



Sun Cluster 3.0 U1 Release Notes Supplement

Sun Microsystems, Inc.
901 San Antonio Road
Palo Alto, CA 94303-4900 U.S.A.
650-960-1300

Part No. 816-2032-10
October 2001, Revision A

Copyright 2001 Sun Microsystems, Inc., 901 San Antonio Road, Palo Alto, CA 94303, U.S.A. All rights reserved.

Sun Microsystems, Inc. has intellectual property rights relating to technology embodied in the product that is described in this document. In particular, and without limitation, these intellectual property rights may include one or more of the U.S. patents listed at <http://www.sun.com/patents> and one or more additional patents or pending patent applications in the U.S. and in other countries.

This product or document is 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.

Sun, Sun Microsystems, the Sun logo, Java, Netra, Solaris, Sun StorEdge, iPlanet, Sun Cluster, SunPlex, Answerbook2, docs.sun.com, Solstice DiskSuite, Sun Enterprise, Sun Enterprise SyMON, Solaris JumpStart, JumpStart, Sun Management Center, OpenBoot, Sun Fire, SunSolve, SunSwift, the 100% Pure Java logo, the AnswerBook logo, the Netra logo, the Solaris logo and the iPlanet logo are trademarks or registered trademarks 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 architecture developed by Sun Microsystems, Inc.

ORACLE® is a registered trademark of Oracle Corporation. Netscape™ is a trademark or registered trademark of Netscape Communications Corporation in the United States and other countries. The Adobe® logo is a registered trademark of Adobe Systems, Incorporated.

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

This product includes software developed by the Apache Software Foundation (<http://www.apache.org/>).

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 2001 Sun Microsystems, Inc., 901 San Antonio Road, Palo Alto, CA 94303, Etats-Unis. Tous droits réservés.

Sun Microsystems, Inc. a les droits de propriété intellectuels relatants à la technologie incorporée dans le produit qui est décrit dans ce document. En particulier, et sans la limitation, ces droits de propriété intellectuels peuvent inclure un ou plus des brevets américains énumérés à <http://www.sun.com/patents> et un ou les brevets plus supplémentaires ou les applications de brevet en attente dans les Etats - Unis et dans les autres pays.

Ce produit ou document est 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.

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.

Sun, Sun Microsystems, le logo Sun, Java, Netra, Solaris, Sun StorEdge, iPlanet, Sun Cluster, Sun Plex, Answerbook2, docs.sun.com, Solstice DiskSuite, Sun Enterprise, Sun Enterprise SyMON, Solaris JumpStart, JumpStart, Sun Management Center, OpenBoot, Sun Fire, SunSolve, SunSwift, le logo 100% Pure Java, le logo AnswerBook, le logo Netra, le logo Solaris et le logo iPlanet sont des marques de fabrique ou des marques déposées 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.

ORACLE® est une marque déposée registre de Oracle Corporation. Netscape™ est une marque de Netscape Communications Corporation aux Etats-Unis et dans d'autres pays. Le logo Adobe® est une marque déposée de Adobe Systems, Incorporated.

Ce produit inclut le logiciel développé par la base de Apache Software Foundation (<http://www.apache.org/>).



Adobe PostScript

LA DOCUMENTATION EST FOURNIE "EN L'ETAT" ET TOUTES AUTRES CONDITIONS, DECLARATIONS ET GARANTIES EXPRESSES OU TACITES SONT FORMELLEMENT EXCLUES, DANS LA MESURE AUTORISEE PAR LA LOI APPLICABLE, Y COMPRIS NOTAMMENT TOUTE GARANTIE IMPLICITE RELATIVE A LA QUALITE MARCHANDE, A L'APTITUDE A UNE UTILISATION PARTICULIERE OU A L'ABSENCE DE CONTREFAÇON.

Sun Cluster 3.0 U1 Release Notes Supplement

This document supplements the standard user documentation, including the *Sun Cluster 3.0 U1 Release Notes* shipped with the Sun™ Cluster 3.0 product. These “online release notes” provide the most current information on the Sun Cluster 3.0 product. This document includes the following information:

- [Revision Record](#)
- [New Features](#)
- [Restrictions and Requirements](#)
- [Known Problems](#)

Revision Record

The following table lists the information contained in this document and provides the revision date for this information.

TABLE 1 Sun Cluster 3.0 U1 Release Notes Supplement (Sheet 1 of 2)

Revision Date	Information Added
October 2001	Added special requirement for Data Services. See “Data Service Special Requirements” on page 20.
September 2001	Revised <code>/etc/serialports</code> setup instruction to support Sun Fire™ servers. See “Cluster Control Panel Support for Sun Fire Servers” on page 11. Support for HA Oracle 9i. See “Support for Oracle 9i” on page 12. Guideline for using VxVM 3.2 Enclosure-Based Naming in a Sun Cluster environment. See “Volume Manager Restrictions and Requirements” on page 18. Support for Sun StorEdge™ Network Data Replicator (Sun SNDR) 3.0 and Sun StorEdge Instant Image 3.0. See “Sun Cluster 3.0 Update 1 With Sun StorEdge 3.0 Services Software Restrictions and Requirements” on page 18. Updated Sun StorEdge T3 partner-group cluster procedures to fix problems and to support the new T3+ product. See Appendix B. Updated Sun StorEdge T3 single-controller cluster procedures to fix problems and to support the new T3+ product. See Appendix C. Updated Netra™ D130 cluster procedures to include support for the StorEdge S1 product. See Appendix D.

TABLE 1 Sun Cluster 3.0 U1 Release Notes Supplement (Sheet 2 of 2)

Revision Date	Information Added
July 2001	<p data-bbox="379 239 1300 262">Support for iPlanet™ Web Server 6.0. See “Support for iPlanet Web Server 6.0” on page 7.</p> <p data-bbox="379 284 1205 307">“Bug ID 4349995” on page 21, restriction for using Solstice DiskSuite metatool.</p> <p data-bbox="379 329 1272 383">“Bug ID 4480277” on page 21, a panicked volume manager slave node that attempts to rejoin the cluster might cause the volume manager master node to crash.</p> <p data-bbox="379 406 1300 460">“Bug ID 4388265” on page 21, errors while replacing a SCSI cable in a Sun StorEdge A3500 disk array.</p> <p data-bbox="379 482 1300 536">“Bug ID 4410535” on page 22, workaround for Sun Management Center if you cannot add a previously deleted resource group.</p> <p data-bbox="379 558 1300 612">Description of a known documentation error in the <i>Sun Cluster 3.0 U1 Hardware Guide</i>. See “Known Documentation Problems” on page 23.</p> <p data-bbox="379 635 1300 689">Appendix about installing and using the Sun Cluster module with the Sun Management Center 3.0 graphical user interface. See Appendix A.</p> <p data-bbox="379 711 1300 765">Appendix to document Sun StorEdge T3 partner-group configurations for clustering. See Appendix B.</p> <p data-bbox="379 788 1300 841">Appendix to document Netra D130 storage unit procedures for clustering. See Appendix D.</p>

New Features

In addition to features documented in *Sun Cluster 3.0 U1 Release Notes*, this release now includes support for the following features.

Support for Sun SNDR 3.0 and Sun StorEdge Instant Image 3.0

Sun Cluster 3.0 Update 1 now supports Sun StorEdge Network Data Replicator (Sun SNDR) 3.0 and Sun StorEdge Instant Image 3.0 as cluster-aware products. These software products are part of the Sun StorEdge Version 3.0 software package. Sun StorEdge Fast Write Cache, which is also a part of the Sun StorEdge Version 3.0 software package, is not supported in any Sun Cluster environment.

The Sun SNDR software is a data replication application that provides access to data as part of business continuance and disaster recovery plans. The Sun StorEdge Instant Image software is a point-in-time copy application that enables you to create copies of application or test data. These software products are now cluster aware and will fail over and switch back in a Sun Cluster environment. For more information, see the following documents, available online at <http://docs.sun.com>.

- *Sun Cluster 3.0 U1 and Sun StorEdge Software 3.0 Integration Guide*
- *Sun StorEdge Network Data Replicator 3.0 System Administrator's Guide*
- *Sun StorEdge Instant Image 3.0 System Administrator's Guide*

There are some restrictions that apply to Sun Cluster 3.0 Update 1 only when using Sun SNDR or Sun StorEdge Instant Image. See “[Sun Cluster 3.0 Update 1 With Sun StorEdge 3.0 Services Software Restrictions and Requirements](#)” on page 18 for more information.

Sun SNDR 3.0 and Sun StorEdge Instant Image 3.0 Support Patches

The following patches are required to support Sun SNDR 3.0 and Sun StorEdge Instant Image 3.0 with Sun Cluster 3.0 Update 1. Patches are available from SunSolve at <http://sunsolve.sun.com/>.

111945-xx	Storage Cache Manager
111946-xx	Storage Volume Driver
111947-xx	Sun StorEdge Instant Image
111948-xx	Sun SNDR

Support for iPlanet Web Server 6.0

Sun Cluster 3.0 Update 1 now supports iPlanet Web Server 6.0.

Two procedures have changed for iPlanet Web Server 6.0.

- [“Installing Certificates on Secure Instances of iPlanet Web Server 6.0” on page 7](#)
- [“Specifying the Location of the Access Logs” on page 9](#)

▼ Installing Certificates on Secure Instances of iPlanet Web Server 6.0

The procedure for installing certificates on secure instances of iPlanet Web Server has changed for version 6.0. If you plan to run secure instances of iPlanet Web Server 6.0, complete the following steps when you install security certificates. This installation procedure requires that you create a certificate on one node, and then create symbolic links to that certificate on the other cluster nodes.

- 1. Run the administrative server on node1.**
- 2. From your Web browser, connect to the administrative server as**
`http://node1.domain:port`.

For example, `http://phys-schost-1.eng.sun.com:8888`. Use whatever port number you specified as the administrative server port during installation. The default port number is 8888.

- 3. Install the certificate on node1.**

This installation creates three certificate files. One file, `secmod.db`, is common to all nodes, and the other two are specific to node1. These files are located in the `alias` subdirectory, under the directory in which the iPlanet Web Server files are installed.

4. If you installed iPlanet Web Server on a global file system, complete the following tasks. If you installed iPlanet Web Server on a local file system, go to [Step 5](#).

a. Note the location and file names for the three files created when installing the certificate in [Step 3](#).

For example, if you installed iPlanet Web Server in `/global/iws/servers`, and you used the IP address “IPx” when installing the certificate, then the paths to the files on node1 would be

```
/global/iws/servers/alias/secmod.db  
/global/iws/servers/alias/https-IPx-node1-cert7.db  
/global/iws/servers/alias/https-IPx-node1-key3.db
```

b. Create symbolic links for all other cluster nodes to the node-specific files for node1.

In the following example, substitute the appropriate file paths for your system.

```
# ln -s /global/iws/servers/alias/https-IPx-node1-cert7.db  
/global/iws/servers/alias/https-IPx-node2-cert7.db  
# ln -s /global/iws/servers/alias/https-IPx-node1-key3.db  
/global/iws/servers/alias/https-IPx-node2-key3.db
```

5. If you installed iPlanet Web Server on a local file system, complete the following tasks.

a. Note the location and file names for the three files created on node1 when installing the certificate in [Step 3](#).

For example, if you installed iPlanet Web Server in `/local/iws/servers`, and you used the IP address “IPx” when installing the certificate, then the paths to the files on node1 would be

```
/local/iws/servers/alias/secmod.db  
/local/iws/servers/alias/https-IPx-node1-cert7.db  
/local/iws/servers/alias/https-IPx-node1-key3.db
```

b. Move the three certificate files to a location on the global file system.

In the following example, substitute the appropriate file paths for your system

```
# mv /local/iws/servers/alias/secmod.db  
/global/secure/secmod.db  
# mv /local/iws/servers/alias/https-IPx-node1-cert7.db  
/global/secure/https-IPx-node1-cert7.db  
# mv /local/iws/servers/alias/https-IPx-node1-key3.db  
/global/secure/https-IPx-node1-key3.db
```

c. **Create symbolic links between the local and global paths of the three certificate files.**

Create the symbolic links on each node in the cluster.

In the following example, substitute the appropriate file paths for your system.

```
# Symbolic links for node1
# ln -s /global/secure/secmod.db
    /local/iws/servers/alias/secmod.db
# ln -s /global/secure/https-IPx-node1-cert7.db
    /local/iws/servers/alias/https-IPx-node1-cert7.db
# ln -s /global/secure/https-IPx-node1-key3.db
    /local/iws/servers/alias/https-IPx-node1-key3.db

# Symbolic links for node2
# ln -s /global/secure/secmod.db
    /local/iws/servers/alias/secmod.db
# ln -s /global/secure/https-IPx-node1-cert7.db
    /local/iws/servers/alias/https-IPx-node2-cert7.db
# ln -s /global/secure/https-IPx-node1-key3.db
    /local/iws/servers/alias/https-IPx-node2-key3.db
```

▼ Specifying the Location of the Access Logs

The procedure for specifying the location of the access logs while configuring an iPlanet Web Server has changed for iPlanet Web Server 6.0. To specify the location of the access logs while configuring an iPlanet Web Server, complete the following steps.

This change replaces Step 6 through Step 8 in the procedure “How to Configure an iPlanet Web Server” in Chapter 3, “Installing and Configuring Sun Cluster HA for iPlanet Web Server,” in the *Sun Cluster 3.0 U1 Data Services Installation and Configuration Guide*.

- 1. Edit the `ErrorLog`, `PidLog`, and access log entries in the `magnus.conf` file to reflect the directory created in Step 5 of the “How to Configure an iPlanet Web Server” procedure in Chapter 3 of the *Sun Cluster 3.0 U1 Data Services Installation and Configuration Guide*, and synchronize the changes from the administrator’s interface.**

The `magnus.conf` file specifies the locations for the error, access, and PID files. Edit this file to change the error, access, and PID file locations to the directory that you created in Step 5 of the “How to Configure an iPlanet Web Server” procedure in Chapter 3 of the *Sun Cluster 3.0 U1 Data Services Installation and Configuration Guide*.

The `magnus.conf` file is located in the config directory of the iPlanet server instance. If the instance directory is located on the local file system, you must modify the `magnus.conf` file on each of the nodes.

Change the following entries:

```
ErrorLog /global/data/netscape/https-schost-1/logs/error
PidLog /global/data/netscape/https-schost-1/logs/pid
...
Init fn=flex-init access="$accesslog" ...
```

to

```
ErrorLog /var/pathname/http-instance/logs/error
PidLog /var/pathname/http-instance/logs/pid
...
Init fn=flex-init access="/var/pathname/http-instance/logs/access" ...
```

As soon as the administrator's interface detects your changes, the interface displays a warning message, as follows.

```
Warning: Manual edits not loaded
Some configuration files have been edited by hand. Use the "Apply"
button on the upper right side of the screen to load the latest
configuration files.
```

2. Click Apply as prompted.

The administrator's interface displays a new web page.

3. Click Load Configuration Files.

Cluster Control Panel Support for Sun Fire Servers

If you intend to use Cluster Control Panel software tools, such as `cconsole`, to connect to a Sun Fire system, use the following instruction to create the `/etc/serialports` file on your administrative console. This instruction replaces Step 8 of the procedure “How to Install Cluster Control Panel Software on the Administrative Console” in the *Sun Cluster 3.0 U1 Installation Guide*.

8. Create an `/etc/serialports` file.

Add an entry for each node in the cluster to the file. Specify the physical node name, the terminal concentrator (TC), System Service Processor (SSP), or Sun Fire system controller hostname, and the port number.

- For the Sun Fire 15K system controller, use the `telnet(1)` port number 23 for the serial port number of each entry.
- For all other servers, use the `telnet(1)` serial port number, not the physical port number. To determine the serial port number, add 5000 to the physical port number. For example, if a physical port number is 6, the serial port number should be 5006.
- For Sun Enterprise E10000 servers, also see the `/opt/SUNWcluster/bin/serialports(4)` man page for details and special considerations.

```
# vi /etc/serialports
node1 TC-hostname port
node2 TC-hostname port
```

<i>node1, node2</i>	Physical names of the cluster nodes
<i>TC-hostname</i>	Hostname of the TC, SSP, or Sun Fire system controller
<i>port</i>	Serial port number

Support for Oracle 9i

Sun Cluster 3.0 Update 1 now supports Oracle 9i.

Download patch T110651 (dash level 04 or higher) “Sun Cluster 3.0: HA-Oracle Patch” from the SunSolve EarlyNotifier page for the Sun Cluster product, at <http://sunsolve.sun.com>.

The procedure for setting up Oracle database permissions has changed to accommodate both Oracle 8i and Oracle 9i. The changed procedure follows.

▼ How to Set Up Oracle Database Permissions

When completing [Step 2](#) or [Step 3](#) of this procedure, select and configure either the Oracle authentication method or the Solaris authentication method for fault monitoring access.

1. Determine the Oracle release you are using.

If you are using Oracle 8i, go to [Step 2](#).

If you are using Oracle 9i, skip to [Step 3](#).

2. Enable access for the user and password to be used for fault monitoring for Oracle 8i.

Note – If you are using Oracle 9i, go to [Step 3](#).

To complete this step, perform **one** of the following tasks, then skip to [Step 4](#).

- **Oracle authentication method for Oracle 8i** – For all supported Oracle releases, enter the following script into the screen that the `svrmgr1` command displays to enable access.

```
# svrmgr1

connect internal;
grant connect, resource to user identified by passwd;
alter user user default tablespace system quota 1m on
system;
grant select on v_$sysstat to user;
grant create session to user;
grant create table to user;
disconnect;

exit;
```

- **Solaris authentication method for Oracle 8i** – Grant permission for the database to use Solaris authentication.

Note – The user for whom you enable Solaris authentication is the user who owns the files under the \$ORACLE_HOME directory. The following code sample shows that the user `oracle` owns these files.

```
# svrmgr1

connect internal;
  create user ops$oracle identified by externally
    default tablespace system quota 1m on system;
  grant connect, resource to ops$oracle;
  grant select on v_$sysstat to ops$oracle;
  grant create session to ops$oracle;
  grant create table to ops$oracle;
disconnect;

exit;
```

3. Enable access for the user and password to be used for fault monitoring for Oracle 9i.

Note – If you are using Oracle 8i, go to [Step 2](#).

- **To use the Oracle authentication method for Oracle 9i** – For all supported Oracle releases, enter the following script into the screen that the `sqlplus` command displays to enable access.

```
# sqlplus "/as sysdba"

grant connect, resource to user identified by passwd;
alter user user default tablespace system quota 1m on
  system;
grant select on v_$sysstat to user;
grant create session to user;
grant create table to user;

exit;
```

- **To use the Solaris authentication method for Oracle 9i** – Grant permission for the database to use Solaris authentication.

Note – The user for which you enable Solaris authentication is the user who owns the files under the \$ORACLE_HOME directory. The following code sample shows that the user oracle owns these files.

```
# sqlplus "/as sysdba"

      create user ops$oracle identified by externally
      default tablespace system quota 1m on system;
      grant connect, resource to ops$oracle;
      grant select on v_$sysstat to ops$oracle;
      grant create session to ops$oracle;
      grant create table to ops$oracle;

exit;
```

4. Configure NET8 for the Sun Cluster software.

The listener.ora and tnsnames.ora files must be accessible from all the nodes in the cluster. Place these files either under the cluster file system or in the local file system of each node that can potentially run the Oracle resources.

Note – If you place the listener.ora and tnsnames.ora files in a location other than the /var/opt/oracle directory or the \$ORACLE_HOME/network/admin directory, then you must specify TNS_ADMIN or an equivalent Oracle variable (see the Oracle documentation for details) in a user-environment file. You must also run the scrgadm(1M) command to set the resource extension parameter User_env, which will source the user-environment file.

Sun Cluster HA for Oracle imposes no restrictions on the listener name—it can be any valid Oracle listener name.

The following code sample identifies the lines in `listener.ora` that are updated.

```
LISTENER =
  (ADDRESS_LIST =
    (ADDRESS =
      (PROTOCOL = TCP)
      (HOST = logical-hostname) <- use logical hostname
      (PORT = 1527)
    )
  )
.
.
SID_LIST_LISTENER =
.
.
(SID_NAME = SID) <- Database name, default is ORCL
```

The following code sample identifies the lines in `tnsnames.ora` that are updated on client machines.

```
service_name =
.
.
  (ADDRESS =
    (PROTOCOL = TCP)
    (HOST = logicalhostname)<- logical hostname
    (PORT = 1527) <- must match port in LISTENER.ORA
  )
)
(CONNECT_DATA =
  (SID = <SID>)) <- database name, default is ORCL
```

The following example shows how to update the `listener.ora` and `tnsnames.ora` files given the following Oracle instances.

Instance	Logical Host	Listener
ora8	hadbms3	LISTENER-ora8
ora7	hadbms4	LISTENER-ora7

The corresponding `listener.ora` entries are the following entries.

```
LISTENER-ora7 =
  (ADDRESS_LIST =
    (ADDRESS =
      (PROTOCOL = TCP)
      (HOST = hadbms4)
      (PORT = 1530)
    )
  )
SID_LIST_LISTENER-ora7 =
  (SID_LIST =
    (SID_DESC =
      (SID_NAME = ora7)
    )
  )
LISTENER-ora8 =
  (ADDRESS_LIST =
    (ADDRESS= (PROTOCOL=TCP) (HOST=hadbms3)(PORT=1806))
  )
SID_LIST_LISTENER-ora8 =
  (SID_LIST =
    (SID_DESC =
      (SID_NAME = ora8)
    )
  )
```

The corresponding `tnsnames.ora` entries are the following entries.

```
ora8 =
(DESCRIPTION =
  (ADDRESS_LIST =
    (ADDRESS = (PROTOCOL = TCP)
      (HOST = hadbms3)
      (PORT = 1806))
    )
  (CONNECT_DATA = (SID = ora8))
)
ora7 =
(DESCRIPTION =
  (ADDRESS_LIST =
    (ADDRESS =
      (PROTOCOL = TCP)
      (HOST = hadbms4)
      (PORT = 1530))
    )
  (CONNECT_DATA = (SID = ora7))
)
```

5. Verify that the Sun Cluster software is installed and running on all nodes.

```
# scstat
```

Restrictions and Requirements

The following restrictions and requirements have been added or updated since the Sun Cluster 3.0 U1 release.

Hardware Restrictions and Requirements

- If you are using a Sun Enterprise 420R server with a PCI-card in slot J4701, the motherboard must be at dash-level 15 or higher (501-5168-15 or higher). To find the motherboard part number and revision level, look at the edge of the board closest to PCI slot 1.

Volume Manager Restrictions and Requirements

- If you use Enclosure-Based Naming of devices (a feature introduced in VxVM version 3.2), ensure that you use consistent device names on all cluster nodes that share the same storage. VxVM does not coordinate these names, so it is up to the administrator to make sure that VxVM assigns the same names to the same devices from different nodes. While failure to assign consistent names will not interfere with correct cluster behavior, it will greatly complicate cluster administration and greatly increase the possibility of configuration errors, potentially leading to loss of data.

Sun Cluster 3.0 Update 1 With Sun StorEdge 3.0 Services Software Restrictions and Requirements

Note – Sun SNDR 3.0 and Sun StorEdge Instant Image 3.0 are not supported with the Sun Cluster 3.0 general release. They are supported only with Sun Cluster 3.0 Update 1 and compatible versions.

The Sun StorEdge 3.0 services software includes Sun StorEdge Network Data Replicator (Sun SNDR) 3.0, Sun StorEdge Instant Image 3.0, and Sun StorEdge Fast Write Cache 3.0. Following are configuration restrictions and requirements that are unique to Sun Cluster 3.0 Update 1 when you run with Sun StorEdge 3.0 services software.

More information about these configuration restrictions and requirements can be found in the *Sun Cluster 3.0 U1 and Sun StorEdge Software 3.0 Integration Guide*. Be sure to read the integration guide before you install and administer Sun SNDR 3.0 or Sun StorEdge Instant Image 3.0 with Sun Cluster software.

■ **Hardware Configuration Restrictions and Requirements**

- The Sun Cluster 3.0 Update 1 and Sun StorEdge 3.0 services software with patches are supported in a two-node cluster environment only.
- You cannot use the Sun StorEdge Fast Write Cache (FWC) product (all versions, including the SUNWnvm Version 3.0 software) in any Sun Cluster environment because cached data is inaccessible from other machines in a cluster. To compensate, you can use a caching array such as the Sun StorEdge A3500 disk array.

■ **Software Installation, Configuration, and Administration Restrictions and Requirements**

- Sun StorEdge 3.0 services software is supported as a cluster-aware product only with Sun Cluster 3.0 Update 1 and compatible versions.
- With the Sun StorEdge Instant Image and Sun SNDR software, each instance of a volume set (master, shadow, bitmap, overflow) and replication set (primary or secondary bitmap) must be in the same disk device group. This allows for all constituent volumes in an instance to fail over and switch back in their entirety. Within a Sun Cluster environment there can be more than one disk device group, that contains instances of Sun StorEdge Instant Image and Sun SNDR constituent volumes.
- The resource group created for Sun StorEdge services software is a lightweight resource group that contains only the associated disk device group and a logical failover host. Do not add additional resources to this resource group or the StorEdge services software might not properly fail over or switch back.
- The resource group name you specify must consist of the disk device group name appended with `-stor-rg`.
- When you add the `SUNW.HAStorage` resource to the resource group, be sure the `AffinityOn` property is set to `True`.
- When you install Sun StorEdge 3.0 services software on a cluster, you must enter the first cluster node's configuration volume location for the cluster when you install on the second cluster node. You will get a warning message that the configuration is already initialized. You can ignore this message.
- When you shut down and restart cluster nodes, your system might experience a panic condition on the node you are restarting. The node panic is expected behavior in the cluster and is part of the cluster software's failfast mechanism. The *Sun Cluster 3.0 U1 Concepts* manual describes this mechanism.
- **Only one system administrator or root user at a time should create and configure Sun StorEdge volume sets** from a single server or node, to avoid corrupting the Sun StorEdge services configuration.

- **Two administrators should not write to the Sun StorEdge services configuration at the same time.** The operations that access the configuration include, but are not limited to, the following list.
 - Creating and deleting volume sets
 - Adding and removing volume sets from I/O groups
 - Assigning new bitmap volumes to a volume set
 - Updating the disk device group or resource name
 - Any operation that changes the Sun StorEdge services and related volume set configuration
- After you synchronize the primary and secondary Sun SNDR volumes, a directory named `._` (period underscore) is created by the cluster file system. You can ignore or delete this directory. When you unmount the cluster file system the directory disappears.
- In addition to the Sun Cluster commands, you can use the command options `C tag` and `-C tag` with the `sndradm` and `iiadm` commands to control the Sun StorEdge services in a Sun Cluster environment. These are described in the Sun SNDR and Sun StorEdge Instant Image system administration documentation.

Data Service Special Requirements

Identify requirements for all data services **before** you begin Solaris and Sun Cluster installation. Failure to do so might result in installation errors that require that you completely reinstall the Solaris and Sun Cluster software.

For example, the Oracle Parallel Fail Safe/Real Application Clusters Guard option of Oracle Parallel Server/Real Application Clusters has special requirements for the hostnames/node names that you use in the cluster. You must accommodate these requirements before you install Sun Cluster software because you cannot change hostnames after you install Sun Cluster software. For more information on the special requirements for the hostnames/node names, see the Oracle Parallel Fail Safe/Real Application Clusters Guard documentation.

Known Problems

In addition to known problems documented in *Sun Cluster 3.0 U1 Release Notes*, the following known problems affect the operation of the Sun Cluster 3.0 U1 release.

Bug ID 4349995

Problem Summary: The DiskSuite Tool (`metatool`) graphical user interface is incompatible with Sun Cluster 3.0.

Workaround: Use command line interfaces to configure and manage shared disksets.

Bug ID 4388265

Problem Summary: You might encounter I/O errors while replacing a SCSI cable from the Sun StorEdge A3500 controller board to the disk tray. These errors are temporary and should disappear when the cable is securely in place.

Workaround: After replacing a SCSI cable in a Sun StorEdge A3500 disk array, use your volume management recovery procedure to recover from I/O errors.

Bug ID 4480277

Problem Summary: For VERITAS Volume Manager, during volume recovery a volume manager (VM) slave node that attempts to join the cluster might cause the VM master node to crash. “VM slave node” means any node that leaves the cluster as the result of a panic.

Workaround: To prevent the VM slave node from causing the VM master node to crash while the master node is synchronizing the shared volumes, perform the following two steps.

1. Update the `/etc/system` file of each cluster nodes to include the following line, then reboot each node to make the change take effect.

```
set halt_on_panic=1
```

2. On all cluster nodes, change the `eeeprom auto-boot?` flag to `false`.

```
eeeprom auto-boot?=false
```

Some shell systems use slightly different syntax. See the man page for the `eeeprom(1M)` command for the exact syntax to use for your shell system.

Bug ID 4410535

Problem Summary: When using the Sun Cluster module to Sun Management Center 3.0, you cannot add a previously deleted resource group.

Workaround:

1. Click on Resource Group->Status->Failover Resource Groups.
2. Right click on the resource group name to be deleted and select Delete Selected Resource Group.
3. Click on the Refresh icon and make sure the row corresponding to the deleted resource group is gone.
4. Right click on Resource Groups in the left pane and select Create New Resource Group.
5. Give the same resource group name that was deleted before and click Next. A dialog box pops up saying that the resource group name is already in use.

Known Documentation Problems

This section discusses documentation errors you might encounter and steps to correct these problems. This information is in addition to known documentation problems documented in the *Sun Cluster 3.0 U1 Release Notes*.

Hardware Guide

Step 3 of the procedure, “How to Upgrade Controller Module Firmware in a Running Cluster,” in the *Sun Cluster 3.0 U1 Hardware Guide* (page 164) describes erroneous syntax for the `scshutdown` command. The `scshutdown` command should *not* use the `-i` option. The correct syntax for the command should be as shown below.

```
# scshutdown -y -g0
```


Sun Management Center 3.0

This section describes undocumented information about the Sun Cluster 3.0 module of the Sun Management Center GUI. It includes information about installation of the Sun Management Center, as well as new functionality.

Accessing the New Sun Management Center Features

There are new locations within the Sun Cluster module console from which you can create, change the state of, or delete resources and resource groups. You can access them by opening the Sun Cluster Details window and choosing the options from the hierarchy (tree) or topology views.

- You can now access the resource and resource group creation wizards from the Resource Group subtree in the hierarchy (tree) view.

- You can now create, change the state of, or delete resources and resource groups from a pop-up menu available at each table shown below.

New Pop-Up Menu Items and Associated Tables

Access from the Resource Group Status table and Resource Group

Properties table:

- Bring Online
- Take Offline
- Delete Selected Resource Group
- Create New Resource Group
- Create New Resource

Access from the Resource Status table and Resource Configuration table:

- Enable
- Disable
- Delete Resource
- Create New Resource Group
- Create New Resource

How to Access the Creation Wizards From the Tree View

Complete the following steps to access the wizards to create a resource or resource group.

1. In either the hierarchy (tree) or topology view, double-click Operating System>Sun Cluster.
2. Click the right mouse button on the Resource Groups item, or on any item within the Resource Groups subtree.
3. Choose “Create New Resource Group” or “Create New Resource” from the pop-up menu.

How to Access the Resource and Resource Group Tables

Complete the following steps to access the resource and resource group tables indicated below.

Resource Group Status Table and Resource Status Table

1. In either the hierarchy (tree) or topology view, double-click Operating System>Sun Cluster>Resource Groups>Status.
2. Double-click either Failover Resource Groups or Scalable Resource Groups, to choose the type of resource group you want.

Resource Group Properties Tables

1. In either the hierarchy (tree) or topology view, double-click Operating System>Sun Cluster>Resource Groups>Configuration.
2. Double-click either Failover Resource Groups or Scalable Resource Groups, to choose the type of resource group you want.
3. Double-click either the failover or scalable Resource Group Configuration item to display the Resource Group Properties table for the desired type of resource group.

Resources Configuration Tables

1. In either the hierarchy (tree) or topology view, double-click Operating System>Sun Cluster>Resource Groups>Configuration.
2. Double-click either Failover Resource Groups or Scalable Resource Groups, to choose the type of resource group you want.
3. Double-click either the failover or scalable Resource Configuration item.
4. Double-click General to display the Resources Configuration table for the desired type of resource group

How to Access the New Menu Options

The new menu options allow you to create, change the state of, or delete resources and resource groups.

Creating Resources and Resource Groups

Complete the following procedure to use the creation wizard on the pop-up menus, accessible from the resource and resource group tables.

1. Display either the resource or resource group table.
2. Point to any cell entry in the table, excluding the header row.
3. Click the right mouse button.
4. Choose the action you want from the pop-up menu.

Deleting or Modifying Resources and Resource Groups

Complete the following procedure to alter the state of or to delete a resource or resource group using the pop-up menus from the resource and resource group tables.

1. Display either the resource or resource group table.
2. Select the item that you want to modify.
 - a. To delete an entry, select the resource or resource group to be deleted
 - b. To change the state of an entry, select the state cell in the row of the resource or resource group to be changed.
3. Click the right mouse button.
4. Choose from the pop-up menu one of the following tasks to perform.

- Bring Online
- Take Offline
- Enable
- Disable
- Delete Selected Resource Group
- Delete Resource

Note – When you delete or edit status of a resource or resource group, the Sun Cluster module launches a Probe Viewer window. If the Sun Cluster module successfully performs the task that you choose, the Probe Viewer window displays the message, “Probe command returned no data.” If the task is not completed successfully, the window displays an error message.

See Sun Management Center documentation and online help for more information about Sun Management Center.

Installing and Maintaining a Sun StorEdge T3 and T3+ Disk Tray Partner-Group Configuration

This chapter contains the procedures for installing, configuring, and maintaining Sun StorEdge™ T3 and Sun StorEdge T3+ disk trays in a partner-group (interconnected) configuration. Differences between the StorEdge T3 and StorEdge T3+ procedures are noted where appropriate.

This chapter contains the following procedures:

- “How to Install StorEdge T3/T3+ Disk Tray Partner Groups” on page 30
- “How to Create a Logical Volume” on page 37
- “How to Remove a Logical Volume” on page 39
- “How to Upgrade StorEdge T3/T3+ Disk Tray Firmware in a Running Cluster” on page 44
- “How to Replace a Failed Disk Drive in a Running Cluster” on page 62
- “How to Add StorEdge T3/T3+ Disk Tray Partner Groups to a Running Cluster” on page 46
- “How to Remove StorEdge T3/T3+ Disk Trays From a Running Cluster” on page 59
- “How to Replace a Node-to-Switch Component in a Running Cluster” on page 63
- “How to Replace a Controller Card, FC Switch, or Disk Tray-to-Switch Component in a Running Cluster” on page 64
- “How to Replace a Disk Tray Chassis in a Running Cluster” on page 66
- “How to Replace a Node’s Host Adapter in a Running Cluster” on page 68
- “How to Migrate From a Single-Controller Configuration to a Partner-Group Configuration” on page 70

For conceptual information on multihost disks, see the *Sun Cluster 3.0 U1 Concepts* document.

Installing Sun StorEdge T3/T3+ Disk Trays

Note – This section contains the procedure for an *initial installation* of StorEdge T3 or StorEdge T3+ disk tray partner groups in a new Sun Cluster that is not running. If you are adding partner groups to an existing cluster, use the procedure in “[How to Add StorEdge T3/T3+ Disk Tray Partner Groups to a Running Cluster](#)” on page 46.

▼ How to Install StorEdge T3/T3+ Disk Tray Partner Groups

Perform the steps in this procedure in conjunction with the procedures in the *Sun Cluster 3.0 U1 Installation Guide* and your server hardware manual.

1. Install the host adapters in the nodes that will be connected to the disk trays.

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

2. Install the Fibre Channel (FC) switches.

Note – You must use FC switches when installing disk trays in a partner-group configuration.

For the procedure on installing a Sun StorEdge network FC switch-8 or switch-16, see the *Sun StorEdge network FC switch-8 and switch-16 Installation and Configuration Guide*.

3. (Skip this step if you are installing StorEdge T3+ disk trays)

Install the media interface adapters (MIAs) in the StorEdge T3 disk trays you are installing as shown in [FIGURE B-1](#).

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

4. If necessary, install GBICs in the FC switches, as shown in [FIGURE B-1](#).

For the procedure on installing a GBIC to an FC switch, see the *Sun StorEdge network FC switch-8 and switch-16 Installation and Configuration Guide*.

5. Set up a Reverse Address Resolution Protocol (RARP) server on the network you want the new disk trays to reside on.

This RARP server enables you to assign an IP address to the new disk trays using the disk tray's unique MAC address. For the procedure on setting up a RARP server, see the *Sun StorEdge T3 and T3+ Array Installation, Operation, and Service Manual*.

6. Cable the disk trays (see [FIGURE B-1](#)):

- a. Connect the disk trays to the FC switches using fiber optic cables.**
- b. Connect the Ethernet cables from each disk tray to the LAN.**
- c. Connect interconnect cables between the two disk trays of each partner group.**
- d. Connect power cords to each disk tray.**

For the procedure on installing fiber optic, Ethernet, and interconnect cables, see the *Sun StorEdge T3 and T3+ Array Installation, Operation, and Service Manual*.

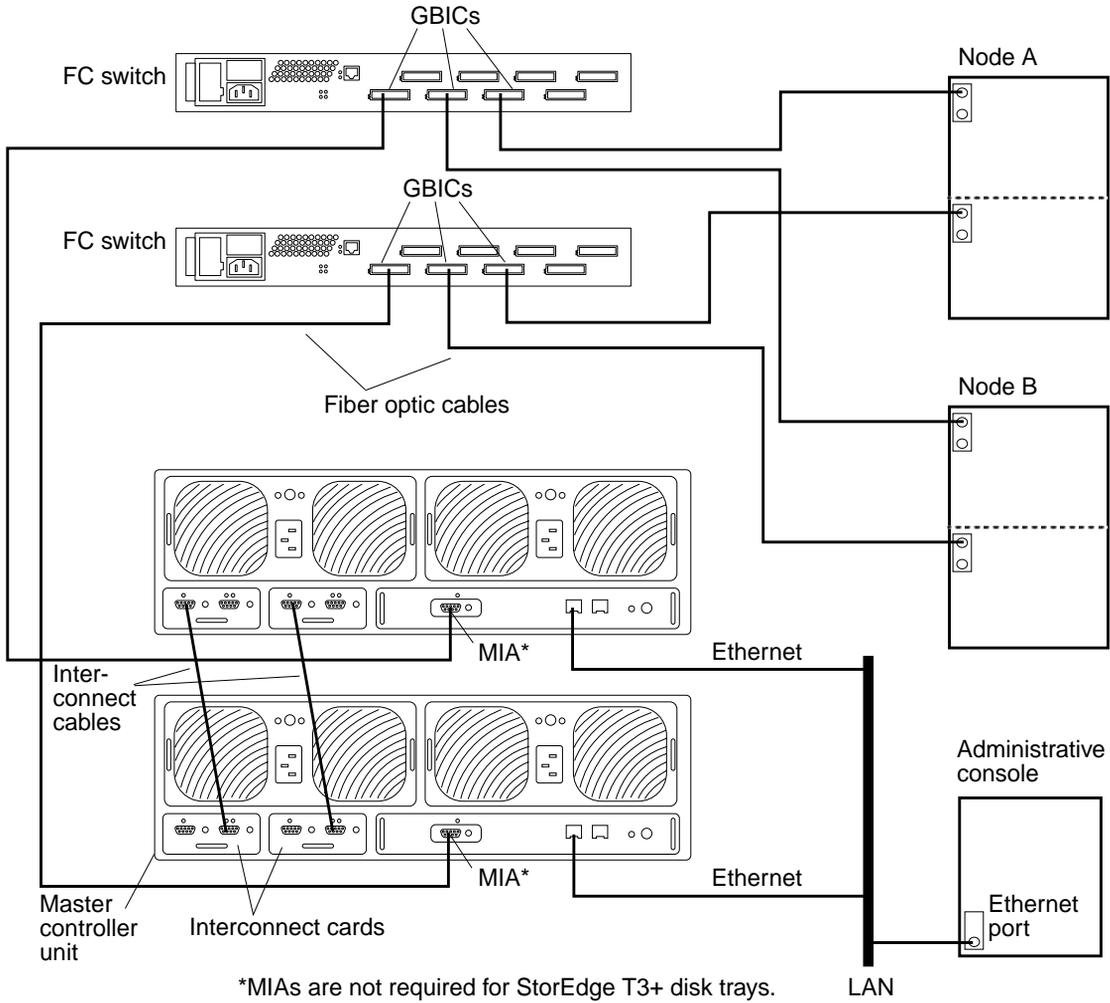


FIGURE B-1 StorEdge T3/T3+ Disk Tray Partner-Group (Interconnected) Controller Configuration

7. Power on the disk trays and verify that all components are powered on and functional.

For the procedure on powering on the disk trays and verifying the hardware configuration, see the *Sun StorEdge T3 and T3+ Array Installation, Operation, and Service Manual*.

8. Administer the disk trays' network settings:

Telnet to the *master* controller unit and administer the disk trays. For the procedure on administering the disk tray network addresses and settings, see the *Sun StorEdge T3 and T3+ Array Installation, Operation, and Service Manual*.

The master controller unit is the disk tray that has the interconnect cables attached to the right-hand connectors of its interconnect cards (when viewed from the rear of the disk trays). For example, [FIGURE B-1](#) shows the master controller unit of the partner-group as the lower disk tray. Note in this diagram that the interconnect cables are connected to the right-hand connectors of both interconnect cards on the master controller unit.

9. Install any required disk tray controller firmware:

For partner-group configurations, telnet to the *master* controller unit and install the required controller firmware.

See the *Sun Cluster 3.0 U1 Release Notes* for information about accessing Sun's EarlyNotifier web pages, which list information about any required patches or firmware levels that are available for download. For the procedure on applying any host adapter firmware patch, see the firmware patch `README` file.

10. At the master disk tray's prompt, use the `port list` command to ensure that each disk tray has a unique target address:

```
t3:/:<#> port list
```

If the disk trays do not have unique target addresses, use the `port set` command to set the addresses. For the procedure on verifying and assigning a target address to a disk tray, see the *Sun StorEdge T3 and T3+ Array Configuration Guide*. For more information about the `port` command see the *Sun StorEdge T3 and T3+ Array Administrator's Guide*.

- 11. At the master disk tray's prompt, use the `sys list` command to verify that the `cache` and `mirror` settings for each disk tray are set to `auto`:**

```
t3://:<#> sys list
```

If the two settings are not already set to `auto`, set them using the following commands:

```
t3://:<#> sys cache auto
t3://:<#> sys mirror auto
```

For more information about the `sys` command see the *Sun StorEdge T3 and T3+ Array Administrator's Guide*.

- 12. At the master disk tray's prompt, use the `sys list` command to verify that the `mp_support` parameter for each disk tray is set to `mpxio`:**

```
t3://:<#> sys list
```

If `mp_support` is not already set to `mpxio`, set it using the following command:

```
t3://:<#> sys mp_support mpxio
```

For more information about the `sys` command see the *Sun StorEdge T3 and T3+ Array Administrator's Guide*.

- 13. At the master disk tray's prompt, use the `sys stat` command to verify that both disk tray controllers are online, as shown in the following example output.**

```
t3://:<#> sys stat
Unit   State   Role   Partner
-----
 1     ONLINE  Master  2
 2     ONLINE  AlterM  1
```

For more information about the `sys` command and how to correct the situation if both controllers are not online, see the *Sun StorEdge T3 and T3+ Array Administrator's Guide*.

14. Configure the disk trays with the desired logical volumes.

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

15. Reset the disk trays.

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

16. Install to the cluster nodes the Solaris operating environment and apply the required Solaris patches for Sun Cluster software and StorEdge T3/T3+ disk tray support.

For the procedure on installing the Solaris operating environment, see the *Sun Cluster 3.0 U1 Installation Guide*.

See the *Sun Cluster 3.0 U1 Release Notes* for information about accessing Sun's EarlyNotifier web pages, which list information about any required patches or firmware levels that are available for download. For the procedure on applying any host adapter firmware patch, see the firmware patch README file.

17. Install to the cluster nodes any required patches or software for Sun StorEdge Traffic Manager software support from the Sun Download Center Web site,

<http://www.sun.com/storage/san/>

For instructions on installing the software, see the information on the web site.

18. Activate the Sun StorEdge Traffic Manager software functionality in the software you installed to the cluster nodes in [Step 17](#).

To activate the Sun StorEdge Traffic Manager software functionality, manually edit the `/kernel/drv/scsi_vhci.conf` file that is installed to change the `mpxio-disable` parameter to `no`:

```
mpxio-disable="no"
```

19. Perform a reconfiguration boot on all nodes to create the new Solaris device files and links.

```
{0} ok boot -r
```

20. On all nodes, update the `/devices` and `/dev` entries:

```
# devfsadm -C
```

Where to Go From Here

To continue with Sun Cluster software installation tasks, see the *Sun Cluster 3.0 U1 Installation Guide*.

Configuring StorEdge T3/T3+ Disk Trays in a Running Cluster

This section contains the procedures for configuring a StorEdge T3 or StorEdge T3+ disk tray in a running cluster. [TABLE B-1](#) lists these procedures.

TABLE B-1 Task Map: Configuring a StorEdge T3/T3+ Disk Tray

Task	For Instructions, Go To...
Create a logical volume	“How to Create a Logical Volume” on page 37
Remove a logical volume	“How to Remove a Logical Volume” on page 39

▼ How to Create a Logical Volume

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

1. Telnet to the disk tray that is the master controller unit of your partner-group.

The master controller unit is the disk tray that has the interconnect cables attached to the right-hand connectors of its interconnect cards (when viewed from the rear of the disk trays). For example, [FIGURE B-1](#) shows the master controller unit of the partner-group as the lower disk tray. Note in this diagram that the interconnect cables are connected to the right-hand connectors of both interconnect cards on the master controller unit.

2. Create the logical volume.

Creating a logical volume involves adding, initializing, and mounting the logical volume.

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

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

```
# devfsadm
```

4. On one node connected to the partner-group, use the `format` command to verify that the new logical volume is visible to the system.

```
# format
```

See the `format` command man page for more information about using the command.

5. Are you running VERITAS Volume Manager?

- If not, go to [Step 6](#)
- If you are running VERITAS Volume Manager, update its list of devices on all cluster nodes attached to the logical volume you created in [Step 2](#).

See your VERITAS Volume Manager documentation for information about using the `vxctl enable` command to update new devices (volumes) in your VERITAS Volume Manager list of devices.

6. If needed, partition the logical volume.

7. From any node in the cluster, update the global device namespace by using the `scgdevs` command.

```
# scgdevs
```

Note – 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.

Note – Do not configure StorEdge T3/T3+ logical volumes as quorum devices in partner-group configurations. The use of StorEdge T3/T3+ logical volumes as quorum devices in partner-group configurations is not supported.

Where to Go From Here

To create a new resource or reconfigure a running resource to use the new logical volume, see the *Sun Cluster 3.0 U1 Data Services Installation and Configuration Guide*.

▼ How to Remove a Logical Volume

Use this procedure to remove a StorEdge T3/T3+ disk tray logical volume. This procedure assumes all cluster nodes are booted and attached to the disk tray that hosts the logical volume you are removing.

This procedure defines “Node A” as the node you begin working with, and “Node B” as the other node.



Caution – This procedure removes all data from the logical volume you are removing.

- 1. If necessary, migrate all data and volumes off the logical volume you are removing.**
- 2. Are you running VERITAS Volume Manager?**
 - If not, go to [Step 3](#).
 - If you are running VERITAS Volume Manager, update its list of devices on all cluster nodes attached to the logical volume you are removing.

See your VERITAS Volume Manager documentation for information about using the `vxdisk rm` command to remove devices (volumes) in your VERITAS Volume Manager device list.

- 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 disk tray that is the master controller unit of your partner-group.**

The master controller unit is the disk tray that has the interconnect cables attached to the right-hand connectors of its interconnect cards (when viewed from the rear of the disk trays). For example, [FIGURE B-1](#) shows the master controller unit of the partner-group as the lower disk tray. Note in this diagram that the interconnect cables are connected to the right-hand connectors of both interconnect cards on the master controller unit.

- 5. Remove the logical volume.**

For the procedure on removing a logical volume, see the *Sun StorEdge T3 and T3+ Array Administrator's Guide*.

6. Use the `scstat` command to identify the resource groups and device groups running on all nodes.

Record this information because you will use it in [Step 15](#) of this procedure to return resource groups and device groups to these nodes.

```
# scstat
```

7. Move all resource groups and device groups off of Node A:

```
# scswitch -s -h nodename
```

8. Shut down and reboot Node A by using the `shutdown` command with the `i6` option.

The `-i6` option with the `shutdown` command causes the node to reboot after it shuts down to the `ok` prompt.

```
# shutdown -y -g0 -i6
```

For more information, see the *Sun Cluster 3.0 U1 System Administration Guide*.

9. On Node A, remove the obsolete device IDs (DIDs):

```
# devfsadm -C  
# scdidadm -C
```

10. On Node A, use the `format` command to verify that the logical volume you removed is no longer visible to the system.

```
# format
```

See the `format` command man page for more information about using the command.

11. Move all resource groups and device groups off Node B:

```
# scswitch -s -h nodename
```

12. **Shut down and reboot Node B by using the `shutdown` command with the `i6` option.**

The `-i6` option with the `shutdown` command causes the node to reboot after it shuts down to the `ok` prompt.

```
# shutdown -y -g0 -i6
```

For more information, see the *Sun Cluster 3.0 U1 System Administration Guide*.

13. **On Node B, remove the obsolete DIDs:**

```
# devfsadm -C
# scdidadm -C
```

14. **On Node B, use the `format` command to verify that the logical volume you removed is no longer visible to the system.**

```
# format
```

See the `format` command man page for more information about using the command.

15. **Return the resource groups and device groups you identified in [Step 6](#) to Node A and Node B:**

```
# scswitch -z -g resource-group -h nodename
# scswitch -z -D device-group-name -h nodename
```

For more information, see the *Sun Cluster 3.0 U1 System Administration Guide*.

Where to Go From Here

To create a logical volume, see [“How to Create a Logical Volume” on page 37](#).

Maintaining StorEdge T3/T3+ Disk Trays

This section contains the procedures for maintaining StorEdge T3 and StorEdge T3+ disk trays. [TABLE B-2](#) lists these procedures. This section does not include a procedure for adding a disk drive or a procedure for removing a disk drive because a StorEdge T3/T3+ disk tray operates only when fully configured.



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

TABLE B-2 Task Map: Maintaining a StorEdge T3/T3+ Disk Tray

Task	For Instructions, Go To...
Upgrade StorEdge T3/T3+ disk tray firmware.	“How to Upgrade StorEdge T3/T3+ Disk Tray Firmware in a Running Cluster” on page 44
Add a StorEdge T3/T3+ disk tray.	“How to Add StorEdge T3/T3+ Disk Tray Partner Groups to a Running Cluster” on page 46
Remove a StorEdge T3/T3+ disk tray.	“How to Remove StorEdge T3/T3+ Disk Trays From a Running Cluster” on page 59
Replace a disk drive in a disk tray.	“How to Replace a Failed Disk Drive in a Running Cluster” on page 62
Replace a node-to-switch fiber optic cable.	“How to Replace a Node-to-Switch Component in a Running Cluster” on page 63
Replace a gigabit interface converter (GBIC) on a node's host adapter.	“How to Replace a Node-to-Switch Component in a Running Cluster” on page 63
Replace a GBIC on an FC switch, connecting to a node.	“How to Replace a Node-to-Switch Component in a Running Cluster” on page 63
Replace a disk tray-to-switch fiber optic cable.	“How to Replace a Controller Card, FC Switch, or Disk Tray-to-Switch Component in a Running Cluster” on page 64
Replace a GBIC on an FC switch, connecting to a disk tray.	“How to Replace a Controller Card, FC Switch, or Disk Tray-to-Switch Component in a Running Cluster” on page 64

TABLE B-2 Task Map: Maintaining a StorEdge T3/T3+ Disk Tray (Continued)

Task	For Instructions, Go To...
Replace a StorEdge network FC switch-8 or switch-16.	“How to Replace a Controller Card, FC Switch, or Disk Tray-to-Switch Component in a Running Cluster” on page 64
Replace a StorEdge network FC switch-8 or switch-16 power cord.	“How to Replace a Controller Card, FC Switch, or Disk Tray-to-Switch Component in a Running Cluster” on page 64
Replace a media interface adapter (MIA) on a StorEdge T3 disk tray (not applicable for StorEdge T3+ disk trays).	“How to Replace a Controller Card, FC Switch, or Disk Tray-to-Switch Component in a Running Cluster” on page 64
Replace interconnect cables.	“How to Replace a Controller Card, FC Switch, or Disk Tray-to-Switch Component in a Running Cluster” on page 64
Replace a StorEdge T3/T3+ disk tray controller.	“How to Replace a Controller Card, FC Switch, or Disk Tray-to-Switch Component in a Running Cluster” on page 64
Replace a StorEdge T3/T3+ disk tray chassis.	“How to Replace a Disk Tray Chassis in a Running Cluster” on page 66
Replace a host adapter in a node.	“How to Replace a Node’s Host Adapter in a Running Cluster” on page 68
Migrate from a single-controller configuration to a partner-group configuration.	“How to Migrate From a Single-Controller Configuration to a Partner-Group Configuration” on page 70
Upgrade a StorEdge T3 disk tray controller to a StorEdge T3+ disk tray controller.	<i>Sun StorEdge T3 Array Controller Upgrade Manual</i>
Replace a Power and Cooling Unit (PCU). Follow the same procedure used in a non-cluster environment.	<i>Sun StorEdge T3 and T3+ Array Installation, Operation, and Service Manual</i>
Replace a unit interconnect card (UIC). Follow the same procedure used in a non-cluster environment.	<i>Sun StorEdge T3 and T3+ Array Installation, Operation, and Service Manual</i>
Replace a StorEdge T3/T3+ disk tray power cable. Follow the same procedure used in a non-cluster environment.	<i>Sun StorEdge T3 and T3+ Array Installation, Operation, and Service Manual</i>
Replace an Ethernet cable. Follow the same procedure used in a non-cluster environment.	<i>Sun StorEdge T3 and T3+ Array Installation, Operation, and Service Manual</i>

▼ How to Upgrade StorEdge T3/T3+ Disk Tray Firmware in a Running Cluster

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



Caution – This procedure requires that you reset the StorEdge T3/T3+ array in which you are upgrading firmware. If the partner groups contain data mirrored between separate arrays, upgrade only one array submirror at one time, as described in this procedure. If you upgrade corresponding submirrors simultaneously, you will lose access to the data. If your partner-group configurations do not have data mirrored between separate arrays, you must shutdown the cluster when upgrading firmware, as described in this procedure. If you upgrade the firmware in an array that does not contain mirrored data, you will lose access to the data.

Note – For all firmware, always read any README files that accompany the firmware for the latest information and special notes.

1. **Determine whether the partner-group is a submirror of a node's volume-manager volume.**
 - If it is a submirror, go to [Step 2](#).
 - If it is not a submirror, shut down the entire cluster, then go to [Step 4](#):

```
# scshutdown -y -g0
```

For the full procedure on shutting down a cluster, see the *Sun Cluster 3.0 U1 System Administration Guide*.

2. **On one node attached to the disk tray you are upgrading, detach that disk tray's submirrors.**

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

3. **Disconnect both disk tray-to-switch fiber optic cables from the two disk trays of the partner-group.**
4. **Apply the controller, disk drive, and UIC firmware patches.**

For the list of required StorEdge T3/T3+ disk tray patches, see the *Sun StorEdge T3 and T3+ Array Release Notes*. For the procedure on applying firmware patches, see the firmware patch README file. For the procedure on verifying the firmware level, see the *Sun StorEdge T3 and T3+ Array Installation, Operation, and Service Manual*.

5. Reset the disk trays.

For the procedure on resetting a disk tray, see the *Sun StorEdge T3 and T3+ Array Installation, Operation, and Service Manual*.

6. Use the StorEdge T3/T3+ `disable uencidctr` command to disable one of the disk tray controllers so that all logical volumes come under the control of the remaining controller.

```
t3:/:<#> disable uencidctr
```

See the *Sun StorEdge T3 and T3+ Array Administrator's Guide* for more information about the `disable` command.

7. Perform this step only if you disconnected cables in [Step 3](#): Reconnect both disk tray-to-switch fiber optic cables to the two disk trays of the partner-group.

8. On one node connected to the partner-group, use the `format` command to verify that the disk tray controllers are rediscovered by the node.

```
# format
```

9. Use the StorEdge T3/T3+ `enable uencidctr` command to enable the disk tray controller that you disabled in [Step 6](#).

```
t3:/:<#> enable uencidctr
```

10. Perform this step only if you were directed to detach the submirrors in [Step 2](#): Reattach the disk tray's submirrors to resynchronize them.

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

11. Did you shut down the cluster in [Step 1](#)?

- If not, you are done with this procedure.
- If you did shut down the cluster, boot all nodes back into the cluster.

```
ok boot
```

For the full procedure on booting nodes into the cluster, see the *Sun Cluster 3.0 U1 System Administration Guide*.

▼ How to Add StorEdge T3/T3+ Disk Tray Partner Groups to a Running Cluster

Note – Use this procedure to add new StorEdge T3/T3+ disk tray partner groups to a running cluster. To install partner groups to a new Sun Cluster that is not running, use the procedure in [“How to Install StorEdge T3/T3+ Disk Tray Partner Groups” on page 30](#).

This procedure defines “Node A” as the node you begin working with, and “Node B” as the second node.

1. **Set up a Reverse Address Resolution Protocol (RARP) server on the network you want the new disk trays to reside on, then assign an IP address to the new disk trays.**

Note – Assign an IP address to the *master* controller unit only. The master controller unit is the disk tray that has the interconnect cables attached to the right-hand connectors of its interconnect cards (see [FIGURE B-2](#)).

This RARP server lets you assign an IP address to the new disk trays using the disk tray’s unique MAC address. For the procedure on setting up a RARP server, see the *Sun StorEdge T3 and T3+ Array Installation, Operation, and Service Manual*.

2. **Install the Ethernet cable between the disk trays and the local area network (LAN) (see [FIGURE B-2](#)).**
3. **If not already installed, install interconnect cables between the two disk trays of each partner group (see [FIGURE B-2](#)).**

For the procedure on installing interconnect cables, see the *Sun StorEdge T3 and T3+ Array Installation, Operation, and Service Manual*.

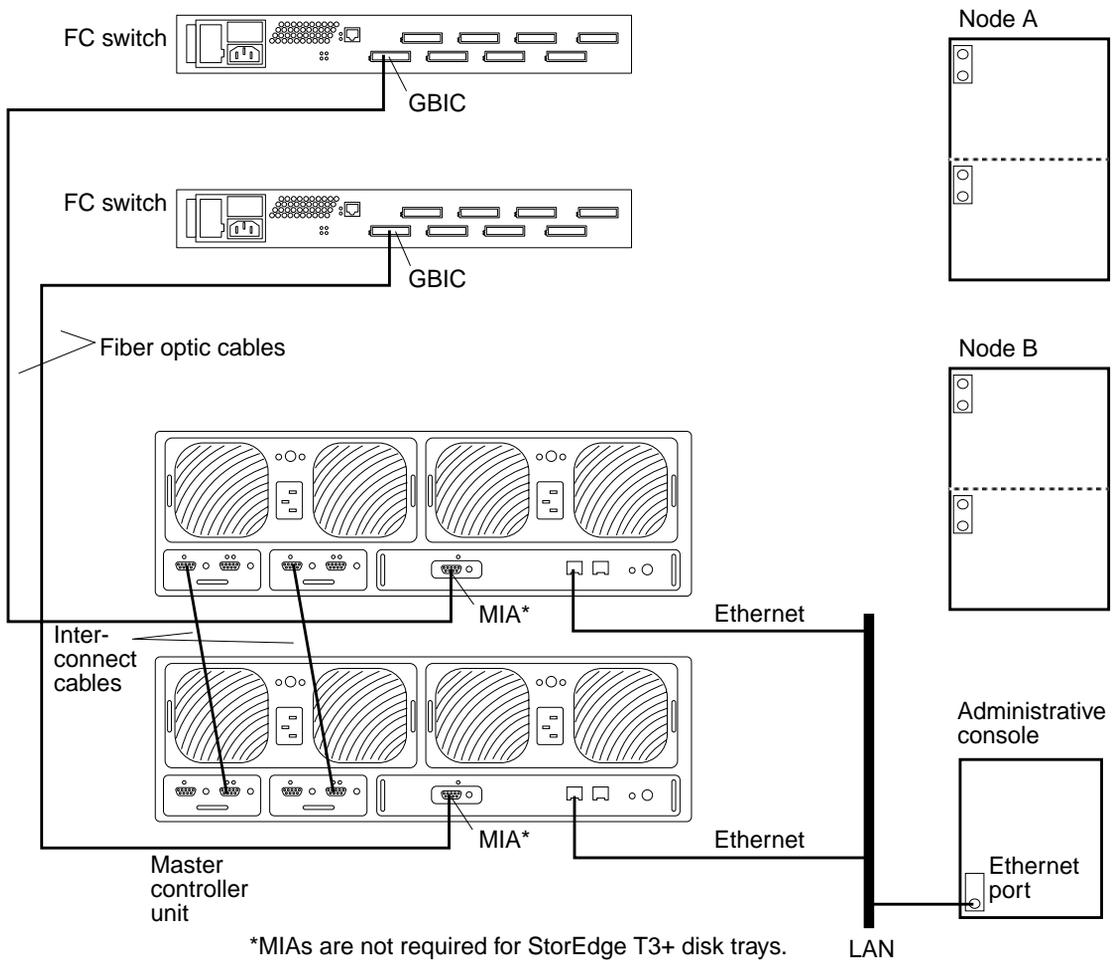


FIGURE B-2 Adding Sun StorEdge T3/T3+ Disk Trays, Partner-Group Configuration

4. Power on the disk trays.

Note – The disk trays might take several minutes to boot.

For the procedure on powering on disk trays, see the *Sun StorEdge T3 and T3+ Array Installation, Operation, and Service Manual*.

5. Administer the disk trays' network addresses and settings.

Telnet to the StorEdge T3/T3+ *master* controller unit and administer the disk trays.

For the procedure on administering disk tray network address and settings, see the *Sun StorEdge T3 and T3+ Array Installation, Operation, and Service Manual*.

6. Install any required disk tray controller firmware upgrades:

For partner-group configurations, telnet to the StorEdge T3/T3+ *master* controller unit and if necessary, install the required disk tray controller firmware.

For the required disk tray controller firmware revision number, see the *Sun StorEdge T3 and T3+ Array Release Notes*.

7. At the master disk tray's prompt, use the `port list` command to ensure that each disk tray has a unique target address:

```
t3:/:<#> port list
```

If the disk trays do not have unique target addresses, use the `port set` command to set the addresses. For the procedure on verifying and assigning a target address to a disk tray, see the *Sun StorEdge T3 and T3+ Array Configuration Guide*. For more information about the `port` command see the *Sun StorEdge T3 and T3+ Array Administrator's Guide*.

8. At the master disk tray's prompt, use the `sys list` command to verify that the `cache` and `mirror` settings for each disk tray are set to `auto`:

```
t3:/:<#> sys list
```

If the two settings are not already set to `auto`, set them using the following commands at each disk tray's prompt:

```
t3:/:<#> sys cache auto  
t3:/:<#> sys mirror auto
```

For more information about the `sys` command see the *Sun StorEdge T3 and T3+ Array Administrator's Guide*.

9. Use the `StorEdge T3/T3+ sys list` command to verify that the `mp_support` parameter for each disk tray is set to `mpxio`:

```
t3:/:<#> sys list
```

If `mp_support` is not already set to `mpxio`, set it using the following command at each disk tray's prompt:

```
t3:/:<#> sys mp_support mpxio
```

For more information about the `sys` command see the *Sun StorEdge T3 and T3+ Array Administrator's Guide*.

10. Configure the new disk trays with the desired logical volumes.

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

11. Reset the disk trays.

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

12. (Skip this step if you are adding StorEdge T3+ disk trays.)
Install the media interface adapter (MIA) in the StorEdge T3 disk trays you are adding, as shown in [FIGURE B-2](#).

For the procedure on installing an MIA, see the *Sun StorEdge T3 and T3+ Array Configuration Guide*.

13. If necessary, install GBICs in the FC switches, as shown in [FIGURE B-2](#).

Note – There are no FC switch port-assignment restrictions. You can connect your disk trays and nodes to any FC switch port.

For the procedure on installing a GBIC to an FC switch, see the *Sun StorEdge network FC switch-8 and switch-16 Installation and Configuration Guide*.

14. Install a fiber optic cable between the FC switch and the new disk tray as shown in [FIGURE B-2](#).

For the procedure on installing a fiber optic cable, see the *Sun StorEdge T3 and T3+ Array Installation, Operation, and Service Manual*.

15. Determine the resource groups and device groups running on all nodes.

Record this information because you will use it in [Step 54](#) of this procedure to return resource groups and device groups to these nodes.

```
# scstat
```

16. Move all resource groups and device groups off Node A.

```
# scswitch -s -h nodename
```

17. Do you need to install host adapters in Node A?

- If not, go to [Step 24](#).
- If you do need to install host adapters to Node A, continue with [Step 18](#).

18. Is the host adapter you are installing the first host adapter on Node A?

- If not, go to [Step 20](#).
- If it is the first host adapter, use the `pkginfo` command as shown below to determine whether the required support packages are already installed on this node. The following packages are required:

```
# 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
system    SUNWluxox  Sun Enterprise Network Array libraries (64-bit)
```

19. Are the required support packages already installed?

- If they are already installed, go to [Step 20](#).
- If not, install the required support packages that are missing.

The support packages are located in the `Product` directory of the Solaris CD-ROM. Use the `pkgadd` command to add any missing packages.

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

20. Shut down and power off Node A.

```
# shutdown -y -g0 -i0
```

For the procedure on shutting down and powering off a node, see the *Sun Cluster 3.0 U1 System Administration Guide*.

21. Install the host adapters in Node A.

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

22. Power on and boot Node A into non-cluster mode.

```
{0} ok boot -x
```

For more information, see the *Sun Cluster 3.0 U1 System Administration Guide*.

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

See the *Sun Cluster 3.0 U1 Release Notes* for information about accessing Sun's EarlyNotifier web pages, which list information about any required patches or firmware levels that are available for download. For the procedure on applying any host adapter firmware patch, see the firmware patch README file.

24. If necessary, install GBICs to the FC switches, as shown in [FIGURE B-3](#).

For the procedure on installing a GBIC to an FC switch, see the *Sun StorEdge network FC switch-8 and switch-16 Installation and Configuration Guide*.

25. Connect fiber optic cables between Node A and the FC switches, as shown in [FIGURE B-3](#).

Note – There are no FC switch port-assignment restrictions. You can connect your StorEdge T3/T3+ disk tray and node to any FC switch port.

For the procedure on installing a fiber optic cable, see the *Sun StorEdge T3 and T3+ Array Installation, Operation, and Service Manual*.

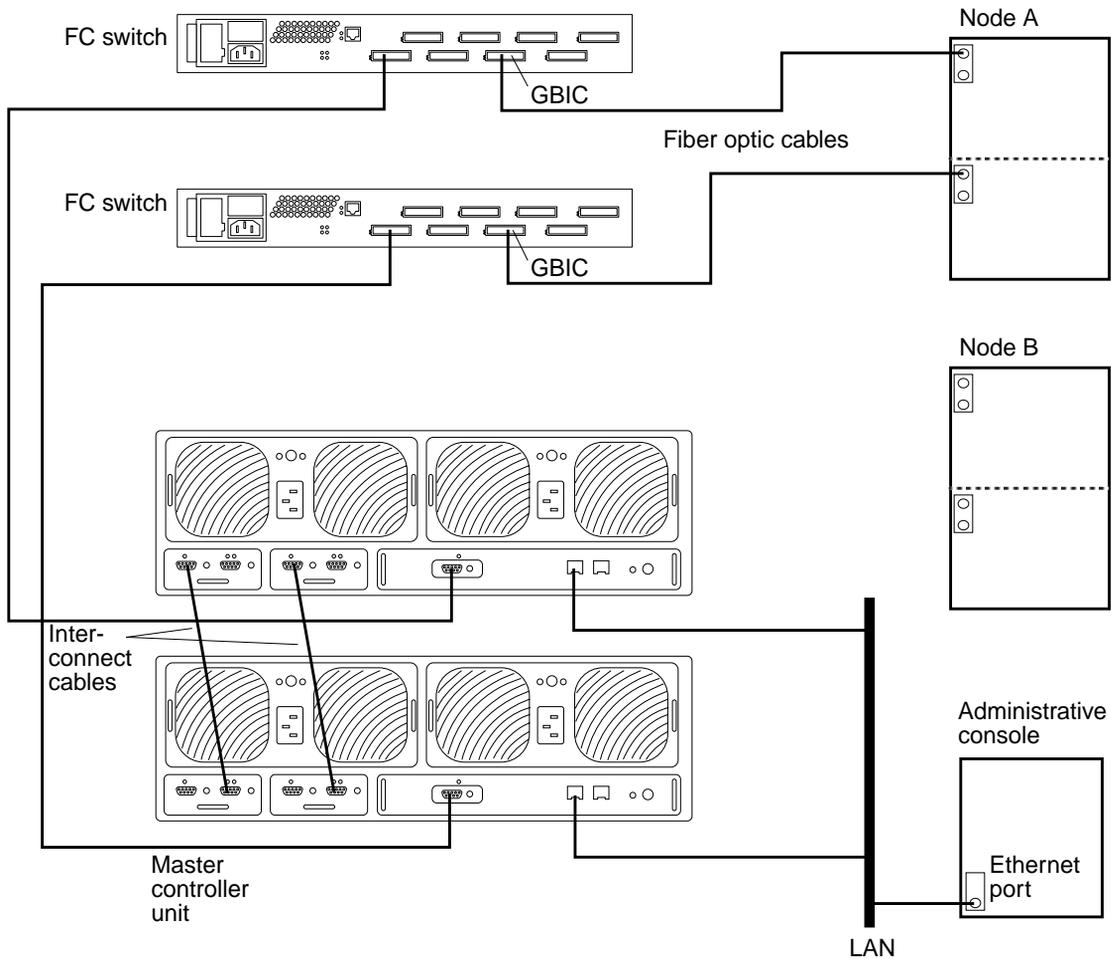


FIGURE B-3 Adding Sun StorEdge T3/T3+ Disk Trays, Partner-Group Configuration

26. If necessary, install the required Solaris patches for StorEdge T3/T3+ disk tray support on Node A.

See the *Sun Cluster 3.0 U1 Release Notes* for information about accessing Sun's EarlyNotifier web pages, which list information about any required patches or firmware levels that are available for download. For the procedure on applying any host adapter firmware patch, see the firmware patch README file.

27. Install any required patches or software for Sun StorEdge Traffic Manager software support to Node A from the Sun Download Center Web site,

<http://www.sun.com/storage/san/>

For instructions on installing the software, see the information on the web site.

28. Activate the Sun StorEdge Traffic Manager software functionality in the software you installed in [Step 27](#).

To activate the Sun StorEdge Traffic Manager software functionality, manually edit the `/kernel/drv/scsi_vhci.conf` file that is installed to change the `mpxio-disable` parameter to `no`:

```
mpxio-disable="no"
```

29. Shut down Node A.

```
# shutdown -y -g0 -i0
```

30. Perform a reconfiguration boot on Node A to create the new Solaris device files and links.

```
{0} ok boot -r
```

31. On Node A, update the `/devices` and `/dev` entries:

```
# devfsadm -C
```

32. On Node A, update the paths to the DID instances:

```
# scdidadm -C
```

33. Label the new disk tray logical volume.

For the procedure on labeling a logical volume, see the *Sun StorEdge T3 and T3+ Array Administrator's Guide*.

34. (Optional) On Node A, verify that the device IDs (DIDs) are assigned to the new disk tray.

```
# scdidadm -l
```

35. Do you need to install host adapters in Node B?

- If not, go to [Step 43](#).
- If you do need to install host adapters to Node B, continue with [Step 36](#).

36. Is the host adapter you are installing the first host adapter on Node B?

- If not, go to [Step 38](#).
- If it is the first host adapter, determine whether the required support packages are already installed on this node. The following packages are required.

```
# 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
system      SUNWluxox   Sun Enterprise Network Array libraries (64-bit)
```

37. Are the required support packages already installed?

- If they are already installed, go to [Step 38](#).
- If not, install the missing support packages.

The support packages are located in the `Product` directory of the Solaris CD-ROM. Use the `pkgadd` command to add any missing packages.

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

38. Move all resource groups and device groups off Node B.

```
# scswitch -s -h nodename
```

39. Shut down and power off Node B.

```
# shutdown -y -g0 -i0
```

For the procedure on shutting down and powering off a node, see the *Sun Cluster 3.0 U1 System Administration Guide*.

40. Install the host adapters in Node B.

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

41. Power on and boot Node B.

```
{0} ok boot -x
```

For more information, see the *Sun Cluster 3.0 U1 System Administration Guide*.

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

See the *Sun Cluster 3.0 U1 Release Notes* for information about accessing Sun's EarlyNotifier web pages, which list information about any required patches or firmware levels that are available for download. For the procedure on applying any host adapter firmware patch, see the firmware patch README file.

43. If necessary, install GBICs to the FC switches, as shown in [FIGURE B-4](#).

For the procedure on installing GBICs to an FC switch, see the *Sun StorEdge network FC switch-8 and switch-16 Installation and Configuration Guide*

44. Connect fiber optic cables between the FC switches and Node B as shown in [FIGURE B-4](#).

For the procedure on installing fiber optic cables, see the *Sun StorEdge T3 and T3+ Array Installation, Operation, and Service Manual*.

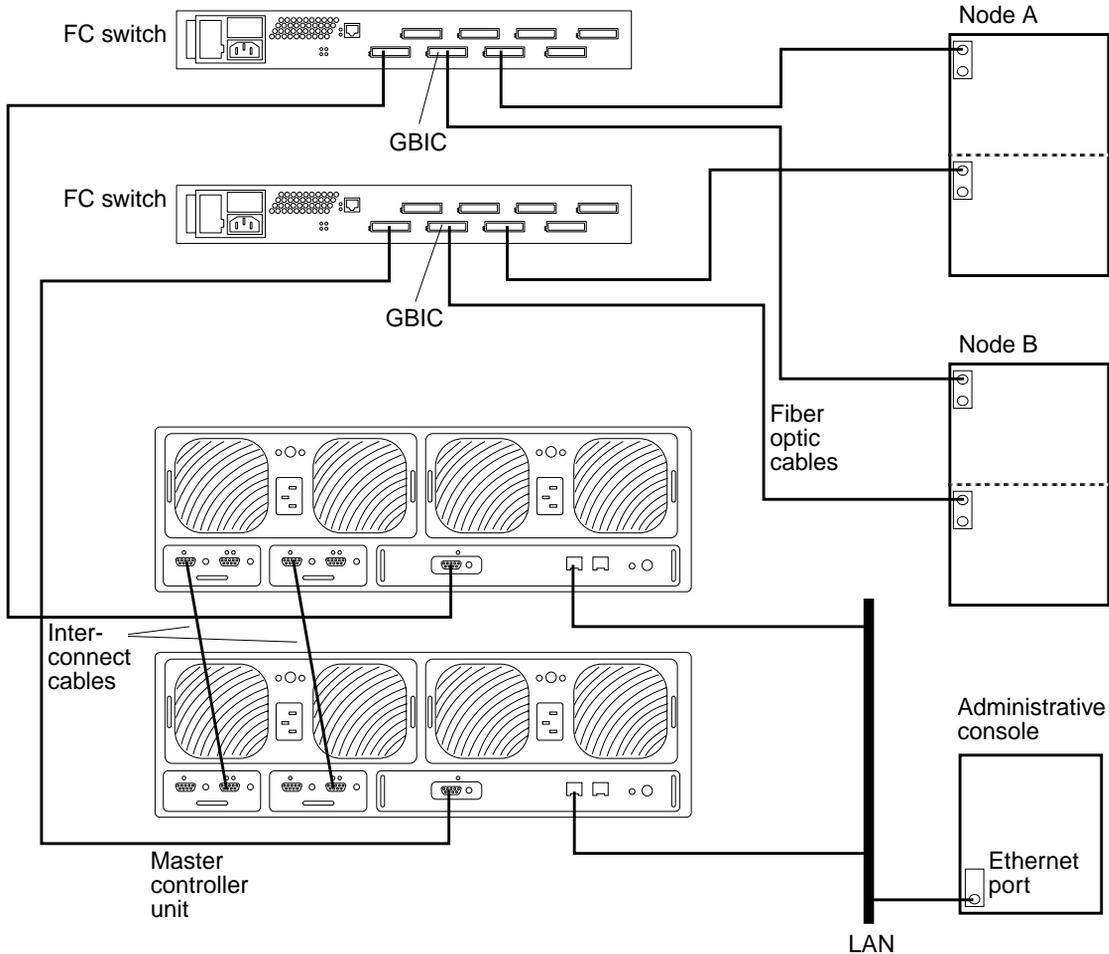


FIGURE B-4 Adding a Sun StorEdge T3/T3+ Disk Tray, Partner-Pair Configuration

45. If necessary, install the required Solaris patches for StorEdge T3/T3+ disk tray support on Node B.

For a list of required Solaris patches for StorEdge T3/T3+ disk tray support, see the *Sun StorEdge T3 and T3+ Array Release Notes*.

46. If you are installing a partner-group configuration, install any required patches or software for Sun StorEdge Traffic Manager software support to Node B from the Sun Download Center Web site, <http://www.sun.com/storage/san/>

For instructions on installing the software, see the information on the web site.

47. Activate the Sun StorEdge Traffic Manager software functionality in the software you installed in [Step 46](#).

To activate the Sun StorEdge Traffic Manager software functionality, manually edit the `/kernel/drv/scsi_vhci.conf` file that is installed to change the `mpxio-disable` parameter to `no`:

```
mpxio-disable="no"
```

48. Shut down Node B.

```
# shutdown -y -g0 -i0
```

49. Perform a reconfiguration boot to create the new Solaris device files and links on Node B.

```
{0} ok boot -r
```

50. On Node B, update the `/devices` and `/dev` entries:

```
# devfsadm -C
```

51. On Node B, update the paths to the DID instances:

```
# scdidadm -C
```

52. (Optional) On Node B, verify that the DIDs are assigned to the new disk trays:

```
# scdidadm -l
```

53. On one node attached to the new disk trays, reset the SCSI reservation state:

```
# scdidadm -R n
```

Where *n* is the DID instance of a disk tray LUN you are adding to the cluster.

Note – Repeat this command on the same node for each disk tray LUN you are adding to the cluster.

54. Return the resource groups and device groups you identified in [Step 15](#) to all nodes.

```
# scswitch -z -g resource-group -h nodename  
# scswitch -z -D device-group-name -h nodename
```

For more information, see the *Sun Cluster 3.0 U1 System Administration Guide*.

55. Perform volume management administration to incorporate the new logical volumes into the cluster.

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

▼ How to Remove StorEdge T3/T3+ Disk Trays From a Running Cluster

Use this procedure to permanently remove StorEdge T3/T3+ disk tray partner groups and their submirrors from a running cluster.

This procedure defines “Node A” as the cluster node you begin working with, and “Node B” as the other node.



Caution – During this procedure, you lose access to the data that resides on each disk tray partner-group you are removing.

- 1. If necessary, back up all database tables, data services, and volumes associated with each partner-group you are removing.**
- 2. If necessary, run the appropriate Solstice DiskSuite or VERITAS Volume Manager commands to detach the submirrors from each disk tray or partner-group that you are removing to stop all I/O activity to the disk tray or partner-group.**

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 references to each LUN that belongs to the disk tray or partner-group that you are removing.**

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

- 4. Determine the resource groups and device groups running on all nodes.**

Record this information because you will use it in [Step 21](#) of this procedure to return resource groups and device groups to these nodes.

```
# scstat
```

- 5. Move all resource groups and device groups off Node A.**

```
# scswitch -s -h nodename
```

6. Shut down Node A.

```
# shutdown -y -g0 -i0
```

For the procedure on shutting down a node, see the *Sun Cluster 3.0 U1 System Administration Guide*.

7. Disconnect from both disk trays the fiber optic cables connecting to the FC switches, then the Ethernet cable(s).

8. Is any disk tray you are removing the last disk tray connected to an FC switch on Node A?

- If not, go to [Step 12](#).
- If it is the last disk tray, disconnect the fiber optic cable between Node A and the FC switch that was connected to this disk tray.

For the procedure on removing a fiber optic cable, see the *Sun StorEdge T3 and T3+ Array Configuration Guide*.

9. Do you want to remove the host adapters from Node A?

- If not, go to [Step 12](#).
- If yes, power off Node A.

10. Remove the host adapters from Node A.

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

11. Without allowing the node to boot, power on Node A.

For more information, see the *Sun Cluster 3.0 U1 System Administration Guide*.

12. Boot Node A into cluster mode.

```
{0} ok boot
```

13. Move all resource groups and device groups off Node B.

```
# scswitch -s -h nodename
```

14. Shut down Node B.

```
# shutdown -y -g0 -i0
```

15. Is any disk tray you are removing the last disk tray connected to an FC switch on Node B?

- If not, go to [Step 16](#).
- If it is the last disk tray, disconnect the fiber optic cable connecting this FC switch to Node B.

For the procedure on removing a fiber optic cable, see the *Sun StorEdge T3 and T3+ Array Configuration Guide*.

16. Do you want to remove the host adapters from Node B?

- If not, go to [Step 19](#).
- If yes, power off Node B.

17. Remove the host adapters from Node B.

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

18. Without allowing the node to boot, power on Node B.

For more information, see the *Sun Cluster 3.0 U1 System Administration Guide*.

19. Boot Node B into cluster mode.

```
{0} ok boot
```

For more information, see the *Sun Cluster 3.0 U1 System Administration Guide*.

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

```
# devfsadm -C  
# sddidadm -C
```

21. Return the resource groups and device groups you identified in [Step 4](#) to all nodes.

```
# scswitch -z -g resource-group -h nodename  
# scswitch -z -D device-group-name -h nodename
```

▼ How to Replace a Failed Disk Drive in a Running Cluster

Use this procedure to replace one failed disk drive in a StorEdge T3/T3+ disk tray 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 StorEdge T3/T3+ disk tray is designed so that an orderly shutdown occurs when you remove a component for longer than 30 minutes. Therefore, a replacement part must be immediately available before starting an FRU replacement procedure. You must replace an FRU within 30 minutes or the StorEdge T3/T3+ disk tray, and all attached StorEdge T3/T3+ disk trays, will shut down and power off.

1. Replace the disk drive in the disk tray.

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

2. Perform volume management administration to configure logical volumes for the new disk tray into the cluster.

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

▼ How to Replace a Node-to-Switch Component in a Running Cluster

Use this procedure to replace the following node-to-switch components in a running cluster:

- Node-to-switch fiber optic cable
- GBIC on an FC switch, connecting to a node

1. On the node connected to the component you are replacing, determine the resource groups and device groups running on the node.

Record this information because you will use it in [Step 4](#) of this procedure to return resource groups and device groups to these nodes.

```
# scstat
```

2. Move all resource groups and device groups to another node.

```
# scswitch -s -h nodename
```

3. Replace the node-to-switch component.

- For the procedure on replacing a fiber optic cable between a node and an FC switch, see the *Sun StorEdge network FC switch-8 and switch-16 Installation and Configuration Guide*.
- For the procedure on replacing a GBIC on an FC switch, see the *Sun StorEdge network FC switch-8 and switch-16 Installation and Configuration Guide*.

4. Return the resource groups and device groups you identified in [Step 1](#) to the node that is connected to the component you replaced.

```
# scswitch -z -g resource-group -h nodename  
# scswitch -z -D device-group-name -h nodename
```

▼ How to Replace a Controller Card, FC Switch, or Disk Tray-to-Switch Component in a Running Cluster

Use this procedure to replace a failed disk tray controller card, an FC switch, or the following disk tray-to-switch components in a running cluster:

- Fiber optic cable that connects an FC switch to a disk tray
- GBIC on an FC switch, connecting to a disk tray
- StorEdge network FC switch-8 or switch-16
- FC switch power cord
- Media Interface Adapter (MIA) on a StorEdge T3 disk tray (not applicable for StorEdge T3+ disk trays)
- Interconnect cables between two interconnected disk trays of a partner-group

1. **Telnet to the disk tray that is connected to the controller card, FC switch, or component that you are replacing.**
2. **Use the T3/T3+ `sys stat` command to view the controller status for the two disk trays of the partner group.**

In the following example, both controllers are ONLINE.

```
t3:/:<#> sys stat
Unit   State      Role      Partner
-----
 1     ONLINE    Master    2
 2     ONLINE    AlterM   1
```

See the *Sun StorEdge T3 and T3+ Array Administrator's Guide* for more information about the `sys stat` command.

3. **Is the controller card, FC switch, or component that you are replacing attached to a disk tray controller that is ONLINE or DISABLED, as determined in [Step 2](#)?**

- If the controller is already DISABLED, go to [Step 5](#).
- If the controller is ONLINE, use the T3/T3+ `disable` command to disable it. Using the example from [Step 2](#), if you want to disable Unit 1, enter the following:

```
t3:/:<#> disable u1
```

See the *Sun StorEdge T3 and T3+ Array Administrator's Guide* for more information about the `disable` command.

4. Use the T3/T3+ `sys stat` command to verify that the controller's state has been changed to DISABLED.

```
t3:::<#> sys stat
Unit   State      Role      Partner
-----
1      DISABLED   Slave
2      ONLINE     Master
```

5. Replace the component using the following references:

- For the procedure on replacing a disk tray controller card, see the *Sun StorEdge T3 and T3+ Array Installation, Operation, and Service Manual*.
- For the procedure on replacing a fiber optic cable between a disk tray and an FC switch, see the *Sun StorEdge network FC switch-8 and switch-16 Installation and Configuration Guide*.
- For the procedure on replacing a GBIC on an FC switch, see the *Sun StorEdge network FC switch-8 and switch-16 Installation and Configuration Guide*.
- For the procedure on replacing a StorEdge network FC switch-8 or switch-16, or an FC switch power cord, see the *Sun StorEdge network FC switch-8 and switch-16 Installation and Configuration Guide*.
- For the procedure on replacing an MIA, see the *Sun StorEdge T3 and T3+ Array Installation, Operation, and Service Manual*.
- For the procedure on replacing interconnect cables, see the *Sun StorEdge T3 and T3+ Array Installation, Operation, and Service Manual*.

6. If necessary, telnet to the disk tray of the partner group that is still online.

7. Use the T3/T3+ `enable` command to reenble the disk tray that you disabled in [Step 3](#).

```
t3:::<#> enable u1
```

See the *Sun StorEdge T3 and T3+ Array Administrator's Guide* for more information about the `enable` command.

8. Use the T3/T3+ `sys stat` command to verify that the controller's state has been changed to ONLINE.

```
t3:::<#> sys stat
Unit   State      Role      Partner
-----
1      ONLINE     AlterM    2
2      ONLINE     Master    1
```

▼ How to Replace a Disk Tray Chassis in a Running Cluster

Use this procedure to replace a StorEdge T3/T3+ disk tray chassis in a running cluster. 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 Sun service providers should use this procedure to replace a StorEdge T3/T3+ disk tray chassis. This procedure requires the *Sun StorEdge T3 and T3+ Array Field Service Manual*, which is available to trained Sun service providers only.

1. Detach the submirrors on the disk tray that is connected to the chassis you are replacing to stop all I/O activity to this disk tray.

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

2. Are the disk trays in your partner-pair configuration made redundant by host-based mirroring?

- If yes, go to [Step 3](#).
- If not, shutdown the cluster.

```
# scshutdown -y -g0
```

3. Replace the chassis/backplane.

For the procedure on replacing a StorEdge T3/T3+ chassis, see the *Sun StorEdge T3 and T3+ Array Field Service Manual*. (This manual is available to trained Sun service providers only.)

4. Did you shut down the cluster in [Step 2](#)?

- If not, go to [Step 5](#).
- If you did shut down the cluster, boot it back into cluster mode.

```
{0} ok boot
```

5. Reattach the submirrors you detached in [Step 1](#) to resynchronize them.

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

6. On all nodes attached to the disk tray, use the following command to make the nodes' volumes recognize the new WWNs.

```
# drvconfig; disks; devlinks
```

7. On one node, use the `format` command to verify that the new logical volume is visible to the system.

```
# format
```

See the `format` command man page for more information about using the command.

▼ How to Replace a Node's Host Adapter in a Running Cluster

Use this procedure to replace a failed host adapter in a running cluster. As defined in this procedure, “Node A” is the node with the failed host adapter you are replacing and “Node B” is the other node.

1. Determine the resource groups and device groups running on all nodes.

Record this information because you will use it in [Step 8](#) of this procedure to return resource groups and device groups to these nodes.

```
# scstat
```

2. Move all resource groups and device groups off Node A.

```
# scswitch -s -h nodename
```

3. Shut down Node A.

```
# shutdown -y -g0 -i0
```

4. Power off Node A.

For more information, see the *Sun Cluster 3.0 U1 System Administration Guide*.

5. Replace the failed host adapter.

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

6. Power on Node A.

For more information, see the *Sun Cluster 3.0 U1 System Administration Guide*.

7. Boot Node A into cluster mode.

```
{0} ok boot
```

For more information, see the *Sun Cluster 3.0 U1 System Administration Guide*.

8. Return the resource groups and device groups you identified in [Step 1](#) to all nodes.

```
# scswitch -z -g resource-group -h nodename  
# scswitch -z -D device-group-name -h nodename
```

For more information, see the *Sun Cluster 3.0 U1 System Administration Guide*.

How to Migrate From a Single-Controller Configuration to a Partner-Group Configuration

Use this procedure to migrate your StorEdge T3/T3+ disk trays from a single-controller (non-interconnected) configuration to a partner-group (interconnected) configuration.

Note – Only trained, qualified Sun service providers should use this procedure. This procedure requires the *Sun StorEdge T3 and T3+ Array Field Service Manual*, which is available to trained Sun service providers only.

1. Remove the non-interconnected disk trays that will be in your partner-group from the cluster configuration.

Follow the procedure in [“How to Remove StorEdge T3/T3+ Disk Trays From a Running Cluster”](#) on page 59.

Note – Backup all data on the disk trays before removing them from the cluster configuration.

Note – This procedure assumes that the two disk trays that will be in the partner-group configuration are correctly isolated from each other on separate FC switches. You must use FC switches when installing disk trays in a partner-group configuration. Do not disconnect the cables from the FC switches or nodes.

2. Connect the single disk trays to form a partner-group.

Follow the procedure in the *Sun StorEdge T3 and T3+ Array Field Service Manual*.

3. Add the new partner-group to the cluster configuration:

- a. At each disk tray’s prompt, use the `port list` command to ensure that each disk tray has a unique target address:**

```
t3:/:<#> port list
```

If the disk trays do not have unique target addresses, use the `port set` command to set the addresses. For the procedure on verifying and assigning a target address to a disk tray, see the *Sun StorEdge T3 and T3+ Array Configuration Guide*. For more information about the `port` command see the *Sun StorEdge T3 and T3+ Array Administrator’s Guide*.

- b. At each disk tray's prompt, use the `sys list` command to verify that the `cache` and `mirror` settings for each disk tray are set to `auto`:**

```
t3://<#> sys list
```

If the two settings are not already set to `auto`, set them using the following commands at each disk tray's prompt:

```
t3://<#> sys cache auto
t3://<#> sys mirror auto
```

- c. Use the StorEdge T3/T3+ `sys list` command to verify that the `mp_support` parameter for each disk tray is set to `mpxio`:**

```
t3://<#> sys list
```

If `mp_support` is not already set to `mpxio`, set it using the following command at each disk tray's prompt:

```
t3://<#> sys mp_support mpxio
```

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

See the *Sun Cluster 3.0 U1 Release Notes* for information about accessing Sun's EarlyNotifier web pages, which list information about any required patches or firmware levels that are available for download. For the procedure on applying any host adapter firmware patch, see the firmware patch `README` file.

- e. If necessary, install the required Solaris patches for StorEdge T3/T3+ disk tray support on Node A.**

See the *Sun Cluster 3.0 U1 Release Notes* for information about accessing Sun's EarlyNotifier web pages, which list information about any required patches or firmware levels that are available for download.

- f. Install any required patches or software for Sun StorEdge Traffic Manager software support from the Sun Download Center Web site,**

<http://www.sun.com/storage/san/>

For instructions on installing the software, see the information on the web site.

g. Activate the Sun StorEdge Traffic Manager software functionality in the software you installed in Step f.

To activate the Sun StorEdge Traffic Manager software functionality, manually edit the `/kernel/drv/scsi_vhci.conf` file that is installed to change the `mpxio-disable` parameter to `no`:

```
mpxio-disable="no"
```

h. Shut down Node A.

```
# shutdown -y -g0 -i0
```

i. Perform a reconfiguration boot on Node A to create the new Solaris device files and links.

```
{0} ok boot -r
```

j. On Node A, update the `/devices` and `/dev` entries:

```
# devfsadm -C
```

k. On Node A, update the paths to the DID instances:

```
# scdidadm -C
```

l. Configure the new disk trays with the desired logical volumes.

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

m. Label the new disk tray logical volume.

For the procedure on labeling a logical volume, see the *Sun StorEdge T3 and T3+ Array Administrator's Guide*.

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

See the *Sun Cluster 3.0 U1 Release Notes* for information about accessing Sun's EarlyNotifier web pages, which list information about any required patches or firmware levels that are available for download. For the procedure on applying any host adapter firmware patch, see the firmware patch `README` file.

- o. If necessary, install the required Solaris patches for StorEdge T3/T3+ disk tray support on Node B.**

For a list of required Solaris patches for StorEdge T3/T3+ disk tray support, see the *Sun StorEdge T3 and T3+ Array Release Notes*.

- p. Install any required patches or software for Sun StorEdge Traffic Manager software support from the Sun Download Center Web site,**

<http://www.sun.com/storage/san/>

For instructions on installing the software, see the information on the web site.

- q. Activate the Sun StorEdge Traffic Manager software functionality in the software you installed in [Step p](#).**

To activate the Sun StorEdge Traffic Manager software functionality, manually edit the `/kernel/drv/scsi_vhci.conf` file that is installed to change the `mpxio-disable` parameter to `no`:

```
mpxio-disable="no"
```

- r. Shut down Node B.**

```
# shutdown -y -g0 -i0
```

- s. Perform a reconfiguration boot to create the new Solaris device files and links on Node B.**

```
{0} ok boot -r
```

- t. On Node B, update the `/devices` and `/dev` entries:**

```
# devfsadm -C
```

- u. On Node B, update the paths to the DID instances:**

```
# scdidadm -C
```

- v. (Optional) On Node B, verify that the DIDs are assigned to the new disk trays:**

```
# scdidadm -l
```

w. On one node attached to the new disk trays, reset the SCSI reservation state:

```
# sctidadm -R n
```

Where *n* is the DID instance of a disk tray LUN you are adding to the cluster.

Note – Repeat this command on the same node for each disk tray LUN you are adding to the cluster.

x. **Perform volume management administration to incorporate the new logical volumes into the cluster.**

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

Installing and Maintaining a Sun StorEdge T3 or T3+ Disk Tray Single-Controller Configuration

This chapter contains the procedures for installing, configuring, and maintaining Sun StorEdge™ T3 and Sun StorEdge T3+ disk trays in a single-controller (non-interconnected) configuration. Differences between the StorEdge T3 and StorEdge T3+ procedures are noted where appropriate.

This chapter contains the following procedures:

- [“How to Install StorEdge T3/T3+ Disk Trays” on page 76](#)
- [“How to Create a Sun StorEdge T3/T3+ Disk Tray Logical Volume” on page 80](#)
- [“How to Remove a Sun StorEdge T3/T3+ Disk Tray Logical Volume” on page 82](#)
- [“How to Upgrade StorEdge T3/T3+ Disk Tray Firmware” on page 87](#)
- [“How to Replace a Disk Drive” on page 88](#)
- [“How to Add a StorEdge T3/T3+ Disk Tray” on page 89](#)
- [“How to Remove a StorEdge T3/T3+ Disk Tray” on page 99](#)
- [“How to Replace a Host-to-Hub Component” on page 102](#)
- [“How to Replace a Sun StorEdge FC-100 Hub and Hub-to-Disk Tray Component” on page 103](#)
- [“How to Replace a StorEdge T3/T3+ Disk Tray Controller” on page 104](#)
- [“How to Replace a StorEdge T3/T3+ Disk Tray Chassis” on page 105](#)
- [“How to Replace a Host Adapter” on page 106](#)

For conceptual information on multihost disks, see the *Sun Cluster 3.0 U1 Concepts* document.

Installing StorEdge T3/T3+ Disk Trays

This section contains the procedure for an initial installation of new StorEdge T3 or StorEdge T3+ disk trays.

▼ How to Install StorEdge T3/T3+ Disk Trays

Use this procedure to install and configure new StorEdge T3 or StorEdge T3+ disk trays in a cluster that is not running. Perform the steps in this procedure in conjunction with the procedures in the *Sun Cluster 3.0 U1 Installation Guide* and your server hardware manual.

1. Install the host adapters in the nodes that are to be connected to the StorEdge T3/T3+ disk trays.

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 Sun StorEdge FC-100 hubs, see the *FC-100 Hub Installation and Service Manual*.

3. Set up a Reverse Address Resolution Protocol (RARP) server on the network you want the new StorEdge T3/T3+ disk trays to reside on.

This RARP server enables you to assign an IP address to the new StorEdge T3/T3+ disk trays by using each StorEdge T3/T3+ disk tray's unique MAC address.

For the procedure on setting up a RARP server, see the *Sun StorEdge T3 Installation, Operation, and Service Manual*.

4. (Skip this step if you are installing a StorEdge T3+ disk tray) Install the media interface adapters (MIAs) in the StorEdge T3 disk trays you are installing, as shown in [FIGURE C-1](#).

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

5. If necessary, install gigabit interface converters (GBICs) in the Sun StorEdge FC-100 hubs, as shown in [FIGURE C-1](#).

The GBICs let you connect the Sun StorEdge FC-100 hubs to the StorEdge T3/T3+ disk trays you are installing. For the procedure on installing an FC-100 hub GBIC, see the *FC-100 Hub Installation and Service Manual*.

6. **Install fiber-optic cables between the Sun StorEdge FC-100 hubs and the StorEdge T3/T3+ disk trays as shown in [FIGURE C-1](#).**

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

7. **Install fiber-optic cables between the Sun StorEdge FC-100 hubs and the cluster nodes as shown in [FIGURE C-1](#).**
8. **Install the Ethernet cables between the StorEdge T3/T3+ disk trays and the Local Area Network (LAN), as shown in [FIGURE C-1](#).**
9. **Install power cords to each disk tray you are installing.**
10. **Power on the StorEdge T3/T3+ disk trays and confirm that all components are powered on and functional.**

Note – The StorEdge T3/T3+ disk trays might require a few minutes to boot.

For the procedure on powering on a StorEdge T3/T3+ disk tray, see the *Sun StorEdge T3 Installation, Operation, and Service Manual*.

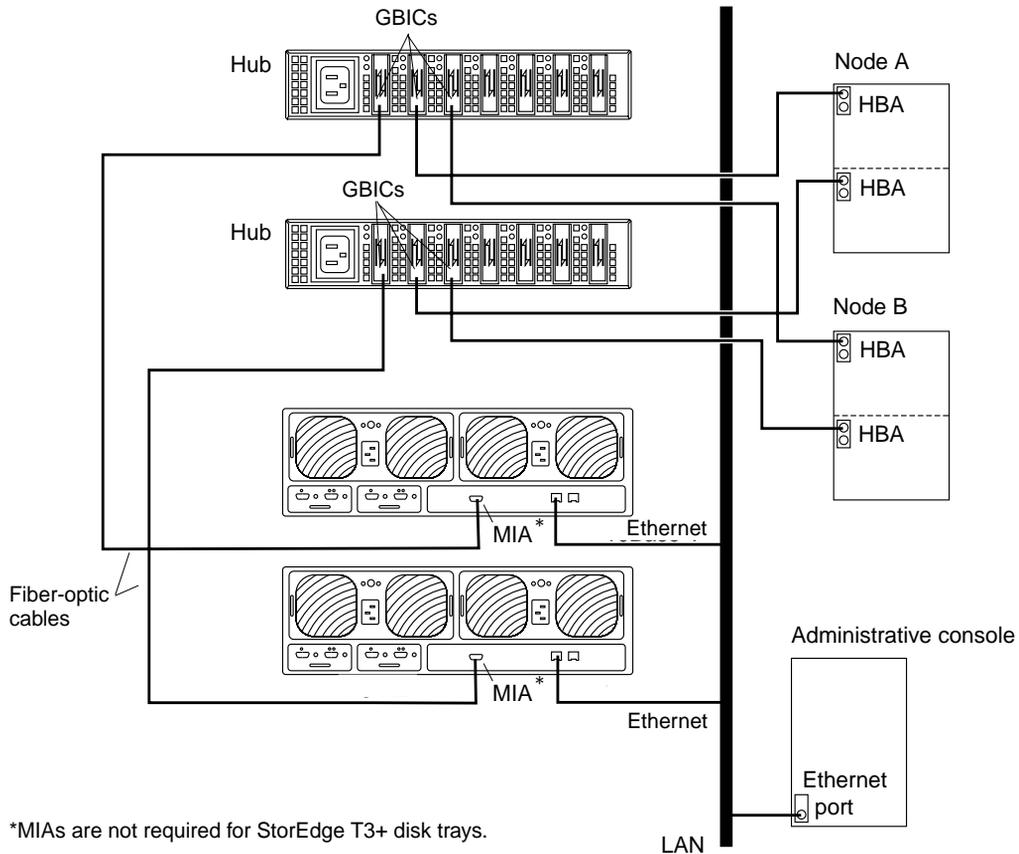


FIGURE C-1 Cabling StorEdge T3/T3+ Disk Trays in a Single-Controller Configuration

Note – Although [FIGURE C-1](#) shows a single-controller configuration, two disk trays are shown to illustrate how two non-interconnected disk trays are typically cabled in a cluster to allow data sharing and host-based mirroring.

11. (Optional) Configure the StorEdge T3/T3+ disk trays with logical volumes.

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

- 12. Telnet to each StorEdge T3/T3+ disk tray you are adding and install the required StorEdge T3/T3+ disk tray controller firmware.**

See the *Sun Cluster 3.0 U1 Release Notes* for information about accessing Sun's EarlyNotifier web pages, which list information about any required patches or firmware levels that are available for download. For the procedure on applying any host adapter firmware patch, see the firmware patch README file.

- 13. Ensure that this new StorEdge T3/T3+ disk tray has a unique target address.**

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

- 14. Reset the StorEdge T3/T3+ disk tray.**

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

- 15. Install to the cluster nodes the Solaris operating environment, and apply any required Solaris patches for Sun Cluster software and StorEdge T3/T3+ disk tray support.**

For the procedure on installing the Solaris operating environment, see the *Sun Cluster 3.0 U1 Installation Guide*. For the location of required Solaris patches and installation instructions for Sun Cluster software support, see the *Sun Cluster 3.0 U1 Release Notes*. For a list of required Solaris patches for StorEdge T3/T3+ disk tray support, see the *Sun StorEdge T3 Disk Tray Release Notes*.

Where to Go From Here

To continue with Sun Cluster software installation tasks, see the *Sun Cluster 3.0 U1 Installation Guide*.

Configuring a StorEdge T3/T3+ Disk Tray

This section contains the procedures for configuring a StorEdge T3 or StorEdge T3+ disk tray in a running cluster. The following table lists these procedures.

TABLE C-1 Task Map: Configuring a StorEdge T3/T3+ Disk Tray

Task	For Instructions, Go To
Create a disk tray logical volume	“How to Create a Sun StorEdge T3/T3+ Disk Tray Logical Volume” on page 80
Remove a disk tray logical volume	“How to Remove a Sun StorEdge T3/T3+ Disk Tray Logical Volume” on page 82

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

Use this procedure to create a logical volume. This procedure assumes all cluster nodes are booted and attached to the StorEdge T3/T3+ disk tray that is to host the logical volume you are creating.

- 1. Telnet to the StorEdge T3/T3+ disk tray that is to host the logical volume you are creating.**
- 2. Create the logical volume.**

The creation of a logical volume involves adding, mounting, and initializing the logical volume.

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

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

```
# devfsadm
```

4. On one node connected to the partner-group, use the `format` command to verify that the new logical volume is visible to the system.

```
# format
```

See the `format` command man page for more information about using the command.

5. Are you running VERITAS Volume Manager?

- If not, go to [Step 6](#)
- If you are running VERITAS Volume Manager, update its list of devices on all cluster nodes attached to the logical volume you created in [Step 2](#).

See your VERITAS Volume Manager documentation for information about using the `vxdctl enable` command to update new devices (volumes) in your VERITAS Volume Manager list of devices.

6. If necessary, partition the logical volume.
7. From any node in the cluster, update the global device namespace.

```
# scgdevs
```

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

Where to Go From Here

To create a new resource or reconfigure a running resource to use the new StorEdge T3/T3+ disk tray logical volume, see the *Sun Cluster 3.1 Data Services Installation and Configuration Guide*.

To configure a logical volume as a quorum device, see the *Sun Cluster 3.0 U1 System Administration Guide* for the procedure on adding a quorum device.

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

Use this procedure to remove a logical volume. This procedure assumes all cluster nodes are booted and attached to the StorEdge T3/T3+ disk tray that hosts the logical volume you are removing.

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



Caution – This procedure removes all data on the logical volume you are removing.

1. **If necessary, migrate all data and volumes off the logical volume you are removing. Otherwise, proceed to [Step 2](#).**
2. **Is the logical volume you are removing a quorum device?**

```
# scstat -q
```

- If yes, remove the quorum device before you proceed.
- If no, go to [Step 3](#).

For the procedure on removing a quorum device, see the *Sun Cluster 3.0 U1 System Administration Guide*.

3. **Are you running VERITAS Volume Manager?**

- If not, go to [Step 4](#).
- If you are running VERITAS Volume Manager, update its list of devices on all cluster nodes attached to the logical volume you are removing.

See your VERITAS Volume Manager documentation for information about using the `vxdisk rm` command to remove devices (volumes) in your VERITAS Volume Manager device list.

4. **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.

5. **Telnet to the disk tray and remove the logical volume.**

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

6. Determine the resource groups and device groups that are running on Node A and Node B.

Record this information because you will use it in [Step 13](#) of this procedure to return resource groups and device groups to these nodes.

```
# scstat
```

7. Move all resource groups and device groups off Node A.

```
# scswitch -s -h nodename
```

8. Shut down and reboot Node A by using the `shutdown` command with the `i6` option.

The `-i6` option with the `shutdown` command causes the node to reboot after it shuts down to the `ok` prompt.

```
# shutdown -y -g0 -i6
```

For more information, see the *Sun Cluster 3.0 U1 System Administration Guide*.

9. On Node A, remove the obsolete device IDs (DIDs).

```
# devfsadm -C  
# scdidadm -C
```

10. Move all resource groups and device groups off Node B.

```
# scswitch -s -h nodename
```

11. Shut down and reboot Node B by using the `shutdown` command with the `i6` option.

The `-i6` option with the `shutdown` command causes the node to reboot after it shuts down to the `ok` prompt.

```
# shutdown -y -g0 -i6
```

For more information, see the *Sun Cluster 3.0 U1 System Administration Guide*.

12. On Node B, remove the obsolete DIDs.

```
# devfsadm -C
# sctidadm -C
```

13. Return the resource groups and device groups you identified in [Step 6](#) to Node A and Node B.

```
# scswitch -z -g resource-group -h nodename
# scswitch -z -D device-group-name -h nodename
```

For more information, see the *Sun Cluster 3.0 U1 System Administration Guide*.

Where to Go From Here

To create a logical volume, see [“How to Create a Sun StorEdge T3/T3+ Disk Tray Logical Volume”](#) on page 80.

Maintaining a StorEdge T3/T3+ Disk Tray

This section contains the procedures for maintaining a StorEdge T3 or StorEdge T3+ disk tray. The following table lists these procedures. This section does not include a procedure for adding a disk drive and a procedure for removing a disk drive because a StorEdge T3/T3+ disk tray only operates when fully configured.



Caution – If you remove any field-replaceable unit (FRU) for an extended period of time, thermal complications might result. To prevent this complication, the StorEdge T3/T3+ disk tray 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 StorEdge T3/T3+ disk tray, and all attached StorEdge T3/T3+ disk trays, will shut down and power off.

TABLE C-2 Task Map: Maintaining a StorEdge T3/T3+ Disk Tray

Task	For Instructions, Go To
Upgrade StorEdge T3/T3+ disk tray firmware.	“How to Upgrade StorEdge T3/T3+ Disk Tray Firmware” on page 87
Replace a disk drive.	“How to Replace a Disk Drive” on page 88
Add a StorEdge T3/T3+ disk tray.	“How to Add a StorEdge T3/T3+ Disk Tray” on page 89
Remove a StorEdge T3/T3+ disk tray.	“How to Remove a StorEdge T3/T3+ Disk Tray” on page 99
Replace a host-to-hub fiber-optic cable.	“How to Replace a Host-to-Hub Component” on page 102
Replace an FC-100/S host adapter GBIC.	“How to Replace a Host-to-Hub Component” on page 102
Replace an FC-100 hub GBIC that connects a FC-100 hub to a host.	“How to Replace a Host-to-Hub Component” on page 102
Replace a hub-to-disk tray fiber optic cable.	“How to Replace a Sun StorEdge FC-100 Hub and Hub-to-Disk Tray Component” on page 103

TABLE C-2 Task Map: Maintaining a StorEdge T3/T3+ Disk Tray (Continued)

Task	For Instructions, Go To
Replace an FC-100 hub GBIC that connects the FC-100 hub to a StorEdge T3 disk tray.	"How to Replace a Sun StorEdge FC-100 Hub and Hub-to-Disk Tray Component" on page 103
Replace a Sun StorEdge FC-100 hub.	"How to Replace a Sun StorEdge FC-100 Hub and Hub-to-Disk Tray Component" on page 103
Replace a Sun StorEdge FC-100 hub power cord.	"How to Replace a Sun StorEdge FC-100 Hub and Hub-to-Disk Tray Component" on page 103
Replace a media interface adapter (MIA) on a StorEdge T3 disk tray (not applicable for StorEdge T3+ disk trays).	"How to Replace a Sun StorEdge FC-100 Hub and Hub-to-Disk Tray Component" on page 103
Replace a StorEdge T3 disk tray controller.	"How to Replace a StorEdge T3/T3+ Disk Tray Chassis" on page 105
Replace a StorEdge T3 disk tray chassis.	"How to Replace a StorEdge T3/T3+ Disk Tray Chassis" on page 105
Replace a host adapter.	"How to Replace a Host Adapter" on page 106
Upgrade a StorEdge T3 disk tray controller to a StorEdge T3+ disk tray controller.	<i>Sun StorEdge T3 Array Controller Upgrade Manual</i>
Replace a Power and Cooling Unit (PCU). 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). Follow the same procedure used in a non-cluster environment.	<i>Sun StorEdge T3 Installation, Operation, and Service Manual</i>
Replace a StorEdge T3 disk tray power cable. Follow the same procedure used in a non-cluster environment.	<i>Sun StorEdge T3 Installation, Operation, and Service Manual</i>
Replace an Ethernet cable. Follow the same procedure used in a non-cluster environment.	<i>Sun StorEdge T3 Installation, Operation, and Service Manual</i>

▼ How to Upgrade StorEdge T3/T3+ Disk Tray Firmware

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



Caution – Perform this procedure on one StorEdge T3/T3+ disk tray at a time. This procedure requires that you reset the StorEdge T3/T3+ disk tray you are upgrading. If you reset more than one StorEdge T3/T3+ disk tray, your cluster will lose access to data if the StorEdge T3/T3+ disk trays are submirrors of each other.

- 1. On one node attached to the StorEdge T3/T3+ disk tray you are upgrading, detach that StorEdge T3/T3+ disk tray's submirrors.**

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

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

See the *Sun Cluster 3.0 U1 Release Notes* for information about accessing Sun's EarlyNotifier web pages, which list information about any required patches or firmware levels that are available for download. For the procedure on applying any host adapter firmware patch, see the firmware patch README file.

- 3. Reset the StorEdge T3/T3+ disk tray, if you have not already done so.**

For the procedure on rebooting a StorEdge T3/T3+ disk tray, 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 StorEdge T3/T3+ disk tray 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 StorEdge T3/T3+ disk tray 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 StorEdge T3/T3+ disk tray, and all attached StorEdge T3/T3+ disk trays, will shut down and power off.

- 1. If the failed disk drive impacted the logical volume's availability, remove the 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 StorEdge T3/T3+ Disk Tray

Use this procedure to add a new StorEdge T3/T3+ disk tray to a running cluster.

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

1. **Set up a Reverse Address Resolution Protocol (RARP) server on the network the new StorEdge T3/T3+ disk tray is to reside on, and then assign an IP address to the new StorEdge T3/T3+ disk tray.**

This RARP server enables you to assign an IP address to the new StorEdge T3/T3+ disk tray by using the StorEdge T3/T3+ disk tray's unique MAC address.

For the procedure on setting up a RARP server, see the *Sun StorEdge T3 Installation, Operation, and Service Manual*.

2. **(Skip this step if you are adding a StorEdge T3+ disk tray) Install the media interface adapter (MIA) in the StorEdge T3 disk tray you are adding as shown in [FIGURE C-2](#).**

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

3. **If necessary, install a gigabit interface converter (GBIC) in the Sun StorEdge FC-100 hub as shown in [FIGURE C-2](#).**

This GBIC enables you to connect the Sun StorEdge FC-100 hub to the StorEdge T3/T3+ disk tray you are adding.

For the procedure on installing an FC-100 hub GBIC, see the *FC-100 Hub Installation and Service Manual*.

4. **Install the Ethernet cable between the StorEdge T3/T3+ disk tray and the Local Area Network (LAN), as shown in [FIGURE C-2](#).**
5. **Power on the StorEdge T3/T3+ disk tray.**

Note – The StorEdge T3/T3+ disk tray might require a few minutes to boot.

For the procedure on powering on a StorEdge T3/T3+ disk tray, see the *Sun StorEdge T3 Installation, Operation, and Service Manual*.

6. **Telnet to the StorEdge T3/T3+ disk tray you are adding, and, if necessary, install the required StorEdge T3/T3+ disk tray controller firmware.**

See the *Sun Cluster 3.0 U1 Release Notes* for information about accessing Sun's EarlyNotifier web pages, which list information about any required patches or firmware levels that are available for download. For the procedure on applying any host adapter firmware patch, see the firmware patch README file.

7. Does this new StorEdge T3/T3+ disk tray have a unique target address?

- If yes, proceed to [Step 8](#).
- If no, change the target address for this new StorEdge T3/T3+ disk tray.

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

8. Install a fiber-optic cable between the Sun StorEdge FC-100 hub and the StorEdge T3/T3+ disk tray as shown in [FIGURE C-2](#).

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

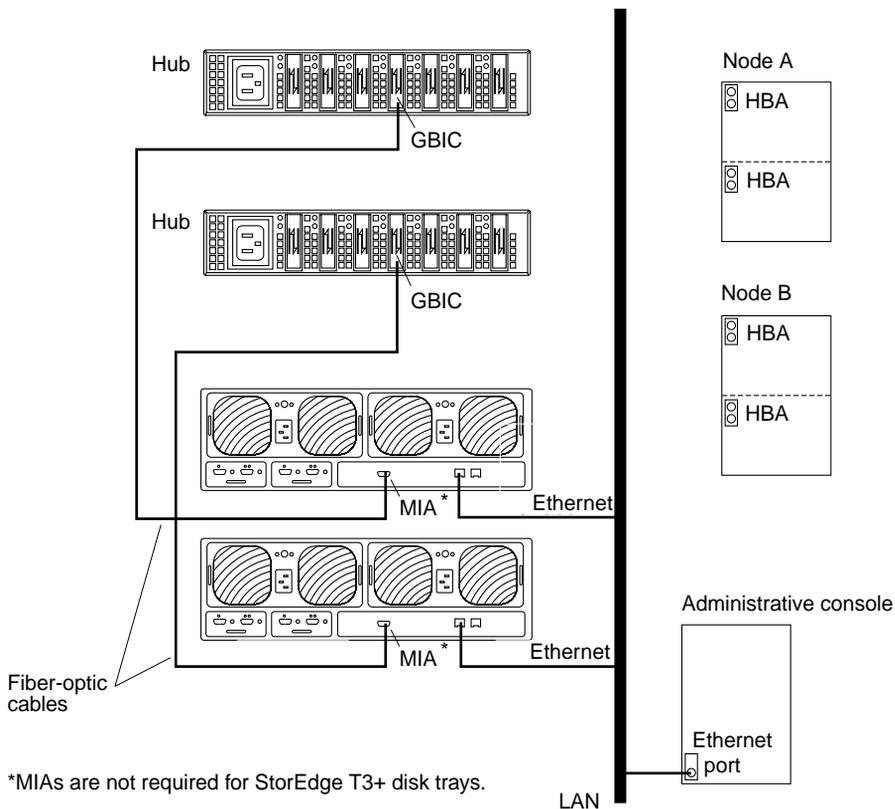


FIGURE C-2 Adding a StorEdge T3/T3+ Disk Tray in a Single-Controller Configuration

Note – Although [FIGURE C-2](#) shows a single-controller configuration, two disk trays are shown to illustrate how two non-interconnected disk trays are typically cabled in a cluster to allow data sharing and host-based mirroring.

9. Configure the new StorEdge T3/T3+ disk tray.

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

10. Determine the resource groups and device groups that are running on Node A and Node B.

Record this information because you will use it in [Step 42](#) of this procedure to return resource groups and device groups to these nodes.

```
# scstat
```

11. Move all resource groups and device groups off Node A.

```
# scswitch -s -h nodename
```

12. Do you need to install a host adapter in Node A?

- If yes, proceed to [Step 13](#).
- If no, skip to [Step 20](#).

13. Is the host adapter you are installing the first FC-100/S host adapter on Node A?

- If no, skip to [Step 15](#).
- 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
```

14. Are the Fibre Channel support packages installed?

- If yes, proceed to [Step 15](#).
- If no, install them.

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

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

15. Stop the Sun Cluster software on Node A and shut down Node A.

```
# shutdown -y -g0 -i0
```

For the procedure on shutting down a node, see the *Sun Cluster 3.0 U1 System Administration Guide*.

16. Power off Node A.

17. 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.

18. If necessary, power on and boot Node A.

```
{0} ok boot -x
```

For more information, see the *Sun Cluster 3.0 U1 System Administration Guide*.

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

See the *Sun Cluster 3.0 U1 Release Notes* for information about accessing Sun's EarlyNotifier web pages, which list information about any required patches or firmware levels that are available for download. For the procedure on applying any host adapter firmware patch, see the firmware patch `README` file.

20. If necessary, install gigabit interface converters (GBIC), as shown in [FIGURE C-3](#).

For the procedure on installing an FC-100 hub GBIC, see the *FC-100 Hub Installation and Service Manual*.

21. If necessary, connect a fiber-optic cable between the Sun StorEdge FC-100 hub and Node A as shown in [FIGURE C-3](#).

For the procedure on installing an FC-100/S host adapter GBIC, see your host adapter documentation. For the procedure on installing a fiber-optic cable, see the *Sun StorEdge T3 Configuration Guide*.

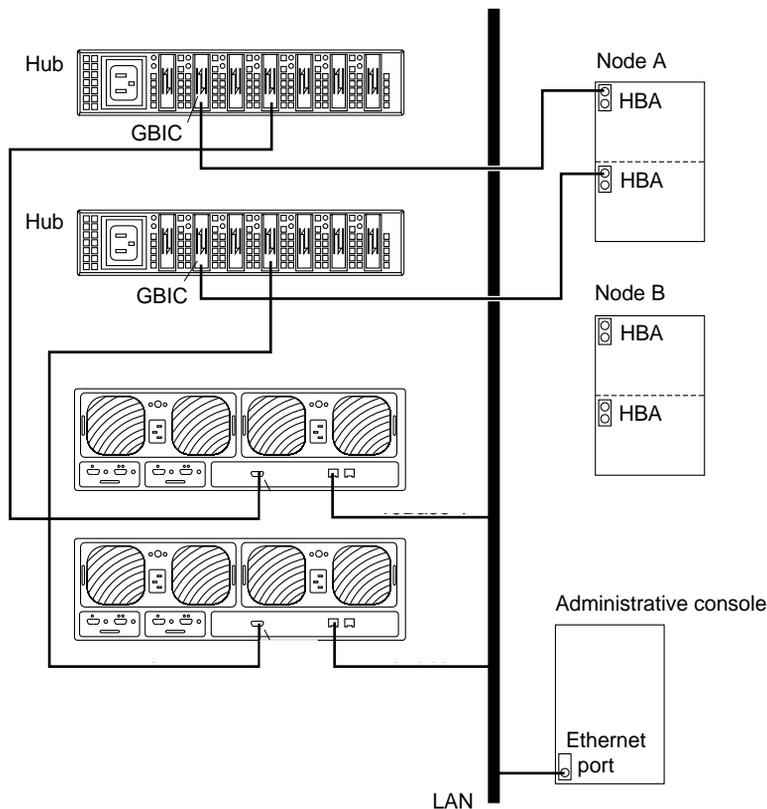


FIGURE C-3 Adding a StorEdge T3/T3+ Disk Tray in a Single-Controller Configuration

- 22. If necessary, install the required Solaris patches for StorEdge T3/T3+ disk tray support on Node A.**

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

- 23. Shut down Node A.**

```
# shutdown -y -g0 -i0
```

- 24. Perform a reconfiguration boot to create the new Solaris device files and links on Node A.**

```
{0} ok boot -r
```

25. Label the new logical volume.

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

26. (Optional) On Node A, verify that the device IDs (DIDs) are assigned to the new StorEdge T3/T3+ disk tray.

```
# scdidadm -l
```

27. Do you need to install a host adapter in Node B?

- If yes, proceed to [Step 28](#).
- If no, skip to [Step 36](#).

28. Is the host adapter you are installing the first FC-100/S host adapter on Node B?

- If no, skip to [Step 30](#).
- 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
```

29. Are the Fibre Channel support packages installed?

- If yes, proceed to [Step 30](#).
- If no, install them.

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

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

30. Move all resource groups and device groups off Node B.

```
# scswitch -s -h nodename
```

31. Stop the Sun Cluster software on Node B, and shut down the node.

```
# shutdown -y -g0 -i0
```

For the procedure on shutting down a node, see the *Sun Cluster 3.0 U1 System Administration Guide*.

32. Power off Node B.

For more information, see the *Sun Cluster 3.0 U1 System Administration Guide*.

33. 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.

34. If necessary, power on and boot Node B.

```
{0} ok boot -x
```

For more information, see the *Sun Cluster 3.0 U1 System Administration Guide*.

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

See the *Sun Cluster 3.0 U1 Release Notes* for information about accessing Sun's EarlyNotifier web pages, which list information about any required patches or firmware levels that are available for download. For the procedure on applying any host adapter firmware patch, see the firmware patch `README` file.

36. If necessary, install gigabit interface converters (GBIC) as shown in [FIGURE C-4](#).

For the procedure on installing an FC-100 hub GBIC, see the *FC-100 Hub Installation and Service Manual*.

37. If necessary, connect a fiber-optic cable between the Sun StorEdge FC-100 hub and Node B as shown in [FIGURE C-4](#).

For the procedure on installing a FC-100/S host adapter GBIC, see your host adapter documentation. For the procedure on installing a fiber-optic cable, see the *Sun StorEdge T3 Configuration Guide*.

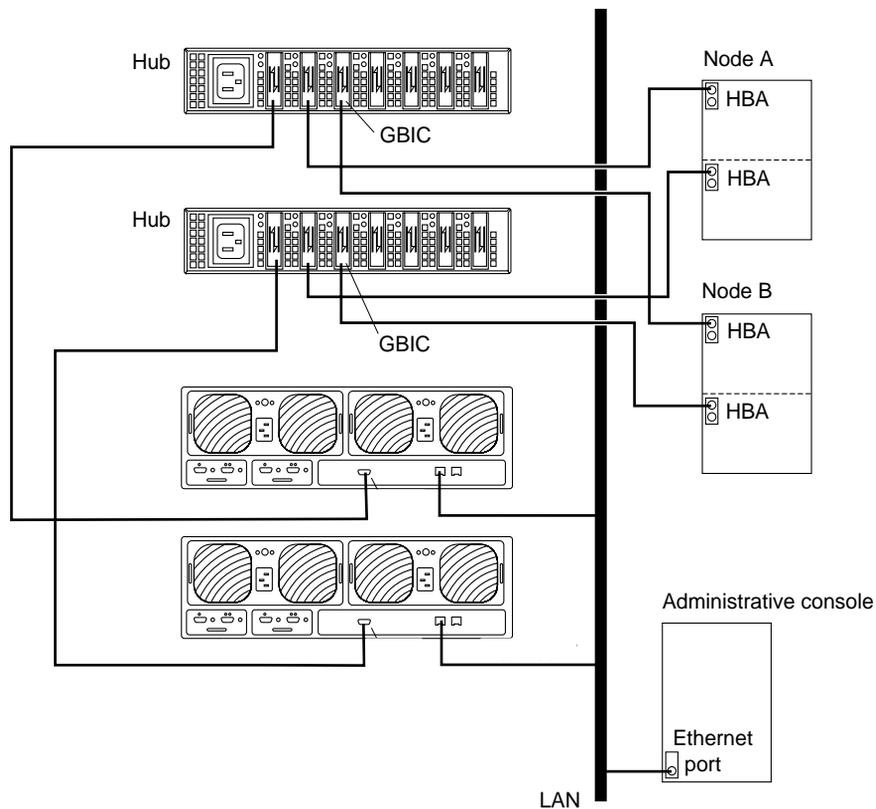


FIGURE C-4 Adding a StorEdge T3/T3+ Disk Tray in a Single-Controller Configuration

38. If necessary, install the required Solaris patches for StorEdge T3/T3+ disk tray support on Node B.

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

39. Shut down Node B.

```
# shutdown -y -g0 -i0
```

40. Perform a reconfiguration boot to create the new Solaris device files and links on Node B.

```
{0} ok boot -r
```

- 41. (Optional) On Node B, verify that the device IDs (DIDs) are assigned to the new StorEdge T3/T3+ disk tray.**

```
# sctidadm -l
```

- 42. Return the resource groups and device groups you identified in [Step 10](#) to Node A and Node B.**

```
# scswitch -z -g resource-group -h nodename  
# scswitch -z -D device-group-name -h nodename
```

For more information, see the *Sun Cluster 3.0 U1 System Administration Guide*.

- 43. Perform volume management administration to incorporate the new logical volumes into the cluster.**

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

▼ How to Remove a StorEdge T3/T3+ Disk Tray

Use this procedure to permanently remove a StorEdge T3/T3+ disk tray and its submirrors from a running cluster. This procedure provides the flexibility to remove the host adapters from the nodes for the StorEdge T3/T3+ disk tray you are removing.

This procedure defines Node A as the node 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 StorEdge T3/T3+ disk tray you are removing.

1. **Back up all database tables, data services, and volumes that are associated with the StorEdge T3/T3+ disk tray that you are removing.**

2. **Detach the submirrors from the StorEdge T3/T3+ disk tray you are removing in order to stop all I/O activity to the StorEdge T3/T3+ disk tray.**

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. **Determine the resource groups and device groups that are running on Node B.**

```
# scstat
```

5. **Move all resource groups and device groups off Node A.**

```
# scswitch -s -h nodename
```

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

```
# shutdown -y -g0 -i0
```

For the procedure on shutting down a node, see the *Sun Cluster 3.0 U1 System Administration Guide*.

7. Is the StorEdge T3/T3+ disk tray you are removing the last StorEdge T3/T3+ disk tray that is connected to Node A?

- If yes, disconnect the fiber-optic cable between Node A and the Sun StorEdge FC-100 hub that is connected to this StorEdge T3/T3+ disk tray, then disconnect the fiber-optic cable between the Sun StorEdge FC-100 hub and this StorEdge T3/T3+ disk tray.
- If no, proceed to [Step 8](#).

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

8. Do you want to remove the host adapter from Node A?

- If yes, power off Node A.
- If no, skip to [Step 11](#).

9. Remove the host adapter from Node A.

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

10. Without allowing the node to boot, power on Node A.

For more information, see the *Sun Cluster 3.0 U1 System Administration Guide*.

11. Boot Node A into cluster mode.

```
{0} ok boot
```

12. Move all resource groups and device groups off Node B.

```
# scswitch -s -h nodename
```

13. Stop the Sun Cluster software on Node B, and shut down Node B.

```
# shutdown -y -g0 -i0
```

14. Is the StorEdge T3/T3+ disk tray you are removing the last StorEdge T3/T3+ disk tray that is connected to the Sun StorEdge FC-100 hub.

- If yes, disconnect the fiber-optic cable that connects this Sun StorEdge FC-100 hub and Node B.
- If no, proceed to [Step 15](#).

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

15. Do you want to remove the host adapter from Node B?

- If yes, power off Node B.
- If no, skip to [Step 18](#).

16. Remove the host adapter from Node B.

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

17. Without allowing the node to boot, power on Node B.

For more information, see the *Sun Cluster 3.0 U1 System Administration Guide*.

18. Boot Node B into cluster mode.

```
{0} ok boot
```

For more information, see the *Sun Cluster 3.0 U1 System Administration Guide*.

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

```
# devfsadm -C  
# sctidadm -C
```

20. Return the resource groups and device groups you identified in [Step 4](#) to Node A and Node B.

```
# scswitch -z -g resource-group -h nodename  
# scswitch -z -D device-group-name -h nodename
```

▼ 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 GBIC
- FC-100 hub GBIC that connects a Sun StorEdge FC-100 hub to a node

1. **On the node that is connected to the host-to-hub connection you are replacing, determine the resource groups and device groups that are running on this node.**

```
# scstat
```

2. **Move all resource groups and device groups to another node.**

```
# scswitch -s -h nodename
```

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 an FC-100 hub GBIC, see the *FC-100 Hub Installation and Service Manual*. For the procedure on replacing a FC-100/S host adapter GBIC, see your host adapter documentation.

4. **Return the resource groups and device groups you identified in [Step 1](#) to the node that is connected to the host-to-hub connection you replaced.**

```
# scswitch -z -g resource-group -h nodename  
# scswitch -z -D device-group-name -h nodename
```

▼ 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 tray components.

- Fiber-optic cable that connects a Sun StorEdge FC-100 hub to a StorEdge T3/T3+ disk tray
- FC-100 hub GBIC that connects a Sun StorEdge FC-100 hub to a StorEdge T3/T3+ disk tray
- Sun StorEdge FC-100 hub
- Sun StorEdge FC-100 hub power cord
- Media interface adapter (MIA) on a StorEdge T3 disk tray (not applicable for StorEdge T3+ disk trays)

- 1. Detach the submirrors on the StorEdge T3/T3+ disk tray that is connected to the hub-to-disk tray fiber-optic cable you are replacing in order to stop all I/O activity to this StorEdge T3/T3+ disk tray.**

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

- 2. Replace the hub-to-disk tray 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 StorEdge T3/T3+ Disk Tray Controller

Use this procedure to replace a StorEdge T3/T3+ disk tray controller.

- 1. Detach the submirrors on the StorEdge T3/T3+ disk tray that is connected to the controller you are replacing in order to stop all I/O activity to this StorEdge T3/T3+ disk tray.**

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

- 2. Replace the controller.**

For the procedure on replacing a StorEdge T3/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 StorEdge T3/T3+ Disk Tray Chassis

Use this procedure to replace a StorEdge T3/T3+ disk tray chassis. This procedure assumes that you are retaining 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 StorEdge T3/T3+ disk tray chassis.

1. **Detach the submirrors on the StorEdge T3/T3+ disk tray that is connected to the chassis you are replacing in order to stop all I/O activity to this StorEdge T3/T3+ disk tray.**

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

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

For the procedure on replacing a StorEdge T3/T3+ chassis, see the *Sun StorEdge T3 Field 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. Node A in this procedure refers to the node with the failed host adapter you are replacing. Node B is a backup node.

1. Determine the resource groups and device groups that are running on Node A and Node B.

Record this information because you will use it in [Step 9](#) of this procedure to return resource groups and device groups to these nodes.

```
# scstat
```

2. Move all resource groups and device groups off Node A.

```
# scswitch -s -h nodename
```

3. Shut down Node A.

```
# shutdown -y -g0 -i0
```

4. Power off Node A.

For more information, see the *Sun Cluster 3.0 U1 System Administration Guide*.

5. Replace the failed host adapter.

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

6. Power on Node A.

For more information, see the *Sun Cluster 3.0 U1 System Administration Guide*.

7. Boot Node A into cluster mode.

```
{0} ok boot
```

For more information, see the *Sun Cluster 3.0 U1 System Administration Guide*.

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

See the *Sun Cluster 3.0 U1 Release Notes* for information about accessing Sun's EarlyNotifier web pages, which list information about any required patches or firmware levels that are available for download. For the procedure on applying any host adapter firmware patch, see the firmware patch README file.

9. Return the resource groups and device groups you identified in [Step 1](#) to Node A and Node B.

```
# scswitch -z -g resource-group -h nodename  
# scswitch -z -D device-group-name -h nodename
```

For more information, see the *Sun Cluster 3.0 U1 System Administration Guide*.

Installing and Maintaining the Netra D130 and StorEdge S1 Enclosures

This chapter provides the procedures for installing and maintaining the Netra™ D130 and StorEdge S1 storage enclosures.

This chapter contains the following procedures:

- [“How to Install a Netra D130/StorEdge S1 Enclosure” on page 110](#)
- [“How to Add a Netra D130/StorEdge S1 Disk Drive to a Running Cluster” on page 119](#)
- [“How to Replace a Netra D130/StorEdge S1 Disk Drive in a Running Cluster” on page 122](#)
- [“How to Remove a Netra D130/StorEdge S1 Disk Drive From a Running Cluster” on page 126](#)
- [“How to Add a Netra D130/StorEdge S1 Enclosure to a Running Cluster” on page 127](#)
- [“How to Replace a Netra D130/StorEdge S1 Enclosure in a Running Cluster” on page 133](#)
- [“How to Remove a Netra D130/StorEdge S1 Enclosure From a Running Cluster” on page 135](#)

For conceptual information on multihost disks, see the *Sun Cluster 3.0 U1 Concepts* document.

Installing Netra D130/StorEdge S1 Enclosures

This section describes the procedure for an initial installation of a Netra D130/StorEdge S1 storage enclosures.

▼ How to Install a Netra D130/StorEdge S1 Enclosure

Use this procedure for an initial installation of a Netra D130/StorEdge S1 enclosures, prior to installing the Solaris operating environment and Sun Cluster software. Perform this procedure in conjunction with the procedures in the *Sun Cluster 3.0 U1 Installation Guide* and your server hardware manual.

Multihost storage in clusters uses the multi-initiator capability of the SCSI (Small Computer System Interface) specification. For conceptual information on multi-initiator capability, see the *Sun Cluster 3.0 U1 Concepts* document.

1. Ensure that each device in the SCSI chain has a unique SCSI address.

The default SCSI address for host adapters is 7. Reserve SCSI address 7 for one host adapter in the SCSI chain. This procedure refers to node that has SCSI address 7 as the “second node.”

To avoid conflicts, in [Step 6](#) you will change the `scsi-initiator-id` of the remaining host adapter in the SCSI chain to an available SCSI address. This procedure refers to the node that has an available SCSI address as the “first node.”

For a full list of commands, see the *OpenBoot 3.x Command Reference Manual*.

Note – Even though a slot in the Netra D130/StorEdge S1 enclosures might not be in use, do not set the `scsi-initiator-id` for the first node to the SCSI address for that disk slot. This precaution minimizes future complications if you install additional disk drives.

2. Install the host adapters and (if used) Netra E1 Expanders in the nodes that will be connected to the Netra D130/StorEdge S1 enclosures.

For the procedure on installing a host adapter, see the documentation that shipped with your host adapter, Netra E1, and node hardware.

Note – If your host has only one SCSI port see [“Single SCSI Port Hosts” on page 111](#). If your host has two SCSI ports see [“Dual SCSI Port Hosts” on page 114](#).

Single SCSI Port Hosts

When installing the Netra D130/StorEdge S1 storage enclosures on single SCSI port hosts use the Netra E1 PCI Expander for the second host SCSI port. [FIGURE D-1](#) shows an overview of the installation. The storage devices are cabled so that there is no single point of failure in the cluster. Netra E1 PCI Expanders provide the second SCSI port for the 1RU form factor hosts such as the Netra t1, x1 or t1 200.

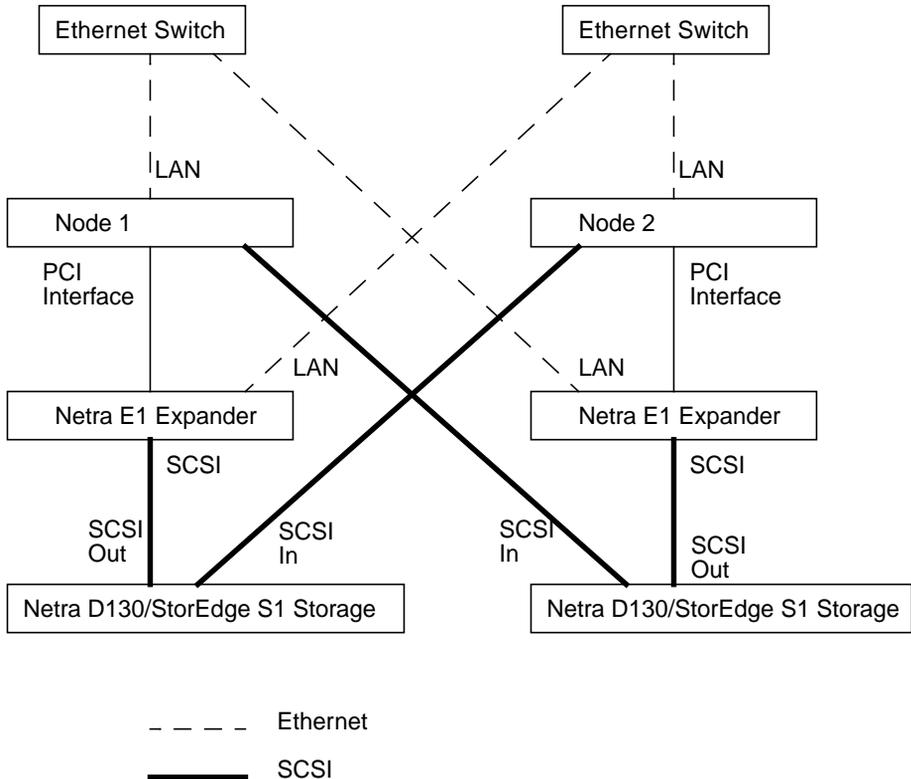


FIGURE D-1 Overview Example of a Enclosure Mirrored Pair Using E1 Expanders

1. Connect the cables to the Netra D130/StorEdge S1 enclosures, as shown in [FIGURE D-2](#).

Make sure that the *entire* SCSI bus length to each Netra D130 enclosures is less than 6 m. The maximum SCSI bus length for the StorEdge S1 enclosure is 12 m. This measurement includes the cables to both nodes, as well as the bus length internal to each enclosure, node, and host adapter. Refer to the documentation that shipped with the enclosures for other restrictions regarding SCSI operation.

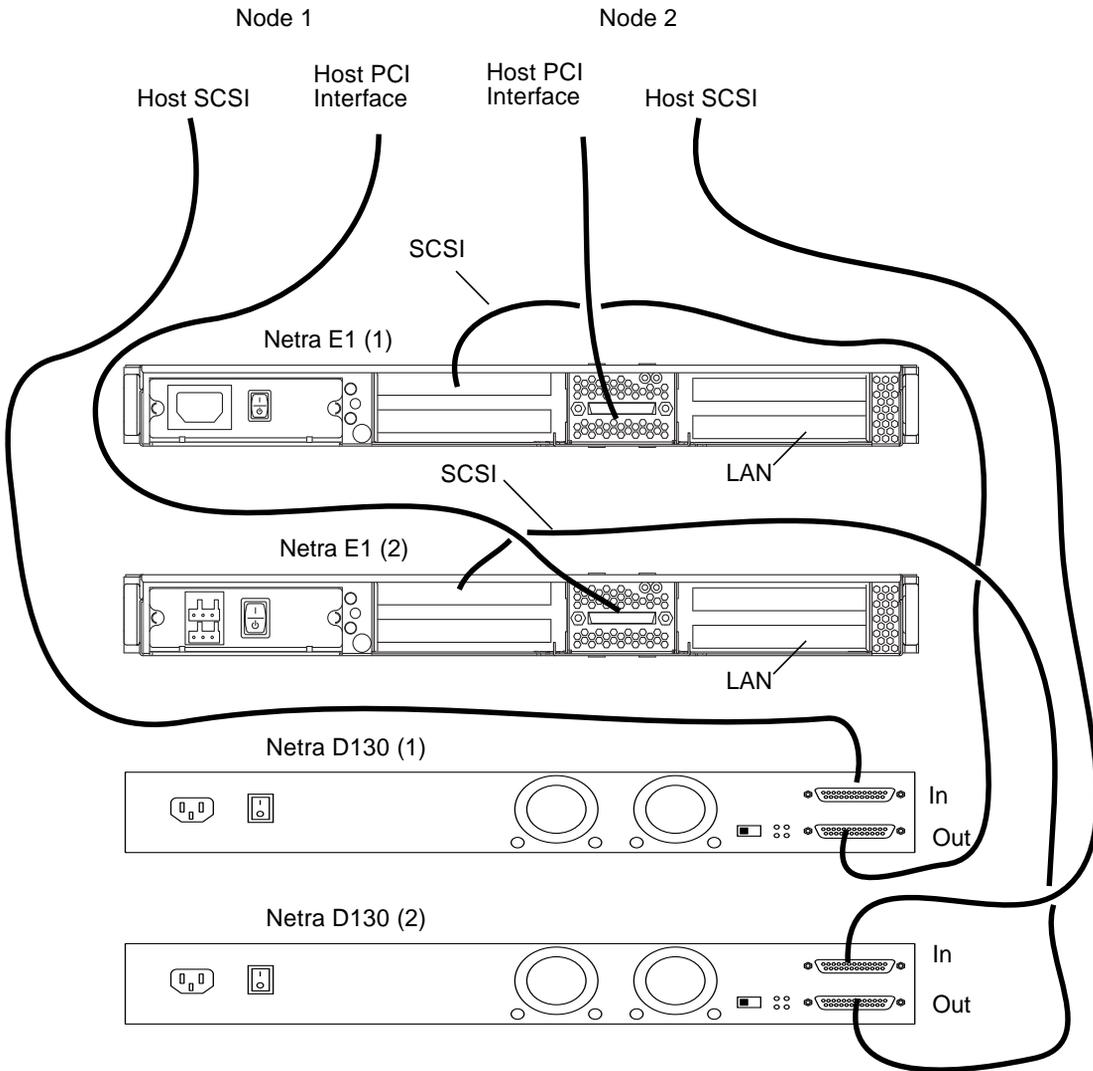


FIGURE D-2 Example of SCSI Cabling for an Enclosure Mirrored Pair

2. Connect the Ethernet cables between the host enclosures, Netra E1 PCI expanders, and Ethernet switches, as shown in [FIGURE D-3](#).

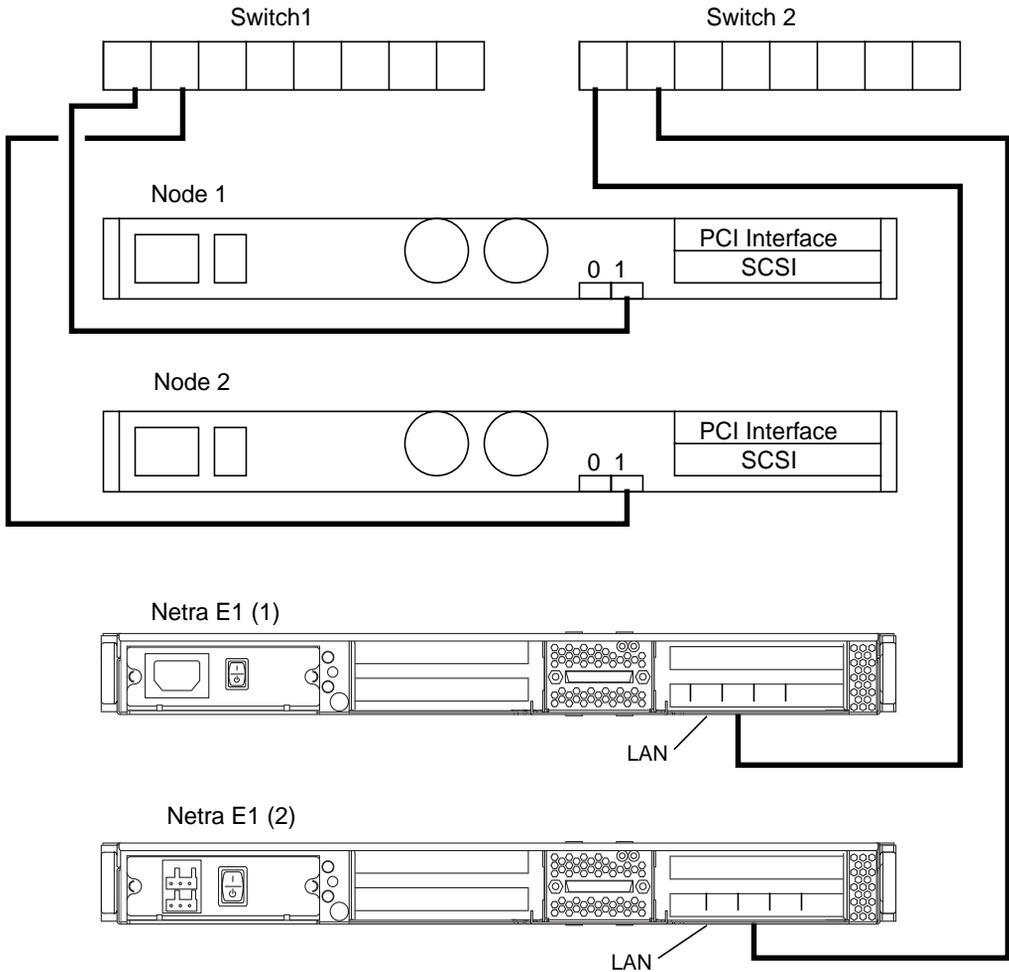


FIGURE D-3 Example of Ethernet cabling for an Mirrored Pair Using E1 Expanders

Dual SCSI Port Hosts

1. **Connect the cables to the Netra D130/StorEdge S1 enclosures, as shown in [FIGURE D-4](#).**

Make sure that the *entire* SCSI bus length to each Netra D130 enclosures is less than 6 m. The maximum SCSI bus length for the StorEdge S1 is 12 m. This measurement includes the cables to both nodes, as well as the bus length internal to each Netra D130/StorEdge S1 enclosures, node, and host adapter. Refer to the documentation that shipped with the Netra D130/StorEdge S1 enclosures for other restrictions regarding SCSI operation.

2. **Connect the cables to the Netra D130/StorEdge S1 enclosures, as shown in [FIGURE D-4](#).**

Make sure that the *entire* SCSI bus length to each Netra D130 enclosures is less than 6 m (12 m for the StorEdge S1). This measurement includes the cables to both nodes, as well as the bus length internal to each Netra D130/StorEdge S1 enclosures, node, and host adapter. Refer to the documentation that shipped with the Netra D130/StorEdge S1 enclosures for other restrictions regarding SCSI operation.

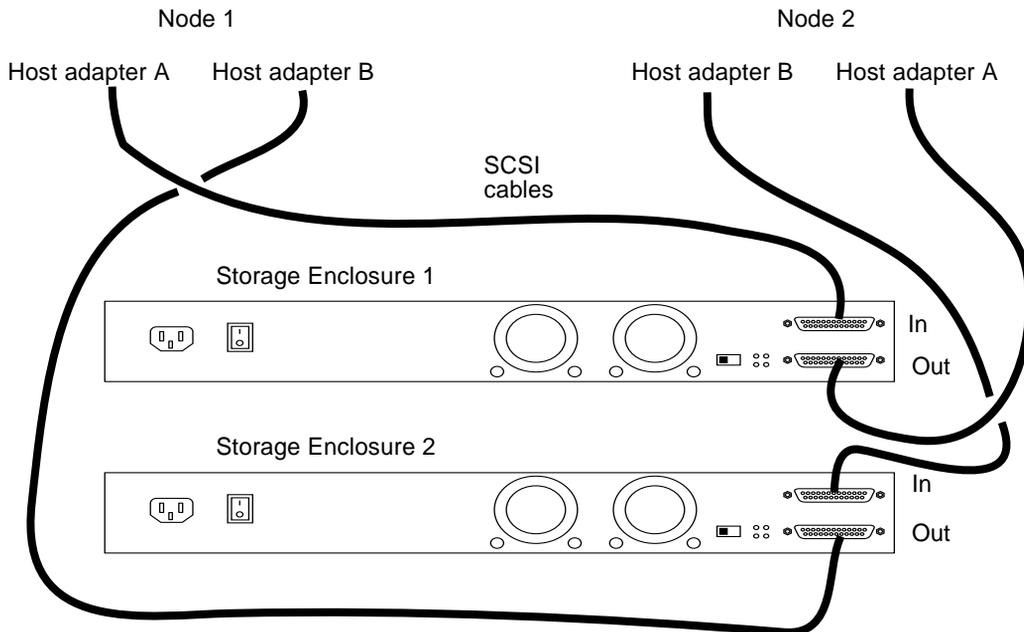


FIGURE D-4 Example of SCSI Cabling for an Enclosure Mirrored Pair

3. **Connect the AC or DC power cord for each Netra D130/StorEdge S1 enclosures of the mirrored pair to a different power source.**

4. Power on the *first* node but do not allow it to boot. If necessary, halt the node to continue with OpenBoot™ PROM (OBP) Monitor tasks (the first node is the node with an available SCSI address).
5. Find the paths to the host adapters.

```
{0} ok show-disks
a) /pci@1f,4000/pci@4/SUNW,isptwo@4/sd
b) /pci@1f,4000/pci@2/SUNW,isptwo@4/sd
```

Identify and record the two controllers that will be connected to the storage devices, and record these paths. Use this information to change the SCSI addresses of these controllers in the `nvrामrc` script in [Step 6](#). Do not include the `/sd` directories in the device paths.

6. Edit the `nvrामrc` script to set the `scsi-initiator-id` for the host adapters on the first node.

For a full list of `nvrामrc` editor and `nvedit` keystroke commands, see the *OpenBoot 3.x Command Reference Manual*.

The following example sets the `scsi-initiator-id` to 6. The OpenBoot PROM Monitor prints the line numbers (0:, 1:, and so on).

Note – Insert exactly one space after the first quotation mark and before `scsi-initiator-id`.

```
{0} ok nvedit
0: probe-all
1: cd /pci@1f,4000/pci@4/SUNW,isptwo@4
2: 6 " scsi-initiator-id" integer-property
3: device-end
4: cd /pci@1f,4000/pci@2/SUNW,isptwo@4
5: 6 " scsi-initiator-id" integer-property
6: device-end
7: install-console
8: banner <Control C>
{0} ok
```

7. Store the changes.

The changes you make through the `nvedit` command are done on a temporary copy of the `nvrामrc` script. You can continue to edit this copy without risk. After you complete your edits, save the changes. If you are not sure about the changes, discard them.

- To store the changes, type:

```
{0} ok nvstore
{0} ok
```

- To discard the changes, type:

```
{0} ok nvquit
{0} ok
```

8. Verify the contents of the `nvrामrc` script you created in [Step 6](#), as shown in the following example.

If the contents of the `nvrामrc` script are incorrect, use the `nvedit` command to make corrections.

```
{0} ok printenv nvrामrc
nvrामrc =
    probe-all
    cd /pci@1f,4000/pci@4/SUNW,isptwo@4
    6 " scsi-initiator-id" integer-property
    device-end
    cd /pci@1f,4000/pci@2/SUNW,isptwo@4
    6 " scsi-initiator-id" integer-property
    device-end
    install-console
    banner
{0} ok
```

9. Instruct the OpenBoot PROM Monitor to use the `nvrामrc` script.

```
{0} ok setenv use-nvrामrc? true
use-nvrामrc? = true
{0} ok
```

10. Power on the *second* node but do not allow it to boot. If necessary, halt the node to continue with OpenBoot PROM Monitor tasks (the second node is the node that has SCSI address 7).

11. Verify that the `scsi-initiator-id` for the host adapter on the second node is set to 7.

Use the `show-disks` command to find the paths to the host adapters connected to these enclosures (as in [Step 5](#)). Select each host adapter's device tree node, and display the node's properties to confirm that the `scsi-initiator-id` for each host adapter is set to 7, as shown in the following example.

```
{0} ok cd /pci@1f,4000/pci@4/SUNW,isptwo@4
{0} ok .properties
...
scsi-initiator-id      00000007
...
{0} ok cd /pci@1f,4000/pci@2/SUNW,isptwo@4
{0} ok .properties
...
scsi-initiator-id      00000007
```

12. Continue with the Solaris operating environment, Sun Cluster software, and volume management software installation tasks.

For software installation procedures, see the *Sun Cluster 3.0 U1 Installation Guide*.

Maintaining a Netra D130/StorEdge S1

This section provides the procedures for maintaining a Netra D130/StorEdge S1 enclosures. The following table lists these procedures.

TABLE D-1 Task Map: Maintaining a Netra D130/StorEdge S1 Enclosure

Task	For Instructions, Go To...
Add a disk drive	“How to Add a Netra D130/StorEdge S1 Disk Drive to a Running Cluster” on page 119
Replace a disk drive	“How to Replace a Netra D130/StorEdge S1 Disk Drive in a Running Cluster” on page 122
Remove a disk drive	“How to Remove a Netra D130/StorEdge S1 Disk Drive From a Running Cluster” on page 126
Add a Netra D130/StorEdge S1 enclosures	“How to Add a Netra D130/StorEdge S1 Enclosure to a Running Cluster” on page 127
Replace a Netra D130/StorEdge S1 enclosures	“How to Replace a Netra D130/StorEdge S1 Enclosure in a Running Cluster” on page 133
Remove a Netra D130/StorEdge S1 enclosures	“How to Remove a Netra D130/StorEdge S1 Enclosure From a Running Cluster” on page 135

▼ How to Add a Netra D130/StorEdge S1 Disk Drive to a Running Cluster

Use this procedure to add a disk drive to a running cluster. Perform the steps in this procedure in conjunction with the procedures in the *Sun Cluster 3.0 U1 System Administration Guide* and your server hardware manual. “[Example—Adding a Netra D130/StorEdge S1 Disk Drive](#)” on page 121 shows how to apply this procedure.

For conceptual information on quorums, quorum devices, global devices, and device IDs, see the *Sun Cluster 3.0 U1 Concepts* document.

1. **Locate an empty disk slot in the Netra D130/StorEdge S1 enclosures for the disk drive you want to add.**

Identify the empty slots either by observing the disk drive LEDs on the front of the Netra D130/StorEdge S1 enclosures, or by removing the left side-cover of the unit. The target address IDs corresponding to the slots appear on the middle partition of the drive bay.

2. **Install the disk drive.**

For detailed instructions, see the documentation that shipped with your Netra D130/StorEdge S1 enclosures.

3. **On all nodes attached to the Netra D130/StorEdge S1 enclosures, configure the disk drive.**

```
# cfgadm -c configure cN
# devfsadm
```

4. **On all nodes, ensure that entries for the disk drive have been added to the /dev/rdisk directory.**

```
# ls -l /dev/rdisk
```

5. **If needed, use the `format(1M)` command or the `fmthard(1M)` command to partition the disk drive.**

6. **From any node, 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 an expected behavior.

```
# scgdevs
```

7. On all nodes, verify that a device ID (DID) has been assigned to the disk drive.

```
# sctidadm -l
```

Note – As shown in [“Example—Adding a Netra D130/StorEdge S1 Disk Drive” on page 121](#), the DID 35 assigned to the new disk drive might not be in sequential order in the Netra D130/StorEdge S1 enclosures.

8. Perform volume management administration to add the new disk drive to the configuration.

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

Example—Adding a Netra D130/StorEdge S1 Disk Drive

The following example shows how to apply the procedure for adding a Netra D130/StorEdge S1 enclosure disk drive.

```
# scdidadm -l
16     phys-circinus-3:/dev/rdisk/c2t0d0 /dev/did/rdsk/d16
17     phys-circinus-3:/dev/rdisk/c2t1d0 /dev/did/rdsk/d17
18     phys-circinus-3:/dev/rdisk/c2t2d0 /dev/did/rdsk/d18
19     phys-circinus-3:/dev/rdisk/c2t3d0 /dev/did/rdsk/d19
...
26     phys-circinus-3:/dev/rdisk/c2t12d0 /dev/did/rdsk/d26
30     phys-circinus-3:/dev/rdisk/clt2d0 /dev/did/rdsk/d30
31     phys-circinus-3:/dev/rdisk/clt3d0 /dev/did/rdsk/d31
32     phys-circinus-3:/dev/rdisk/clt10d0 /dev/did/rdsk/d32
33     phys-circinus-3:/dev/rdisk/c0t0d0 /dev/did/rdsk/d33
34     phys-circinus-3:/dev/rdisk/c0t6d0 /dev/did/rdsk/d34
8190   phys-circinus-3:/dev/rmt/0       /dev/did/rmt/2
# cfgadm -c configure c1
# devfsadm
# scgdevs
Configuring DID devices
Could not open /dev/rdsk/c0t6d0s2 to verify device id.
    Device busy
Configuring the /dev/global directory (global devices)
obtaining access to all attached disks
reservation program successfully exiting
# scdidadm -l
16     phys-circinus-3:/dev/rdisk/c2t0d0 /dev/did/rdsk/d16
17     phys-circinus-3:/dev/rdisk/c2t1d0 /dev/did/rdsk/d17
18     phys-circinus-3:/dev/rdisk/c2t2d0 /dev/did/rdsk/d18
19     phys-circinus-3:/dev/rdisk/c2t3d0 /dev/did/rdsk/d19
...
26     phys-circinus-3:/dev/rdisk/c2t12d0 /dev/did/rdsk/d26
30     phys-circinus-3:/dev/rdisk/clt2d0 /dev/did/rdsk/d30
31     phys-circinus-3:/dev/rdisk/clt3d0 /dev/did/rdsk/d31
32     phys-circinus-3:/dev/rdisk/clt10d0 /dev/did/rdsk/d32
33     phys-circinus-3:/dev/rdisk/c0t0d0 /dev/did/rdsk/d33
34     phys-circinus-3:/dev/rdisk/c0t6d0 /dev/did/rdsk/d34
35     phys-circinus-3:/dev/rdsk/c2t13d0 /dev/did/rdsk/d35
8190   phys-circinus-3:/dev/rmt/0       /dev/did/rmt/2
#
```

Where to Go From Here

To configure a disk drive as a quorum device, see the Sun Cluster 3.0 U1 System Administration Guide for the procedure on adding a quorum device.

▼ How to Replace a Netra D130/StorEdge S1 Disk Drive in a Running Cluster

Use this procedure to replace a Netra D130/StorEdge S1 enclosure's disk drive. “[Example—Replacing a Netra D130/StorEdge S1 Disk Drive](#)” on page 125 shows how to apply this procedure. Perform the steps in this procedure in conjunction with the procedures in *Sun Cluster 3.0 U1 System Administration Guide* and your server hardware manual. Use the procedures in your server hardware manual to identify a failed disk drive.

For conceptual information on quorums, quorum devices, global devices, and device IDs, see the *Sun Cluster 3.0 U1 Concepts* document.

1. Identify the disk drive that needs replacement.

If the disk error message reports the drive problem by device ID (DID), use the `scdidadm -l` command to determine the Solaris logical device name. If the disk error message reports the drive problem by the Solaris physical device name, use your Solaris documentation to map the Solaris physical device name to the Solaris logical device name. Use this Solaris *logical* device name and DID throughout this procedure.

```
# scdidadm -l deviceID
```

2. Determine if the disk drive you want to replace is a quorum device.

```
# scstat -q
```

- If the disk drive you want to replace is a quorum device, put the quorum device into maintenance state before you go to [Step 3](#). For the procedure on putting a quorum device into maintenance state, see the *Sun Cluster 3.0 U1 System Administration Guide*.
- If the disk is not a quorum device, go to [Step 3](#).

3. If possible, back up the metadevice or volume.

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

4. Perform volume management administration to remove the disk drive from the configuration.

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

5. Identify the failed disk drive's physical DID.

Use this physical DID in [Step 12](#) to verify that the failed disk drive has been replaced with a new disk drive.

```
# scdidadm -o diskid -l cNtXdY
```

6. If you are using Solstice DiskSuite as your volume manager, save the disk partitioning for use when partitioning the new disk drive.

If you are using VERITAS Volume Manager, skip this step and go to [Step 7](#).

```
# prtvtoc /dev/rdisk/cNtXdYsZ > filename
```

Note – Do not save this file under /tmp because you will lose this file when you reboot. Instead, save this file under /usr/tmp.

7. Replace the failed disk drive.

For more information, see the *Sun StorEdge MultiPack Storage Guide*.

8. On one node attached to the Netra D130/StorEdge S1 enclosures, run the devfsadm(1M) command to probe all devices and to write the new disk drive to the /dev/rdisk directory.

Depending on the number of devices connected to the node, the devfsadm command can take at least five minutes to complete.

```
# devfsadm
```

9. If you are using Solstice DiskSuite as your volume manager, from any node connected to the Netra D130/StorEdge S1 enclosures, partition the new disk drive, using the partitioning you saved in [Step 6](#).

If you are using VERITAS Volume Manager, skip this step and go to [Step 10](#).

```
# fmthard -s filename /dev/rdisk/cNtXdYsZ
```

10. One at a time, shut down and reboot the nodes connected to the Netra D130/StorEdge S1 enclosures.

```
# scswitch -S -h nodename
# shutdown -y -g0 -i6
```

For more information, see the *Sun Cluster 3.0 U1 System Administration Guide*.

11. From any node connected to the disk drive, update the DID database.

```
# scdidadm -R deviceID
```

12. From any node, confirm that the failed disk drive has been replaced by comparing the new physical DID to the physical DID identified in [Step 5](#).

If the new physical DID is different from the physical DID in [Step 5](#), you successfully replaced the failed disk drive with a new disk drive.

```
# scdidadm -o diskid -l cNtXdY
```

13. On all connected nodes, upload the new information to the DID driver.

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 an expected behavior.

```
# scdidadm -ui
```

14. Perform volume management administration to add the disk drive back to its diskset or disk group.

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

15. If you want this new disk drive to be a quorum device, add the quorum device.

For the procedure on adding a quorum device, see the *Sun Cluster 3.0 U1 System Administration Guide*.

Example—Replacing a Netra D130/StorEdge S1 Disk Drive

The following example shows how to apply the procedure for replacing a Netra D130/StorEdge S1 enclosure disk drive.

```
# scdidadm -l d20
20      phys-schost-2:/dev/rdisk/c3t2d0 /dev/did/rdisk/d20
# scdidadm -o diskid -l c3t2d0
5345414741544520393735314336343734310000
# prtvtoc /dev/rdisk/c3t2d0s2 > /usr/tmp/c3t2d0.vtoc
...
# devfsadm
# fmthard -s /usr/tmp/c3t2d0.vtoc /dev/rdisk/c3t2d0s2
# scswitch -S -h node1
# shutdown -y -g0 -i6
...
# scdidadm -R d20
# scdidadm -o diskid -l c3t2d0
5345414741544520393735314336363037370000
# scdidadm -ui
```

▼ How to Remove a Netra D130/StorEdge S1 Disk Drive From a Running Cluster

Use this procedure to remove a disk drive from a Netra D130/StorEdge S1 enclosures. Perform the steps in this procedure in conjunction with the procedures in the *Sun Cluster 3.0 U1 System Administration Guide* and your server hardware manual.

For conceptual information on quorum, quorum devices, global devices, and device IDs, see the *Sun Cluster 3.0 U1 Concepts* document.

1. Determine if the disk drive you want to remove is a quorum device.

```
# scstat -q
```

- If the disk drive you want to replace is a quorum device, put the quorum device into maintenance state before you go to [Step 2](#). For the procedure on putting a quorum device into maintenance state, see the *Sun Cluster 3.0 U1 System Administration Guide*.
- If the disk is not a quorum device, go to [Step 2](#).

2. Perform volume management administration to remove the disk drive from the configuration.

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

3. Identify the disk drive that needs to be removed and the slot from which the disk drive needs to be removed.

If the disk error message reports the drive problem by DID, use the `scdidadm -l` command to determine the Solaris device name.

```
# scdidadm -l deviceID  
# cfgadm -al
```

4. Remove the disk drive.

For more the procedure on removing a disk drive, see the *Sun StorEdge MultiPack Storage Guide*.

5. On all nodes, remove references to the disk drive.

```
# cfgadm -c unconfigure cN::dsk/cNtXdY  
# devfsadm -C  
# scdidadm -C
```

▼ How to Add a Netra D130/StorEdge S1 Enclosure to a Running Cluster

Use this procedure to install a Netra D130/StorEdge S1 enclosures in a running cluster. Perform the steps in this procedure in conjunction with the procedures in the *Sun Cluster 3.0 U1 Installation Guide* and your server hardware manual.

For conceptual information on multi-initiator SCSI and device IDs, see the *Sun Cluster 3.0 U1 Concepts* document.

1. Ensure that each device in the SCSI chain has a unique SCSI address.

The default SCSI address for host adapters is 7. Reserve SCSI address 7 for one host adapter in the SCSI chain. This procedure refers to the node with SCSI address 7 as the “second node.”

To avoid conflicts, in [Step 9](#) you will change the `scsi-initiator-id` of the remaining host adapter in the SCSI chain to an available SCSI address. This procedure refers to the node with an available SCSI address as the “first node.”

For a full list of commands, see the *OpenBoot 3.x Command Reference Manual*.

Note – Even though a slot in the Netra D130/StorEdge S1 enclosures might not be in use, do not set the `scsi-initiator-id` for the first node to the SCSI address for that disk slot. This precaution minimizes future complications if you install additional disk drives.

2. Shut down and power off the *first* node.

```
# scswitch -s -h nodename
# shutdown -y -g0 -i0
```

For more information, see the *Sun Cluster 3.0 U1 System Administration Guide*.

3. Install the host adapters in the first node.

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

4. Connect the appropriate SCSI cable between the node and the Netra D130/StorEdge S1 enclosures, as shown in [FIGURE D-5](#).

Make sure that the *entire* SCSI bus length to each Netra D130/StorEdge S1 enclosures is less than 6 m. This measurement includes the cables to both nodes, as well as the bus length internal to each Netra D130/StorEdge S1 enclosures, node, and host adapter. Refer to the documentation that shipped with the Netra D130/StorEdge S1 enclosures for other restrictions regarding SCSI operation.

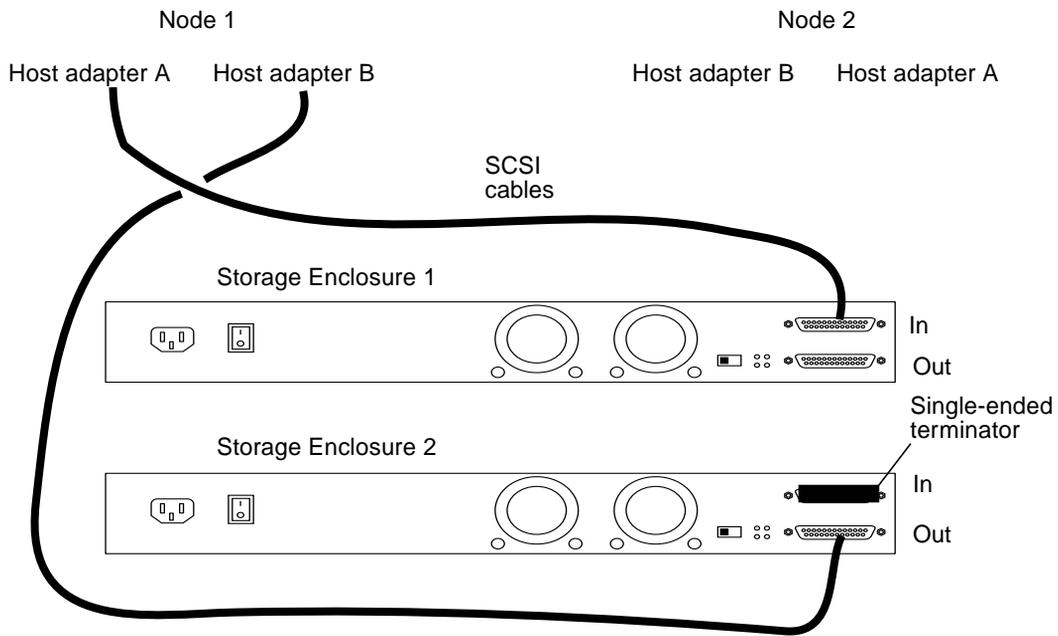


FIGURE D-5 Example of a Netra D130/StorEdge S1 enclosures mirrored pair

5. Temporarily install an appropriate terminator on the SCSI IN port of the *second* Netra D130/StorEdge S1 enclosures, as shown in [FIGURE D-5](#).
6. Connect each Netra D130/StorEdge S1 enclosures of the mirrored pair to different power sources.
7. Power on the first node and the Netra D130/StorEdge S1 enclosures.
8. Find the paths to the host adapters.

```
{0} ok show-disks
a) /pci@1f,4000/pci@4/SUNW,ispw@4/sd
b) /pci@1f,4000/pci@2/SUNW,ispw@4/sd
```

Identify and record the two controllers that will be connected to the storage devices, and record these paths. Use this information to change the SCSI addresses of these controllers in the `nvrampc` script in [Step 9](#). Do not include the `/sd` directories in the device paths.

9. Edit the `nvrामrc` script to set the `scsi-initiator-id` for the host adapters on the first node.

For a full list of commands, see the *OpenBoot 3.x Command Reference Manual*.

The following example sets the `scsi-initiator-id` to 6. The OpenBoot PROM Monitor prints the line numbers (0:, 1:, and so on).



Caution – Insert exactly one space after the first quotation mark and before `scsi-initiator-id`.

```
{0} ok nvedit
0: probe-all
1: cd /pci@1f,4000/pci@4/SUNW,isptwo@4
2: 6 " scsi-initiator-id" integer-property
3: device-end
4: cd /pci@1f,4000/pci@2/SUNW,isptwo@4
5: 6 " scsi-initiator-id" integer-property
6: device-end
7: install-console
8: banner <Control C>
{0} ok
```

10. Store the changes.

The changes you make through the `nvedit` command are done on a temporary copy of the `nvrामrc` script. You can continue to edit this copy without risk. After you complete your edits, save the changes. If you are not sure about the changes, discard them.

- To store the changes, type:

```
{0} ok nvstore
{0} ok
```

- To discard the changes, type:

```
{0} ok nvquit
{0} ok
```

11. Verify the contents of the `nvrामrc` script you created in [Step 9](#), as shown in the following example.

If the contents of the `nvrामrc` script are incorrect, use the `nvedit` command to make corrections.

```
{0} ok printenv nvrामrc
nvrामrc =
    probe-all
    cd /pci@1f,4000/pci@4/SUNW,isptwo@4
    6 " scsi-initiator-id" integer-property
    device-end
    cd /pci@1f,4000/pci@2/SUNW,isptwo@4
    6 " scsi-initiator-id" integer-property
    device-end
    install-console
    banner
{0} ok
```

12. Instruct the OpenBoot PROM Monitor to use the `nvrामrc` script, as shown in the following example.

```
{0} ok setenv use-nvrामrc? true
use-nvrामrc? = true
{0} ok
```

13. Boot the first node and wait for it to join the cluster.

```
{0} ok boot -r
```

For more information, see the *Sun Cluster 3.0 U1 System Administration Guide*.

14. On all nodes, verify that the DIDs have been assigned to the disk drives in the Netra D130/StorEdge S1 enclosures.

```
# scdidadm -l
```

15. Shut down and power off the *second* node.

```
# scswitch -S -h nodename
# shutdown -y -g0 -i0
```

16. Install the host adapters in the second node.

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

17. Connect the Netra D130/StorEdge S1 enclosures to the host adapters as shown in [FIGURE D-6](#) using the appropriate SCSI cables. (Remove the SCSI terminator you installed in [Step 5](#).)

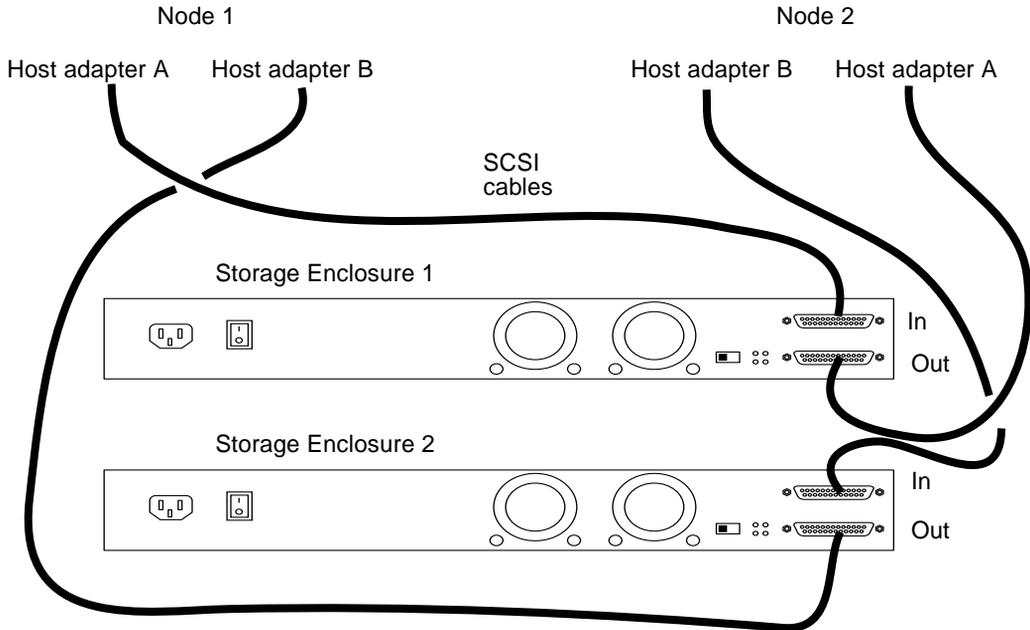


FIGURE D-6 Example of a Netra D130/StorEdge S1 enclosures mirrored pair

18. Power on the *second* node but do not allow it to boot. If necessary, halt the node to continue with OpenBoot PROM Monitor tasks.

19. Verify that the second node sees the new host adapters and disk drives.

```
{0} ok show-disks
```

- 20. Verify that the `scsi-initiator-id` for the host adapter on the second node is set to 7.**

Use the `show-disks` command to find the paths to the host adapters connected to these enclosures. Select each host adapter's device tree node, and display the node's properties to confirm that the `scsi-initiator-id` for each host adapter is set to 7.

```
{0} ok cd /pci@1f,4000/pci@4/SUNW,isptwo@4
{0} ok .properties
...
scsi-initiator-id      00000007
...
{0} ok cd /pci@1f,4000/pci@2/SUNW,isptwo@4
{0} ok .properties
...
scsi-initiator-id      00000007
```

- 21. Boot the second node and wait for it to join the cluster.**

```
{0} ok boot -r
```

- 22. On all nodes, verify that the DIDs have been assigned to the disk drives in the Netra D130/StorEdge S1 enclosures.**

```
# scdidadm -l
```

- 23. Perform volume management administration to add the disk drives in the Netra D130/StorEdge S1 enclosures to the volume management configuration.**

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

▼ How to Replace a Netra D130/StorEdge S1 Enclosure in a Running Cluster

Use this procedure to replace a Netra D130/StorEdge S1 enclosures. This procedure assumes that you want to retain the disk drives in the Netra D130/StorEdge S1 enclosures you are replacing and to retain the references to these same disk drives.

If you want to replace your disk drives, see [“How to Replace a Netra D130/StorEdge S1 Disk Drive in a Running Cluster” on page 122.](#)

- 1. If possible, back up the metadevices or volumes that reside in the disk array.**
For more information, see your Solstice DiskSuite or VERITAS Volume Manager documentation.
- 2. Perform volume management administration to remove the disk array from the configuration.**
For more information, see your Solstice DiskSuite or VERITAS Volume Manager documentation.
- 3. Disconnect the SCSI cables from the Netra D130/StorEdge S1 enclosures, disconnecting the cable on the SCSI OUT connector first, then the cable on the SCSI IN connector second (see [FIGURE D-7](#)).**

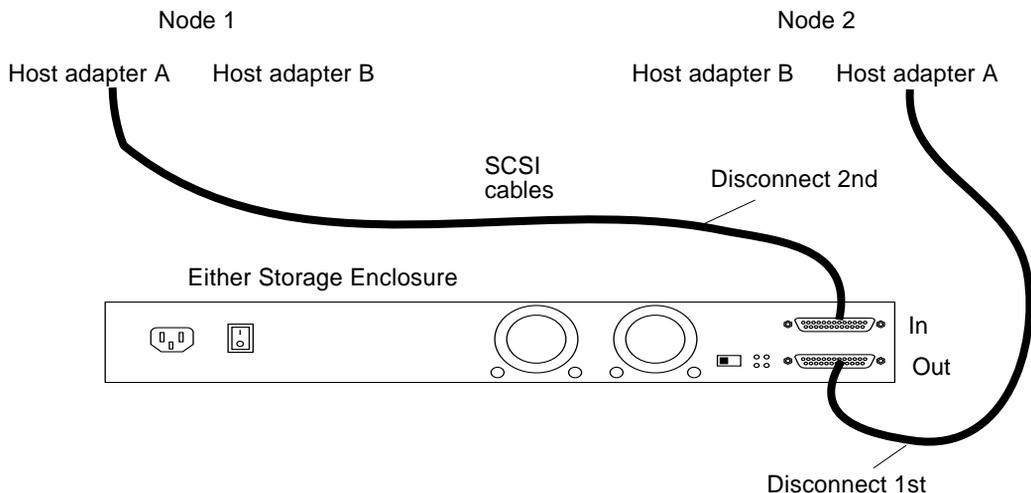


FIGURE D-7 Disconnecting the SCSI cables

4. Power off, and disconnect the Netra D130/StorEdge S1 enclosures from the AC power source.

For more information, see the documentation that shipped with your Netra D130/StorEdge S1 enclosures and the labels inside the lid of the Netra D130/StorEdge S1 enclosures.

5. Connect the new Netra D130/StorEdge S1 enclosures to an AC power source.

Refer to the documentation that shipped with the Netra D130/StorEdge S1 enclosures and the labels inside the lid of the Netra D130/StorEdge S1 enclosures.

6. Connect the SCSI cables to the new Netra D130/StorEdge S1 enclosures, reversing the order in which you disconnected them (connect the SCSI IN connector first, then the SCSI OUT connector second). See [FIGURE D-7](#).

7. Move the disk drives one at time from the old Netra D130/StorEdge S1 enclosures to the same slots in the new Netra D130/StorEdge S1 enclosures.

8. Power on the Netra D130/StorEdge S1 enclosures.

9. On all nodes attached to the Netra D130/StorEdge S1 enclosures, run the `devfsadm(1M)` command.

```
# devfsadm
```

10. One at a time, shut down and reboot the nodes connected to the Netra D130/StorEdge S1 enclosures.

```
# scswitch -s -h nodename  
# shutdown -y -g0 -i6
```

For more information on `shutdown(1M)`, see the *Sun Cluster 3.0 U1 System Administration Guide*.

11. Perform volume management administration to add the Netra D130/StorEdge S1 enclosures to the configuration.

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

▼ How to Remove a Netra D130/StorEdge S1 Enclosure From a Running Cluster

Use this procedure to remove a Netra D130/StorEdge S1 enclosures from a cluster. This procedure assumes that you want to remove the references to the disk drives in the enclosure.

- 1. Perform volume management administration to remove the Netra D130/StorEdge S1 enclosures from the configuration.**

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

- 2. Disconnect the SCSI cables from the Netra D130/StorEdge S1 enclosures, disconnecting them in the order shown in [FIGURE D-8](#).**

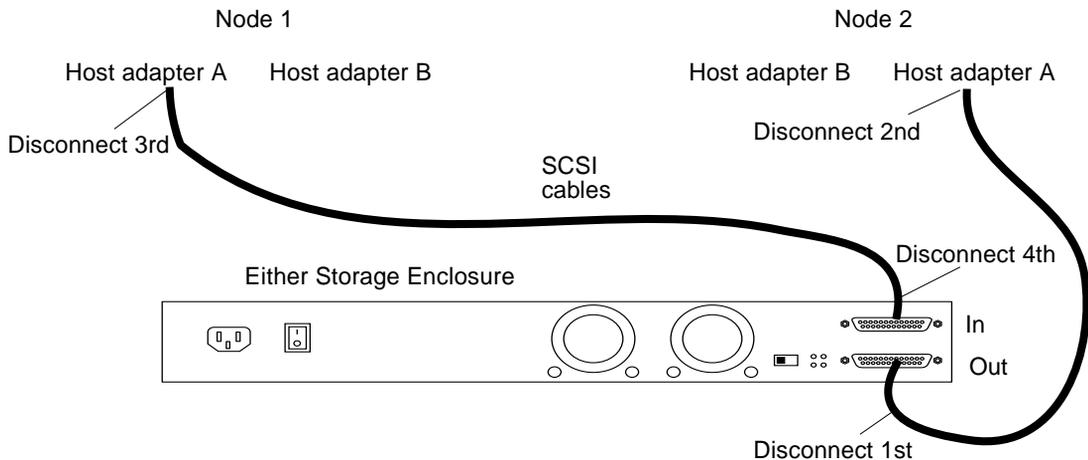


FIGURE D-8 Disconnecting the SCSI cables

- 3. Power off, and disconnect the Netra D130/StorEdge S1 enclosures from the AC power source.**

For more information, see the documentation that shipped with the Netra D130/StorEdge S1 enclosures and the labels inside the lid of the Netra D130/StorEdge S1 enclosures.

- 4. Remove the Netra D130/StorEdge S1 enclosures.**

For the procedure on removing an enclosures, see the *Sun StorEdge MultiPack Storage Guide*.

5. Identify the disk drives you need to remove from the cluster.

```
# cfgadm -al
```

6. On all nodes, remove references to the disk drives that were in the Netra D130/StorEdge S1 enclosures you removed in [Step 4](#).

```
# cfgadm -c unconfigure cN::dsk/cNtXdY  
# devfsadm -C  
# sddidadm -C
```

7. If needed, remove any unused host adapters from the nodes.

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