



Sun Cluster 2.2 7/00 Release Notes

Updated February 2002

Sun Microsystems, Inc.
4150 Network Circle
Santa Clara, CA 95054 U.S.A.
650-960-1300

Part No. 806-5345-10
February 2002, Revision A

Copyright 2002 Sun Microsystems, Inc., 901 San Antonio Road, Palo Alto, California 94303-4900 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. Third-party software, including font technology, is copyrighted and licensed from Sun suppliers.

Parts of the product may be derived from Berkeley BSD systems, licensed from the University of California. UNIX is a registered trademark in the U.S. and other countries, exclusively licensed through X/Open Company, Ltd. For Netscape Communicator™, the following notice applies: (c) Copyright 1995 Netscape Communications Corporation. All rights reserved.

Sun, Sun Microsystems, the Sun logo, AnswerBook2, docs.sun.com, Solstice DiskSuite, Solaris, OpenBoot, SunSolve Online, SunSolve EarlyNotifier, Ultra Enterprise, and StorEdge are trademarks, registered trademarks, or service marks of Sun Microsystems, Inc. in the U.S. and other countries. All SPARC trademarks are used under license and are trademarks or registered trademarks of SPARC International, Inc. in the U.S. and other countries. Products bearing SPARC trademarks are based upon an architecture developed by Sun Microsystems, Inc.

The OPEN LOOK and Sun™ Graphical User Interface was 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 GUIs and otherwise comply with Sun's written license agreements.

PostScript is a trademark or registered trademark of Adobe Systems, Incorporated, which may be registered in certain jurisdictions.

Federal Acquisitions: Commercial Software-Government Users Subject to Standard License Terms and Conditions.

DOCUMENTATION IS PROVIDED "AS IS" AND ALL EXPRESS OR IMPLIED CONDITIONS, REPRESENTATIONS AND WARRANTIES, INCLUDING ANY IMPLIED WARRANTY OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE OR NON-INFRINGEMENT, ARE DISCLAIMED, EXCEPT TO THE EXTENT THAT SUCH DISCLAIMERS ARE HELD TO BE LEGALLY INVALID.

Copyright 2002 Sun Microsystems, Inc., 901 San Antonio Road, Palo Alto, Californie 94303 Etats-Unis. Tous droits réservés.

Ce produit ou document est protégé par un copyright et distribué avec des licences qui en restreignent l'utilisation, la copie, la distribution, et la décompilation. Aucune partie de ce produit ou document ne peut être reproduite 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 y en a. 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.

Des parties de ce produit pourront être dérivées des systèmes Berkeley BSD licenciés par l'Université de Californie. UNIX est une marque déposée aux Etats-Unis et dans d'autres pays et licenciée exclusivement par X/Open Company, Ltd. La notice suivante est applicable à Netscape Communicator™: (c) Copyright 1995 Netscape Communications Corporation. Tous droits réservés.

Sun, Sun Microsystems, le logo Sun, AnswerBook2, docs.sun.com, Solstice DiskSuite, Solaris, OpenBoot, SunSolve Online, SunSolve EarlyNotifier, Ultra Enterprise, et StorEdge sont des marques de fabrique ou des marques déposées, ou marques de service, de Sun Microsystems, Inc. aux Etats-Unis et dans d'autres pays. Toutes les marques SPARC sont utilisées sous licence et sont des marques de fabrique ou des marques déposées de SPARC International, Inc. aux Etats-Unis et dans d'autres pays. Les produits portant les marques SPARC sont basés sur une architecture développée par Sun Microsystems, Inc.

L'interface d'utilisation graphique OPEN LOOK et Sun™ a été développée 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 Xerox, cette licence couvrant également les licenciés de Sun qui mettent en place l'interface d'utilisation graphique OPEN LOOK et qui en outre se conforment aux licences écrites de Sun.

PostScript est une marque de fabrique d'Adobe Systems, Incorporated, laquelle pourrait être déposée dans certaines juridictions.

CETTE PUBLICATION EST FOURNIE "EN LETAT" ET AUCUNE GARANTIE, EXPRESSE OU IMPLICITE, N'EST ACCORDEE, Y COMPRIS DES GARANTIES CONCERNANT LA VALEUR MARCHANDE, L'APTITUDE DE LA PUBLICATION A REPENDRE A UNE UTILISATION PARTICULIERE, OU LE FAIT QU'ELLE NE SOIT PAS CONTREFAISANTE DE PRODUIT DE TIERS. CE DENI DE GARANTIE NE S'APPLIQUERAIT PAS, DANS LA MESURE OU IL SERAIT TENU JURIDIQUEMENT NUL ET NON AVENU.



Adobe PostScript

Sun Cluster 2.2 7/00 Release Notes

This document provides information about the following topics.

- ["Revision Record" on page 4](#)
- ["Features" on page 5](#)
- ["Command Changes From Previous Releases" on page 7](#)
- ["Restrictions" on page 8](#)
- ["Licensing" on page 9](#)
- ["Patches" on page 9](#)
- ["Volume Managers" on page 12](#)
- ["Data Services" on page 16](#)
- ["Installation Notes" on page 18](#)
- ["Upgrade Notes" on page 19](#)
- ["AnswerBooks" on page 20](#)
- ["Known Problems" on page 23](#)
- ["Undocumented Error Messages" on page 40](#)
- ["Future Changes" on page 46](#)
- ["Notes and Issues for Localized Versions" on page 48](#)

Revision Record

The following table lists recently discovered software or documentation problems, new data service support, and new hardware qualifications. See the referenced section in this document for details.

TABLE 0-1 Sun Cluster 2.2 Release Notes Revision Record

Revision Date	Information Added
February 2002	Restriction on dynamic reconfiguration for cluster nodes with Sbus-SCI cards. See “Hardware Notes” on page 17 .
June 2001	Clarification to BugId 4345750 that the documentation change applies only to clusters that run on Solaris 7 11/99 and later. See “Documentation Errata” on page 30 .
May 2001	Support for Instant Image 2.0 coexistence in “Features” on page 5 . Support for Solaris Resource Manager versions 1.0 and 1.2 coexistence in “Features” on page 5 . See the Solaris Resource Manager 1.2 Documentation Errata in the Solaris Resource Manager 1.2 Collection on http://sunsolve.sun.com for information about this coexistence feature. List of patches required for support of Instant Image 2.0 coexistence and Solaris Resource Manager 1.2 coexistence in “Patches” on page 9 . BugId 4448815 and 4448860 to “Documentation Errata” on page 30 .
April 2001	BugId 4393512 to “Hardware Qualification Bugs” on page 25 . Removed the procedure on “How to Replace a Sun StorEdge T3 Disk Array.” This is not a supported procedure and was included in the Sun StorEdge T3 disk array documentation in error.
March 2001	How to install Apache software from the Solaris 8 CD-ROM for the Sun Cluster HA for Apache data service. See the information following the table of supported data services in “Data Services” on page 16 . Restriction about setting <code>local-mac-address?=true</code> to “Restrictions” on page 8 .
February 2001	BugId 4367622, 4374280, and 4399132 in “Hardware Qualification Bugs” on page 25 . BugId 4405556 in “Data Service Bugs” on page 28 . Support for Sun StorEdge™ T3 disk trays in “Features” on page 5 and Appendix A .

Features

Sun™ Cluster 2.2 7/00 Release includes the following features:

- Support for Solaris™ 2.6, Solaris 7, and Solaris 8.
- Support for Solstice DiskSuite™ on Solaris 2.6, Solaris 7, and Solaris 8. See [“Volume Managers” on page 12](#) for details.
- Support for VERITAS Volume Manager 3.0.4 and VERITAS Volume Manager 3.0.4 cluster functionality (formerly known as Cluster Volume Manager). See [“Volume Managers” on page 12](#) for more information.
- Support for both Solstice DiskSuite and VERITAS Volume Manager in a single cluster configuration, when you use Solstice DiskSuite to manage the local disks and VERITAS Volume Manager to control the multihost disks. See [“Volume Managers” on page 12](#) for more information.
- New data service, Sun Cluster HA for NetBackup, supported on Solaris 2.6 only. See [“Data Services” on page 16](#) for the complete list of supported data services.
- Japanese, Korean, and French localization of system messages, end-user documentation, and online help.
- Online documentation in AnswerBook2™ format for the English locale, and AnswerBook2, derived HTML (DHTML), and PostScript™ formats for Japanese, French, and Korean locales. See [“AnswerBooks” on page 20](#) for more information.
- Upgrade from Solstice HA 1.3, Sun Cluster 2.0, Sun Cluster 2.1, and earlier versions of Sun Cluster 2.2. See [“Upgrade Notes” on page 19](#) for details.
- Support for Scalable Coherent Interface/Peripheral Component Interconnect (SCI PCI) for private networks, on Solaris 8 only.
- Support for Sun StorEdge T3 disk trays. 1.16a level firmware, patch number 110760-01 is required to use a Sun StorEdge T3 disk tray in a cluster. See [Appendix A](#) for hardware installation and maintenance procedures.
- Support for Instant Image 2.0 coexistence with Sun Cluster 2.2 in Solaris 2.6 and Solaris 8 environments. See [“Patches” on page 9](#) for a list of patches required to implement this feature.
- Support for Solaris Resource Manager 1.0 coexistence with Sun Cluster 2.2 in a Solaris 2.6 environment, and support for Solaris Resource Manager 1.2 coexistence with Sun Cluster 2.2 in a Solaris 8 environment. See [“Patches” on page 9](#) for a list of patches required to implement this feature. Read the Solaris Resource Manager 1.2 Documentation Errata in the Solaris Resource Manager 1.2 Collection on <http://docs.sun.com> for information about this coexistence feature. This Errata is applicable to both the 1.0 and 1.2 versions of Solaris Resource Manager.

Enhancements to Process Monitoring

The process monitoring daemon (`pmfd`) in Sun Cluster has been enhanced to provide greater control of process monitoring:

- `pmfd` can now be used to restrict monitoring to a subset of all processes. Therefore, any performance impacts associated with monitoring can be limited to only those processes being monitored.
- `pmfd` is less likely to reach resource limits when the number of processes monitored is reduced.
- The new features can be used to temporarily switch off monitoring to allow processes to be traced (with `truss` or `pstack`, for example) for diagnostic purposes.
- `pmfd` can now correctly identify a failed process, even if its descendants are still running.

The new features are implemented through the `-C` option to `pmfadm(1M)`, and are described in a revised `pmfadm(1M)` man page. The new option and revised man page are available for Sun Cluster 2.2 in cluster framework patch 109208, and for Sun Cluster 2.1 in patch 105458-15. Obtain the patches from your service provider or from the Sun patch web site, <http://sunsolve.sun.com>.

Command Changes From Previous Releases

The following Solstice HA 1.3 commands have been replaced in or removed from Sun Cluster 2.2. See the associated man pages for more information.

	Solstice HA 1.3	Sun Cluster 2.2
Replaced	hainstall	scinstall
	hainetconfig	hadsconfig
	haremove	scinstall
	hasetup	scconf and scinstall
	hastart	scadmin startcluster (first node) scadmin startnode (remaining nodes)
	hastop	scadmin stopnode
Removed	hacheck	
	hafstab	
	halicense	
	haload	

The syntax and usage of the Sun Cluster 2.1 `scinstall(1M)` command has been changed for Sun Cluster 2.2. For current syntax and usage, refer to the `scinstall(1M)` man page and to Chapter 3 in the *Sun Cluster 2.2 Software Installation Guide*.

Restrictions

The following restrictions apply to Sun Cluster 2.2 7/00 Release software at time of initial release. See your service provider for the latest information about supported products and features.

- The layered volumes feature of VERITAS Volume Manager 3.0.4 is not supported on clusters using the volume manager cluster feature (formerly known as Cluster Volume Manager).
- Use of SCI PCI for the Sun Cluster private network is supported on all hardware platforms, but only on Solaris 8. See Appendix B of the *Sun Enterprise Cluster System Site Preparation, Planning, and Installation Guide* for more information about using SCI PCI components with Sun Cluster.
- Sun Cluster HA for NetBackup is supported only on Solaris 2.6 at introduction of Sun Cluster 2.2 7/00 Release. Support on Solaris 8 is expected to follow shortly. See your service provider for the latest support information.
- The Remote System Control (RSC) feature of the Ultra Enterprise™ 250 server is not supported with Sun Cluster 2.2 7/00 Release.
- On Solaris 8, you cannot run Sun Cluster Manager as an application (viewed through a browser). To use Sun Cluster Manager with Solaris 8, you must run it as a standalone application, as described in Chapter 2 of the *Sun Cluster 2.2 Software Installation Guide*. You can run Sun Cluster Manager as either an application or applet with the Solaris 2.6 and 7 operating environments.
- You cannot place SCI and UDWIS cards on the same SBus in Ultra Enterprise 3x00/4x00/5x00/6x00/10000 configurations. SCI and UDWIS cards must be placed on separate SBuses when these servers are used.
- Although the shared CCD is an optional feature for two-node clusters running VERITAS Volume Manager, a shared CCD cannot be used in Solstice DiskSuite configurations.
- Cross-mounting of Sun Cluster HA for NFS resources is not supported.
- The user-defined script `clustername.reconfig.user_script` is not supported in Sun Cluster 2.2.
- The `+D` option to `scconf(1M)` cannot be used with Solstice DiskSuite.
- It is not supported to set `local-mac-address?` to `true` in Sun Cluster 2.2.

Licensing

You will receive a paper license for the Sun Cluster 2.2 framework, one for each hardware platform on which Sun Cluster 2.2 will run. You will also receive a paper license for each Sun Cluster data service, one per node. The Sun Cluster 2.2 framework does not enforce these licenses, but you should retain the paper licenses as proof of ownership when you need technical support or other support services.

You do not need licenses to run Solstice DiskSuite with a licensed Sun Cluster 2.2 configuration. However, you need a license for VERITAS Volume Manager (VxVM) and optionally for VERITAS Volume Manager cluster functionality (formerly called Cluster Volume Manager). The base VxVM license certificates are included with Sun Cluster Server license kits, and VxVM cluster functionality license certificates are bundled with Oracle Parallel Server Right-To-Use license kits. The Sun Cluster and Oracle Parallel Server license kits are available from Sun. Follow the instructions printed on the license certificates to obtain active license keys.

You might need to obtain licenses for DBMS products and other third-party products. Contact your third-party service provider for third-party product licenses. See <http://www.sun.com/licensing/> for more information.

Patches

A set of patches is included on the Sun Cluster 2.2 7/00 product CD. There are also patches you may need to install from the SunSolve Online web site. This section describes how to locate both types of patches.

Installing Patches From CD

You install the patches from the Sun Cluster 2.2 7/00 product CD during the installation process, using the `install_scpatches` command as documented in Chapter 3 of the *Sun Cluster 2.2 Software Installation Guide*.

The patches are divided by operating environment and are located in the directory `/cdrom/multi_suncluster_sc_2_2/Sun_Cluster_2_2/Sol_2.x/Patches`.

Installing Patches From SunSolve

In addition to the patches included on the product CD, there are other patches that might be recommended or required for your Sun Cluster configuration. Obtain the patches from your service provider or from the Sun patch web site <http://sunsolve.sun.com>. Follow the instructions in the patch README files to install the patches.

See your service provider for the latest information about required and recommended patches.

TABLE 0-2 Additional Patches for Sun Cluster 2.2 7/00 Release

Patch Number	Description
108508	Solaris 2.6: VxVM 3.0.4 layered volumes
108509	Solaris 7: VxVM 3.0.4 layered volumes
109210	Solaris 8: VxVM 3.0.4 layered volumes
108423: hadconfig(1M) 108450: hadconfig(1M) 109214: scinstall(1M)	Solaris 2.6: Sun Cluster HA for NetBackup patch
109208	Framework mini-jumbo patch for Sun Cluster 2.2
110871-02 (or later): II support patch 109208-06 (or later) 109967 (or later): Core patch 109975 (or later): II patch 109983 (or later): STE patch	Solaris 2.6: Sun Cluster 2.2 and Instant Image 2.0 patches for coexistence support. See the README file for patch 110871-02 (or later) for installation and use instructions.
110871-02 (or later): II support patch 109210-07 (or later) 109970 (or later): Core patch 109978 (or later): II patch 109986 (or later): STE patch	Solaris 8: Sun Cluster 2.2 and Instant Image 2.0 patches for coexistence support. See the README file for patch 110871-02 (or later) for installation and use instructions.
110653-01 (or later): SRM support patch 109208-07 (or later): Framework Patch 109211-03 (or later): Comm Jumbo Patch 107996-09 (or later): HA Oracle Patch (if using HA-Oracle)	Solaris 2.6: Sun Cluster 2.2 and Solaris Resource Manager 1.0 patches for coexistence support.
110655-01 (or later): SRM support patch 109210-06 (or later): Framework Patch 109213-03 (or later): Comm Jumbo Patch 109426-04 (or later): HA Oracle Patch (if using HA-Oracle)	Solaris 8: Sun Cluster 2.2 and Solaris Resource Manager 1.2 patches for coexistence support.

Finding Patches Online Using SunSolve

The SunSolve OnlineSM Web site provides 24-hour access to the most up-to-date information regarding patches, software, and firmware for Sun products. Access the SunSolve Online site at <http://sunsolve.sun.com> for the most current matrixes of supported software, firmware, and patch revisions.

Note – You must have a SunSolve account registered to view and download the required patches for the Sun Cluster product. If you don't have an account registered, contact your Sun service representative or sales engineer, or register through the SunSolve Online Web site.

You can find Sun Cluster 2.2 patch information using the SunSolve EarlyNotifierSM page. To view the EarlyNotifier information, follow these steps.

1. **Log into the SunSolve Online site at <http://sunsolve.sun.com>.**
2. **Access the Simple Search selection from the top of the main page.**
3. **In the Simple Search page, scroll down to Step 2 to enter your search criteria.**
4. **Type `sun cluster 2.2` in the Entire document text box or type `19224` in the Document id text box.**
Make any other changes to the search options that you want to make.
5. **Click Go in Step 9.**
6. **When Simple Search displays your results, click the highlighted number 19224 in the list.**

This will bring up the EarlyNotifier page for Sun Cluster 2.2.

Before installing Sun Cluster 2.2 and applying patches to a cluster component (Solaris operating system, Sun Cluster software, volume manager or data services software, or disk hardware), review the EarlyNotifier information and any README files that accompany the patches. All cluster nodes must have the same patch level for proper cluster operation.

For specific patch procedures and tips on administering patches, see the *Sun Cluster 2.2 System Administration Guide*.

Volume Managers

Sun Cluster 2.2 7/00 Release supports the following volume managers:

Volume Manager	Release	Solaris Version		
		2.6	7	8
Solstice DiskSuite	4.2	X	X	
	4.2.1			X
VERITAS Volume Manager	3.0.4	X	X	X
VERITAS Volume Manager cluster feature	3.0.4	X		X

Note – VERITAS Volume Manager 3.0.4 includes the product formerly known as Cluster Volume Manager. This functionality is now called VERITAS Volume Manager cluster feature, cluster capability, or cluster functionality. This new terminology is used in both VERITAS and Sun Cluster documentation.

Note – You can use Solstice DiskSuite and VERITAS Volume Manager together only if you use Solstice DiskSuite to manage the local disks and VERITAS Volume Manager to control the multihost disks. In such a configuration, plan your physical disk needs accordingly. You might need additional disks for the VERITAS Volume Manager root disk group, for example. See your volume manager documentation for more information.

VERITAS Volume Manager Notes

Sun Cluster 2.2 7/00 Release supports VERITAS Volume Manager 3.0.4 on Solaris 2.6, 7, and 8.

Sun Cluster 2.2 7/00 Release supports the VERITAS Volume Manager 3.0.4 cluster feature (formerly called Cluster Volume Manager), only when used with Oracle Parallel Server. This combination is supported on Solaris 2.6 and Solaris 8.

Command Notes

The following (1M) commands and options are supported only with VxVM. See the associated man pages for more information.

- `confccdssa`
- `scconf -q`
- `scconf -F`
- `scconf -S`
- `scadmin reldisks`
- `scadmin resdisks`

Dynamic Multi-Pathing Feature

The VERITAS Volume Manager Dynamic Multipathing (DMP) feature is not supported with Sun Cluster 2.2 7/00 Release. You must explicitly disable DMP using the procedure documented in the “DMP Issues” section of the VERITAS Volume Manager 3.0.4 Release Notes.

Layered Volumes Feature

The VERITAS Volume Manager layered volumes feature is supported on Sun Cluster 2.2 7/00 Release, but a patch is required to alleviate a problem that impacts creation and switchover of logical hosts associated with disk groups that contain subvolumes. See [“Patches” on page 9](#) for patch details.

Note – The layered volumes feature is not supported on clusters using the VERITAS Volume Manager cluster feature (formerly called Cluster Volume Manager).

Solstice DiskSuite Notes

Sun Cluster 2.2 7/00 Release supports Solstice DiskSuite 4.2 on Solaris 2.6 and 7, and Solstice DiskSuite 4.2.1 on Solaris 8.

Command Notes

The following command changes apply to Solstice DiskSuite when used with Sun Cluster 2.2 7/00 Release. See the associated man pages for more information.

- The `reserve` option to the `scadmin(1M)` command is used only with Solstice DiskSuite.
- The `+D` option to `scconf(1M)` cannot be used with Solstice DiskSuite.
- The `scdidadm(1M)` command is used in Solstice DiskSuite configurations to create mappings between disk ID instance numbers and remote paths to the disk, and to check disk ID major numbers.

Location of Solstice DiskSuite Packages

The Solstice DiskSuite 4.2 product CD is co-packaged with Sun Cluster 2.2 7/00 Release. Note that the Solstice DiskSuite 4.2 mediators package (`SUNWwmdm`) is included on the Sun Cluster 2.2 product CD, in the directory `/cdrom/multi_suncluster_sc_2_2/Sun_Cluster_2_2/Sol2_x/Product`.

The Solstice DiskSuite 4.2.1 product is included on the Solaris 8 CD. Note that the Solstice DiskSuite 4.2.1 mediators package (`SUNWwmdm`) is included on the Sun Cluster 2.2 product CD, in the directory `/cdrom/multi_suncluster_sc_2_2/Sun_Cluster_2_2/Sol2_x/Packages`.

Installing or Upgrading Solstice DiskSuite

To install or upgrade Solstice DiskSuite, use the detailed instructions included in your Solstice DiskSuite documentation. To access the Solstice DiskSuite documentation, perform these steps:

1. **Open the `README` file contained on the Solstice DiskSuite CD (for Solstice DiskSuite 4.2) or Solaris 8 CD (for Solstice DiskSuite 4.2.1), using a browser to access the menu options that enable you to read an HTML file.**

For example, in Netscape, do the following:

- a. **From the Netscape browser menu bar, choose `File>Open Page>Choose File`.**

This opens the File Browser dialog box.

- b. **Choose the file `/cdrom/cdrom0/README.html`.**

The browser brings up the `README.html` file.

2. **Install the AnswerBook2 server and the Solstice DiskSuite AnswerBook using the `README` file instructions.**

3. Access the Solstice DiskSuite AnswerBook and follow the online instructions found in the *Solstice DiskSuite 4.x Installation and Product Notes* to install Solstice DiskSuite.

Note – The latest version of Patch 106627 is required for Solstice DiskSuite 4.2 running on either Solaris 2.6 or Solaris 7. The patch is available from all Sun service providers and from the Sun patch web site, <http://sunsolve.sun.com>.

Configuring Mediators When Migrating From Solstice HA 1.3 to Sun Cluster 2.2 7/00 Release

This section is relevant only for clusters that were originally set up under Solstice HA 1.3 using Solstice DiskSuite mediators (two-string configurations). It describes changes that are automatically made to a mediator configuration when you upgrade from Solstice HA 1.3 to Sun Cluster 2.2. There is no direct user impact, but you should note the changes in any configuration information you keep on the cluster.

The documented Solstice HA 1.3-to-Sun Cluster 2.2 upgrade procedure changes the Solstice HA 1.3 mediator configuration. In Solstice HA 1.3, the hosts referred to the private links by physical names, whereas in Sun Cluster 2.2, the private link IP addresses are used. The original Solstice HA 1.3 mediator configuration resembles the following:

Mediator Host(s)	Aliases
ha-red	ha-red-priv1, ha-red-priv2
ha-green	ha-green-priv1, ha-green-priv2

After running the Sun Cluster 2.2 upgrade procedure, this configuration is converted to one that resembles the following:

Mediator Host(s)	Aliases
ha-red	204.152.65.34
ha-green	204.152.65.33

For more information about configuring mediators for Sun Cluster 2.2, see Chapter 9, "Using Dual-String Mediators," in the *Sun Cluster 2.2 System Administration Guide*.

Data Services

Sun Cluster 2.2 7/00 Release supports the following data services.

Note – The data services and data service versions supported by Sun Cluster 2.2 7/00 Release are updated frequently. Consider the following table a checkpoint. See your service provider for the most current information about which data services and versions are supported.

HA Data Service	Application and Version	Solaris Version		
		2.6	7	8
N/A	Informix-Online XPS 8.2.1	X		
N/A	Oracle Parallel Server 7.3.4, 8.0.4, 8.0.5 and 8i (8.1.5)	X		
N/A	Oracle Parallel Server 8.0.6 and 8i (8.1.6)	X		X
Sun Cluster HA for Oracle	7.3.4	X		
	8.0.4, 8.0.5, 8.0.6, and 8i (8.1.5)	X	X	
	8i 8.1.6	X	X	X
Sun Cluster HA for Sybase	Sybase 11.5	X		
	Sybase 11.9.2, 12.0	X	X	
Sun Cluster HA for Informix	Informix 7.23	X		
	Informix 7.30	X		
Sun Cluster HA for NFS	NFS 2.0, 3.0	X	X	X
Sun Cluster HA for DNS	DNS	X	X	X
Sun Cluster HA for Netscape	Netscape Mail 3.5	X		
	Netscape Messaging Server 4.1	X	X	
	Netscape Directory Server (LDAP) 4.1	X		X
	Netscape HTTP Server 3.5,	X		
	Netscape HTTP Server 3.6, 4.0, HTTP Secure 3.6	X		X
	iPlanet Web Server 4.1			X
Sun Cluster HA for SAP	SAP 3.1h, 3.1i, 4.0b, 4.5b with Oracle	X		
	SAP 4.6b with Oracle 8.0.6	X		
	SAP 4.5b with Informix	X		
Sun Cluster HA for Lotus	Lotus 4.6, 4.6.1, 4.6.3	X		
Sun Cluster HA for Tivoli	Tivoli 3.2, 3.6	X		
Sun Cluster HA for NetBackup	NetBackup 3.2	X		
Sun Cluster HA for Apache	Apache Web Server 1.39			X

The Sun Cluster 2.2 Data Services Update: Apache Web Server describes the procedure for installing the Apache Web Server from the Apache web site (<http://www.apache.org>). However, you can also install the Apache Web Server from the Solaris 8 operating environment CD-ROM.

The Apache binaries are included in three packages—`SUNWapchr`, `SUNWapchu`, and `SUNWapchd`—that form the `SUNWCapache` package metacluster. You must install `SUNWapchr` before `SUNWapchu`.

Place the Web server binaries on the local file system on each of your cluster nodes.

Hardware Notes

Sun Cluster 2.2 software does not support the dynamic reconfiguration (DR) features of the Sun Enterprise E3x00 through E6x00 servers when these servers use SBus-SCI communications interface cards. An attempt to replace a server component while the server is still a running, active cluster member can cause unplanned loss of service.

To replace a server component in a cluster environment, first check the *Sun Enterprise Cluster Hardware Service Manual* (805-6512) to determine what special tasks, if any, you must perform. In general, to replace server components you must first switch over the data services to another functioning node of the cluster, halt the node to be serviced, and power down the node. Then you are ready to perform the hardware procedure to replace the component. After the procedure is complete, rejoin the node to the cluster and switch back the logical hosts to the default masters.

Installation Notes

The Sun Cluster 2.2 7/00 Release media kit consists of the following CDs:

- Sun Cluster 2.2 7/00 Release software, which includes Sun Cluster HA data services.
- Solstice DiskSuite 4.2 software and Solstice DiskSuite online documentation in AnswerBook2 format. (Note that Solstice DiskSuite 4.2.1 software is included on the Solaris 8 CD.)
- VERITAS Volume Manager 3.0.4 software, which includes cluster functionality formerly known as Cluster Volume Manager.

For more information about installing Solstice DiskSuite software and documentation, see [“Installing or Upgrading Solstice DiskSuite” on page 14](#).

For instructions to install and configure VxVM, see your VERITAS Volume Manager documentation and the *Sun Cluster 2.2 Software Installation Guide*.

See also [“Installation Bugs” on page 25](#).

Overview of Installation Procedures

The Sun Cluster installation procedures have changed significantly from Solstice HA 1.3 and Sun Cluster 2.1. In Sun Cluster 2.2, the interactive command `scinstall(1M)` is used to install the software and to set up cluster components such as logical hosts and network interfaces.

For detailed installation procedures, see Chapter 3, “Installing and Configuring Sun Cluster Software,” in the *Sun Cluster 2.2 Software Installation Guide*.

Generally, the steps to install and configure Sun Cluster are grouped into three procedures:

1. Preparing the administrative workstation and installing the client software.

This entails installing the Solaris operating environment and Sun Cluster 2.2 client software on the administrative workstation.

2. Installing the server software.

This includes using the Cluster Console to install the Solaris operating environment and Sun Cluster 2.2 software on all cluster nodes; using `scinstall(1M)` to set up network interfaces, logical hosts, and quorum devices; and selecting data services and volume manager support packages.

3. Configuring and bringing up the cluster.

This includes setting up paths; installing patches; installing and configuring your volume manager, SCI, PNM backup groups, logical hosts, and data services; and bringing up the cluster.

Upgrade Notes

You can upgrade to Sun Cluster 2.2 7/00 Release from the following platforms:

- Solstice HA 1.3
- Sun Cluster 2.0
- Sun Cluster 2.1
- Sun Cluster 2.2 Base Release
- Sun Cluster 2.2 4/00 Release (running Solaris 2.6 or 7)

These upgrade scenarios are documented in the *Sun Cluster 2.2 Software Installation Guide*.

No upgrade is necessary to migrate to Sun Cluster 2.2 7/00 Release on Solaris 8 from earlier versions of Sun Cluster 2.2 on Solaris 8. Instead, just update all cluster nodes with any Sun Cluster or Solaris 8 patches, available from your service provider or from the Sun patch web site, <http://sunsolve.sun.com>.

See also “[Installing or Upgrading Solstice DiskSuite](#)” on page 14 and “[Upgrade Bugs](#)” on page 26.



Caution – You cannot upgrade to Sun Cluster 2.2 7/00 Release on Solaris 8 versions 4/00 or later from configurations that do not use disk IDs. (Use of disk IDs is optional in Solstice HA 1.3 and Sun Cluster 2.2.) This limitation is caused by behavior changes in Solstice DiskSuite on Solaris 8 versions 4/00 or later.

If your pre-upgrade cluster configuration does not use disk IDs, you will have to reinstall completely instead of performing the upgrade procedures.

AnswerBooks

The Sun Cluster 2.2 7/00 Release user documentation is available in AnswerBook2 format for use with AnswerBook2 documentation servers. The Sun Cluster 2.2 7/00 Release AnswerBook2 documentation set consists of the following books, in the languages indicated:

- *Sun Cluster 2.2 Software Installation Guide* (English, Japanese)
- *Sun Cluster 2.2 System Administration Guide* (English, Japanese, French, Korean)
- *Sun Cluster 2.2 API Developer's Guide* (English, Japanese)

In addition to the books listed above, the Sun Cluster 2.2 7/00 Release hardcopy documentation kits also include the following books:

- *Sun Cluster 2.2 Hardware Site Preparation, Planning, and Installation Guide* (English)
- *Sun Cluster 2.2 Hardware Service Manual* (English)
- *Sun Cluster 2.2 Error Messages Guide* (English)

Setting Up the AnswerBook2 Documentation Server

AnswerBook2 documentation server software is included as part of the Solaris operating system release. The documentation server software is included on a Solaris documentation CD that is separate from the Solaris operating environment CD. You need this CD to install an AnswerBook2 documentation server.

If you have an AnswerBook2 documentation server installed at your site, you can use the same server for the Sun Cluster 2.2 AnswerBooks. If you do not have an AnswerBook2 documentation server installed, install a documentation server on a machine at your site. The administrative console that you will use as the administrative interface to your cluster is a good choice for the documentation server. A cluster node is not a good choice.

For complete information about installing an AnswerBook2 documentation server, load the Solaris documentation CD on a server and view the `README` files.

Viewing Sun Cluster AnswerBooks

Use the following procedure to view Sun Cluster 2.2 AnswerBooks from your AnswerBook2 documentation server. Install the Sun Cluster AnswerBook2 documents on a file system on the same server on which the documentation server is installed. The Sun Cluster 2.2 AnswerBook2 packages include a post-install script that will automatically add the AnswerBooks to your existing AnswerBook2 library.

You can also view the Japanese, French, or Korean Sun Cluster 2.2 documentation in PostScript or derived HTML (DHTML) formats directly from the Sun Cluster 2.2 CD. The documents are located in the directory
`/Sun_Cluster_2_2/Sol_2.x/Docs/locale/locale.`

Install the AnswerBooks using the following procedure. You will need:

- Root access to the server on which the AnswerBook2 documentation server is installed. This server must have a CD drive. If your site does not have an AnswerBook2 documentation server installed already, you need the Solaris operating environment documentation CD. It contains the software and instructions you need to install an AnswerBook2 documentation server.
- The Sun Cluster 2.2 7/00 Release CD, which includes the Sun Cluster 2.2 7/00 Release AnswerBook2 collection.

▼ How to Install the Sun Cluster AnswerBooks

Use this procedure to install the Sun Cluster 2.2 7/00 AnswerBook2 packages.

1. **Become super user on the server on which the AnswerBook2 documentation server is installed.**
2. **Load the Sun Cluster 2.2 7/00 Release CD into a CD drive attached to your documentation server.**

The Volume Management daemon `vold(1M)` should mount the CD automatically.

3. Change directory to the location on the CD that contains the Sun Cluster AnswerBook2 packages and install the packages.

From the `pkgadd` installation options menu, choose `heavy` to add the complete package to the system and to update the AnswerBook2 catalog.

For Sun Cluster 2.2 7/00 Release, the packages are located in
`/cdrom/multi_suncluster_sc_2_2/Sun_Cluster_2_2/Sol_2.x/Product`.

```
# cd /cdrom/multi_suncluster_sc_2_2/Sun_Cluster_2_2/Sol_2.x/Product
# pkgadd -d . SUNWscab (English)
# pkgadd -d . SUNWfrabh (French)
# pkgadd -d . SUNWjabha (Japanese)
# pkgadd -d . SUNWkabha (Korean)
```

The AnswerBook2 packages included on the Sun Cluster CD include a post-install script that adds the collection to the documentation server's database and restarts the server. You should now be able to view the Sun Cluster AnswerBooks using your documentation server.

Known Problems

The following known problems affect the operation of Sun Cluster 2.2 7/00 Release.

Framework Bugs

4132195 – Clusters that include dual-Ultra-2, dual-fas-SunSwift, and dual-Sun StorEdge MultiPack MI-SCSI devices can experience a bug in the fas chip that causes the fas SCSI bus to hang when it is selected by more than one host. This can occur in a number of situations—for example, when a SCSI target device driver is attached or after a dormant detached device is re-attached.

If an active cluster node with SCSI ID target 6 is running while another node using SCSI ID target 7 is rebooted, timeouts or resets might result. Note that Solaris reboots can cause a probe to all possible devices.

To prevent timeouts or resets caused by this bug, remove target 6 (and target 7 if present, and/or any other SCSI ID coinciding with SCSI IDs being used by SunSwift/fas initiators in an MI-SCSI configuration) from both the `st.conf` and `sd.conf` files.

4217658 – The public network monitor (PNM) daemon, `pnmd`, does not perform failover of multicast groups after a network adapter failover occurs. Sun Cluster supports failover of only the default multicast route (224.0.0.0). When a network adapter failover or switchover occurs, the default multicast route is switched to the appropriate adapter, but any client application that had established a multicast group will no longer work. You must restart any client applications in this condition.

4233956 – Cluster fails to come up and displays error messages if IP addresses are not assigned to logical hosts. The error messages might indicate that `ifconfig` failed. To prevent the problem, make sure all logical hosts have entries in the `/etc/hosts` files or name service maps indicating their associated IP addresses, before you attempt to bring up the cluster.

4270573 – The `confccdssa(1M)` command displays error messages and hangs if the disk name you specify contains the default suffix of a subdisk name. Work around the problem by creating disk groups (`sc_dg`) manually or by renaming any disks that contain a numeric suffix of the format `-XX`, so that they do not contain the suffix.

4286442 – In a cluster environment with shared single-ended or Differential SCSI devices, the SCSI chain can be broken when a node is powered off incorrectly or when the SCSI cable is disconnected before the bus is quiesced. This can cause data access errors on the node that is still active. Prevent this problem by following

instructions exactly as documented in the *Sun Enterprise Cluster System Site Preparation, Planning, and Installation Guide* and the *Sun Enterprise Cluster Hardware Service Manual* when powering off a node or disconnecting SCSI cables.

4291427 – In Sun Cluster 2.2 running on Solaris 7, using the `scinstall(1M)` command to remove the client packages can fail with the following error message:

```
Patch 108400 is required to be installed by patch 108446
it cannot be backed out until patch 108446 is backed out.
```

This occurs because of dependencies between patches 108446 and 108400. Work around the problem by removing patches 108446 and 108400 manually and then re-starting the package removal process using `scinstall`.

4296706 – If a connection is lost from a differential SCSI device (A/D1000, A3x00), or if termination is lost due to one cluster node being powered off, the storage device can become inaccessible to the surviving host and the surviving host can panic. Prevent this problem by following instructions exactly as documented in the *Sun Enterprise Cluster System Site Preparation, Planning, and Installation Guide* and the *Sun Enterprise Cluster Hardware Service Manual* when powering off a node or disconnecting SCSI cables.

4299187 – The cluster console does not accept non-ascii characters, for example, Japanese characters or French (accented) characters. Work around the problem by inputting such characters through the individual terminal windows on each cluster node, instead of through the cluster console.

4319412 – Killing `clustd` on the master node panics both the master node and backup node. Prevent or work around the problem by applying a Solstice DiskSuite patch, available from your service provider.

4321549 – Cannot switch over logical host while database instance is running on single-CPU nodes, on clusters using Oracle 8.1.6 and Solstice DiskSuite. Work around the problem by applying patch 108508 (Solaris 2.6) or 108509 (Solaris 7), available from your service provider or from the Sun patch web site, <http://sunsolve.sun.com>.

4326020 – Layered volumes feature of VERITAS Volume Manager 3.0.x: Problems can occur when you create or switch over a logical host associated with a disk group containing layered volumes. Prevent the problem by installing a Sun Cluster patch. See “Patches” on page 9 for patch details.

4326276 – Node failover or removal is prevented on clusters using Instant Image 2.0. Because the volume manager is overlaid with the Instant Image `sv` driver, the Sun Cluster software cannot unmount disk group volumes during failover. Prevent the problem by applying the relevant Sun Cluster and Instant Image patches, available from your service provider or from the Sun patch web site, <http://sunsolve.sun.com>.

Hardware Qualification Bugs

4367622 – Upgrading Sun StorEdge T3 disk tray firmware from 1.16 to 1.16a can result in a hung telnet window. The controller firmware upgrade command `boot -i` with telnet can hang due to lack of memory. Work around the problem by upgrading the controller firmware using a serial port connection to the Sun StorEdge T3 disk tray or by resetting the Sun StorEdge T3 disk tray and trying it again.

4374280 – If you are running a RAID-0 volume on a Sun StorEdge T3 disk tray and you lose a disk drive in this Sun StorEdge T3 disk tray, the Sun StorEdge T3 disk tray continues to make the volume available to the host, resulting in VERITAS Volume Manager delays and overall system performance issues. Work around this problem by using RAID-0 volumes with host-based mirroring configurations.

4399132 – During volume reconstruction using the Sun StorEdge T3 disk tray, if `recon_rate` is set to `high`, nodes cannot join the cluster. Work around this problem by using the factory default (`medium`) for `recon_rate`.

4393512 – SCSI-reservations failures have been observed when clustering StorEdge MultiPack enclosures that contain a particular model of Quantum disk drive: SUN4.2G VK4550J. It is recommended that you do not use this particular model of Quantum disk drive for clustering with StorEdge MultiPack enclosures. If you do use this model of disk drive, you must set the `scsi-initiator-id` of the “first node” to 6. If you are using a six-slot StorEdge MultiPack enclosure, this also requires that you set it for the 9-through-14 SCSI target address range (for more information, see the *Sun StorEdge MultiPack Storage Guide*).

Installation Bugs

4336171 – During initial cluster software installation, the `scinstall(1M)` command displays volume manager choices as follows:

- ```
1) Cluster Volume Manager (CVM)
2) Sun StorEdge Volume Manager (SSVM)
3) Solstice DiskSuite (SDS)
```

Sun Cluster 2.2 7/00 Release supports VERITAS Volume Manager 3.0.4, which includes the functionality formerly called Cluster Volume Manager. However, `scinstall` has not been updated to reflect the new product names. If you want to install VERITAS Volume Manager 3.0.4, select option 2. If you need the cluster functionality formerly known as Cluster Volume Manager (if your cluster includes Oracle Parallel Server, for example), select option 1.

**4359807** – When installing Netscape Messaging Server, the name you type for the server instance name should not start with the prefix `msg-`. The installation software automatically adds that prefix to the base name you specify. You should also specify that same base name as the data service instance name when you run the `hadsconfig(1M)` utility. The `hadsconfig` utility automatically adds the prefix `SUNWscnsm_` to the base name you specify. For example, if you specify the base name `my_mail`, the resulting server instance name would be `msg-my_mail`, and the resulting data service instance name would be `SUNWscnsm_my_mail`.

## Upgrade Bugs

**4218613** – During upgrade to Sun Cluster 2.2 from HA 1.3, instance configuration information for the HA-DBMS data services is not propagated to the new cluster. This prevents the database instances from starting when the new cluster is started. This bug affects the Sun Cluster HA for Oracle, Sun Cluster HA for Sybase, and Sun Cluster HA for Informix data services.

Work around the problem by manually recreating the database instance after completing the upgrade. Use the appropriate `hadbms insert` command (`haoracle insert`, `hasybase insert`, or `hainformix insert`) as described in the associated man pages, and in the appropriate data service chapters in the *Sun Cluster 2.2 Software Installation Guide*.

After you recreate the database instances, start the instances by using the appropriate `hadbms start` command.

**4218823** – During upgrade from HA 1.3 to Sun Cluster 2.2, only two of three required IP addresses are added to the `.rhosts` file on each node. The address lost is the highly available IP address for the private interconnects. Utilities such as `hadsconfig(1M)` will not work without this entry. The user must manually add the required entries to the `.rhosts` file. The procedure is documented in Chapter 3 of the *Sun Cluster 2.2 Software Installation Guide*.

**4327771** – When upgrading from Sun Cluster 2.2 on Solaris 2.6 to Sun Cluster 2.2 7/00 on Solaris 8, the `SUNWdidx` package is not installed. This occurs only when Solaris 8 is booted in 64-bit mode. This causes initialization of disk IDs to fail, leaving the upgrade incomplete. Work around the problem by installing the `SUNWdidx` package manually, after installing the upgraded Solaris and Sun Cluster packages. Then re-initialize disk IDs, using the `scdidadm(1M)` command, as documented in Chapter 4 of the *Sun Cluster 2.2 Software Installation Guide*.

# Administrative Command Bugs

**4204883** – The `confccdssa(1M)` command will fail when you select the controller that contains the boot disk, and will display the misleading error message: “First RE may not be NULL. WARNING: All disks on this SSA (ctrl: *mm*) are either already in disk groups, have already been selected as one of the devices for the shared CCD or are otherwise unavailable.” To prevent this problem, do not select the controller that contains the boot disk.

**4235744** – The `scconf clustername -F logicalhost` command creates the primary and mirror of the HA administrative volume `dg-stat` on two different disks in the same storage device. If that storage device fails, or connection to that storage device is lost, automatic volume recovery is not possible. You must manually fix the volume and restart the volume.

To diagnose and correct this problem, perform the following steps.

1. Check whether your existing administrative file system is created with the mirrors on the same controller. If not, then no further action is needed.

If the administrative file system volumes are mirrored on disks on same controller, proceed with the following steps to rebuild the administrative file system so that the volumes are correctly mirrored *across* controllers.

2. Back up any data that is in the administrative file system (`/logicalhost`) directory.
3. Put the logical host in maintenance mode.
4. Using VERITAS Volume Manager commands, manually import the disk group to where the administrative file system resides, remove the `dg-stat` volume, then create the volume using the same name `dg-stat`, specifying a mirror layout across controllers.
5. Recreate the administrative file system.

```
scconf clustername -F logicalhost
```

The command will find that an administrative file system volume (`dg-stat`) already exists, and will use that volume to create the administrative file system.

6. Unmount the newly-created file system.
7. Deport the disk group.
8. Bring up the logical host by using the `haswitch(1M)` command.
9. Restore any data to the `/logicalhost` directory.

**4240225** – A `umount` operation will fail during a switchover if the `df` command is run before the partition is unmounted. This causes the cluster to attempt to re-master the logical host on the original node, which fails, leaving the logical host in a partially mastered state. The error message produced in this situation is cryptic: “ID[SUNWcluster.scnfs.4010]: unmount /mail/spool failed.” To work around the problem, switch the logical host into maintenance mode by using `haswitch` or `scconf(1M)`, and then re-master the logical host correctly, using the `scconf` command. See the `scconf(1M)` man page for details.

## Data Service Bugs

**4262913** – The Sun Cluster HA for Oracle `ksh` script in `/opt/SUNWcluster/bin` has problems if it encounters a non-Sun-supplied data service with the string “oracle” in its name. Therefore, do not include the string “oracle” in any names of data services you create using the Sun Cluster Application Programming Interface (API).

**4336343** – Inclusion of “child-level monitoring” features in `pmfd` for Sun Cluster 2.2. See “[Features](#)” on [page 5](#) for more information.

**4338556** – The Sun Cluster HA for NetBackup activity monitor does not display the correct status after switchover. This means that you cannot detect the status of backups after a switchover. No workaround exists for this problem currently. See your service provider for the latest status.

**4345031** – If a switchover or failover from a NetBackup client cluster occurs during a restore operation, the restore process continues to write to the root disk and might fill up that disk, thus stalling the switchover of the cluster. Simultaneously, the NetBackup Progress Report utility does not report the correct status of the restore operation.

To correct this situation, include the following parameter in the `bp.conf` file for the NetBackup client cluster:

```
REQUIRED_INTERFACE=logicalhostname
```

For example:

```
REQUIRED_INTERFACE=lh-schost-1
```

This parameter ensures that the restore process (a tar operation) stops writing to the root disk via the shared disk mount point after a switchover or failover. However, the tar process may persist for a while; you can just delete it and reexecute the restore.

**4387527** – When installing the ha-oracle agent in Sun Cluster 2.2, root does not need to be listed as a member of the database administrator group in the `/etc/group` file as previously documented. The entry can now be

```
dba:*:520:oracle
```

For more information on installing and configuring Sun Cluster HA for Oracle, see the *Sun Cluster 2.2 Software Installation Guide*.

**4405556** – Missing information on installing Sun Cluster HA for Oracle on multihost disks. The following note should be included in Chapter 5 of the *Sun Cluster 2.2 Software Installation Guide*.

---

**Note** – If you install the Oracle binaries on a multihost disk, you must install the `SQL*PLUS` option from Oracle on the local nodes as well. The Sun Cluster HA for Oracle fault monitor only works correctly when you install the `SQL*PLUS` option on the local nodes.

---

## Sun Cluster Manager Bugs

**Running SCM with the HotJava browser** – If you choose to use the HotJava browser shipped with your Solaris 2.6 or Solaris 7 operating environment to run SCM, there may be problems such as:

- Menu errors. For example, after making your menu selection, the menu selection might remain visible on the browser.
- Use of swap space. If you choose to use the HotJava browser with SCM, you should have at least 40 Megabytes of free swap space. If you find that swap space gets low, restarting the HotJava browser might help.
- Cannot access online help. When running HotJava on a cluster node and displaying it remotely, online help may freeze when attempting to access online help.

**4221612** – Sun Cluster Manager sometimes incorrectly reports that the Sun Cluster HA for Netscape HTTP data service is down when it is actually up. Work around the problem by checking the status of the data service with `hareg(1M)` or `hastat(1M)` instead. See the `hareg(1M)` and `hastat(1M)` man pages for details.

**4312093** – On Solaris 8, you cannot run Sun Cluster Manager as an applet with the Netscape browser and Java Development Toolkit (JDK) version 1.2. Instead, you can either run Sun Cluster Manager as a standalone application, or change the default JDK to version 1.1 (if your cluster is not running any applications that depend upon JDK 1.2).

To run Sun Cluster Manager as a standalone application, follow the detailed instructions in Chapter 2 of the *Sun Cluster 2.2 System Administration Guide*.

If your cluster is not hosting any applications that depend upon JDK 1.2, you can choose to change the JDK default to version 1.1. Do this by modifying the `JAVA_HOME` fields in the Sun Cluster Manager start-up script to specify version 1.1:

```
cd /opt/SUNWcluster/scmgr/lib
vi scm_server_start
...
JAVA_HOME=/usr/java1.1
PATH=${JAVA_HOME}/bin:/bin:/etc:/sbin:/usr/sbin
...
```

**4316289** – Sun Cluster Manager, when invoked from the command line, does not display data services associated with a logical host. Work around the problem by obtaining data service information from the Properties > Registered HA Services menu of the Sun Cluster Manager GUI.

**4332639** – Sun Cluster Manager calls the HotJava browser by default, but HotJava is not supported on Solaris 8. Work around the problem on Solaris 8 configurations by specifying the Netscape browser. Use the following command:

```
/opt/SUNWcluster/bin/scmgr -b /usr/dt/bin/netscape
```

**4333246** – Sun Cluster Manager displays `qfe` private network connections as Unknown. Work around the problem by refreshing the cluster view in the Sun Cluster Manager GUI, using the menus Help > Cluster > Refresh Current View.

## Documentation Errata

**4233113** – Sun Cluster documentation omits information regarding logical host timeout values and how they are used. When you configure the cluster, you set a timeout value for the logical host. This timeout value is used by the CCD when you bring a data service up or down using the `hareg(1M)` command. The CCD

operation occurs in two steps; half of the timeout value is used for each step. Therefore, when configuring *START* and *STOP* methods for data services, make sure each method uses no more than half of the timeout value set for the logical host.

**4330501** – The *Sun Cluster 2.2 System Administration Guide*, section 4.4, “Disabling Automatic Switchover” indicates that you can disable automatic switchover of logical hosts by using the `scconf -m` command. This is misleading. You can use `scconf -m` to disable automatic switchover of logical hosts *only* if you issue the command when you create the logical hosts initially.

If the logical host already exists, you must remove the logical host and then re-create it using `scconf -m`, in order to disable automatic switchover.

**4336091** – Sun Cluster documentation omits information regarding how to set logical unit numbers (LUNs) for A1000 and A3x00 storage devices.

When you add A1000 and A3x00s to a Sun Cluster configuration, you must set the LUNs so that they survive switchover or failover of a cluster node, without loss of pseudo-device information. Use the following procedure to ensure that LUNs are set correctly and permanently for these disk types.

1. On both nodes, install or verify the existence of the RAID manager packages, `SUNWosafw`, `SUNWosamn`, `SUNWosar`, and `SUNWosau`.
2. (Solaris 8 only) Install or verify the existence of the RAID manager patch 108553.

Obtain the patch from your service provider or from the patch web site <http://sunsolve.sun.com>.

3. Use the `RM6` tool to set up the LUNs on the first node.

Using the tool’s GUI, click on “Configuration,” then on “Module Name,” and then on “Create LUN icon.”

4. Compare the `/etc/osa/rdac_address` files on both nodes.

In Step 3, LUNs were assigned to either controller A or B, and the `rdac_address` file records this assignment. If necessary, modify the `rdac_address` file on the second node so that the controller assignments match those on the first node.

Run the following RAID manager command on both nodes.

```
/usr/lib/osa/bin/hot_add
```

**4341222** – Chapter 1, sections 1.3.2 and 1.5.6 of the *Sun Cluster 2.2 Software Installation Guide* do not accurately describe the behavior of CCD quorum during cluster configuration. The documentation should reflect that it is possible to modify the cluster configuration database even when CCD quorum conditions are not met (that is, when greater than half of all cluster nodes do not have a valid CCD).

Typically, the cluster software requires a quorum before updating the CCD. This requirement is highly restrictive in configurations using logical hosts. Therefore, to overcome this limitation in Sun Cluster 2.2, all administrative and configuration commands related to logical hosts and data services that update the CCD database can be executed without CCD quorum. Such commands include `hareg(1M)` and `scconf(1M)` operations, for example.

To prevent loss of any CCD updates, you should always make sure that the last node to leave the cluster during cluster shutdown is the first node to rejoin the cluster upon start up.

**4342066** – The procedure “How to Change the Name of a Cluster Node” in section 3.2 of the *Sun Cluster 2.2 System Administration Guide* is incorrect and should not be used. The correct procedure involves changing various framework files which should not be altered manually by anyone other than your service representative. If you need to change the name of a cluster node, contact your service provider for assistance.

**4342236** – The `abort_net` method is described incorrectly in the *Sun Cluster 2.2 API Developer's Guide*. The documentation states that the `abort_net` method can be used to execute “last wishes” cleanup code before a cluster is stopped. This is incorrect.

Instead, the `abort_net` method is called by the `clustd` daemon when a node is about to abort from the cluster, typically in a split-brain situation when the node in question is the loser in the race for the quorum device (see Chapter 1 in the *Sun Cluster 2.2 Software Installation Guide* for more information about quorum devices). In such a case, first `abort_net` is called, then the network is taken down, and finally `abort` is called. However, these methods are executed on the node that is aborting *only* if the node owns the data service associated with the methods. (That is, if the aborting node does not own any logical host, then it will not execute any of the `abort` methods associated with a data service.) The aborting node will stop the cluster software, but the node itself will remain up.

Note that `stop` and `stop_net` methods are called each time the cluster reconfigures itself (due to nodes joining or leaving the cluster), as part of normal cluster operation.

**4343021** – In Chapter 14 in the *Sun Cluster 2.2 Software Installation Guide*, the documented installation procedure for Sun Cluster HA for NetBackup is incorrect. The correct procedure is as follows.

1. **Install Sun Cluster 2.2 7/00 Release using the procedures documented in Chapter 3 of the *Sun Cluster 2.2 Software Installation Guide*.**
2. **Stop the cluster by running the following command on all nodes, sequentially.**

```
scadmin stopnode
```

3. On all nodes, install VERITAS NetBackup, using the procedures documented in Chapter 14 of the *Sun Cluster 2.2 Software Installation Guide*.
4. On all nodes, install Sun Cluster patch 109214, which enhances the `scinstall(1M)` command to recognize Sun Cluster HA for NetBackup. The patch is available from your service provider or from the Sun patch web site <http://sunsolve.sun.com>.
5. On all nodes, re-run the `scinstall` command and install the Sun Cluster HA for NetBackup data service.

```
scinstall
```

6. On all nodes, install Sun Cluster patches 108450 and 108423, which enhance the `hadsconfig(1M)` command to recognize Sun Cluster HA for NetBackup.
7. Start the cluster. On the first node, run the following command:

```
scadmin startcluster
```

Sequentially, on all other nodes, run the following command:

```
scadmin startnode
```

8. Register the data service by running the following command on one node only.

```
hareg -s -r netbackup
```

9. Configure the Sun Cluster HA for NetBackup data service by running the `hadsconfig` command on one node only. See Chapter 14 in the *Sun Cluster 2.2 Software Installation Guide* for configuration parameters to supply to `hadsconfig`.

```
hadsconfig
```

10. Activate the data service by running the following command on one node only.

```
hareg -y netbackup
```

**4343093** – Chapter 2, section 2.6.5, in the *Sun Cluster 2.2 Software Installation Guide* states that Sun Cluster 2.2 must run in `C` locale. This is incorrect. Sun Cluster 2.2 7/00 Release can run in `C`, `fr` (French), `ko` (Korean), and `ja` (Japanese) locales.

**4344711** – Appendix C in the *Sun Cluster 2.2 Software Installation Guide* contains incorrect or incomplete information about configuring VERITAS Volume Manager. These errors are described in more detail below.

The document mentions only VxFS file systems, and omits information about UFS file systems. In Sun Cluster configurations, UFS file systems can be created in a similar fashion to VxFS file systems. See your system administration documentation for more information about creating and administering UFS file systems.

When using the `mkfs(1M)` command to administer VxFS file systems, use the fully qualified path to the command, such as `/usr/lib/fs/vxfs/mkfs`. The documentation omits this information wherever the `mkfs` command is described.

In section C.3, “Configuring VxFS File Systems on the Multihost Disks,” the procedure contains erroneous steps. The correct procedure follows. Use this procedure after creating logical hosts as described in the `scconf(1M)` man page or in Chapter 3, “Installing and Configuring Sun Cluster Software.”

- 1. Take ownership of the disk group containing the volume by using the `vxvg(1M)` command to import the disk group to the active node.**

```
phys-hahost1# vxvg import diskgroup
```

- 2. Run the following `scconf(1M)` command on each cluster node.**

This `scconf` command will create a volume for the administrative file system, create a file system within that volume, create mount points for that volume in the root file system (“/”), create `dfstab.logicalhost` and `vfstab.logicalhost` files in `/etc/opt/SUNWcluster/conf/hanfs`, and create an appropriate entry in the `vfstab.logicalhost` file for the administrative file system.

```
phys-hahost1# scconf clustername -F logicalhost
```

- 3. Create file systems for all volumes. These volumes will be mounted by the logical hosts.**

```
phys-hahost1# mkfs -F vxfs /dev/vx/rdisk/diskgroup/volume
```

- 4. Update the `vfstab.logicalhost` file to include entries for the file systems created in Step 3.**

## 5. Create mount points for the file systems created in Step 3.

```
phys-hahost1# mkdir /logicalhost/volume
```

## 6. Import the disk groups to their default masters.

It is most convenient to create and populate disk groups from the active node that is the default master of the particular disk group.

Import each disk group onto the default master node using the `-t` option. The `-t` option is important, as it prevents the import from persisting across the next boot.

```
phys-hahost1# vxvg -t import diskgroup
```

## 7. (Optional) To make file systems NFS-sharable, refer to Chapter 11, “Installing and Configuring Sun Cluster HA for NFS.”

**4345750** – In Chapter 14 of the *Sun Cluster 2.2 System Administration Guide*, the procedure “How to Replace a Sun StorEdge A5000 Disk (VxVM),” Steps 3 and 4 are not valid for a cluster that runs on the Solaris 7 11/99 operating environment and later. In Step 3, you should run the `luxadm remove_device` command on only one of the nodes connected to the array. Performing the command on additional nodes is unnecessary and will generate error messages. In Step 4, after you physically replace the disk, do not run the `luxadm insert_device` command. This command is not necessary.

If your cluster runs on a Solaris operating environment earlier than the Solaris 7 11/99 release, Steps 3 and 4 are still valid as documented.

**4356674** – In Chapter 14 of the *Sun Cluster 2.2 System Administration Guide*, the procedure “How to Replace a Sun StorEdge A5000 Disk (Solstice DiskSuite)” contains errors in Steps 2, 3, 11, 12, and 13. In these steps, the directory `/tmp` should be replaced with `/var/tmp`, and physical device names should be replaced with `did` device names. The corrected procedure, in its entirety, is as follows.

### 1. Identify all metadevices or applications that use the failing disk.

If the metadevices are mirrored or RAID 5, the disk can be replaced without stopping the metadevices. Otherwise all I/O to the disk must be stopped using the appropriate commands. For example, use the `umount(1M)` command to unmount a file system on a stripe or concatenation.

### 2. Preserve the disk label, if necessary. For example:

```
prtvtoc /dev/rdisk/c1t3d0s2 > /var/tmp/c1t3d0.vtoc
```

3. (Optional) Use `metareplace` to replace the disk slices if the disk has not been hot-spared. For example:

```
metareplace d1 /dev/did/dsk/d23 /dev/did/dsk/d88
d1: device d23 is replaced with d88
```

4. Use `luxadm -F` to remove the disk.

The `-F` option is required because Solstice DiskSuite does not offline disks. Repeat the command for all hosts, if the disk is multihosted. For example:

```
luxadm remove -F /dev/rdsk/clt3d0s2
WARNING!!! Please ensure that no filesystems are mounted on these
device(s). All data on these devices should have been backed
up. The list of devices which will be removed is:
1: Box Name "macs1" rear slot 1
Please enter 'q' to Quit or <Return> to Continue: stopping: Drive
in "macs1" rear slot 1...Done
offlining: Drive in "macs1" rear slot 1...Done
Hit <Return> after removing the device(s).
```

---

**Note** – The FPM icon for the disk drive to be removed should be blinking. The amber LED under the disk drive should also be blinking.

---

5. Remove the disk drive and enter `Return`. The output should look similar to the following:

```
Hit <Return> after removing the device(s).
Drive in Box Name "macs1" rear slot 1
Removing Logical Nodes:
Removing clt3d0s0 Removing clt3d0s1 Removing clt3d0s2 Removing
clt3d0s3 Removing clt3d0s4 Removing clt3d0s5 Removing clt3d0s6
Removing clt3d0s7 Removing c2t3d0s0 Removing c2t3d0s1 Removing
c2t3d0s2 Removing c2t3d0s3 Removing c2t3d0s4 Removing c2t3d0s5
Removing c2t3d0s6 Removing c2t3d0s7
#
```

6. Repeat Step 4 for all nodes, if the disk array is in a multi-host configuration.

7. Use the `luxadm insert` command to insert the new disk. Repeat for all nodes. The output should be similar to the following:

```
luxadm insert macs1,r1
The list of devices which will be inserted is:
1: Box Name "macs1" rear slot 1
Please enter 'q' to Quit or <Return> to Continue: Hit <Return>
after inserting the device(s).
```

8. Insert the disk drive and enter Return. The output should be similar to the following:

```
Hit <Return> after inserting the device(s). Drive in Box Name
"macs1" rear slot 1 Logical Nodes under /dev/dsk and /dev/rdisk:
c1t3d0s0 c1t3d0s1 c1t3d0s2 c1t3d0s3 c1t3d0s4 c1t3d0s5 c1t3d0s6
c1t3d0s7 c2t3d0s0 c2t3d0s1 c2t3d0s2 c2t3d0s3 c2t3d0s4 c2t3d0s5
c2t3d0s6 c2t3d0s7
#
```

---

**Note** – The FPM icon for the disk drive you replaced should be lit. In addition, the green LED under the disk drive should be blinking.

---

9. On all nodes connected to the disk, use `scdidadm(1M)` to update the DID pseudo device information.

In this command, *DID\_instance* is the instance number of the disk that was replaced. Refer to the `scdidadm(1M)` man page for more information.

```
scdidadm -R DID_instance
```

10. Reboot all nodes connected to the new disk.

To avoid down time, use the `haswitch(1M)` command to switch ownership of all logical hosts that can be mastered by the node to be rebooted. For example,

```
haswitch phys-hahost2 hahost1 hahost2
```

11. Label the disk, if necessary. For example:

```
cat /var/tmp/c1t3d0.vtoc | fmthard -s - /dev/rdisk/c1t3d0s2
fmthard: New volume table of contents now in place.
```

**12. Replace the metadb, if necessary. For example:**

```
metadb -s setname -d /dev/did/rdisk/d23s7;
metadb -s setname -a -c 3 /dev/did/rdisk/d23s7
```

**13. Enable the new disk slices with metareplace -e. For example:**

```
metareplace -e d1 /dev/did/rdisk/d23s0
d1: device d23s0 is enabled
```

This completes the disk replacement procedure.

**4448815** – In the `cports(1M)` man page, there is a typo in a file name. The man page currently says: “If an entry for “serialports” has been made in the `/etc/nisswitch.conf` file, then the order of lookups is ...” The correct file name is `/etc/nsswitch.conf`.

**4448860** – In the `chosts(1)` man page, there is a typo in a file name. The man page currently says: “If an entry for “clusters” has been made in the `/etc/nisswitch.conf` file, then the order of lookups is ...” The correct file name is `/etc/nsswitch.conf`.

## Other Known Issues

### Oracle Parallel Server 8.1.6 UDLM Requirements

To run Oracle Parallel Server 8.1.6 with Sun Cluster 2.2 7/00 Release on Solaris 8, you must download an Oracle patch that provides fixes to the UNIX Dynamic Lock Manager (UDLM), version 3.3.4.4, allowing it to recognize and install on Solaris 8. The patch is available from your Oracle or Sun service provider.

### Failover/Switchover When Logical Host File System Is Busy

If a failover or switchover occurs while a logical host’s file system is busy, the logical host fails over only partially; part of the disk group remains on the original target physical host. Do *not* attempt a switchover if a logical host’s file system is busy.

## Displaying LOG\_DB\_WARNING Messages for the SAP Probe

The Sun Cluster HA for SAP parameter LOG\_DB\_WARNING determines whether warning messages should be displayed if the Sun Cluster HA for SAP probe cannot connect to the database. When LOG\_DB\_WARNING is set to `-y` and the probe cannot connect to the database, a message is logged at the warning level in the `local0` facility. By default, the `syslogd(1M)` daemon does not display these messages to `/dev/console` or to `/var/adm/messages`. To see these warnings, you must modify the `/etc/syslog.conf` file to display messages of `local0.warning` priority. For example:

```
...
*.err;kern.notice;auth.notice;local0.warning /dev/console
*.err;kern.debug;daemon.notice;mail.crit;local0.warning /var/adm/messages
...
```

After modifying the file, you must restart `syslogd`. See the `syslog.conf(1M)` and `syslogd(1M)` man pages for more information.

---

# Undocumented Error Messages

The following error messages may be generated by Sun Cluster 2.2 7/00 Release, but are not included in the *Sun Cluster 2.2 Error Messages Manual*.

## Sun Cluster HA for SAP Error Messages

The following error messages for Sun Cluster HA for SAP were omitted from the *Sun Cluster 2.2 Error Messages Manual*.

```
SUNWcluster.ha.sap.stop_net.2076: proha:SUNWscsap_PRO: Found 2
leftover IPC objects for SAP instance, removing via cleanipc
```

This message indicates that during shutdown of the SAP central instance by the `stop_net` method, two IPC segments from the central instance were found. The `stop_net` code uses the SAP-supplied utility `cleanipc` to remove all IPC segments of the central instance during shutdown (and also before startup). This is to ensure a thorough shutdown as well as a clean startup. The error message is an informational message only, and is expected. No user action is required.

```
Graceful shutdown failed for oracle instance PRO, starting abort
```

This message indicates that the HA-Oracle `oracle_db_shutdown` script did not complete a graceful shutdown of the database within the timeout limit (30 seconds, by default). If the normal shutdown does not complete during the allowed time, then a shutdown abort is issued. This is an informational message and no user action is required.

```
SUNWcluster.ccd.ccdctl.4403: (error) checkpoint, ccdd, ticlts:
RPC: Program not registered
```

This message indicates that the `ccdadm` command could not contact the `ccdd` demon for the requested operation—the RPC call `clnt_create()` failed. Verify that the cluster has been started on the current node, and the `ccdd` daemon is running.

```
SUNWcluster.clustd.transition.4010: cluster aborted on this node
nodename
```

This message indicates that the current node is being aborted. Other error messages should indicate why this is occurring; check the `scadmin.log` log file in `/var/opt/SUNWcluster`.

```
reconf.pnm.3009: pnminit faced problems
```

This message is generated by the script `/opt/SUNWcluster/bin/pnm`. This script is called during step 1 of cluster reconfiguration, when PNM is initialized with `pnminit`. The error message appears if the execution of `pnminit` resulted in a non-zero exit. Reasons for a non-zero exit of `pnminit` include:

- Invalid command line arguments
- Environment variables not set (`localnodeid`, `clustname`, `currnodes`, `numnodes`)
- `pnminit` could not communicate with the `ccdd` daemon correctly

Check for any error messages logged to `/var/opt/SUNWcluster/ccd/ccd.log`, then restart the cluster reconfiguration.

```
SUNWcluster.reconfig.4018: Aborting--received abort request from
nodename
```

This message indicates a request from a remote node to abort the current node. Use checksum to verify that the `/etc/opt/SUNWcluster/conf/clustname.cdb` files are identical on all nodes. If necessary, manually copy the most recent `clustname.cdb` file to all nodes, and then restart the cluster.

## monitor\_rpcbind Error Messages

The following error messages potentially produced by `monitor_rpcbind` were omitted from the *Sun Cluster 2.2 Error Messages Manual*.

```
SUNWcluster.monitor_rpcbind.1001: Invalid daemon:
```

This message indicates that the daemon name is set incorrectly. To remedy this, contact your Sun representative.

```
SUNWcluster.monitor_rpcbind.3001: Failed to restart rpcbind -w.
Aborting this node.
```

This message indicates that `rpcbind` is not running on this node, and the system attempted unsuccessfully to restart it. The system will be aborted automatically.

```
SUNWcluster.monitor_rpcbind.4502: rpcbind is not running -- manual
reboot may be needed
```

This message indicates that `rpcbind` is not running on this node and could not be restarted automatically by the system. The system will be aborted automatically.

```
SUNWcluster.monitor_rpcbind.5001: rpcbind is not running but warm
restart seems to be possible. Will attempt to restart.
```

This message indicates that `rpcbind` is not running on this node and an attempt will be made by the system to restart it. This is an informational message only; no user action is necessary.

```
SUNWcluster.monitor_rpcbind.5002: rpcinfo failed - no rpcbind.
```

This message indicates that the test for an active `rpcbind` failed, for whatever reason is specified in the message. This is an informational message only; no user action is necessary.

```
SUNWcluster.monitor_rpcbind.5003: rpcbind in process list but has not responded.
```

This message indicates that although `rpcbind` appears in the process table for the system, it has failed to respond to the fault monitor in the required time. The fault probe will be retried automatically. This is an informational message only; no user action is required.

```
SUNWcluster.monitor_rpcbind.5010: rpcbind is not running on this node and cannot be restarted. This node will be aborted.
```

This message indicates that `rpcbind` is not running on this node and an unsuccessful attempt was made by the system to restart it. As a result this node will abort automatically. No user action is required.

```
SUNWcluster.monitor_rpcbind.5011: rpcbind is not running on this node and cannot be restarted. Selected action is to continue operation.
```

This message indicates that `rpcbind` is not running on this node and the system was unable to restart it. Because the fault monitor has been told not to abort the node, operation will continue. However, the Sun Cluster framework will not be able to reconfigure without operator intervention. Reboot the node manually to ensure correct operation.

```
SUNWcluster.monitor_rpcbind.6000: Restarted the daemon rpcbind,
pid= <pid>
```

This message indicates that rpcbind was not running on this node and was successfully restarted by Sun Cluster. No user action is required.

## Framework Error Messages

The following error messages are potentially produced by the Sun Cluster process monitor facility. These messages were omitted from the *Sun Cluster 2.2 Error Messages Manual*.

```
SUNWcluster.pmf.1030: failfast_open: running with failfast in
debug/disabled mode
```

This message indicates that the pmf daemon, pmfd, is running in debug mode. A non-responsive pmfd will not trigger a failfast panic while running in this mode. This is a notification message only. No action is required.

```
SUNWcluster.pmf.1031: pmfd_failfast_thread: re-armed in %lld ms,
was expecting %lld ms with variance of %lld ms
```

The rpc.pmf daemon registers with the failfast timer on startup, and then a reset thread is spawned to rearm the failfast timeout continuously. This warning message is printed when this reset thread is scheduled past the expected time plus some padding or variance time. The variance is set at 10% of rearm time initially (5.5 seconds), and then is incremented to twice the rearm time (10 seconds). This only affects the rate at which messages are printed, not the rearm time or the timeout. This warning message indicates an excessive workload on this node, which in turn is causing a delay in the scheduling of the pmfd failfast reset thread. Further delay of this thread could result in a failfast timeout.

```
in.rdiscd[517]: setsockopt (IP_DROP_MEMBERSHIP): Cannot assign
requested address
```

This error message might be displayed when you stop a cluster node. The error is caused by a timing issue between the `in.rdiscd` daemon and the IP module. It is harmless and can be ignored safely.

```
WARNING: lockd: cannot contact statd (error 4), continuing.
```

On clusters using Sun Cluster HA for NFS on Solaris 7, this error message is displayed if the `lockd` daemon is killed before the `statd` daemon is fully running. This error message can be ignored safely.

---

# Future Changes

This section describes Sun Cluster features that might be changed or discontinued after Sun Cluster 2.2.

## Sun Cluster 2.2 Commands To Be Replaced or Made Obsolete

The following commands will be changed or discontinued after Sun Cluster 2.2, as noted.

### Commands with options or interfaces to be changed:

- `scconf(1M)`
- `scinstall(1M)`
- `scswitch(1M)`

### Commands to be renamed:

- `hastat(1M)` will be renamed `scstat(1M)`
- `haswitch(1M)` will be renamed `scswitch(1M)`

### Commands to be removed:

- `ccdadm(1M)`
- `ccdctl(1M)`
- `ccdinstall(1M)`
- `cconsole(1)`
- `ccp(1)`
- `chosts(1)`
- `confccdssa(1M)`
- `cports(1)`
- `crlogin(1)`
- `ctelnet(1)`
- `db_check(1M)`
- `finddevices(1M)`
- `get_ci_status(1M)`
- `get_node_status(1M)`
- `ha_fault_mon_config(1M)`
- `hadsconfig(1M)`
- `hainformix(1M)`

- hainformix\_support(4)
- haoracle(1M)
- haoracle\_support(4)
- hareg(1M)
- hasybase(1M)
- hasybase\_support(4)
- scadmin(1M)
- sm\_config(1M)
- sma\_configd(1M)

## API Commands or Command Options To Be Replaced or Made Obsolete

The following commands and command options might not be available in future Sun Cluster releases.

- haget(1M) query on names\_on\_subnets

```
haget -f names_on_subnets -h physical_host
```

- haget query on vfstab\_file

```
haget -f vfstab_file -h logical_hostname
```

- -r reconfiguration option to hactl(1M)

```
hactl -r ...
```

- The cluster reconfiguration process described in the hareg(1M) man page might change in future releases. The reconfiguration process might no longer be a lock step process. You should not develop highly available data services that depend on the steps occurring in lock step.

# Internal Programs To Be Retired in Future Releases

The Sun Cluster implementation contains many programs that are used internally by the implementation and are not intended for use by customers. Any program that does not have a man page in Sun Cluster 2.2 7/00 Release falls into this category. These programs will not exist in their current form in subsequent releases of the product. Some examples include `clustm`, `scccd`, and `ccdmatch`.

---

## Notes and Issues for Localized Versions

This section describes installation requirements, patches, and issues applicable when installing localized versions of the Sun Cluster 2.2 7/00 Release.

### Supported Locales

Sun Cluster 2.2 7/00 Release supports the following locales.

| Locale | Language |
|--------|----------|
| ja     | Japanese |
| fr     | French   |
| ko     | Korean   |

### Locale Installation Overview

Follow these general guidelines to install localized Sun Cluster 2.2 7/00 Release. See the *Sun Cluster 2.2 Software Installation Guide* for detailed installation and configuration instructions.

- 1. Install the Solaris operating environment, using your Solaris documentation.**

Sun Cluster 2.2 requires the Solaris 2.6, Solaris 7, or Solaris 8 operating environment. Install the Solaris software and all required and recommended Solaris patches. Solaris patches are available from your service provider or from the patch web site, <http://sunsolve.sun.com>.

2. (Optional) Install the Sun Cluster 2.2 7/00 Release AnswerBooks onto an existing AnswerBook2 server or non-cluster node, using the instructions in “AnswerBooks” on page 20.
3. (Optional) View or print the PostScript or derived HTML (DHTML) versions of the Sun Cluster 2.2 documentation from the Sun Cluster 2.2 CD.
4. Install the English version of Sun Cluster 2.2, using the `scinstall(1M)` script found on the Sun Cluster 2.2 CD.

```
cd Sun_Cluster_2_2/Sol_2.x/Tools
./scinstall
```

5. Install the Sun Cluster 2.2 patches, which are included on the Sun Cluster 2.2 product CD.

Also check with your service provider or the Sun patch web site, <http://sunsolve.sun.com>, for any additional patches.

```
cd Sun_Cluster_2_2/Sol_2.x/Patches
./install_scpatches
```

6. Install the localized version of the Sun Cluster 2.2 software, using the `pkgadd(1M)` utility and the appropriate package name from the table below.

```
cd Sun_Cluster_2_2/Sol_2.x/Product
pkgadd -d . package_name
```

The locale package names for Sun Cluster 2.2 7/00 are as follows:

| Package Name | Locale | Description                                   |
|--------------|--------|-----------------------------------------------|
| SUNWfrscc    | fr     | French localized Sun Cluster client           |
| SUNWfrscs    | fr     | French localized Sun Cluster server           |
| SUNWjecmc    | ja     | Japanese localized Sun Cluster client         |
| SUNWjescc    | ja     | Japanese man pages for Sun Cluster client     |
| SUNWjescs    | ja     | Japanese localized Sun Cluster server         |
| SUNWjecms    | ja     | Japanese man pages for Sun Cluster server     |
| SUNWjemdm    | ja     | Japanese Solstice DiskSuite Mediator software |
| SUNWkscc     | ko     | Korean localized Sun Cluster client           |
| SUNWkscs     | ko     | Korean localized Sun Cluster server           |

## Setting the System Default Locale

You must set the system default locale correctly on all cluster nodes in order to see localized messages. Set the system default locale by editing the `/etc/default/init` files on all cluster nodes so that the `LANG` or `LC_MESSAGES` fields are set to the appropriate locale (`fr`, `ko`, or `ja`). For example, to see French messages, specify `LANG=fr` or `LC_MESSAGES=fr`. The system default locale settings must be identical on all nodes.

## Dependencies on Packages for Localized Error Messages

Localized Solaris packages are required on all cluster nodes to enable localized messages. These packages are installed by default during installation of the localized Solaris operating environment. These packages are `SUNWfros` for French messages, `SUNWjeuc` for Japanese messages, and `SUNWkleu` for Korean messages.

## Locale-Specific Known Problems

**4299187** – The cluster console does not accept non-ascii characters, for example, Japanese characters or French (accented) characters. Work around the problem by inputting such characters through the individual terminal windows on each cluster node, instead of through the cluster console.

**4277778** – The Help messages about the Cluster Control Panel (CCP) are displayed in English, regardless of the locale in which Sun Cluster is running. CCP Help messages are translated to Japanese, however, and can be accessed through a browser, with the following URL:

```
file:/opt/SUNWcluster/helpfiles/ja/sc/home_page
```

# Installing, Configuring, and Maintaining a Sun StorEdge T3 Disk Array

---

This appendix the procedures for installing, configuring, and maintaining a Sun StorEdge™ T3 disk array.

This appendix contains the following procedures:

- [“How to Install a Sun StorEdge T3 Disk Array” on page 52](#)
- [“How to Create a Sun StorEdge T3 Disk Tray Logical Volume” on page 55](#)
- [“How to Remove a Sun StorEdge T3 Disk Tray Logical Volume” on page 57](#)
- [“How to Upgrade Sun StorEdge T3 Disk Array Firmware” on page 61](#)
- [“How to Replace a Disk Drive” on page 62](#)
- [“How to Add a Sun StorEdge T3 Disk Array” on page 62](#)
- [“How to Remove a Sun StorEdge T3 Disk Array” on page 70](#)
- [“How to Replace a Host-to-Hub Component” on page 72](#)
- [“How to Replace a Sun StorEdge FC-100 Hub and Hub-to-Disk Tray Component” on page 73](#)
- [“How to Replace a Sun StorEdge T3 Disk Array Controller” on page 74](#)
- [“How to Replace a Sun StorEdge T3 Disk Array Chassis” on page 74](#)
- [“How to Replace a Host Adapter” on page 75](#)

For conceptual information on multihost disks, see the *Sun Cluster 2.2 Software Installation Guide*.

---

# Installing a Sun StorEdge T3 Disk Array

This section provides the procedure for an initial installation of a Sun StorEdge T3 disk array.

## ▼ How to Install a Sun StorEdge T3 Disk Array

Use this procedure for an initial installation and configuration of a Sun StorEdge T3 disk array. Perform the steps in this procedure in conjunction with the procedures in the *Sun Cluster 2.2 Software Installation Guide* and your server hardware manual.

**1. Install the host adapters in the nodes that will be connected to the Sun StorEdge T3 disk arrays.**

For the procedure on installing host adapters, see the documentation that shipped with your host adapters and nodes.

**2. Install the Sun StorEdge FC-100 hubs.**

For the procedure on installing a Sun StorEdge FC-100 hub, see the *FC-100 Hub Installation and Service Manual*.

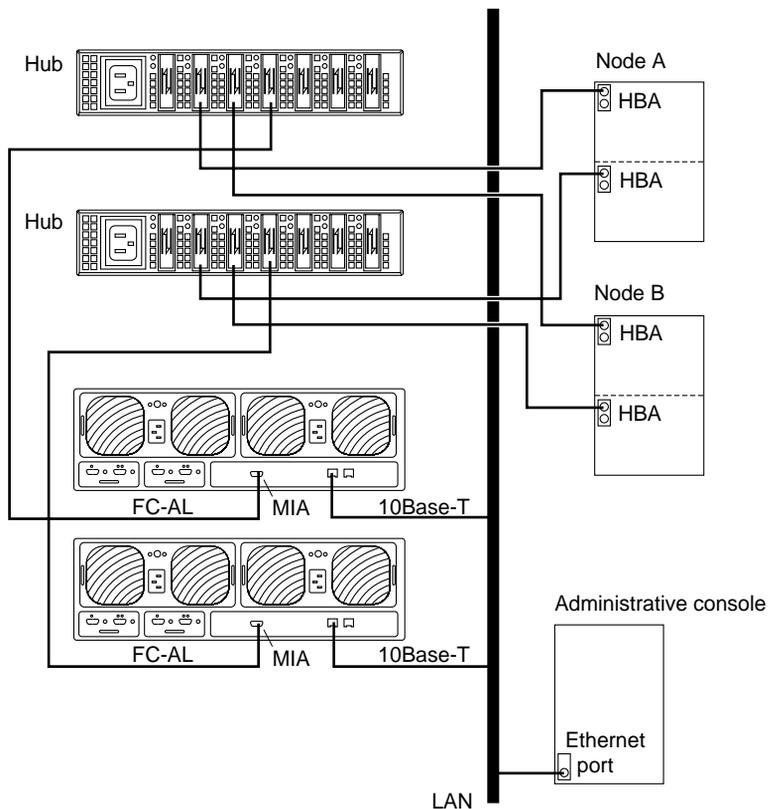
**3. Cable and power on the Sun StorEdge T3 disk arrays as shown in [FIGURE A-1](#).**

---

**Note** – There are no hub port assignment restrictions. You can connect your Sun StorEdge T3 disk array and node to any hub port.

---

For the procedure on installing fiber optic cables, see the *Sun StorEdge T3 Configuration Guide*. For the procedure on powering on the Sun StorEdge T3 disk array, see the *Sun StorEdge T3 Installation, Operation, and Service Manual*.



**FIGURE A-1** Cabling a Sun StorEdge T3 disk array in a single-tray configuration

**4. (Optional) Configure the Sun StorEdge T3 disk array with the desired Sun StorEdge T3 disk tray logical volumes.**

For the procedure on configuring the Sun StorEdge T3 disk array with Sun StorEdge T3 disk tray logical volumes, see the *Sun StorEdge T3 Disk Tray Administrator's Guide*.

**5. Install the Solaris operating environment, and apply the required Solaris patches for Sun Cluster software and Sun StorEdge T3 disk array support.**

For the procedure on installing the Solaris operating environment, see the *Sun Cluster 2.2 Software Installation Guide*. For the location of required Solaris patches and installation instructions for Sun Cluster software support, see the *Sun Cluster 2.2 Release Notes*. For a list of required Solaris patches for Sun StorEdge T3 disk array support, see the *Sun StorEdge T3 Disk Tray Release Notes*.

**6. Telnet to the Sun StorEdge T3 disk array you want to add, and install the required Sun StorEdge T3 disk array controller firmware.**

For the required Sun StorEdge T3 disk array controller firmware revision number, see the *Sun Cluster 2.2 Release Notes*. For the procedure on upgrading firmware, see the firmware patch README.

**7. Ensure that this new Sun StorEdge T3 disk array has a unique target address.**

For the procedure on verifying and assigning a target address, see the *Sun StorEdge T3 Configuration Guide*.

**8. Reset the Sun StorEdge T3 disk array.**

For the procedure on rebooting or resetting a Sun StorEdge T3 disk array, see the *Sun StorEdge T3 Installation, Operation, and Service Manual*.

## Where to Go From Here

To continue with Sun Cluster software installation tasks, see the *Sun Cluster 2.2 Software Installation Guide*.

---

# Configuring a Sun StorEdge T3 Disk Array

This section provides the procedures for configuring a Sun StorEdge T3 disk array in a running cluster. The following table lists these procedures.

**TABLE A-1** Task Map: Configuring a Sun StorEdge T3 Disk Array

| Task                                               | For Instructions, Go To...                                                                                                                                                     |
|----------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Create a Sun StorEdge T3 disk tray logical volume  | <a href="#">“How to Create a Sun StorEdge T3 Disk Tray Logical Volume” on page 55</a>                                                                                          |
| Replace a Sun StorEdge T3 disk tray logical volume | <a href="#">“How to Create a Sun StorEdge T3 Disk Tray Logical Volume” on page 55</a><br><a href="#">“How to Remove a Sun StorEdge T3 Disk Tray Logical Volume” on page 57</a> |
| Remove a Sun StorEdge T3 disk tray logical volume  | <a href="#">“How to Remove a Sun StorEdge T3 Disk Tray Logical Volume” on page 57</a>                                                                                          |

## ▼ How to Create a Sun StorEdge T3 Disk Tray Logical Volume

Use this procedure to create a Sun StorEdge T3 disk tray logical volume. This procedure assumes all cluster nodes are booted and attached to the Sun StorEdge T3 disk array that will host the Sun StorEdge T3 disk tray logical volume you want to create.

1. **Telnet to the Sun StorEdge T3 disk array that will host the Sun StorEdge T3 disk tray logical volume you want to create.**

## 2. Create the Sun StorEdge T3 disk tray logical volume.

Creating a Sun StorEdge T3 disk tray logical volume involves adding, mounting, and initializing a Sun StorEdge T3 disk tray logical volume.

For the procedure on creating and initializing a Sun StorEdge T3 disk tray logical volume, see the *Sun StorEdge T3 Disk Tray Administrator's Guide*. For the procedure on mounting a Sun StorEdge T3 disk tray logical volume, see the *Sun StorEdge T3 Installation, Operation, and Service Manual*.

## 3. On all cluster nodes, update the `/devices` and `/dev` entries.

Solaris 2.6:

```
drvconfig
disks -C
devlinks
```

Solaris 8:

```
devfsadm -C
```

After this process, a Solaris logical device name for the new Sun StorEdge T3 disk tray logical volume appears in the `/dev/rdisk` and `/dev/dsk` directories on all cluster nodes attached to the Sun StorEdge T3 disk array.

## 4. If needed, partition the Sun StorEdge T3 disk tray logical volume.

## 5. From any node in the cluster, update the global device namespace.

If a volume management daemon such as `vold` is running on your node, and you have a CD-ROM drive connected to the node, a `device busy` error might be returned even if no disk is in the drive. This error is expected behavior.

```
scgdevs
```

## Where to Go From Here

To create a new logical host or reconfigure a running logical host to use the new Sun StorEdge T3 disk tray logical volume, see the *Sun Cluster 2.2 Software Installation Guide*.

## ▼ How to Remove a Sun StorEdge T3 Disk Tray Logical Volume

Use this procedure to remove a Sun StorEdge T3 disk tray logical volume. This procedure also assumes all cluster nodes are booted and attached to the Sun StorEdge T3 disk array that hosts the Sun StorEdge T3 disk tray logical volume you want to remove.



---

**Caution** – This procedure removes all data on the Sun StorEdge T3 disk tray logical volume you want to remove.

---

1. **If necessary, migrate all data and volumes off the Sun StorEdge T3 disk tray logical volume you want to remove. Otherwise, proceed to [Step 3](#).**
2. **Determine if the Sun StorEdge T3 disk tray logical volume you want to remove is a quorum device.**

```
scstat -q
```

If the disk drive you want to remove is a quorum device, put the quorum device into maintenance state before you proceed. Otherwise, proceed to [Step 1](#).

For the procedure on putting a quorum device into maintenance state, see the *Sun Cluster 3.0 U1 System Administration Guide*.

3. **Run the appropriate Solstice DiskSuite or VERITAS Volume Manager commands to remove the reference to the logical unit number (LUN) from any diskset or disk group.**

For more information, see your Solstice DiskSuite or VERITAS Volume Manager documentation.

4. **Telnet to the Sun StorEdge T3 disk array that hosts the Sun StorEdge T3 disk tray logical volume you want to remove.**
5. **Remove the Sun StorEdge T3 disk tray logical volume.**

For the procedure on deleting a Sun StorEdge T3 disk tray logical volume, see the *Sun StorEdge T3 Disk Tray Administrator's Guide*.

**6. On all cluster nodes, update the /devices and /dev entries.**

Solaris 2.6:

```
drvconfig
disks -C
devlinks
```

Solaris 8:

```
devfsadm -C
```

## Where to Go From Here

To create a Sun StorEdge T3 disk tray logical volume, see the [“How to Create a Sun StorEdge T3 Disk Tray Logical Volume”](#) on page 55.

---

# Maintaining a Sun StorEdge T3 Disk Array

This section provides the procedures for maintaining a Sun StorEdge T3 disk array. The following table lists these procedures. This section does not include a procedure for adding a disk drive because a Sun StorEdge T3 disk array only operates when fully configured. Moreover, this section does not include a procedure for removing a disk drive because you must replace, within 30 minutes, any FRUs you remove.



---

**Caution** – If you remove any field replaceable unit (FRU) for an extended period of time, thermal complications might result. To prevent this complication, the Sun StorEdge T3 disk array is designed so an orderly shutdown occurs when you remove a component for longer than 30 minutes. A replacement part must be immediately available before starting a FRU replacement procedure. You must replace a FRU within 30 minutes or the Sun StorEdge T3 disk array, and all attached Sun StorEdge T3 disk arrays, will shut down and power off.

---

**TABLE A-2** Task Map: Maintaining a Sun StorEdge T3 Disk Array

| Task                                                                                    | For Instructions, Go To...                                                                           |
|-----------------------------------------------------------------------------------------|------------------------------------------------------------------------------------------------------|
| Upgrade Sun StorEdge T3 disk array firmware                                             | <a href="#">“How to Upgrade Sun StorEdge T3 Disk Array Firmware” on page 61</a>                      |
| Replace a disk drive                                                                    | <a href="#">“How to Replace a Disk Drive” on page 62</a>                                             |
| Add a Sun StorEdge T3 disk array                                                        | <a href="#">“How to Add a Sun StorEdge T3 Disk Array” on page 62</a>                                 |
| Remove a Sun StorEdge T3 disk array                                                     | <a href="#">“How to Remove a Sun StorEdge T3 Disk Array” on page 70</a>                              |
| Replace a host-to-hub fiber optic cable                                                 | <a href="#">“How to Replace a Host-to-Hub Component” on page 72</a>                                  |
| Replace an FC-100 hub GBIC that connects a FC-100 hub to a host                         | <a href="#">“How to Replace a Host-to-Hub Component” on page 72</a>                                  |
| Replace a hub-to-disk array fiber optic cable                                           | <a href="#">“How to Replace a Sun StorEdge FC-100 Hub and Hub-to-Disk Tray Component” on page 73</a> |
| Replace an FC-100 hub GBIC that connects the FC-100 hub to a Sun StorEdge T3 disk array | <a href="#">“How to Replace a Sun StorEdge FC-100 Hub and Hub-to-Disk Tray Component” on page 73</a> |

**TABLE A-2** Task Map: Maintaining a Sun StorEdge T3 Disk Array (*Continued*)

| <b>Task</b>                                                                                                      | <b>For Instructions, Go To...</b>                                                                    |
|------------------------------------------------------------------------------------------------------------------|------------------------------------------------------------------------------------------------------|
| Replace a Sun StorEdge FC-100 hub                                                                                | <a href="#">"How to Replace a Sun StorEdge FC-100 Hub and Hub-to-Disk Tray Component" on page 73</a> |
| Replace a Sun StorEdge FC-100 hub power cord                                                                     | <a href="#">"How to Replace a Sun StorEdge FC-100 Hub and Hub-to-Disk Tray Component" on page 73</a> |
| Replace a media interface adapter (MIA)                                                                          | <a href="#">"How to Replace a Sun StorEdge FC-100 Hub and Hub-to-Disk Tray Component" on page 73</a> |
| Replace a Sun StorEdge T3 disk array controller                                                                  | <a href="#">"How to Replace a Sun StorEdge T3 Disk Array Controller" on page 74</a>                  |
| Replace a Sun StorEdge T3 disk array chassis                                                                     | <a href="#">"How to Replace a Sun StorEdge T3 Disk Array Chassis" on page 74</a>                     |
| Replace a host adapter                                                                                           | <a href="#">"How to Replace a Host Adapter" on page 75</a>                                           |
| Replace a Power and Cooling Unit (PCU)<br>Follow the same procedure used in a non-cluster environment.           | <i>Sun StorEdge T3 Installation, Operation, and Service Manual</i>                                   |
| Replace a unit interconnect card (UIC)<br>Follow the same procedure used in a non-cluster environment.           | <i>Sun StorEdge T3 Installation, Operation, and Service Manual</i>                                   |
| Replace a Sun StorEdge T3 disk array power cable<br>Follow the same procedure used in a non-cluster environment. | <i>Sun StorEdge T3 Installation, Operation, and Service Manual</i>                                   |
| Replace an Ethernet cable<br>Follow the same procedure used in a non-cluster environment.                        | <i>Sun StorEdge T3 Installation, Operation, and Service Manual</i>                                   |

## ▼ How to Upgrade Sun StorEdge T3 Disk Array Firmware

Use this procedure to upgrade Sun StorEdge T3 disk array firmware in a running cluster. Sun StorEdge T3 disk array firmware includes controller firmware, unit interconnect card (UIC) firmware, and disk drive firmware.



---

**Caution** – Perform this procedure on one Sun StorEdge T3 disk array at a time. This procedure requires that you reset the Sun StorEdge T3 disk array you want to upgrade. If you reset more than one Sun StorEdge T3 disk array, your cluster will lose access to data if the Sun StorEdge T3 disk array are submirrors of each other.

---

1. **On one node attached to the Sun StorEdge T3 disk array you want to upgrade, detach that Sun StorEdge T3 disk array's submirrors.**

For more information, see your Solstice DiskSuite or VERITAS Volume Manager documentation.

2. **Apply the controller, disk drive, and UIC firmware patches.**

For the list of required Sun StorEdge T3 disk array patches, see the *Sun StorEdge T3 Disk Tray Release Notes*. For the procedure on applying firmware patches, see the firmware patch README. For the procedure on verifying the firmware level, see the *Sun StorEdge T3 Installation, Operation, and Service Manual*.

3. **Reset the Sun StorEdge T3 disk array, if you have not already done so.**

For the procedure on rebooting a Sun StorEdge T3 disk array, see the *Sun StorEdge T3 Installation, Operation, and Service Manual*.

4. **Reattach the submirrors to resynchronize them.**

For more information, see your Solstice DiskSuite or VERITAS Volume Manager documentation.

## ▼ How to Replace a Disk Drive

Use this procedure to replace one failed disk drive in a Sun StorEdge T3 disk array in a running cluster.



---

**Caution** – If you remove any field replaceable unit (FRU) for an extended period of time, thermal complications might result. To prevent this complication, the Sun StorEdge T3 disk array is designed so an orderly shutdown occurs when you remove a component for longer than 30 minutes. A replacement part must be immediately available before starting a FRU replacement procedure. You must replace a FRU within 30 minutes or the Sun StorEdge T3 disk array, and all attached Sun StorEdge T3 disk arrays, will shut down and power off.

---

1. **If the failed disk drive impacted the Sun StorEdge T3 disk tray logical volume's availability, remove the Sun StorEdge T3 disk tray logical volume from volume management control. Otherwise, proceed to [Step 2](#).**

For more information, see your Solstice DiskSuite or VERITAS Volume Manager documentation.

2. **Replace the disk drive.**

For the procedure on replacing a disk drive, see the *Sun StorEdge T3 Installation, Operation, and Service Manual*.

3. **If you removed a LUN from volume management control in [Step 1](#), return the LUN(s) to volume management control. Otherwise, [Step 2](#) completes this procedure.**

For more information, see your Solstice DiskSuite or VERITAS Volume Manager documentation.

## ▼ How to Add a Sun StorEdge T3 Disk Array

Use this procedure to add a Sun StorEdge T3 disk array to a running cluster. This procedure provides flexibility so you can either add a Sun StorEdge T3 disk array, a Sun StorEdge FC-100 hub and host adapters, or add a Sun StorEdge T3 disk array to an existing Sun StorEdge FC-100 hub and host adapters.

This procedure defines Node A as the node you want to begin working with, and Node B as the remaining node.

1. **If you have to install a host adapter in Node A, proceed to [Step 2](#). Otherwise, skip to [Step 8](#).**

**2. Is the host adapter you want to install the first FC-100/S host adapter on the nodes that will be connected to the Sun StorEdge T3 disk array you want to add?**

- If no, skip to [Step 4](#).
- If yes, determine whether the Fibre Channel support packages are already installed on these nodes. This product requires the following packages.

```
pkginfo | egrep Wlux
system SUNWluxd Sun Enterprise Network Array sf Device Driver
system SUNWluxdx Sun Enterprise Network Array sf Device Driver (64-bit)
system SUNWluxl Sun Enterprise Network Array social Device Driver
system SUNWluxlx Sun Enterprise Network Array social Device Driver (64-bit)
system SUNWluxop Sun Enterprise Network Array firmware and utilities
```

**3. Are the Fibre Channel support packages installed?**

- If yes, proceed to [Step 4](#).
- If no, install them.

The Sun StorEdge T3 disk array packages are located in the `Product` directory of the CD-ROM. Use the `pkgadd` command to add any necessary packages.

```
pkgadd -d path_to_Solaris/Product Pkg1 Pkg2 Pkg3 ... PkgN
```

**4. If necessary, move Node A's services to Node B.**

```
haswitch nodename logicalhost
```

**5. Stop the Sun Cluster software on Node A, and then shut down Node A.**

```
scadmin stopnode
shutdown -i0 -g0 -y
```

For the procedure on shutting down a node, see the *Sun Cluster 2.2 System Administration Guide*.

**6. Power off Node A.**

**7. Install the host adapter in Node A.**

For the procedure on installing a host adapter, see the documentation that shipped with your host adapter and node.

**8. Install the gigabit interface converters (GBIC) in the host adapter and the Sun StorEdge FC-100 hub.**

For the procedure on installing a GBIC, see the *FC-100 Hub Installation and Service Manual*. For the procedure on powering on a Sun StorEdge T3 disk array, see the *Sun StorEdge T3 Installation, Operation, and Service Manual*.

**9. Install the media interface adapter (MIA) in the Sun StorEdge T3 disk array you want to add.**

For the procedure on installing a media interface adapter (MIA), see the *Sun StorEdge T3 Configuration Guide*.

**10. Install the fiber optic cables and the Ethernet cable as shown in [FIGURE A-2](#).**

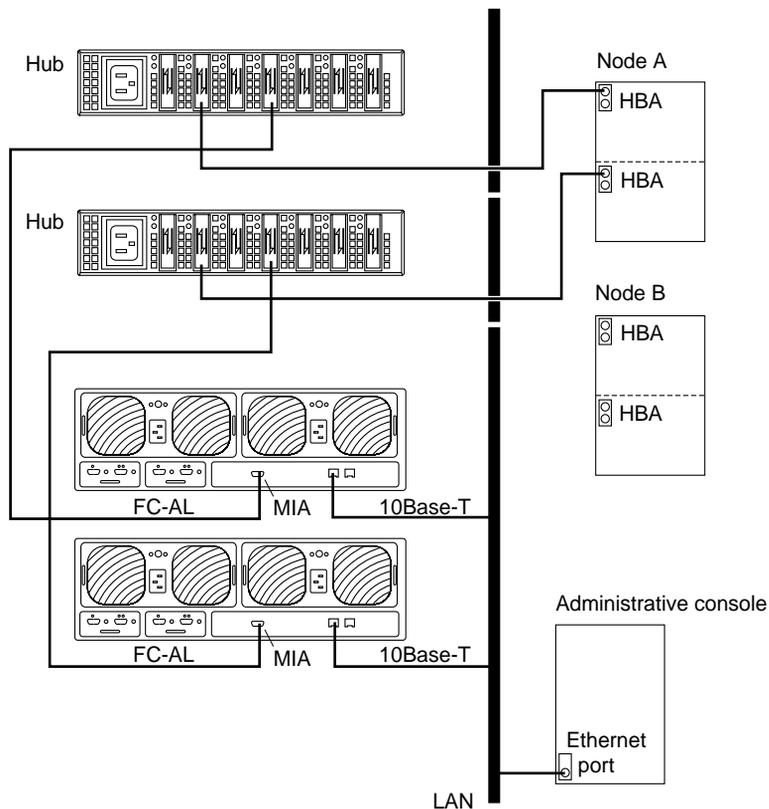
Connect the fiber optic cable between the Sun StorEdge FC-100 hub and the Sun StorEdge T3 disk array, connect the fiber optic cable between the Sun StorEdge FC-100 hub and Node A, and install the Ethernet cable.

---

**Note** – There are no hub port assignment restrictions. You can connect your Sun StorEdge T3 disk array and node to any hub port.

---

For the procedure on installing a fiber optic cable, see the *Sun StorEdge T3 Configuration Guide*.



**FIGURE A-2** Adding a Sun StorEdge T3 disk array in a single-tray configuration

11. Assign an IP address to the Sun StorEdge T3 disk array you want to add.
12. Telnet to the Sun StorEdge T3 disk array you want to add, and install the required Sun StorEdge T3 disk array controller firmware.  
For the required Sun StorEdge T3 disk array controller firmware revision number, see [“Features” on page 5](#). For the procedure on upgrading firmware, see the firmware patch README.
13. Ensure that this new Sun StorEdge T3 disk array has a unique target address.  
For the procedure on verifying and assigning a target address, see the *Sun StorEdge T3 Configuration Guide*.
14. Configure the new Sun StorEdge T3 disk array.  
For the procedure on creating a Sun StorEdge T3 disk tray logical volume, see the *Sun StorEdge T3 Disk Tray Administrator’s Guide*.

**15. Power on and boot Node A.**

```
{0} ok boot
```

**16. If necessary, upgrade the host adapter firmware on Node A.**

For the list of required host adapter firmware patches, see the *Sun StorEdge T3 Disk Tray Release Notes*. For the procedure on applying host adapter firmware patches, see the firmware patch README.

**17. On Node A, install the required Solaris patches for Sun StorEdge T3 disk array support.**

For a list of required Solaris patches for Sun StorEdge T3 disk array support, see the *Sun StorEdge T3 Disk Tray Release Notes*.

**18. Reboot Node A.**

```
reboot -- -r
```

**19. Start the Sun Cluster software on Node A.**

```
scadmin startnode
```

For more information, see the *Sun Cluster 2.2 System Administration Guide*.

**20. (Optional) On Node A, verify that the device names have been assigned to the new Sun StorEdge T3 disk array.**

```
format
```

**21. If you have to install a host adapter in Node B, proceed to [Step 22](#). Otherwise, skip to [Step 26](#).**

**22. If necessary, move Node B's services to Node A.**

```
haswitch nodename logicalhost
```

**23. Stop the Sun Cluster software on Node B, then shut down the node.**

```
scadmin stopnode
shutdown -i0 -g0 -y
```

For the procedure on shutting down a node, see the *Sun Cluster 2.2 System Administration Guide*.

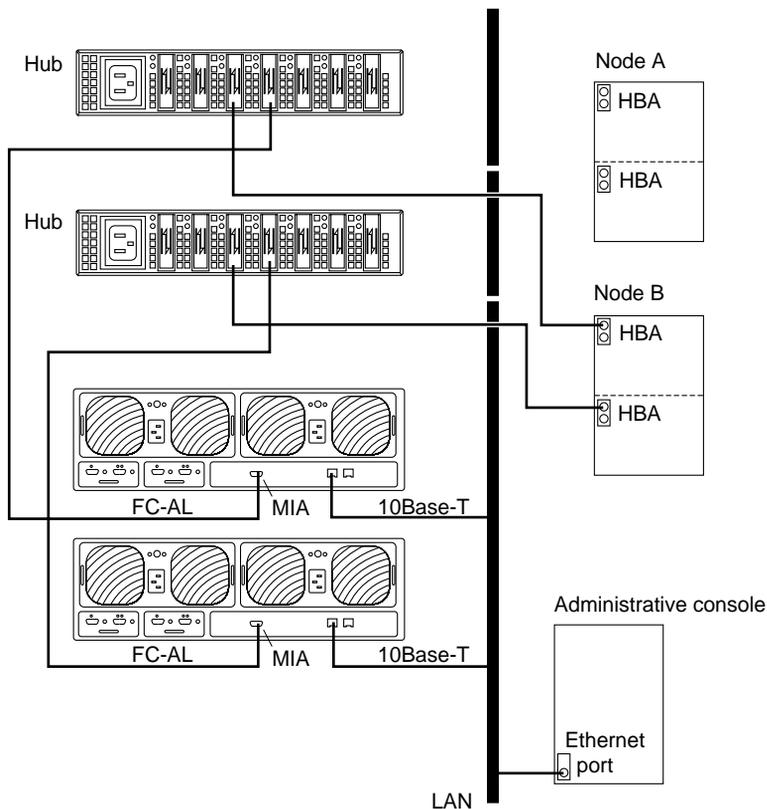
**24. Power off Node B.**

**25. Install the host adapter in Node B.**

For the procedure on installing a host adapter, see the documentation that shipped with your host adapter and node.

**26. Install the gigabit interface converters (GBIC), and connect the fiber optic cable between the Sun StorEdge FC-100 hub and Node B as shown in [FIGURE A-3](#).**

For the procedure on installing a GBIC or Sun StorEdge FC-100 hub, see the *FC-100 Hub Installation and Service Manual*. For the procedure on installing a fiber optic cable, see the *Sun StorEdge T3 Configuration Guide*.



**FIGURE A-3** Adding a Sun StorEdge T3 disk array in a single-tray configuration

**27. Power on and boot Node B.**

```
{0} ok boot
```

**28. Power on Node B, and, without allowing it to join the cluster, boot Node B.**

```
{0} ok boot
```

**29. If necessary, upgrade the host adapter firmware on Node B.**

For the list of required host adapter firmware patches, see the *Sun StorEdge T3 Disk Tray Release Notes*. For the procedure on applying host adapter firmware patches, see the firmware patch README.

**30. On Node B, install the required Solaris patches for Sun StorEdge T3 disk array support.**

For a list of required Solaris patches for Sun StorEdge T3 disk array support, see the *Sun StorEdge T3 Disk Tray Release Notes*.

**31. Reboot Node B.**

```
reboot -- -r
```

**32. Start the Sun software on Node B.**

```
scadmin startnode
```

For more information, see the *Sun Cluster 2.2 System Administration Guide*.

**33. (Optional) On Node B, verify that the device names have been assigned to the new Sun StorEdge T3 disk array.**

```
format
```

**34. If you moved all services off Node B in [Step 22](#), return these services to Node B.**

```
haswitch nodename logicalhost
```

For more information, see the *Sun Cluster 2.2 System Administration Guide*.

**35. Perform volume management administration to incorporate the new Sun StorEdge T3 disk tray logical volumes into the cluster.**

For more information, see your Solstice DiskSuite or VERITAS Volume Manager documentation.

## ▼ How to Remove a Sun StorEdge T3 Disk Array

Use this procedure to permanently remove a Sun StorEdge T3 disk array and its submirrors from a running cluster. This procedure provides the flexibility to remove the host adapters from the nodes for the Sun StorEdge T3 disk array you want to remove.

This procedure defines Node A as the node you want to begin working with, and Node B as the remaining node.



---

**Caution** – During this procedure, you will lose access to the data that resides on the Sun StorEdge T3 disk array you want to remove.

---

1. **If necessary, back up all database tables, data services, and volumes associated with the Sun StorEdge T3 disk array you want to remove.**
2. **Detach the submirrors from the Sun StorEdge T3 disk array you want to remove to stop all I/O activity to the Sun StorEdge T3 disk array.**  
For more information, see your Solstice DiskSuite or VERITAS Volume Manager documentation.
3. **Run the appropriate Solstice DiskSuite or VERITAS Volume Manager commands to remove the references to the LUN(s) from any diskset or disk group.**  
For more information, see your Solstice DiskSuite or VERITAS Volume Manager documentation.
4. **On one node connected to the Sun StorEdge T3 disk array you want to remove, run the `luxadm remove` command, then, when prompted, disconnect the Sun StorEdge T3 disk array.**

```
luxadm remove -F pathname
```

5. **Disconnect the fiber optic cables from the Sun StorEdge T3 disk array you want to remove.**  
Your volume manager might report disk drive failures. This reporting is expected behavior.
6. **If you want to remove the host adapter from Node A, proceed to [Step 7](#). Otherwise, skip to [Step 20](#).**
7. **Move Node A's services to Node B.**

```
haswitch nodename logicalhost
```

**8. Stop the Sun Cluster software on Node A, and then shut down Node A.**

```
scadmin stopnode
shutdown -i0 -g0 -y
```

**9. Power off Node A.**

**10. Remove the host adapter from Node A.**

For the procedure on removing host adapters, see the documentation that shipped with your nodes.

**11. Power on and boot Node A.**

```
{0} ok boot
```

**12. Start the Sun Cluster software on Node A.**

```
scadmin startnode
```

For more information, see the *Sun Cluster 2.2 System Administration Guide*.

**13. If you have to remove a host adapter in Node B, proceed to [Step 14](#). Otherwise, skip to [Step 20](#).**

**14. Move Node B's services to Node B.**

```
haswitch nodename logicalhost
```

**15. Stop the Sun Cluster software on Node B, and then shut down Node B.**

```
scadmin stopnode
shutdown -i0 -g0 -y
```

**16. Power off Node B.**

**17. Remove the host adapter from Node B.**

For the procedure on removing host adapters, see the documentation that shipped with your nodes.

## 18. Power on and boot Node B.

```
{0} ok boot
```

## 19. Start the Sun Cluster software on Node B.

```
scadmin startnode
```

For more information, see the *Sun Cluster 2.2 System Administration Guide*.

## 20. On all cluster nodes, update the /devices and /dev entries.

Solaris 2.6:

```
drvconfig
disks -C
devlinks
```

Solaris 8:

```
devfsadm -C
```

## ▼ How to Replace a Host-to-Hub Component

Use this procedure to replace the following host-to-hub components.

- Host-to-hub fiber optic cable
- FC-100/S host adapter Gigabit Interface Converter (GBIC)
- FC-100 hub GBIC that connects a Sun StorEdge FC-100 hub to a node

1. On the node connected to the host-to-hub connection you want to replace, determine the logical hosts owned by this node.

```
hastat
```

2. Move the resource groups you identified in [Step 1](#) to another node.

```
haswitch nodename logicalhost
```

### 3. Replace the host-to-hub component.

For the procedure on replacing a fiber optic cable, see the *Sun StorEdge T3 Configuration Guide*. For the procedure on replacing a FC-100/S host adapter GBIC, see the *FC-100 Hub Installation and Service Manual*.

### 4. Return the resource groups you moved in [Step 2](#).

```
haswitch nodename logicalhost
```

## ▼ How to Replace a Sun StorEdge FC-100 Hub and Hub-to-Disk Tray Component

Use this procedure to replace the following hub-to-disk array components.

- Fiber optic cable that connects a Sun StorEdge FC-100 hub to a Sun StorEdge T3 disk array
- FC-100 hub GBIC that connects a Sun StorEdge FC-100 hub to a Sun StorEdge T3 disk array
- Sun StorEdge FC-100 hub
- Sun StorEdge FC-100 hub power cord
- Media Interface Adapter (MIA)

### 1. Detach the submirrors on the Sun StorEdge T3 disk array connected to the hub-to-disk array fiber optic cable you want to replace to stop all I/O activity to this Sun StorEdge T3 disk array.

For more information, see your Solstice DiskSuite or VERITAS Volume Manager documentation.

### 2. Replace the hub-to-disk array component.

For the procedure on replacing a fiber optic cable, see the *Sun StorEdge T3 Configuration Guide*. For the procedure on replacing an FC-100 hub GBIC, a Sun StorEdge FC-100 hub, or a Sun StorEdge FC-100 hub power cord, see the *FC-100 Hub Installation and Service Manual*. For the procedure on replacing an MIA, see the *Sun StorEdge T3 Installation, Operation, and Service Manual*.

### 3. Reattach the submirrors to resynchronize them.

For more information, see your Solstice DiskSuite or VERITAS Volume Manager documentation.

## ▼ How to Replace a Sun StorEdge T3 Disk Array Controller

Use this procedure to replace a Sun StorEdge T3 disk array controller.

1. **Detach the submirrors on the Sun StorEdge T3 disk array connected to the controller you want to replace to stop all I/O activity to this Sun StorEdge T3 disk array.**

For more information, see your Solstice DiskSuite or VERITAS Volume Manager documentation.

2. **Replace the controller.**

For the procedure on replacing a Sun StorEdge T3 controller, see the *Sun StorEdge T3 Installation, Operation, and Service Manual*.

3. **Reattach the submirrors to resynchronize them.**

For more information, see your Solstice DiskSuite or VERITAS Volume Manager documentation.

## ▼ How to Replace a Sun StorEdge T3 Disk Array Chassis

Use this procedure to replace a Sun StorEdge T3 disk array chassis. This procedure assumes that you want to retain all FRUs other than the chassis and the backplane. To replace the chassis, you must replace both the chassis and the backplane because these components are manufactured as one part.

---

**Note** – Only trained, qualified service providers should use this procedure to replace a Sun StorEdge T3 disk array chassis.

---

1. **Detach the submirrors on the Sun StorEdge T3 disk array connected to the chassis you want to replace to stop all I/O activity to this Sun StorEdge T3 disk array.**

For more information, see your Solstice DiskSuite or VERITAS Volume Manager documentation.

2. **Replace the chassis/backplane.**

For the procedure on replacing a Sun StorEdge T3 chassis, see the *Sun StorEdge T3 Installation, Operation, and Service Manual*.

### 3. Reattach the submirrors to resynchronize them.

---

**Note** – Account for the change in the World Wide Number (WWN).

---

For more information, see your Solstice DiskSuite or VERITAS Volume Manager documentation.

## ▼ How to Replace a Host Adapter

Use this procedure to replace a failed host adapter in a running cluster. The node in this procedure refers to the node with the failed host adapter you want to replace.

1. **Move the node's services to the other node within the cluster, stop Sun Cluster software from running, and then shut down the node.**

```
haswitch nodename logicalhost
scadmin stopnode
shutdown -y -g 0
```

For the procedure on shutting down a node, see the *Sun Cluster 2.2 System Administration Guide*.

2. **Power off the node.**
3. **Replace the failed host adapter.**

For the procedure on removing and adding host adapters, see the documentation that shipped with your nodes.

4. **Power on the node.**
5. **Boot the node to prepare it for cluster mode.**

```
{0} ok boot
```

**6. Log on and bring the node into cluster mode.**

If this is the first node in the cluster:

```
scadmin startcluster hostname clustername
```

If this is not the first node in the cluster:

```
scadmin startnode
```

For more information, see the *Sun Cluster 2.2 System Administration Guide*.