2. In Navisphere Manager or Supervisor, click the **Storage** tab and expand the view of the SPs.
3. Start failover by either
  - disconnecting the cable to the I/O port of one SP from either the SP or the HBA (do this with Fibre Channel SPs only).
  - pulling an SP with active I/O about 1 inch (2.5 cm) from its enclosure (this is allowed with all SP types but is less convenient than disconnecting the cable). To do this, you will need to open the storage-system chassis, pull the SP, and close the chassis within a 2-minute period to avoid thermal shutdown. Removing an SP is explained in the installing and maintaining manual shipped with the storage system.

---

After you disconnect the cable(s) or pull out the SP, I/O may slow for a short time.

---

4. Wait a few moments, and then manually poll the storage system from Manager using the **Poll** button on the toolbar. The **Storage** tree will show that all the LUNs are under the operational SP.
5. After you are satisfied that failover worked correctly, reconnect the cables or push the SP fully into its slot.

After the SP becomes ready (indicated by the SP ready light, described in the storage-system installation manual), restore the LUNs to the original SP as shown in *Restoring the Original Path*, on page 5-7.

6. After a few moments, manually poll the storage system again. Verify that the affected LUNs are back under their original SP.

## Understanding Device Names

The Sun and ATF naming conventions differ. They use different names to refer to the same device. When you boot, the system logs messages that let you match each Sun and ATF name to the appropriate unit.

Sun software identifies each disk using the form: **/dev/dsk/cDtSdLsP**

where

- D* is the controller number assigned by the Solaris operating system.
- S* is the target SCSI ID or FC-AL address ID of the SP.
- L* is the ID of a LUN owned by the SP identified by *tS*. This ID can range from 0 through 7.
- P* is the partition (slice) number.

A sample name is **c0t0d1s1**.

ATF identifies each disk in its system log messages using the form: *[n1 / n2 / n3 / n4]*

where

- n1* identifies the adapter. (This identifier is unique to each adapter and may not be consistent across reboots of the system.)
- n2* identifies the SCSI bus or FC loop number (not related to Sun's **cD**).
- n3* identifies the target SCSI ID (same as Sun's **tS**).
- n4* identifies the LUN (same as Sun's **dL**).

A sample name is **0/ 0/ 0/1**. A failover can change any field except the LUN field, *n4*.

## Adding Devices to a System That Has ATF Installed

To add new storage systems, SPs, host bus adapter and driver, and/or LUNs to your configuration, follow these steps.

---

If this is a SCSI storage system, complete steps 1 and 2. If this is a Fibre Channel storage system, proceed to step 3.

---

1. Shut down the server system - both servers if there are two.
2. Shut down any storage systems connected to the server(s).
3. Install the new hardware: storage systems, SPs, adapter with adapter driver, and/or disk modules.
4. Power up the storage systems and verify that the storage-system disks are ready.
5. Power up each server and when you see the operating system banner, immediately enter **Stop-A** (hold the **Stop** key down, and press **A**).
6. At the prompt, boot with the reconfigure option:

**boot -r ↵**

---

The next step is to run the **atf\_configure** script. (The install program does this automatically at ATF installation, but you need to do it manually if you add devices to a system with ATF installed.) The **atf\_configure** script creates new **sd.conf** and **atf.conf** file. The disk-controller information is divided between the two files as follows.

---

- The SPs of all storage systems that are not in 7-slot (series 100) storage systems are defined in **atf.conf**
- The SPs in 7-slot storage systems, and controllers of all SCSI disks not in a storage system, are defined in **sd.conf**

7. Run the **atf\_configure** program:

```
# /usr/sbin/atf_configure ↵
```

The **atf\_configure** program creates new **sd.conf** and **atf.conf** files in directory, **/kernel/drv**. If these files already exist, **atf\_configure** asks whether it should save the current versions before creating the new files. Any saved files have the suffix **.old**. If you have any custom information in the original files (for example, you might have inserted lines that remove a LUN from ATF control), be sure to save the original files so that you can copy the custom information in them to the new ones.

Now you should make sure that **atf.conf** contains entries for the SPs in all storage systems that are not 7-slot systems. The SPs are summarized near the beginning of the file, so you can find them easily.

The following is a sample SP summary in an **atf.conf** file:

```
# The following SPs have been found:
# parent = /sbus@1f,0/fcaw@1,0
#       target = 0
# parent = /sbus@1f,0/fcaw@2,0
#       target = 2
# Total: 2
```

8. Reboot with the reconfigure option (on all connected servers):

```
# reboot -- -r ↵
```

You have added the new storage system SPs and/or adapters to your configuration.

## Failover and Other Messages

### ATF Automatic Failover

If a device fails and ATF automatically fails over to the other path, you can tell by any of these events:

- Applications continue running after an SP or adapter failure.
- The system console displays messages that resemble the following. These messages are also recorded in the file `/var/adm/messages`.

```

mo day yr hr min sec hostname unix WARNING: atf[1/0/0/0][Id:0x2000c77]:FAILOVER
STARTED

mo day yr hr min sec hostname unix WARNING: atf[1/0/0/0][Id:0x2000c77]: Second
path
    active:[0/0/2/0]

mo day yr hr min sec hostname unix WARNING: atf[1/0/0/0][Id:0x2000c77]:FAILOVER
SUCCESSFUL

mo day yr hr min sec hostname unix WARNING: atf[1/0/0/0]:To restore, execute
    /usr/sbin/atf_restore/dev/atf_sp:c2t0

```

The system displays a message for each LUN on the failed path as the system tries to send I/O to the device. Write down the last line, which will help you restore the original path later.

### Detecting Multiple Paths to LUN

ATF can detect and handle multiple path-to-LUN situations, such as when using cascaded switches. When ATF detects a multiple path, the following message displays:

```
WARNING!! MULTIPLE PATH TO CLARIION SP DETECTED
```

A multiple path is defined as seeing the same bound LUN via one of the storage system's SPs more than once.

## ATF Log Messages

ATF logs messages in the system log file `/var/adm/messages`. Messages in the system log file have different priorities, depending on the word that introduces the message text (after the date, time, server hostname, and origination point of the error). The introductory words are

**WARNING**: Warning messages indicate a serious conflict or error. You should investigate them.

**DEBUG**: Normal operation, this is informational and you can ignore it.

**INFO**: Informational messages often appear at system startup and provide useful information on the ATF configuration and driver revisions.

The message `Saved configuration is invalid` should appear for each LUN after you reconfigure the system. If such a message appears and you have not reconfigured the system, you should investigate the cause.

## Restoring the Original Path

If a device failed and ATF automatically switched the path, you can restore the original path as explained here. You can tell that ATF switched paths by any of these events:

- Applications continue running after an SP or adapter failure
- The system console displays a message similar to the following:

```
WARNING: atf[ 0/ 0/ /0 1]: FAILOVER STARTED
WARNING: atf[ 1/ 0/ 1/ 1]: Second path active:[0/0/0/1]
WARNING: atf[ 1/ 0/ 0/ 0]: FAILOVER SUCCESSFUL
WARNING: atf[ 1/ 0/ 1/ 1]: To restore, execute:
                        /usr/sbin/atf_restore /dev/atf_sp:c3t1
```

After the failed device has been replaced, you can restore the original path by either of the following procedures:

- Shut down the server whose path has failed; turn off power to the storage system, and then turn it on again; and reboot the server system; or
- Enter the following command:

```
/usr/sbin/atf_restore argument
```

where *argument* is the identifier shown in the system console message indicating that a failover occurred - above, **/dev/atf\_sp:c3t1**. The console message displayed at failover shows the entire command line.

If the **atf\_restore** command does not restore the transferred LUNs to their primary path, stop I/O to the LUNs, and then reissue the **atf\_restore** command.

If the communication path is interrupted by a cable failure, the HBA and SP may be left in unknown states causing restore to fail. If this occurs, reset the SP associated with the failed path by pulling out the SP and re-seating it. This will establish proper communication between the HBA and SP. Restore should then succeed when rerun.

When running in a multihost configuration, **atf\_restore** should be performed on all hosts. Restore operations can cause all LUNs associated with a storage processor to be “enabled” on that storage processor. In a configuration using multiple hosts connected to one storage system, restore operations must be performed on all systems to ensure that they will have valid paths to their respective LUNs. The **atf\_restore** operations should be performed on one host at a time. The operation should be allowed to complete on each host before issuing the command on another host.

---

**atf\_restore\_all** restores all connections.

---

**What Next?** If the server will run the FirstWatch failover software or the Veritas Cluster Server (VCS) software - Continue to Chapter 6, *Modifying VERITAS FirstWatch Scripts*, before continuing with the SAN installation.

---

This chapter describes how to modify VERITAS® FirstWatch® scripts so that FirstWatch works properly in the split-bus asymmetrical configuration and the split-bus symmetrical configuration.

If using the VERITAS Cluster Software (VCS), refer to the *VERITAS Cluster Server Installation Guide* that shipped with your VCS software.

The topics discussed in this chapter are:

- Modifying Firstwatch Scripts in a Split-Bus Asymmetrical Configuration ..... 6-2
- Modifying FirstWatch Scripts in a Split-Bus Symmetrical Configuration ..... 6-8

## Modifying Firstwatch Scripts in a Split-Bus Asymmetrical Configuration

The *split-bus asymmetrical* configuration consists of a *primary host* that is on-line and a *secondary host* that is a backup system. With this configuration, the secondary host *might* have *no* LUNs assigned to its SP until it takes over the services for the primary host. For Navisphere and FirstWatch to work together in this environment, you must modify the following FirstWatch scripts:

- FirstWatch startup script
- FirstWatch takeover script
- FirstWatch surrender script

---

Navisphere Manager must be restarted after each failover to ensure that the Manager can properly connect to the devices by the device name currently connected to the host. If Navisphere Manager is not restarted, then the secondary host takes over the services of the primary host. In this case, Manager reports the arrays as faulted and responds with messages that it cannot locate or manage the subsystem with the given device name.

---

### What Next?

Continue with the next section to modify FirstWatch scripts on the primary host.

---

### Modifying FirstWatch Scripts on the Primary Host

On the primary host, you must modify the following scripts as described in this section:

- the startup script, `/opt/VRTSfw/bin/startup.d/S05Clariion`
- the surrender script, `/opt/VRTSfw/bin/surrender.script`

## Modifying the S05Clariion Startup Script

To ensure that the LUNs have been properly trespassed prior to starting the Agent, modify the *startup script*, `/opt/VRTSfw/bin/startup.d/S05Clariion`, to stop the Navisphere Agent, if running, and then to restart the agent. To do this, add the lines shown in **bold lettering** in the following sample code after the execution of the `clariion_online` routine.

```

if [ "$1" = "start" ]
then
    clariion_online $CLARIION_PRIMARY

    ps -ef | grep -v grep | grep
"/opt/Navisphere/bin/naviagent -f
/etc/Navisphere/agent.config" >> /dev/null
if [ $? -eq 0 ]
then
    id=`ps -ef|grep -v grep| grep
"/opt/Navisphere/bin/naviagent
-f /etc/Navisphere/agent.config" | awk
'{print $2}'`
    kill -9 $id
    /opt/Navisphere/bin/naviagent -f
/etc/Navisphere/agent.config
else
    pkginfo | grep NAVIAGENT >> /dev/null
    if [ $? -eq 0 ]
then

    /opt/navisphere/bin/naviagent-f/etc/Navisphere/agen
t.config
    fi
fi
    exit 0
else
    $ECHO "Usage: $BNAME { start | stop }"
    exit 1
fi

```

**What Next?** Continue with the next section to modify the surrender script.

## Modifying the Surrender Script

Modify the `/opt/VRTSfwb/bin/surrender.script` to stop the Navisphere Agent when the primary host surrenders and goes IDLE. To do this, add the following lines shown in **bold lettering** to the end of the surrender script prior to the echoing of the `SUCCESSFUL EXIT` message.

```

ps -ef | grep -v grep | grep
"/opt/Navisphere/bin/naviagent -f
/etc/Navisphere/agent.config" >> /dev/null
  if [ $? -ne 0 ]
  then
    id=` ps -ef | grep -v grep | grep
"/opt/Navisphere/bin/naviagent -f
/etc/Navisphere/agent.config" | awk '{print $2}'`
    kill -9 $id
  fi

  $ECHO ''
  $ECHO
  =====
  $ECHO "FirstWatch $BNAME FOR $OSVER - SUCCESSFUL
EXIT"
  $ECHO ` $DT `
  $ECHO
  =====
  =====

  exit 0

```

**What Next?** Continue with the next section to prepare the FirstWatch scripts on the secondary host.

## Modifying the FirstWatch Scripts on the Secondary Host

FirstWatch determines when the secondary host takes over the LUNS from the primary host. Set-up the secondary host as follows.

1. Move the **S80agent** file by entering the following command:

```
mv /etc/rc2.d/S80agent /etc/rc2.d/s80agent
```

Move the **/etc/rc2.d/S80agent** file that starts Navisphere on boot so that Navisphere will not start at boot time.

2. Modify the *surrender script*, **/opt/VRTSfw/bin/surrender.script**, to kill the Navisphere agent when the secondary host surrenders its takeover. To do this, add the following lines shown in **bold lettering** to the end of the surrender script prior to echoing the **SUCCESSFUL EXIT** message.

```
ps -ef | grep -v grep | grep
"/opt/Navisphere/bin/naviagent -f
/etc/Navisphere/agent.config" >> /dev/null
if [ $? -ne 0 ]
then
    id=` ps -ef | grep -v grep | grep
"/opt/Navisphere/bin/naviagent
-f /etc/Navisphere/agent.config" | awk '{print
$2}'`
    kill -9 $id
fi
$ECHO ''
$ECHO
=====
$ECHO "FirstWatch $BNAME FOR $OSVER - SUCCESSFUL
EXIT"
$ECHO '$DT'
$ECHO
=====
=====
exit 0
```

3. To ensure that the LUNs have been properly trespassed prior to starting the Agent, modify the *takeover script*, `/opt/VRTSfw/bin/takeover.d/S05Clariion`, to stop the Navisphere Agent when the secondary host takes over the services of the primary host. To do this, add the lines shown in **bold lettering** in the following sample code after the execution of the `clariion_online` routine.

```

if [ "$1" = "start" ]
then
    clariion_online $CLARIION_TAKEOVER
    ps -ef | grep -v grep | grep
"/opt/Navisphere/bin/naviagent -f
/etc/Navisphere/agent.config" >> /dev/null
    if [ $? -eq 0 ]
    then
        id=`ps -ef|grep -v grep| grep
"/opt/Navisphere/bin/naviagent
-f /etc/Navisphere/agent.config" | awk
'{print $2}'`
        kill -9 $id
        /opt/Navisphere/bin/naviagent -f
/etc/Navisphere/agent.config
    else
        pkginfo | grep NAVIAGENT >> /dev/null
        if [ $? -eq 0 ]
        then
            /opt/Navisphere/bin/naviagent -f
/etc/Navisphere/agent.config
        fi
    fi
    exit 0
else
    $ECHO "Usage: $BNAME { start | stop }"
    exit 1
fi

```

4. To ensure that the LUNs have been properly trespassed prior to starting the Agent, modify the *startup script*, `/opt/VRTSfw/bin/startup.d/S05Clariion`, to stop the Navisphere agent, if running, and then restart the agent again. To do this, add the lines shown in **bold lettering** in the following sample code after the execution of the `clariion_online` routine.

```

if [ "$1" = "start" ]
then
    clariion_online $CLARIION_PRIMARY

    ps -ef | grep -v grep | grep
"/opt/Navisphere/bin/naviagent -f
/etc/Navisphere/agent.config" >> /dev/null
    if [ $? -eq 0 ]
    then
        id=` ps -ef|grep -v grep| grep
"/opt/Navisphere/bin/naviagent
-f /etc/Navisphere/agent.config" | awk
'{print $2}'`
        kill -9 $id
        /opt/Navisphere/bin/naviagent -f
/etc/Navisphere/agent.config
    else
        pkginfo | grep NAVIAGENT >> /dev/null
        if [ $? -eq 0 ]
        then
            /opt/Navisphere/bin/naviagent -f
/etc/Navisphere/agent.config
        fi
    fi

    exit 0
else
    $ECHO "Usage: $BNAME { start | stop }"
    exit 1
fi

```

### What Next?

You have modified the FirstWatch files and you can therefore run FirstWatch to handle host failover if you have more than one host connected to the storage system.

In addition, you can reassign ownership of LUNs as described in the next chapter, *Reassigning LUN Ownership*.

## Modifying FirstWatch Scripts in a Split-Bus Symmetrical Configuration

A *split-bus symmetrical* configuration consists of two hosts that are providing services and that are backup servers for each other. In this instance, the `/opt/VRTSfw/bin/startup.d/S05Clariion` startup script on both hosts must be modified as per the Asymmetrical Configuration. Navisphere starts automatically during the boot and the change to the S05Clariion script allows the agent to be restarted once all the devices have been trespassed back to the host when going online as a primary host.

When one host fails and the second host takes over (dual services), if the Navisphere Manager was connected to both agents, the following messages will be returned by the Navisphere Manager during a polling cycle:

### Polling Cycle Messages

Error	Description
Agent Not Running	Navisphere Manager was not able to access the <i>failed</i> host while the <i>takeover</i> host was going Dual-Services.
Unable to locate managed storage system with given device name	The agent running on the takeover host does not have the LUN configuration expected for the failed host. Therefore Navisphere Manager cannot manage the failed hosts.

These messages are reported until either the failed host is back online as the primary host or Navisphere Manager is restarted.

When the failed host is back on-line as the primary host, Navisphere may fail to reconnect to the LUNs. This is because the Navisphere Agent starts before the FirstWatch daemon has started and no LUNs have been trespassed to the host. To solve this, you must modify the **S05Clariion** startup script as described in the next section.

## Modifying the S05Clariion Startup Script

To ensure that the LUNs have been properly trespassed prior to starting the Agent, modify the **/opt/VRTSfw/bin/startup.d/S05Clariion** script to stop the Navisphere Agent, if running, and then restart the agent again. Add the lines shown in **bold lettering** in the following sample code after the execution of the `clariion_online` routine.

```
if [ "$1" = "start" ]
then
    clariion_online $CLARIION_PRIMARY

    ps -ef | grep -v grep | grep
"/opt/Navisphere/bin/naviagent -f
/etc/Navisphere/agent.config" >> /dev/null
if [ $? -eq 0 ]
then
    id=` ps -ef|grep -v grep| grep
"/opt/Navisphere/bin/naviagent
-f /etc/Navisphere/agent.config" | awk '{print
$2}'`
    kill -9 $id
/opt/Navisphere/bin/naviagent -f
/etc/Navisphere/agent.config
fi

    exit 0
else
    $ECHO "Usage: $BNAME { start | stop }"
    exit 1
fi
```

### What Next?

If you have more than one host connected to the storage system, you are ready to run FirstWatch.

In addition, you can reassign ownership of LUNs as described in the next chapter, *Reassigning LUN Ownership*.



## Reassigning LUN Ownership

This chapter describes the following:

- Methods for Reassigning LUN Ownership.....7-2
- Updating Disk Names After Reassigning LUNs.....7-6
- Examples — Manually Reassigning LUNs .....7-7

## Methods for Reassigning LUN Ownership

You can reassign LUN ownership from one SP to another SP by either

- manually using Navisphere Manager, or **trespass\_array**.
- automatically using the Application-Transparent Failover (ATF) software (optional for direct connect; required for switch connect).

When you use Navisphere Manager to reassign LUN ownership, you change the default SP owner of the LUN. Neither **trespass\_array** nor ATF change the default SP owner of the LUN. When you change ownership of LUNs using Navisphere Manager, the reassignment of the LUN to the other SP does not take effect until you power the storage system off and then on again. For this reason, we recommend that you use Navisphere Manager to reassign LUN ownership in non-failure situations only when

- You add a second SP and you want to assign LUNs to the new SP.
- You add LUNs and you want to balance your LUNs between two SPs.

When an SP or a component (cable or adapter) in the path to an SP fails, the process of reassigning LUNs from the failed SP to the working SP is called a failover. If ATF is running and a failure occurs, ATF will automatically execute a failover. If ATF is not installed, and the working SP owns at least one LUN, you can fail over LUNs to the working SP using **trespass\_array**. When you replace the failed SP or component, manually restore the LUN that failed over to their original SP using either the **atf\_restore\_all** command or another **trespass\_array** command.

---

## Failover Using Application Transparent Failover (ATF)

---

ATF is required if your server has two HBAs connected to a storage system.

ATF automatically reassigns LUNs when a failure occurs. This lets applications continue to run with minimal interruption after an FC-AL or fabric route fails. ATF is the easiest form of failover.

---

If your system is booted with a path removed, Navisphere ATF will fail over the affected LUNs and allow I/O to take place to these LUNs. However, when the path is restored, Navisphere ATF will not be able to restore these LUNs until the system is rebooted.

---

---

## Failover Using `trespass_array` and `rescan_array`

Using `/usr/bin/trespass_array` and `/usr/bin/rescan_array`, you can manually reassign LUNs. When using `trespass_array`, you must allocate an unshared LUN to each SP.

---

To install the `trespass_array` and `rescan_array` utilities, refer to either of the following manuals that ships with your HBA and HBA driver.

*Fibre Channel Sbus HBA and Driver for Solaris Installation Guide P/N 014003031*  
*Fibre Channel PCI HBA and Driver for Solaris Installation Guide P/N 014003032*

---

### `trespass_array`

The `trespass_array` utility does not require you to turn off storage-system power. Typically, you use `trespass_array` if a path to a LUN through an SP fails. For example, if an adapter is connected to two SPs, and one of the SPs fails, you can use `trespass_array` to transfer control of LUNs from the failed SP to the working SP.

For you to transfer control of LUNs to an SP using `trespass_array`, that SP must already own at least one LUN. The transferred LUNs belong to the new SP until you transfer them back, which you can do with another `trespass_array`. To create entries in `/dev` for the transferred LUNs, you must run the `rescan_array` utility (see the next section), or reboot the server using the `init 6` or `shutdown -y -i6 -g0` command.

The format for the **trespass\_array** command is as follows:

**trespass\_array** *device option*

where

*device* is the Solaris name of the disk that you want trespassed.

*option* is one of the following:

- A** reassigns all LUNs. This is the default.
- H** prints help information for the **trespass\_array** command.
- L *n*** reassigns LUN *n* only, where *n* is the LUN number.
- O** transfers control of all the original LUNs of this SP.

For example, to transfer control of only LUN 3 to device `c1t0d0s0`, you would issue the following command:

```
trespass_array c1t0d0s0 -L 3
```

### **rescan\_array**

Use the **rescan\_array** utility after **trespass\_array** to size the new devices without having to reboot the operating system.

The **rescan\_array** command has no arguments; the format is

**rescan\_array**

When you use **trespass\_array**, the physical addresses of the LUNs change to the FC-AL address ID of the SP that now owns the LUNs. You use **rescan\_array** to configure the new physical addresses.

The **rescan\_array** utility restarts the disk driver, which rescans for devices. Since LUNs are named in the order that they are found, disk names might change after running **trespass\_array** and **rescan\_array**.

**Direct or hub connection** — With Solaris 2.3 or higher, disk device filenames are derived from the controller number, FC-AL address ID, and LUN number of the device. Since a trespassed disk will have a different FC-AL address ID, the disk device filename will always be different after running **trespass\_array** and **rescan\_array**.

**Switch connection** — Disk device filenames are derived from the controller number, SP target ID, and LUN number of the device. Since a trespassed disk will have a different SP target ID, the disk device filename will always be different after running **trespass\_array** and **rescan\_array**.

---

## Failover Using FirstWatch or VERITAS Cluster Server (VCS) Software

FirstWatch and VCS are host failover utilities. In a split-bus dual-server configuration, if a host or path to a host fails, FirstWatch or VCS will fail the array over to the other host, and automatically reassign LUNs. They both work with any adapter but require a split-bus dual-server configuration. See the FirstWatch or the VCS documentation for installation information.

---

## Updating Disk Names After Reassigning LUNs

For Solaris, a change in LUN ownership using **trespass\_array** and **rescan\_array** affects disk names according to the type of SP failure that occurs; that is, if an SP fails or if you change SP ownership manually.

---

If ATF reassigns the LUNs, disk names do not change.

---

---

### Updating Disk Names if an SP Fails

If an SP fails, you cannot access the LUNs on the failed SP. At this point, disk names remain unchanged. However, the next time you boot the server, you must use the **boot -r** command, which changes the disk names to reflect the new target ID (the **tS** portion of the disk name). In a dual-adapter configuration, the disk names also reflect the new adapter (the **cD** portion of the disk name).

---

### Updating Disk Names if You Manually Change LUN Ownership

Manually changing LUN ownership affects disk names. You must power down and reboot the server after manually changing LUN ownership.

#### To Power Down and Reboot the Server

1. Shut down the server's operating system.
2. Power off and power on the storage system.
3. Shut down the server.
4. Reboot the server using the **boot -r** command.

Now the disk names reflect the new target (the **tS** portion of the disk name). In a dual-adapter configuration, the disk names reflect the new adapter (the **cD** portion of the disk name).

## Examples — Manually Reassigning LUNs

When you have more than one route from the server to a LUN, you can reassign LUN ownership from one SP to the other SP. This section describes a sample unshared direct single-server configuration, and a sample unshared direct dual-server configuration, and then uses these configurations to describe how you can either reassign the default SP owner using Navisphere Manager, or fail over the LUN using `trespass_array`.

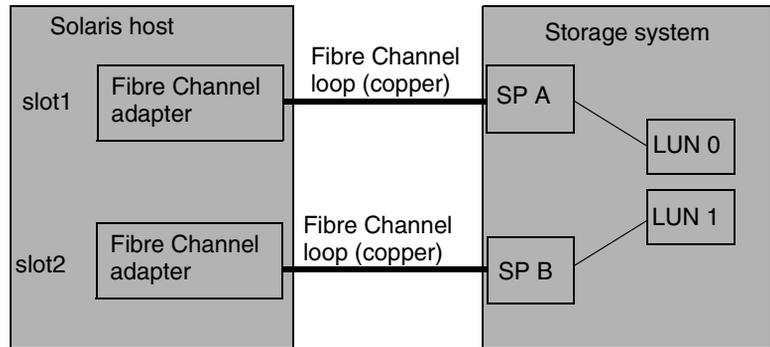
### Unshared Direct Single Server Configuration

This example starts with a Sun Solaris server and one storage system with 10 unbound disk modules. The disk modules were bound into LUNs and the LUNs were made available to Solaris.

The sample system is illustrated as follows.

Table 7-1 Unshared Direct Single-Server Configuration Before Assigning LUN 0

Adapter	SP and FC-AL address ID	LUNs	Solaris Name
Controller 1 in slot 1	SP A FC_AL address ID 6	LUN 0 RAID 5	<code>/dev/dsk/c1t6d0s&lt;0-7&gt;</code>
Controller 2 in slot 2	SP B FC_AL address ID 0	LUN 1 RAID 5	<code>/dev/dsk/c2t0d1s&lt;0-7&gt;</code>



**Figure 7-1 Unshared Direct Single-Server Configuration Before Assigning LUN 0**

The procedures that follow describe how to reassign LUN 0 from SP A to SP B. After reassigning the LUN, the configuration will appear as follows.

**Table 7-2 Unshared Direct Single-Server Configuration After Assigning LUN 0**

Adapter	SP and FC-AL address ID	LUNs	Solaris Name
Controller 1 in slot 1	SP A FC_AL address ID 6		
Controller 2 in slot 2	SP B FC_AL address ID 0	LUN 0 RAID 5	<code>/dev/dsk/c1t6d0s&lt;0-7&gt;</code>
		LUN 1 RAID 5	<code>/dev/dsk/c2t0d1s&lt;0-7&gt;</code>

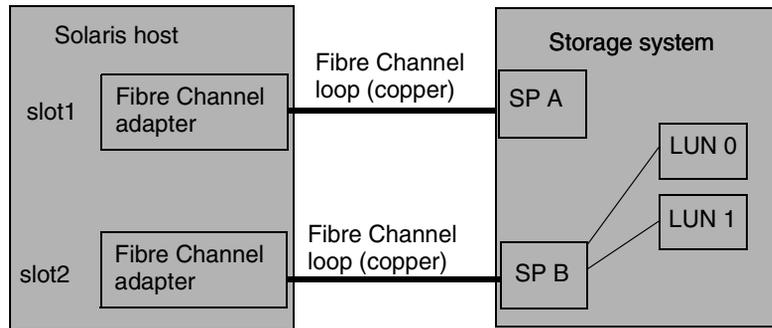


Figure 7-2 Unshared Direct Single-Server Configuration After Assigning LUN 0

## Reassigning LUN Ownership Using Navisphere Manager

If you have added LUNs, you might want to use Navisphere Manager to balance the load and reassign the default SP ownership for some LUNs. The new configuration will take effect only after you turn the power to the storage system off and on.

This procedure is not intended to temporarily reassign LUNs when a problem occurs. If a path or SP fails, use either ATF to automatically fail over LUNs or `trespass_array` to manually fail over LUNs.

This section describes how to reassign the default ownership of a LUN from one SP to another SP using the unshared direct configuration. The example reassigns LUN 0, `/dev/dsk/c1t6d0s0` mounted at `/mount1a`, from SP A to SP B.

1. Unmount the file system on the LUN to be reassigned.

To unmount the file system, use the `umount` command. For example, if the mount point is `/mount1a`, you would enter the command

```
umount /mount1a
```

If you cannot unmount the file system, run `fsck` on your file system.

2. Change the SP ownership as described in the Navisphere Manager manual.

3. Halt the server by entering  
**shutdown -y -i0 -g0**
4. Power the storage system off and on.
5. At the `ok` prompt, reboot the server by entering  
**boot -r**

Rebooting enables the operating system to recognize that a device exists.

6. If necessary, run **fsck** on your file system. You may not need to run **fsck** if you were able to unmount the file system in step 1.
7. Mount and use the file system on the new LUN.

In step 1, we used the example mount point **/mount1a**. Although the new LUN has a different name, mount the partition on **/mount1a** so that users can still use the pathnames to which they are accustomed.

LUN 0 is reassigned to SP B.

---

## Failing Over a LUN Using `Trespass_array`

This section describes how to fail over a LUN using **trespass\_array**. In this example, SP A fails and we reassign RAID 5 LUN **/dev/dsk/c1t6d0s0** mounted at **/mount1a** from SP A to SP B.

To install the **trespass\_array** and **rescan\_array** utilities, refer to either of the following manuals that ships with your HBA and HBA driver.

*Fibre Channel Sbus HBA and Driver for Solaris Installation Guide - 014003031*  
*Fibre Channel PCI HBA and Driver for Solaris Installation Guide - 014003032*

1. Unmount the file system on the LUN whose ownership you are reassigning.

To unmount the file system, use the **umount** command. For example, if the mount point is **/mount1a**, you would enter the command

```
umount /mount1a
```

2. Issue the **trespass\_array** command as follows to reassign all LUNs from SP A to SP B:

```
trespass_array c1t6d0s0 -A
```

Or issue the following command to reassign a specific LUN (in this instance LUN 0) to SP B:

```
trespass_array c1t6d0s0 -L 0
```

3. Issue the **rescan\_array** command as follows:

```
rescan_array
```

If you do not issue the **rescan\_array** command, you must reboot the server for Solaris to recognize the reassigned LUN.

The LUN that you just reassigned can now be remounted. For example:

```
mount /dev/dsk/c2t0d0s0 /mount1a
```

4. Run **fsck** on the file system if **rescan\_array** fails and displays a message to run **fsck**. This will occur if you could not unmount the file system in step 1.
5. Mount and use the file system on the new LUN.

In step 1, we used the example mount point **/mount1a**. Although the new LUN has a different name, use the **mount** command with the new FC-AL address ID, so that users can still use the pathnames to which they are accustomed.

LUN 0 is reassigned to SP B.

---

## Sample Unshared Direct Dual-Server Configuration

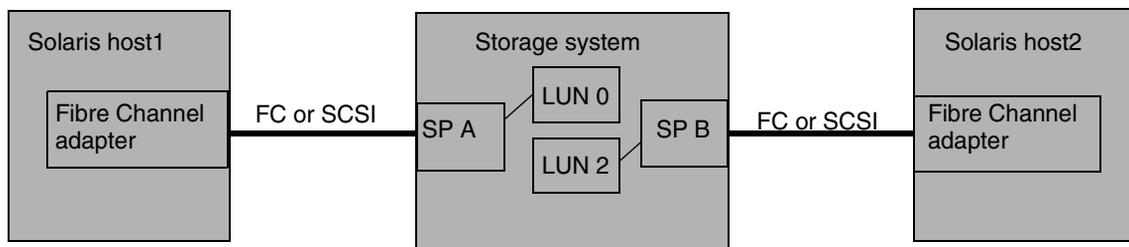
The unshared direct dual-server configuration has two servers, each with an adapter connected to an SP through a Fibre Channel cable. Each server independently uses its own LUNs in the storage system. For detailed configuration information, see the hardware reference or installation manual that ships with the storage system.

This example starts with two Sun Solaris servers and one storage system with 10 unbound disk modules that are bound into LUNs and are available to the operating system.

The sample system is illustrated as follows.

**Table 7-3 Unshared Direct Dual-Server Configuration Before Assigning LUN 0**

Servers	SP and FC-AL address ID	LUNs	Solaris Name
Solaris host 1	SP A FC_AL address ID 6	LUN 0 RAID 5	<code>/dev/dsk/c1t6d0s&lt;0-7&gt;</code>
Solaris host 2	SP B FC_AL address ID 0	LUN 1 RAID 5	<code>/dev/dsk/c2t0d1s&lt;0-7&gt;</code>

**Figure 7-3 Unshared Direct Dual-Server Configuration Before Assigning LUN 0**

The procedures that follow describe how to reassign LUN 0 from SP A to SP B. After reassigning LUN 0, the configuration will appear as follows.

**Table 7-4 Unshared Direct Dual-Server Configuration After Assigning LUN 0**

Servers	SP and FC-AL address ID	LUNs	Solaris Name
Solaris host 1	SP A FC_AL address ID 6		
Solaris host 2	SP B FC_AL address ID 0	LUN 0 RAID 5	<code>/dev/dsk/c1t6d0s&lt;0-7&gt;</code>
		LUN 1 RAID 5	<code>/dev/dsk/c2t0d1s&lt;0-7&gt;</code>

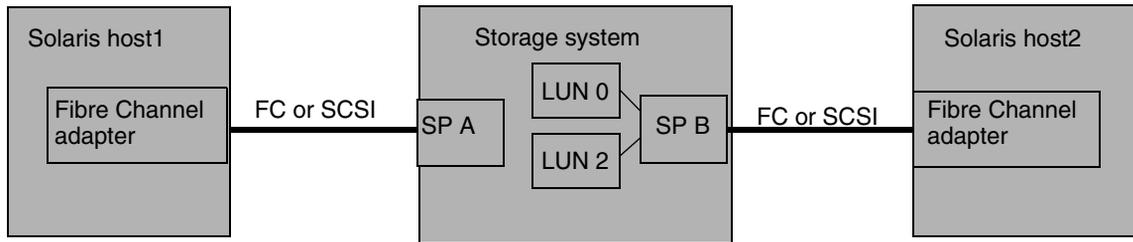


Figure 7-4 Unshared Direct Dual-Server Configuration After Assigning LUN 0

## Reassigning LUN Ownership Using Navisphere Manager

If you have added LUNs, you might want to use Navisphere Manager to balance the load and to reassign the default SP ownership for some LUNs. The new configuration will take effect only after you turn the power to the storage system off and on.

This procedure is not intended to temporarily reassign LUNs when a problem occurs. If a path or SP fails, use `trespass_array` to manually fail over LUNs.

From the sample unshared direct dual server configuration, we reassign the Solaris host1 RAID 5 LUN `/dev/dsk/c1t6d0s0` mounted at `/mount1a` from SP A to SP B, which is connected to server Solaris host2. Before being reassigned, the LUN is owned by SP A at FC-AL address ID 6, and after it is reassigned, it is owned by SP B at FC-AL address ID 0.

1. On Solaris host1, unmount the mount points for one LUN. In this example we transfer ownership of LUN 0, which has partition device entries `/dev/dsk/c1t6d0s0`. We unmount file systems on these partitions.

Unmount the file system using the `umount` command. For example, if the mount point is `/mount1a`, the command is

```
umount /mount1a ↵
```

If you cannot unmount the file system, run `fsck` on your file system.

2. Change the SP ownership as described in the Navisphere Manager manual.
3. Halt both servers by entering on each:  
**shutdown -y -i0 -g0**
4. Power the storage system off and on.
5. At the `ok` prompt, reboot each server by entering  
**boot -r**  
Rebooting enables the operating system to recognize that a device exists.
6. Mount and use the file systems on controller1 on Solaris host2.

In step 1, we used the example mount point **/mount1a**. Although the new LUN has a different name, mount the partition on **/mount1a** on the new server so that users can still use the pathnames to which they are accustomed.

---

## Failing Over a LUN Using `trespass_array`

This section describes how to reassign a LUN using `trespass_array`. From the sample unshared direct dual server configuration, we reassign the Solaris host1 RAID 5 LUN `/dev/dsk/c1t6d0s0` mounted at **/mount1a** from SP A to SP B, which is connected to server Solaris host2. Before being reassigned, the LUN is owned by SP A at FC-AL address ID 6, and after it is reassigned, it is owned by SP B at FC-AL address ID 0.

1. On Solaris host1, unmount the file systems on all LUNs that you are reassigning.

To unmount a file system, use the `umount` command. For example, if the mount point is **/mount1a**, you would enter the command

```
umount /mount1a
```

2. On Solaris host2, issue the `trespass_array` command as follows to reassign all LUNs from SP A to SP B:

```
trespass_array c1t6d0s0 -A
```

Or issue the following command to reassign a specific LUN (in this instance LUN 0) to SP B:

```
trespass_array c1t6d0s0 -L 0
```

3. On Solaris host2, issue the `rescan_array` command:

```
rescan_array
```

If you do not issue the `rescan_array` command, you must reboot the server for Solaris to recognize the reassigned LUN.

4. On Solaris host2, run `fsck` on the file systems if necessary.
5. On Solaris host2, mount and use the file systems.

In step 1, we used the example mount point **/mount1a**. Although the new LUN has a different name, mount the partition on **/mount1a** on the new server so that users can still use the pathnames to which they are accustomed.



---

This appendix describes how to troubleshoot problems with either the server or the storage system.

## Only LUN 0 Is Available to the Sun Host, or the Sun Host Cannot See New LUNs

For storage systems with SPs, edit the Sun `/kernel/drv/sd.conf` file as described in Chapter 3, *Making LUNs Available to Solaris*. The file has default entries for LUN 0 on each target. You must add entries for each LUN available under each target (SP FC-AL ID) for Solaris to access it.



### CAUTION

**The `sd.conf` file is very sensitive to hidden format characters. A misplaced tab or other formatting character can make all subsequent entries unreadable.**

## Cannot Find LUNs

If you cannot find the LUNs after a **boot -r**, run the **probe-scsi-all** command as a diagnostic tool. This command tells you if the hardware sees the LUNs. If it does not see them, then check the cables and arrays. If it does see them, then you have an operating system problem such as the driver or configuration file.

### To Test for LUNs

At the `ok>` prompt, enter the **probe-scsi-all** command with no arguments to list all of the LUNs that exist.

The term *disk* refers to a LUN.

## Configuring the Host Agent

This appendix describes how to manually reconfigure the Host **agent.config** file.

You need to edit the Host Agent configuration file only if you have non-FC4700 storage systems connected to the server.

## Configuring the Host Agent Using a Text Editor

Before Navisphere Agent can communicate with the storage system, you must edit the Agent's configuration file to include specific information. Before you are able to edit the Agent's configuration file or configure a managed storage system, you must have write access to the Agent. To secure write access you must be entered as a privileged user in the Host Agent's configuration file.

The Host Agent file does not require any device entries for FC4700 storage systems, and none of the other parameters defined in the file affect FC4700 storage systems.

This appendix describes how to edit the agent configuration file using a text editor, such as vi, emacs. or textedit.

---

Any user who can log in to a host that is a Navisphere management station can monitor the status of any of the managed storage systems.

---

---

### Host Agent Configuration File Pathname

The place-name to the Agent configuration file is `/etc/Navisphere/.naviagent.config.xxxxxx` <link> , where `xxxxxx` represents the time-date stamp.

The configuration file contains the following information for each storage system to be managed:

- Agent description and contact
- Communication channel(s) to the storage system(s)
- Users
- Polling frequency
- SP event log entries to transfer

The rest of this section describes how to specify the configuration information.

## Sample Agent Configuration File (agent.config)

```
# This file is used by Navisphere Agent to configure site specific values.
# All entries in this file have this form:      object value-list White space
# (tabs, space, line-feeds, carriage-returns) and blank lines are ignored;
# To continue a line onto multiple lines use "\".
#####
# Description specification
# The following is a list of objects which may be specified in this file, and
# the default values used by the Agent otherwise. Where object is one of
# {clarDescr, clarContact, device, user}, and value is a valid value for the
# given object.
clarDescr Navisphere Agent 4.2
clarContact John Smith, 800-555-1212
#
# CLARiiON specification.
# The CLARiiON device format is: device device-spec name "description"
# where the device-specification is the device name to open to access the
# clariion device. It can be a SCSI-2 file name or a tty name.This can NOT be
# defaulted. The name is the per clariion name that can be used to identify
# each CLARiiON to the manager application. It defaults to "" if not
# specified. The description is a quoted string that contains a brief
# description of the CLARiiON chassis being managed. It defaults to "" if not
# specified.
#
# SAMPLES:
device c1t0d0s2 NAVISPHERE-1 "NAVISPHERE-1"
device c2t1d1s2 NAVISPHERE-2 "NAVISPHERE-2"
#
# When an array has no bound LUNs, use device clspn,(n being the SP
# number) to specify the SP with unbound LUNs. This works if the CLSP driver
# has been loaded. Refer to the server set up manual for more information.
#
# User Specification.
# This is a list of users who are allowed to perform WRITE operations to the
# CLARiiON devices managed by this agent. By default, only root on this
# machine has WRITE access. Others can be granted access by including their
# login name and host which they can access from. A null list prevents anyone,
# including root, from having write access. In order to get write access, this
# file would have to be modified and the Agent signaled to re-read the
# configuration file.#
user root
user bcello@sunfish      # individual user
user afulton@bluejay    # individual user
#
# Poll rate
poll 60
#
# Event log entries to transfer
# Specifies the number of log entries to transfer. If "0", read the whole log.
# If a positive number, read that number of entries.
#
eventlog 100
# end of configuration file
```

---

## Specifying the Agent Description and Contact

The Agent description (**clarDescr**) and contact (**clarContact**) entries are optional informational text strings that Navisphere Manager does not use. You can use the **clarDescr** entry to include information about the Agent, such as its revision, and the **clarContact** entry to include information about the person to contact concerning any Agent questions or problems that might occur.

### To Specify the Agent Description

Either modify the **clarDescr** entry or add a new **clarDescr** entry.

A **clarDescr** entry has this format:

**clarDescr** [*description*]

where

*description* is an optional description of the Agent, such as the name of the Agent and its revision; for example, **Disk-Array Agent, revision *x***, where *x* is the revision of the Agent you are running.

### To Specify the Contact Description

Either modify the **clarContact** entry or add a new **clarContact** entry.

The **clarContact** entry has this format:

**clarContact** [*contact\_info*]

where

*contact\_info* is any optional information you want to include about whom to contact should any questions or problems with the Agent arise. You could include the name of a person to contact and the person's electronic mail address and phone number.

---

## Specifying the Communication Channel

Each SP in a storage system is connected to the server by a Fibre Channel arbitrated loop (FC-AL) or fabric connection to a Fibre Channel adapter, sometimes called a host bus adapter (HBA), or by a SCSI bus connected to a SCSI-2 adapter, also sometimes referred to as an HBA. The path from the adapter through the SP to one LUN in the storage system is a communication channel to the storage system for the Agent. The device entries in the agent configuration file define the communication channels. The fibre controllers use SCSI mapping.

If the storage system has two SPs that are both connected to host bus adapters in the same server, as they are in a dual-adapter/dual-SP configuration, the server's Agent configuration file must have a device entry (**device**) for the communication channel through each SP. If the SPs are connected to host bus adapters in different servers, the Agent configuration file in each server must have a device entry for the communication channel through the SP connected to that server.

If you add another storage system to the server in the future, you must make sure that the Agent configuration file has a device entry for the communication channel through each SP in the new storage system.

You need to know the FC-AL address IDs or SCSI target IDs of the storage system's SP(s) and of the host's HBAs to which the SP(s) are connected. The FC-AL address ID or the SCSI target ID of an SP is set by switches on the back of the storage system, as described in the storage-system installation and maintenance manual. In addition, you need to know the number for the LUN owned by each SP in the storage system.

---

## To Specify a Communication Channel

---

Do not add a communication channel for an FC4700 SP. Doing so will cause Manager to display an unnecessary unmanaged storage-system icon for the FC4700 SP.

---

Either modify the device entry or add a new device entry.

The device entry has this format:

**device** *channel name* [*"description"*]

where  
*channel*

is a required field for the communication channel (device path, device name, or disk name) that the Agent uses to communicate with the storage system. The format is *cDtSdLsP* where:

**For SCSI disks:**

- D is the controller number assigned by the Solaris operating system.
- S is the SCSI ID of the storage system's SP that is connected to *cD*. If the adapter identified by *cD* provides a narrow SCSI bus, this SCSI ID can range from 0 through 7; and if the adapter provides a wide SCSI bus, this SCSI ID can range from 0 through 15.
- L is the ID of a LUN owned by the SP identified by *tS*. This ID can range from 0 through 7.
- P is the partition (slice) number.

**For Fibre Channel disks:**

- D is the controller number assigned by the Solaris operating system.
- S is the FC-AL or target address ID (decimal) of the storage system's SP that is connected to fibre bus *cD*.
- L is the ID of a LUN owned by the SP identified by *tS*.
- P is the partition (slice) number.

- name* is a required description for identifying the communication channel. We recommend that this name include a name for the storage system, the SP, and the LUN. For example, if the name of the storage system is `array1`, the SP is `SP B`, and the LUN is `0`, you could use the name `array1_spb_lun0`. The Navisphere Manager does not use this information, but the Navisphere CLI does.
- `["description"]` is an optional description of the connection to the chassis. Be sure to enclose the description in quotation marks. For example, `"Path from host img01 to SP A to LUN0 in array1"`. The Navisphere Manager does not use this information, but the Navisphere CLI does.

To find out which devices Solaris sees, execute:

### format

Solaris returns output similar to the following.

```
Searching for disks...done
c3t2d0: configured with capacity of 32.58GB
AVAILABLE DISK SELECTIONS:
  0. c0t0d0 <SUN4.2G cyl 3880 alt 2 hd 16 sec 135>
    /pci@1f,4000/scsi@3/sd@0,0
  1. c3t2d0 <DGC-RAID3-0203 cyl 65437 alt 2 hd 4 sec
    261>/pci@4,4000/fibre-channel@2/sd@2,0
...
Specify disk (enter its number):
```

Any disk with DGC in its description can be specified in the Agent configuration file, as a communication channel to the storage system. Note that disk 0 does not have DGC in the description and cannot be a communication channel to the storage system.

---

If the list of available disks does not include all the disks with DGC descriptors (for example, you know that the system is configured with four SPs but **format** only lists two), verify that the `sd.conf` file contains the required target and LUN numbers; if not, add them, **reboot**, and rerun **format**. If the list still does not contain all the DGC disks, use the `clsp` driver to open a communication channel to those disks, and then bind one LUN on each SP.

---

To determine which disk(s) should be specified as a communication channel in the Agent configuration file, observe the device names of the DGC disks.

---

If the storage system has more than one LUN per SP connected to the same server, include one (*and only one*) **device** entry for each SP. Use a LUN that you do not expect to reconfigure.

---

In the sample there is only one SP. This could be either SP A or SP B. You must know the FC-AL address or target IDs. For any configuration but a SAN (shared storage that uses switches), the address IDs are manually set with a switch (hard coded) on each SP.

For this sample, we will assume that **c3t2d0** is SP A. Because there is only one SP, you will specify one communication channel in the Agent configuration file. You will also need to determine which partition (slice) number to specify in the device name. In **format**, enter the selection number for **c3t2d0**.

```
Specify disk (enter its number): 1 ↵
```

If the disk is not labeled you will see the following query; otherwise the FORMAT MENU appears.

```
selecting c3t2d0
[disk formatted]
Disk not labeled. Label it now? y ↵
```

The FORMAT prompt appears.

```
format> verify ↵
```

The current partition table appears. The sample shows the default partitioning that Solaris performed when it labeled the disk.

Current partition table (original):

Total disk cylinders available: 65437 + 2 (reserved cylinders)

Part	Tag	Flag	Cylinders	Size	Blocks
0	root	wm	0-251	28.46MB	252/0/0)263088
1	swap	wu	252-503	128.46MB	(252/0/0)263088
2	backup	wu	0-65436	32.58GB	(65437/0/0)68316228
3	unassign	wm	0	0	(0/0/0) 0
4	unassign	wm	0	0	(0/0/0) 0
5	unassign	wm	0	0	(0/0/0) 0
6	usr	wm	504-65436	32.32GB	(6493/0/0)6779052
7	unassign	wm	0	0	(0/0/0) 0

partition> **format**

Observe the partition table and determine a valid slice number to use. In the sample, valid slice numbers are, s0, s1, s2, s6. For this sample, we choose to use slice 2 (because it often exists). Therefore, the channel portion of the device name entry for the Agent configuration file would be **c3t2d0s2**.

If, in the future, you reformat your disk, eliminating or over-writing the slice (slice 2 in the sample), you will lose contact with the device.

## Communicating When a Storage System Has No Bound LUNs

If no bound LUNs exist, you must configure the Agent to communicate with each SP. To do this, you enter the **clsp** entries for the LUNs on each SP in the Agent configuration file.

If your server has CDE (CLARiiON Driver Extensions) software or ATF (Application-Transparent Failover) software, then CDE or ATF generate the **clsp** entries for the LUNs. The entries are listed in the format **clspn** in directory **/dev**.

If your server does not have CDE or ATF, then you must use the CLSP driver. (Do not use the CLSP driver with CDE or ATF.) If the CLSP driver has been loaded, **clspn** entries for each SP are listed in **/dev**.

After you determine which **clspn** entries to use, perform the following steps:

1. Edit the **agent.config** file. Add the required **clspn** device names and comment out any other device names.
2. Save the **agent.config** file.

3. Restart the Navisphere Agent (page 2-34).
4. Using Navisphere Manager, Supervisor, or CLI storage management software, bind one or more LUNs on that storage system.
5. Exit the Navisphere storage management software.
6. Stop the Navisphere Agent (page 2-34).
7. Reboot the storage-system server.
8. Execute **format** to see the device names.
9. Edit the **agent.config** file. Comment out the **clspn** entry, and add the new device entries (**cDtSdLsP**) for the LUN(s) you just bound.
10. Save the **agent.config** file.
11. Restart the Navisphere Agent.

Continue to the next section and resume editing the Agent configuration file.

---

## Specifying Who Can Configure a Managed Storage System

Any user who can log in to a host that is a Navisphere management station can monitor the status of any of the managed storage systems. However, a user *cannot* change a managed storage system's configuration without write access to the Agent. To give a user write access to the Agent, you must add to the host's Agent configuration file an entry for the user for each host.

Add the following entry to the `/etc/Navisphere/agent.config` file:

**user** *name*

where

*name* is the person's username. The format of this name differs depending on whether the person will be using Manager on a local or remote host.

**For a local host** - The format is *user*

where

*user* is the person's user account name.

**For a remote host** - The format is *user@hostname*

where

*user* is the person's user account name.

*hostname* is the name of the remote host.

For example, if you want to allow user **anne** to configure and monitor a host's storage system using the Navisphere Manager running on either remote host **img01** or remote host **img02**, you must add the following entries to the server's Agent configuration file:

**user anne@img01**

**user anne@img02**

---

To let a user configure a server's storage system(s) from the Navisphere Manager on different hosts, you must add to the Agent configuration file one **user** entry for each different host.

---

---

## Specifying a Polling Frequency for the Agent

The Agent does not automatically poll the SP. It polls only in response to a command from a client application such as Navisphere Manager. Once the poll request is received from the client application, the Agent determines when the last poll occurred and compares that to the polling frequency. If the polling frequency time has expired, the Agent will poll the hardware before returning status to the client application. Otherwise, the Agent will return status without performing a poll.

Add the following line to the `/etc/Navisphere/agent.config` file:

**poll *n***

where

*n* is 0 or an integer specifying the number of seconds between each poll of the Agent's SP. The minimum value, other than 0, is 10. We recommend that *n* be equal to or greater than 60 (the default value). If *n* is omitted, Agent polling is disabled and the SP is inaccessible. If *n* equals 0, the Agent will poll every time it gets a request.

---

## Specifying the Number of SP Events to Read

The Navisphere Agent reads event entries from the SP unsolicited event log at startup. You specify the number of the most current event to read. Since the complete log can contain as many as 16,800 entries, you may not want to burden Navisphere Agent with reading all of them.

---

If you enter a small value, Navisphere Manager and Navisphere CLI users will not be able to view the entire contents of the log without first stopping all applications, changing the value in the configuration file, saving the configuration file, restarting the Agent, and restarting the applications.

Add the following line to the `/etc/Navisphere/agent.config` file:

**eventlog *n***

where

$n$

is the number of log entries to transfer. For a serial connection;  $100 < n < 16,800$ . For a fibre connection;  $2048 < n < 16,800$ . If  $n$  is 0, the entire log is read. If  $n$  is a positive integer in the specified range, the specified number of entries are read.



---

This glossary contains terms related to disk storage subsystems. Many of these terms are used in this manual.

### A

**Agent** See Host Agent.

**arbitration** The process of selecting one respondent from a group requesting service at the same time.

**ATF** Application-Transparent Failover that provides applications with access to the LUNs via the alternate path if the path to the SP that owns the LUN fails.

### B

**bind** In the context of disk-array storage systems, the procedure by which you assign one or more disks to one LUN #. You can bind modules as one of the following RAID types: RAID 5 (independent access array), RAID 3 (parallel access array), RAID 1 (mirrored pair), RAID 0 (nonredundant array), RAID 1/0 (mirrored RAID 0), disk (individual disk), or hot spare. Before the LUN can store data, you must make it available to the operating system. Unbinding reverses the bind process, changing a LUN into its original disks or disk parts. Unbinding destroys all user information on the LUN. You bind (and unbind) disks using the Navisphere Manager, Supervisor, or CLI as described in the Navisphere Manager, Supervisor, or CLI manual.

**C**

**cache** See storage-system caching.

**CLI** A command line interface that lets you bind disks into LUNs on a non-RAID Group storage system, create RAID Groups and bind LUNs on them on a RAID Group storage system, unbind LUNs, destroy RAID Groups, set storage system properties, and ascertain storage system status. Navisphere CLI lets you do that for multiple storage systems on multiple servers running Navisphere Agent. For more information on Navisphere CLI, refer to the Navisphere CLI Reference manual.

**Core Software** Microcode that controls the operation of the storage system. You can update Core Software using Navisphere Manager or Supervisor.

**D**

**disk** A self-contained disk drive that slides into one slot in the front of the storage system. It consists of the carrier assembly, which holds the disk drive and the regulator board. Also called disk module.

**disk unit** A short name for LUN.

**F**

**FC-AL (Fibre Channel arbitrated loop)** An arrangement of Fibre Channel stations such that messages pass from one to the next in a ring.

**FC-AL address ID** A number that identifies a device as a node on a Fibre Channel loop. You select the FC-AL address ID for an SP using switches on the back of the storage system. The default FC-AL address ID for SP A is 0; the default FC-AL address ID for SP B is 1.

**fibre** General name for physical media supported by Fibre Channel, either optical or copper.

**Fibre Channel** A high-speed, serial, bidirectional, topology-independent, multi-protocol, highly scalable interconnection between computers, networks, and peripherals.

**Fibre Channel adapter (controller)** The name for the printed-circuit board within the computer chassis that allows the host to access the storage system through the Fibre Channel.

**H**

**host** See server.

**Host Agent** Software that runs on the server that is connected to the storage system and that communicates with the Navisphere client applications and with the Base and Core Software on the storage system.

**hot repair** See replace under power.

**hot spare** A disk module bound as a global hot spare that can replace any failed disk module in a RAID 1, RAID 1/0, RAID 3, or RAID 5 LUN.

**I**

**individual (disk) unit** A disk module bound as an individual unit, independent of any other disk modules in the cabinet. An individual unit has no inherent high-availability feature, but since the operating system supports software mirroring, you can make it highly available by software mirroring it with another individual unit, preferably on a different internal bus.

**L**

**LCC (link control card)** A CRU in an enclosure that connects Fibre Channel signalling to the disk modules. The LCC receives and electrically terminates the incoming FC-AL signal. It then passes the input signal to the disk modules in the enclosure, and drives the output signal via cables to the next device (if any) in the loop.

**logical unit (LUN)** A logical unit is one or more disks or parts of disks bound into a single entity, accessible by logical unit (LUN) number. Logical unit is a SCSI term. This manual generally uses the term LUN. The RAID types are: RAID 5 (independent access array), RAID 3 (parallel access array), RAID 1 (mirrored pair), RAID 0 (nonredundant array), RAID 1/0 (mirrored RAID 0), disk (individual disk), or hot spare. The operating system sees the LUN, which might include more than one disk, as one contiguous span of disk space.

**LUN** See logical unit (LUN).

**M**

**memory module** See SP memory modules.

**mirroring** Maintenance of a second (and optionally, via software, a third) copy of a logical volume image that provides continuous access if an image becomes inaccessible. The system and user applications continue running on the good image without interruption. There are two kinds of mirroring: hardware mirroring, in which the storage system maintains synchronization of the disk images, and software mirroring, in which the operating system maintains synchronization. Mirroring is further explained in the Fibre Channel configuration planning guide.

**Manager** A program with a graphical user interface that lets you bind disks into LUNs on a non-RAID Group storage system, create RAID Groups and bind LUNs on them on a RAID Group storage system, unbind LUNs, destroy RAID Groups, set storage system properties, and ascertain storage system status. Navisphere Manager lets you do that for multiple storage systems on multiple servers running Navisphere Agent. For more information on Navisphere Manager, refer to the Navisphere Manager manual.

**N**

**non-RAID Group storage system** A storage system with SPs running a revision of Core Software that does not support RAID Groups.

**P**

**path** See primary storage system route.

**primary storage-system route or primary route** In a storage system that has two SPs, each SP connects to each LUN through a route indicated by the device name. The primary route belongs to the SP that owns the LUN, and the secondary route belongs to the other SP. The SP on which a LUN is bound is the default owner of the unit. Knowing the routes is important because you may need to transfer control to the secondary route in the event of SP or Fibre Channel adapter failure.

**R**

- RAID (redundant array of independent disks)** A technology with its own set of definitions. See RAID 0, RAID 0, RAID 1/0, RAID 3, and RAID 5.
- RAID 0** Three or more disk modules bound as striped disks (the storage system reads and writes file information with more than one disk at a time). RAID 0 offers enhanced performance via simultaneous I/O to different modules, but does not intrinsically offer high availability. For high availability, you can software mirror the striped disks.
- RAID 1/0** Four, six, eight, ten, twelve, fourteen, or sixteen disk modules bound as mirrored RAID 0 group. These disk modules make up two mirror images of two, three, four, or five modules each. A RAID 1/0 group combines the speed advantage of RAID 0 with the redundancy advantage of mirroring.
- RAID 1** Two bound disk modules that the storage-system hardware will mirror.
- RAID 3** Five modules that use disk striping (as with RAID 5). The hardware maintains parity information that lets the group continue running and be rebuilt after a disk-module failure. With RAID 3, I/O occurs in smaller blocks than with RAID 5, and parity information is stored on one module, not distributed among all of them. RAID 3 works well for single-task applications that use I/Os of one or more 2-Kbyte blocks, aligned to start at disk addresses that are multiples of 2 Kbytes from the beginning of the logical disk.
- RAID 5** Three to sixteen disk modules that use disk striping, in which the hardware writes to or reads from multiple modules simultaneously, with high availability provided by parity information on each module. The ideal number of disk modules in a RAID 5 group is five.
- RAID Group** A group of disks on which you bind one or more LUNs of a specific RAID type. The RAID type is that of the first LUN bound on the group. Storage systems must run a specific revision of Core Software to support RAID Groups.
- RAID Group storage system** A storage system with SPs running a revision of Core Software that supports RAID Groups.

**replace under power** The storage system provides replace under power capability, allowing you to replace, for example, a disk or a fan module without powering down the storage system. Applications continue while you replace the failed module.

**route** See primary storage system route.

## S

**secondary storage-system route or secondary route** See primary storage system route.

**server** In the context of disk-array storage systems, a processor that runs an operating system and uses a disk-array storage system for data storage and retrieval.

**SP** See storage processor (SP).

**SP log (of storage-system events)** A time-ordered list of messages about storage-system events (such as errors) maintained by the SP. You can view this log using Navisphere Manager or Supervisor.

**SP memory modules** Memory modules (SIMMs or DIMMs) that provide the local storage for an SP. An SP must have at least two 4-Mbyte memory modules to support the storage-system cache.

**SPS** See standby power supply (SPS).

**standby power supply (SPS)** A unit that provides backup power for a Fibre Channel storage system with SPs in case of a power outage. An SPS is required for the storage-system write caching. If power fails, the SPS allows the SP to write the cache image to the cache vault area on disk. You can replace an SPS under power, without interrupting applications.

**Storage Group** A group of LUNs on a storage system with the Access Logix option. A storage system can have multiple Storage Groups. A Storage Group can be connected to multiple servers, but a server can be connected to only one Storage Group in each storage system.

**storage processor (SP)** A printed-circuit board with memory modules that control the disks in the storage-system chassis. The SP runs Core Software that controls the operation of the storage system. For higher availability, a site can use a second SP in a storage-system chassis.

**storage-system caching** For write caching, two SPs and an SPS are required. The SP stores (caches) information temporarily in its local memory, which can save time when data is written to and/or read from a LUN. Write caching is mirrored because each SP mirrors the cached information stored by the other SP.

**striping (disk striping)** The arrangement of a LUN such that reads and writes can occur with multiple disk modules simultaneously and independently. By allowing multiple sets of read/write heads to work on the same task at once, disk striping can enhance performance. You can implement hardware disk striping by configuring disk modules as a RAID 5 group, RAID 3 group, RAID 0, or RAID 1/0 group. The size of the stripe is the stripe element size (area on each disk) multiplied by the number of data disks in the group. For RAID 5, RAID 0, RAID 1/0 groups, you can select the stripe element size. Also see RAID.

## U

**unbind** See bind.



**A**

- adapters
  - adding 5-4
- adding devices 5-4
- agent configuration file
  - editing 2-23
  - user entry 2-22
  - verifying device entries 2-23
- agent, see Host Agent
- application-transparent failover (ATF)
  - defined 1-2
  - failover messages 5-6
  - installing
    - new revision 2-12
  - introduced 1-3
  - log file 5-7
  - reassigning LUNs after failure 7-3
  - using 7-3
- atf.conf configuration file 5-4
- atf\_configure utility 2-26

**B**

- balancing load by reassigning disks 7-2
- boot disk, installing ATF on, new revision 2-12

**C**

- clarContact entry in agent configuration file
  - on Solaris host B-4
- clarDescr entry in agent configuration file
  - on Solaris host B-4
- CLI
  - installing on a Solaris host 2-18

- removing
  - from Solaris host 2-37
- using
  - on Solaris host 2-37
- commands, failover
  - trespass\_array 7-3
  - rescan-array 7-3
- communication channel, specifying in agent configuration file
  - on Solaris host B-5
- configuration file 5-4
- configurations, storage system, with ATF 1-6
- C-series storage system, defined 1-2

**D**

- device entry in agent configuration file
  - on Solaris host B-5
- devices, adding 5-4
- disk names for
  - LUNs 3-2
- driver, CLSP B-9
- dual-adapter/dual-SP configuration, features 1-6
- dual-initiator/dual-bus configuration, features 1-7

**E**

- eventlog entry in agent configuration file
  - on Solaris host B-12

**F**

- failover
  - ATF 5-6
- failover using
  - application-transparent failover 7-3
  - FirstWatch 7-5
  - rescan\_array 7-4
  - trespass 7-3
  - VERITAS Cluster Server (VCS) 7-5
- FC-AL address ID B-5
- FC-series storage system
  - defined 1-2
- file systems, creating on LUN partitions 3-3
- file systems, mounting on LUN partitions 3-4
- FirstWatch 7-5
  - agent file, secondary host split-bus
    - asymmetrical configuration 6-5
  - failover using 7-5
  - modifying scripts 6-1
    - split-bus asymmetrical configuration 6-2
  - startup script
    - primary host, split-bus asymmetrical configuration 6-3
    - split-bus symmetrical configuration 6-9
  - surrender script
    - primary host, split-bus asymmetrical configuration 6-4
    - secondary host, split-bus asymmetrical configuration 6-5

**H**

- Host Agent
  - defined 1-2
  - installing on Solaris host 2-18
  - logging status
    - on Solaris host 2-35
  - removing
    - from Solaris host 2-37
  - starting
    - on Solaris host 2-34, B-13
  - stopping
    - on Solaris host 2-34, B-13

- Host Agent configuration file
  - on Solaris host
    - clarContact entry B-4
    - clarDescr entry B-4
    - device entry B-5
    - eventlog entry B-12
    - poll entry B-12
    - user entry B-11

**J**

- JBOD storage system, defined 1-2

**L**

- load balancing by reassigning disks 7-2
- login scripts, modifying 2-21
- LUNs (logical units)
  - creating file system on partitions 3-3
  - labeling and partitioning 3-2
  - mounting file systems 3-4
  - reassigning ownership using Manager 7-13
  - Solaris cannot find A-2
  - specifying disk names for 3-2
  - testing for existing LUNs A-2
  - verifying server can see 2-33

**M**

- managed Agent, defined 1-2
- managed storage system, defined 1-2

**N**

- non-RAID Group storage system, defined 1-2

**P**

- path
  - restoring 5-7
- persistent bindings
  - adding 2-28
  - cautions 2-25
  - defined 2-25
  - editing existing 2-30
  - modifying target IDs 2-28
- poll entry in agent configuration file, on Solaris host B-12

**R**

RAID Group storage system, defined 1-2  
RAID groups, testing for existing LUNs A-2  
restoring the original path 5-7

**S**

scripts, login, modifying 2-21  
sd.conf configuration file 5-4  
shared storage system, defined 1-2  
Solaris host  
    configuring the agent B-2  
    removing  
        Agent from 2-37  
        CLI from 2-37  
SP (storage processor), adding 5-4  
SP Agent, defined 1-2  
Specifying B-4  
starting Agent, on Solaris host 2-34, B-13

status, Agent, on Solaris host 2-35  
stopping Agent, on Solaris host 2-34, B-13  
storage systems, configuring 2-24

**T**

terminology 1-2  
testing for LUNs A-2  
trespass, reassigning LUNs 7-3  
trespass\_array command 7-3

**U**

unshared storage systems, defined 1-2  
user entry in agent configuration file  
    on Solaris host 2-22, B-11

**V**

VERITAS Cluster Server (VCS) using 7-5

