

SPARCcluster™ High Availability™ Server Software Planning and Installation Guide



The Network Is the Computer™

Sun Microsystems Computer Company

2550 Garcia Avenue
Mountain View, CA 94043 USA
415 960-1300 fax 415 969-9131

Part No.: 802-3509-10
Revision A, October 1995

© 1995 Sun Microsystems, Inc. 2550 Garcia Avenue, Mountain View, California 94043-1100 U.S.A.

All rights reserved. This product or document is protected by copyright and distributed under licenses restricting its use, copying, distribution and decompilation. No part of this product or document may be reproduced in any form by any means without prior written authorization of Sun and its licensors, if any.

Portions of this product may be derived from the UNIX® system and from the Berkeley 4.3 BSD system, licensed from the University of California. Third-party software, including font technology in this product, is protected by copyright and licensed from Sun's Suppliers.

RESTRICTED RIGHTS LEGEND: Use, duplication, or disclosure by the government is subject to restrictions as set forth in subparagraph (c)(1)(ii) of the Rights in Technical Data and Computer Software clause at DFARS 252.227-7013 and FAR 52.227-19.

The product described in this manual may be protected by one or more U.S. patents, foreign patents, or pending applications.

TRADEMARKS

Sun, Sun Microsystems, the Sun logo, Solaris, SunSoft, the SunSoft logo, Solaris, SunOS, Solstice, OpenWindows, DeskSet, SunFastEthernet, Solstice DiskSuite, Solstice AutoClient, ONC, ONC+, Online: Backup, JumpStart, and NFS are trademarks or registered trademarks of Sun Microsystems, Inc. in the United States and may be protected as trademarks in other countries. UNIX is a registered trademark in the United States and other countries, exclusively licensed through X/Open Company, Ltd. OPEN LOOK is a registered trademark of Novell, Inc. PostScript and Display PostScript are trademarks of Adobe Systems, Inc. ORACLE® and SQL*NET® are registered trademarks of Oracle Corporation. Prestoserve is a trademark of Legato Systems, Inc. All other product, service, or company names mentioned herein are claimed as trademarks and trade names by their respective companies.

All SPARC trademarks, including the SCD Compliant Logo, are trademarks or registered trademarks of SPARC International, Inc. in the United States and may be protected as trademarks in other countries. SPARCcenter, SPARCcluster, SPARCcompiler, SPARCdesign, SPARC811, SPARCengine, SPARCprinter, SPARCserver, SPARCstation, SPARCstorage, SPARCworks, microSPARC, microSPARC-II, and UltraSPARC are licensed exclusively to Sun Microsystems, Inc. Products bearing SPARC trademarks are based upon an architecture developed by Sun Microsystems, Inc.

The OPEN LOOK™ and Sun™ Graphical User Interfaces were developed by Sun Microsystems, Inc. for its users and licensees. Sun acknowledges the pioneering efforts of Xerox in researching and developing the concept of visual or graphical user interfaces for the computer industry. Sun holds a non-exclusive license from Xerox to the Xerox Graphical User Interface, which license also covers Sun's licensees who implement OPEN LOOK GUI's and otherwise comply with Sun's written license agreements.

X Window System is a trademark of X Consortium, Inc.

THIS PUBLICATION IS PROVIDED "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE, OR NON-INFRINGEMENT.

THIS PUBLICATION COULD INCLUDE TECHNICAL INACCURACIES OR TYPOGRAPHICAL ERRORS. CHANGES ARE PERIODICALLY ADDED TO THE INFORMATION HEREIN, THESE CHANGES WILL BE INCORPORATED IN NEW EDITIONS OF THE PUBLICATION. SUN MICROSYSTEMS, INC. MAY MAKE IMPROVEMENTS AND/OR CHANGES IN THE PRODUCT(S) AND/OR THE PROGRAM(S) DESCRIBED IN THIS PUBLICATION AT ANY TIME.



Contents

Preface	xv
1. Introduction	1-1
1.1 Features Provided.....	1-1
1.2 Data Protection Provided	1-2
1.3 Supported Hardware	1-2
1.4 Getting Help	1-3
1.5 Glossary.....	1-3
2. Planning and Installation Overview	2-1
2.1 Planning Overview	2-1
2.2 Functional Diagram	2-2
2.3 Installation Procedure Overview	2-6
3. Installation Planning	3-1
3.1 Planning Your Configuration	3-1
3.1.1 Network Configuration	3-3
3.1.2 Physical and Logical Host Names	3-8

3.1.3	Data Services	3-9
3.1.4	Planning Your Disk Requirements	3-9
3.1.5	Mirroring	3-10
3.1.6	Trans Devices.	3-10
3.1.7	Hot Spares	3-11
3.1.8	Disk Space Growth	3-11
3.1.9	Total Storage Requirement.	3-12
3.1.10	Size and Number of Disk Drives.	3-12
3.1.11	Number of SPARCstorage Arrays.	3-13
3.1.12	Size of Servers	3-13
3.1.13	Power Sources and Location for Configurations . .	3-15
3.1.14	Type of Configuration	3-17
3.1.15	Data on Local Disks	3-17
3.1.16	Planning Your Multi-host File System Hierarchy . .	3-18
3.1.17	File System Size and Disk Layout.	3-19
3.1.18	Size of Disksets	3-19
3.1.19	Metadevice Naming and Creation	3-21
3.1.20	Migrating Existing Data	3-23
3.1.21	Backing Up Multi-host Data	3-23
3.2	Configuration Rules That Improve Reliability	3-23
3.3	Configuration Restrictions	3-26
4.	Hardware Installation Information	4-1
4.1	SPARCcluster High Availability Hardware Installation .	4-1
4.2	SPARCstation 4 Installation	4-2

4.3	SPARCcenter 2000 Installation	4-2
4.4	SPARCserver 1000 Installation	4-3
4.5	SPARCstorage Arrays Installation	4-3
4.6	Terminal Concentrator Installation	4-4
5.	Solaris 2.4 Installation	5-1
5.1	Installation Rules	5-1
5.2	Selecting the Size of the Root Slice	5-2
5.3	Sizing File Systems for Install	5-6
	▼ How to Install Solaris on the Administration Workstation	5-7
	▼ How to Install the Terminal Concentrator	5-8
	▼ How to Install Solaris on the Solstice HA Servers	5-10
6.	Installing the Packages and Patches	6-1
6.1	Packages to Be Installed	6-1
	▼ How to Add the Required Packages and Patches	6-2
6.2	Path Requirements	6-30
6.3	Viewing AnswerBook Online Documentation	6-31
7.	Setting Up Name Services and Networks	7-1
7.1	Overview of Tasks	7-1
7.2	Updates To Name Service Entries	7-3
7.3	Automatic Changes to /etc/nsswitch.conf	7-3
7.4	Updates to /etc Files	7-5
7.5	Setting Up Private Networks	7-7
7.6	Setting Up / .rhosts Files	7-8

7.7	Editing <code>/etc/syslog.conf</code>	7-8
8.	Creating the Configuration	8-1
8.1	Overview of Tasks	8-1
8.2	Preparing to Run <code>hasetup</code>	8-2
▼	How to Run <code>hasetup</code>	8-3
9.	Metadevice Creation on Local and Multi-host Disks	9-1
9.1	Overview of Tasks	9-1
9.2	Metadevice Considerations	9-2
9.2.1	Diversity of the Configuration	9-2
9.2.2	Availability of the Configuration	9-3
9.2.3	Performance Issues	9-5
9.2.4	Takeover and Metadevice State Database Issues ..	9-6
9.2.5	Managing the Metadevice Name Space	9-6
9.3	Creation of Metadevices on Multi-host Disks	9-12
▼	How to Create Metadevices Using <code>metatool</code>	9-12
▼	How to Create Metadevices Using the <code>md.tab</code> File	9-13
9.4	Logging File Systems on the Local Disk	9-13
9.5	Setting Up Slices on Multi-host Disks	9-14
9.6	Setting Up Metadevices for HA-ORACLE	9-14
9.7	Setting Up Metadevices for HA-NFS	9-14
9.8	Setting Up Hot Spares	9-14
10.	Setting Up Data Service Classes	10-1
10.1	Overview of Tasks	10-1
10.2	Setting Up HA-ORACLE	10-2

▼	How to Prepare Solstice HA Servers for ORACLE Installation.....	10-2
10.2.1	Creating an ORACLE Database and Setting Up HA-ORACLE.....	10-5
▼	How to Prepare the Logical Hosts for ORACLE Databases.....	10-5
▼	How to Create an ORACLE Database and Set Up HA-ORACLE.....	10-6
10.2.2	HA-ORACLE Client Set Up.....	10-10
10.3	Setting Up HA-NFS.....	10-10
▼	How to Share HA-NFS File Systems.....	10-11
▼	How to Set Up a Multi-host UFS File System.....	10-13
10.4	Bringing the Solstice HA Servers On Line.....	10-16
11.	Verification and Validation.....	11-1
11.1	Overview of Tasks.....	11-1
11.2	Running the <code>hacheck</code> Command.....	11-2
11.3	Running the <code>haswitch</code> Command.....	11-2
11.4	Performing Physical and Manual Tests.....	11-3
11.5	Performing the Fault Detection Test.....	11-4
11.6	Verify HA-ORACLE Setup.....	11-4
11.7	Verify SPARCstorage Array Firmware.....	11-4
11.8	Power Sequence Testing a SPARCcluster 1000.....	11-5
11.9	Power Sequence Testing a SPARCcluster 2000.....	11-6
A.	Error Messages.....	A-1
A.1	<code>hacheck</code> Error Messages.....	A-1
A.2	<code>hasetup</code> Error Messages.....	A-9

B. Terminal Concentrator Security	B-1
B.1 Security Concepts	B-1
B.2 Using Security	B-2
B.2.1 Enabling Security	B-3
B.2.2 Setting Up Host-based Security	B-3
B.2.3 Local Virtual CLI Password	B-5
B.2.4 Concentrator Administrative Password	B-6
B.3 Setting Port Security	B-7
B.4 Port Server and CLI Security	B-8
B.5 Security on Virtual CLI Connections	B-9
B.6 PPP Security	B-10
B.7 Host-based Security	B-10
B.7.1 Defining a Security Server	B-11
B.7.2 Creating the <code>acp_dialup</code> File	B-12
B.7.3 Creating User Password Files	B-13
B.7.4 Setting Up Connection Security	B-14
B.7.5 Creating the ACP Encryption Keys File	B-16
B.7.6 Logging Security Events	B-18
B.7.7 Modifying the Supplied Security Application	B-18
B.8 Securing the Network	B-21
B.8.1 The Concentrator Administrative Password	B-22
B.8.2 Protecting Ports from Unauthorized Access	B-22
B.8.3 Protecting the Superuser CLI	B-22
B.8.4 Preventing Unauthorized Access to <code>na</code>	B-23

Index	I-1
-------------	-----

Figures

Figure 2-1	Basic Hardware Layout of a SPARCcluster HA Configuration	2-3
Figure 2-2	Eample Diskset Layout With Three SPARCstorage Arrays	2-5
Figure 2-3	Solstice HA Installation Procedure	2-6
Figure 2-4	Solstice HA Installation Procedure (continued)	2-7
Figure 3-1	Interface Names of Logical Network Interfaces	3-5
Figure 3-2	Example Diskset Allocation	3-21
Figure 3-3	Solstice HA Systems With Multi-host Disks	3-25
Figure 9-1	Example Trans Metadevice Naming Convention	9-8
Figure 9-2	Example Raw Mirrored Metadevice Naming Convention	9-9

Tables

Table 1-1	Minimum and Maximum Number of SPARCstorage Arrays	1-3
Table 3-1	Primary Network Naming	3-6
Table 3-2	Secondary Network Naming	3-6
Table 3-3	Example Private Network Naming	3-7
Table 3-4	Physical Limits of SPARCcluster Components	3-14
Table 3-5	Determining Drives Needed for a Configuration	3-20
Table 3-6	Division of Disksets	3-20
Table 4-1	SPARCcluster High Availability Hardware Documentation	4-1
Table 4-2	SPARCstation 4 Documentation	4-2
Table 4-3	SPARCcenter 2000 Documentation	4-2
Table 4-4	SPARCserver 1000 Documentation	4-3
Table 4-5	SPARCstorage Array Documentation	4-3
Table 4-6	Terminal Concentrator Documentation	4-4
Table 5-1	Root Slice Calculation Template	5-3
Table 5-2	Example Root Slice Calculation Template	5-4

Table 5-3	Example of a Root Slice for Maximum Solstice HA Configuration.	5-4
Table 5-4	Required Root File System Sizes and <code>md.conf</code> Values.	5-5
Table 5-5	Sizing File Systems for Install	5-7
Table 8-1	Solstice HA Configuration Worksheet.	8-2
Table B-1	Entries in the <code>acp_dialup</code> File	B-12
Table B-2	Arguments in the <code>acp_restrict</code> File Entries	B-14

Preface

Solstice™ High Availability™ 1.0 (Solstice HA) is a software product that supports specific dual-server hardware configurations. It is compatible with the Solaris™ 2.4 software environment. When configured properly, the hardware and software together provide highly available data services. Solstice HA depends upon the mirroring and diskset capabilities and other functionality provided by Solstice DiskSuite™ 4.0, which is shipped with Solstice High Availability.

The *SPARCcluster HA Server Software Planning and Installation Guide* documents the procedures for setting up hardware and installing and configuring the software. This manual is intended to be used with other hardware and software manuals which are listed under “Related Documentation” on page xvii.

Who Should Use This Book

The information in this manual is for Sun representatives who are performing the initial installation of Solstice High Availability configurations. The instructions and discussions in this manual are complex and intended for a technically advanced audience.

The instructions in this manual and the *SPARCcluster High Availability Software Planning and Installation Guide* assume the reader has expertise with the Solstice DiskSuite product.

System administrators with at least six years of UNIX® system experience will find this book useful when learning to administer Solstice High Availability configurations. This manual should be used in conjunction with the *SPARCcluster High Availability Software Administration Guide*.

Note – Junior and less experienced system administrators should not attempt to install, configure, or administer Solstice HA configurations.

How This Book Is Organized

This document has 11 chapters and two appendixes, as follows:

Chapter 1, “Introduction,” provides a high-level overview of Solstice HA and a glossary of the terminology associated with the product.

Chapter 2, “Planning and Installation Overview,” presents an overview of the planning and installation process and explains the supported Solstice HA configurations.

Chapter 3, “Installation Planning,” discusses how to plan for the installation of Solstice HA systems at your site.

Chapter 4, “Hardware Installation Information,” provides pointers to the appropriate hardware installation documentation.

Chapter 5, “Solaris 2.4 Installation,” gives the supplementary instructions necessary for installing the Solaris 2.4 operating system on the local disks of both servers in the Solstice HA configuration.

Chapter 6, “Installing the Packages and Patches,” explains the steps necessary to install the unbundled software packages and patches associated with the Solstice HA distribution.

Chapter 7, “Setting Up Name Services and Networks,” presents an overview of the steps you must follow to configure the name services and networks for the Solstice HA systems.

Chapter 8, “Creating the Configuration,” provides step-by-step instructions that must be followed when creating Solstice HA configurations.

Chapter 9, “Metadevice Creation on Local and Multi-host Disks,” gives directions for creating concatenations, stripes, UNIX file system logs, and mirrors on multi-host disks. It also describes steps for creating UFS logs on local disks.

Chapter 10, “Setting Up Data Service Classes,” describes the steps necessary to set up HA-NFS and HA-ORACLE.

Chapter 11, “Verification and Validation,” tells the steps necessary to ensure the Solstice HA configuration you created provides the desired high availability of data services.

Appendix A, “Error Messages,” explains error messages displayed by `hasetup(1M)` and `hacheck(1M)`.

Appendix B, “Terminal Concentrator Security,” gives directions for setting up security on the terminal concentrator.

Related Documentation

Use the following manuals with the *SPARCcluster HA Server Software Planning and Installation Guide*:

- *SPARC: Installing Solaris Software*
- *SPARCcluster High Availability Server Hardware Planning and Installation* (part number 802-3510)
- *SPARCcluster High Availability Software Administration Guide* (part number 802-3511)
- *SPARCcluster High Availability Server Service Manual* (part number 802-3512)
- *Solstice DiskSuite 4.0 Administration Guide*
- *Solstice DiskSuite Tool 4.0 User’s Guide*
- *Solaris 2.4 Administering File Systems*
- *NFS Administration Guide*
- *SMCC NFS Server Performance and Tuning Guide* (part number 801-7289)
- *TCP/IP Network Administration Guide*
- *Name Services Administration Guide*

-
- *Name Services Configuration Guide*
 - *SPARCserver 1000 System Installation Manual* (part number 801-2893)
 - *SPARCserver 1000 System Service Manual* (part number 801-2895)
 - *SPARCcenter 2000 Installation Manual* (part number 801-6975)
 - *SPARCcenter 2000 Service Manual* (part number 801-2007)
 - *SPARCstorage Array Model 100 Installation Manual* (part number 801-2205)
 - *SPARCstorage Array Model 100 Service Manual* (part number 801-2206)
 - *Multi-Disk Pack Installation and Service Manual* (part number 801-6119)
 - *Solaris 2.x Handbook for SMCC Peripherals* (part number 801-5488)
 - *FSBE/S SBus Card Manual* (part number 800-7508)
 - *Installing SBus Cards in Deskside and Data Cabinet Systems* (part number 800-6366)
 - *SunFast Ethernet Adapter User Guide* (part number 801-6109)
 - *Fibre Channel SBus Card Installation Manual* (part number 801-6313)
 - *Fibre Channel Optical Module Installation Manual* (part number 801-6326)
 - *Oracle7[®] Server Administrator's Guide* (part number 6694-70-1292)
 - *Oracle7 for Sun SPARC Solaris 2.x Installation and Configuration Guide*
 - *Terminal Concentrator Binder Set*

What Typographic Changes Mean

The following table describes the typographic changes used in this book.

Table P-1 Typographic Conventions

Typeface or Symbol	Meaning	Example
AaBbCc123	The names of commands, files, and directories; on-screen computer output	Edit your <code>.login</code> file. Use <code>ls -a</code> to list all files. machine_name% You have mail.
AaBbCc123	What you type, contrasted with on-screen computer output	machine_name% su Password:
<i>AaBbCc123</i>	Command-line placeholder: replace with a real name or value	To delete a file, type <code>rm filename</code> .
<i>AaBbCc123</i>	Book titles, new words or terms, or words to be emphasized	Read Chapter 6 in <i>User's Guide</i> . These are called <i>class</i> options. You <i>must</i> be root to do this.

Shell Prompts in Command Examples

The following table shows the default system prompt and superuser prompt for the C shell, Bourne shell, and Korn shell.

Table P-2 Shell Prompts

Shell	Prompt
C shell prompt	machine_name%
C shell superuser prompt	machine_name#
Bourne shell and Korn shell prompt	\$
Bourne shell and Korn shell superuser prompt	#

Introduction



Solstice High Availability 1.0 (Solstice HA) is a software product that provides fault-tolerant support and automatic data service failover for specific dual-server hardware configurations. The configurations will recover from server, disk, and network failures, as well as software failures.

Solstice HA uses Solstice DiskSuite 4.0 software to provide the diskset, mirroring, concatenation, striping, hot spare disks, file system growing, and UNIX file system logging capabilities.

Solstice HA allows configurations to be either symmetric or asymmetric. In the asymmetric configuration, one of the systems acts as a hot standby for the other system. In symmetric configurations, both servers can be actively offering data services.

The data services provided by Solstice High Availability are:

- Highly available NFS[®], Sun's distributed computing file system (Solstice High Availability/HA-NFS)
- Highly available ORACLE[®] database management system (Solstice High Availability/HA-ORACLE)

1.1 Features Provided

The key features of the Solstice HA software package include:

- Tolerance of single-point software failures or crashes

- Tolerance of single-point hardware faults
- Appearance of continuous availability of data service (NFS clients will only see “Server Not Responding” messages during a takeover)
- Automatic detection of system and data service failure
- Automatic takeover, recovery, and service restoration
- Automatic post-takeover redirection of HA-NFS clients
- Automatic restart of fault monitoring on a server after it is repaired

Additionally, Solstice HA provides online serviceability which enables administrators to take an appropriately configured server offline for repair or routine maintenance while the data services remain available from the other server in the configuration.

1.2 Data Protection Provided

The disks holding the data exported by the data services are multi-host and mirrored. This allows Solstice HA configurations to tolerate the following types of single-point failures:

- Server operating system failure because of either a crash or a panic
- Data service application failure
- Server hardware failure
- Network interface failure
- Disk media failure

1.3 Supported Hardware

Several hardware configurations are supported by Solstice HA. The supported servers are the SPARCserver™ 1000 and SPARCcenter™ 2000. These two types of servers are configured with SPARCstorage™ Arrays Model 100 Series.

Each SPARCstorage Array can contain up to 30 SCSI disks. Both 1.05-Gbyte and 2.1-Gbyte disks can be used in SPARCstorage Arrays. Table 1-1 shows the minimum and maximum number of SPARCstorage Arrays supported in the SPARCserver 1000 and SPARCcenter 2000 configurations.

Table 1-1 Minimum and Maximum Number of SPARCstorage Arrays

Servers Used (2 each)	Minimum Number of SPARCstorage Arrays	Maximum Number of SPARCstorage Arrays
SPARCserver 1000	3	8
SPARCcenter 2000	3	20

Each SPARCserver 1000 must have a minimum of four CPUs (two in each of the two System Boards). Each SPARCcenter 2000 must have a minimum of six CPUs (two in each of three System Boards).

1.4 Getting Help

If you have problems installing or using Solstice HA, call the distributor from which you purchased the software and hardware and provide the following information:

- Your name and electronic mail address (if available)
- Your company name, address, and phone number
- The model and serial numbers of your systems (factory assembled configurations come with the same serial number for each server)
- The release number of the operating system (for example, Solaris 2.4)
- Any information that will help diagnose the problem

1.5 Glossary

The terminology used in this manual includes the following:

Asymmetric configuration – A configuration that contains a single diskset. In an asymmetric configuration, one server acts as the default master of the diskset and the other server acts as a hot standby.

Concatenation – A metadvice created by sequentially mapping blocks on several physical slices (partitions) to a logical device. Two or more physical components can be concatenated. The slices are accessed sequentially rather than interlaced (as with stripes).

Data service – A network service that implements read-write access to disk-based data from clients on a network. Examples of data services include NFS and ORACLE*SERVER. The data service may be composed of multiple processes that work together.

Default master – The server that is master of a diskset if both servers rebooted simultaneously. The default master is specified when the system is initially configured.

Diskset – A group of disks that move as a unit between the two servers in a Solstice HA configuration.

DiskSuite – See Solstice DiskSuite.

Fault detection – Solstice HA provides programs that detect two types of failures. The first type includes low-level failures such as system panics and hardware faults (that is, failures that cause the entire server to be inoperable). These failures can be detected quickly. The second type of failures are related to data service, such as HA-ORACLE or HA-NFS. These types of failures take longer to detect.

HA-NFS – Highly available NFS. HA-NFS provides highly available remote mount service, status monitor service, and network locking service.

HA-ORACLE – HA-DBMS for ORACLE*SERVER.

Heartbeat – A periodic message sent between the two membership monitors to each other. Lack of a heartbeat after a specified interval and number of retries may trigger a takeover.

Highly available data service – A data service that appears to remain continuously available, despite single-point failures of server hardware or software components.

Hot standby – In an asymmetric (single diskset) configuration, the machine that is not the default master of the diskset. If both servers reboot simultaneously, the hot standby will not master the diskset and thus will not be running any Solstice HA data services.

Local disks – Disks attached to a single Solstice HA server. The local disks contain the Solaris 2.4 distribution and the Solstice HA and DiskSuite software packages. Local disks must not contain data exported by any Solstice HA data service.

Logical host – A diskset and its collection of logical host names and their associated IP addresses. Each logical host is mastered by one physical host at a time.

Logical host name – The name assigned to one of the logical network interfaces. A logical host name is used by clients on the network to refer to the location of data and data services. The logical host name is the name for a path to the logical host. Because a host may be on multiple networks, there may be multiple logical host names for a single logical host.

Logical network interface – In the Internet architecture, a host may have one or more IP addresses. Solstice HA configures up additional logical network interfaces to establish a mapping between several logical network interfaces and a single physical network interface. This allows a single physical network interface to respond to multiple logical network interfaces. This also enables the IP address to move from one Solstice HA server to the other in the event of a takeover or `haswitch(1M)`, without requiring additional hardware interfaces.

Master – The server with exclusive read and write access to a diskset. The current master host for the logical host runs the data service and has the logical IP addresses mapped to its Ethernet address.

Membership monitor – A process running on both Solstice HA servers that monitors the servers. The membership monitor sends and receives heartbeats to its sibling host. The monitor can initiate a takeover if the heartbeat stops. It also keeps track of which servers are active.

Metadevice – A group of components accessed as a single logical device by concatenating, striping, mirroring, or logging the physical devices. Metadevices are sometimes called pseudo devices.

Metadevice state database – Information kept in nonvolatile storage (on disk) for preserving the state and configuration of metadevices.

Metadevice state database replica – A copy of the state database. Keeping multiple copies of the state database protects against the loss of state and configuration information. This information is critical to all metadevice operations.

Mirroring – Replicating all writes made to a single logical device (the mirror) to multiple devices (the submirrors), while distributing read operations. This provides data redundancy in the event of a failure.

Multi-host disk – A disk configured for potential accessibility from multiple servers via multiple ports. Solstice HA software enables data on a multi-host disk to be exported to network clients via a highly available data service.

Sibling host – One of the two physical servers in a Solstice HA configuration.

Solstice HA – See Solstice High Availability.

Solstice High Availability – A software package that enables two machines to act as read-write data servers while acting as backups for each other.

Solstice DiskSuite – A software product that provides data reliability through disk striping, concatenation, mirroring, UFS logging, dynamic growth of metadevices and file systems, and metadevice state database replicas.

Stripe – Similar to concatenation, except the addressing of the component blocks is interlaced on the slices (partitions), rather than placed sequentially. Striping is used to gain performance. By striping data across disks on separate controllers, multiple controllers can access data simultaneously.

Submirror – A metadevice that is part of a mirror. See also mirroring.

Switchover – The coordinated moving of a logical host (diskset) from one operational Solstice HA server to the other. A switchover is initiated by an administrator using the `haswitch(1M)` command.

Symmetric configuration – A Solstice HA configuration that contains two disksets. In a symmetric configuration, each server is the default master for one diskset.

Takeover – The automatic moving of a logical host from one Solstice HA server to the other after a failure has been detected. The Solstice HA server that has the failure is forced to give up mastery of the logical host.

Trans device – A pseudo device responsible for managing the contents of a UFS log.

UFS – An acronym for the UNIX file system.

UFS logging – Recording UFS updates to a log (the logging device) before the updates are applied to the UFS (the master device).

UFS logging device – The component of a trans device that contains the UFS log.

UFS master device – The component of a trans device that contains the UFS file system.

Planning and Installation Overview



This chapter presents an overview of the planning and installation process and explains the supported Solstice HA configurations. Use the following table to locate specific information in this chapter.

<i>Planning Overview</i>	<i>page 2-1</i>
<i>Functional Diagram</i>	<i>page 2-2</i>
<i>Installation Procedure Overview</i>	<i>page 2-6</i>

2.1 Planning Overview

Solstice HA supports specific dual-server hardware configurations. When properly configured, the hardware and software together provide highly available data services.

Solstice HA provides an environment in which data services remain available after one of the servers in the configuration fails.

In Chapter 3, “Installation Planning,” you will find more specific planning information. Be sure to read these two chapters before setting up your Solstice HA configuration.

Any machine on the network can be a client of a Solstice HA system without modifications to the client system or user programs.

Note – A Solstice HA server cannot be a client of the sibling Solstice HA server.

2.2 *Functional Diagram*

The Solstice HA configuration includes two servers, a minimum of three SPARCstorage Arrays, Solstice DiskSuite software, and Solstice HA software. Currently the data services available include Solstice HA-NFS and Solstice HA-ORACLE.

Solstice High Availability requires a dual-server configuration, as shown in Figure 2-1. A list of the specific hardware supported in Solstice HA configurations is given in Section 1.3, “Supported Hardware,” on page 1-2.

Configurations must be identical in the following ways:

- Number and size of local disks
- Number and size of multi-host disks
- Number of network connections
- Number of system boards and CPUs
- Number and location of other optional SBus cards
- Same version of Solaris software and all Solstice HA packages installed on each server’s local disk

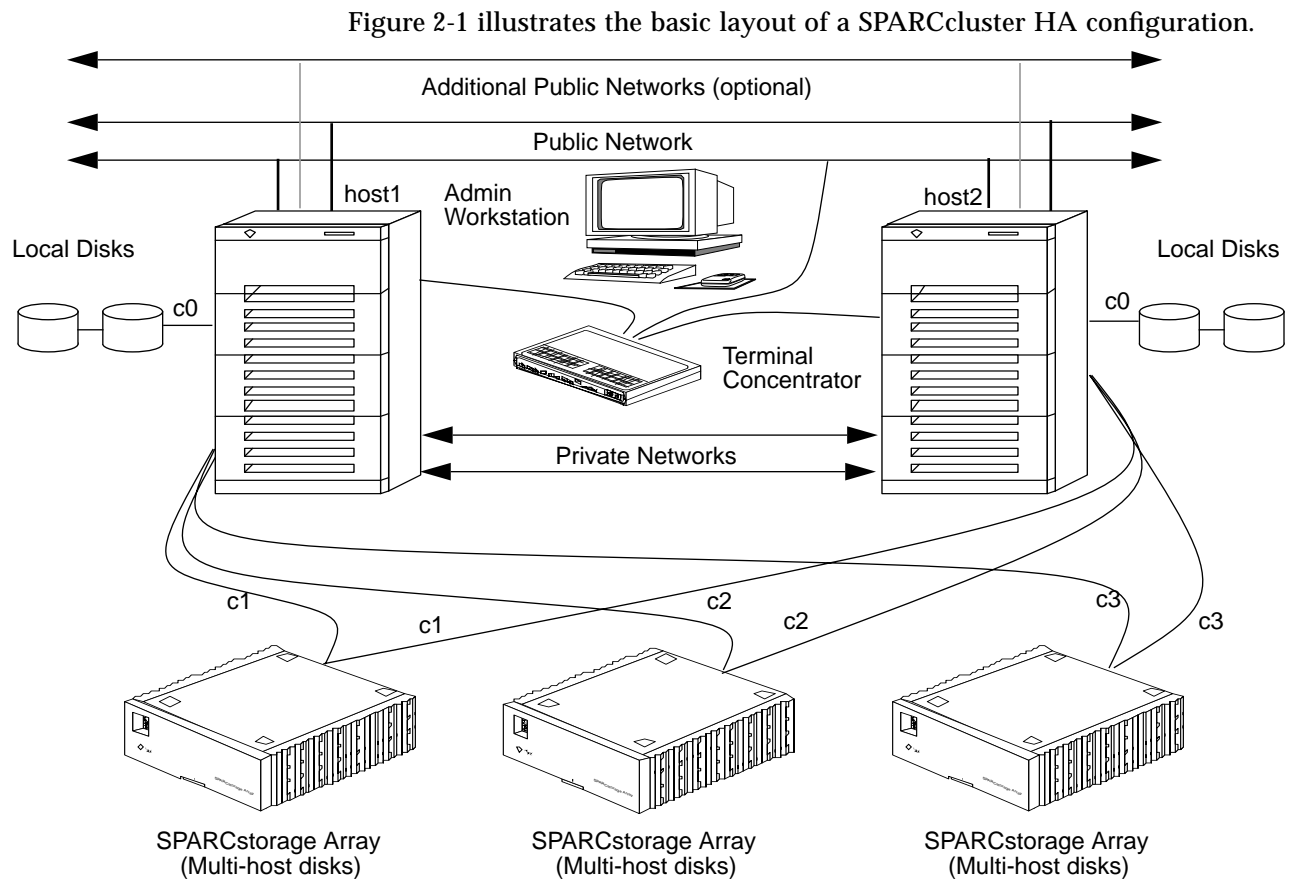


Figure 2-1 Basic Hardware Layout of a SPARCcluster HA Configuration

The parts of the SPARCcluster HA configuration shown in the illustration are:

- There are two SPARCserver 1000s or SPARCcenter 2000s servers. The servers are connected in a Solstice HA configuration provide an environment for highly available services.
- You must have at least one public network must be attached. Additional public networks can be added.
- Two private networks are required. The private network connections are used to exchange Solstice HA status and configuration information. The two private networks must use SunFastEthernet™ SBus cards.

- Fibre Channel SBus cards attached to the same SPARCstorage Array from each of the servers must be connected symmetrically. In Figure 2-1, the strings labeled `c2` from each server are connected to the same controller number and the same SPARCstorage Array. Therefore a disk in this SPARCstorage Array will be accessible as `c2tndn` from each system.
- The hardware configuration is the same regardless of the software configuration that is used. That is, multi-host disks are connected the same in both symmetric and asymmetric configurations. The difference is that in the asymmetric configuration, only one of the servers is actually offering data services to clients.
- Each server must have at least two local disks. The local disks contain the Solaris operating environment, all Solstice HA packages, local metadevice state database replicas, and possibly other site-specific data.
- Multi-host disks must be mirrored. The mirrors must be across separate SPARCstorage Arrays to provide controller fault-tolerance.

- You can create one or two disksets from the multi-host disks. In Figure 2-2 two disksets were created. `Diskset_1` contains 26 drives split as evenly as possible among the three SPARCstorage Arrays. `Diskset_2` has 28 drives, which are also spread fairly evenly across the three SPARCstorage Arrays. If you create only one diskset, the configuration will be asymmetric.

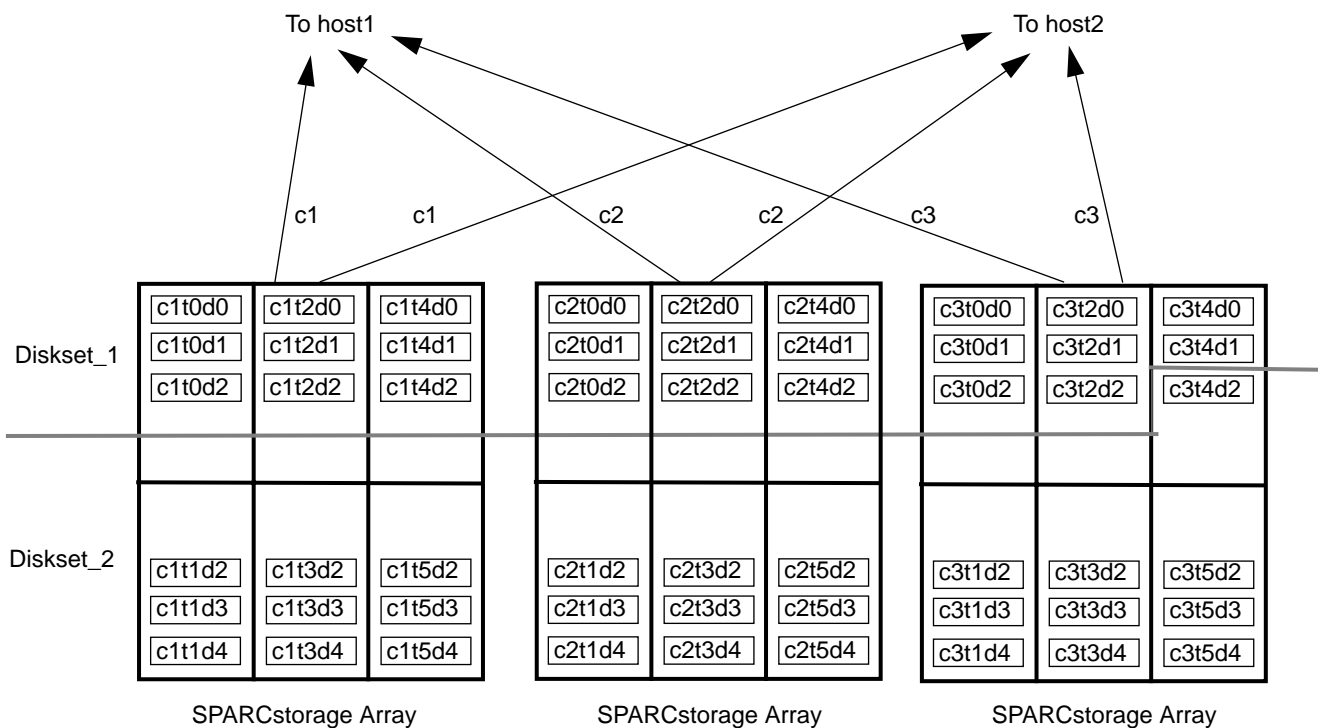


Figure 2-2 Example Diskset Layout With Three SPARCstorage Arrays

- Multi-host disks can be automatically repartitioned and populated with metadvice state database replicas when you assign them to disksets using the `hsetup(1M)` command (see Chapter 8, “Creating the Configuration”).
- Any UFS file system within a diskset must use UFS logging. This includes HA-NFS as well as any file systems used by HA-ORACLE. HA-ORACLE may also operate on raw mirrored metadvice.

2.3 Installation Procedure Overview

Solstice HA installation and configuration involves a complex procedure. You must follow these steps in the order shown below to complete the installation, configuration, and validation of the Solstice HA configuration.

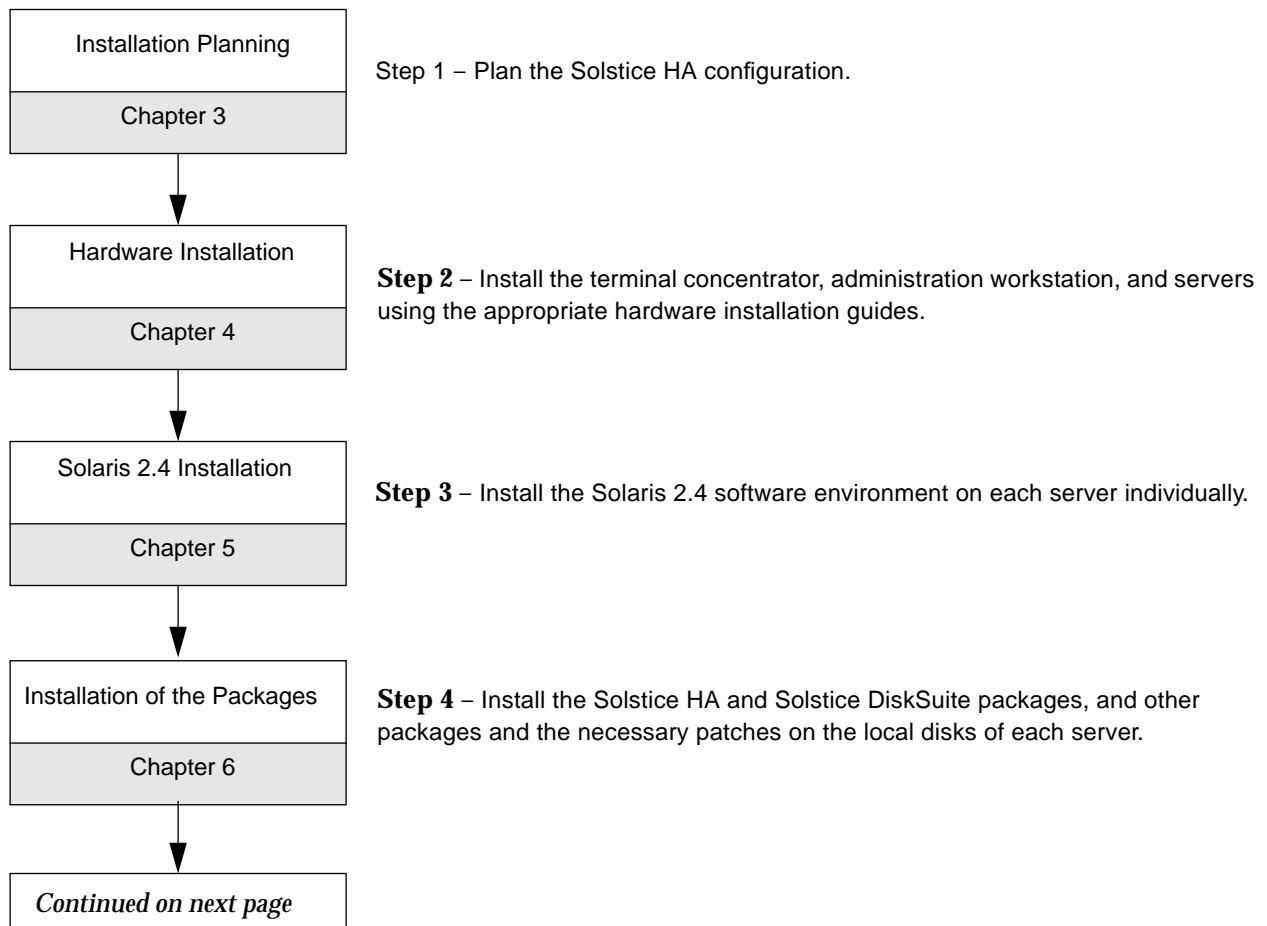


Figure 2-3 Solstice HA Installation Procedure

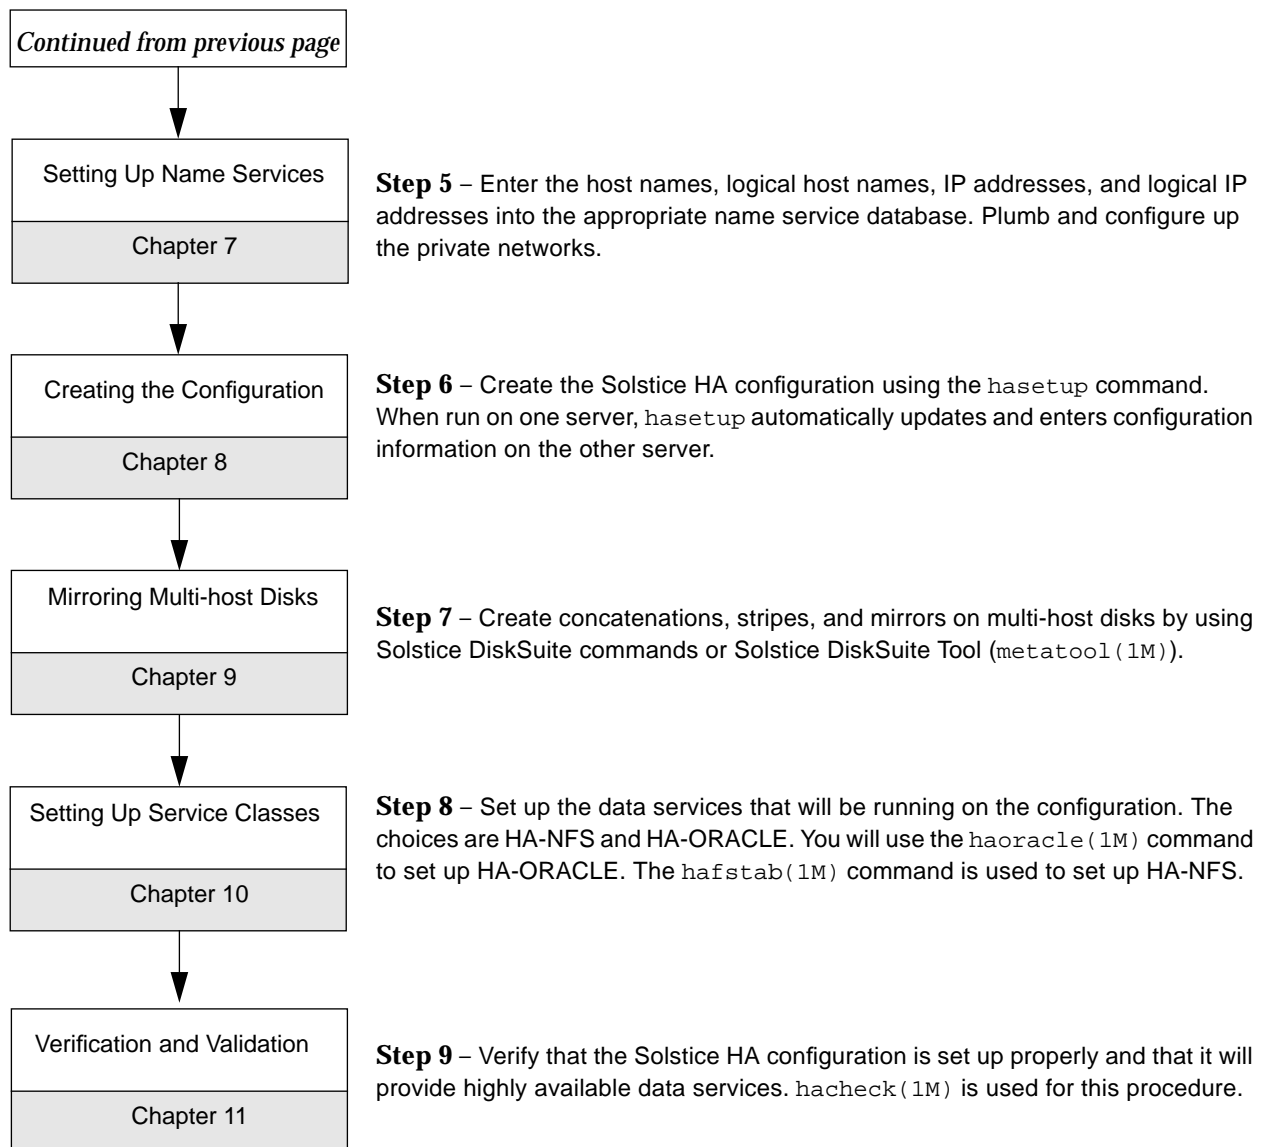


Figure 2-4 Solstice HA Installation Procedure (continued)

Installation Planning



This chapter discusses how you should plan for the installation of Solstice HA servers. Use the following table to locate specific information in this chapter.

<i>Planning Your Configuration</i>	<i>page 3-1</i>
<i>Configuration Rules That Improve Reliability</i>	<i>page 3-23</i>
<i>Configuration Restrictions</i>	<i>page 3-26</i>

3.1 Planning Your Configuration

Solstice HA configuration involves planning for hardware as well as software. The hardware planning includes decisions about network connections, disk space requirements, disk sizes, and the systems that will be used. You will make decisions about the type of configuration (symmetric or asymmetric), makeup of the disksets, and file system layout.

This subsection first presents a list of the questions you will consider when planning your configuration. You are then given detailed information about each step in the planning. The planning questions and information about planning may not be useful in every situation. Your site's configuration may involve special planning.

You should also become familiar with “Configuration Rules That Improve Reliability” on page 3-23, and “Configuration Restrictions” on page 3-26 before planning the configuration.

The following is a list of questions to consider when planning a configuration. A detailed discussion of each question follows the list.

1. How many public network connections will be made and what will be the names of the connections? See Section 3.1.1, “Network Configuration,” on page 3-3.
2. What are the names of the physical hosts and logical hosts? See Section 3.1.2, “Physical and Logical Host Names,” on page 3-8.
3. What data services will be provided from the configuration? See Section 3.1.3, “Data Services,” on page 3-9.
4. How much existing data will be migrating to the multi-host disks? See Section 3.1.4, “Planning Your Disk Requirements,” on page 3-9.
5. What effect will mirroring have on the amount of multi-host data? See Section 3.1.5, “Mirroring,” on page 3-10.
6. What will be the use of trans devices in the configuration? See Section 3.1.6, “Trans Devices,” on page 3-10.
7. How many hot spares will be allocated for use by multi-host disks? See Section 3.1.7, “Hot Spares,” on page 3-11.
8. How much extra disk space will be added for growth? See Section 3.1.8, “Disk Space Growth,” on page 3-11.
9. How much disk space is required for multi-host disks? See Section 3.1.9, “Total Storage Requirement,” on page 3-12.
10. What size and number of disk drives will be used for multi-host data? See Section 3.1.10, “Size and Number of Disk Drives,” on page 3-12.
11. How many SPARCstorage Arrays are needed for the multi-host disks? See Section 3.1.11, “Number of SPARCstorage Arrays,” on page 3-13.
12. What size servers are necessary for your configuration? See Section 3.1.12, “Size of Servers,” on page 3-13.
13. Where will the servers be located and what power resources are necessary? See Section 3.1.13, “Power Sources and Location for Configurations,” on page 3-15.

14. Will the configuration be symmetric or asymmetric? See Section 3.1.14, “Type of Configuration,” on page 3-17.
15. How much data will be placed on the local disks? See Section 3.1.15, “Data on Local Disks,” on page 3-17.
16. What file system hierarchy are you planning for the configuration? See Section 3.1.16, “Planning Your Multi-host File System Hierarchy,” on page 3-18.
17. What is the size and layout of the file system and the metadevices that will be placed on multi-host disks? See Section 3.1.17, “File System Size and Disk Layout,” on page 3-19.
18. How many multi-host disks will be placed in each of the disksets? See Section 3.1.18, “Size of Disksets,” on page 3-19.
19. What will your naming convention be for metadevices and in which manner will they be created? See Section 3.1.19, “Metadvice Naming and Creation,” on page 3-21.
20. How will the data be migrated to the multi-host disks? See Section 3.1.20, “Migrating Existing Data,” on page 3-23.
21. What is the backup strategy for the data on the multi-host disks? See Section 3.1.21, “Backing Up Multi-host Data,” on page 3-23.

3.1.1 Network Configuration

You must have at least one public network connection to a local area network and two private network connections between the systems. Additional public local area networks can be added.

Solstice HA uses logical network interfaces to establish a mapping between several logical host names and a single physical network interface. This enables one of the physical interfaces to respond to multiple logical host names. The physical network interface on the host that currently has the logical interface configured is the one that services packets destined for that logical interface.

In a symmetric configuration, when each diskset is mastered by its respective default master host, there is only one logical host name associated with a physical network connection. However, when a takeover occurs, two logical host names will be associated with one physical network connection.

If another physical network connection is added, a unique logical host name must also be assigned to the connection on each of the machines. These will also move to the sibling during a switchover or takeover.

A unique host name must be assigned for each logical host on each public network. This requires two logical host names per network for symmetric configurations but only one logical host name per network for asymmetric configurations.

This logical naming may lead to many host names. Thus, you may find it useful to adopt a suffix naming convention in assigning the host names for multiple networks. In Figure 3-1, the last octet of the network number was

appended to the logical host name to distinguish it from the other logical host name (201). You could also use the Ethernet interface name such as `qe0` as a naming suffix.

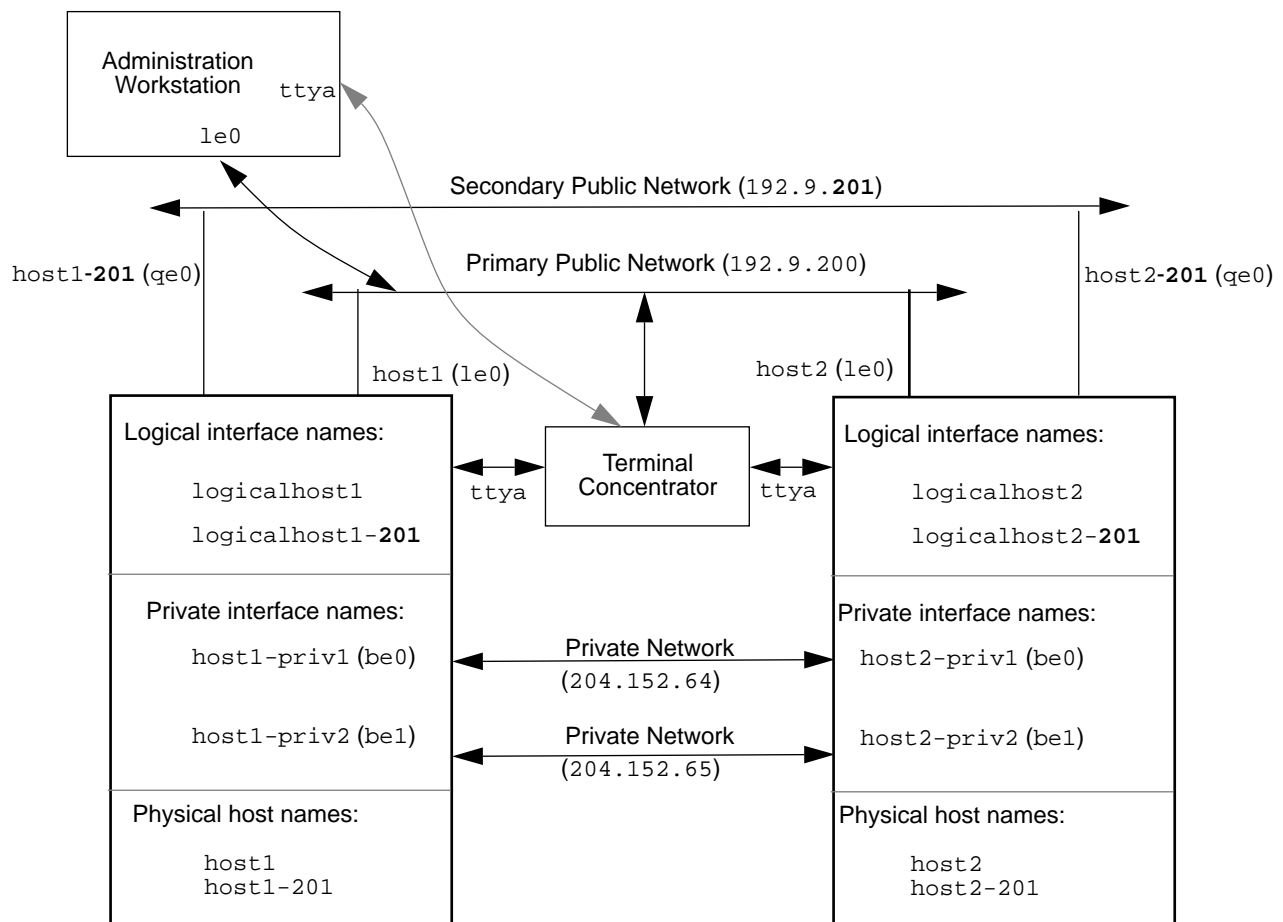


Figure 3-1 Interface Names of Logical Network Interfaces

The following worksheets show the information for the primary and secondary public networks for the configuration shown in Figure 3-1.

Table 3-1 Primary Network Naming

Primary Network Name: primary-net1	Primary Network Number: 192.9.200	
Host Name: host1 host2	Physical Host Name IP Address: 192.9.200.100 192.9.200.101	Network Interface: 1e0 1e0
Host Name: logicalhost1 logicalhost2	Logical Host Name IP Address: 192.9.200.110 192.9.200.111	

In Table 3-1 the physical host name of a server on the primary network is by convention the nodename reported by the `uname -n` command. Similarly, a logical host name on the primary network is used as the name of the logical host associated with specific Solstice HA data services, and as the name of the diskset containing the data.

Table 3-2 Secondary Network Naming

Secondary Network Name: secondary-net2	Secondary Network Number: 192.9.201	
Host Name: host1-201 host2-201	Physical Host Name IP Address: 192.9.201.100 192.9.201.101	Network Interface: qe0 qe0
Host Name: logicalhost1-201 logicalhost2-201	Logical Host Name IP Address: 192.9.201.110 192.9.201.111	

The information in the worksheets in Table 3-1 and Table 3-2 can be used to make entries in the `/etc/inet/networks`, `/etc/inet/host`, and `/etc/inet/netmasks` files as well as the corresponding name services tables and maps. It can also be useful when configuring your Solstice HA systems with `hasetup(1M)`.

Solstice HA systems require two private networks for normal operation. These networks are implemented on SunFastEthernet cards with point-to-point cables between the two Solstice HA systems. Only Solstice HA private traffic appears on these networks.

Two class C network numbers (204.152.64 and 204.152.65) have been reserved for private network use by the servers. The same network numbers are used by all Solstice HA systems.

One way you may wish to form the host names used on these networks is by appending a suffix of `-priv1` (for 204.152.64) and `-priv2` (204.152.65) to the physical host name. An example of this is illustrated in the Table 3-3.

Table 3-3 Example Private Network Naming

IP Address	Host Name	Network Interface
204.152.64.1	host1-priv1	be0
204.152.65.1	host1-priv2	be1
204.152.64.2	host2-priv1	be0
204.152.65.2	host2-priv2	be1

These names must be entered in the `/etc/inet/hosts` file of both of the Solstice HA servers, but not in the host table of the network name service because they are private.

In addition to `/etc/inet/hosts` these private host names must appear in the `/.rhosts` file of both Solstice HA systems to allow `hasetup(1M)` to communicate between the two systems during configuration.

Finally, you must assign IP addresses and host names for the terminal concentrator and the administration workstation.

Some points to consider when planning network connections are:

- You must create a mapping of the network name to the hardware. For example, a sample mapping would include the IP address (192.9.201.2), the host name (`host2-221`), and hardware network interface name (`qe0`). This information is used in the network configuration files. Refer to Chapter 7, "Setting Up Name Services and Networks," for additional information.
- The bandwidth requirements at your site may affect the number of public networks you configure.
- There are network administration issues to consider, including those for data-intensive applications. Refer to the *SMCC NFS Server Performance and Tuning Guide* (part number 801-7289) for additional information.

- By default the two private networks will be assigned class C network numbers 204.152.64 and 204.152.65. These network numbers can be changed during configuration.

3.1.2 *Physical and Logical Host Names*

You must select host names for the two Solstice HA servers and a logical host name for each of the logical hosts. In an asymmetric configuration, you will only need one logical host name.

Some points to consider when selecting physical and logical host names are:

- You must select a primary host name for each server. These host names are entered as part of the Solaris 2.4 installation.
- You must select a logical host name for each logical host. In a symmetric configuration there are two logical hosts, while there is only one in an asymmetric configuration. These names are entered when you create the configuration using `hasetup(1M)`.

Note – The primary logical host name and the diskset name must be the same string value.

- There must be two private host names for each server for the Solstice HA software to communicate across the two private network connections. A possible naming convention is to use suffixes to show the association between the private host names and the host names. An example of the suffixes used would be `host1-priv1` and `host1-priv2`. It is not necessary to follow such a host name convention. The names must be manually entered in the `/etc/inet/host` file and you must create `/etc/hostname.be[01]` files on each server.
- If additional public networks are used, you must create a secondary host name for each server and for each of the logical hosts for each such network. For example, if you are setting up a symmetric configuration with four public networks you will need to assign 20 host names and IP addresses for the Solstice HA configuration. This includes two physical host names and two logical host names for each network. In addition you will need four host names and IP addresses for the private networks.
- If you change the host name of a machines with existing disksets, the Solstice DiskSuite configuration (all your data) will become inaccessible.

3.1.3 Data Services

Planning a configuration requires deciding which data services to run. The two types of data services supported are HA-NFS and HA-DBMS.

Some points to consider when deciding which data services to use are:

- Both of these services can be run on the same configuration.
- HA-NFS provides highly available NFS file service. This means the failure of one of the systems in a configuration causes the HA-NFS client to experience a brief delay while the takeover is in progress.
- HA-DBMS provides highly available DBMS service. Currently HA-ORACLE is the only supported HA-DBMS. HA-ORACLE client applications experience a database recovery interruption in the service during either switchover or takeover. Because database transactions are “stateful,” the client applications may need to re-establish the state of their transactions after either a switchover or takeover.

3.1.4 Planning Your Disk Requirements

You should add up the amount of data that you want to move to the Solstice HA configuration. Double that amount to allow space for mirroring. You may also want to consider growth and hot spare capacity.

Some points to consider when planning the amount of data are:

- When you are calculating the amount of data to migrate to Solstice HA, you should keep in mind the sizes of disks available for the SPARCstorage Arrays (currently, 1.05 Gbytes and 2.1 Gbytes). If the file systems or table spaces you are moving do not fit evenly on these disks, you should plan to use an extra disk and concatenate the slices with Solstice DiskSuite.
- Under some circumstances, there may be an advantage to merging several smaller file systems into a single larger file system. This will reduce the number of file systems to administer and may help speed up takeovers if one of the hosts fails.
- The size of the dump media (backup system) may influence the size of the file systems in Solstice HA configurations.

3.1.5 Mirroring

All multi-host disks must be mirrored in Solstice HA configurations. This enables the configuration to tolerate single-component failures.

Some points to consider when mirroring multi-host disks are:

- Each submirror must be in a different SPARCstorage Array.
- With two-way mirroring, the amount of disk space needed will double.
- Three-way mirroring is supported by Solstice DiskSuite. However only two-way mirroring is required for Solstice HA. If you perform three-way mirroring, the amount of disk space required will triple.
- Under Solstice DiskSuite mirrors are made up of other metadevices such as concatenations or stripes. In large configurations there may be a large number of metadevices. For example, seven metadevices are created for each logging UFS file system. Refer to the *SPARCcluster High Availability Software Administration Guide* for additional information about accommodating the large number of metadevices you may have.
- If you mirror metadevices of a different size, your mirror capacity will only be the size of the smallest submirror.
- When a new device is added to an existing diskset, another device of the same size must be added to mirror the data.

For additional information about mirroring, refer to the *Solstice DiskSuite 4.0 Administration Guide*.

3.1.6 Trans Devices

Trans devices are composed of a UFS log and a UFS master. Under Solstice HA both the log and master must be mirrored. The mirrors should be constructed following the mirroring guidelines given in Section 3.1.5, “Mirroring.”

Some points to consider about trans devices are:

- Because the log and master devices often perform I/O at the same time, you would want to place them on separate drives in order to minimize the disk seek activity.
- Logs should reserve one Mbyte of space for every 100 Mbytes of the trans master device space.

3.1.7 Hot Spares

Solstice DiskSuite's hot spare facility automatically replaces failed submirror components, provided that a suitable spare component is available and reserved. Hot spares are temporary fixes, allowing you to schedule a time when the systems are less busy to perform maintenance on the failed component.

Some points to consider when planning for hot spares are:

- If you are going to use hot spares, at least one should be assigned to each diskset.
- Hot spares are useful for assuring the availability of critical data.
- New devices can be added to an existing diskset at any time to be used as hot spares.
- When deciding on the location of hot spares you must consider whether you are attempting to protect your data from media, tray, system board, or SPARCstorage Array failure.

For additional information about hot spares, refer to the *Solstice DiskSuite 4.0 Administration Guide*.

3.1.8 Disk Space Growth

When planning your data needs when moving to the Solstice HA platform, you may want to consider growth.

Some points to consider when planning for growth are:

- It will require less administration time to add disks during initial configuration than when the system is in service.
- You may want to leave empty slots in the SPARCstorage Arrays during initial configuration. This will allow the addition of disks later. If additional slots are not available, you may be in a situation where you must move SPARCstorage Arrays or reconfigure all the metadevices.
- When your site needs additional SPARCstorage Arrays, you may have to reconfigure your data to prevent mirroring within a single SPARCstorage Array chassis. The easiest way to add SPARCstorage Arrays without reorganizing data is to add them in pairs, if you are in a situation where all the existing SPARCstorage Array chassis are full.

3.1.9 Total Storage Requirement

To calculate the total storage requirement for Solstice HA you should compute the amount of existing data that you are migrating. That number must be multiplied by two for mirroring.

Some points to consider when calculating the total storage requirement are:

- A small fraction of each disk will be used by Solstice DiskSuite for metadvice state database replicas and for UFS logging. Typically this is less than 10 to 20 Mbytes per disk. The UFS logging space will only be necessary if you are creating UNIX file systems.
- The number of hot spares that will be used in each diskset counts towards the total amount of storage required.
- You must round up on a per file system basis. For example, if you have four file systems that contain three Gbytes each, you will need two 2.1 Gbyte drives for each file system, or a total of 16 Gbytes, eight 2.1 Gbyte drives. If you simply add the amount of file system space, you would conclude incorrectly that only 12 Gbytes of disk space, or six drives are needed.
- To simplify administration, use the default disk slicing performed by `hasetup(1M)`. The default slicing gives you a single large slice in addition to small slices for a metadvice state database and a UFS log. While you can repartition the disk to create additional smaller slices, you may find the administrative overhead of tracking the partition tables and metadvice time consuming.

3.1.10 Size and Number of Disk Drives

Two sizes of disks are supported in SPARCstorage Arrays: 1.05 Gbyte and 2.1 Gbyte. Calculate the number of drives needed by dividing the total amount of storage required by the size of disk selected.

Some points to consider when deciding which size drives to use are:

- If you use lower capacity drives you can have more spindles, which increases potential I/O bandwidth. This assumes the disks have the same I/O rates.
- Using higher capacity disks means fewer devices are required in the configuration, which can help speed up takeovers. This is because takeover time is dependent on the number of drives being taken over.

- Using higher capacity disks can also reduce the chances of a disk failure, because fewer drives translates to lower aggregate failure rate.
- You should keep in mind that during configuration, space may automatically be allocated on the multi-host disks for Solstice DiskSuite metadevice state database replicas and at your option for possible UFS logs. These disk slices are very small, roughly a few Mbytes.
- To determine the number of disks needed, the total disk capacity you have selected is divided by the disk size. For example, if you need 40 Gbytes of total multi-host data storage, you would use 20 2.1-Gbyte disks.

3.1.11 Number of SPARCstorage Arrays

The minimum number of SPARCstorage Arrays required in a Solstice HA configuration is three. This minimum configuration means three sets of metadevice state database replicas are available to protect the configuration information in the event one SPARCstorage Array fails.

Some points to consider when deciding how many SPARCstorage Arrays to use are:

- When additional SPARCstorage Arrays are in a configuration, administration flexibility is available when one of the arrays fails.
- It is recommended the disks placed in the SPARCstorage Arrays be divided evenly between chassis, trays, and buses within the trays.
- Future growth of the multi-host data is easier to perform if there are empty slots in the SPARCstorage Arrays. You may want to consider planning the number of SPARCstorage Arrays so that growth can easily be accommodated.

3.1.12 Size of Servers

After you have decided the data services, number and size of each instance of data to be exported by the data service type, and the amount of data to be stored on the local and multi-host disks, selection of the hardware platform becomes easy.

Some points to consider when deciding which size servers should be used are:

- SPARCserver 1000s can be used if the configuration includes three to eight SPARCstorage Arrays. Otherwise, use SPARCcenter 2000s, which can have up to 20 SPARCstorage Arrays attached.
- Determining whether to use two SPARCserver 1000 or two SPARCcenter 2000 configurations will depend on your requirements for SPARCstorage Array storage, CPU cycles, memory, and other site-specific needs. The following table shows the physical limits and some other factors to consider.

Table 3-4 Physical Limits of SPARCcluster Components

Components	SPARCserver 1000 Minimum	SPARCserver 1000 Maximum	SPARCcenter 2000 Minimum	SPARCcenter 2000 Maximum	Notes
System Boards	2	4	3	10	
CPUs	4	8	6	20	Maximum of 2 per System Board
Memory	128 Mbytes	2 Gbytes	256 Mbytes	5 Gbytes	512 Mbytes maximum per System Board
SBus Slots	6	12	12	40	See Note 1
SPARCstorage Arrays	3	8	3	20	See Note 2
Storage Capacity					See Note 3
– Number of Disks	14	240	20	600	30 per SPARCstorage Array
– 1.05 Gbyte Disks	14.7 Gbytes	252 Gbytes	21 Gbytes	630 Gbytes	
– 2.1 Gbyte Disks	29.4 Gbytes	504 Gbytes	42 Gbytes	1.26 Tbytes	

Note 1 – Each server requires two SunFastEthernet SBus cards to be installed for the private networks. The SPARCcenter 2000-based configurations also require two FSBE/S cards to be installed on each server for the private disks, tape drives, and CD-ROM drives, along with at least one SBus card for the client network. Depending on your client network needs, you may require additional client network interfaces.

Note 2 – Each SPARCstorage Array uses one SBus slot on each server. Each FC/S SBus card only supports a single FC/OM in Solstice HA configurations. You must calculate these SBus cards into the total SBus slot requirement.

Note 3 – The storage capacity minimum for the SPARCserver 1000 is 14 disks. The minimum for the SPARCcenter 2000 is 20 disks, Availability is improved if the disks are divided as evenly as possible across the three SPARCstorage Arrays.

Note – The maximums of 240 disks in the SPARCserver 1000 and 600 disks in the SPARCcenter 2000 configurations are based on fully-populated SPARCstorage Arrays. This may not be suitable for your needs from a SCSI bus loading perspective.

3.1.13 Power Sources and Location for Configurations

You must decide what kind of power sources will be used to provide power to the servers and SPARCstorage Arrays. You may choose to use an uninterruptable power supply (UPS). If the power system at your site is reliable you may not need to take precautions.

Highest availability will be obtained if the key components of the HA system have independent power sources. For this discussion, independent power sources means that power failure of one source will not cause power failure in another source. For example, multiple UPS systems will have independent power failures since they have no critical component in common. Separate branch circuits are typically not independent since they share a common main circuit breaker somewhere along the way.

You might want to use multiple UPS systems to protect your system against failures in any one UPS. Those with lower availability requirements may find one UPS satisfactory. Those unconcerned with power failures may find commercial power acceptable.

The requirement for ongoing system operation in light of partial power failure is that at least one processor and a majority of the disk drives survive the partial power failure. This is most easily done by ensuring that power to each

processor is from an independent source and that power to the SPARCstorage Arrays is such that a majority of the drives survive failure of one power source. Three power sources are sufficient for the typical Solstice HA configuration.

- The first for a server and one SPARCstorage Array.
- The second for the other server and another SPARCstorage Array.
- The third for the last SPARCstorage Array.

Additional SPARCstorage Arrays should be distributed evenly across the three power sources or additional independent power can be supplied.

The two processors are each attached to separate power sources while the three SPARCstorage Arrays are also attached to separate power sources (it is acceptable for one storage array and one processor to share a power source). Configurations with more than three SPARCstorage Arrays can attach them to the three power sources in turn, keeping the distribution even so no single power source controls half or more of the disks. This ensures a majority of the disks will survive partial power failure, which is important because of the availability of the metadevice state database replicas.

Now if these sources are truly independent it should be unlikely that more than one of them should fail. However, in the unlikely event that all power sources fail it is difficult to control the sequence in which power is restored. Lacking this control it is possible for the processors to boot before the SPARCstorage Arrays and so be unable to access the disks. In this case the processor must reboot to access the disks. In some rare circumstances the processors may be unable to obtain access to the disks and manual intervention may be required.

Some points to consider when deciding the location are:

- The two Solstice HA servers can be located a maximum of 100 meters apart.
- You may want to separate the servers and connect them to different power sources.
- You may want to separate the servers so they are not both located under the same sprinkler system zone (if you have more than one computer room or more than one sprinkler zone).
- You may want to separate the servers into different air conditioning cooling zones (if you have more than one computer room or more than one cooling zone).

3.1.14 *Type of Configuration*

Solstice HA allows configurations to be either symmetric or asymmetric. In the symmetric configuration, each system masters one of the two disksets and makes one or more data services available from that diskset. In the event of a failure, the remaining system takes control of both disksets and provides all the data services.

In an asymmetric configuration there is only one diskset and all data services run from that diskset on only one of the hosts. The second host is a hot standby that is ready to assume control of the diskset and provide the data services in the event of a failure.

Some points to consider when deciding whether to have a symmetric or asymmetric configuration are:

- In an asymmetric configuration, one host will be idle until the other host goes down and a failover occurs.
- An asymmetric configuration is more likely to avoid an overload after a switchover.
- In a symmetric configuration, both hosts are actively providing data services until a failure of one of the systems occurs.
- In a symmetric configuration you may experience overload of a server following a failover. You must plan and monitor usage to avoid an overload.

3.1.15 *Data on Local Disks*

Each of the Solstice HA systems must have the same software available on the local disks. This includes the Solaris operating environment, all Solstice HA software, all Solstice DiskSuite software and all ORACLE software if you are running HA-DBMS.

Some points to consider when planning for data on the local disks are:

- If you have sufficient disk space, you can optionally mirror the data on the local disks. However mirroring the local disks may decrease availability because you are not mirroring across controllers. The mirroring is not necessary because the local data is replicated on the local disks of the two servers in the Solstice HA configuration.

- During installation of the Solaris operating environment you will be instructed to leave adequate space for three Solstice DiskSuite metadvice state database replicas and for a UFS file system log. Sun recommends you use Slice 4 of the boot disk for the three replicas. See Chapter 5, “Solaris 2.4 Installation” for instructions on setting up the metadvice state database replicas.
- The root slice (/) of the local disk must be large enough to accommodate the Solaris 2.4 system components that reside in root along with any components of any additional software packages that are installed. The root slice must also have enough inodes for the various files and directories as well as the device inodes in /devices and symbolic links in /dev. Refer to Section 5.2, “Selecting the Size of the Root Slice,” on page 5-2 for detailed information about the size of the root slice.
- Solstice HA is supported only on the Solaris 2.4 Hardware Release 395 operating system. It will not run on earlier or later versions of Solaris.

3.1.16 Planning Your Multi-host File System Hierarchy

Regardless of whether you are running HA-ORACLE or HA-NFS you are required to have multi-host UNIX file systems. These file systems must be mirrored file systems set up with UFS logging. This will affect your metadvice configuration.

You must design a hierarchy of mount points for these file systems. These will be entered in the /etc/opt/SUNWhadf/hadf/vfstab.*logicalhost* file. Refer to Section 10.3, “Setting Up HA-NFS,” on page 10-10 for detailed instructions.

Solstice HA conventions require that a small file system be created and mounted on /*logicalhost*. The size this file system is a minimum of 10 Mbytes. This file system will contain the directory named *statmon*.

Some points to consider when planning your multi-host file system layout are:

- The file system that contains /*logicalhost*/*statmon* cannot be grown.
- You must create mount points (file systems) other multi-host file systems at the /*logicalhost* level. You can grow these file systems using the Solstice DiskSuite `growfs(1M)` command.

- Your application may dictate a file system hierarchy and naming convention. Solstice HA has no restrictions about the file system naming, provided it begins with */logicalhost* and it does not conflict with the */logicalhost/statmon* directory.

3.1.17 File System Size and Disk Layout

After deciding on data services, you must decide the number and size of each instance of these types. For example, if you decide to run only HA-NFS, you must determine the number and size of the file systems that will be exported.

If you plan to use the Solstice HA configuration for existing data that is not mirrored, you will need twice the amount of currently used disk space. If the data is already mirrored, the space required is approximately equal. The DiskSuite requirement for metadevice state database replicas and for UNIX file system (UFS) logging is about one percent of each of the multi-host disks.

Ideally, the load on each diskset in a symmetric Solstice HA configuration should be balanced. The data services should be balanced on each of the disksets. This makes configuration easier and the performance of each system would be similar. Balancing the load may not be possible at every site.

Some points to consider when deciding on file system size and disk layout are:

- Ideally, the minimum file system size is the size of the disk. This will make general administration and Solstice DiskSuite administration easier and will allow greater flexibility in meeting space requirements for users.
- If you share UFS logs among several file systems, keep in mind that if a UFS panic or a log I/O problem occurs each of the file systems that are sharing the log must be checked by `fsck(1M)` before a takeover can be completed.

3.1.18 Size of Disksets

During initial configuration, the `hasetup` command automatically attempts to place the same number of disks in each diskset of a symmetric configuration.

The following example helps to explain the process you should use to determine the number of disks to place in each diskset. In this example there are existing applications running over NFS (two file systems of five Gbytes each) and several ORACLE databases (one five Gbytes and one 10 Gbytes).

Table 3-5 Determining Drives Needed for a Configuration

Use	Data	Disk Storage Needed	Drives Needed
nfs1	5 Gbytes	3x2.1 Gbyte disks * 2 (Mirror)	6
nfs2	5 Gbytes	3x2.1 Gbyte disks * 2 (Mirror)	6
oracle1	5 Gbytes	3x2.1 Gbyte disks * 2 (Mirror)	6
oracle2	10 Gbytes	5x2.1 Gbyte disks * 2 (Mirror)	10

Table 3-5 shows the calculations used to determine the number of drives needed in the example configuration. The number needed is 28, which would be divided as evenly as possible among each of the three SPARCstorage Arrays. Note that the five Gbyte file systems were given an additional Gbyte of disk space because the number of disks needed was rounded up.

Each of the HA-NFS file systems and the HA-ORACLE databases will be placed on separate servers. This is because of the processing power necessary for the applications, the correlation in time of access (I/O), and possibly other site-specific factors. Because the two servers are made up of identical hardware, it may be beneficial to balance the loads evenly between the servers.

Table 3-6 Division of Disksets

Logical host (diskset)	Data Services	Disks	SPARCstorage Array 1	SPARCstorage Array 2	SPARCstorage Array 3
logicalhost1	nfs1/oracle1	12	4	4	4
logicalhost2	nfs2/oracle2	16	5	6	5

Initially, four disks on each SPARCstorage Array will be assigned to logicalhost1 and five or six disks on each will be assigned to logicalhost2. The four disks in logicalhost1 will be distributed evenly across the three trays (and six buses) in the SPARCstorage Arrays. Similarly, the five or six drives will be spread across the six buses of each tray (using slots 2 and 7 of each tray and slot 3 of tray 1 and 2. This distribution gives the

maximum number of buses for each diskset. In Figure 3-2 the disk allocation is illustrated, as the disks are labeled with the name of the diskset (1 for logicalhost1 and 2 for logicalhost2.)

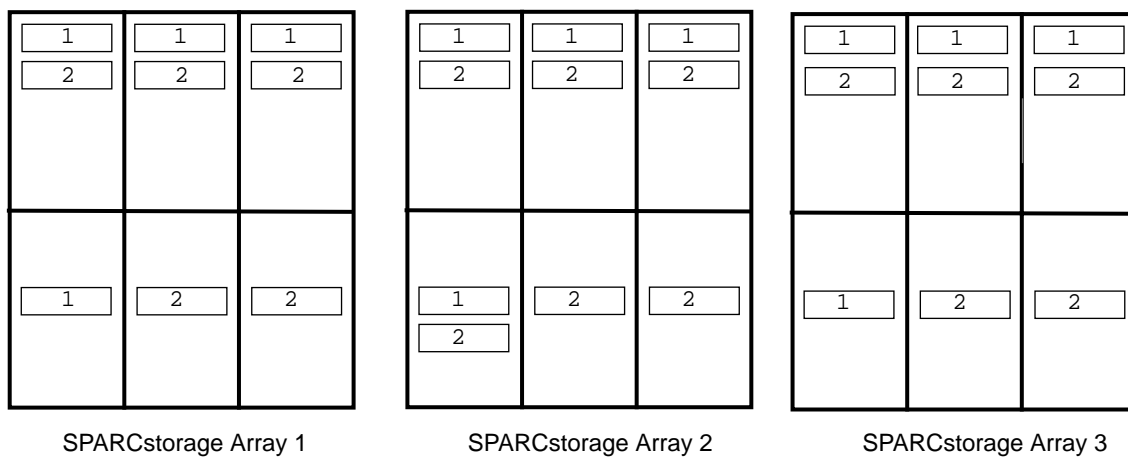


Figure 3-2 Example Diskset Allocation

No hot spares have been assigned to either diskset. A minimum of one hot spare for each diskset would allow one drive to be hot spared (restoring full two-way mirroring). However, there would be a short interval of one-way mirroring on at least some data when a tray is pulled from the SPARCstorage Array to make a repair. Larger configurations that use additional SPARCstorage Arrays might allow a tray to contain drives for only one diskset. This would provide a better opportunity to use hot spares.

A point to consider when deciding the size of disksets is:

- In an asymmetric configuration there is only one diskset. All multi-host disks will be placed in the single diskset.

3.1.19 Metadevice Naming and Creation

It is important to determine the number of metadevice names needed for your configuration before setting up the configuration because if the number needed exceeds 128 you must edit the `/kernel/drv/md.conf` file.

Changes to the `/kernel/drv/md.conf` file do not take effect until a reconfiguration reboot is performed. The `md.conf` files on each server must be identical.

Note – The *Solstice DiskSuite 4.0 Administration Guide* tells you the only modifiable field in the `/kernel/drv/md.conf` file is the `nmd` field. You can also modify the `md_nsets` field from 4 to 3.

To calculate the number of metadevice names needed, determine the larger of the number of metadevice names to be used in either diskset (do not add the requirements of the disksets).

If you follow the naming convention suggested “Managing the Metadevice Name Space” on page 9-6, you should allocate 10 metadevice names per UFS logging file system. Do not forget the small file system mounted on `/logicalhost`. If you are using raw mirrored metadevices (for example, for HA-ORACLE), you should allocated 10 metadevice names per three mirrored devices.

You should set `nmd` in `/kernel/drv/md.conf` to the larger of either diskset. If you modify `md.conf` you should also set `md_nsets` to 3 (rather than the default 4, Solstice HA uses at most 3 disksets, the local set and two multi-host disksets). This will help conserve on inodes in the root file system (and so the space that must be allocated in the root file system). If you are going to change the `nmd` value, it should be raised in factors of two (that is, 128, 256, 512, or 1024).

These values will be used to help determine the size of your root file system. See Table 5-4 for a list of the root file system sizes.

Some points to consider when naming and creating metadevices include:

- You must read and use all information in Chapter 12 of the *Solstice DiskSuite 4.0 Administration Guide*. This chapter covers all the configuration guidelines and considerations for performance, availability, capacity, security, and compatibility.
- Metadevices can be created using either the Solstice DiskSuite commands or the Solstice DiskSuite Tool graphical user interface, `metatool(1M)`. The `metatool` has certain limitations which are documented in the *Solstice DiskSuite Tool 4.0 User's Guide*.

3.1.20 Migrating Existing Data

During initial configuration, the `hasetup` command will optionally repartition all drives. All data on the drives will be lost. Thus, you cannot simply move data into the Solstice HA configuration by connecting existing disks that contain data.

A point to consider when planning to move existing data to the multi-host disks is:

- You can use `ufsdump(1M)` and `ufsrestore(1M)` to migrate data to Solstice HA servers.
- When migrating ORACLE databases to a Solstice HA configuration, use the method recommended by ORACLE.

3.1.21 Backing Up Multi-host Data

Before loading data on the multi-host disks in a Solstice HA configuration you should have a plan in place for backing up the data. Sun recommends you use Solstice Backup or `ufsdump` and `ufsrestore` to backup your SPARCcluster High Availability configuration.

A point to consider when planning a strategy for backup of multi-host data is:

- If you are converting your backup method from Online: Backup™ to Solstice Backup™, special considerations exist because there is no interoperability between the products. The primary decision for the system administrator is whether or not the files backed up with Online: Backup will be available online after Solstice Backup is in use. Refer to the *Solstice Backup 4.1.2 Administration Guide* for additional details about conversion.

3.2 Configuration Rules That Improve Reliability

The rules discussed in this section help ensure that your Solstice HA configuration is highly available. These rules also help determine what number of disks and controllers are appropriate for your configuration.

- The two systems must have identical local hardware. This means if one SPARCserver 1000 is configured with two local disks, the sibling SPARCserver 1000 must also have two local disks.
- Multi-host disks must be connected identically on the two servers.

- The two systems must have identical network interface configurations. That means the systems must have the same number and type of private and public network connections. For example, if one node in your configuration has two SunFastEthernet Adapter cards, the sibling systems must have the same.
- At least three SPARCstorage Arrays are required to build a highly available system. That means each of the servers in the configuration must have enough system boards and Fibre Channel SBus cards installed to support a minimum of three multi-host SPARCstorage Arrays.
- In a configuration that has more than three SPARCstorage Arrays, the metadevice state database replicas will be distributed among the multi-host disk strings so that half of the replicas plus one are accessible when a single hardware failure occurs (for example, disk, controller, power supply, or cable).

Note – During initial configuration, metadevice state database replicas are automatically placed on multi-host disks by the `hasetup` command. Placement of replicas on the local boot disks is a manual procedure.

- Do not mirror the root (`/`), `swap`, `/usr`, `/var`, or `/opt` file systems, which reside on local disks. The second server in the configuration acts as a replication for the data on these file systems. Other file systems that reside on local disks and contain important data can be mirrored.
- The `/usr`, `/var`, and `/opt` file systems can be logged to improve boot performance.
- At least two local disks must be attached to each system. One disk will be used for the Solaris 2.4 distribution, Solstice HA, and Solstice DiskSuite. The second local disk can be used for the ORACLE software distribution. The boot disk must also have three metadevice state database replicas.

- You must have the same number of multi-host Fibre Channel SBus cards attached to each system in the Solstice HA configuration. You must also ensure that the Fibre Channel SBus cards are cabled symmetrically.




Figure 3-3 Solstice HA Systems With Multi-host Disks

- You must have two private networks. Interfaces be0 and be1 are reserved for these private network interfaces. During initial configuration, the hasetup command will look for be0 and be1.

3.3 Configuration Restrictions

The following are not supported in Solstice HA configurations:

- Do not configure the Solstice HA servers as mail servers, because `sendmail(1M)` is not supported in a Solstice HA environment. No mail directories should reside on Solstice HA servers.
- Do not configure Solstice AutoClient™ or diskless clients that are served by Solstice HA servers (that is, the directory for diskless clients should not exist on Solstice HA systems).
- Do not configure Solstice HA systems as routers (gateways). If the system goes down, the clients cannot find an alternate router and recover.
- Do not configure Solstice HA systems as NIS or NIS+ servers. Solstice HA servers can be NIS or NIS+ clients.
- On Solstice HA systems, users should not locally access Solstice HA file systems that are NFS exported. The reason for this restriction is:
 - Local locking interferes with the ability to `kill(1M)` and restart `lockd(1M)`. Between the `kill` and the restart, a blocked local process would be granted the lock, which prevents the client machine that owns that lock from reclaiming it.
- Solstice HA does not support either Secure NFS or Kerberos.
- Do not run any applications on either server that access the HA-NFS file systems.
- The file system containing the Solstice HA `/logicalhost/statmon` directory cannot be grown using the DiskSuite `growfs(1M)` command. Refer to “Planning Your Multi-host File System Hierarchy” on page 3-18 for additional information.
- The network time protocol (NTP), which is not shipped by Sun Microsystems, cannot be run on Solstice HA systems.
- A Solstice HA configuration cannot be used to provide a highly available boot or install service to client systems.
- A diskset must have disks in each of the three SPARCstorage Arrays. No SPARCstorage Array can contain 50 percent or greater of the disks in a diskset.
- File system quotas are not supported.

-
- Solstice HA does not support the use of the loopback file system (`lofs`) on the Solstice HA servers.
 - Since Solstice HA requires mirroring, the RAID 5 feature in the Solstice DiskSuite product is not supported.
 - The SPARCstorage Array NVRAM feature is not supported in Solstice HA configurations.
 - HA-NFS requires that all NFS client mounts be “hard” mounts.
 - Because Solstice DiskSuite does not support Sun Prestoserve™, it cannot be used in a Solstice HA configuration.

Hardware Installation Information



This chapter provides pointers to the appropriate hardware installation instructions for your SPARCcluster High Availability configuration. Use the following table to locate specific information in this chapter.

<i>SPARCcluster High Availability Hardware Installation</i>	<i>page 4-1</i>
<i>SPARCstation 4 Installation</i>	<i>page 4-2</i>
<i>SPARCcenter 2000 Installation</i>	<i>page 4-2</i>
<i>SPARCserver 1000 Installation</i>	<i>page 4-3</i>
<i>SPARCstorage Arrays Installation</i>	<i>page 4-3</i>
<i>Terminal Concentrator Installation</i>	<i>page 4-4</i>

4.1 SPARCcluster High Availability Hardware Installation

After planning your configuration using the information in Chapter 3, “Installation Planning,” you are ready to install the hardware. The documentation necessary to set up and install the SPARCcluster High Availability configuration is listed in Table 4-1.

Table 4-1 SPARCcluster High Availability Hardware Documentation

Title	Part Number
<i>SPARCcluster High Availability Server Hardware Planning and Installation</i>	802-3510
<i>SPARCcluster High Availability Server Service Manual</i>	802-3512

4.2 SPARCstation 4 Installation

The documentation necessary to set up and configure the SPARCstation 4 Administration Workstation is listed in Table 4-2.

Table 4-2 SPARCstation 4 Documentation

Title	Part Number
<i>SPARCstation 4 Hardware Setup Instructions</i>	802-1900
<i>SPARCstation 4 Installation Guide</i>	802-1523
<i>Twisted-Pair Ethernet Link Test (10Base-T)</i>	801-6184
<i>Specifications - SPARCstation 4 System</i>	802-1526
<i>SPARCstation 4 Software installation Note</i>	802-2052

4.3 SPARCcenter 2000 Installation

If your configuration includes SPARCcenter 2000s, you may need some or all of the documentation listed in Table 4-3.

Table 4-3 SPARCcenter 2000 Documentation

Title	Part Number
<i>SPARCcenter 2000 Binder Set</i>	825-1509
– <i>SPARCcenter 2000 Installation Manual</i>	800-6975
– <i>SPARCcenter 2000 Storage Devices User's Guide</i>	800-7009
– <i>SPARCcenter 2000 Service Manual</i>	801-2007
– <i>SPARCcenter 2000 Regulatory Compliance Manual</i>	801-3051

4.4 SPARCserver 1000 Installation

If your configuration includes SPARCserver 1000s, you may need some or all of the documentation listed in Table 4-4.

Table 4-4 SPARCserver 1000 Documentation

Title	Part Number
<i>SPARCserver 1000 Binder Set</i>	825-1725
– <i>SPARCserver 1000 System Installation Manual</i>	801-2893
– <i>SPARCserver 1000 Storage Devices User's Guide</i>	801-2198
– <i>SPARCserver 1000 System Service Manual</i>	801-2895
– <i>SPARCserver 1000 Regulatory Compliance Manual</i>	801-2892

4.5 SPARCstorage Arrays Installation

A listing of other documentation you may need to configure SPARCstorage Arrays is listed in Table 4-5.

Table 4-5 SPARCstorage Array Documentation

Title	Part Number
<i>SPARCstorage Array Model 100 Hardware Binder Set</i>	825-2513
– <i>SPARCstorage Array Model 100 Installation Manual</i>	801-2205
– <i>SPARCstorage Array Model 100 Service Manual</i>	801-2206
– <i>SPARCstorage Array Model 100 Compliance Manual</i>	801-7103

4.6 *Terminal Concentrator Installation*

A listing of the documentation needed to set up and configure the terminal concentrator is listed in Table 4-6.

Table 4-6 Terminal Concentrator Documentation

Title	Part Number
<i>SPARCcluster High Availability Server Hardware Planning and Installation</i>	802-3510
<i>Terminal Concentrator Binder Set</i>	

The procedure for setting up security on the terminal concentrator is documented in Appendix B, “Terminal Concentrator Security.”

Solaris 2.4 Installation



The instructions in this chapter must be used to install systems as part of a Solstice High Availability configuration.

Note – Solstice HA 1.0 is supported only on the Solaris 2.4 Hardware: 3/95 operating system. It will not run on earlier or later versions of Solaris.

Use the following table to locate specific information in this chapter.

<i>Installation Rules</i>	<i>page 5-1</i>
<i>Selecting the Size of the Root Slice</i>	<i>page 5-2</i>
<i>Sizing File Systems for Install</i>	<i>page 5-6</i>

5.1 Installation Rules

When you install Solaris 2.4 on Solstice HA servers, you must follow these general rules:

- Set up each of the Solstice HA systems as standalone machines.
- Install the same Solaris software cluster on both servers. For example, if the entire distribution is installed on one server, it must be installed on the other Solstice HA server.
- Use identical partitioning on the local disks of each of the Solstice HA servers where Solaris 2.4 is installed.

5.2 *Selecting the Size of the Root Slice*

The root slice on your local disk must be large enough to hold the following:

- The Solaris 2.4 system components
- Any components from Solstice DiskSuite, Solstice HA, and other third party software packages
- Data space for symbolic links in `/dev` for the SPARCstorage Arrays and Solstice DiskSuite metadevices

The root slice must also have enough inodes for the various files and directories as well as the device inodes in `/devices` and symbolic links in `/dev`. By default one inode is allocated per two Kbytes of file system capacity when `newfs(1M)` is run on the file system during Solaris installation.

The Solaris install process automatically determines the file system size required for the operating system components placed on the root slice. For Solaris 2.4 this is approximately 14 Mbytes (assuming separate files systems for `/opt`, `/var`, and `/usr`).

While the install process does not determine the number of inodes required, the number of Solaris inodes is approximately 3,000. These do not include SPARCstorage Array and Solstice DiskSuite inodes. The default ratio of inodes to storage (one inode per two Kbytes of UFS capacity) is sufficient for the Solaris root slice in configurations that do not have SPARCstorage Arrays or Solstice DiskSuite. The demand for inodes by the SPARCStorage Arrays and Solstice DiskSuite may exceed the number available. The worksheet in Table 5-1 will help you plan your root slice size. Table 5-4 tells you the required root file system sizes for various configurations.

Each disk partition requires two inodes in `/devices` (one inode for each raw and block device). Each partition also requires two symbolic links in `/dev` that point to the device inodes in `/devices`. Each symbolic link requires one inode and one fragsize block of disk storage (typically this is one Kbyte). Since each disk drive has eight slices you will need 32 inodes and 16 Kbytes on the root file system for each disk.

You should allocate resources for the maximum number of SPARCstorage Array chassis you will configure into your system (for the SPARCserver 2000 systems 20 SPARCstorage Arrays are supported, while the SPARCcenter 1000s

support eight SPARCstorage Arrays). While you may initially only partially populate your SSA chassis you should allocate enough inodes and space for a full complement of disks to allow for future growth and reconfiguration.

The SPARCstorage Array 100 chassis accommodates 30 drives, and so 960 inodes and 480 Kbytes should be allocated on the root file system for each chassis.

Each Solstice DiskSuite metadvice consumes two inodes (raw and block) in `/devices` and two inodes as well as two Kbytes in `/dev` (symbolic links to `/devices`). The default Solstice DiskSuite configuration has `md_nsets=4` and `nmd=128` (refer to the *Solstice DiskSuite 4.0 Administration Guide*) for a total of 512 metadvicees.

Note – The Solstice DiskSuite documentation refers to hot spare pools as metadvicees, but they do not consume inodes in `/dev` or `/devices`, so they do not enter the computation.

Under Solstice HA, a maximum Solstice DiskSuite configuration would have `md_nsets=3` and `nmd=1024` for a total of 3,072 metadvicees. Thus a maximum Solstice DiskSuite configuration would require 12,288 inodes and 6,144 Kbytes of disk storage on the root file system.

The worksheet and examples in Table 5-1, and Table 5-2 show the computation. The number of Solstice DiskSuite metadvicees is the product of the values for `nmd` and `md_nsets` that are found in the `/kernel/drv/md.conf` file.

Table 5-1 Root Slice Calculation Template

Components	Inode Calculations	Kbyte Calculations
SPARCstorage Array 100 with 30 disks number of SPARCstorage Arrays (SSA)	$SSAs * 960 = \text{___} \text{ inodes}$	$SSAs * 480 \text{ Kbytes} = \text{___} \text{ Kbytes}$
Solstice DiskSuite metadvicees (<code>md_nsets</code> times <code>nmd</code>)	$nmd * md_nsets * 4 = \text{___} \text{ inodes}$	$nmd * md_nsets * 2 = \text{___} \text{ Kbytes}$
Solaris 2.4	3,000 inodes	= 15,000 Kbytes
Totals	$total \text{ inodes}$	$\text{___} \text{ Kbytes}$
	$total \text{ inodes times } 2 \text{ Kbytes} =$ minimum UFS size $\text{___} \text{ Kbytes}$	

You should take the larger of Inode Calculations and Kbyte Calculations from the last line of Table 5-1 to use as your required root file system size. You also need to allow for additional free space.

The example in Table 5-2 shows how the template could be used in a typical configuration that has six SPARCstorage Arrays and 512 metadevices. In this example the entries are made in bold font.

Table 5-2 Example Root Slice Calculation Template

Components	Inode Calculations	Kbyte Calculations
SPARCstorage Array 100 with 30 disks	6 * 960 = 5,760 inodes	6 * 480 Kbytes = 2,880 Kbytes
Solstice DiskSuite metadevices (md_nsets times nmd)	4 * 128 * 4 = 2,048 inodes	4 * 128 * 2 Kbytes = 1,024 Kbytes
Solaris 2.4	3,000 inodes	= 15,000 Kbytes
Totals	10,808 inodes	18,904 Kbytes
	10,808 inodes times 2 Kbytes = minimum UFS size 21,616 Kbytes	

Table 5-3 shows a maximum Solstice HA configuration that has two SPARCserver 2000s and 20 SPARCstorage Arrays and uses the maximum number of Solstice DiskSuite metadevices.

Table 5-3 Example of a Root Slice for Maximum Solstice HA Configuration

Components	Inode Calculations	Kbyte Calculations
SPARCstorage Array 100 with 30 disks	20 * 960 = 19,200 inodes	20 * 480 Kbytes = 9,600 Kbytes
Solstice DiskSuite metadevices (md_nsets times nmd)	3 * 1024 * 4 = 12,288 inodes	3 * 1024 * 2 Kbytes = 6,144 Kbytes
Solaris 2.4	= 3,000 inodes	= 15,000 Kbytes
Totals	34,488 inodes	30,744 Kbytes
	34,488 inodes times 2 Kbytes = minimum UFS size 68,976 Kbytes	

In the example shown in Table 5-2 the inode and symbolic link requirements for the SPARCstorage Arrays and Solstice DiskSuite devices require a file system more than double the size of that needed for the data alone.

You may want to allocate additional capacity to the root slice to accommodate future growth and package installation. This is especially true in small configurations where the size dictated by inodes is roughly equal to that dictated by data storage requirements and so comparatively little free data space is available.

Table 5-4 shows the required root file system size. The table should be used with the maximum number of drives you ultimately expect to have in the configuration, not the number you have during initial configuration. Similarly the metadevice parameters (`md_nsets` and `nmd`) should reflect your ultimate needs. The table entries are of the form size (free space), where size is the amount of space required in megabytes for the root file system and the amount of free space after Solaris 2.4, Solstice HA, and other required software have been installed. The values in the table are based on the calculations outlined earlier in this section.



Caution – Remember that you cannot enlarge the size of the root file system without reinstalling the Solaris operating environment. So you should calculate your disk needs based on future growth of the Solstice HA configuration.

Table 5-4 Required Root File System Sizes and `md.conf` Values

Change <code>md.conf</code> File Values to:	<code>md_nsets=4</code> <code>nmd=128</code>	<code>md_nsets=3</code> <code>nmd=256</code>	<code>md_nsets=3</code> <code>nmd=512</code>	<code>md_nsets=3</code> <code>nmd=1024</code>
Number of Drives	root (free) in Mbytes	root (free) in Mbytes	root (free) in Mbytes	root (free) in Mbytes
30	17 (0)	17 (0)	20 (1.6)	32 (11)
60	17 (0)	18 (0)	22 (3)	34 (12)
90	18 (0)	18 (0)	24 (4.4)	36 (13)
120	18 (0)	20 (1.3)	26 (5.8)	38 (15)
150	20 (1.2)	22 (2.7)	28 (7.2)	40 (16)
180	22 (2.6)	24 (4.1)	30 (8.6)	42 (18)
210	23 (4.1)	25 (5.6)	31 (10)	43 (19)
240	25 (5.5)	27 (7)	33 (11)	45 (20)
270	27 (6.9)	29 (8.4)	35 (13)	47 (22)
300	29 (8.3)	31 (9.8)	37 (14)	49 (23)

Table 5-4 Required Root File System Sizes and md.conf Values

Change md.conf File Values to:	md_nsets=4 nmd=128	md_nsets=3 nmd=256	md_nsets=3 nmd=512	md_nsets=3 nmd=1024
Number of Drives	root (free) in Mbytes	root (free) in Mbytes	root (free) in Mbytes	root (free) in Mbytes
330	29 (9.7)	33 (11)	39 (16)	51 (25)
360	33 (11)	35 (13)	41 (17)	53 (26)
390	35 (12)	37 (14)	43 (18)	55 (27)
420	37 (14)	39 (15)	45 (20)	57 (29)
450	38 (15)	40 (17)	46 (21)	58 (30)
480	40 (17)	42 (18)	48 (23)	60 (32)
510	42 (18)	44 (20)	50 (24)	62 (33)
540	44 (20)	46 (21)	52 (26)	64 (35)
570	46 (21)	48 (22)	54 (27)	66 (36)
600	48 (22)	50 (24)	56 (28)	68 (37)

In Table 5-4 you must determine which of the columns you want to use based on the number of metadevice names needed in the largest diskset. You must configure both servers with the same size root file system and the same values entered in the `/kernel/drv/md.conf` file. Refer to Section 3.1.19, “Metadevice Naming and Creation,” on page 3-21 for help in determining which of the columns you should use.

5.3 Sizing File Systems for Install

The following table estimates the space required for the various file systems after Solaris2.4 Hardware: 3/95 release, the Solaris 2.4 Hardware: 3/95 patches, Solstice DiskSuite, Solstice HA and required patches have been installed. Note that the size of the root partition should be obtained from Table 5-4 for a given number of SPARCstorage Array drives and Solstice DiskSuite metadevices that you want to allocate.

Note – The Solstice HA space requirements in Table 5-5 include no free space. You should provide additional free space for future additions and short term needs. If you exhaust the free space you must reinstall the operating system (on both servers) to obtain additional free space. You should error on the side of extra free space; 20 percent or more is suggested. The `/var` file system will contain system messages and logs and can fill quickly, even more free space is recommended for `/var`.

In addition to the calculations in Table 5-5, you must allocate six Mbytes on slice 4 of the boot disk on each of the Solstice HA servers. This space is used for metadvice state database replicas and UFS logs.

Table 5-5 Sizing File Systems for Install

File System Name	Solstice HA Space Requirements	Your Additional Requirements	Total Space Needed
/	Value from Table 5-4	+ _____	= _____
/opt	110 Mbytes	+ _____	= _____
/usr	180 Mbytes	+ _____	= _____
/var	20 Mbytes	+ _____	= _____

▼ How to Install Solaris on the Administration Workstation

This section contains an example of the appropriate instructions for installing Solaris 2.4 on a Solstice HA system, using the character user interface (CUI) Solaris install program.

1. **Turn on the CD-ROM, the SPARCstation 4 monitor, and the SPARCstation 4.**
2. **Insert the Solaris 2.4 CD.**
3. **Power up the system.**
4. **Use the JumpStart™ default selections.**
JumpStart instructions are included in the *SPARC: Installing Solaris Software* manual.

5. The system will reboot.

When the system comes back up, you will be prompted for a root password.

6. Log in to the system as root.**7. Eject the Solstice CD.**

```
# eject
```

8. Insert the Solstice 2.4 CD.

By ejecting and reinserting the CD, Volume Management will recognize the CD.

9. When the installation completes, change to the directory on the CD-ROM to where the Solaris 2.4 patches reside. Refer to the Solaris 2.4 Latest News document for a listing of the patches on the Solaris 2.4 Hardware: 3/95 CD-ROM. Then enter the following command:

```
# cd /cdrom/cdrom0/s0/Patches
```

10. Install the Solaris 2.4 patches by entering the following command:

```
# ./install_patches

Installing Patches (using install_patches 1.25)

Installation will be logged in
/var/sadm/install_data/Patches_log

install_patches completed successfully
#
```

11. Reboot the system.

▼ How to Install the Terminal Concentrator

1. Configure the terminal concentrator.

Use the instructions in Chapter 11 of the *SPARCcluster High Availability Server Hardware Planning and Installation Manual* to configure the terminal concentrator.

Enter the IP address, subnet mask, and the network broadcast address for the terminal concentrator.

2. Enable terminal concentrator security.

At this point, the terminal concentrator is operating with limited security. Terminal concentrator configuration is protected by an administrative password (the terminal concentrator IP address) which is documented in the terminal concentrator manuals and widely known. Because the port access is unrestricted, any user who can connect to the terminal concentrator can also connect to the Solstice HA server console and has access to a root login or can get into OpenBoot command interpreter by sending a break signal.

Unless access to the public network serving the terminal concentrator is adequately restricted you will probably want to enhance the access security of the terminal concentrator.

Using the command line interface (CLI), you can set separate internal passwords for administrative access and port access, using either specific port numbers or using the port selector. All port access is protected using the single port password. A more complex approach that allows individual user and password access is possible using the security host feature. For more information on this approach see Appendix B, “Terminal Concentrator Security” or the *Terminal Concentrator Binder Set* delivered with your system.

To enable basic password security on your terminal concentrator open a `telnet(1)` connection and enter the following commands.

Note – There will be about a 30 delay after you enter the admin command `reset annex all`.

```
# telnet terminal_concentrator
...
Enter Annex port name or number: cli

Annex Command Line Interpreter * Copyright 1991 Xylogics, Inc.

annex: su
Password: terminal_concentrator_password
annex# admin
Annex administration MICRO-XL-UX R7.0.1, 8 ports
admin : set annex password new_terminal_concentrator_password
        You may need to reset the appropriate port, Annex subsystem
        or reboot the Annex for changes to take effect.
admin : set port port_password port_password
        You may need to reset the appropriate port, Annex subsystem
        or reboot the Annex for changes to take effect.
admin : set annex enable_security y
        You may need to reset the appropriate port, Annex subsystem
        or reboot the Annex for changes to take effect.
admin : reset annex all
admin : quit
annex# su
annex: hangup

Resetting line and disconnecting.

Connection closed by foreign host.
#
```

▼ How to Install Solaris on the Solstice HA Servers

1. Log in to the administration workstation as root.
2. Start OpenWindows.

```
# /usr/openwin/bin/openwin
```

3. Open two `shelltool(1)` windows on the administration workstation.
4. Open a `telnet` connection in each of the windows to one of the Solstice HA servers through the terminal concentrator. Follow each of the steps in this procedure in both windows.

Ports 2 and 3 of the terminal concentrator are connected to the respective server console serial ports. These two ports are numbered 5002 and 5003. Run the following command to open the `telnet` connection to a server in one of the windows.

```
# telnet terminal_concentrator_name 5002
```

Run the following command to open a `telnet` connection in the other window to the other server.

```
# telnet terminal_concentrator_name 5003
```

5. Power on the two Solstice HA servers.
You can view the power on self tests via the two `shelltool` windows.
6. If either server attempts to boot, you will want to send a break to enter the OpenBoot PROM.
To sent a break, you must first enter the `telnet` escape character `Control-]`.

```
^]  
telnet> send brk  
ok
```

7. Run the following commands to set the variables from the OpenBoot PROM.

```
ok set-defaults  
Setting NVRAM parameters to default values.  
ok setenv watchdog-reboot? true  
watchdog-reboot?= true  
ok
```

8. Insert a Solaris 2.4 Hardware: 3/95 CD-ROM in the CD of each Solstice HA servers.

9. In each of the two windows, boot the install image from the CD-ROM.
For example if you have inserted the Solaris 2.4 CD-ROM, enter:

```
ok boot cdrom
```

Some or all of the following system identification screens may not appear if the installation software can get the information from the name service.

10. Select option 8 in each of the shelltool windows.

If you are not using shelltool windows, select the appropriate terminal type. The screens used in this example may appear different depending on the terminal type you select.

```
What type of terminal are you using?
1) ANSI Standard CRT
2) DEC VT52
3) DEC VT100
4) Heathkit 19
5) Lear Siegler ADM31
6) PC Console
7) Sun Command Tool
8) Sun Workstation
9) Televideo 910
10) Televideo 925
11) Wyse Model 50
12) X Terminal Emulator (xterms)
13) Other
Type the number of your choice and press Return: 8
starting rpc services: rpcbind sysidnis ypbind done.
```


11. When the following screen is displayed, press F2.

```
- The Solaris Installation Program -----  
  
You are now interacting with the Solaris installation program.  
The program is divided into a series of short sections. At the  
end of each section, you will see a summary of the choices you've  
made, and be given the opportunity to make changes.  
  
As you work with the program, you will complete one or more of the  
following tasks:  
  
    1 - Identify peripheral devices  
    2 - Identify your system  
    3 - Install Solaris software  
  
About navigation...  
  
    - The mouse cannot be used  
  
    - If your keyboard does not have function keys, or they do  
      not respond, press ESC; the legend at the bottom of the  
      screen will change to show the ESC keys to use for navigation.  
-----  
    F2_Continue    F6_Help
```

12. If the following screen displays, press F2. Otherwise, skip to Step 15.

```
- Identify This System -----  
  
On the next screens, you must identify this system as networked or  
non-networked, and set the default time zone and date/time.  
  
If this system is networked, the software will try to find the  
information it needs to identify your system; you will be prompted  
to supply any information it cannot find.  
  
    > To begin identifying this system, press F2.  
  
-----  
    F2_Continue    F6_Help
```

13. Answer any questions about the network configuration.

The types of questions asked will depend on the name service you are using at your site. If the installation software can find all of the network information you will not see any questions.

14. Answer any questions about the time zone and date.

15. If the following screen is displayed, press F4 to perform an initial install. Otherwise, go to the next step.



Caution – The instructions in this chapter are for initial install. Never choose the “upgrade option” on a Solstice HA system.

```
- Upgrade System? -----  
  
This system is upgradable. Choosing the upgrade option means any  
bundled Solaris software will be updated to the new release, and  
as many local modifications as possible will be saved.  
  
While your system is upgradable, you can choose the initial  
option; however, files on your disk will be overwritten and data  
will be lost.  
  
      CAUTION: If you choose the upgrade option, it is especially  
important to back up your system. However, backing up is also  
recommended for the initial option if there is any data on  
the  
      disk that you want to save.  
  
> To start the upgrade option, choose F2.  
  
> To start the initial option, choose F4.  
  
-----  
      F2_Upgrade    F4_Initial    F5_Exit    F6_Help
```

16. When the following screen is displayed, press F2.

```
- Install Solaris Software - Initial -----  
  
The screens that follow let you tailor how Solaris is installed  
on your system. The main tasks you must perform to install Solaris  
software on your system are:  
  
    - Selecting a system type  
    - Selecting Solaris software  
    - Selecting disks to hold the software you've selected  
    - Specifying how files systems are laid out on the disks  
  
After completing these tasks, a summary of your selections  
(called  
a profile) is displayed. You can go back and change selections as  
many times as you like before starting to install Solaris.  
  
> To continue, press F2.  
  
-----  
F2_Continue    F3_Go Back    F5_Exit    F6_Help
```

17. Select the Standalone system type then press F2.

```
- System Type -----  
  
On this screen you must specify one of the following system types.  
A system type determines where a system will get its directories  
and file systems, and whether it provides portions of Solaris  
software to other systems.  
  
> To make a selection, use the arrow keys to highlight the option  
and press Return to mark it [X].  
  
                [X] Standalone  
                [ ] Server  
                [ ] Dataless Client  
  
-----  
F2_Continue    F3_Go Back    F5_Exit    F6_Help
```

18. Select the Entire Distribution plus OEM support then press F2.

```

- Software -----
On this screen you must select the Solaris software to install on
your system. The software groups shown below are bundled and
defined by SunSoft. You can accept the default [X], or select
another.

NOTE: After selecting a software group, you can add or
remove software from it by selecting F4 to customize.
However, this requires understanding of software
dependencies and how Solaris software is packaged.

> To make a selection, use the arrow keys to highlight the option
and press Return to mark it [X]; then press F4 to customize or
F2 to continue.

[X] Entire Distribution plus OEM support .. 448.33 MB (F4 to
Customize)
[ ] Entire Distribution ..... 448.27 MB
[ ] Developer System Support ..... 360.24 MB
[ ] End User System Support ..... 267.21 MB
[ ] Core System Support ..... 164.50 MB

-----
F2_Continue   F3_Go Back   F4_Customize  F5_Exit     F6_Help

```

Optionally, you may choose to only install the software required to run Solstice HA and Solstice DiskSuite. The required software includes the End User Support software group in addition to the SUNWmfrun, SUNWssadv, SUNWssaop, and SUNWssamn packages.

If you are planning to use HA-ORACLE, you must install software required for Solstice HA in addition to the SUNWbtool, SUNwsprot, SUNwarc, and SUNwlibm packages.



Caution – Use caution when you are partitioning the boot disk if you choose to load only the minimally required parts of Solaris 2.4. If you fail to leave adequate space for future additions to the software you may need to repartition the disk and reload the Solaris distribution.

19. Select the local disk where Solaris 2.4 will be installed, along with the other local disk where you plan to install the ORACLE distribution if you are planning to use HA-ORACLE.



Caution – The bootdrive you select on one host must be the same size and have the same device number on the other server. Otherwise, when you invoke `hasetup(1M)`, data could be lost and the metadevice state database replicas may not be set up correctly.

```

- Disks -----

On this screen you must select the disks for installing Solaris
software. Start by looking at the Suggested Minimum field; this
value is the approximate space needed to install the software
you've selected. Keep selecting disks until the Total Selected
value exceeds the Suggested Minimum value.

> To make selections, use the arrow keys to highlight options
and press Return to mark them [X].

          Disk Device (Size)          Available Space
          =====
-   [X] c0t0d0 (1002 MB) bootdrive    1002 MB
|   [X] c0t1d0 (1002 MB)              1002 MB
|   [ ] c1t0d0 (1002 MB)              1002 MB
|   [ ] c1t0d1 (1002 MB)              1002 MB
|   [ ] c1t0d2 (1002 MB)              1002 MB
v   [ ] c1t0d3 (1002 MB)              1002 MB

                                Total Selected:  1002 MB
-----

          F2_Continue    F3_Go Back    F5_Exit    F6_Help
  
```



Caution – You must select the correct local disk. In Chapter 3, “Installation Planning,” you decided which local disk would contain Solaris 2.4. You must select that disk now. On the previous screen, `c0t0d0` was selected as the disk that will contain Solaris 2.4 (that is, the bootdrive).

20. If the following screen is displayed, press F2.

```
- Preserve Data? -----  
  
At least one of the disks you've selected has file systems or  
unnamed slices that you may want to save.  
  
> To save file systems or unnamed slices, press F4.  
  
> To allow current file systems and unnamed slices to be  
   overwritten when Solaris software is installed, press F2.  
  
-----  
F2_Continue   F3_Go Back   F4_Preserve   F5_Exit   F6_Help
```

21. Select the auto-layout feature by pressing F2.

```
- Automatically Layout File Systems? -----  
  
Do you want to use the auto-layout feature to automatically layout  
file systems on your disks? Manually laying out file systems on  
disks requires advanced system administration skills.  
  
> To use the auto-layout feature, press F2.  
  
> To manually layout file systems, press F4.  
  
-----  
F2_Auto Layout   F3_Go Back   F4_Manual Layout   F5_Exit   F6_Help
```

22. On the Auto-Layout File Systems screen, select the /opt and /var file systems in addition to the default selections for automatic layout. Then press F2.

```
- Auto-layout File Systems -----  
  
On this screen you must select all the file system you want  
auto-layout to create, or accept the defaults shown.  
  
NOTE: For small disks (less than 200 MB), it may be necessary  
for auto-layout to break up some of the file systems you  
request into smaller file systems to fit the available disk  
space. So, after auto-layout completes, you may find file  
systems in the layout that you did not select from the list  
below.  
  
File Systems for Auto-layout  
=====
```

<input checked="" type="checkbox"/>	/
<input checked="" type="checkbox"/>	swap
<input checked="" type="checkbox"/>	/usr
<input checked="" type="checkbox"/>	/opt
<input type="checkbox"/>	/usr/openwin
<input checked="" type="checkbox"/>	/var

```
-----  
F2_Continue    F5_Cancel    F6_Help
```

Note – Only select auto-layout for /opt and /var. Do not select auto-layout for /usr/openwin, because Slice 4 will be used for another purpose.

23. When the following screen is displayed, press F4 to customize the disk layout.

```
- File System and Disk Layout -----  
  
The summary below is your current file system and disk layout,  
based on the information you've supplied.  
  
NOTE: If you press F4 to customize, you should understand file  
systems, their intended purpose on the disk, and how changing  
them may affect the operation of the system.  
  
> To accept the layout shown, press F2.  
  
> To customize the layout, press F4.  
  
File system/Mount point          Disk/Slice          Size  
=====
```

-	/	c0t0d0s0	14 MB
	swap	c0t0d0s1	128 MB
	/usr	c0t0d0s6	227 MB
	/opt	c0t0d0s5	173 MB
	/var	c0t0d0s3	110 MB
	/export/home	c0t0d0s7	346 MB
v	/export/home0	c0t1d0s7	1002 MB

```
-----  
F2_Continue  F3_Go Back  F4_Customize  F5_Exit  F6_Help
```

24. Select the bootdrive (c0t0d0) to customize by pressing F4.

```
Select Disk to Customize

> To select a disk, use the arrow keys to highlight it,
and press F4.

Disk      (Size)                Configured File Systems
=====
c0t0d0    (1002 MB)    /, swap, /var, /opt, /usr, /export/home
c0t1d0    (1002 MB)    /export/home0

-----
F2_OK      F4_Customize  F6_Help
```

25. Scroll through the list of file systems and remove the 346 Mbytes allocated for /export/home. Also remove the /export/home mount point.
You can use either the arrow keys on the key pad or the Tab key to move the cursor around the Customize Disk screen.

Note – The slice sizes entered in the following steps are for an example configuration. Refer to Table 5-5 on page 5-7 to calculate the sizes for your configuration.

```
- Customize Disk: c0t0d0 -----
Entry:                Recommended:      MB      Minimum:      MB
=====
Slice  Mount Point                Size (MB)
  0    /                          14
  1    swap                       128
  2    overlap                    1002
  3    /var                       110
  4                                0
  5    /opt                       173
  6    /usr                       227
  7                                0
=====
                          Capacity:      1002 MB
                          Allocated:      652 MB
                          Rounding Error:  4 MB
                          Free:          346 MB
-----
F2_OK      F4_Options  F5_Cancel  F6_Help
```

26. Refer to Table 5-5 for the appropriate amount of space to allocate to the root (/) file system.
 In this example 35 Mbytes is allocated.

```

- Customize Disk: c0t0d0 -----
Entry: /                Recommended:  14 MB      Minimum:    12 MB

=====
Slice  Mount Point          Size (MB)
  0    /                    35
  1    swap                 128
  2    overlap             1002
  3    /var                 110
  4                    0
  5    /opt                 173
  6    /usr                 227
  7                    0

=====
                          Capacity:    1002 MB
                          Allocated:    673 MB
                          Rounding Error: 3 MB
                          Free:        326 MB

-----
F2_OK      F4_Options    F5_Cancel    F6_Help
  
```

27. Refer to Table 5-5 for the appropriate amount of space to allocate to the /var file system.

```
- Customize Disk: c0t0d0 -----
Entry: /usr           Recommended: 227 MB       Minimum: 193 MB

=====
Slice  Mount Point           Size (MB)
  0    /                   35
  1    swap                128
  2    overlap            1002
  3    /var                130
  4
  5    /opt               173
  6    /usr               227
  7
=====
                          Capacity:      1002 MB
                          Allocated:      693 MB
                          Rounding Error:   3 MB
                          Free:            306 MB
-----
F2_OK      F4_Options    F5_Cancel   F6_Help
```

28. Place 6 Mbytes in slice 4 for metadvice state databases and local UFS logs. Do not give this slice a mount point.

Note – The amount of free space may vary by a small amount on different disks.

```

- Customize Disk: c0t0d0 -----
Entry:                Recommended:      MB      Minimum:      MB
=====
  Slice  Mount Point                Size (MB)
    0    /                          35
    1    swap                        128
    2    overlap                      1002
    3    /var                         130
    4                                6
    5    /opt                         173
    6    /usr                         227
    7                                0
=====
                          Capacity:      1002 MB
                          Allocated:      699 MB
                          Rounding Error:  3 MB
                          Free:           300 MB
-----
F2_OK      F4_Options    F5_Cancel  F6_Help
  
```

29. Refer to Table 5-5 for the appropriate amount of space to allocate to the /opt file system.

```
- Customize Disk: c0t0d0 -----
Entry: /opt           Recommended: 173 MB       Minimum: 150 MB

=====
Slice  Mount Point          Size (MB)
  0    /                  35
  1    swap               128
  2    overlap           1002
  3    /var              130
  4
  5    /opt              283
  6    /usr              227
  7
=====
                          Capacity:      1002 MB
                          Allocated:      809 MB
                          Rounding Error:   3 MB
                          Free:            190 MB
-----
F2_OK      F4_Options    F5_Cancel  F6_Help
```

30. Refer to Table 5-5 for the appropriate amount of space to allocate to the /usr file system. After you have allocated space for every file system, press F2.

```

- Customize Disk: c0t0d0 -----
Entry: /usr           Recommended:  227 MB       Minimum:  193 MB

=====
Slice  Mount Point           Size (MB)
  0    /                   35
  1    swap                128
  2    overlap            1002
  3    /var               130
  4
  5    /opt               283
  6    /usr               418
  7
=====
                          Capacity:      1002 MB
                          Allocated:      1000 MB
                          Rounding Error:   2 MB
                          Free:            0 MB
-----
F2_OK      F4_Options  F5_Cancel  F6_Help

```


**31. (Optional) Press F4 to customize the second local disk for ORACLE.
Otherwise, press F2 and skip to Step 34.**

```
- Select Disk to Customize -----  
  
> To select a disk, use the arrow keys to highlight it,  
and press F4.  
  
          Disk (Size)                Configured File Systems  
=====
```

c0t0d0	(1002 MB)	/, swap, /var, /opt, /usr
c0t1d0	(1002 MB)	/export/home0

```
-----  
F2_OK      F4_Customize    F6_Help
```

32. (Optional) If you are going to use the HA-ORACLE data service, the entire second local disk can be assigned to the ORACLE distribution. After changing the /export/home0 mount point to /oracle, press F2.

```

- Customize Disk: c0t1d0 -----
Entry:                Recommended:      MB      Minimum:      MB

=====
  Slice  Mount Point                Size (MB)
    0
    1
    2  overlap                    1002
    3
    4
    5
    6
    7  /oracle                      1002
=====
                                Capacity:      1002 MB
                                Allocated:     1002 MB
                                Free:           0 MB

-----
F2_OK      F4_Options    F5_Cancel    F6_Help

```

33. (Optional) Press F2 to continue.

```
- Select Disk to Customize -----  
  
> To select a disk, use the arrow keys to highlight it,  
and press F4.  
  
          Disk (Size)                Configured File Systems  
=====
```

c0t0d0	(1002 MB)	/, swap, /var, /opt, /usr
c0t1d0	(1002 MB)	/oracle

```
-----  
F2_OK    F4_Customize    F6_Help
```

34. Verify the file system and disk layout for the two local disks then Press F2 to continue with Solaris installation.

```

- File System and Disk Layout -----

The summary below is your current file system and disk layout,
based on the information you've supplied.

NOTE: If you press F4 to customize, you should understand file
systems, their intended purpose on the disk, and how changing
them may affect the operation of the system.

> To accept the layout shown, press F2.

> To customize the layout, press F4.

File system/Mount point           Disk/Slice           Size
=====
/                                 c0t0d0s0             35 MB
swap                              c0t0d0s1             128 MB
/usr                              c0t0d0s6             407 MB
/opt                              c0t0d0s5             283 MB
/var                              c0t0d0s3             130 MB
/oracle                           c0t1d0s7             1002 MB

-----
-   F2_Continue   F3_Go Back   F4_Customize   F5_Exit   F6_Help

```

35. Press F2 to bypass mounting remote file systems.

```
- Mount Remote File Systems? -----  
  
Do you want to mount software from a remote file server? This may  
be necessary if you had to remove software because of disk space  
problems.  
  
> To go to the Mount Remote File System screen, press F4.  
  
> To bypass mounting remote file systems, press F2.  
  
-----  
F2_Continue    F3_Go Back    F4_Remote Mounts    F6_Help
```

36. Examine the Profile screen. To change any of the installation information, press F4. To use the file system and disk layout as presented, press F2.

```
- Profile -----  
  
The information below is your profile for installing Solaris  
software. It is a summary of the choices you've made on previous  
screens.  
  
> To go back to the beginning to make changes, Press F4.  
  
=====
```

System Type:	Standalone
Software:	Solaris 2.4, Entire Distribution plus OEM's

```
File System and Disk Layout: /                c0t0d0s0    35 MB  
                               swap           c0t0d0s1    128 MB  
                               /usr           c0t0d0s6    407 MB  
                               /opt          c0t0d0s5    283 MB  
                               /var          c0t0d0s3    130 MB  
                               /oracle       c0t1d0s7    1002 MB  
  
-----  
F2_Continue    F4_Change    F5_Exit     F6_Help
```

37. Select Do Not Reboot when the following screen is displayed. Press F2 to continue.

```
- Begin Installing Solaris -----  
  
By default, your system will reboot after Solaris is installed.  
However, you can bypass rebooting and access a UNIX shell to make  
changes before rebooting.  
  
> To begin installing Solaris, press F2; to go back and make  
changes press F5.  
  
      [ ] Reboot  
      [X] Do not reboot  
  
-----  
      F2_Begin Installation      F5_Cancel
```

38. When Solaris 2.4 Hardware: 3/95 installation begins, you will see the following types of messages:

```
Preparing system to install Solaris. Please wait.  
  
Setting up disk c0t0d0:  
  - Creating Solaris disk label (VTOC)  
  
Setting up disk c0t1d0:  
  - Creating Solaris disk label (VTOC)  
  
Creating and checking UFS file systems:  
  - Creating / (c0t0d0s0)  
  - Creating /var (c0t0d0s3)  
  - Creating /opt (c0t0d0s5)  
  - Creating /usr (c0t0d0s6)  
  - Creating /oracle (c0t1d0s7)  
  
Beginning Solaris package installation...
```

39. When the installation completes, change to the directory on the CD-ROM to where the Solaris 2.4 Hardware: 3/95 patches reside.

Refer to the Solaris 2.4 Latest News document for a listing of the patches on the Solaris 2.4 Hardware: 3/95 CD-ROM. Then enter the following command:

```
# cd /cdrom/Patches
```

40. Install the Solaris 2.4 Hardware: 3/95 patches by entering the following command:

```
# ./install_patches /a

Installing Patches (using install_patches 1.25)

Installation will be logged in
                        /a/var/sadm/install_data/Patches_log

install_patches completed successfully
#
```

41. Edit the /a/etc/system file on each of the servers.

You must edit the /a/etc/system file to exclude the loopback file system. The following two lines should be added to the end of the /a/etc/system file on each of the servers.

```
* Solstice HA excludes the loopback file system
exclude: lofs
```

42. (Optional) Make other changes to the /a/etc/system file on each of the servers.

If you are planning to use the HA-ORACLE data service, you may need to make additional changes to the /a/etc/system file. These changes could be made at this time. Consult your ORACLE documentation for the appropriate changes to this file.

43. Reboot both Solstice HA servers.

When Solaris installation completes, the system will reboot with the following sequence.

```
# reboot  
...
```

44. Log in to each of the Solstice HA servers as root and eject the CD using the following command:

```
# eject
```

45. Turn to Chapter 6, “Installing the Packages and Patches,” and follow the procedures given there.

Installing the Packages and Patches



After the Solaris 2.4 software is installed, you are ready to install the Solstice High Availability packages and patches, Solstice DiskSuite packages, and SPARCstorage Array patches. The procedures in this chapter must be followed on both systems in the Solstice HA configuration.

Use the following table to locate specific information in this chapter.

<i>Packages to Be Installed</i>	<i>page 6-1</i>
<i>Path Requirements</i>	<i>page 6-30</i>
<i>Viewing AnswerBook Online Documentation</i>	<i>page 6-31</i>

6.1 Packages to Be Installed

Packages and patches must be installed in the correct sequence. All the packages and the patches must be installed on each of the Solstice HA servers. The correct sequence of installation is:

1. The SunVTS validation test suite package (SUNWvts)
2. All DiskSuite packages from the Solstice DiskSuite 4.0 CD (in any order)
3. All packages from the Solstice High Availability 1.0 CD. The sequence for installation is:
 - a. SUNWff – SPARCcluster FailFast Device Driver
 - b. SUNWcmm – SPARCcluster Membership Monitor

- c. SUNWhagen – Solstice High Availability
 - d. SUNWhanfs – Solstice High Availability – HA-NFS component
 - e. SUNWabha – Solstice High Availability AnswerBook
4. All patches on the Solstice High Availability 1.0 CD (in any order)
 5. (Optional) Load the HA-ORACLE package (SUNWhaor)

▼ How to Add the Required Packages and Patches

You must repeat the step-by-step instructions in this procedure on each of the systems in the Solstice HA configuration.

To install all packages and patches follow these steps:

- 1. Open two `shelltool(1)` windows on the Administration Workstation (SPARCstation 4). Open a `telnet(1)` session in each of the two windows and log in to each of the Solstice HA servers as root.**
- 2. Insert the Solstice High Availability 1.0 CD into the CD-ROM drive.**
There is a brief delay while Solaris scans the CD and mounts the proper file systems.

3. Enter the following command to install the SunVTS package. Enter y at any prompts about changing modes on directories.

pkgadd installs the SunVTS files in the /opt/SUNWvts directory.

```
# pkgadd -d /cdrom/cdrom0 SUNWvts
Processing package instance <SUNWvts> from </cdrom/cdrom0>

Online Validation Test Suite
(sparc) 1.8.0,REV=0.1.6
    Copyright 1994 Sun Microsystems, Inc. All Rights Reserved.
    Printed in the United States of America.
2550 Garcia Avenue, Mountain View, California, 94043-1100 U.S.A.

... (miscellaneous copyright information)...

This product incorporates technology used under license from
Fulcrum
Technologies, Inc.
Using </opt> as the package base directory.
## Processing package information.
## Processing system information.
## Verifying package dependencies.
## Verifying disk space requirements.
## Checking for conflicts with packages already installed.
## Checking for setuid/setgid programs.

Installing Online Validation Test Suite as <SUNWvts>

## Installing part 1 of 1.
/opt/SUNWvts/bin/.customtest
/opt/SUNWvts/bin/57fonts.300
/opt/SUNWvts/bin/57fonts.400
/opt/SUNWvts/bin/audio
/opt/SUNWvts/bin/bpptest
/opt/SUNWvts/bin/cdtest
/opt/SUNWvts/bin/cg14test
/opt/SUNWvts/bin/cg6
/opt/SUNWvts/bin/fbtest
/opt/SUNWvts/bin/ffb_checksum.data
```

(continued on following page)

```
/opt/SUNWvts/bin/ffbtst
/opt/SUNWvts/bin/ffbtst.data
/opt/SUNWvts/bin/fputest
/opt/SUNWvts/bin/fstest
/opt/SUNWvts/bin/isdntest
/opt/SUNWvts/bin/leotest
/opt/SUNWvts/bin/leotest.data
/opt/SUNWvts/bin/lpvitest
/opt/SUNWvts/bin/mpstest
/opt/SUNWvts/bin/music.au
/opt/SUNWvts/bin/netttest
/opt/SUNWvts/bin/pcsertest
/opt/SUNWvts/bin/plntest
/opt/SUNWvts/bin/pmem
/opt/SUNWvts/bin/pstest
/opt/SUNWvts/bin/rawtest
/opt/SUNWvts/bin/spdtest
/opt/SUNWvts/bin/spif
/opt/SUNWvts/bin/sptest
/opt/SUNWvts/bin/sunbuttons
/opt/SUNWvts/bin/sundials
/opt/SUNWvts/bin/sunlink
/opt/SUNWvts/bin/sunvts
/opt/SUNWvts/bin/sxtest
/opt/SUNWvts/bin/tapetest
/opt/SUNWvts/bin/tcxttest
/opt/SUNWvts/bin/test_name
/opt/SUNWvts/bin/vmem
/opt/SUNWvts/bin/vtsinfo.so
/opt/SUNWvts/bin/vtsk
/opt/SUNWvts/bin/vtsprobe
/opt/SUNWvts/bin/vtstty
/opt/SUNWvts/bin/vtstty.help
/opt/SUNWvts/bin/vtsui
/opt/SUNWvts/bin/xbttest
[ verifying class <none> ]

Installation of <SUNWvts> was successful.
```

4. Eject the Solstice HA 1.0 CD.

5. Turn to Chapter 11 of the *SPARCcluster High Availability Server Hardware Planning and Installation* manual for instructions on running the SunVTS diagnostic tools.

6. After completion of the SunVTS tests, remove the SUNWvts package.

```
# pkgrm SUNWvts
The following package is currently installed:
  SUNWvts           Online Validation Test Suite
                   (sparc) 1.8.0,REV=0.1.6

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

7. Insert the Solstice DiskSuite 4.0 CD into the CD-ROM drive.

8. Enter the following command to install the three Solstice DiskSuite packages. Enter y at any prompts about changing modes on directories. pkgadd installs the DiskSuite files in the /usr/opt/SUNWmd directory, as well as in other system directories.

```
# pkgadd -d /cdrom/cdrom0 SUNWabmd SUNWmd SUNWmdg

Processing package instance <SUNWabmd> from </cdrom/cdrom0>
DiskSuite 4.0 AnswerBook
(all) 68.1.5
    Copyright 1995 Sun Microsystems, Inc. All Rights Reserved.
    Manufactured in the United States of America.
    2550 Garcia Avenue, Mountain View, California, 94043-1100 U.S.A.
... (miscellaneous copyright information)...

This product incorporates technology used under license from Fulcrum
Technologies, Inc.
Using </opt> as the package base directory.

The installation options are as follows:
Option: Description:
-----
1. nil:    less than 1 Megabyte disk space required [slowest performance].
2. heavy:  14.21 Megabytes disk space required [best performance].

Note: If the install option which you choose below fails
      due to lack of space, try another location, or
      choose a lower install option number.

      (continued on following page)
```

```
Enter the number of an installation option from the list above (1 or 2).
Select an installation option: 2
Installation option: heavy selected.
```

The next request for input asks you to specify the parent directory of AnswerBook. Make sure to choose a parent directory on a file system big enough to accommodate all the files to be moved for the INSTALL OPTION you selected.

```
Specify the parent of the AnswerBook home directory: /opt
For the heavy option all files will be placed under /opt/SUNWabmd.
## Processing package information.
## Processing system information.
## Verifying package dependencies.
## Verifying disk space requirements.
## Checking for conflicts with packages already installed.
## Checking for setuid/setgid programs.
```

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 this package [y,n,?] y
```

```
Installing DiskSuite 4.0 AnswerBook as <SUNWabmd>
```

```
## Installing part 1 of 1.
/opt/SUNWabmd/index
/opt/SUNWabmd/index/Keys
/opt/SUNWabmd <implied directory>
/opt/SUNWabmd/index <implied directory>
/opt/SUNWabmd/index/index.cat
/opt/SUNWabmd/index/index.cfg
/opt/SUNWabmd/index/index.cix
/opt/SUNWabmd/index/index.dct
/opt/SUNWabmd/index/index.log
/opt/SUNWabmd/index/index.ref
[ verifying class <Index> ]
/opt/SUNWabmd/toc/DISKSUITEADMIN.ind
/opt/SUNWabmd/toc/DISKSUITEADMIN.lock
/opt/SUNWabmd/toc/DISKSUITEADMIN.rec
/opt/SUNWabmd/toc/DISKSUITEUG.ind
/opt/SUNWabmd/toc/DISKSUITEUG.lock
/opt/SUNWabmd/toc/DISKSUITEUG.rec
```

(continued on following page)

```
/opt/SUNWabmd/toc/SUNWab_68_1.ind
/opt/SUNWabmd/toc/SUNWab_68_1.lock
/opt/SUNWabmd/toc/SUNWab_68_1.rec
[ verifying class <ContentsDB> ]
/opt/SUNWabmd/ps/DISKSUITEADMIN/01.Introduction
/opt/SUNWabmd/ps/DISKSUITEADMIN/02.Installation_and_Setup
/opt/SUNWabmd/ps/DISKSUITEADMIN/03.Overview_of_Solstice_DiskSuite
/opt/SUNWabmd/ps/DISKSUITEADMIN/04.Concatenating_and_Striping
/opt/SUNWabmd/ps/DISKSUITEADMIN/05.Mirroring
/opt/SUNWabmd/ps/DISKSUITEADMIN/06.UFS_Logging
/opt/SUNWabmd/ps/DISKSUITEADMIN/07.Hot_Spares
/opt/SUNWabmd/ps/DISKSUITEADMIN/08.Disksets
/opt/SUNWabmd/ps/DISKSUITEADMIN/09.RAID_Devices
/opt/SUNWabmd/ps/DISKSUITEADMIN/10.State_Database_Replicas
/opt/SUNWabmd/ps/DISKSUITEADMIN/11.Expanding_a_File_System
/opt/SUNWabmd/ps/DISKSUITEADMIN/12.Configuration_Guidelines
/opt/SUNWabmd/ps/DISKSUITEADMIN/A.Solstice_DiskSuite_Files
/opt/SUNWabmd/ps/DISKSUITEADMIN/B.Solstice_DiskSuite_Messages
/opt/SUNWabmd/ps/DISKSUITEADMIN/C.Recovery_From_Failed_Boots
/opt/SUNWabmd/ps/DISKSUITEADMIN/Contents
/opt/SUNWabmd/ps/DISKSUITEADMIN/Cover
/opt/SUNWabmd/ps/DISKSUITEADMIN/Credits
/opt/SUNWabmd/ps/DISKSUITEADMIN/D.Upgrading_to_Other_Solaris_Versions
/opt/SUNWabmd/ps/DISKSUITEADMIN/E.Using_Solstice_DiskSuite_4.0_with_the
/opt/SUNWabmd/ps/DISKSUITEADMIN/F.Man_Pages
/opt/SUNWabmd/ps/DISKSUITEADMIN/Figures
/opt/SUNWabmd/ps/DISKSUITEADMIN/Index
/opt/SUNWabmd/ps/DISKSUITEADMIN/Preface
/opt/SUNWabmd/ps/DISKSUITEADMIN/files_to_print
/opt/SUNWabmd/ps/DISKSUITEUG/01.Introduction
/opt/SUNWabmd/ps/DISKSUITEUG/02.Installation_and_Setup
/opt/SUNWabmd/ps/DISKSUITEUG/03.Overview_of_DiskSuite_Tool
/opt/SUNWabmd/ps/DISKSUITEUG/04.Concatenating_and_Striping
/opt/SUNWabmd/ps/DISKSUITEUG/05.Mirroring
/opt/SUNWabmd/ps/DISKSUITEUG/06.UFS_Logging
/opt/SUNWabmd/ps/DISKSUITEUG/07.Hot_Spares
/opt/SUNWabmd/ps/DISKSUITEUG/08.Administering_Metadevices_in_Disksets
/opt/SUNWabmd/ps/DISKSUITEUG/09.RAID_Devices
/opt/SUNWabmd/ps/DISKSUITEUG/10.State_Database_Replicas
/opt/SUNWabmd/ps/DISKSUITEUG/11.Expanding_a_File_System
/opt/SUNWabmd/ps/DISKSUITEUG/A.Example_Use_of_DiskSuite_Tool
/opt/SUNWabmd/ps/DISKSUITEUG/B.DiskSuite_Tool_Messages
/opt/SUNWabmd/ps/DISKSUITEUG/C.Man_Page
```

(continued on following page)

```
/opt/SUNWabmd/ps/DISKSUITEUG/Contents
/opt/SUNWabmd/ps/DISKSUITEUG/Cover
/opt/SUNWabmd/ps/DISKSUITEUG/Credits
/opt/SUNWabmd/ps/DISKSUITEUG/figures
/opt/SUNWabmd/ps/DISKSUITEUG/Index
/opt/SUNWabmd/ps/DISKSUITEUG/Preface
/opt/SUNWabmd/ps/DISKSUITEUG/Tables
/opt/SUNWabmd/ps/DISKSUITEUG/files_to_print
[ verifying class <PostScript> ]
## Executing postinstall script.

Installation of <SUNWabmd> was successful.

Processing package instance <SUNWmd> from </cdrom/cdrom0>

Solstice DiskSuite
(all) 4.0,REV=1.0

        Copyright (c) 1995 Sun Microsystems, Inc. All Rights Reserved.
        Manufactured in the United States of America.
2550 Garcia Avenue, Mountain View, California, 94043-1100 U.S.A.

... (miscellaneous copyright information)...

Using </> as the package base directory.
## Processing package information.
## Processing system information.
    10 package pathnames are already properly installed.
## Verifying package dependencies.
## Verifying disk space requirements.
## Checking for conflicts with packages already installed.
## Checking for setuid/setgid programs.

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 this package [y,n,?] y

Installing Solstice DiskSuite as <SUNWmd>

## Executing preinstall script.
## Installing part 1 of 1.
```

(continued on following page)


```
/etc/init.d/SUNWmd.init
/etc/init.d/SUNWmd.sync
/etc/opt/SUNWmd/lock
/etc/rc2.d/S95SUNWmd.sync <symbolic link>
/etc/rcS.d/S35SUNWmd.init <symbolic link>
/usr/opt/SUNWmd/etc <symbolic link>
/usr/opt/SUNWmd/locale/C/LC_MESSAGES/SUNWmd.po
/usr/opt/SUNWmd/man/man1m/growfs.1m
/usr/opt/SUNWmd/man/man1m/metaclear.1m
/usr/opt/SUNWmd/man/man1m/metadb.1m
/usr/opt/SUNWmd/man/man1m/metadetch.1m
/usr/opt/SUNWmd/man/man1m/metahs.1m
/usr/opt/SUNWmd/man/man1m/metainit.1m
/usr/opt/SUNWmd/man/man1m/metaoffline.1m
/usr/opt/SUNWmd/man/man1m/metaonline.1m
/usr/opt/SUNWmd/man/man1m/metaparam.1m
/usr/opt/SUNWmd/man/man1m/metareplace.1m
/usr/opt/SUNWmd/man/man1m/metaroot.1m
/usr/opt/SUNWmd/man/man1m/metastat.1m
/usr/opt/SUNWmd/man/man1m/metasync.1m
/usr/opt/SUNWmd/man/man1m/metattach.1m
/usr/opt/SUNWmd/man/man4/md.cf.4
/usr/opt/SUNWmd/man/man4/md.tab.4
/usr/opt/SUNWmd/man/man4/mddb.cf.4
/usr/opt/SUNWmd/man/man7/md.7
/usr/opt/SUNWmd/sbin/growfs
[ verifying class <none> ]
cp /cdrom/cdrom0/SUNWmd/reloc/etc/opt/SUNWmd/md.cf /etc/opt/SUNWmd/md.cf
cp /cdrom/cdrom0/SUNWmd/reloc/etc/opt/SUNWmd/md.ctrlmap
/etc/opt/SUNWmd/md.ctrlmap
cp /cdrom/cdrom0/SUNWmd/reloc/etc/opt/SUNWmd/md.tab /etc/opt/SUNWmd/md.tab
cp /cdrom/cdrom0/SUNWmd/reloc/etc/opt/SUNWmd/mddb.cf /etc/opt/SUNWmd/mddb.cf
cp /cdrom/cdrom0/SUNWmd/reloc/kernel/drv/md.conf /kernel/drv/md.conf
[ verifying class <preserve> ]
/kernel/drv/md
/kernel/misc/md_hotspares
/kernel/misc/md_mirror
/kernel/misc/md_raid
/kernel/misc/md_stripe
/kernel/misc/md_trans
/usr/lib/drv/preen_md.so.1
/usr/opt/SUNWmd/sbin/metaclear
/usr/opt/SUNWmd/sbin/metadb
```

(continued on following page)

```

/usr/opt/SUNWmd/sbin/metadetach
/usr/opt/SUNWmd/sbin/metahs
/usr/opt/SUNWmd/sbin/metainit
/usr/opt/SUNWmd/sbin/metaoffline
/usr/opt/SUNWmd/sbin/metaonline
/usr/opt/SUNWmd/sbin/metaparam
/usr/opt/SUNWmd/sbin/metareplace
/usr/opt/SUNWmd/sbin/metaroot
/usr/opt/SUNWmd/sbin/metastat
/usr/opt/SUNWmd/sbin/metasync
/usr/opt/SUNWmd/sbin/metattach
[ verifying class <sparc> ]
/usr/opt/SUNWmd/man/man1m/rpc.metamhd.1m
[ verifying class <noneOs54> ]
/usr/opt/SUNWmd/man/man1m/metaset.1m
/usr/opt/SUNWmd/man/man1m/rpc.metad.1m
/usr/opt/SUNWmd/sbin/metaset
/usr/opt/SUNWmd/sbin/rpc.metad
/usr/opt/SUNWmd/sbin/rpc.metamhd
[ verifying class <sparcOs54> ]
## Executing postinstall script.

Installation of <SUNWmd> was successful.

Processing package instance <SUNWmdg> from </cdrom/cdrom0>

Solstice DiskSuite Tool
(all) 4.0,REV=1.0

      Copyright (c) 1995 Sun Microsystems, Inc. All Rights Reserved.
      Manufactured in the United States of America.
2550 Garcia Avenue, Mountain View, California, 94043-1100 U.S.A.

... (miscellaneous copyright information)...

Using </> as the package base directory.
## Processing package information.
## Processing system information.
      8 package pathnames are already properly installed.
## Verifying package dependencies.
## Verifying disk space requirements.
## Checking for conflicts with packages already installed.
## Checking for setuid/setgid programs.
      (continued on following page)

```

```
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 this package [y,n,?] y
```

```
Installing Solstice DiskSuite Tool as <SUNWmdg>
```

```
## Executing preinstall script.
## Installing part 1 of 1.
/usr/opt/SUNWmd/lib/X11/app-defaults/Metatool
/usr/opt/SUNWmd/lib/X11/bitmaps/browser_icon.xpm
/usr/opt/SUNWmd/lib/X11/bitmaps/browser_icon_mask.xbm
/usr/opt/SUNWmd/lib/X11/bitmaps/concat_16.xpm
/usr/opt/SUNWmd/lib/X11/bitmaps/concat_16_insens.xpm
/usr/opt/SUNWmd/lib/X11/bitmaps/concat_32.xpm
/usr/opt/SUNWmd/lib/X11/bitmaps/concat_32_insens.xpm
/usr/opt/SUNWmd/lib/X11/bitmaps/concat_drag.xbm
/usr/opt/SUNWmd/lib/X11/bitmaps/concat_drag_mask.xbm
/usr/opt/SUNWmd/lib/X11/bitmaps/controller_12.xpm
/usr/opt/SUNWmd/lib/X11/bitmaps/controller_16.xpm
/usr/opt/SUNWmd/lib/X11/bitmaps/controller_32.xpm
/usr/opt/SUNWmd/lib/X11/bitmaps/diskview_icon.xpm
/usr/opt/SUNWmd/lib/X11/bitmaps/diskview_icon_mask.xbm
/usr/opt/SUNWmd/lib/X11/bitmaps/drop_none.xbm
/usr/opt/SUNWmd/lib/X11/bitmaps/drop_none_mask.xbm
/usr/opt/SUNWmd/lib/X11/bitmaps/drop_not_ok.xbm
/usr/opt/SUNWmd/lib/X11/bitmaps/drop_not_ok_mask.xbm
/usr/opt/SUNWmd/lib/X11/bitmaps/drop_ok.xbm
/usr/opt/SUNWmd/lib/X11/bitmaps/drop_ok_mask.xbm
/usr/opt/SUNWmd/lib/X11/bitmaps/drop_site_selection.xbm
/usr/opt/SUNWmd/lib/X11/bitmaps/editor_icon.xpm
/usr/opt/SUNWmd/lib/X11/bitmaps/editor_icon_mask.xbm
/usr/opt/SUNWmd/lib/X11/bitmaps/help_icon.xpm
/usr/opt/SUNWmd/lib/X11/bitmaps/help_icon_mask.xbm
/usr/opt/SUNWmd/lib/X11/bitmaps/hsp_16.xpm
/usr/opt/SUNWmd/lib/X11/bitmaps/hsp_16_insens.xpm
/usr/opt/SUNWmd/lib/X11/bitmaps/hsp_32.xpm
/usr/opt/SUNWmd/lib/X11/bitmaps/hsp_32_insens.xpm
/usr/opt/SUNWmd/lib/X11/bitmaps/hsp_drag.xbm
/usr/opt/SUNWmd/lib/X11/bitmaps/hsp_drag_mask.xbm
/usr/opt/SUNWmd/lib/X11/bitmaps/log_icon.xpm
/usr/opt/SUNWmd/lib/X11/bitmaps/log_icon_mask.xbm
```

(continued on following page)

```

/usr/opt/SUNWmd/lib/X11/bitmaps/metadb_16.xpm
/usr/opt/SUNWmd/lib/X11/bitmaps/metadb_16_insens.xpm
/usr/opt/SUNWmd/lib/X11/bitmaps/metadb_32.xpm
/usr/opt/SUNWmd/lib/X11/bitmaps/metadb_32_insens.xpm
/usr/opt/SUNWmd/lib/X11/bitmaps/metadb_drag.xbm
/usr/opt/SUNWmd/lib/X11/bitmaps/metadb_drag_mask.xbm
/usr/opt/SUNWmd/lib/X11/bitmaps/mirror_16.xpm
/usr/opt/SUNWmd/lib/X11/bitmaps/mirror_16_insens.xpm
/usr/opt/SUNWmd/lib/X11/bitmaps/mirror_32.xpm
/usr/opt/SUNWmd/lib/X11/bitmaps/mirror_32_insens.xpm
/usr/opt/SUNWmd/lib/X11/bitmaps/mirror_drag.xbm
/usr/opt/SUNWmd/lib/X11/bitmaps/mirror_drag_mask.xbm
/usr/opt/SUNWmd/lib/X11/bitmaps/multi_drag.xbm
/usr/opt/SUNWmd/lib/X11/bitmaps/multi_drag_mask.xbm
/usr/opt/SUNWmd/lib/X11/bitmaps/raid_16.xpm
/usr/opt/SUNWmd/lib/X11/bitmaps/raid_16_insens.xpm
/usr/opt/SUNWmd/lib/X11/bitmaps/raid_32.xpm
/usr/opt/SUNWmd/lib/X11/bitmaps/raid_32_insens.xpm
/usr/opt/SUNWmd/lib/X11/bitmaps/raid_drag.xbm
/usr/opt/SUNWmd/lib/X11/bitmaps/raid_drag_mask.xbm
/usr/opt/SUNWmd/lib/X11/bitmaps/slice_16.xpm
/usr/opt/SUNWmd/lib/X11/bitmaps/slice_16_insens.xpm
/usr/opt/SUNWmd/lib/X11/bitmaps/slice_drag.xbm
/usr/opt/SUNWmd/lib/X11/bitmaps/slice_drag_mask.xbm
/usr/opt/SUNWmd/lib/X11/bitmaps/target_12.xpm
/usr/opt/SUNWmd/lib/X11/bitmaps/target_16.xpm
/usr/opt/SUNWmd/lib/X11/bitmaps/target_32.xpm
/usr/opt/SUNWmd/lib/X11/bitmaps/trans_16.xpm
/usr/opt/SUNWmd/lib/X11/bitmaps/trans_16_insens.xpm
/usr/opt/SUNWmd/lib/X11/bitmaps/trans_32.xpm
/usr/opt/SUNWmd/lib/X11/bitmaps/trans_32_insens.xpm
/usr/opt/SUNWmd/lib/X11/bitmaps/trans_drag.xbm
/usr/opt/SUNWmd/lib/X11/bitmaps/trans_drag_mask.xbm
/usr/opt/SUNWmd/lib/X11/uid/Metatool/Editor.uid
/usr/opt/SUNWmd/lib/X11/uid/Metatool/EditorTpl.uid
/usr/opt/SUNWmd/lib/X11/uid/Metatool/concat.uid
/usr/opt/SUNWmd/lib/X11/uid/Metatool/hsp.uid
/usr/opt/SUNWmd/lib/X11/uid/Metatool/logs.uid
/usr/opt/SUNWmd/lib/X11/uid/Metatool/md.uid
/usr/opt/SUNWmd/lib/X11/uid/Metatool/mirror.uid
/usr/opt/SUNWmd/lib/X11/uid/Metatool/physicalview.uid
/usr/opt/SUNWmd/lib/X11/uid/Metatool/popups.uid
/usr/opt/SUNWmd/lib/X11/uid/Metatool/raid.uid

```

(continued on following page)

```
/usr/opt/SUNWmd/lib/X11/uid/Metatool/slice.uid
/usr/opt/SUNWmd/lib/X11/uid/Metatool/slicer.uid
/usr/opt/SUNWmd/lib/X11/uid/Metatool/stripe.uid
/usr/opt/SUNWmd/locale/C/help/metatool/howto/Howto
/usr/opt/SUNWmd/locale/C/help/metatool/howto/cat.create.h.hlp
/usr/opt/SUNWmd/locale/C/help/metatool/howto/catinfo.h.hlp
/usr/opt/SUNWmd/locale/C/help/metatool/howto/catstripe.del.h.hlp
/usr/opt/SUNWmd/locale/C/help/metatool/howto/cflog.h.hlp
/usr/opt/SUNWmd/locale/C/help/metatool/howto/cflog.show.h.hlp
/usr/opt/SUNWmd/locale/C/help/metatool/howto/clnup.h.hlp
/usr/opt/SUNWmd/locale/C/help/metatool/howto/clpse.h.hlp
/usr/opt/SUNWmd/locale/C/help/metatool/howto/col.assign.h.hlp
/usr/opt/SUNWmd/locale/C/help/metatool/howto/col.map.h.hlp
/usr/opt/SUNWmd/locale/C/help/metatool/howto/commit.h.hlp
/usr/opt/SUNWmd/locale/C/help/metatool/howto/confirm.h.hlp
/usr/opt/SUNWmd/locale/C/help/metatool/howto/delete.h.hlp
/usr/opt/SUNWmd/locale/C/help/metatool/howto/diskinfo.h.hlp
/usr/opt/SUNWmd/locale/C/help/metatool/howto/diskview.displaying.h.hlp
/usr/opt/SUNWmd/locale/C/help/metatool/howto/diskview.h.hlp
/usr/opt/SUNWmd/locale/C/help/metatool/howto/duplicate.h.hlp
/usr/opt/SUNWmd/locale/C/help/metatool/howto/eval.h.hlp
/usr/opt/SUNWmd/locale/C/help/metatool/howto/exit.h.hlp
/usr/opt/SUNWmd/locale/C/help/metatool/howto/expand.h.hlp
/usr/opt/SUNWmd/locale/C/help/metatool/howto/fail.h.hlp
/usr/opt/SUNWmd/locale/C/help/metatool/howto/find.h.hlp
/usr/opt/SUNWmd/locale/C/help/metatool/howto/help.help.h.hlp
/usr/opt/SUNWmd/locale/C/help/metatool/howto/hsp.about.h.hlp
/usr/opt/SUNWmd/locale/C/help/metatool/howto/hsp.addspares.h.hlp
/usr/opt/SUNWmd/locale/C/help/metatool/howto/hsp.associating.h.hlp
/usr/opt/SUNWmd/locale/C/help/metatool/howto/hsp.browser.h.hlp
/usr/opt/SUNWmd/locale/C/help/metatool/howto/hsp.browsing.h.hlp
/usr/opt/SUNWmd/locale/C/help/metatool/howto/hsp.create.h.hlp
/usr/opt/SUNWmd/locale/C/help/metatool/howto/hsp.deletespares.h.hlp
/usr/opt/SUNWmd/locale/C/help/metatool/howto/hsp.enable.h.hlp
/usr/opt/SUNWmd/locale/C/help/metatool/howto/hsp.getinfo.h.hlp
/usr/opt/SUNWmd/locale/C/help/metatool/howto/hsp.replacing.h.hlp
/usr/opt/SUNWmd/locale/C/help/metatool/howto/hsp.setfilters.h.hlp
/usr/opt/SUNWmd/locale/C/help/metatool/howto/hsp.statuscheck.h.hlp
/usr/opt/SUNWmd/locale/C/help/metatool/howto/hspinfo.h.hlp
/usr/opt/SUNWmd/locale/C/help/metatool/howto/info.h.hlp
/usr/opt/SUNWmd/locale/C/help/metatool/howto/interlace.h.hlp
/usr/opt/SUNWmd/locale/C/help/metatool/howto/load.h.hlp
/usr/opt/SUNWmd/locale/C/help/metatool/howto/log.to.file.h.hlp
```

(continued on following page)

```

/usr/opt/SUNWmd/locale/C/help/metatool/howto/maps.show.h.hlp
/usr/opt/SUNWmd/locale/C/help/metatool/howto/mapselect.h.hlp
/usr/opt/SUNWmd/locale/C/help/metatool/howto/md.browser.h.hlp
/usr/opt/SUNWmd/locale/C/help/metatool/howto/md.filters.h.hlp
/usr/opt/SUNWmd/locale/C/help/metatool/howto/mddbinfo.h.hlp
/usr/opt/SUNWmd/locale/C/help/metatool/howto/mded.h.hlp
/usr/opt/SUNWmd/locale/C/help/metatool/howto/mdinfo.h.hlp
/usr/opt/SUNWmd/locale/C/help/metatool/howto/mds.browsing.h.hlp
/usr/opt/SUNWmd/locale/C/help/metatool/howto/mirror.create.h.hlp
/usr/opt/SUNWmd/locale/C/help/metatool/howto/mirror.existing.fsys.h.hlp
/usr/opt/SUNWmd/locale/C/help/metatool/howto/mirror.onlinebackup.h.hlp
/usr/opt/SUNWmd/locale/C/help/metatool/howto/mirror.root.h.hlp
/usr/opt/SUNWmd/locale/C/help/metatool/howto/mirror.statuscheck.h.hlp
/usr/opt/SUNWmd/locale/C/help/metatool/howto/mirror.swap.h.hlp
/usr/opt/SUNWmd/locale/C/help/metatool/howto/mirror.usr.h.hlp
/usr/opt/SUNWmd/locale/C/help/metatool/howto/mirrorinfo.h.hlp
/usr/opt/SUNWmd/locale/C/help/metatool/howto/mousing.h.hlp
/usr/opt/SUNWmd/locale/C/help/metatool/howto/pan.h.hlp
/usr/opt/SUNWmd/locale/C/help/metatool/howto/panics.h.hlp
/usr/opt/SUNWmd/locale/C/help/metatool/howto/phy.map.h.hlp
/usr/opt/SUNWmd/locale/C/help/metatool/howto/problist.display.h.hlp
/usr/opt/SUNWmd/locale/C/help/metatool/howto/problist.h.hlp
/usr/opt/SUNWmd/locale/C/help/metatool/howto/putaway.h.hlp
/usr/opt/SUNWmd/locale/C/help/metatool/howto/raid.create.h.hlp
/usr/opt/SUNWmd/locale/C/help/metatool/howto/raidinfo.h.hlp
/usr/opt/SUNWmd/locale/C/help/metatool/howto/save.h.hlp
/usr/opt/SUNWmd/locale/C/help/metatool/howto/sdb.create.h.hlp
/usr/opt/SUNWmd/locale/C/help/metatool/howto/sdb.rem.h.hlp
/usr/opt/SUNWmd/locale/C/help/metatool/howto/sdb.status.h.hlp
/usr/opt/SUNWmd/locale/C/help/metatool/howto/selectall.h.hlp
/usr/opt/SUNWmd/locale/C/help/metatool/howto/slicebrowser.h.hlp
/usr/opt/SUNWmd/locale/C/help/metatool/howto/slicefilters.h.hlp
/usr/opt/SUNWmd/locale/C/help/metatool/howto/sliceinfo.h.hlp
/usr/opt/SUNWmd/locale/C/help/metatool/howto/slices.add.h.hlp
/usr/opt/SUNWmd/locale/C/help/metatool/howto/slices.browsing.h.hlp
/usr/opt/SUNWmd/locale/C/help/metatool/howto/slices.enable.h.hlp
/usr/opt/SUNWmd/locale/C/help/metatool/howto/slices.getinfo.h.hlp
/usr/opt/SUNWmd/locale/C/help/metatool/howto/slices.replace.h.hlp
/usr/opt/SUNWmd/locale/C/help/metatool/howto/slices.showinfo.h.hlp
/usr/opt/SUNWmd/locale/C/help/metatool/howto/spares.adding.h.hlp
/usr/opt/SUNWmd/locale/C/help/metatool/howto/spares.deleting.h.hlp
/usr/opt/SUNWmd/locale/C/help/metatool/howto/spares.enabling.h.hlp
/usr/opt/SUNWmd/locale/C/help/metatool/howto/spares.replacing.h.hlp

```

(continued on following page)

```
/usr/opt/SUNWmd/locale/C/help/metatool/howto/spares.status.h.hlp
/usr/opt/SUNWmd/locale/C/help/metatool/howto/stripeinfo.h.hlp
/usr/opt/SUNWmd/locale/C/help/metatool/howto/stripes.add.h.hlp
/usr/opt/SUNWmd/locale/C/help/metatool/howto/stripes.create.h.hlp
/usr/opt/SUNWmd/locale/C/help/metatool/howto/stripes.info.h.hlp
/usr/opt/SUNWmd/locale/C/help/metatool/howto/stripes.rem.h.hlp
/usr/opt/SUNWmd/locale/C/help/metatool/howto/submirr.attach.h.hlp
/usr/opt/SUNWmd/locale/C/help/metatool/howto/submirr.detach.h.hlp
/usr/opt/SUNWmd/locale/C/help/metatool/howto/submirr.enable.h.hlp
/usr/opt/SUNWmd/locale/C/help/metatool/howto/submirr.offline.h.hlp
/usr/opt/SUNWmd/locale/C/help/metatool/howto/submirr.online.h.hlp
/usr/opt/SUNWmd/locale/C/help/metatool/howto/submirr.rep.comp.h.hlp
/usr/opt/SUNWmd/locale/C/help/metatool/howto/submirr.setoptions.h.hlp
/usr/opt/SUNWmd/locale/C/help/metatool/howto/trans.create.h.hlp
/usr/opt/SUNWmd/locale/C/help/metatool/howto/trans.rem.logging.h.hlp
/usr/opt/SUNWmd/locale/C/help/metatool/howto/trans.sharelog.h.hlp
/usr/opt/SUNWmd/locale/C/help/metatool/howto/trans.status.h.hlp
/usr/opt/SUNWmd/locale/C/help/metatool/howto/trans.umtfilesys.logging.h.hlp
/usr/opt/SUNWmd/locale/C/help/metatool/howto/transinfo.h.hlp
/usr/opt/SUNWmd/locale/C/help/metatool/howto/ufslog.setup.h.hlp
/usr/opt/SUNWmd/locale/C/help/metatool/howto/undo.h.hlp
/usr/opt/SUNWmd/locale/C/help/metatool/howto/unmirror.filesys.h.hlp
/usr/opt/SUNWmd/locale/C/help/metatool/reference/Reference
/usr/opt/SUNWmd/locale/C/help/metatool/reference/catinfo.r.hlp
/usr/opt/SUNWmd/locale/C/help/metatool/reference/cflog.r.hlp
/usr/opt/SUNWmd/locale/C/help/metatool/reference/confirm.r.hlp
/usr/opt/SUNWmd/locale/C/help/metatool/reference/dev.suppl.r.hlp
/usr/opt/SUNWmd/locale/C/help/metatool/reference/dialogs.r.hlp
/usr/opt/SUNWmd/locale/C/help/metatool/reference/diskinfo.r.hlp
/usr/opt/SUNWmd/locale/C/help/metatool/reference/diskview.about.r.hlp
/usr/opt/SUNWmd/locale/C/help/metatool/reference/diskview.areas.r.hlp
/usr/opt/SUNWmd/locale/C/help/metatool/reference/diskview.menus.r.hlp
/usr/opt/SUNWmd/locale/C/help/metatool/reference/find.r.hlp
/usr/opt/SUNWmd/locale/C/help/metatool/reference/fs.log.r.hlp
/usr/opt/SUNWmd/locale/C/help/metatool/reference/gloss.r.hlp
/usr/opt/SUNWmd/locale/C/help/metatool/reference/growfs.cautions.r.hlp
/usr/opt/SUNWmd/locale/C/help/metatool/reference/growfs.r.hlp
/usr/opt/SUNWmd/locale/C/help/metatool/reference/guide.avail.r.hlp
/usr/opt/SUNWmd/locale/C/help/metatool/reference/guide.capacity.r.hlp
/usr/opt/SUNWmd/locale/C/help/metatool/reference/guide.labels.r.hlp
/usr/opt/SUNWmd/locale/C/help/metatool/reference/guide.perform.r.hlp
/usr/opt/SUNWmd/locale/C/help/metatool/reference/help.help.r.hlp
/usr/opt/SUNWmd/locale/C/help/metatool/reference/hspbrowser.r.hlp
```

(continued on following page)

```
/usr/opt/SUNWmd/locale/C/help/metatool/reference/hspfilters.r.hlp
/usr/opt/SUNWmd/locale/C/help/metatool/reference/hspinfo.r.hlp
/usr/opt/SUNWmd/locale/C/help/metatool/reference/load.r.hlp
/usr/opt/SUNWmd/locale/C/help/metatool/reference/log.r.hlp
/usr/opt/SUNWmd/locale/C/help/metatool/reference/mapselect.r.hlp
/usr/opt/SUNWmd/locale/C/help/metatool/reference/mbrowser.r.hlp
/usr/opt/SUNWmd/locale/C/help/metatool/reference/mddbinfo.r.hlp
/usr/opt/SUNWmd/locale/C/help/metatool/reference/mded_areas.r.hlp
/usr/opt/SUNWmd/locale/C/help/metatool/reference/mded_buttons.r.hlp
/usr/opt/SUNWmd/locale/C/help/metatool/reference/mded_menus.r.hlp
/usr/opt/SUNWmd/locale/C/help/metatool/reference/mdfilters.r.hlp
/usr/opt/SUNWmd/locale/C/help/metatool/reference/mdname.r.hlp
/usr/opt/SUNWmd/locale/C/help/metatool/reference/messages_dialog.r.hlp
/usr/opt/SUNWmd/locale/C/help/metatool/reference/messages_diskview.r.hlp
/usr/opt/SUNWmd/locale/C/help/metatool/reference/messages_logging.r.hlp
/usr/opt/SUNWmd/locale/C/help/metatool/reference/messages_mdedit.r.hlp
/usr/opt/SUNWmd/locale/C/help/metatool/reference/mirror_examples.r.hlp
/usr/opt/SUNWmd/locale/C/help/metatool/reference/mirror_naming.r.hlp
/usr/opt/SUNWmd/locale/C/help/metatool/reference/mirrorinfo.r.hlp
/usr/opt/SUNWmd/locale/C/help/metatool/reference/pan.r.hlp
/usr/opt/SUNWmd/locale/C/help/metatool/reference/phy_to_log.r.hlp
/usr/opt/SUNWmd/locale/C/help/metatool/reference/problist.r.hlp
/usr/opt/SUNWmd/locale/C/help/metatool/reference/raidinfo.r.hlp
/usr/opt/SUNWmd/locale/C/help/metatool/reference/save.r.hlp
/usr/opt/SUNWmd/locale/C/help/metatool/reference/security.r.hlp
/usr/opt/SUNWmd/locale/C/help/metatool/reference/slicebrowser.r.hlp
/usr/opt/SUNWmd/locale/C/help/metatool/reference/slicefilters.r.hlp
/usr/opt/SUNWmd/locale/C/help/metatool/reference/sliceinfo.r.hlp
/usr/opt/SUNWmd/locale/C/help/metatool/reference/spares_cautions.r.hlp
/usr/opt/SUNWmd/locale/C/help/metatool/reference/spares_naming.r.hlp
/usr/opt/SUNWmd/locale/C/help/metatool/reference/spares_pools.r.hlp
/usr/opt/SUNWmd/locale/C/help/metatool/reference/spares_states.r.hlp
/usr/opt/SUNWmd/locale/C/help/metatool/reference/stripeinfo.r.hlp
/usr/opt/SUNWmd/locale/C/help/metatool/reference/stripes_cautions.r.hlp
/usr/opt/SUNWmd/locale/C/help/metatool/reference/sysfiles.r.hlp
/usr/opt/SUNWmd/locale/C/help/metatool/reference/targets.r.hlp
/usr/opt/SUNWmd/locale/C/help/metatool/reference/trans_naming.r.hlp
/usr/opt/SUNWmd/locale/C/help/metatool/reference/transinfo.r.hlp
/usr/opt/SUNWmd/locale/C/help/metatool/reference/ufs_space_reqmts.r.hlp
/usr/opt/SUNWmd/locale/C/help/metatool/topics/Topics
/usr/opt/SUNWmd/locale/C/help/metatool/topics/cat_stripe.t.hlp
/usr/opt/SUNWmd/locale/C/help/metatool/topics/cflog.t.hlp
/usr/opt/SUNWmd/locale/C/help/metatool/topics/colors.t.hlp
```

(continued on following page)


```
/usr/opt/SUNWmd/locale/C/help/metatool/topics/commit.t.hlp
/usr/opt/SUNWmd/locale/C/help/metatool/topics/dialogs.t.hlp
/usr/opt/SUNWmd/locale/C/help/metatool/topics/diskinfo.t.hlp
/usr/opt/SUNWmd/locale/C/help/metatool/topics/diskview.t.hlp
/usr/opt/SUNWmd/locale/C/help/metatool/topics/find.t.hlp
/usr/opt/SUNWmd/locale/C/help/metatool/topics/gethelp.t.hlp
/usr/opt/SUNWmd/locale/C/help/metatool/topics/gloss.t.hlp
/usr/opt/SUNWmd/locale/C/help/metatool/topics/growfs.t.hlp
/usr/opt/SUNWmd/locale/C/help/metatool/topics/help.help.t.hlp
/usr/opt/SUNWmd/locale/C/help/metatool/topics/intro.t.hlp
/usr/opt/SUNWmd/locale/C/help/metatool/topics/key.t.hlp
/usr/opt/SUNWmd/locale/C/help/metatool/topics/load.t.hlp
/usr/opt/SUNWmd/locale/C/help/metatool/topics/log.t.hlp
/usr/opt/SUNWmd/locale/C/help/metatool/topics/mdd.t.hlp
/usr/opt/SUNWmd/locale/C/help/metatool/topics/mded.t.hlp
/usr/opt/SUNWmd/locale/C/help/metatool/topics/mds.t.hlp
/usr/opt/SUNWmd/locale/C/help/metatool/topics/mirror.t.hlp
/usr/opt/SUNWmd/locale/C/help/metatool/topics/object.create.t.hlp
/usr/opt/SUNWmd/locale/C/help/metatool/topics/phy.map.t.hlp
/usr/opt/SUNWmd/locale/C/help/metatool/topics/problist.t.hlp
/usr/opt/SUNWmd/locale/C/help/metatool/topics/putaway.t.hlp
/usr/opt/SUNWmd/locale/C/help/metatool/topics/raid.t.hlp
/usr/opt/SUNWmd/locale/C/help/metatool/topics/save.t.hlp
/usr/opt/SUNWmd/locale/C/help/metatool/topics/sdb.t.hlp
/usr/opt/SUNWmd/locale/C/help/metatool/topics/slicebrowser.t.hlp
/usr/opt/SUNWmd/locale/C/help/metatool/topics/slicefilters.t.hlp
/usr/opt/SUNWmd/locale/C/help/metatool/topics/slices.t.hlp
/usr/opt/SUNWmd/locale/C/help/metatool/topics/spares.about.t.hlp
/usr/opt/SUNWmd/locale/C/help/metatool/topics/stripes.t.hlp
/usr/opt/SUNWmd/locale/C/help/metatool/topics/templates.t.hlp
/usr/opt/SUNWmd/locale/C/help/metatool/topics/trans.t.hlp
/usr/opt/SUNWmd/locale/C/help/metatool/topics/warning.t.hlp
/usr/opt/SUNWmd/man/man1m/metatool.1m
[ verifying class <none> ]
/usr/opt/SUNWmd/sbin/metatool
[ verifying class <sparc> ]
```

Installation of <SUNWmdg> was successful.

*** IMPORTANT NOTICE ***

```
    This machine must now be rebooted in order to ensure
    sane operation.  Execute
        shutdown -y -i6 -g0
    and wait for the "Console Login:" prompt.
```

9. Eject the Solstice DiskSuite 4.0 CD and remove the CD from the caddy. Do not reboot the system at this time.

10. Insert the Solstice HA 1.0 CD-ROM into the drive and enter the command shown below. Enter *y* at any prompts about changing modes on directories.

It is essential that the packages on the Solstice HA CD-ROM be installed in the correct order. Pay close attention to the following instructions. The `pkgadd` command installs the Solstice HA files in the `/opt/SUNWhadf` and other directories. `pkgadd` does not overwrite any software included with the standard Solaris 2.4 release.

```
# pkgadd -d /cdrom/cdrom0 SUNWff SUNWcmm SUNWhagen SUNWhanfs SUNWabha

Processing package instance <SUNWff> from </cdrom/cdrom0>

SPARCcluster FailFast Device Driver
(sparc) 1.0,REV=1.0

    Copyright 1995 Sun Microsystems, Inc. All Rights Reserved.
    Manufactured in the United States of America.
    2550 Garcia Avenue, Mountain View, California, 94043-1100 U.S.A.

... (miscellaneous copyright information)...

Using </> as the package base directory.
## Processing package information.
## Processing system information.
    3 package pathnames are already properly installed.
## Verifying disk space requirements.
## Checking for conflicts with packages already installed.
## Checking for setuid/setgid programs.

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 this package [y,n,?] y

Installing SPARCcluster FailFast Device Driver as <SUNWff>

## Installing part 1 of 1.
/usr/kernel/drv/ff
    (continued on following page)
```

```
/usr/kernel/drv/ff.conf
[ verifying class <none> ]
Modifying /etc/devlink.tab
[ verifying class <sed> ]
## Executing postinstall script.

Installation of <SUNWff> was successful.

Processing package instance <SUNWcmm> from </cdrom/cdrom0>

SPARCcluster Cluster Membership Monitor
(sparc) 1.0,REV=1.0

... (miscellaneous copyright information)...

Using </opt> as the package base directory.
## Processing package information.
## Processing system information.
## Verifying disk space requirements.
## Checking for conflicts with packages already installed.
## Checking for setuid/setgid programs.

Installing SPARCcluster Cluster Membership Monitor as <SUNWcmm>

## Installing part 1 of 1.
/opt/SUNWcluster/bin/clustd
/opt/SUNWcluster/bin/clustm
/opt/SUNWcluster/lib/libclustm.a
[ verifying class <none> ]

Installation of <SUNWcmm> was successful.

Processing package instance <SUNWhagen> from </cdrom/cdrom0>

... (miscellaneous copyright information)...

Using </> as the package base directory.
## Processing package information.
## Processing system information.
    10 package pathnames are already properly installed.
## Verifying package dependencies.
## Verifying disk space requirements.
## Checking for conflicts with packages already installed.
## Checking for setuid/setgid programs.
    (continued on following page)
```

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 this package [y,n,?] **y**

Installing Solstice High Availability as <SUNWhagen>

```
## Installing part 1 of 1.
/etc/init.d/SUNWhadf
/etc/opt/SUNWhadf/hadf/cmm_confcdb.tmpl
/etc/opt/SUNWhadf/hadf/hadfconfig.tmpl
/etc/opt/SUNWhadf/hadf/hafmconfig
/etc/opt/SUNWhadf/hadf/hafmsuspend
/etc/opt/SUNWhadf/hadf/hasetup_sed.tmpl
/etc/opt/SUNWhadf/hadf/nsswitch.files
/etc/opt/SUNWhadf/hadf/nsswitch.files_dns
/etc/opt/SUNWhadf/hadf/nsswitch.files_nis
/etc/opt/SUNWhadf/hadf/nsswitch.files_nisplus
/etc/opt/SUNWhadf/hadf/nsswitch.files_nisplus_dns
/etc/opt/SUNWhadf/hadf/vfstab.tmpl
/etc/rc0.d/K40SUNWhadf <symbolic link>
/etc/rc1.d/K40SUNWhadf <symbolic link>
/etc/rc2.d/K40SUNWhadf <symbolic link>
/opt/SUNWhadf/bin/hacheck
/opt/SUNWhadf/bin/haconfig
/opt/SUNWhadf/bin/hafstab
/opt/SUNWhadf/bin/haload
/opt/SUNWhadf/bin/hasetup
/opt/SUNWhadf/bin/hastat
/opt/SUNWhadf/bin/haswitch
/opt/SUNWhadf/clust_progs/cltime
/opt/SUNWhadf/clust_progs/cltrans_fault_init_all
/opt/SUNWhadf/clust_progs/cltrans_fault_start_all
/opt/SUNWhadf/clust_progs/cltrans_fault_stop_all
/opt/SUNWhadf/clust_progs/cltrans_foreign_ipaddrs_down
/opt/SUNWhadf/clust_progs/cltrans_fsck_mount
/opt/SUNWhadf/clust_progs/cltrans_hadf_init
/opt/SUNWhadf/clust_progs/cltrans_ipaddrs_up
/opt/SUNWhadf/clust_progs/cltrans_release_foreign_disks
/opt/SUNWhadf/clust_progs/cltrans_start_svcs
/opt/SUNWhadf/clust_progs/cltrans_stop_all
/opt/SUNWhadf/clust_progs/cltrans_stop_foreign_svcs
/opt/SUNWhadf/clust_progs/cltrans_tkown_disks
```

(continued on following page)

```
/opt/SUNWhadf/clust_progs/haclksyn
/opt/SUNWhadf/clust_progs/hafs_mount
/opt/SUNWhadf/clust_progs/hafs_umount
/opt/SUNWhadf/clust_progs/havemaj
/opt/SUNWhadf/clust_progs/killmetatool
/opt/SUNWhadf/clust_progs/majmddb
/opt/SUNWhadf/clust_progs/mhictl
/opt/SUNWhadf/clust_progs/netctl
/opt/SUNWhadf/clust_progs/new_hadflog
/opt/SUNWhadf/clust_progs/opensleep
/opt/SUNWhadf/clust_progs/release_disks
/opt/SUNWhadf/clust_progs/remhastat <symbolic link>
/opt/SUNWhadf/clust_progs/runclocksyc
/opt/SUNWhadf/clust_progs/scsirstd
/opt/SUNWhadf/clust_progs/switchoversub
/opt/SUNWhadf/clust_progs/tkown_disks
/opt/SUNWhadf/clust_progs/utilities
/opt/SUNWhadf/fault_progs/faultd
/opt/SUNWhadf/fault_progs/faultdloop
/opt/SUNWhadf/fault_progs/faultm
/opt/SUNWhadf/fault_progs/fdl_checkload
/opt/SUNWhadf/fault_progs/fdl_checknameservice
/opt/SUNWhadf/fault_progs/fdl_cleanupfiles
/opt/SUNWhadf/fault_progs/fdl_consider_takeover
/opt/SUNWhadf/fault_progs/fdl_consider_takeover_crit
/opt/SUNWhadf/fault_progs/fdl_epoch
/opt/SUNWhadf/fault_progs/fdl_epochchanged
/opt/SUNWhadf/fault_progs/fdl_fault_suspend
/opt/SUNWhadf/fault_progs/fdl_fault_suspend_do
/opt/SUNWhadf/fault_progs/fdl_faultd_stop_all
/opt/SUNWhadf/fault_progs/fdl_faultmstart
/opt/SUNWhadf/fault_progs/fdl_filemtime
/opt/SUNWhadf/fault_progs/fdl_headc
/opt/SUNWhadf/fault_progs/fdl_host2ipaddr
/opt/SUNWhadf/fault_progs/fdl_load
/opt/SUNWhadf/fault_progs/fdl_lockedrun
/opt/SUNWhadf/fault_progs/fdl_lookuphost
/opt/SUNWhadf/fault_progs/fdl_mddb_latch
/opt/SUNWhadf/fault_progs/fdl_merge_err_on_out
/opt/SUNWhadf/fault_progs/fdl_printstatus
/opt/SUNWhadf/fault_progs/fdl_renamefile
/opt/SUNWhadf/fault_progs/fdl_request_relinquish
/opt/SUNWhadf/fault_progs/fdl_rshstatus
```

(continued on following page)

```

/opt/SUNWhadf/fault_progs/fdl_setpgidme
/opt/SUNWhadf/fault_progs/fdl_timedrun
/opt/SUNWhadf/fault_progs/fdl_timelabel
/opt/SUNWhadf/fault_progs/fdl_timesecs
/opt/SUNWhadf/fault_progs/ha_dbms_call
/opt/SUNWhadf/fault_progs/ha_dbms_serv
/opt/SUNWhadf/fault_progs/ha_fault_mon_config
/opt/SUNWhadf/fault_progs/hastat_net
/opt/SUNWhadf/fault_progs/load.faultd
/opt/SUNWhadf/fault_progs/load_fault_init
/opt/SUNWhadf/fault_progs/load_fault_start
/opt/SUNWhadf/fault_progs/net.faultd
/opt/SUNWhadf/fault_progs/net_arp_retrans
/opt/SUNWhadf/fault_progs/net_check_my_ok
/opt/SUNWhadf/fault_progs/net_diagnose_comm
/opt/SUNWhadf/fault_progs/net_diagnose_comm_alone
/opt/SUNWhadf/fault_progs/net_diagnose_comm_takeover
/opt/SUNWhadf/fault_progs/net_donetstati
/opt/SUNWhadf/fault_progs/net_fault_init
/opt/SUNWhadf/fault_progs/net_fault_start
/opt/SUNWhadf/fault_progs/net_hosts_check
/opt/SUNWhadf/fault_progs/net_hosts_interfaces_table
/opt/SUNWhadf/fault_progs/net_hosts_table
/opt/SUNWhadf/fault_progs/net_hosts_table_check
/opt/SUNWhadf/fault_progs/net_ipaddrs_arp
/opt/SUNWhadf/fault_progs/net_periodic_netstat
/opt/SUNWhadf/fault_progs/net_periodic_ping_other
/opt/SUNWhadf/fault_progs/net_pinghost
/opt/SUNWhadf/fault_progs/net_pingnet
/opt/SUNWhadf/fault_progs/net_probe_brother
/opt/SUNWhadf/fault_progs/net_probe_brother_private
/opt/SUNWhadf/fault_progs/net_probe_brother_public
/opt/SUNWhadf/fault_progs/net_rpcinfo
/opt/SUNWhadf/fault_progs/net_rpcprobe
/opt/SUNWhadf/fault_progs/net_transition_check_comm
/opt/SUNWhadf/fault_progs/netinterface
/opt/SUNWhadf/fault_progs/netstat_analyze
/opt/SUNWhadf/fault_progs/netstat_analyze_cmp
/opt/SUNWhadf/fault_progs/sarpioc
/opt/SUNWhadf/fault_progs/sendarp
/opt/SUNWhadf/man/man1m/hacheck.1m
/opt/SUNWhadf/man/man1m/hafstab.1m
/opt/SUNWhadf/man/man1m/haload.1m

```

(continued on following page)

```
/opt/SUNWhadf/man/man1m/hasetup.1m
/opt/SUNWhadf/man/man1m/hastat.1m
/opt/SUNWhadf/man/man1m/haswitch.1m
[ verifying class <none> ]
Modifying /etc/group
[ verifying class <sed> ]
## Executing postinstall script.

Installation of <SUNWhagen> was successful.

Processing package instance <SUNWhanfs> from </cdrom/cdrom0>

Solstice High Availability -- HA-NFS component

... (miscellaneous copyright information)...

Using </> as the package base directory.
## Processing package information.
## Processing system information.
    8 package pathnames are already properly installed.
## Verifying package dependencies.
## Verifying disk space requirements.
## Checking for conflicts with packages already installed.
## Checking for setuid/setgid programs.

Installing Solstice High Availability -- HA-NFS component as <SUNWhanfs>

## Installing part 1 of 1.
/etc/opt/SUNWhadf/nfs/dfstab.tmpl
/opt/SUNWhadf/bin/nfs_hacheck
/opt/SUNWhadf/clust_progs/lm_shutdown
/opt/SUNWhadf/clust_progs/nfs_svc_abort
/opt/SUNWhadf/clust_progs/nfs_svc_start
/opt/SUNWhadf/clust_progs/nfs_svc_stop
/opt/SUNWhadf/clust_progs/share_to_mountp
/opt/SUNWhadf/fault_progs/nfs.faultd
/opt/SUNWhadf/fault_progs/nfs_check_my_lockd_statd
/opt/SUNWhadf/fault_progs/nfs_fm_check_this_host_ok
/opt/SUNWhadf/fault_progs/nfs_fm_init
/opt/SUNWhadf/fault_progs/nfs_fm_start
/opt/SUNWhadf/fault_progs/nfs_fm_stop
/opt/SUNWhadf/fault_progs/nfs_hastat
/opt/SUNWhadf/fault_progs/nfs_islockedfs
```

(continued on following page)

```
/opt/SUNWhadf/fault_progs/nfs_just_umount
/opt/SUNWhadf/fault_progs/nfs_mon
/opt/SUNWhadf/fault_progs/nfs_monitor_pids
/opt/SUNWhadf/fault_progs/nfs_mounttouchfile
/opt/SUNWhadf/fault_progs/nfs_probe_local
/opt/SUNWhadf/fault_progs/nfs_probe_local_restart
/opt/SUNWhadf/fault_progs/nfs_probe_once
/opt/SUNWhadf/fault_progs/nfs_probe_one_common
/opt/SUNWhadf/fault_progs/nfs_probe_remote
/opt/SUNWhadf/fault_progs/nfs_probe_umount_all
/opt/SUNWhadf/fault_progs/nfs_share2mp
/opt/SUNWhadf/fault_progs/nfs_touchfile
[ verifying class <none> ]
```

Installation of <SUNWhanfs> was successful.

Processing package instance <SUNWabha> from </cdrom/cdrom0>

High Availability AnswerBook
(all) 83.1.4

... (*miscellaneous copyright information*)...

The installation options are as follows:

Option: Description:

- 1. nil: less than 1 Megabyte disk space required [slowest performance].
- 2. heavy: 13.13 Megabytes disk space required [best performance].

Note: If the install option which you choose below fails
due to lack of space, try another location, or
choose a lower install option number.

Enter the number of an installation option from the list above (1 or 2).

Select an installation option: **2**

Installation option: heavy selected.

The next request for input asks you to specify the parent directory of AnswerBook
Make sure to choose a parent directory on a file system big enough to
accommodate all the files to be moved for the INSTALL OPTION you selected.

Specify the parent of the AnswerBook home directory: **/opt**

For the heavy option all files will be placed under /opt/SUNWabha.

(continued on following page)


```
## Processing package information.
## Processing system information.
## Verifying package dependencies.
## Verifying disk space requirements.
## Checking for conflicts with packages already installed.
## Checking for setuid/setgid programs.
```

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 this package [y,n,?] **y**

Installing High Availability AnswerBook as <SUNWabha>

```
## Installing part 1 of 1.
/opt/SUNWabha/index
/opt/SUNWabha/index/Keys
/opt/SUNWabha <implied directory>
/opt/SUNWabha/index <implied directory>
/opt/SUNWabha/index/index.cat
/opt/SUNWabha/index/index.cfg
/opt/SUNWabha/index/index.cix
/opt/SUNWabha/index/index.dct
/opt/SUNWabha/index/index.log
/opt/SUNWabha/index/index.ref
[ verifying class <Index> ]
/opt/SUNWabha/toc/HAADMIN.ind
/opt/SUNWabha/toc/HAADMIN.lock
/opt/SUNWabha/toc/HAADMIN.rec
/opt/SUNWabha/toc/HAHWINSTALL.ind
/opt/SUNWabha/toc/HAHWINSTALL.lock
/opt/SUNWabha/toc/HAHWINSTALL.rec
/opt/SUNWabha/toc/HAINSTALL.ind
/opt/SUNWabha/toc/HAINSTALL.lock
/opt/SUNWabha/toc/HAINSTALL.rec
/opt/SUNWabha/toc/HASERVICE.ind
/opt/SUNWabha/toc/HASERVICE.lock
/opt/SUNWabha/toc/HASERVICE.rec
/opt/SUNWabha/toc/SUNWab_83_1.ind
/opt/SUNWabha/toc/SUNWab_83_1.lock
/opt/SUNWabha/toc/SUNWab_83_1.rec
[ verifying class <ContentsDB> ]
```

(continued on following page)

```

/opt/SUNWabha/ps/HAADMIN/01.Introduction
/opt/SUNWabha/ps/HAADMIN/02.Preparing_for_Administration
/opt/SUNWabha/ps/HAADMIN/03.Monitoring_the_Solstice_HA_Servers
/opt/SUNWabha/ps/HAADMIN/04.Hardware_Replacement_and_Repair
/opt/SUNWabha/ps/HAADMIN/05.Adding_Hardware
/opt/SUNWabha/ps/HAADMIN/06.HA-NFS_Maintenance
/opt/SUNWabha/ps/HAADMIN/07.HA-DBMS_Maintenance
/opt/SUNWabha/ps/HAADMIN/08.Metadevice_and_Diskset_Administration
/opt/SUNWabha/ps/HAADMIN/09.General_Solstice_HA_Maintenance
/opt/SUNWabha/ps/HAADMIN/A.Error_Messages
/opt/SUNWabha/ps/HAADMIN/B.Man_Pages
/opt/SUNWabha/ps/HAADMIN/Contents
/opt/SUNWabha/ps/HAADMIN/Cover
/opt/SUNWabha/ps/HAADMIN/Credits
/opt/SUNWabha/ps/HAADMIN/Figures
/opt/SUNWabha/ps/HAADMIN/Index
/opt/SUNWabha/ps/HAADMIN/Preface
/opt/SUNWabha/ps/HAADMIN/files_to_print
/opt/SUNWabha/ps/HAHWINSTALL/01.Checklist
/opt/SUNWabha/ps/HAHWINSTALL/02.System_Description
/opt/SUNWabha/ps/HAHWINSTALL/03.Specifications
/opt/SUNWabha/ps/HAHWINSTALL/04.Site_Preparation_and_Planning
/opt/SUNWabha/ps/HAHWINSTALL/05.SPARCcluster_1000_High_Availability_C
/opt/SUNWabha/ps/HAHWINSTALL/06.SPARCcluster_2000_High_Availability_C
/opt/SUNWabha/ps/HAHWINSTALL/07.Powering_Off_and_On
/opt/SUNWabha/ps/HAHWINSTALL/08.Access
/opt/SUNWabha/ps/HAHWINSTALL/09.SPARCcluster_1000_High_Availability_H
/opt/SUNWabha/ps/HAHWINSTALL/10.SPARCcluster_2000_High_Availability_S
/opt/SUNWabha/ps/HAHWINSTALL/11.Installing_the_Software
/opt/SUNWabha/ps/HAHWINSTALL/A.Rack_Mount_Rail_Installation
/opt/SUNWabha/ps/HAHWINSTALL/Contents
/opt/SUNWabha/ps/HAHWINSTALL/Cover
/opt/SUNWabha/ps/HAHWINSTALL/Credits
/opt/SUNWabha/ps/HAHWINSTALL/Figures
/opt/SUNWabha/ps/HAHWINSTALL/Index
/opt/SUNWabha/ps/HAHWINSTALL/Part_1_-t.Site_Preparation_and_Planning
/opt/SUNWabha/ps/HAHWINSTALL/Part_2_-t.Hardware_Installation
/opt/SUNWabha/ps/HAHWINSTALL/Part_3_-t.Software_Installation
/opt/SUNWabha/ps/HAHWINSTALL/Part_4_-t.Appendix_and_Index
/opt/SUNWabha/ps/HAHWINSTALL/Preface
/opt/SUNWabha/ps/HAHWINSTALL/Revision_History
/opt/SUNWabha/ps/HAHWINSTALL/Tables
/opt/SUNWabha/ps/HAHWINSTALL/files_to_print
      (continued on following page)

```

```
/opt/SUNWabha/ps/HAHWINSTALL/zz-card
/opt/SUNWabha/ps/HAINSTALL/01.Introduction
/opt/SUNWabha/ps/HAINSTALL/02.Planning_and_Installation_Overview
/opt/SUNWabha/ps/HAINSTALL/03.Configuration_Planning
/opt/SUNWabha/ps/HAINSTALL/04.Hardware_Installation_Information
/opt/SUNWabha/ps/HAINSTALL/05.Solaris_2.4_Installation
/opt/SUNWabha/ps/HAINSTALL/06.Installing_the_Packages_and_Patches
/opt/SUNWabha/ps/HAINSTALL/07.Setting_Up_Name_Services__and_Network
/opt/SUNWabha/ps/HAINSTALL/08.Creating_the_Configuration
/opt/SUNWabha/ps/HAINSTALL/09.Metadevice_Creation_on_Local_and_Mult
/opt/SUNWabha/ps/HAINSTALL/10.Setting_Up_Data_Service_Classes
/opt/SUNWabha/ps/HAINSTALL/11.Verification_and_Validation
/opt/SUNWabha/ps/HAINSTALL/A.Error_Messages
/opt/SUNWabha/ps/HAINSTALL/B.Planning_Template
/opt/SUNWabha/ps/HAINSTALL/C.Terminal_Concentrator_Security
/opt/SUNWabha/ps/HAINSTALL/Contents
/opt/SUNWabha/ps/HAINSTALL/Cover
/opt/SUNWabha/ps/HAINSTALL/Credits
/opt/SUNWabha/ps/HAINSTALL/Figures
/opt/SUNWabha/ps/HAINSTALL/Index
/opt/SUNWabha/ps/HAINSTALL/Preface
/opt/SUNWabha/ps/HAINSTALL/Tables
/opt/SUNWabha/ps/HAINSTALL/files_to_print
/opt/SUNWabha/ps/HASERVICE/01.Product_Description
/opt/SUNWabha/ps/HASERVICE/02.Troubleshooting_Overview
/opt/SUNWabha/ps/HASERVICE/03.Hardware_Troubleshooting
/opt/SUNWabha/ps/HASERVICE/04.Software_Troubleshooting
/opt/SUNWabha/ps/HASERVICE/05.Diagnostics
/opt/SUNWabha/ps/HASERVICE/06.Saftey__and_Tools_Requirements
/opt/SUNWabha/ps/HASERVICE/07.Shutdown_and_Restart_Procedures
/opt/SUNWabha/ps/HASERVICE/08.Internal_Access
/opt/SUNWabha/ps/HASERVICE/09._Major_Subassemblies
/opt/SUNWabha/ps/HASERVICE/10.Illustrated_Parts_Breakdown
/opt/SUNWabha/ps/HASERVICE/A.Product_Specifications
/opt/SUNWabha/ps/HASERVICE/B.Connector_Pinouts_and_Cabling
/opt/SUNWabha/ps/HASERVICE/C.SCSI__Targeting
/opt/SUNWabha/ps/HASERVICE/Contents
/opt/SUNWabha/ps/HASERVICE/Cover
/opt/SUNWabha/ps/HASERVICE/Credits
/opt/SUNWabha/ps/HASERVICE/D.SPARCstorage_Array_Firmware_and_Device
/opt/SUNWabha/ps/HASERVICE/Figures
/opt/SUNWabha/ps/HASERVICE/Index
/opt/SUNWabha/ps/HASERVICE/Part_1_-t.System_Information
      (continued on following page)
```

```

/opt/SUNWabha/ps/HASERVICE/Part_2_-t.Troubleshooting
/opt/SUNWabha/ps/HASERVICE/Part_3_-t.Preparing_for_Service
/opt/SUNWabha/ps/HASERVICE/Part_4_-t.Subassembly_Removal_and_Repla
/opt/SUNWabha/ps/HASERVICE/Part_5_-t.Illustrated_Parts_Breakdown
/opt/SUNWabha/ps/HASERVICE/Part_6_-t.Appendixes_Glossary_and_Ind
/opt/SUNWabha/ps/HASERVICE/Preface
/opt/SUNWabha/ps/HASERVICE/Revision_History
/opt/SUNWabha/ps/HASERVICE/Tables
/opt/SUNWabha/ps/HASERVICE/comment_card
/opt/SUNWabha/ps/HASERVICE/files_to_print
[ verifying class <PostScript> ]
## Executing postinstall script.
Installation of <SUNWabha> was successful.

```

11. Install all the Solstice HA patches by running the following script.
 Installation of the patches may take some time.

```

# cd /cdrom/cdrom0
# ls
101945-34/      102432-07/      102769-01/      SUNWhagen/
101977-04/      102446-01/      SUNWabha/        SUNWhanfs/
102216-02/      102580-02/      SUNWcmm/         SUNWvts/
102283-01/      102768-01/      SUNWff/          load_patches*
# ./load_patches

Installing 101945-34
Installing 101977-04
Installing 102216-02
Installing 102283-01
Installing 102432-07
Installing 102446-01
Installing 102580-02
Installing 102768-01
Installing 102769-01
#

```

12. Change to the root directory.

```
# cd /
```

13. Eject the Solstice HA 1.0 CD.

14. Insert the Solstice HA ORACLE 1.0 CD-ROM into the drive and enter the command shown below. Enter *y* at any prompts about changing modes on directories.

```
# pkgadd -d /cdrom/cdrom0 SUNWhaor
Processing package instance <SUNWhaor> from </cdrom/cdrom0>

Solstice High Availability -- HA-Oracle component
(all) 1.0,REV=1.0
... (miscellaneous copyright information)...
Using </> as the package base directory.

Installing Solstice High Availability -- HA-Oracle component as <SUNWhaor>

## Installing part 1 of 1.
/etc/opt/SUNWhadf/hadf/haoracle_config_V1
/etc/opt/SUNWhadf/hadf/haoracle_databases.tmpl
/etc/opt/SUNWhadf/hadf/haoracle_support
/opt/SUNWhadf/bin/haoracle
/opt/SUNWhadf/bin/oracle_hacheck
/opt/SUNWhadf/clust_progs/oracle_config_check
/opt/SUNWhadf/clust_progs/oracle_db_shutdown
/opt/SUNWhadf/clust_progs/oracle_status
/opt/SUNWhadf/clust_progs/oracle_status_svcs
/opt/SUNWhadf/clust_progs/oracle_svc_abort
/opt/SUNWhadf/clust_progs/oracle_svc_start
/opt/SUNWhadf/clust_progs/oracle_svc_stop
/opt/SUNWhadf/fault_progs/haoracle_fmon
/opt/SUNWhadf/fault_progs/haoracle_fmon_start
/opt/SUNWhadf/fault_progs/oracle_db_restart
/opt/SUNWhadf/fault_progs/oracle_fm_check_this_host_ok
/opt/SUNWhadf/fault_progs/oracle_fm_init
/opt/SUNWhadf/fault_progs/oracle_fm_start
/opt/SUNWhadf/fault_progs/oracle_fm_stop
/opt/SUNWhadf/fault_progs/oracle_hastat
/opt/SUNWhadf/man/man1m/haoracle.1m
/opt/SUNWhadf/man/man1m/haoracle_config.1m
/opt/SUNWhadf/man/man4/haoracle_config.4
/opt/SUNWhadf/man/man4/haoracle_databases.4
/opt/SUNWhadf/man/man4/haoracle_support.4
[ verifying class <none> ]

Installation of <SUNWhaor> was successful.
```

15. (Optional) Edit the `/kernel/drv/md.conf` file if you need more than 128 metadvice names.

Refer to “Metadvice Naming and Creation” on page 3-21 for information about the number of metadvice names needed. Additional information about this topic can be found in Section 5.2, “Selecting the Size of the Root Slice.”

16. Check the accuracy of the packages that have been installed.

The following command should be run for each of the packages installed.

```
# pkgchk -n package_name
```

17. Reboot both servers.

The `telnet(1)` connection to each of the servers will survive the reboot because the connections are actually made to the `/dev/ttya` console port of the terminal concentrator and not directly to the servers.

18. Update the SPARCstorage Array firmware, if required.

Follow the “Special Install Instructions” found in the `/var/sadm/patch/102432-07/README.102432-07` file to update the firmware. You must update the firmware if the revision level is lower than 2.3. This procedure only needs to be performed once per SPARCstorage Array. Typically you will do all SPARCstorage Arrays from one of the servers.

19. Create metadvice state database replicas on the boot disk of each server by entering the following command on each server.

```
# /usr/opt/SUNWmd/sbin/metadb -afc 3 c0t0d0s4
```

20. Set the required paths in the following section.

6.2 Path Requirements

After installing the software, you must set your environment variables `PATH` and `MANPATH`.

The general requirements are as follows:

- Insert `/usr/opt/SUNWmd/sbin` and `/opt/SUNWhadf/bin` either before or after `PATH`.
- Insert `/usr/opt/SUNWmd/man` and `/opt/SUNWhadf/man` either before or after your `MANPATH`.

You can set these paths by using an editor to change your `.profile` or `.login` file, as follows.

If you installed the software into the default directory, then:

- If you're using the Bourne shell (`sh(1)`), your `.profile` file should have lines like this:

```
PATH=/usr/opt/SUNWmd/sbin:/opt/SUNWhadf/bin:$PATH
MANPATH=/usr/opt/SUNWmd/man:/opt/SUNWhadf/man:/usr/man
export PATH MANPATH
```

- If you're using the C shell (`csh(1)`), the `.login` file should have lines that look something like this:

```
set path = (/usr/opt/SUNWmd/sbin /opt/SUNWhadf/bin $path)
setenv MANPATH /usr/opt/SUNWmd/man:/opt/SUNWhadf/man:/usr/man
```

6.3 Viewing AnswerBook Online Documentation

To view the AnswerBook online documentation use the `answerbook(1)` command. Be certain that `/usr/openwin/bin` is in the search path or the `docviewer(1)` will not be found. You must set the `DISPLAY` environment variable before running `answerbook`. For example:

```
$ DISPLAY=workstation:0
$ EXPORT DISPLAY
$ /usr/openwin/bin/answerbook
```

Note – Because Solstice HA servers allow NFS sharing only from multi-host file systems and the AnswerBook is installed on the local disks, you can only access the AnswerBook from the local server. If you need network access to the AnswerBooks, you should copy the AnswerBook files to another NFS server.

Setting Up Name Services and Networks



This chapter discusses how to set up the name service databases on Solstice HA systems. Use the following table to locate specific information in this chapter.

<i>Overview of Tasks</i>	<i>page 7-1</i>
<i>Updates To Name Service Entries</i>	<i>page 7-3</i>
<i>Automatic Changes to /etc/nsswitch.conf</i>	<i>page 7-3</i>
<i>Updates to /etc Files</i>	<i>page 7-5</i>
<i>Setting Up Private Networks</i>	<i>page 7-7</i>
<i>Setting Up /.rhosts Files</i>	<i>page 7-8</i>
<i>Editing /etc/syslog.conf</i>	<i>page 7-8</i>

Note – The same network name service that is used throughout the site should be used for the Solstice HA systems.

7.1 Overview of Tasks

When setting up the name service and networks for Solstice HA systems, several /etc files must be edited. Use the information gathered in Chapter 3, “Installation Planning,” to update the files.

The basic steps you will perform to update the name service and networks include:

1. Updating your name service information throughout the network
2. Updating the `/etc` name service files on each Solstice HA system to contain the private network host names
3. Creating `.rhost` files on each Solstice HA server
4. Editing the `/etc/syslog.conf` file on each Solstice HA server

Setting up the NIS name service for Solstice HA configurations is the same as for any other system on the network. Entries for all Solstice HA host names and all IP addresses (whether they are for physical or logical machines) should be made in the local `/etc/inet/hosts` file. Adding the Solstice HA host names and IP addresses adds an extra level of integrity.



Caution – Solstice HA configurations must not be set up as routers. The Solstice HA machines can only be clients of routers. Also, the Solstice HA servers cannot be set up as name servers because the NIS or NIS+ functionality does not failover to the other machine in the Solstice HA configuration.

During installation, the primary network host name and IP address for each physical Solstice HA machine were entered into the `/etc/inet/hosts` file. If you have secondary host names and IP addresses for the physical Solstice HA machines these names and addresses must also be present in the `/etc/inet/hosts` file. `hasetup` will prompt for the needed information. `hasetup` will consult your network name service for the appropriate IP addresses before prompting for them.

Because there may be multiple public network interfaces on each of the Solstice HA machines, you must decide which of the interfaces the logical machine will use. Ideally, the logical machines use all of the public interfaces. This implies that when adding the logical and physical IP addresses and host names to the name service, you must ensure that all entries are added to the appropriate host maps.

7.2 Updates To Name Service Entries

If you are using a network name service you will find it convenient to add the entries for the logical and physical hosts (but not the private hosts) into your name service as well as the associated `networks` and `netmasks` entries. You can use standard administrative procedures to update your name service tables associated with `/etc/inet/hosts`, `/etc/inet/networks`, and `/etc/inet/netmasks` with the same information entered in the local `/etc` files.

If you do not make the entries in the network name service, you may optionally create entries in the local `/etc` files on both hosts or `hasetup` will prompt for the information and make the entries automatically.

7.3 Automatic Changes to `/etc/nsswitch.conf`

When you run the `hasetup(1M)` command as part of the Solstice HA installation (see Chapter 8, “Creating the Configuration”), the `/etc/nsswitch.conf` file on each of the Solstice HA server will be overwritten. The new `/etc/nsswitch.conf` file will be copied from the `/etc/opt/SUNWhadf/hadf` directory. For instance, if you are using NIS+ and DNS, the `nsswitch.conf` file is overwritten with the `nsswitch.files.nisplus_dns` from that directory.

The following example `/etc/nsswitch.nis` file shows the type of changes made by the Solstice HA software. The changes made to the example `/etc/nsswitch.conf` file include moving files in front of nis in every field except netgroup.

```
#ident "@(#)nsswitch.files_nis 1.1 95/05/18 SMI"
#
# /etc/nsswitch.files_nis:
#
# This file is installed as part of the SolsticeHA installation
# procedures. SolsticeHA requires that "files" are consulted
# first for most items.
#
# "hosts:" and "services:" in this file are used only if the
# /etc/netconfig file has a "-" for nametoaddr_libs of "inet"
# transports.

# the following two lines obviate the "+" entry in /etc/passwd
# and /etc/group.
passwd: files nis
group: files nis

# always consult /etc "files" before consulting nis
hosts: files nis
networks: files nis
protocols: files nis
rpc: files nis
ethers: files nis
netmasks: files nis
bootparams: files nis
publickey: files nis

# At present there isn't a "files" backend for netgroup.
netgroup: nis

automount: files nis
aliases: files nis

services: files nis
sendmailvars: files
```

7.4 Updates to `/etc` Files

To add IP addresses to the `/etc` name service, the `/etc/inet/hosts` file must be manually edited on each of the Solstice HA servers. Optionally, you can also edit the `/etc/inet/networks` file. You may also want to verify that the name service entries in the `/etc/inet/netmasks` are correct.

The IP address and hostnames for the four private interfaces must be entered in the `/etc/inet/hosts` files on both Solstice HA servers. The IP address for the Solstice HA hosts were added during installation.

Note – Use either NIS or NIS+ name service for Solstice HA systems because of the support provided by administration tools that can propagate the information from one Solstice HA system to the other Solstice HA system.

Follow these steps to add the host names and IP addresses. The changes must be made in the local `/etc` files and in the network name service.

1. Edit the `/etc/inet/hosts` file on each Solstice HA server.

Add entries to the `/etc/inet/hosts` file on each server for the private Solstice HA interfaces. The following is an example `/etc/inet/hosts` file. Optionally, all physical and logical hosts may be added or deferred until you are running `hasetup(1M)`.

```
# Internet host table
#
192.9.201      localhost
192.9.200      host1 loghost

# private links
204.152.64.1  host1-priv1
204.152.65.1  host1-priv2
204.152.64.2  host2-priv1
204.152.65.2  host2-priv2
```

2. (Optional) Make sure name service entries are correct.

The `hsetup` command will automatically edit the `/etc/inet/netmasks` file on each of the Solstice HA servers. If you have multiple networks and use non-default `netmasks` you should make sure that your `netmasks` name service or the local `/etc` file has the appropriate netmask.

```
#
# The netmasks file associates Internet Protocol (IP) address
# masks with IP network numbers.
#
#      network-number  netmask
#
# Both the network-number and the netmasks are specified in
# "decimal dot" notation, e.g:
#
#           192.9.200.0 255.255.255.0
#
192.9.200      255.255.255.0
192.0.0.0     255.0.0.0
192           255.0.0.0
192.9.200.0   255.255.255.0
192.9.200     255.255.255.0
204.152.64.0  255.255.255.0
204.152.64    255.255.255.0
204.152.65.0  255.255.255.0
204.152.65    255.255.255.0
```

3. (Optional) Edit the `/etc/inet/networks` file.

Make sure there are entries for each physically attached network. The entries are made so a network name is displayed when the `netstat(1M)` command is run. These entries enable `netstat(1M)` to print the names of

machines rather than the network numbers. The entries also make `netstat` run faster because it eliminates the need to query NIS or NIS+. The following is an example `/etc/inet/networks` file.

```
#ident "@(#)networks 1.4 92/07/14 SMI" /* SVr4.0 1.1 */
#
# The networks file associates Internet Protocol (IP) network
# numbers with network names. The format of this file is:
#
#     network-name    network-number  nicnames . . .
#
#
# The loopback network is used only for intra-machine
communication
#
loopback            192

# Example Networks
#
example-net1       192.9.200.0
example-net2       192.9.201.0
priv-net1          204.152.64.0
priv-net2          204.152.65.0
...
```

7.5 Setting Up Private Networks

On each of the Solstice HA servers, you must create two new files. These files are named `/etc/hostname.be0` and `/etc/hostname.be1`, by default. The `be[01]` suffix corresponds to the private network names. These may be different on your system.

Each file contains the private network host name associated with the connection. Following the example in Table 3-3 on page 3-7 the contents of the `/etc/hostname.be0` on `host1` would be `host1-priv1` and the contents of the sibling server `/etc/hostname.be1` would be `host2-priv2`.

When all four files are created, you must either reboot the systems or run the following `ifconfig(1M)` commands on each host. For example:

```
# ifconfig be0 plumb
# ifconfig be0 ip_address netmask + broadcast + -trailers up
```

The above commands must be run for both the `be0` and `be1` interfaces on each of the servers. Alternatively, you can reboot both Solstice HA servers.

7.6 Setting Up `.rhosts` Files

Both systems in the Solstice HA configuration must have entries for private network interfaces in their respective `.rhosts` files.

The reason for adding the entries to the `.rhosts` files is to allow Solstice HA software to update configuration files on both physical hosts using the `rsh(1)` command.

The following example `.rhosts` file shows the entries required on each of the Solstice HA servers:

```
host1-priv1 root
host1-priv2 root
host2-priv1 root
host2-priv2 root
```

Note – In the above example, only root access is allowed between the two systems within the Solstice HA configuration.

7.7 Editing `/etc/syslog.conf`

Add information to `/etc/syslog.conf` to enable Solstice HA error reporting. The following two lines should be added exactly as shown. Note that there are no spaces anywhere in the two lines.

Note – A single tab separates `info` and the `/var/adm/messages` path in the first entry. Also, a single tab separates `info` and `/dev/console` in the second entry. All the entries are comma separated. Do not use spaces.

```
local7.emerg,local7.alert,local7.crit,local7.warning,local7.notice,local7.info    /var/adm/messages

local7.emerg,local7.alert,local7.crit,local7.warning,local7.notice,local7.info    /dev/console
```

The default version of the `/etc/syslog.conf` file delivered with Solaris 2.4 contains the string, `*.err`. If this entry has been removed, add an entry for `local7.err` to both the `/var/adm/messages` and `/dev/console`.

Creating the Configuration



This chapter provides step-by-step instructions for creating a Solstice High Availability configuration using the `hasetup(1M)` command. `hasetup` provides an easy-to-use interface that allows initial configuration of the Solstice HA servers.

Use the following table to locate specific information in this chapter.

<i>Overview of Tasks</i>	<i>page 8-1</i>
<i>Preparing to Run <code>hasetup</code></i>	<i>page 8-2</i>

8.1 Overview of Tasks

`hasetup` prompts or checks for the following information:

- Host names and logical hostnames
- Private network connections
- Additional public network connections
- Metadevice state database replicas on local disks
- The data services being used
- Allocating space for UNIX file system logs
- Metadevice state database replicas on multi-host disks

8.2 Preparing to Run `hasetup`

The `hasetup` command prompts for considerable information about your Solstice HA configuration. During execution, the program checks to verify the information you enter, then writes the information to the Solstice HA configuration files.

`hasetup` provides default answers for most questions, providing the information is available from the network name services. The default answers are inside brackets. To accept the default, press Return. If no default is provided, you must provide an entry.



Caution – The Solstice HA configuration files should never be edited by hand. The `hasetup` program adds the necessary information to these files. If the information is inconsistent or incorrect in these files, both Solstice HA machines may crash and services may be unavailable.

Use the worksheet provided in Table 8-1 to make a list of the hosts on each public network that are part of your Solstice HA configuration.

Table 8-1 Solstice HA Configuration Worksheet

	Server1 Public			Server2 Public		
	Physical Host Name	Logical Host Name	Network Interface	Physical Host Name	Logical Host Name	Network Interface
1st Network	_____	_____	_____	_____	_____	_____
2nd Network	_____	_____	_____	_____	_____	_____
3rd Network	_____	_____	_____	_____	_____	_____
4th Network	_____	_____	_____	_____	_____	_____
	Server1 Priv			Server2 Priv		
	Private Link Host Name		Network Interface	Private Link Host Name		Network Interface
204.152.64	_____		_____	_____		_____
204.152.65	_____		_____	_____		_____

▼ How to Run hasetup

1. **As root, invoke `hasetup` from either of the servers in the configuration by entering the following:**

```
# hasetup
Checking local /etc/nsswitch.conf file for Solstice HA compliance ... failed
Updating local /etc/nsswitch.conf file for Solstice HA compliance ... done
```

When started, `hasetup` verifies the `/etc/nsswitch.conf` file is properly configured and if the file is not, it overwrites the file with a new version.

2. **You will be asked if you want to set up additional public (client) network interfaces.**

While discovering network connections, `hasetup` may generate messages on the system console.

The following example accepts the default answer `no`. When you respond `no`, the network controllers on the additional public networks found by `hasetup` are turned off. If you respond `yes`, you will be prompted for the names of the network interfaces.

```
...
Attempting to discover local network controller config ... done
There are 5 unused network controllers on the "local" host
Do you want to setup any of those interfaces now (yes/no) [no]? no
Turning off network controller "qe0" on local host ... done
Turning off network controller "qe1" on local host ... done
Turning off network controller "qe2" on local host ... done
Turning off network controller "qe3" on local host ... done
Turning off network controller "le1" on local host ... done
...
```

3. **Enter the host name of the first private network.**

This is the name of the current host that will be used on the private network. By convention, this name is generated by using the primary host name and adding a `-priv1` suffix. `hasetup` verifies the private host name you entered is in the `/etc/hosts` file, then probes the network to discover the corresponding private host name on the sibling host.



Caution – Never set up private network connections on one server that do not exist on the sibling.

```
...
Local hostname for the 1st private net [host1-priv1]? Return
Broadcast from "host1-priv1" to discover sibling's private hostname ... done
Local hostname for the 1st private net is set to "host1-priv1"
Sibling hostname for the 1st private net is set to "host2-priv1"
...
```

4. Enter the host name of the second private network.

The process followed in Step 3 is repeated. In this example the default was again accepted. hasetup verifies the private connectivity, which allows the command to perform remote command execution on the sibling host. This allows for a single point of administration from this point forward.

```
...
Local hostname for the 2nd private net [host1-priv2]? Return
Broadcast from "host1-priv2" to discover the sibling's private hostname.done
Local hostname for the 2nd private net is set to "host1-priv2"
Sibling hostname for the 2nd private net is set to "host2-priv2"
Checking access to "host1-priv1" from "host2-priv1" ... done
Checking access to "host2-priv1" from "host1-priv1" ... done
Checking access to "host1-priv2" from "host2-priv2" ... done
Checking access to "host2-priv2" from "host1-priv2" ... done
...
```

The `/etc/nsswitch.conf` file is copied to the sibling host.

```
...
Copying /etc/nsswitch.conf to host "host2-priv1" ... done
...
```

Network information about the sibling host is now verified. You are then asked if you want to set up additional network interfaces on the sibling host. The default answer is no.

```
...
Attempting to discover sibling's network controller config ... done
There are 5 unused network controllers on the "sibling" host
Do you want to setup any of those interfaces now (yes/no) [no]? Return
Turning off network controller "qe0" on the sibling host ... done
Turning off network controller "qe1" on the sibling host ... done
Turning off network controller "qe2" on the sibling host ... done
Turning off network controller "qe3" on the sibling host ... done
Turning off network controller "le1" on the sibling host ... done
Attempting to rediscover sibling's network controller config ... done
Finding the name of the sibling host for the main public net ... done
Verifying that the two primary hostnames are on the same net ... done
...
```

5. Select the type of configuration.

You are now asked if the configuration is symmetric. If you answer no, the configuration is created with one logical host. You are then asked if the current host is the default master of the single logical host. Also, if you are creating an asymmetric configuration, you will not see Step 9, diskset allocation. In this example, the default (symmetric) configuration was selected.

```
...
Is this configuration symmetric (yes/no) [yes]? yes
...
```

6. Enter the logical host name for each server. (Only symmetric configurations will see this prompt.)

The logical host name entered is the name used by the server to communicate with the logical hosts. When the names are entered, the name service is checked for the associated IP address. If the IP address is found, it is used. If it is not found, you will be prompted for the IP address and it will

be entered in the local and remote host's `/etc/inet/host` file only. The network name service will not be updated. You must perform this at a later time manually.

```
...
Logical hostname whose default master is "host1"? logicalhost1
Logical hostname whose default master is "host2"? logicalhost2
...
```

hsetup then checks the host information for consistency and copies the network files as shown below to the sibling.

```
...
Final checking of /etc/inet/hosts ... done
Updating /etc/inet/netmasks ... done
Copying /tmp/hosts.hsetup.1960 to host "host2" as /etc/inet/hosts ...done
Copying /etc/inet/netmasks to host "host2" ... done
...
```

hsetup prints the following four informational lines.

```
...
Diskset assigned to logical host "logicalhost1" [logicalhost1]? logicalhost1
Diskset assigned to logical host "logicalhost2" [logicalhost2]? logicalhost2
Pathprefix for logical host "logicalhost1" [/logicalhost1]? /logicalhost1
Pathprefix for logical host "logicalhost2" [/logicalhost2]? /logicalhost2
...
```

7. Select the highly available data services that will be used in the configuration.

You are asked first if HA-NFS is being used in the configuration, then about HA-ORACLE if you have installed the SUNWhaor package. The default answer is yes.

```
...
Will the HA-NFS data service be used (yes/no) [yes]? Return
Will the HA-ORACLE data service be used (yes/no) [yes]? Return
...
```


The configuration files are updated and copied to the sibling host with all the information you have previously entered.

```
...
Updating "/etc/opt/SUNWhadf/hadf/hadfconfig" ... done
Updating "/etc/opt/SUNWhadf/hadf/cmm_confcdb" ... done
Copying /etc/opt/SUNWhadf/hadf/hadfconfig to host "host2" ... done
Copying /etc/opt/SUNWhadf/hadf/cmm_confcdb to host "host2" ... done
...
```

hasetup now checks and verifies the existence of Solstice DiskSuite metadvice state database replicas and compiles a list of the multi-host disks.

```
...
Checking for Solstice DiskSuite replicas on the local host ... done
Checking for Solstice DiskSuite replicas on the sibling host ... done
Verifying that disksets do not already exist ... done
Compiling the list of multi-host disks ... done
Checking the list of multi-host disks ... done
...
```

The remainder of the procedure deals with setting up the diskset.

8. Reserve space for UFS logs on multi-host disks.

hasetup provides an easy way to allocate space for UFS logs. By default, hasetup reserves space on slice 6 of each disk in the diskset for UFS logs.

UFS logging speeds reboots, provides faster local directory operations, and decreases synchronous disk writes. The UFS updates are safely recorded in a log before they are applied to the server.

For information about how large to make the UFS log, see Section 3.1.17, "File System Size and Disk Layout," on page 3-19 for a complete discussion of UFS logging.

```
...
Do you want to reserve UFS log space at this time (yes/no) [yes]? Return
How large will the UFS logs be (Mbytes) [10]? Return
...
```

9. Assign disks to disksets.

By default, the `hasetup` program automatically allocates the disks as evenly as possible between the two disksets (in a symmetric configuration). This allocation is subject to the majority drive requirements of Solstice DiskSuite. You have the option of editing the allocations of controllers to the two disksets. To accept the default allocation, enter a `y`.



Caution – If you bypass diskset creation in this step (Step 9), you must create disksets manually before the Solstice HA configuration is functional. This will require that you use the `metaset(1M)` command and enter all the device names. When you do this, make sure the diskset names match the names given to the logical hosts. Refer to the *Solstice DiskSuite 4.0 Administration Guide* for additional information.

```

===== Default Diskset Allocation =====

                relo-host2                                relo-host1

c1t0d0, c1t0d1, c1t1d0, c1t1d1                c1t2d0, c1t2d1, c1t3d0, c2t0d0
c1t4d0, c1t4d1, c1t5d0, c2t2d0                c2t0d1, c2t0d2, c2t1d0, c2t1d1
c2t2d1, c2t2d2, c2t3d0, c2t3d1                c2t4d0, c2t4d1, c2t4d2, c2t5d0
c3t0d0, c3t0d1, c3t1d0, c3t1d1                c2t5d1, c2t5d2, c3t2d0, c3t2d1
c3t4d0, c3t5d0                                c3t3d0, c3t4d1

Total number of disks: 18                        Total number of disks: 18

=====

Do you accept the default allocation of disks into disksets [y|n]? y
    
```

If you respond with an `n` at the above prompt, you are allowed to edit the allocation of disks to each of the disksets, as shown in the following screen.

You interact with the following screen using the up and down arrows on the keypad. To add disks to a different diskset, enter the drive name (*cntndn*). Shell syntax wild carding is accepted. The “Total number of disks” fields reflect any changes made during editing. When editing is complete, move the cursor to the last line and enter a *y*.

```

===== Diskset Allocation =====

                relo-host1                                relo-host2

c1t0d0, c1t0d1, c1t1d0, c1t1d1        c1t2d0, c1t2d1, c1t3d0, c2t0d0
c1t4d0, c1t4d1, c1t5d0, c2t2d0        c2t0d1, c2t0d2, c2t1d0, c2t1d1
c2t2d1, c2t2d2, c2t3d0, c2t3d1        c2t4d0, c2t4d1, c2t4d2, c2t5d0
c3t0d0, c3t0d1, c3t1d0, c3t1d1        c2t5d1, c2t5d2, c3t2d0, c3t2d1
c3t4d0, c3t5d0                        c3t3d0, c3t4d1

                Total number of disks: 18                Total number of disks: 18
===== Edit Disksets =====

Add to relo-host1:
Add to relo-host2:
Revert back to default diskset allocation [y]?

=====

Is disk allocation editing complete [y|n]? y

```

If you create a diskset that violates the configuration rules, *hassetup* will not allow the configuration to be saved. The following message will be displayed:

```

>> WARNING: disk configuration rules for one or both disksets violated! <<
Is disk allocation editing complete [y|n]?

```

If you decide to ignore the message and use the diskset allocation that you entered, the following message will display:

```

=====
Each diskset must be balanced so that the number of disks
assigned to each set from any given controller is less than
half of the total number of disks assigned to that set.

This requirement is necessary in order to guarantee that a simple
majority of disks will always be present within a diskset if any
one controller fails.

The disk assignments for controller "<controller>"
in diskset "<diskset>" violated this requirement.

Please restart "hasetup" and try again.
=====

Exiting due to fatal error ...

```

10. Repartition and populate disks with metadvice state databases.

When editing is complete, you are asked whether you are ready to have hasetup repartition and populate the disks with metadvice state database replicas. On large configurations, repartitioning the disks can take awhile.



Caution – When disks are repartitioned, existing data will be destroyed. Also, repartitioning of the disks and creation of the replicas may take some time.

```

...
All multi-host disks will now be repartitioned and
populated with metadvice state database replicas.
Is this okay (yes/no)? y
Repartitioning all multi-host disks ... done
Creating diskset "relo-host1" ... done
Populating diskset "relo-host1" ... done
Releasing diskset "relo-host1" ... done
Creating diskset "relo-host2" ... done
Populating diskset "relo-host2" ... done
Releasing diskset "relo-host2" ... done
#

```

If you respond with `y`, the disksets are automatically created and populated with metadvice state database replicas. If you respond with `n`, metadvice state database placement on multi-host disks and diskset creation are bypassed and `hasetup` configures the servers with only the data previously entered.

When the `metaset` command is invoked by `hasetup`, slice 7 on every disk in the diskset is set up to contain only metadvice state database replicas. The remainder of the disk is partitioned as slice 0. It is important to keep in mind that when multi-host disks are mirrored, the metadvice state database replicas are not mirrored, only slice 0 is mirrored. Also, if you elected to have space allocated for UFS logging, a slice of the specified size will be created on slice 6.

You may modify the disksets at any time after configuration using the `metaset` command.

Metadevice Creation on Local and Multi-host Disks



This chapter contains the information necessary for creating metadevices on multi-host and local disks.

Use the following table to locate specific information in this chapter.

<i>Overview of Tasks</i>	<i>page 9-1</i>
<i>Metadevice Considerations</i>	<i>page 9-2</i>
<i>Logging File Systems on the Local Disk</i>	<i>page 9-13</i>
<i>Setting Up Slices on Multi-host Disks</i>	<i>page 9-14</i>
<i>Setting Up Metadevices for HA-ORACLE</i>	<i>page 9-14</i>
<i>Setting Up Metadevices for HA-NFS</i>	<i>page 9-14</i>
<i>Setting Up Hot Spares</i>	<i>page 9-14</i>

9.1 Overview of Tasks

After configuring the Solstice HA systems and disksets in Chapter 8, “Creating the Configuration,” you are ready to create metadevices on the local and multi-host disks.

Refer to the *Solstice DiskSuite 4.0 Administration Guide* and *Solstice DiskSuite Tool 4.0 User’s Guide* to perform the procedures described in this chapter. These manuals are delivered as part of the SPARCcluster High Availability documentation set.

Note – When `hasetup(1M)` completed the Solstice HA configuration described in Chapter 8, it released ownership of the disksets. You must take ownership of each diskset using `metaset(1M)` before running `metatool(1M)` or the Solstice DiskSuite command line utilities to set up metadevices.

For instance, if you are on `host1` and want to take ownership of the diskset named `logicalhost1`, you would enter the following:

```
host1# metaset -s logicalhost1 -t
```

After taking ownership, if you are using the `metatool` graphical user interface, you would enter the following:

```
host1# metatool -s logicalhost1
```

9.2 Metadevice Considerations

Before creating the metadevices in your Solstice HA configuration, you should read the following subsections. These subsections offer a discussion of the following considerations:

- Diversity
- Availability
- Performance
- Takeover and metadevice state database replica issues
- Managing the metadevice name space

9.2.1 Diversity of the Configuration

In order to create a Solstice HA configuration that will provide the most highly available service the configuration must be diverse. To achieve this diversity, consider the points where interconnections can cause a failure. These include:

- Connections of the disks to the strings
- Connections from the strings to the SPARCstorage Array chassis
- Connections of the SPARCstorage Array chassis through the fibre channel to the SBus card

- Connections made between the SBus card into the System Board
- Connections from the System Board into the processor can be considered as a “tree of interconnections” that would have the processor at the highest level and the disk at the lowest level

A group of components, such as disk slices, could be said to have diversity up to but not including the level where their subtrees intersect. For example, two drives on the same string have drive diversity, but they do not have string diversity.

If two drives inside two different SPARCstorage Array chassis have drive diversity, string diversity, chassis diversity, they may or may not have SBus diversity depending on whether they share the same SBus card. If they are different SBus cards they again may or may not have System Board diversity.

Remember that the same interconnection tree usually defines another relationship related to robustness, that being fault dominance. In particular, a fault on a specific piece of hardware in the tree typically dominates all other pieces of hardware below it in the tree. For example, the piece of hardware representing a SPARCstorage Array chassis dominates the six disk strings within the chassis as well as the disks on those strings. Thus a failure of the SPARCstorage Array chassis (such as a power supply) will cause a failure of either a string or a disk. If the server has sufficient availability to tolerate the failure of a dominant component it will also tolerate the failure of the dominated components.

You can often focus your attention on the larger components in a configuration such as the SPARCstorage Array chassis. If the failure of that component can be tolerated then in general failure of all components below it in the tree can also be tolerated.

9.2.2 Availability of the Configuration

In a Solstice HA configuration the primary goal should be availability. The basis of availability is that the data must be mirrored to survive disk failures. This is because disk failures could result in permanent loss of access to the data. The submirrors should be placed in such a way that other failures such as controllers or power supplies, which result in loss of access to the data, will be tolerated.

A disk failure usually means the data on that piece of media is lost while a controller failure means that only access to the data is lost until the controller is replaced. Data generally becomes available when the access failure is repaired. That is because the controller, string, and other associated hardware do not contain any essential user data. They serve as conduits for the user data. However, this is not true of the drive itself which contains the user data.

There are situations where a controller failure (ordinarily considered an access failure) results in lost data. For instance, scribbling can occur on the disk. Usually these types of data losses are rare.

The key to protecting against loss of data is mirroring. The two submirrors should be separated sufficiently in the System Board, SBus, and string space to survive all failures. For access to a mirror to survive the failure of a specific component the submirrors must have diversity at the level of the failed component.

Sufficient availability should be provided if the submirrors connected to a Solstice HA server are on separate SPARCstorage Array chassis. You could carry the analysis to the SBus card or System Board level. But these components are dominated by power supply failures, either internal or through power loss. You are be protected against that type of failure by the failover mechanism.

It is important to remember that the faults protected by the Solstice DiskSuite software are not usually visible to the applications or the Solstice HA fault monitor. Thus, you must monitor and promptly repair the problems in order to retain availability.

When the components (slices) of one submirror (as well as its hot spares) are compared with the components of the other submirror for the common parents in the diversity tree, a failure at the junction where the two submirrors meet will result in the loss of access to the data. If that common junction is a drive (mirroring two slices on the same drive) then loss of that drive will result in loss of data, not just loss of access. This is the fundamental rule guiding metadvice availability.

It is important to consider the faults of interest. The easiest are disk failures and SPARCstorage Array chassis failures.

When a SPARCstorage Array tray fails, the online repair involves removing the tray, which introduces at least an artificial access failure while the tray is removed.

9.2.3 Performance Issues

You can enhance the disk I/O bandwidth by properly setting up metadevices. Disk striping can also increase performance by distributing sequential I/O operations across several disk devices. The key to performance here is also the “distance” between the disk devices in the System Board, SBus, and string space (see “Diversity of the Configuration” on page 9-2).

Maximizing the diversity by pushing it to the highest “level” in the tree between the slices within a stripe is the way to improve performance for a single thread sequential I/O. The result will be that most of the pieces of the server will be working in parallel. This is based on a single thread sequential reference model, typically with large transfer sizes. If you know your application’s typical transfer length you should plan your stripes so that the number of stripes multiplied by the interleave factor roughly equals the typical transfer length. For example, if the typical transfer length is 64 Kbytes, an 8-way stripe with an 8 Kbyte interleave value might be best. Alternatively, you could create a 4-way stripe with a 16 Kbyte interleave value.

Naturally, if your site is serious about performance you will perform benchmark studies, where the parameters can be varied and an ideal value for the expected reference pattern could be found.

If you have multiple threads of reference, as is the case with most time sharing systems like UNIX, you may have some inherent concurrency that operates independently of the striping of the disks. You may end up with several threads of reference, some of which may be sequential, but taken together appear somewhat random over time. Still, for things like UFS, these references do have a typical transfer size that should be factored into the striping parameters. This usually depends on the UFS block size, which is often the page size of the host architecture, and the `MAXCONTIG` UFS parameter set via the `tunefs(1M)` command.

Database applications probably have similar typical transfer lengths that will dictate the striping parameters you use.

An analysis of the common parent (see “Diversity of the Configuration” on page 9-2) will review potential performance bottlenecks. All the hardware components of the tree below the common parent presumably can operate independently. Only at the common parent does interference begin. For example two drives on the same string can operate independently so far as the disk arm position is concerned. This is because a seek operation on one disk

does not affect a seek on another disk, but transfer operations on the string can interfere and affect performance. This interaction may not reduce performance but reduced performance is the typical result.

9.2.4 *Takeover and Metadevice State Database Issues*

The most important issue to consider about metadevice state database replicas is to ensure that enough replicas are available after the various possible faults have occurred. There must be a majority of replicas which survive a particular fault or the disksets cannot be taken over. This applies at boot time also, when a takeover is uncontested. There is also a requirement that a majority of the drives in the diskset are available.

The easiest way to ensure there are an adequate number of replicas is to evenly distribute the disks across a set of three or more SPARCstorage Array chassis. If that is done for N chassis then in the limit $(N-1)/N$ of the drives will be available if one of the N chassis should fail. This results in a majority for all $N > 2$. This represents the default configuration of three SPARCstorage Arrays.

When the replicas are created by the `metaset(1M)` command, they are fairly evenly distributed across the SPARCstorage Array strings, up to a limit of 50 replicas, so the replica survival scenario parallels the drive survival scenario quite closely.

The three rules that dictate metadevice state database replica operation are:

1. A server will panic if less than one half of the replicas are in service at any point in time.
2. A server will continue operation if exactly one half of the replicas are in service. However, a reconfiguration or takeover will fail.
3. A successful takeover requires successful disk reservation on the majority of the drives in a diskset.

9.2.5 *Managing the Metadevice Name Space*

If you are configuring a large Solstice HA configuration you should take time to plan the number of metadevices to use. By default there are 128 unique metadevice names provided by Solstice DiskSuite in each diskset.

Metadevice names begin with “d” and are followed by a number in the range 0 to 127. Each UFS logging (trans) device you create will use at least seven metadevice names. Thus in a large Solstice HA configuration it is possible that you will need additional metadevice names. Refer to Appendix A of the *Solstice DiskSuite 4.0 Administration Guide* for instructions on changing the number.

During the planning you should:

- Determine the number of metadevice names needed
- Ensure there are enough names for growth
- Develop a naming convention that will allow for easy identification of devices

Ideally you should create a metadevice naming scheme for your site that will allow you to immediately recognize the slices that are part of a certain trans device. This will help you when you are performing administration on the Solstice HA configuration.

The following is an example naming convention could be used for trans devices. In the following convention, *n* can be either an empty string or a number in the range 1 to 12. Larger *n* could be used if you have edited the `/kernel/drv/md.conf` file to increase the value of `nmd`. This convention will allow for easy identification of metadevices.

- `dn0` – The trans device
- `dn1` – The UFS master mirror
- `dn2` – First submirror of UFS master mirror
- `dn3` – Second submirror of UFS master mirror
- `dn4` – The UFS log mirror
- `dn5` – First submirror of UFS log mirror
- `dn6` – Second submirror of UFS log mirror
- `dn7` – Not used (available for future mirror use)
- `dn8` – Not used (available for future mirror use)
- `dn9` – Not used (available for future mirror use)

The example trans naming convention is illustrated in Figure 9-1

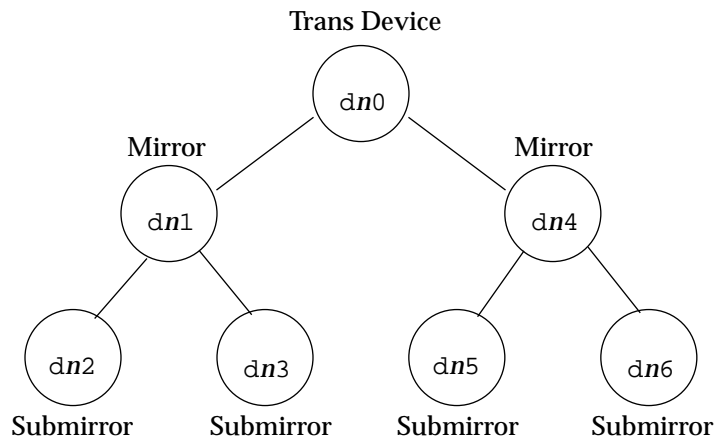


Figure 9-1 Example Trans Metadevice Naming Convention

In a configuration that is running the HA-ORACLE data service, the following example naming convention could be used for raw mirrored metadevices. In the following convention, *n* can be either an empty string or a number in the range 1 to 11. Larger *n* could be used if you have edited the `/kernel/drv/md.conf` file to increase the value of `nmd`. This convention will allow for easy identification of metadevices.

- `dn0` – Not used
- `dn1` – First mirror device
- `dn2` – First submirror of first mirror
- `dn3` – Second submirror of first mirror
- `dn4` – Second mirror device
- `dn5` – First submirror of second mirror
- `dn6` – Second submirror of second mirror
- `dn7` – Third mirror device
- `dn8` – First submirror of third mirror
- `dn9` – Second submirror of third mirror

The example raw mirrored metadvice naming convention is illustrated in Figure 9-2.




Figure 9-2 Example Raw Mirrored Metadvice Naming Convention

Using the suggested metadvice naming convention shown in the Figure 9-2, you can create an `/etc/opt/SUNWmd/md.tab` file to set up the metadvice. The file will be referred to as the `md.tab` file.



Caution – The `md.tab` file can be used safely to create the initial Solstice HA configuration. However, the `md.tab` file should not be used when adding metadevices to the configuration. It is also important to note that the `md.tab` file is not automatically updated when the configuration is changed using Solstice DiskSuite commands. Refer to the *Solstice DiskSuite 4.0 Administration Guide* for detailed information.

The ordering of lines in the `md.tab` file is not important but it is suggested you construct your file in a “top down” fashion following the diagram. The following example defines the metadevices for the diskset named `logicalhost1`, as required by Solstice HA conventions.

An example `md.tab` file would appear as follows. Note that the `#` character can be used to add annotation to the file.

```
# Example md.tab file
logicalhost1/d10 -t logicalhost1/d11 logicalhost1/d14
  logicalhost1/d11 -m logicalhost1/d12 logicalhost1/d13
    logicalhost1/d12 1 2 c1t0d0s0 c1t0d1s0
    logicalhost1/d13 1 2 c2t0d0s0 c2t0d1s0
  logicalhost1/d14 -m logicalhost1/d15 logicalhost1/d16
    logicalhost1/d15 1 1 c1t1d0s6
    logicalhost1/d16 1 1 c2t1d0s6
```

The first line defines the trans device `d10` to consist of a master (UFS) metadvice `d11` and a log device `d14`. The `-t` signifies this is a trans device, the master and log device are implied by their position after the `-t` flag.

```
logicalhost1/d10 -t logicalhost1/d11 logicalhost1/d14
```

The second line defines the master device as a mirror of the metadvice. The `-m` in this definition signifies a mirror device.

```
logicalhost1/d11 -m logicalhost1/d12 logicalhost1/d13
```

The fifth line similarly defines the log device.

```
logicalhost1/d14 -m logicalhost1/d15 logicalhost1/d16
```

The fourth line defines the first submirror of the master device as a two way stripe on two disks. Note that the usual `/dev/dsk` or `/dev/rdisk` prefix may be omitted in the `md.tab` file.

```
logicalhost1/d12 1 2 c1t0d0s0 c1t0d1s0
```


The next line defines the other master submirror.

```
logicalhost1/d13 1 2 c2t0d0s0 c2t0d1s0
```

Finally, the submirrors are defined. In this example, simple metadevices for each submirror are created. No stripes or concatenations are used.

```
logicalhost1/d15 1 1 c1t1d0s6  
logicalhost1/d16 1 1 c2t1d0s6
```

Assuming the diskset is owned appropriately, the metadevices can now be set up using the `metainit(1M)` command.

```
# metainit -s logicalhost1 -a
```

The `metainit` command will read the `md.tab` file and construct the metadevices in a “bottom up” fashion. As `metainit` constructs the mirror devices it will issue the following warning.

```
WARNING: This form of metainit is not recommended.  
The submirrors may not have the same data.  
Please see ERRORS in metainit(1M) for additional information.
```

When setting up the initial Solstice HA configuration, the form used is acceptable because we have no existing data on the submirrors. After `metainit` finishes, you must run `newfs(1M)` to create a UFS file system before using the disks.

If you have existing data on the disks that will be used for the submirrors you must dump the data before metadevice setup and restore it onto the mirror.

Note – It is possible to create a one-way mirror that contains the existing data and then use the `metattach(1M)` command to attach the other submirror. This approach ensures that both sides of the mirror contain the preexisting data. This procedure is a multistep operation. The procedure cannot be done with `metainit` alone. You must also run `metattach` after `metainit` completes. Refer to the *Solstice DiskSuite 4.0 Administration Guide* for additional details.

9.3 Creation of Metadevices on Multi-host Disks

You can either use the `metatool` command or the `md.tab` file to create metadevices on the multi-host disks.

▼ How to Create Metadevices Using `metatool`

- 1. Take ownership of the logical host where the metadevices will be set up.**
You will use the `metaset(1M)` command to take ownership of the disksets.

```
# metaset -s logicalhost -t
```

- 2. Bring up the `metatool` graphical user interface on the Solstice HA server with the display redirected to your workstation.**

You may have to use the `xhost(1)` command on your workstation to enable `metatool` to display there.

```
host1# metatool -display workstation:0 -s logicalhost
Initializing metatool... Done.
Looking up metaset relo-host1... Done.
Checking metaset ownership... Done.
Discovering drives and slices... Done.
Discovering database replicas... Done.
Discovering hot spare pools... Done.
Discovering concat/stripes...Done.
Discovering RAID devices... Done.
Discovering mirrors... Done.
Discovering trans devices... Done.
Updating mount and swap information... Done.
```

3. Use the instructions in the *Solstice DiskSuite Tool 4.0 User's Guide* to create the desired metadvice configuration on the logicalhost.

▼ How to Create Metadevices Using the `md.tab` File

1. **Create the `/etc/opt/SUNWmd/md.tab` file.**
Use the example in “Managing the Metadvice Name Space” on page 9-6 to create the `md.tab` file.
2. **Take ownership of the logical host where the metadevices will be set up.**
You will use the `metaset` command to take ownership of the disksets.

```
# metaset -s logicalhost -t
```

3. **Run the `metainit` command to set up the devices.**
The `metainit` command will use the `md.tab` file to create the metadevices.

```
# metainit -s logicalhost -a
```

9.4 Logging File Systems on the Local Disk

If you followed the instructions in Chapter 5, “Solaris 2.4 Installation,” and left six Mbytes of space on slice 4, you can now use that space to log all or some of the `/usr`, `/opt`, and `/var` file systems. Logging these file systems is not mandatory.

You cannot log the root file system or swap partition.

Note – When determining whether or not to log these file systems, keep in mind that you must disable the logging for `/usr`, `/var`, `/opt`, or any other file system used during a Solaris upgrade or installation.

To create the UFS logging for these file systems, you must make a trans device for each file system with the file system as the master and slice 4 as the log. The log will be shared by these file systems. After you have created trans devices for the three file systems you must reboot the machine.

Unlike the multi-host file systems, the local file systems need not be mirrored to be logged.

For complete instructions on creating UFS logs, refer to Chapter 6 of the *Solstice DiskSuite Tool 4.0 User's Guide*.

9.5 *Setting Up Slices on Multi-host Disks*

When you created the disksets using `hasetup`, slice 7 was allocated for a metadvice state database replica, slice 6 was set aside for UFS logging (if you selected this option during configuration), and the remainder of the disk was set aside for slice 0.

You can use the `format(1M)` command to repartition the disks, however, do not delete, resize, or move slice 7.

If you responded to queries from `hasetup(1M)` during the configuration procedure in Chapter 8, “Creating the Configuration,” slice 6 has been set aside for UFS logging.

9.6 *Setting Up Metadevices for HA-ORACLE*

HA-ORACLE may be configured to use UFS file system logging or raw mirrored metadevices.

9.7 *Setting Up Metadevices for HA-NFS*

If you are running HA-NFS on the Solstice HA systems, you must create one or more trans devices that contains a mirrored log and a mirrored master. The submirrors can consist of either concatenations or stripes.

If you decide to use `metatool(1M)` to create the metadevices, refer to the *Solstice DiskSuite Tool 4.0 User's Guide* for details on using this utility.

9.8 *Setting Up Hot Spares*

In Chapter 3, “Installation Planning,” you were asked if you wanted to use hot spares in the configuration.

For complete instructions on creating hot spares, refer to Chapter 7 of the *Solstice DiskSuite Tool 4.0 User's Guide*. For additional information about the use of hot spares in Solstice HA configuration, see Section 3.1.7, "Hot Spares," on page 3-11.

Setting Up Data Service Classes

This chapter gives instructions for setting up the Solstice HA data service classes. Use the following table to locate specific information in this chapter.

<i>Overview of Tasks</i>	<i>page 10-1</i>
<i>Setting Up HA-ORACLE</i>	<i>page 10-2</i>
<i>Setting Up HA-NFS</i>	<i>page 10-10</i>
<i>Bringing the Solstice HA Servers On Line</i>	<i>page 10-16</i>

10.1 Overview of Tasks

When the Solstice HA configuration was created by `hasetup(1M)`, you were asked which data services you were going to be running on the configuration. The valid data services available for Solstice HA include:

- HA-NFS only
- HA-ORACLE only
- Both HA-NFS and HA-ORACLE

Note – If you did not configure the desired data services when you were setting up the configuration, you must remove the `/etc/opt/SUNWhadf/hadf/hadfconfig` file and rerun `hasetup`. `hasetup` will exit when you reach the diskset allocation part of the program.

If you are running both HA-NFS and HA-ORACLE in your Solstice HA configuration, the two data services can be set up in any order.

Note – Regardless of what data servers are run in your configuration, you must perform the tasks in Section 10.4, “Bringing the Solstice HA Servers On Line,” to start Solstice HA.

10.2 Setting Up HA-ORACLE

Before setting up HA-ORACLE you must have performed the following procedures on each Solstice HA server:

1. Install the Solaris 2.4 Hardware: 3/95 distribution
2. Install the Solstice HA packages and patches
3. Reboot

Since some of the changes made during HA-ORACLE setup require a reboot, you may want to install ORACLE now before rebooting. Alternatively the ORACLE installation may be deferred until just before HA-ORACLE database setup. However, this may require an additional reboot.



Caution – In the “How to Prepare Solstice HA Servers for ORACLE Installation” on page 10-2, all the steps should be performed on both Solstice HA servers.

▼ How to Prepare Solstice HA Servers for ORACLE Installation

You should consult your ORACLE documentation before performing this procedure.

1. **Create a /etc/passwd entry for the ORACLE user ID (*oracle_id*).**
This entry is normally `oracle` for the user ID. For example, add the following line:

```
oracle:*:35:520:Oracle Database:/oracle/home:
```


2. Update the `/etc/shadow` file from `/etc/passwd` by using the `pwconv(1M)` command.

```
# pwconv
```

3. Make a `/etc/group` entry for the database administrator group. The group is normally named `dba`. You must ensure that `root` and `oracle` are members of the `dba` group. For example, add the following line:

```
dba:*:520:root,oracle
```

While the name service entries may be made in a network name service (for example, NIS or NIS+) so the information is available to HA-ORACLE clients, entries should also be made in the local `/etc` files to eliminate dependency on the network name service.

Note – This duplicate information must be kept consistent when changes are made.

4. Prepare the environment for ORACLE installation.

Choose a location for the `ORACLE_HOME` and `ORACLE_DOC` directories. If possible, the ORACLE distribution should be installed on a separate disk and mounted as `/oracle`. This is the location where the ORACLE distribution should be installed.

When they are created, the `ORACLE_HOME` and `ORACLE_DOC` directories should be owned by `oracle_id` and be in the `dba_id` group.

5. Requirements for ORACLE installation.

When first installing ORACLE, you have the option of doing a `COMPLETE SOFTWARE/DATABASE FRESH INSTALL` or `Install/Upgrade/Patch Software Only`. You should select the `Install/Upgrade/Patch Software Only` option. This is necessary because database initialization and configuration files must be modified to reflect the multi-host disks as the location for the database. Having the database reside on the multi-host disks ensures high availability of the database during a cluster reconfiguration.

Note – If you choose the `COMPLETE SOFTWARE/DATABASE FRESH INSATALLATION` to install the software, you will not have the opportunity to modify the `init.ora` file before the database is created.

6. Install the ORACLE software.

For complete instructions on installing ORACLE, refer to the *ORACLE7 Installation and Configuration Guide* and the *Oracle7 for Sun SPARC Solaris 2.x Installation and Configuration Guide*.

Note – If possible, the ORACLE software distribution should occupy its own separate file system. This would prevent ORACLE from being overwritten if the operating system has to be reinstalled.

7. Remember to create the `/opt/bin` directory as root.

This is the directory where some ORACLE binaries created by the `root.sh` script will reside. For example:

```
# mkdir -p /opt/bin
```

8. Remember to run the ORACLE `root.sh` script.

The `ORACLE_HOME/orainst/root.sh` script is created during the ORACLE installation procedure. This script makes certain programs `setuid` and sets necessary permissions for ORACLE products. Log in to the server as `oracle`, then become `root` before running the script.

```
# $ORACLE_HOME/orainst/root.sh
```

9. Set up the SQL*Net V1/TCP listener by making an entry in

`/etc/inet/services` **on both Solstice HA servers.**

Modify the `/etc/inet/services` to include an entry for the SQL*Net V1/TCP listener. An example entry would be:

```
oraserv 1525/tcp oracle
```

Note – This entry is required on Solstice HA servers even if you are using the SQL*Net V2 listener.

10. Establish means to start the ORACLE SQL*Net V1/TCP listener.

One method that can be used is to create a script in the `/etc/rc3.d` directory to start the listener. You must start the listener with the `opsrooton` option if you want to use Solaris authentication.

10.2.1 *Creating an ORACLE Database and Setting Up HA-ORACLE*

If Solstice HA is not yet running on your servers, you must now manually mount all file systems before creating an ORACLE database and setting up HA-ORACLE.

Follow both the following procedures to create and set up the initial ORACLE database in a Solstice HA configuration. If you are creating and setting up subsequent databases, only perform the procedure in “How to Create an ORACLE Database and Set Up HA-ORACLE” on page 10-6.

▼ How to Prepare the Logical Hosts for ORACLE Databases

1. Take ownership of the disksets.

`hasetup` releases the disksets upon completion so before you can perform actions on the disksets you must take ownership.

2. Set up raw mirrored metadevices on both servers.

If you will be using raw mirrored metadevices to contain the databases you will want to change the owner, group, and mode of each of the raw mirrored metadevices. If you are not using raw mirrored metadevices, skip to Step 3. Instructions for creating mirrored metadevices are provided in Chapter 9, “Metadevice Creation on Local and Multi-host Disks.” If you are creating raw mirrored metadevices, you must run the following commands for each metadevice.

```
# chown oracle_id /dev/md/logicalhost/rdisk/dn
# chgrp dba_id /dev/md/logicalhost/rdisk/dn
# chmod 600 /dev/md/logicalhost/rdisk/dn
```



Caution – While ORACLE supports raw I/O to both raw physical devices and raw metadevices (mirrored or non-mirrored), Solstice HA only supports raw ORACLE I/O on raw mirrored metadevices. You may not use devices such as `/dev/rdisk/c1t1d1s2` to contain ORACLE data under Solstice HA.

3. Configure the logging UFS file systems.

Regardless of whether the databases will reside on raw mirrored metadevices or on UFS file systems, at least one logging UFS file system must be created to contain the ORACLE configuration files associated with the databases. You may want to create additional logging UFS file systems to contain your ORACLE databases. Each of these UFS file systems must have an entry in the `/etc/opt/SUNWhadf/hadf/vfstab.logicalhost` file.

Use the procedure in “How to Set Up a Multi-host UFS File System” on page 10-13 to set up the UFS file systems and `vfstab.logicalhost` files.

4. Mount the UFS file systems.

Use the `mount (1M)` command to mount all UFS file systems in Step 3.

▼ **How to Create an ORACLE Database and Set Up HA-ORACLE**

1. Prepare database configuration files.

By default database configuration files are created in `$ORACLE_HOME/dbs`. It is essential that these files do not reside in the default location; they must reside on the multi-host file system. Configuration file location can be forced during database creation by explicitly assigning the various `ORACLE_SID` specific files and directories in the `init.ora` file used to build the database. The same assignments should be used in all `init.ora` files used with the database (for example, those used for database building and running).

Within the `init.ora` file, the assignments include the file name assignments for `control_files`, `background_dump_dest`, `core_dump_dest`, `log_archive_dest`, and `user_dump_dest`. The only `ORACLE_SID` specific file that should reside in `$ORACLE_HOME/dbs` is the `sgadef$ORACLE_SID.dbf` file which is recreated at each database start up. HA-ORACLE will start the database using the `pfile` specified in the `haoracle_databases` file; this `pfile` is on the multi-host file system and it in turn assigns the other configuration files to multi-host paths.

Note – If you are using Solaris authentication for database logins, the `remote_os_authent` variable in the `init.ora` file must be set to `TRUE`.

2. Create the database.

During creation, ensure that all configuration and database files are placed on the multi-host disks.

Start the ORACLE installer (`orainst`) and select the `Create New Database Objects` option. During the `orainst` session, place all the database files on the multi-host disks. Be sure to override the default file locations provided by the ORACLE installer.

Ensure that the file names of the control files you specify match the file names in your configuration files.

Alternatively, the database can be created through the ORACLE `sqldba` command. Two points to consider if you are using `sqldba` are:

- All of the `*.ora` files must be placed on the `logicalhost`.
- You must also run the catalog scripts that create the `v_$sysstat` table. This table is used by the Solstice HA fault monitoring scripts.

3. Make entries for the SIDs of all database instances.

The `SID` of the instance associated with your database must be in the file `/var/opt/oracle/oratab` on both Solstice HA servers. This file must be kept current on both Solstice HA servers for a failover to succeed (the HA-ORACLE databases, that is `SIDs`, must exist in `/var/opt/oracle/oratab` on both servers). This must be done manually as `SIDs` are added or removed.

All entries in the `/var/opt/oracle/oratab` file should have the `:N` option specified to ensure that the instance will not automatically start.

Note – The HA-ORACLE fault monitoring scripts parse the `init.ora` files. This script assumes a white space appears on either side of the `=` sign. While ORACLE does not require the spaces, HA-ORACLE does.

4. Enable access for the user and password to be used for fault monitoring.

In the following examples the user is `scott` and the password is `tiger`. Note that the user and password pair must agree with those used in Step 6 if you are using ORACLE authentication.

You can enable access by entering the following script into the `sqldba(1M)` command.

```
# sqldba lmode=y <<EOF

connect internal;
  grant connect, resource to scott identified by tiger;
  grant select on v_$sysstat to scott;
  grant create session to scott;
  grant create table to scott;
  alter user scott quota 1M on system;
disconnect;

exit;
EOF
#
```

5. Optionally grant permission for the database to use Solaris authentication.

The following example entry will enable you to use Solaris authentication.

```
# sqldba lmode=y <<EOF

connect internal;
  create user ops$root identified by externally;
  grant select on v_$sysstat to ops$root;
  grant create session to ops$root;
  grant create table to ops$root;
  alter user ops$root quota 1M on system;
disconnect;

exit;
EOF
#
```

6. Update the `haoracle_databases(4)` file.

The `haoracle_databases` files must be updated using the `haoracle insert` command, after database creation. This file is owned by root and has a mode of 600 to protect the information.

```
# haoracle insert $ORACLE_SID logicalhost 60 10 120 300 scott/tiger \  
/logicalhost1/.../init$ORACLE_SID.ora
```

The above command line includes the following:

- `haoracle insert` - The command and subcommand
- `$ORACLE_SID` - The name of the ORACLE database instance.
- `logicalhost1` - The logical host serving `$ORACLE_SID` (not the physical host)
- `60 10 120 300` - These parameters are for probe cycle time of 60 seconds, a connectivity probe cycle count of 10 seconds, a probe time out of 120 seconds, and a restart delay of 300 seconds.
- `scott/tiger` - The user and password to be used for fault monitoring. This must agree with the permissions granted in Step 4. To use Solaris authentication enter a / instead of the user name and password.
- `/logicalhost1/.../init$ORACLE_SID.ora` - The pfile to be used to start up the database. This must be on a multi-host logging UFS file system. It should not be available via HA-NFS. The `init$ORACLE_SID.ora` file is the `init.ora` file referred to in Step 1.

7. Verify the changes.

Run the `hacheck(1M)` command on one of the servers to check the changes. This command will validate the changes made to the `haoracle_databases` file.

```
# hacheck
```

8. Bring the HA-ORACLE database into service.

If the `hacheck(1M)` command reports no problems you can bring the HA-ORACLE database into service. You can bring the database into service by running `haoracle(1M)`.

```
# haoracle start database_instance
```



Caution – From the time the ORACLE database instance is brought into service, the monitoring should be enabled.

If you have not started Solstice HA but have manually taken ownership of the diskset and mounted the file systems, you should unmount the file systems and release the disksets before starting Solstice HA.

HA-ORACLE is now ready to begin servicing HA-ORACLE clients.

10.2.2 HA-ORACLE Client Set Up

No special configuration should be necessary for HA-ORACLE clients.

Note – ORACLE client-server connections are stateful and will not survive an HA-ORACLE switchover. The client application must be prepared to cope with disconnects and reconnect or recover as appropriate. A transaction monitor may simplify the application. Further, HA-ORACLE server recovery time is application dependent.

Clients must always refer to the database using the logical hostname and not the physical hostname. Except for start-up times the database should always be available if the logical host is responding on the network.

10.3 Setting Up HA-NFS

The instructions in this subsection assume you have performed all file system planning (see Chapter 3, “Installation Planning”), created the disksets, and set up the trans (UFS logging) devices. These instructions show how to set up HA-NFS in a configuration where both logical hosts offer HA-NFS services.

Note – When setting up HA-NFS, plan to have an extra multi-host file system that is not shared under NFS. This file system will be used for the Solstice HA locking and status monitor data. This file system must be mounted on */logicalhost* and should have at least 10 Mbytes of space, which is sufficient to provide locking for at least 1,000 NFS clients.

You will use the `hafstab(1M)` command to edit the logical host's `vfstab.logicalhost` and `dfstab.logicalhost` files. You run this command on only one server. The files are automatically distributed to the other server.

▼ How to Share HA-NFS File Systems

1. Create your multi-host UFS file systems.

Use the procedure in “How to Set Up a Multi-host UFS File System” on page 10-13 to create the multi-host file system.

2. Use the `hafstab` command on one Solstice HA server to edit the `dfstab.logicalhost` files.

The `dfstab.logicalhost` file is in `dfstab` format. The `hafstab` command allows editing the `dfstab` files then distributes the file to both servers in the configuration. The editor defined as the `EDITOR` environment variable is used. A template of the `dfstab` file is presented for editing if no version of the file exists.

Start `hafstab` with the following argument and respond to the prompts.

```
host1# hafstab dfstab.logicalhost1
Failed to copy host2 version.
Type 'y' to edit host1 version; type 'n' to exit. [y|n] y
```

In the following example, we will not export the file system on `d0`.

After the editor begins, the following `dfstab` template file is displayed:

```
#
#ident "@(#)dfstab.tmpl 1.10 95/07/26 SMI"
#
# Copyright 26 Jul 1995 Sun Microsystems, Inc. All Rights Reserved.
#
# This file must be replicated on both servers of the HADF configuration;
# use hafstab(1M) to edit and distribute.
#
# Place the share(1M) commands at the end of this file.
#
# share [-F fstype] [ -o options] [-d "<text>"] <pathname> [resource]
# .e.g,
# share -F nfs -o rw=engineering -d "home dirs" /export/home2
#
# Notes: If you use the "rw", "rw=", "ro", and/or "ro=" options to the
# share command, then HA-NFS fault monitoring will work best if you
# grant access to both of the HA server hosts, to ALL of their
# hostnames. Look at your hadfconfig file,
# /etc/opt/SUNWhadf/hadf/hadfconfig. For all of the HOSTNAME
# lines, all of the physical hostnames, which are the 2nd column and
# 4th column of the HOSTNAME line(s), should be granted access.
# If you use netgroups in the share command (rather than names of
# individual hosts) please add all of those HA server hostnames to
# the appropriate netgroup.
# Also, ideally, both read and write access should be granted
# to all of the HA server hosts' hostnames, to enable the NFS fault
# probes do a more thorough job.
#
# Example share command lines:
# share -F nfs -d "description" /<pathprefix>/1
# share -F nfs -d "description" /<pathprefix>/2
```

3. Make the following changes to the `dfstab` file:

- a. Remove the comment character (#) from the share lines.
- b. Change the *description* to be a description of the resources being shared. In this example, the field includes the logical host name, the use (`fs`), and the mount point name.
- c. Change the `<pathprefix>` entries to be the logical host name, by convention.

d. Correct the mount point names.

The resulting file will look like the following.

```
share -F nfs -d "logicalhost1 fs 1" /logicalhost1/1
share -F nfs -d "logicalhost1 fs 2" /logicalhost1/2
```

4. Save this file and quit the editor.

5. Respond to the following question.

If you respond with a `y`, the edited version of the `dfstab` file will be distributed to the sibling host.

```
Would you like to save these changes locally and distribute them
to host1 now? [y|n] y

changes to dfstab.logicalhost1 are saved on host1.
Copying dfstab.logicalhost1 to spin ... copy complete.
host1#
```

6. If you have created a symmetric configuration, repeat all the previous steps for the second logical host.

▼ How to Set Up a Multi-host UFS File System

1. Invoke `newfs(1M)` on all trans devices.

Run `newfs` on each diskset in the configuration before Solstice HA is running. Before you can run `newfs` on the disksets, you must take ownership using the `metaset(1M)` command. For example, the following list of `metatrans` devices will be used for file systems that will be exported to clients on `logicalhost1`. (The `metatrans` devices were created using `hasetup`, `metatool(1M)`, or the `metainit(1M)` command.)

```
host1# metaset -s logicalhost1 -t
host1# newfs /dev/md/logicalhost1/rdsk/d0
host1# newfs /dev/md/logicalhost1/rdsk/d1
host1# newfs /dev/md/logicalhost1/rdsk/d2
host1# metaset -s logicalhost1 -r
```

Note – To save time, the `newfs` commands can be run in parallel.

2. Use the `hafstab(1M)` command to edit the `vfstab.logicalhost` files on each server in the Solstice HA configuration.

Note – `hafstab` automatically updates files on both servers so it need only be run on one of the Solstice HA servers.

The `vfstab.logicalhost` file is in `vfstab` format. The `hafstab` command allows editing a copy of the `vfstab` files, performs a limited sanity check on the `vfstab` files, then distributes the file to both servers in the configuration. The editor defined as the `EDITOR` environment variable is used. A template of the `vfstab` file is presented for editing if no version of the file exists.

3. Start `hafstab` with the correct argument and respond to the prompts.

```
host1# hafstab vfstab.logicalhost1
Failed to copy host2 version.
Type 'y' to edit host1 version; type 'n' to exit. [y|n] y
```

4. After the editor begins, the following template file is displayed:

```
#
#ident "@(#)vfstab.tmpl 1.6 94/07/23 SMI"
#
# Copyright 26 Jul 1995 Sun Microsystems, Inc. All Rights Reserved
#
# This file must be replicated on both servers of the HADF configuration;
# use hafstab(1M) to edit and distribute.

#device          device          mount   FS   fsck  mount mount
#to mount        to fsck          point  type pass all  options
#
#/dev/md/<setname>/dsk/d0 /dev/md/<setname>/rdsk/d0 /<pathprefix>   ufs  -  yes  -
#/dev/md/<setname>/dsk/d1 /dev/md/<setname>/rdsk/d1 /<pathprefix>/1 ufs  -  yes  -
#/dev/md/<setname>/dsk/d2 /dev/md/<setname>/rdsk/d2 /<pathprefix>/2 ufs  -  yes  -
```

Make the following changes to the `vfstab` file:

- a. Remove the comment character (`#`) from the last three lines.
- b. Change `<setname>` to be the diskset name to be the diskset name, which must be the same as the logical hostname.
- c. Correct the metadvice names (`d0`, `d1`, `d2`) as needed.
- d. Correct `<pathprefix>` entries to be the logical hostname, by convention. Also correct the mount points as necessary.

Create additional entries as planned for your configuration.

The resulting file will look like the following:

#device	device	mount	FS	fsck	mount	mount
#to mount	to fsck	point	type	pass	all	options
#						
/dev/md/logicalhost1/dsk/d0	/dev/md/logicalhost1/rdisk/d0	/logicalhost1	ufs	-	yes	-
/dev/md/logicalhost1/dsk/d1	/dev/md/logicalhost1/rdisk/d1	/logicalhost1/1	ufs	-	yes	-
/dev/md/logicalhost1/dsk/d2	/dev/md/logicalhost1/rdisk/d2	/logicalhost1/2	ufs	-	yes	-

5. Save this file and quit the editor.

6. Respond to the following question.

If you respond with a `y`, the edited version of the `vfstab` file will be distributed to the sibling host.

```
Would you like to save these changes locally and distribute them
to host1 now? [y|n] y

changes to vfstab.logicalhost1 are saved on host1.
Copying vfstab.logicalhost1 to spin ... copy complete.
host1#
```

7. If you have created a symmetric configuration, repeat all the previous steps for the second logical host.

10.4 Bringing the Solstice HA Servers On Line

After you have configured the appropriate data services on your Solstice HA servers, you must create a link in the `/etc/rc3.d` directory and start the Solstice HA software.

Note – Unless you complete the steps in this subsection, the data services will not start automatically at system boot time.

The steps to follow to bring the Solstice HA servers online are:

- 1. On each server, create a hard link from `/etc/rc3.d/S20SUNWhadf` to `/etc/init.d/SUNWhadf`.**

This line enables automatic starts Solstice HA at boot time, when the system is brought up multi-user. It will not automatically start Solstice HA when the server is brought up in single user mode. On each server, enter the following:

```
# ln /etc/init.d/SUNWhadf /etc/rc3.d/S20SUNWhadf
```

- 2. On each server, start the Solstice HA services.**

Start the services by entering the following on each server:

```
# /etc/init.d/SUNWhadf start
```



Caution – The above command must be entered simultaneously on both servers to ensure that each logical host is served by its default master. This is very important for symmetric configurations.

Verification and Validation

After the Solstice HA systems are configured, you must verify and validate the configuration to ensure high availability.

Use the following table to locate specific information in this chapter.

<i>Overview of Tasks</i>	<i>page 11-1</i>
<i>Running the hacheck Command</i>	<i>page 11-2</i>
<i>Running the haswitch Command</i>	<i>page 11-2</i>
<i>Performing Physical and Manual Tests</i>	<i>page 11-3</i>
<i>Performing the Fault Detection Test</i>	<i>page 11-4</i>
<i>Verify HA-ORACLE Setup</i>	<i>page 11-4</i>
<i>Verify SPARCstorage Array Firmware</i>	<i>page 11-4</i>
<i>Power Sequence Testing a SPARCcluster 1000</i>	<i>page 11-5</i>
<i>Power Sequence Testing a SPARCcluster 2000</i>	<i>page 11-6</i>

11.1 Overview of Tasks

Verifying and validating the Solstice HA configuration involves:

- Running the `hacheck(1M)` command
- Running the `haswitch(1M)` command (this is done as part of the physical and manual tests)
- Performing physical and manual tests

- Testing the fault detection
- Verify HA-ORACLE setup
- Verifying the power sequencing

11.2 *Running the `hacheck` Command*

You should run the `hacheck` command on both servers in the Solstice HA configuration. The command automatically performs the following on both servers in the Solstice HA configuration:

- Checks that HA-ORACLE is installed properly.
- Checks and verifies all Solstice HA configuration information.

Run `hacheck` by entering the following command on both servers:

```
host1# hacheck
```

Error messages reported by `hacheck` are documented in Appendix A, “Error Messages.”

11.3 *Running the `haswitch` Command*

The `haswitch` command transfers the specified disksets along with the associated data services and logical IP addresses to the specified server.

Run the `haswitch` command from each side of the Solstice HA configuration. The following example command line associated the diskset named `logicalhost1`, along with its data services and logical IP addresses with physical host `host2`.

```
host1# haswitch host2 logicalhost1
```


11.4 Performing Physical and Manual Tests

You must conduct both physical and manual tests of the Solstice HA configuration to make sure one system will take over if the other fails. You must perform each of the tests listed:

1. Make sure diskset ownership moves from one server to the other when the power is turned off on the default master.

Perform the following steps to conduct this test:

- a. **Turn off the power on one of the servers.**
You will begin seeing `be0` error messages on the console of the sever that remains running.
- b. **Verify the other system has taken ownership of the diskset that was mastered by the server you turned off.**
- c. **Turn the power back on. Let the system reboot. The system will automatically start the membership monitor software. The host then rejoins the configuration.**
- d. **Perform a switchover (using `haswitch`) to give ownership of the diskset back to the default master.**
- e. **Repeat the procedure, by turning off the power to the second server.**

2. Ensure a takeover occurs when one system is halted.

Perform the following steps to conduct this test:

- a. **As root, invoke `uadmin(1M)` on one host. For example:**

```
host1# /sbin/uadmin 1 0
Program terminated
Type help for more information
<#0> ok
```

- b. **Make sure the sibling host has taken over the diskset that was mastered by system you halted.**
- c. **Reboot the server.**
- d. **Perform a switchover (using `haswitch`), moving ownership of the diskset back to the default master.**
- e. **Repeat the procedure on the other server.**

11.5 *Performing the Fault Detection Test*

There are several ways to test the fault detection monitor that runs on each system in a Solstice HA configuration.

1. Disconnect one of the private network connections.

You can verify that this action is recognized by the Solstice HA software when error messages are displayed on the Solstice HA consoles on each host or in the `/var/adm/messages` file. This fault does not result in a takeover.

2. Reconnect the private network connection.

3. Disconnect all the public network connections on one of the servers.

It may take several minutes before error messages will be generated and a takeover occurs.

4. Reconnect the network and wait for the server to reboot.

11.6 *Verify HA-ORACLE Setup*

Perform the following verification tests to ensure the HA-ORACLE installation was performed correctly.

- Run `sqlplus(1M)` to connect from one server to the other. For instance:

```
host1# sqlplus user/password@T:logicalhost2:test_db
```

- Run a sample application that connects to the database from a remote system.
- Make sure that a sample application can access the database, regardless of which Solstice HA server owns the logical hosts.

11.7 *Verify SPARCstorage Array Firmware*

Perform a verification test to ensure the SPARCstorage Arrays have the latest version of the firmware installed. This test can be performed on either server. It only needs to be run on one of the servers.

For each SPARCstorage Array *cn*, run the command:

```
host1# ssaadm display cn
```

Check the output for a line documenting the firmware revision. The firmware must be at `Firmware Rev: 2.3` or higher for Solstice HA.

There may be later SPARCstorage Array patches than those shipped with Solstice HA (102432-07). Consult the file `/var/sadm/patch/102432-nn/README.102432.nn` for instructions about upgrading the firmware and the appropriate firmware revision.

11.8 Power Sequence Testing a SPARCcluster 1000

Before putting your Solstice HA configuration into service you should perform a basic validation for the power cabling and sequencing. In a SPARCcluster 1000 configuration, if both SPARCserver 1000s, the three SPARCstorage Arrays, and the terminal concentrator are all powered through the cabinet power sequencer, your system is vulnerable to outage if your power source fails or if some component internal to the power sequencer fails.

It is only possible for the configuration to recover if the power is cabled correctly. Perform the following test to verify the proper power cable connections.

1. Turn off the key switch at the front of the cabinet.
2. Turn off the circuit breaker at the back of the cabinet.
3. Confirm that both the SPARCserver 1000s, all SPARCstorage Arrays, and the terminal concentrator do not have power.
4. Turn on the circuit breaker.
5. Confirm that the terminal concentrator has power even though the key switch is in the off position.

No other component should have power at this step. It is important for the for the terminal concentrator to remain on even if the key switch is in the off position. This ensures that existing open connections to the terminal concentrator are not interrupted when the key switch is turned off. This also allows the power on messages from the SPARCserver 1000s to be visible on

those connections. If the terminal concentrator is not properly connected there may be up to a 90-second period of time when it is booting that messages are not displayed.

- 6. Turn on the key switch.**
- 7. Confirm that all SPARCstorage Arrays power up as soon as the key switch is turned on.**

All SPARCstorage Arrays should be on the “switched 1” outlets. It is also important that the SPARCstorage Arrays be ready before Solaris begins accessing the device. If the SPARCstorage Arrays are not ready, Solaris cannot access them until the next reboot. The 20-second head start helps ensure that the SPARCstorage Arrays are ready before Solaris attempts to access them.

The position of the DIAG switch on the SPARCstorage Arrays is critical to ensuring they are ready before the SPARCserver 1000s begin booting the Solaris operating environment. The correct position of the switch is DIAG, which is opposite the DIAG EXT setting.

- 8. Confirm that the SPARCserver 1000s power up approximately 20 seconds after the SPARCstorage Arrays.**
- 9. Ensure that all fans are powered at or before the time the SPARCserver 1000s obtain power.**

Depending on the cabling, the fans may come on earlier in this procedure. If the fans are not running you should shut down the system now and check the power cabling.

11.9 Power Sequence Testing a SPARCcluster 2000

If the SPARCstorage Arrays and the terminal concentrator are powered from the SPARCcenter 2000 power controller, you should perform the power sequence testing in this section.

- 1. Turn off the key switch at the front of each SPARCcenter 2000.**
- 2. Turn off the circuit breakers at the rear of the SPARCcenter 2000 cabinets.**
- 3. Confirm that there is no power to either SPARCcenter 2000s, the local disks, the SPARCstorage Arrays, or the terminal concentrator.**

- 4. Turn on the circuit breaker in cabinet that has the two SPARCstorage Arrays and the terminal concentrator.**
- 5. Verify that the SPARCstorage Arrays, the terminal concentrator, and the SPARCcenter 2000 local disks power up.**
If these do not power up, ensure that the power controller remote/local switch is in the local position.
- 6. Turn on the key switch on the front of the cabinet.**
The SPARCcenter 2000 should receive power and begin power-on self testing. It is important for the SPARCstorage Arrays and the terminal concentrator power to remain independent for the SPARCcenter 2000 key switch. This ensures that open connections to the server console ports remain even while the SPARCcenter 2000 is powered down for service.
- 7. Turn on the circuit breaker at the rear of the other cabinet.**
- 8. Verify that the SPARCstorage Array and the SPARCcenter 2000 local disks power up.**
If these do not power up, ensure that the power controller remote/local switch is in the local position.
- 9. Turn on the key switch on the front of the other cabinet.**
- 10. Verify that the SPARCcenter 2000s power up correctly.**

Note – The local disks in the SPARCcenter 2000 configuration remain powered even though the processor chassis has no power. If you will be removing the SCSI cables to the local disks they should be individually powered off using the desktop storage pack power switch. Unless the SPARCcenter 2000 power controller must be serviced, key switch power off should suffice for most service needs such as SBus card replacement.

In the event of an external power failure, the SPARCstorage Arrays, SPARCcenter 2000s, and terminal concentrator will all obtain power simultaneously. In this configuration no head start is afforded the SPARCstorage Array before the SPARCcenter 2000s begin to boot. The position of the DIAG switch on the SPARCstorage Array is critical to ensuring the SPARCstorage Array is ready before the SPARCcenter 2000s are booting the Solaris operating environment. The correct position of this switch is DIAG, which is the opposite of DIAG EXT switch.

Error Messages



This appendix contains the error messages displayed by `hasetup(1M)` and `hacheck(1M)` commands. Use the following table to locate specific information in this appendix.

<i>hacheck Error Messages</i>	<i>page A-1</i>
<i>hasetup Error Messages</i>	<i>page A-9</i>

A.1 `hacheck` Error Messages

The `hacheck` command displays the following messages in the console. Use the information here to correct the error.

```
did not find framework configuration file (hadfconfig)
```

`hacheck` could not access the `hadfconfig(4)` file. The `hadfconfig` file should be found in `/etc/opt/SUNWhadf/hadf`. Check for access to the file. The original installation may be incomplete or corrupted. Verify the necessary packages are in place using the `pkgchk(1M)` command.

```
error(s) found in performing "configuration" checks
```

One or more errors or warnings occurred while performing the configuration category of checks. Refer to other messages for the source of the problem.

```
HA_BASEDIR and HA_FILES settings are inconsistent
```

There is an error in the `hadfconfig(4)` file. Contact your service representative for assistance.

```
not enough memory
```

You must reconfigure the system with additional swap space to correct this error.

```
bad keyword value for base directory directory_name
```

There is an error in the `hadfconfig` file. Contact your service representative for assistance.

```
base directory must begin with /
```

There is an error in the `hadfconfig` file. Contact your service representative for assistance.


```
no HOSTNAME entries found in config_file
```

There is an error in the `hadfconfig` file. Contact your service representative for assistance.

```
cannot find local host name in the HOSTNAME entries
```

There is an error in the `hadfconfig` file. Contact your service representative for assistance.

```
cannot find remote host name in the HOSTNAME entries
```

There is an error in the `hadfconfig` file. Contact your service representative for assistance.

```
cannot find address in /etc/etc/hosts for host hostname
```

There is an error in the `hadfconfig` file. Contact your service representative for assistance.

```
netmask in /etc/netmasks conflicts with ifconfig for host hostname
```

There is an error in the `hadfconfig` file. Contact your service representative for assistance.

```
cannot find netmask in /etc/inet/netmasks for host hostname
```

There is an error in the `hadfconfig` file. Contact your service representative for assistance.

```
more than one network interface is assigned to host
```

Solstice HA systems use `/etc/hostname.interface` files at boot time to assign host names to network interfaces. This error message indicates that the same *hostname* is found in more than one of these files. To correct the error, remove or correct the incorrect `/etc/hostname.interface` files and reboot the systems.

```
attempt to set logical network interface (/etc/hostname.name)
```

Solstice HA systems use `/etc/hostname.name` files at boot time to assign host names to network interfaces. If the name position of the file name contains a colon (:), this is interpreted as an attempt to set a logical or virtual address for the host. To correct the error, remove the indicated `hostname.name` file.

```
there is no hostname.name file for local host hostname
```

Solstice HA systems use `/etc/hostname.name` files at boot time to assign host names to network interfaces. There should be a `hostname.name` file for all LOCAL physical public and private interfaces found in the HOSTNAME and PRIVATELINK keyword parameter declarations in the `hadfconfig` file. To correct this error, create the necessary `/etc/hostname.name` file.

```
network interface mismatch in PRIVATELINK entry entry for host host
```

There is an error in the `hadfconfig` file. Contact your service representative for assistance.

```
file hostname.name describes unknown host host
```

Solstice HA systems use `/etc/hostname.name` files at boot time to assign host names to network interfaces. If a host name which is not known as a Solstice HA system is found in any of the `/etc/hostname.name` files, this warning message is issued. You must remove the file and reboot the system.

```
trying to default lookup facility for host host
```

There is an error in the `hadfconfig` file. Contact your service representative for assistance.

```
the default lookup facility failed to get address for host host
```

There is an error in the `hadfconfig` file. Contact your service representative for assistance.

```
network naming service required to find address for host
```

There is an error in the `hadfconfig` file. Contact your service representative for assistance.

```
misplaced - found in HOSTNAME entries entries
```

There is an error in the hadfconfig file. Contact your service representative for assistance.

```
duplicate local host names found in HOSTNAME entries entries
```

There is an error in the hadfconfig file. Contact your service representative for assistance.

```
duplicate remote host names found in HOSTNAME entries entries
```

There is an error in the hadfconfig file. Contact your service representative for assistance.

```
unable to generate list of local physical public hosts
```

There is an error in the hadfconfig file. Contact your service representative for assistance.

```
unable to generate list of remote physical public hosts
```

There is an error in the hadfconfig file. Contact your service representative for assistance.

```
unable to determine local host name from entries
```

There is an error in the hadfconfig file. Contact your service representative for assistance.

```
unexpected local host name found for this configuration
```

There is an error in the hadfconfig file. Contact your service representative for assistance.

```
unable to determine local host names from names
```

There is an error in the hadfconfig file. Contact your service representative for assistance.

```
unexpected local host name found for this configuration
```

There is an error in the hadfconfig file. Contact your service representative for assistance.

```
unable to determine remote host name from names
```

There is an error in the hadfconfig file. Contact your service representative for assistance.

```
unexpected remote host name found for this configuration
```

There is an error in the hadfconfig file. Contact your service representative for assistance.

```
unable to determine remote host name from names
```

There is an error in the hadfconfig file. Contact your service representative for assistance.

```
unexpected remote host name found for this configuration
```

There is an error in the hadfconfig file. Contact your service representative for assistance.

```
no PATHPREFIX entries found in filename
```

There is an error in the hadfconfig file. Contact your service representative for assistance.

```
found number PATHPREFIX entries
```

There is an error in the hadfconfig file. Contact your service representative for assistance.

```
PATHPREFIX - bad number of lines num for asymmetric configuration
PATHPREFIX - bad number of lines num for symmetric configuration
DISKSET - bad number of lines num for asymmetric configuration
DISKSET - bad number of lines num for symmetric configuration
```

There is an error in the `hadfconfig` file. Contact your service representative for assistance.

A.2 `hasetup` *Error Messages*

The following messages are reported by the `hasetup` command.

```
USAGE: hasetup
```

`hasetup` does not accept command line arguments.

```
Initialization errors ...
```

Errors were encountered during the initialization phase of `hasetup`. Such errors are likely the result of a lack of swap space. Check the system for the amount of available swap. If swap space is low, correct the problem and restart `hasetup`. If swap space is not the problem, contact your service provider.

```
Cannot initialize directory names ...
```

Errors were encountered during the initialization phase of `hasetup`. Such errors are likely due to either a lack of swap space or the required `SUNWhadf` directories were not found. To correct this error, check the `SUNWhadf` package using the `pkgchk(1M)` command to ensure it is installed correctly. Verify that none of the following environment variables are set: `HA_BASEDIR`, `HA_BIN`, `HA_FILES`, `HA_VAR`, `CM_BIN`, `FM_PROGS`, or `CT_PROGS`. Also check the server for the amount of swap space. If the swap space is low, correct the problem and restart `hasetup`.

```
Cannot locate bin directory ...
```

This message is displayed when either an internal error has occurred or there is not enough swap space. Check the server for the amount of available swap space. If the swap space is low, correct the problem and restart `hasetup`. If swap space is not the problem, contact your service provider.

```
You must be root to run "hasetup"
```

You must be the superuser in order to run `hasetup`. Use the `su(1M)` command to become root and restart `hasetup`.

```
"hasetup" is an interactive program  
Only keyboard input is accepted
```

`hasetup` can only be run interactively. Input is only accepted from the keyboard. If this message is printed, `hasetup` finds that standard input is not associated with a `tty` device. Restart `hasetup` from your shell prompt. Do not attempt to redirect standard input away from your terminal or terminal window.

```
"hasetup" is an interactive program  
Output must go to your screen
```

`hasetup` is designed to run interactively and must send output to your terminal or terminal window. If this message is displayed, `hasetup` has found that standard output is not associated with a `tty` device. You should restart `hasetup` from the shell prompt. Do not attempt to redirect standard output away from your terminal or terminal window.


```
Exiting before configuration is complete ...
```

You have chosen to quit `hasetup` without allowing the program to complete all its functions. You either entered a Control-D or responded with negative answers to certain questions. This message will not be displayed if `hasetup` is terminated as a result of a signal. Depending on where in the configuration setup you were when you quit the program there may have been additional instructions telling you how to manually complete certain steps. If no additional instructions were given, you may need to restart `hasetup` to complete the necessary set up procedures. Refer to the Chapter 8, “Creating the Configuration,” for additional information.

```
Exiting due to fatal error
```

`hasetup` has encountered an unrecoverable error. This message will always be preceded by other error or warning messages. Check those messages for additional information.

```
Cannot locate files directory ...
```

This error is either a result of a lack of swap space or an internal error. Check the system for the amount of available swap space. If the swap space is low, correct the problem and restart `hasetup`. If swap space is not the problem, contact your service provider.

```
The "/etc/opt/SUNWhadf/hadf/hadfconfig" file is found to already exist.
```

If this warning message is printed, `hasetup` will simply skip over the steps necessary for the creation of the `hadfconfig` file. `hasetup` will never overwrite an already existing `hadfconfig` file on the local system. Instead, it

will check the contents of the file and find the information it needs in order to continue with other parts of the set up. No action is necessary when this message is displayed.

```
The "/etc/opt/SUNWhadf/hadf/hadfconfig" file is found to already
exist
But, "/etc/opt/SUNWhadf/hadf/cmm_confcdb" is NOT found!
You may need to remove the "hadfconfig" file and try again
```

The two files are important to the configuration of Solstice HA. The files are always created at the same time by `hasetup`. This message should not be displayed, unless the `cmm_confcdb` file has been removed. You may need to remove or rename the `hadfconfig` file and restart `hasetup`.

```
Cannot find any installed HA dataservice packages!
```

At least one of the data service packages, either `SUNWhanfs` or `SUNWhaor`, must be installed on the system before `hasetup` can be run. Use the `pkgadd(1M)` command to install the desired package.

```
Attempting to discover local network controller config ... failed
```

One or more errors were encountered while `hasetup` attempted to compile a list of all the local network controllers. You should verify the following:

- The `ifconfig -a` command is working correctly.
- The `SUNWhadf` package is correctly installed on both servers and that the `/opt/SUNWhadf/bin/hadfconfig` file has the correct execute permissions.

```
WARNING: found logical net interface configured as "<ctrl>:n"  
        (<hostname>)  
WARNING: reboot is recommended
```

One or more logical network interface is configured. Reboot the server and restart the `hasetup` command.

```
Primary hostnames used in this environment must NOT be  
qualified with trailing domain names (<hostname>.<domain>)
```

A dot (.) was found in or more of the hostnames. Remove the domain name from the hostname `/etc/inet/hosts` files.

```
Attempting to rediscover local network controller  
config ... failed
```

One or more errors were encountered when `hasetup` attempted to compile a list of all the assigned local network controllers. You should verify the following:

- The `ifconfig -a` command is working correctly.
- The `SUNWhadf` package is correctly installed on both servers and that the `/opt/SUNWhadf/bin/hadfconfig` file has the correct execute permissions.

```
The primary host is not on the network!
```

The hostname for the local machine, as reported by `uname(1)` does not match any of the hosts assigned to the configured network controllers. You should verify the following:

- That you are not attempting to run `hasetup` from a standalone, nonnetworked machine.

- The `ifconfig -a` command is working correctly.
- The `SUNWhadf` package is correctly installed on both servers and that the `/opt/SUNWhadf/bin/hadfconfig` file has the correct permissions.
- The `/etc/inet/hosts` file has the correct IP address for the host.
- The `uname -n` command returns the correct name for your host. If not, update the `/etc/nodename` file and reboot the server.
- The `uname -n` output matches that of `/etc/nodename` and that it also matches one or more of the `/etc/hostname.*` files.
- There were no warning messages complaining about unsupported network controller types.

```
Two private networks are required
```

For your Solstice HA configuration to meet requirements, at least two private networks are required. Verify that the output from `ifconfig` shows at least three assigned and configured networks, one for the primary local host and two for the private networks. Also verify that all required Solaris package are correctly installed and that `ifconfig` is working. Use the `pkgchk(1M)` command to verify that the `SUNWhadf` package is correctly installed and that the `/opt/SUNWhadf/bin/haconfig` file has the correct permissions.

```
Attempting to discover remote network controller config ... failed
```

One or more errors were encountered while `hasetup` attempted to compile a list of all the remote network controllers. You should verify the following:

- The `ifconfig -a` command is working correctly.
- The `SUNWhadf` package is correctly installed on both servers and that the `/opt/SUNWhadf/bin/haconfig` file has the correct permissions.

```
WARNING: found logical net interface configured as "<ctrl>:n"  
        (<hostname>) on remote host  
WARNING: reboot of remote host is recommended
```

One or more logical network interfaces are configured on the sibling host. Reboot the sibling server and restart `hasetup`.

```
Attempting to rediscover remote network controller
config ... failed
```

One or more errors were encountered when `hasetup` attempted to compile a list of all the assigned remote network controllers. You should verify the following:

- That the use of the `rsh` command through the private link interfaces is successful.
- The `ifconfig -a` command is working correctly.
- The `SUNWhadf` package is correctly installed on both servers and that the `/opt/SUNWhadf/bin/haconfig` file has the correct permissions.

```
Not enough space
```

There is not enough swap space for `hasetup` to run. Check the system for the amount of swap space available. Correct the problem and restart `hasetup`.

```
Unable to locate remote host <remote_privatehost> in the
remote network config
```

`hasetup` could not find the remote (sibling) host. You should verify the following:

- That the use of the `rsh` command through the private link interfaces is successful.
- The `ifconfig -a` command is working correctly.
- The `SUNWhadf` package is correctly installed on both servers and that the `/opt/SUNWhadf/bin/haconfig` file has the correct permissions.
- That the `/etc/inet/hosts` file on the sibling server has the correct IP address for the host name.
- That the `/etc/inet/hosts` files on both servers are in agreement.

- That there were no warning messages complaining about unsupported network controller types.

```
Internal error while matching private hosts
```

You should contact your service provider if this message is displayed.

```
Finding the name of the remote host for the main
public net ... failed
```

The hostname for the sibling server, as reported by `uname`, does not match any of the hosts assigned to the configured network controllers. You should verify the following:

- The `/etc/inet/hosts` file on the sibling server has the correct IP address for the host.
- The `/etc/inet/hosts` files on both servers are in agreement.
- The `ifconfig -a` command is working correctly.
- The `SUNWhadf` package is correctly installed on both servers and that the `/opt/SUNWhadf/bin/haconfig` file has the correct permissions.
- The `uname -n` command on the sibling server returns the correct name for your host. If not, update the `/etc/nodename` file and reboot the server.
- The `uname -n` output matches that of `/etc/nodename` and that it also matches one or more of the `/etc/hostname.*` files.
- There were no warning messages complaining about unsupported network controller types.
- The `rsh` command works across the private link interfaces.

```
Verifying that the two primary hostnames are on
the same net ... failed
```

`hasetup` found that the primary host names for the two servers are not found to be on the same networks. You should verify the following:

- The two servers are on the same network.

- The netmasks for the two servers are identical.
- The network portion of the network IP addresses for the two servers are identical.
- All required Solaris packages are correctly installed on both servers.
- The `ifconfig -a` command is working correctly.
- The `SUNWhadf` package is correctly installed on both servers and that the `/opt/SUNWhadf/bin/haconfig` file has the correct permissions.
- The `/etc/inet/hosts` files on both servers are in agreement and correct.
- There were no warning messages complaining about unsupported network controller types.
- The `rsh` command works across the private link interfaces.

```
WARNING: network controller name mismatch for
hostnames "<hostname>" and "<hostname>"
```

The two host names given are either primary or secondary hostnames, which are both found to be assigned to the same networks. However, their network controller names do not match. Solstice HA installation procedures recommend symmetric hardware and software configurations between the two servers. This is not a supported configuration. Correct the problem by matching the network controller names with the networks on both machines. This may involve changes to the hardware configuration or to the `/etc/hostnames.*` files. You must reboot before restarting `hasetup`.

```
At least one data service must be selected
```

You must select at least one data service in order to complete configuration of the servers.

```
Updating "/etc/opt/SUNWhadf/hadf/hadfconfig" ... failed
```

`hasetup` was unable to update or create the `hadfconfig` file. You should verify the following:

- That all required Solaris packages are correctly installed.

- Use `pkgchk` to check on whether the `SUNWhadf` package is correctly installed on both servers.
- If only a partial update was found to occur to the `hadfconfig` file, remove the file before restarting `hasetup`.

```
Updating "/etc/opt/SUNWhadf/hadf/cmm_confcdb" ... failed
```

`hasetup` was unable to update or create the `cmm_confcdb` file. You should verify the following:

- That all required Solaris packages are correctly installed.
- Use `pkgchk` to check on whether the `SUNWhadf` package is correctly installed on both servers.
- If only a partial update was found to occur to the `cmm_confcdb` file, remove the file before restarting `hasetup`.

```
Bad configuration file - /etc/opt/SUNWhadf/hadf/hadfconfig
```

`hasetup` found an existing `hadfconfig` file that was corrupted. You should remove or rename the file and restart `hasetup`. If the problem persists, contact your support representative.

```
Unable to setup necessary variables
```

Errors were discovered during an internal check. These errors may be due to a lack of swap space. Check the server for available swap space. If the swap space is too small, correct the problem and restart `hasetup`. If this is not the problem, contact your support representative.


```
Bad configuration parameters!
```

Errors were discovered during an internal check. These errors may be due to a lack of swap space. Check the server for available swap space. If the swap space is too small, correct the problem and restart `hasetup`. If this is not the problem, contact your support representative.

```
Unable to determine if the cluster monitor is running
on the local host
```

An attempt to check for an instance of the cluster membership monitor on the local host failed. You should verify that all required Solaris packages are correctly installed and that the `ps(1)` command is working. Also check to make sure the `SUNWhadf` package is correctly installed and that the `/opt/SUNWhadf/bin/hadf` file has the correct execute permissions.

```
Unable to determine if the cluster monitor is running
on the remote host
```

An attempt to check for an instance of the cluster membership monitor on the local host failed. You should verify that all required Solaris packages are correctly installed and that the `ps(1)` command is working correctly. Also check to make sure the `SUNWhadf` package is correctly installed and that the `/opt/SUNWhadf/bin/hadf` file has the correct execute permissions. Make sure that the use of the `rsh` command through the private link interfaces is successful.

```
The cluster monitor is active on the local host!
```

The cluster membership monitor is running on the local host. Solstice HA does not allow the cluster membership monitor (`clustd`) to be active on either host while `hasetup` is running. You should remove the

`/etc/rc3.d/S20SUNwhadf` file on both servers. The `S20SUNwhadf` file is a hard link to `/etc/init.d/SUNwhadf`, which must be restored as the final step to installation. You must also run the following command:

```
# /etc/init.d/SUNwhadf stop
```

```
The cluster monitor is active on the remote host!
```

The cluster membership monitor is running on the sibling server. Solstice HA does not allow the cluster membership monitor (`clustd`) to be active on either host while `hasetup` is running. You should remove the `/etc/rc3.d/S20SUNwhadf` file on both servers. The `S20SUNwhadf` file is a hard link to `/etc/init.d/SUNwhadf`, which must be restored as the final step to installation. You must also run the following command:

```
# /etc/init.d/SUNwhadf stop
```

```
Updating local /etc/nsswitch.conf file for Solstice HA
compliance ... failed
```

The operation of updating the `/etc/nsswitch.conf` file failed on the local server. You should verify that all required Solaris packages are correctly installed. Also check that the `SUNwhadf` package is correctly installed and that `/opt/SUNwhadf/bin/haconfig` has the correct execute permissions.

```
Copying <filename> to host "<hostname>" ... failed
```

The indicated copy operation failed. You should verify that all required Solaris packages are correctly installed and that the use of the `rcp(1)` command through the private link interfaces is successful.

```
"<hostname>" is already assigned
```

The hostname you entered is already in use. Enter a different hostname. If the problem is due to an earlier entry, exit `hasetup` with a Control-D and restart the program.

```
Can't find "hostname" in the hosts map  
Do you want to add it now (yes/no) [yes]?
```

The hostname you entered has no IP address. If you answer `yes` you must enter the IP address. If you answer `no`, you must enter a different hostname.

```
Can't find "hostname" in the hosts map  
Try again or type control-D to exit ...
```

At one of the hostname prompts, you entered a hostname for a private link that does not have an IP address. Since the private links must be entered into `/etc/inet/hosts` before running `hasetup`, you are not given the opportunity to enter a network IP address. To correct the problem, exit `hasetup` by entering a Control-D, set up the private links, and restart `hasetup`.

```
Address already in use by host "<hostname>"
```

At one of the prompts for a network IP address, you entered an IP address that is already assigned to the indicated hostname. `hasetup` cannot be used to create hostname aliases. If this was the intent, you must manually edit the `/etc/inet/hosts` file. To correct the problem, exit `hasetup` by entering a Control-D, set up the IP address, and restart `hasetup`.

```
Updating /etc/inet/hosts ... failed
```

Either a local or remote update to the `/etc/inet/hosts` file failed. You should verify that all required Solaris packages are correctly installed on both servers. Also check to make sure the `SUNWhadf` package is correctly installed on both servers and that the `/opt/SUNWhadf/bin/haconfig` file has the correct permissions.

```
Updating /etc/inet/netmasks ... failed
```

An update to the `/etc/inet/netmasks` file failed. You should verify that all required Solaris packages are correctly installed on both servers. Also check to make sure the `SUNWhadf` package is correctly installed on both servers and that the `/opt/SUNWhadf/bin/haconfig` file has the correct permissions.

```
Configuring network controller interface ... failed
```

The configuration of the network controller interface failed. You should verify that all required Solaris packages are correctly installed on both servers and that `ifconfig` is working. Also check to make sure the `SUNWhadf` package is correctly installed and that the `/opt/SUNWhadf/bin/haconfig` file has the correct permissions. Rebooting the servers may solve the problem.

```
Internal error getting private hosts ...
```

You should contact your service provider if this message is displayed.

```
Hostname "<hostname>" is not configured
Try again or type control-D to exit ...
```

At one of the hostname prompts, you entered a hostname for one of the private links that is not configured. All private links must be configured before running `hasetup`. Enter a different hostname. You can also exit `hasetup` by entering a Control-D, set up the private links, and restart `hasetup`.

```
Hostname "<hostname>" is not legal
```

At one of the hostname prompts, you entered an illegal hostname for one of the private links. You may have entered a hostname which is associated with a logical network interface. Enter a different hostname. You can also exit `hasetup` by entering a Control-D, set up the private links, and restart `hasetup`.

```
Hostname "<hostname>" has already been used
```

The hostname you entered has already been used. Select a different one.

```
Interface for host "<hostname>" is not up
```

At one of the hostname prompts you entered a hostname for one of the private links that is configured but not marked as up (see `ifconfig`). The private links must be configured up before running `hasetup`. Enter a different hostname. You can also exit `hasetup` by entering a Control-D, use `ifconfig` to configure up the private links, and restart `hasetup`.

```
Hostname "<hostname>" is not on the same network as "<hostname>"
```

At one of the hostname prompts you entered a hostname for one of the remote private links that does not have a network IP address matching that or the corresponding local private link name. Enter a different hostname. You can also exit `hasetup` by entering a Control-D, correct the problem, and restart `hasetup`.

```
Checking access to "<hostname>" from "<hostname>" ... failed
```

Access to the specified hostname failed. You should verify that the `/.rhosts` files on both systems are set up correctly. Check that all required Solaris packages are correctly installed on both systems and that `/bin/rsh` and `/bin/echo` are working correctly. You can try the `rsh` command with a simple `echo(1)` command to the host that rejects access.

```
Broadcast from "<hostname>" to discover remote private hostname.."
Hostname "<hostname>" is not on a private network
```

The indicated *hostname* for the private link is not on a private link. This is known because more than one other host replied to a broadcast on the link. Enter a different host name. You can also exit `hasetup` by entering a Control-D, correct the problem, and restart `hasetup`.

```
Broadcast from "<hostname>" to discover remote private hostname.."
No response from the remote host on this network
```

The broadcast to *hostname* timed out with no response. Enter a different hostname. You can also exit `hasetup` by entering a Control-D, correct the problem, and restart `hasetup`.

```
WARNING: Ignoring unsupported controller - "<network_controller>"
```

The indicated *network_controller* is not supported by Solstice HA and has been ignored by `hasetup`. You should contact your service provider.

```
Internal error while searching for primary host
```

You should contact your service provider if this message is displayed.

```
Logical hosts MUST be on the same subnets as their masters  
Try again ...
```

The logical hostnames must be on the same network as the master. Enter the correct logical hostname for the master.

```
Internal error - bad filename
```

You should contact your service provider if this message is displayed.

```
How large will the UFS logs be (MB) [10]?  
Error: "<number>" is out of range [1 - 64]
```

UNIX file system logs must be between 1 and 64 Mbytes in size. Enter a value in this range.

```
Please use the unqualified hostname; do not include
trailing domain names!
```

A dot (.) was found in one or more of the hostnames. Remove the domain name portion of the hostname from the `/etc/inet/hosts` files.

```
You will need to use the metaset(lm) command to create
diskset(s) manually!
```

You answered no at the following prompt:

```
All multi-host disks will now be repartitioned and
populated with metadevice state database replicas.
Is this okay? no
```

The necessary operations for repartitioning the disks and creating disksets must be done manually. If you accidentally entered no, you should restart `hasetup`. Otherwise, refer to the *Solstice DiskSuite 4.0 Administration Guide* for instructions on creating disksets.


```
Checking for Solstice DiskSuite replicas on the local
host ... failed
Unable to execute "metadb" on the local host

Checking for Solstice DiskSuite replicas on the local
host ... failed
Unexpected response from "metadb" on the local host

Checking for Solstice DiskSuite replicas on the local
host ... failed
Error running "metadb" on the local host
```

Any of the above messages indicates there was a error when `hasetup` attempted to execute the `metadb(1M)` command on the local host. You should check to make sure the `SUNWmd` package is correctly installed. You can also try running the `/usr/opt/SUNWmd/sbin/metadb` command with no options to make sure it is functioning properly.

```
Checking for Solstice DiskSuite replicas on the remote
host ... failed
Unable to execute "metadb" on the remote host

Checking for Solstice DiskSuite replicas on the remote
host ... failed
Unexpected response from "metadb" on the remote host

Checking for Solstice DiskSuite replicas on the remote
host ... failed
Error running "metadb" on the remote host
```

Any of the above messages indicates there was a error when `hasetup` attempted to execute the `metadb(1M)` command on the sibling host. You should check to make sure the `SUNWmd` package is correctly installed. You can also try running the `/usr/opt/SUNWmd/sbin/metadb` command with no options to make sure it is functioning properly. Verify that the you can communicate with the sibling host using the `rsh` command.

```
Checking for Solstice DiskSuite replicas on the local
host ... failed
No replicas were found on the local host
```

No metadvice state database replicas were found on the private disks of the local server. Solstice HA requires that a minimum of three replicas be created on the local boot disk on each of the servers before running `hasetup`. Refer to Chapter 5, “Solaris 2.4 Installation,” for instructions on setting up the replicas and restart `hasetup`.

```
Checking for Solstice DiskSuite replicas on the remote
host ... failed
No replicas were found on the remote host
```

No metadvice state database replicas were found on the private disks of the sibling server. Solstice HA requires that a minimum of three replicas be created on the local boot disk on each of the servers before running `hasetup`. Refer to Chapter 5, “Solaris 2.4 Installation,” for instructions on setting up the replicas and restart `hasetup`.

```
Checking for Solstice DiskSuite replicas on the local
host ... failed
There must be at least 3 private replicas on each host
But, only <n> replicas were found on the local host
```

Fewer than three metadvice state database replicas were found on the boot disk of the local server. Solstice HA requires that a minimum of three replicas be created on the local boot disk on each of the servers before running `hasetup`. Refer to Chapter 5, “Solaris 2.4 Installation,” for instructions on setting up the replicas and restart `hasetup`.

```
Checking for Solstice DiskSuite replicas on the remote
host ... failed
There must be at least 3 private replicas on each host
But, only <n> replicas were found on the remote host
```

Fewer than three metadatabase state database replicas were found on the boot disk of the sibling server. Solstice HA requires that a minimum of three replicas be created on the local boot disks on each of the servers before running `hasetup`. Refer to Chapter 5, “Solaris 2.4 Installation,” for instructions on setting up the replicas and restart `hasetup`.

```
Verifying that disksets do not already exist ... failed
Unable to execute "metaset"

Verifying that disksets do not already exist ... failed
Error executing "metaset"
```

An error was discovered when `hasetup` attempted to execute the `metaset` command. Use the `pkgchk` command to verify that the `SUNWmd` package is correctly installed. Also try running `/usr/opt/SUNWmd/sbin/metaset` with no options.

```
Verifying that disksets do not already exist ... failed
One or more disksets already exist!
You may need to use metaset(lm) to either remove or update these
disksets!
```

`hasetup` discovered that disksets already existed. The program will not attempt to partition disks or in any way modify existing disksets. You must decide if the removal or modification of the disksets is the action you want to take and perform the appropriate steps. If you do decide not to modify the disksets, you should consider the `hasetup` phase of the installation complete. If you do decide to remove the disksets you can restart `hasetup`.

```
Checking the list of multi hosted disks ...
ERROR: unable to open "<disk>"

Checking the list of multi hosted disks ...
ERROR: unable to get geometry of "<disk>"
```

When `hasetup` checked the list of multi-host disks, it discovered failures occurred for the specified *disks*. Check the console for disk errors and correct the problem before restarting `hasetup`. You may also need to check the `/dev` and `/device` names along with the permissions. Use the `ssaadm(1M)` command with the `display` option to verify that all multi-host disks have a label.

```
Repartitioning all multi hosted disks ...
ERROR: unable to open "<disk>"

Repartitioning all multi hosted disks ...
ERROR: unable to get geometry of "<disk>"

Repartitioning all multi hosted disks ...
ERROR: unable to get geometry of "<disk>"

Repartitioning all multi hosted disks ...
ERROR: unable to read vtoc of "<disk>"

Repartitioning all multi hosted disks ...
ERROR: unable to write vtoc of "<disk>"
```

The above errors are reported when one or more failures occurs on the specified multi-host *disk*. Check the console for disk errors and correct the problem before restarting `hasetup`. You may also need to check the `/dev` and `/device` names along with the permissions. Use the `ssaadm` command with the `display` option to verify that all multi-host disks have a label.

```
Creating diskset "<diskset>" ... failed
Error executing "metaset" command
```

hassetup discovered an error when it attempted to execute the `metaset` command during diskset creation. Check for any error messages printed by the `metaset` command. Also refer to the *Solstice DiskSuite 4.0 Administration Guide* for additional information. Before attempting to restart the `hassetup` command you may need to manually remove any disksets that were created. You should also run `pkgchk` on the sibling host to verify that the `SUNWmd` package is correctly installed.

```
Populating diskset "<diskset>" ... failed
Error executing "metaset" command
```

An error occurred when attempting to execute the `metaset` command. Check for any error messages printed by the `metaset` command. Also refer to the *Solstice DiskSuite 4.0 Administration Guide* for additional information. Before attempting to restart the `hassetup` command, you will need to manually remove any existing diskset. You should also run `pkgchk` on the sibling host to verify that the `SUNWmd` package is correctly installed.

```
Releasing diskset "<diskset>" ... failed
Error executing "metaset" command
```

An error occurred when attempting to execute the `metaset` command. Check for any error messages printed by the `metaset` command. Also refer to the *Solstice DiskSuite 4.0 Administration Guide* for additional information. You should also run `pkgchk` on the sibling host to verify that the `SUNWmd` package is correctly installed. Run the `metaset` command manually to release the disksets. There is no need to restart `hassetup`.

```
Compiling the list of multi hosted disks ... failed
```

hasetup failed to compile the list of multi-host disks. You should verify that all multi-host disks are correctly cabled and that the disk names for these disks are identical on each of the servers. Check for error in the console. Verify that all required Solaris packages are correctly installed. You should also run `pkgchk` to verify that the `SUNWhadf` package is correctly installed and that `/opt/SUNWhadf/bin/haconfig` has the correct execute permissions.

Terminal Concentrator Security



This appendix describes how to install and use the security features available with your terminal concentrator.

The terminal concentrator security features allow you to choose:

- Host-based security
- Terminal concentrator security
- Host-based and terminal concentrator security

For operating software installation instructions, refer to the software installation manual in the *Terminal Concentrator Binder Set*.

For additional security information, refer to the Solaris documentation that came with your system.

B.1 Security Concepts

The terminal concentrator provides comprehensive security features to protect the concentrator and the network from unauthorized access. These features allow you to choose between host-based security, in which one host on the network functions as a security server with password protection, where passwords are stored on the concentrator. Optionally, local password protection can be used as a backup for the host-based security.

You can configure the following security checkpoints:

- CLI (Command Line Interpreter) security
Access to the concentrator by a user at a device attached to a port.
- Port server security
Access to a device attached to a port by a user at another host on the network.
- Connection security
Access to hosts or networks by a user at a concentrator.
- Virtual CLI security
Access to a virtual CLI connection from a host on the network.
- SLIP security
Access to a Serial Line Interface Protocol (SLIP) connection.
- PPP security

The concentrator provides protection using an administrative password that controls access to the root (superuser) CLI commands. This password can also protect access to a concentrator through `na`, a terminal concentrator application. The security system provides audit trails that monitor users and their activities. The concentrator also provides the source code for the Access Control Protocol (ACP) security system, and the flexibility to integrate concentrator security with existing security for a network-wide system.

B.2 Using Security

The terminal concentrator provides a security system that allows you to implement as many security measures as the network requires. You can set up the security subsystem to use host-based security, local password protection, or a combination of the two. In addition to these security mechanisms, the concentrator provides an administrative password that is used to validate access through the administrative tools.

Note – If unauthorized users can access your concentrator, we strongly suggest that you enable the security features after loading the host code and booting the unit.

Host-based security uses the Access Control Protocol (ACP) of the concentrator and requires that a host function as a security server. You can modify the host-based security policy supplied with the concentrator to implement a security policy that fits the needs of your environment. Host-based security allows you to define:

- User validation on access to and from the concentrator
- Connection protection to other hosts and networks
- User activity logs
- Encryption of security messages between the concentrator and the server

Local password protection can be defined for access to and from individual ports and for virtual CLI access. Local password protection does not provide logging of security events to the security server. If event logging is enabled, user activities can be logged to `syslog` with local password protection. See the *Terminal Concentrator General Reference Guide* for more information.

Configuring the concentrator for security involves:

- Enabling security for host-based security, local password protection, or both
- Defining the server if you are using host-based security
- Defining a local virtual CLI password for local password protection
- Defining the administrative password

B.2.1 Enabling Security

To use any security feature, you must enable security for the terminal concentrator by setting the `enable_security` parameter to Y. This parameter is mandatory if you intend to use any concentrator security mechanisms for access to administrative tools except the administrative password.

B.2.2 Setting Up Host-based Security

Host-based security requires at least one host on the network acting as a security server. The security server maintains files that contain user validation and connection permission information. Also, the security server can encrypt security messages.

B.2.2.1 Specifying the Security Hosts

The `pref_secure1_host` and `pref_secure2_host` parameters specify the preferred security hosts. The concentrator first queries the `pref_secure1_host` for user validation or permission connection requests. If a response is not received within a limited amount of time, the concentrator repeats the query several times. If the concentrator still does not receive a response, it queries the host defined with the `pref_secure2_host` parameter. If a response is not received from the second security host within a limited time, the concentrator broadcasts. If the broadcast fails, the concentrator denies the user's request.

The `network_turnaround` parameter specifies the amount of time in seconds in which the concentrator expects a response from the security servers. To reduce the possibility of a retry, the network turnaround time should be sufficient to allow for a network transmission to the security server and transmission back to the concentrator.

B.2.2.2 Disable Broadcasting for Security Servers

The terminal concentrator automatically broadcasts for a security server when the host defined using either the `pref_secure1_host` or `pref_secure2_host` parameter does not respond. You can disable broadcasting for a security server by setting the `security_broadcast` parameter to N.

B.2.2.3 Restricting telnet Access to Certain Ports

The concentrator operational code passes the TCP port number to the ACP when doing a `telnet(1)` on a port with `connect_security` enabled. This feature allows restrictions on connections to certain TCP ports. The port number is available in the `annex_to_net()` parameter in the `acp_policy.c` file.

B.2.2.4 ACP Virtual CLI Protection

You can set up security for virtual CLI connections where users must provide a valid user name and password before they are granted access to a virtual CLI. You can define security on the virtual CLI connections of the terminal concentrator using the `vcli_security` parameter. To set up virtual CLI protection with host-based security:

- 1. Enable security by setting the `enable_security` parameter to `Y`.**
- 2. Enable virtual CLI security by setting the `vcli_security` parameter to `Y`, then create a password file on the security server (see the *Annex Network Administrator's Guide*).**

B.2.2.5 Encrypting Security Messages

With host-based security, passwords are continuously sent across the network. To prevent unwanted access to these passwords, you can request encryption of messages between the terminal concentrator and the security server.

The `acp_key` parameter specifies the encryption key the concentrator uses to exchange messages with the security server. The security server maintains the encryption key for the concentrator in the `acp_keys` file (see the *Annex Network Administrator's Guide*).

The encryption key also validates the security host. In this case the host must know the ACP key of the concentrator for the concentrator to consider the host valid. Without the appropriate key, the concentrator denies the user's request.

Note – The `show annex` command does not display the value of the `acp_key` parameter. Instead, it displays `<set>` or `<unset>`.

B.2.3 Local Virtual CLI Password

Local password security allows you to assign a password that a user must enter before accessing the terminal concentrator. Since this password is stored locally on the concentrator, it does not require a remote security server. Local password security can be used as a back-up security mechanism in case the host-based security servers are unavailable.

Local password security can be implemented for the concentrator in one of two ways: upon virtual CLI connection and upon access through administrative utilities. The `vcli_password` parameter allows you to define a local password for virtual CLI connections. With this security mechanism, the user enters only the password stored on the concentrator, not a user name and password. To configure the concentrator for a local virtual CLI password:

1. **Enable security by setting the `enable_security` parameter to Y.**
2. **Disable virtual CLI security by setting the `vcli_security` parameter to N.**
3. **Define a password using the `vcli_password` parameter.**

You can also use the `vcli_password` as a backup to host-based security. When this local virtual CLI password is used as a backup, the concentrator first accesses the security server to validate a CLI connection request. If no response is received from a security server, the concentrator requests the local virtual CLI password. The user can enter either the virtual CLI password or the concentrator administrative password. To set up the local virtual CLI password for backup security:

1. **Enable security by setting the `enable_security` parameter to Y.**
2. **Enable virtual CLI security by setting the `vcli_security` parameter to Y.**
3. **Define a password with the `vcli_password` parameter.**
4. **Create a password file on the security server (see the *Annex Network Administrator's Guide* for instructions).**

Note – The `show annex` command does not display the value of the `vcli_password` parameter. Instead, it displays `<set>` or `<unset>`.

B.2.4 Concentrator Administrative Password

The administrative password of the terminal concentrator always validates access to superuser CLI commands; it can also validate `na` program access to the concentrator if the `enable_security` parameter is set to Y.

The password parameter modifies the administrative password of the concentrator. The default password is the Internet address of the concentrator.



Warning – The administrative password never displays. If you forget the modified password, you can reset it only by erasing all of the nonvolatile memory in the concentrator by using the ROM Monitor.

The administrative password is required to access the superuser CLI commands. When you enter the `su` command at the CLI prompt, you are asked for a password. If you have not modified the administrative password, enter the Internet address of the concentrator.

The administrative password validates access to the concentrator through `na` only when security is enabled and the password is defined. When issuing the annex command or specifying the concentrator using `na`, you must enter this password. The password also encrypts messages between the concentrator and `na`. The administrative password can be used as the virtual CLI password and to override the password assigned using the CLI lock command.

B.3 Setting Port Security

The concentrator provides a security system allowing you to configure security on a per-port basis. You can use host-based security, local password protection, or a combination of the two.

Host-based security provides access validation for the following access requests:

- From the device, a terminal or modem, attached to the port
- From a user on the network connecting to the port
- From a device attached to the port connecting to a host or a network

Local password protection can be used as a stand-alone security mechanism or as a backup to host-based security. It validates port access from either the device or the network. Local password protection supports CLI, slave, and adaptive ports.

Host-based security also allows you to mask out CLI commands. When a command is masked, it is not displayed with the `help` command; if entered, CLI displays an error message. You can customize the security policy to meet your individual security requirements. The local password protection policy cannot be altered (see the *Annex Network Administrator's Guide* for more information).

Note – For any security mechanism, you must enable security for the concentrator by setting the `enable_security` parameter to `Y`.

B.4 Port Server and CLI Security

The port server security provides the option of configuring host-based security, local password protection, or both. The host-based security for the port server normally requires a user name and password. Local password protection on the port server requires only a password.

As with security on CLI connections, the local password protection can be used as the sole security mechanism or as a backup to host-based security for those occasions when the security servers are not available.

With port server security, the port server invokes the security mechanism when the user requests access to a specific port or rotary at the port server prompt. User validation occurs before the user is connected to the port to ensure that the user is authorized to connect to the selected port. If the user is not authorized, the port server notifies the user and prompts for another port. The following example is an example of the supplied user validation for port server security.

```
annex: telnet annex01
Trying...
Connected to annex01.
Escape character is "^]".
Rotaries Defined:
  modem_1200  1-3,5
  modem_2400  4,7
  cli         -
Enter concentrator port name or number: modem_1200
Annex username: kate
Annex password:
Permission granted
Attached to port 1.
```

For camp-on, user validation occurs when a port becomes available. If the name or password is incorrect, the user is returned to the port identification prompt.

To use host-based security with the port server:

1. **Set the concentrator parameter** `enable_security` **to Y.**
2. **Set the port parameter** `port_server_security` **to Y.**
3. **Create a password file** (`acp_passwd`) **on the security server(s)** (see the *Annex Network Administrator's Guide* for more details).
4. **For local password protection with the port server, define a password for the** `port_password` **parameter.**

Note – The port parameter `port_password` is applicable to both CLI and port server connections.

B.5 Security on Virtual CLI Connections

The concentrator establishes security on virtual CLI (VCLI) connections using host-based security, local password protection, or both. Host-based security validates the user name and user password.

The virtual CLI security mechanism is similar to the port server security mechanism in that user validation is invoked after the user has requested access to the VCLI at the port server prompt. This ensures that the user is authorized to access the VCLI.

To set up host-based security on virtual CLI connections:

1. **Set the** `enable_security` **concentrator parameter to Y.**
2. **Set the** `vcli_security` **concentrator parameter to Y.**
3. **Create a password file** (`acp_passwd`) **on the security server(s)** (see the *Annex Network Administrator's Guide* for more information).

Local password protection requires only a password for validation. To set up this protection for the concentrator:

1. **Set the concentrator parameter** `enable_security` **to Y.**
2. **Set the concentrator parameter** `vcli_security` **to N.**
3. **Define a password for all virtual CLI connections for the concentrator using the** `vcli_password` **parameter.**

Local password protection can function as a backup to host-based security:

1. **Set up host-based security for the virtual CLI connection.**
2. **Define a virtual CLI connection password using the concentrator parameter `vcli_password`. Virtual CLI connections must adhere to any connect security defined for the concentrator.**

B.6 PPP Security

The optional PPP phase occurs after the LCP negotiations are complete. Each half of the connection can require security. The terminal concentrator first tries to use ACP and then falls back to local security; if neither are available, the concentrator closes the link.

To start authentication, the peer sends a PPP Authenticate message to the concentrator. The concentrator sends an ACP `ppp_security_protocol()` call to verify the peer's user name/password. The concentrator allows three attempts before closing the link. If ACP vouches for the user, the concentrator sends a PPP Authenticate-Ack message and continues NCP negotiations. The concentrator uses local security when ACP is unavailable and the `port_password` parameter is set; local security ignores the user name and checks the password against `port_password`.

B.7 Host-based Security

The Access Control Protocol (ACP) provides the following host-based security options:

- User validation for connecting to a CLI
- User validation for access to the port server
- Validation for host connections
- Logging security events
- User validation for access to a PPP port
- User validation for access to dial-up address connections
- Encryption key protection for concentrator access to the security server, and optionally, for access to a concentrator through `na`

ACP requires that at least one host is a security server and that security is enabled on the concentrator. All files that maintain access validation reside on the security server. Using host-based security, the concentrator logs messages to the security server when security events occur and when the concentrator is booted. You can modify this supplied policy to:

- Disable user name and password validation
- Change the name of the password file
- Disable CLI commands

Using the concentrator distribution's library calls, and any library calls available on your system, you can modify the code to provide your own prompts and encoded validation.

B.7.1 Defining a Security Server

The ACP security server software is provided as part of the expedited remote procedure call daemon (`erpcd`) supplied with the file server software. Included with the software is the services file that has two entries: one for the block file server (`bfs`) and one for Access Control Protocol (ACP).

The ACP security server software has the ACP service commented out, as follows:

```
#erpc remote programs
#
# prog no.   verlo   verhi   name
#
# 1          0       99     bfs
# 3          0       99     acp
```

To define a security server, you must install the file server software on a host and delete the `#` symbol in front of the ACP entry. For example:

```
#erpc remote programs
#
# prog no.   verlo   verhi   name
#
# 1          0       99     bfs
  3          0       99     acp
```

You can use the same host to be both file server and security server, or move the servers to two separate hosts. You can put the security server on more than one host. However, you must define one host as a first preferred security server. Optionally, you can define a second host as a second preferred security server. Additional security servers respond to broadcast requests if the preferred servers are unavailable and `security_broadcast` is set to `Y`.

Note – The contents of the files `acp_passwd`, `acp_restrict`, and `acp_keys` should match on all servers.

B.7.2 Creating the `acp_dialup` File

The `acp_dialup` file resides in the install directory. Any ACP dial-up address request that comes from the concentrator has a `username` and an Internet address which is used as keys in this file. Once the keys are matched, the corresponding dial-up addresses are returned to the caller on the concentrator. If no match is found, the concentrator uses the port's `remote_address` and `local_address`.

You can specify the concentrator by name, Internet address, or wild card (*); the wild card means that any incoming address request with that user name will match. Also, you can specify a range of ports. The file format allows one entry per line; the concentrator ignores any data following the comment character (#); a new line character terminates an entry.

The `acp_dialup` file displays the following fields: User, Annex, Remote Address, and Local Address (optional). If a local address is not specified, the local address sent back to the concentrator is that address of the concentrator. Table B-1 shows an example.

Table B-1 Entries in the `acp_dialup` File

User Name	Concentrator/Port_set	Remote Address	Local Address (Optional)
smith	4,5,6-9@100.30.200.39	100.30.200.45	100.30.200.46
green	*	100.30.200.48	
cody	3,7,20-25@jupiter	100.30.200.47	
harris	mars	100.30.200.55	

In the preceding example:

- User **smith** can make a dial-up address request from concentrator 100.30.200.39 on ports 4, 5, and 6 through 9. The remote address is 100.30.200.45; the local address is 100.30.200.46.
- User **green** can make a dial-up address request from any concentrator on any port. The remote address is 100.30.200.48; the local address is the address of green's concentrator.
- User **cody** can make a dial-up address request from concentrator jupiter on ports 3, 7, and 20-25. The remote address is 100.30.200.47; the local address is the address of cody's concentrator.
- User **harris** can make a dial-up address request from concentrator mars on any port. The remote address is 100.30.200.55; the local address 100.30.200.40.

B.7.3 Creating User Password Files

CLI ports are configured for user validation with the `cli_security` port parameter. Virtual CLI connection user validation is enabled with the `vcli_security` concentrator parameter. User validation on accessing ports through the port server is provided with the `port_server_security` port parameter.

Note – User validation is available only if host-based security is enabled at the concentrator using the `enable_security` concentrator parameter.

For user validation with the supplied security policy, the concentrator sends a request to the security server to compare a user name and password against entries in a password file. If a match is found, the user is granted access to the CLI; otherwise, the user is denied access to the CLI. For example:

```
Annex Command Line Interpreter * Copyright 1990 Xylogics, Inc.  
Checking authorization, Please wait...  
Annex username: kate  
Annex password:  
Permission granted  
annex:
```

The supplied policy expects the password file to be the `acp_passwd` file. This file uses the same format as the `/etc/passwd` file. The easiest way to create this file is to copy the `/etc/passwd` file to `acp_passwd`. One advantage for creating the `acp_passwd` file this way is that you can merge `/etc/passwd` files from different hosts into one file on the security server. This allows you to create a network-wide password file.

Under Solaris 2.x the passwords reside in the `/etc/shadow` file.

B.7.4 Setting Up Connection Security

ACP provides connection security, which authorizes connection requests to hosts using the CLI connection commands `telnet` and `rlogin(1)`. The connection security mechanism uses a host-resident file that lists the hosts to which connections are restricted.

Connection security is enabled for individual ports by setting the concentrator parameter `enable_security` and the port parameter `connect_security`. Setting the concentrator parameter `vcli_security` enables connection security for all virtual CLI connections. When a user issues a connection command, the concentrator checks a restrict file for permission to connect to that host. The supplied policy expects the restrict file to be `acp_restrict`, which is an ASCII file that you create with any text editor. Table B-2 describes the arguments in each entry. The entry format is:

```
concentrator: restricted host, restricted host, ...
concentrator~ unrestricted host, unrestricted host, ...
```

Table B-2 Arguments in the `acp_restrict` File Entries

Argument	Description
<code>annex</code>	The name or Internet address of the concentrator for which connection security is enabled
<code>:</code> (colon)	Indicates the hosts listed in the same entry are restricted.

Table B-2 Arguments in the `acp_restrict` File Entries

Argument	Description
~ (tilde)	Indicates the hosts are unrestricted.
restricted host	The name or Internet address of a restricted host (including concentrators). The list of restricted hosts is separated by commas.
unrestricted host	The name or Internet address of an unrestricted host (including concentrators). The list of unrestricted hosts is separated by commas.

You can use an asterisk (*) as a wild card in place of a host name or the host part of an Internet address. Following is an example of two restricted-host entries:

```
annex01: hosta, hostb, hostf, 132.245.6.23
annex02: hostc, 312.245.6.15, hostf, 132.245.6.23, hosth, annex01
```

The first entry prevents users in which connection security is enforced on *annex01* from accessing *hosta*, *hostb*, *hostf*, and the Internet address 132.245.6.23. The second entry prevents users on *annex02* from accessing five hosts and one concentrator. Connection security is enforced on all virtual CLI connections at *annex01* and *annex02*.

Hosts that are not listed in the file are considered unrestricted. Since ACP searches the `acp_restrict` file sequentially, the order of placement in the file is important. The search stops when it finds a host that matches. You can use unrestricted host entries to prevent users on one network from accessing hosts on any other network.

In the following example, the policy finds the unrestricted definition for concentrators and hosts on network 192.17.5 and grants access. It finds the restricted definition for hosts on any other network and does not grant access.

The next example illustrates wild cards. A *publichost* is defined as accessible from all concentrators and a *securehost* is inaccessible from any concentrator.

```
192.17.5.*~ 192.17.5.*  
192.17.5.*: *
```

*~ publichost

*: securehost

If permission is granted to a connection security request, the user follows the normal login procedure. If the request is denied, the message

```
Permission denied
```

is displayed and the session (job) is aborted.

B.7.5 Creating the ACP Encryption Keys File

The host-based security can encrypt messages between a concentrator and the security server. Encryption of messages is based on the ACP key, which is set using the concentrator parameter `acp_key`. This option is available when host-based security is enabled for the concentrator by setting the concentrator parameter `enable_security` to Y.

When ACP encryption is enabled, the supplied policy expects a file called `acp_keys`, which maintains a list of concentrators and their respective encryption key. When the security server receives an encrypted message from a concentrator, the server tries to match that key against the key assigned to the concentrator in the file. If no match exists, the concentrator and the server cannot communicate.

Each entry in the encryption keys file contains a list of concentrator names or Internet addresses separated by commas and an encryption key for those concentrators. The concentrator or the list of concentrators and the key are separated by a colon. The order of placement in the file is important, as the file is sequentially read. Syntax rules for the `acp_keys` file are:

Any part of an Internet address in the list can be specified with an asterisk (*).

- A backslash (\) is used to escape a new line.
- Any ASCII character except spaces and tabs are valid encryption keys.
- Each key can contain a maximum of fifteen characters.

Typically, concentrators that do not expect encryption do not need to be included in the file. However, if a concentrator falls into a group of multiple concentrators that were specified with an * as part of the Internet address and it is not expecting encryption, that concentrator must be included in the file with the key left blank.

In the following example, the first three entries specify `insomniac-1` as the key for the concentrator whose Internet address is `132.245.6.15`, no encryption for the concentrator whose Internet address in `132.245.6.75`, and `Piano` as the key for all other concentrator on the same network. The last entry specifies `g112ch` as the key for `annex01`, `annex02`, and `annex03`. Each `acp_key` parameter for the concentrators listed in the example must be identical to the key included in the `acp_keys` file.

```
132.245.6.15:insomniac-1
132.245.6.75:
132.245.6.*:Piano
annex01,annex02,annex03,g112ch
```

Changing the value of the `acp_key` parameter on any concentrator requires the same change to the `acp_keys` file on the security server. The recommended order for changing the ACP encryption key on a concentrator is:

- 1. Edit the `acp_keys` file on all security server hosts.**
- 2. Use `na` to change the value of the `acp_key` parameter for all affected concentrators.**
- 3. Update the cache by sending the `erpcd` on all security server hosts a HUP signal with `kill`.**
- 4. Reset the security subsystem for all affected concentrators using the `na` command `reset concentrator security`.**

B.7.6 Logging Security Events

Host-based security provides the ability to generate audit trails of user activity. Each time the security server grants or denies a request for user access, the security server logs it. Each event is logged as a message to a file defined as `acp_logfile` (see the *Annex Network Administrator's Guide*). You can change `acp_logfile` in the file `/annex_root/src/erpcd/acp_policy.h` to any suitable pathname including `/dev/null`, `/etc/console`, etc.

Each logged message contains the Internet address of the concentrator, a sequence number, the port number, the date and time, the event, and other information that is protocol dependent. The fields in the file are separated by colons and can be used by UNIX utilities that sort, merge, select, or filter streams.

B.7.7 Modifying the Supplied Security Application

You can modify the supplied security policy to create a security scheme that meets the needs of your network. Some simple modifications involve changing system definitions in the file `/annex_root/src/erpcd/acp_policy.h`. More elaborate security policies may require modifying or replacing functions in the file `/annex_root/src/erpcd/acp_policy.c`.

Note – Do not change the function declarations or the description of the interface; these are fixed by the calls made into this library. Before making even the smallest change, save the base version of the file requiring modification.

If you modify the default policy, you must recompile `erpcd`, `kill(1M)` the current version, and start the new version. Instructions for doing this are provided later in this section.

B.7.7.1 Disabling User Name and Password Validation

When security is enabled, users are requested to provide a user name and password. You can disable this policy by modifying the `/annex_root/src/erpcd/acp_policy.h` file to change the line that defines user validation from:

```
#define USER_VALIDATION 1
```


to:

```
#define USER_VALIDATION 0
```

Messages are logged to the security server host when users access the CLI, but the message does not include a user name.

B.7.7.2 *Changing the Expected Names of Files Used by ACP*

The supplied policy uses names for various files, such as `acp_passwd`, `acp_keys`, `acp_restrict`, and `acp_logfile`. You can change the names of any of these files in the `/annex_root/src/erpcd/acp_policy.h` file. If you decide to use either an existing system or a network-wide password file instead of `acp_passwd` file, change the following lines in the file `/annex_root/src/erpcd/acp_policy.h`:

```
#define ACP_PASSWD (str) \  
    sprintf(str, "%s/acp_passwd",INSTALL_DIR)  
#define ACP_PTMP (str) \  
    sprintf(str, "%s/acp_ptmp",INSTALL_DIR)
```

To change only the filename:

```
#define ACP_PASSWD (str) \  
    sprintf(str, "%s/new_filename",INSTALL_DIR)  
#define ACP_PTMP (str) \  
    sprintf(str, "%s/new_filename",INSTALL_DIR)
```

To change the full pathname:

```
#define ACP_PASSWD (str) \  
    sprintf(str, "new_path/new_filename")  
#define ACP_PTMP (str) \  
    sprintf(str, "new_path/new_tmpfile")
```

The *new_filename* is the name of the new password file, and the *new_tempfile* is a temporary file used by the `ch_passwd` command. Because you do not need the temporary file if you are using an existing system file, comment out the line for the temporary file.

Note – `INSTALL_DIR` is defined in the file `/annex_root/src/make.config` with the leading quote supplied by the makefile. Because the trailing quote is required by the two strings, double quote the names for the new password and temporary files.

B.7.7.3 *Disabling CLI Commands*

When the security subsystem is enabled, you can mask user access to specific CLI commands by changing the following line, which defines user commands as not masked, in the `/annex_root/src/erpcd/acp_policy.h` file:

```
#define CLI_MASK(unsigned log 0)
```

To mask the `hosts`, `netstat`, and `stats` commands, modify that line to read:

```
#define CLI_MASK (MASK_HOSTS | MASK_NETSTAT | MASK_STATS)
```

If the user enters the masked command, CLI displays an error message. The superuser CLI commands cannot be masked individually. They can all be disabled by masking the `su(1M)` command.

B.7.7.4 *Modifying the Code*

You can create a more elaborate security policy application by modifying the code in the files `/annex_root/src/erpcd/acp_policy.c` and `/annex_root/src/erpcd/acp_policy.h`. The program that executes ACP forks itself each time a security request is received from a concentrator. A call is made to an ACP remote procedure, which makes calls to functions in the ACP library to prompt for user names and passwords. When ACP gathers the information required to perform the authorization algorithm, it again calls functions in the library to grant or deny the request. The program then exits.

The distribution policy file `/annex_root/src/erpcd/acp_policy.c` is documented in the form of C programming language comments. The `/annex_root/src/erpcd/policy.doc` file provides a complete description of the available library functions.

B.7.7.5 *Recompiling erpcd*

You must recompile `erpcd` if you modify the supplied policy and the `ch_passwd` utility if you changed the name of the ACP password file from `acp_passwd`. The source files are in `/annex_root/src/erpcd`, where `annex_root` is the directory to which the source code of the concentrator was copied. To recompile:

1. **cd to** `/annex_root/src/erpcd`.

2. **To recompile only erpcd, enter the command:**

```
# make -f ../make.config -f Makefile erpcd
```

3. **To recompile both erpcd and ch_passwd, enter the command:**

```
# make -f ../make.config -f Makefile all
```

4. **To install, enter the command:**

```
# make -f ../make.config -f Makefile install
```

This command saves the old version of `erpcd` as `OLDerpcd` in the installation directory.

5. **Kill the current erpcd and start the new one.**

B.8 *Securing the Network*

The concentrator security system provides a full set of features that allow you to establish any security policy required by the network. You can choose between host-based and local password protection on the concentrator, or you can use local password protection as a backup for host-based security.

B.8.1 The Concentrator Administrative Password

The concentrator administrative password protects the administrative tools; the default administrative password is the Internet address of the concentrator. When the `show annex` command displays the password as `<unset>`, use the default administrative password for:

- Access to superuser CLI commands
- Access to ports locked with the CLI lock command
- Access to a virtual CLI connection through local password protection

Modifying the assigned administrative password enables:

- Password protection on access to concentrators through `na`
- Encryption of messages between a concentrator and `na`

Note – If you lose your administrative password, your only option is to erase the non-volatile memory of the concentrator by using the ROM Monitor erase command, and re-enter all parameters.

B.8.2 Protecting Ports from Unauthorized Access

When terminals are connected to a network, they provide users with the potential for unauthorized access to hosts and resources. In addition to the available security schemes, the concentrator provides timers that can reset a port. The `cli_inactivity` port parameter sets the CLI inactivity timer: when the last session is completed, the concentrator resets the port after the amount of time specified in this parameter has elapsed.

Users can protect their login sessions using the CLI lock command if they do not want to log out when leaving the terminal unattended.

B.8.3 Protecting the Superuser CLI

A concentrator administrative password is required for access to the superuser CLI using the `su` command. The default password is the Internet address of the concentrator. There are two ways to change the password:

- Using the superuser CLI `passwd` command
- Changing the password parameter using `na`

Using either method, the new password takes effect immediately for access to the superuser CLI. Reset the password to the Internet address of the concentrator by:

- Using `na` to set the password parameter to the null string. For example:

```
command: set annex password ""
```

- Using the superuser CLI `passwd` command and pressing Return in response to the prompt for a new password.
- Erasing all parameters using the ROM monitor `erase` command.

B.8.4 Preventing Unauthorized Access to na

When using `na`, users can access concentrator parameters and obtain useful information, or reconfigure and reboot concentrators. Protecting `na` involves UNIX superuser protection and the concentrator administrative password.

Upon installation, the `na` utility is owned by root and executable by all. Only a superuser can execute the `set`, `reset`, `broadcast`, `dumpboot`, `boot`, `read`, and `copy` commands.

Index

A

- additional public networks, 3-8
- administration workstation
 - host name, 3-7
 - IP address, 3-7
- asymmetric configuration, 1-3, 3-17
- audience for this manual, xv
- available disk sizes, 3-9

B

- backing up data
 - Solstice Backup, 3-23
- boot disk
 - replica space, 5-7

C

- calculating size of root slice, 5-2
 - inode usage, 5-2
- changes made to `nsswitch.conf`, 7-3
- concatenation, 1-4
- configuration location
 - maximum separation, 3-16
- configuration planning
 - backing up multi-host data, 3-23
 - data migration, 3-9
 - data on local disks, 3-17

- data services, 3-9
 - disk space growth, 3-11
 - file system size and disk layout, 3-19
 - hot spares, 3-11
 - metadevice naming and creation,
 - 3-21
 - migrating existing data, 3-23
 - mirroring, 3-10
 - network configuration, 3-3
 - number of SPARCstorage Arrays,
 - 3-13
 - overview, 3-1
 - physical and logical host names, 3-8
 - questions to consider, 3-2
 - setting up a power source, 3-15
 - size and number of disk drives, 3-12
 - size of disksets, 3-19
 - size of servers, 3-13
 - total storage requirement, 3-12
 - trans devices, 3-10
 - type of configuration, 3-17
- configuration restrictions, 3-26
 - configuration rules, 3-23
 - creating an ORACLE database, 10-6
 - creating local replicas, 6-30
 - creating the configuration, 8-1
 - creation of multi-host metadevices, 9-12
 - using `md.tab` file, 9-13

using `metatool`, 9-12

D

data protection, 1-2
data service, 1-4
data service class setup, 10-1
 HA-NFS, 10-10
 HA-ORACLE, 10-2
default master, 1-4
definition of terms, 1-3 to 1-7
disk layout
 sharing UFS logs, 3-19
disk space growth
 considerations, 3-11
 planning, 3-11
diskset, 1-4
 example layout, 2-5
disksets
 asymmetric configuration, 3-21
 determining the size, 3-20
 example allocation, 3-21
 number of drives needed, 3-20
DiskSuite, 1-4
dump media size, 3-9

E

editing `/etc/syslog.conf` file, 7-8
error messages
 hacheck, A-1
 hacheck, A-1
 hasetup, A-9

F

fault detection, 1-4
file system size, 3-9
 minimum, 3-19

G

getting help, 1-3
glossary, 1-3

H

hacheck error messages, A-1
HA-NFS, 1-4
HA-ORACLE, 1-4
hardware configuration requirements, 2-4
hardware installation
 manual set, 4-1
 SPARCcenter 2000, 4-2
 SPARCserver 1000, 4-3
 SPARCstation 4, 4-2
 SPARCstorage Arrays, 4-3
 terminal concentrator, 4-4
hasetup
 overview, 8-1
 preparing to run, 8-2
 running, 8-3
 worksheet, 8-2
hasetup error messages, A-9
heartbeat, 1-4
help, 1-3
highly available data service, 1-4
hot spares
 assignment to disksets, 3-11
 considerations, 3-11
 setting up, 9-14
hot standby, 1-4
how to run `hasetup`, 8-3

I

identical configurations, 2-2
inodes required
 for disk slices, 5-2
 for metadevices, 5-3
 for SPARCstorage Arrays, 5-2
inodes required by Solaris, 5-2
installation procedure overview, 2-6

L

local disk
 requirements, 2-4
local disks, 1-5

- mirroring, 3-17
- replicas, 3-18
- root slice size, 3-18
- UFS logs, 3-18
- location
 - separating servers, 3-16
- logical host, 1-5
 - naming, 3-4
 - suffix, 3-5
- logical host name, 3-8
- logical hostname, 1-5
- logical network interface, 1-5
- loopback file system
 - exclusion of, 5-34

M

- MANPATH variable, 6-30 to 6-31
- manual organization, xvi
- master, 1-5
- md.conf file, 3-21
 - changing, 3-22
 - changing values, 5-3
 - values to change, 3-22
 - when to edit, 6-30
- md.tab file
 - building top down, 9-9
 - example, 9-10
 - using to configure, 9-9
- membership monitor, 1-5
- metadevice, 1-5
- metadevice considerations, 9-2
 - availability, 9-3
 - diversity of the configuration, 9-2
 - managing metadevice name space, 9-6
 - performance issues, 9-5
 - state database replica issues, 9-6
 - takeover issues, 9-6
- metadevice creation, 9-1
 - HA-NFS, 9-14
 - HA-ORACLE, 9-14
- metadevice name space example, 9-7
- metadevice state database, 1-5

- metadevice state database replica, 1-6
- metadevices
 - creating on multi-host disks, 9-12
- migrating data
 - ORACLE databases, 3-23
 - using `ufsdump`, 3-23
- mirroring, 1-6
 - three-way, 3-10
- mirroring requirements, 2-4
- modifying existing clients, 2-1
- multi-host disk, 1-6
- multi-host disks
 - determining number needed, 3-13
 - I/O bandwidth, 3-12
 - placement in SPARCstorage Arrays, 3-13
 - replica space, 3-12
 - takeover speed, 3-12
 - UFS log space, 3-12
 - using high capacity, 3-12
 - using low capacity, 3-12
- multi-host file system
 - mount points, 3-18
 - planning, 3-18
 - statmon file system, 3-18
- multi-host file system hierarchy, 3-18

O

- ORACLE installation
 - creating a database, 10-5
 - on local disks, 5-28
 - preparing servers, 10-2
 - setting up HA-ORACLE, 10-5

P

- package installation
 - checking the packages, 6-30
 - correct sequence, 6-1
 - procedure, 6-2
 - Solstice DiskSuite, 6-5
 - Solstice HA, 6-18
 - Solstice HA patches, 6-28
 - Solstice HA-ORACLE, 6-29

- SunVTS, 6-3
- partitioning multi-host disks, 9-14
- path requirements, 6-30 to 6-31
- PATH variable, 6-30 to 6-31
- physical limits of SPARCcluster
 - components, 3-14
- planning and installation overview, 2-1
- planning overview, 2-1
- power sequence testing
 - SPARCcluster 1000, 11-5
 - SPARCcluster 2000, 11-6
- power sources
 - independent, 3-16
 - partial power failure, 3-15
 - power distribution to components, 3-16
 - use of UPS, 3-15
- primary host name, 3-8
- private host names, 3-8
- private network, 3-6
 - IP numbers, 3-7
- private network requirements, 2-3
- private networks
 - naming, 3-7
 - setting up, 7-7
- public network requirements, 2-3
- public networks
 - bandwidth requirements, 3-7

R

- related documentation, xvii
- replica space, 5-7
- replicas
 - creating local, 6-30
- required root file system size, 5-5
- reserving space for UFS logs on local disks, 8-7
- restrictions, 3-26
- root slice
 - calculating the size of, 5-2
- rules to improve reliability, 3-23
 - Fibre Channel SBus cards, 3-25

- identical hardware, 3-23
- identical network interfaces, 3-24
- local disks, 3-24
- multi-host disk connections, 3-23
- private networks, 3-25
- replica distribution, 3-24

S

- servers
 - determining which size, 3-14
- setting required paths, 6-30
- setting up `.rhost` files, 7-8
- setting up HA-ORACLE, 10-6
- setting up hot spares, 9-14
- setting up name services and networks tasks, 7-1
- setting up private networks, 7-7
- sibling host
 - definition, 1-6
- simplifying administration
 - disk space usage, 3-12
- sizing file systems, 5-7
- Solaris installation
 - instructions, 5-7
 - of the administration workstation, 5-7
 - on Solstice HA servers, 5-10
 - patches, 5-8, 5-34
 - rules, 5-1
- Solstice DiskSuite, 1-6
 - introduction, 1-1
- Solstice HA, 1-6
 - functional diagram, 2-2
 - introduction, 1-1
 - online serviceability, 1-2
 - parts of a configuration, 2-3
- Solstice HA data protection, 1-2
- Solstice HA data services, 1-1
- Solstice HA features, 1-1
- Solstice High Availability, 1-6
- SPARCcluster 1000
 - power sequence testing, 11-5

SPARCcluster 2000
 power sequence testing, 11-6
SPARCstorage Arrays supported, 1-3
stripe, 1-6
submirror, 1-6
SunVTS
 running diagnostic tools, 6-4
supported hardware, 1-2
switchover, 1-6
symbolic link requirements, 5-4
symmetric configuration, 1-6
symmetric configurations, 3-17

T

takeover, 1-6
terminal concentrator
 administrative password, B-6
 enabling security, 5-9, B-3
 host name, 3-7
 host-based security, B-10
 installation, 5-8
 IP address, 3-7
 port server and CLI security, B-8
 PPP security, B-10
 securing the network, B-21
 security, B-1
 security concepts, B-1
 security on virtual CLI
 connections, B-9
 setting port security, B-7
 setting up host-based security, B-3
 using security, B-2
trans device, 1-6
 log and master I/O, 3-10
trans devices
 log sizes, 3-10
typographic conventions, xix

U

UFS, 1-7
UFS logging, 1-7
 requirement, 2-5

UFS logging device, 1-7
UFS logging on local disks, 9-13
UFS master device, 1-7
updates to /etc files, 7-5
updates to name service entries, 7-3
updating SPARCstorage Array
 firmware, 6-30

V

verification and validation
 fault detection, 11-4
 hacheck, 11-2
 HA-ORACLE, 11-4
 haswitch, 11-2
 manual tests, 11-3
 physical tests, 11-3
 SPARCstorage Array Firmware, 11-4
 tasks, 11-1
viewing AnswerBook, 6-31

Copyright 1995 Sun Microsystems, Inc., 2550 Garcia Avenue, Mountain View, Californie 94043-1100 USA.

Tous droits réservés. Ce produit ou document est protégé par un copyright et distribué avec des licences qui en restreignent l'utilisation, la copie et la décompilation. Aucune partie de ce produit ou de sa documentation associée ne peuvent être reproduits sous aucune forme, par quelque moyen que ce soit sans l'autorisation préalable et écrite de Sun et de ses bailleurs de licence, s'il en a.

Des parties de ce produit pourront être dérivées du système UNIX®, licencié par UNIX Systems Laboratories Inc., filiale entièrement détenue par Novell, Inc. ainsi que par le système 4.3. de Berkeley, licencié par l'Université de Californie. Le logiciel détenu par des tiers, et qui comprend la technologie relative aux polices de caractères, est protégé par un copyright et licencié par des fournisseurs de Sun.

LEGENDE RELATIVE AUX DROITS RESTREINTS : l'utilisation, la duplication ou la divulgation par l'administration américaine sont soumises aux restrictions visées à l'alinéa (c)(1)(ii) de la clause relative aux droits des données techniques et aux logiciels informatiques du DFAR 252.227- 7013 et FAR 52.227-19.

Le produit décrit dans ce manuel peut être protégé par un ou plusieurs brevet(s) américain(s), étranger(s) ou par des demandes en cours d'enregistrement.

MARQUES

Sun, Sun Microsystems, le logo Sun, Solaris, SunSoft, the SunSoft logo, Solaris, SunOS, Solstice, OpenWindows, DeskSet, SunFastEthernet, Solstice DiskSuite, Solstice AutoClient, ONC, ONC+, Online: Backup, JumpStart, and NFS sont des marques déposées ou enregistrées par Sun Microsystems, Inc. aux Etats-Unis et dans certains autres pays. UNIX est une marque enregistrée aux Etats-Unis et dans d'autres pays, et exclusivement licenciée par X/Open Company Ltd. OPEN LOOK est une marque enregistrée de Novell, Inc., PostScript et Display PostScript sont des marques d'Adobe Systems, Inc., ORACLE® and SQL*NET® are registered trademarks of Oracle Corporation. Prestoserve is a trademark of Legato Systems, Inc. All other product, service, or company names mentioned herein are claimed as trademarks and trade names by their respective companies.

Toutes les marques SPARC sont des marques déposées ou enregistrées de SPARC International, Inc. aux Etats-Unis et dans d'autres pays. SPARCcenter, SPARCcluster, SPARCcompiler, SPARCdesign, SPARC811, SPARCengine, SPARCprinter, SPARCserver, SPARCstation, SPARCstorage, SPARCworks, microSPARC, microSPARC II et UltraSPARC sont exclusivement licenciées à Sun Microsystems, Inc. Les produits portant les marques sont basés sur une architecture développée par Sun Microsystems, Inc.

Les utilisateurs d'interfaces graphiques OPEN LOOK® et Sun™ ont été développés par Sun Microsystems, Inc. pour ses utilisateurs et licenciés. Sun reconnaît les efforts de pionniers de Xerox pour la recherche et le développement du concept des interfaces d'utilisation visuelle ou graphique pour l'industrie de l'informatique. Sun détient une licence non exclusive de Xerox sur l'interface d'utilisation graphique, cette licence couvrant aussi les licences de Sun qui mettent en place OPEN LOOK GUIs et qui en outre se conforment aux licences écrites de Sun.

Le système X Window est un produit du X Consortium, Inc.

CETTE PUBLICATION EST FOURNIE "EN L'ETAT" SANS GARANTIE D'AUCUNE SORTE, NI EXPRESSE NI IMPLICITE, Y COMPRIS, ET SANS QUE CETTE LISTE NE SOIT LIMITATIVE, DES GARANTIES CONCERNANT LA VALEUR MARCHANDE, L'APTITUDE DES PRODUITS A REpondre A UNE UTILISATION PARTICULIERE OU LE FAIT QU'ILS NE SOIENT PAS CONTREFAISANTS DE PRODUITS DE TIERS.

CETTE PUBLICATION PEUT CONTENIR DES MENTIONS TECHNIQUES ERRONEES OU DES ERREURS
TYPOGRAPHIQUES. DES CHANGEMENTS SONT PERIODIQUEMENT APPORTES AUX INFORMATIONS CONTENUES
AUX PRESENTES, CES CHANGEMENTS SERONT INCORPORES AUX NOUVELLES EDITIONS DE LA PUBLICATION.
SUN MICROSYSTEMS INC. PEUT REALISER DES AMELIORATIONS ET/OU DES CHANGEMENTS DANS LE(S)
PRODUIT(S) ET/OU LE(S) PROGRAMME(S) DECRITS DANS DETTE PUBLICATION A TOUS MOMENTS.

Reader Comments

We welcome your comments and suggestions to help improve this manual. Please let us know what you think about the *SPARCcluster High Availability Server Software Planning and Installation Guide*, part number 802-3509-10.

- The procedures were well documented.

Strongly Agree	Agree	Disagree	Strongly Disagree	Not Applicable
<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>

Comments _____

- The tasks were easy to follow.

Strongly Agree	Agree	Disagree	Strongly Disagree	Not Applicable
<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>

Comments _____

- The illustrations were clear.

Strongly Agree	Agree	Disagree	Strongly Disagree	Not Applicable
<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>

Comments _____

- The information was complete and easy to find.

Strongly Agree	Agree	Disagree	Strongly Disagree	Not Applicable
<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>

Comments _____

- Do you have additional comments about the *SPARCcluster High Availability Server Software Planning and Installation Guide*?

Name: _____


Title: _____

Company: _____


Address: _____

Telephone: _____

Email address: _____



NO POSTAGE
NECESSARY
IF MAILED
IN THE
UNITED STATES



BUSINESS REPLY MAIL
FIRST CLASS MAIL PERMIT NO. 1 MOUNTAIN VIEW, CA

POSTAGE WILL BE PAID BY ADDRESSEE



SUN MICROSYSTEMS, INC.
Attn: Manager, Hardware Publications
MS MPK 14-101
2550 Garcia Avenue
Mt. View, CA 94043-9850

