



Sun Cluster 3.x With Sun StorEdge MultiPack Enclosure Manual

Sun Microsystems, Inc.
4150 Network Circle
Santa Clara, CA 95054
U.S.A.

Part No: 817-0169-10
July 2003, Revision A

Copyright 2003 Sun Microsystems, Inc. 4150 Network Circle, Santa Clara, CA 95054 U.S.A. All rights reserved.

This product or document is protected by copyright and distributed under licenses restricting its use, copying, distribution, and decompilation. No part of this product or document may be reproduced in any form by any means without prior written authorization of Sun and its licensors, if any. Third-party software, including font technology, is copyrighted and licensed from Sun suppliers.

Parts of the product may be derived from Berkeley BSD systems, licensed from the University of California. UNIX is a registered trademark in the U.S. and other countries, exclusively licensed through X/Open Company, Ltd.

Sun, Sun Microsystems, the Sun logo, docs.sun.com, AnswerBook, AnswerBook2, and Solaris are trademarks, registered trademarks, or service marks of Sun Microsystems, Inc. in the U.S. and other countries. All SPARC trademarks are used under license and are trademarks or registered trademarks of SPARC International, Inc. in the U.S. and other countries. Products bearing SPARC trademarks are based upon an architecture developed by Sun Microsystems, Inc.

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

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

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

Copyright 2003 Sun Microsystems, Inc. 4150 Network Circle, Santa Clara, CA 95054 U.S.A. Tous droits réservés.

Ce produit ou document est protégé par un copyright et distribué avec des licences qui en restreignent l'utilisation, la copie, la distribution, et la décompilation. Aucune partie de ce produit ou document ne peut être reproduite sous aucune forme, par quelque moyen que ce soit, sans l'autorisation préalable et écrite de Sun et de ses bailleurs de licence, s'il y en a. Le logiciel détenu par des tiers, et qui comprend la technologie relative aux polices de caractères, est protégé par un copyright et licencié par des fournisseurs de Sun.

Des parties de ce produit pourront être dérivées du système 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, docs.sun.com, AnswerBook, AnswerBook2, et Solaris sont des marques de fabrique ou des marques déposées, ou marques de service, de Sun Microsystems, Inc. aux Etats-Unis et dans d'autres pays. Toutes les marques SPARC sont utilisées sous licence et sont des marques de fabrique ou des marques déposées de SPARC International, Inc. aux Etats-Unis et dans d'autres pays. Les produits portant les marques SPARC sont basés sur une architecture développée par Sun Microsystems, Inc.

L'interface d'utilisation graphique OPEN LOOK et Sun™ a été développée par Sun Microsystems, Inc. pour ses utilisateurs et licenciés. Sun reconnaît les efforts de pionniers de Xerox pour la recherche et le développement du concept des interfaces d'utilisation visuelle ou graphique pour l'industrie de l'informatique. Sun détient une licence non exclusive de Xerox sur l'interface d'utilisation graphique Xerox, cette licence couvrant également les licenciés de Sun qui mettent en place l'interface d'utilisation graphique OPEN LOOK et qui en outre se conforment aux licences écrites de Sun.

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



030807 @ 5943



Contents

Preface 5

1 Installing and Maintaining a StorEdge Multipack Enclosure 11

Installing Storage Arrays 11

 ▼ How to Install a Storage Array 11

Maintaining Storage Arrays 16

 ▼ How to Add a Disk Drive 17

 ▼ How to Replace a Disk Drive 19

 ▼ How to Remove a Disk Drive 22

 ▼ How to Add a Storage Array 22

 ▼ How to Replace a Storage Array 28

 ▼ How to Remove a Storage Array 30

 ▼ How to Replace a SCSI Cable 31

 ▼ How to Replace a Host Adapter 32

Index 35

Preface

The *Sun Cluster 3.x With Sun StorEdge Multipack Enclosure Manual* provides procedures specific to Sun StorEdge™ Multipack enclosures that are placed in a Sun™ Cluster environment.

Use this manual with any version of Sun Cluster 3.x software. Unless otherwise noted, procedures are the same for all Sun Cluster 3.x versions. See the “Revision History” on page 6 for a list of changes to this manual.

Who Should Use This Book

This book is for Sun representatives who are performing the initial installation of a Sun Cluster configuration and for system administrators who are responsible for maintaining the system.

This document is intended for experienced system administrators with extensive knowledge of Sun software and hardware. Do not use this document as a planning or presales guide. You should have already determined your system requirements and purchased the appropriate equipment and software before reading this document.

How This Book Is Organized

This book contains one chapter that consists of two major sections.

Section 1 discusses how to install Sun StorEdge Multipack storage array.

Section 2 describes how to maintain Sun StorEdge Multipack storage array in a running cluster.

Revision History

The following table lists the information that has been revised or added since the initial release of this documentation. The table also lists the revision date for these changes.

TABLE P-1 Sun Cluster 3.x With Sun StorEdge Multipack Enclosure Manual

Revision Date	New Information
July 2003	Added Quantum disk drive SUN2.1GB VK2275J to the list of Quantum disk drive models that require a specific <code>scsi-initiator-id</code> . See “How to Install a Storage Array” on page 11.

Related Documentation

The following books provide conceptual information or procedures to administer hardware and applications. If you plan to use this documentation in a hard-copy format, ensure that you have the following books available for your reference.

The following Sun Cluster books support the Sun Cluster 3.1 release. If you are maintaining a different version of Sun Cluster software, refer to the appropriate documentation. All Sun Cluster documentation is available on <http://docs.sun.com>.

Documentation that is not available on <http://docs.sun.com> is listed with the appropriate URL.

TABLE P-2 Hardware Documentation

Title	Part Number
<i>OpenBoot 2.x Command Reference Manual</i>	806-2906
Available on http://docs.sun.com	
<i>Sun StorEdge MultiPack User's Guide</i>	Unavailable online
Shipped with your storage	

TABLE P-3 Sun Cluster Documentation

Application	Title	Part Number
Concepts	<i>Sun Cluster 3.1 Concepts Guide</i>	816-3383
Hardware	<i>Sun Cluster 3.x Hardware Administration Manual</i>	817-0168
	Sun Cluster 3.x Hardware Administration Collection at http://docs.sun.com/db/coll/1024.1/	
Software Installation	<i>Sun Cluster 3.1 Software Installation Guide</i>	816-3388
Data Services	<i>Sun Cluster 3.1 Data Service Planning and Administration Guide</i>	817-1526
	Sun Cluster 3.1 Data Service Collection at http://docs.sun.com/db/coll/573.10/	
API Development	<i>Sun Cluster 3.1 Data Services Developer's Guide</i>	816-3385
Administration	<i>Sun Cluster 3.1 System Administration Guide</i>	816-3384
Error Messages	<i>Sun Cluster 3.1 Error Messages Guide</i>	816-3382
Man Pages	<i>Sun Cluster 3.1 Man Page Reference Manual</i>	816-5251
Release Notes	<i>Sun Cluster 3.1 Release Notes</i>	816-5317
	<i>Sun Cluster 3.1 Release Notes Supplement</i>	816-3381

Using UNIX Commands

This document contains information on commands used to install, configure, or upgrade a Sun Cluster configuration. This document might not contain complete information on basic UNIX[®] commands and procedures such as shutting down the system, booting the system, and configuring devices.

See one or more of the following sources for this information.

- Online documentation for the Solaris software environment
- Other software documentation that you received with your system
- Solaris operating environment man pages

Getting Help

If you have problems installing or using Sun Cluster, contact your service provider and provide the following information.

- Your name and email address (if available)
- Your company name, address, and phone number
- The model number and serial number of your systems
- The release number of the operating environment (for example, Solaris 8)
- The release number of Sun Cluster (for example, Sun Cluster 3.0)

Use the following commands to gather information on your system for your service provider.

Command	Function
<code>prtconf -v</code>	Displays the size of the system memory and reports information about peripheral devices
<code>psrinfo -v</code>	Displays information about processors
<code>showrev -p</code>	Reports which patches are installed
<code>prtdiag -v</code>	Displays system diagnostic information
<code>/usr/cluster/bin/scinstall -pv</code>	Displays Sun Cluster release and package version information

Also have available the contents of the `/var/adm/messages` file.

Accessing Sun Documentation Online

The docs.sun.comSM Web site enables you to access Sun technical documentation online. You can browse the docs.sun.com archive or search for a specific book title or subject. The URL is `http://docs.sun.com`.

Typographic Conventions

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

TABLE P-4 Typographic Conventions

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

Shell Prompts in Command Examples

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

TABLE P-5 Shell Prompts

Shell	Prompt
C shell prompt	<code>machine_name%</code>
C shell superuser prompt	<code>machine_name#</code>
Bourne shell and Korn shell prompt	<code>\$</code>
Bourne shell and Korn shell superuser prompt	<code>#</code>

Installing and Maintaining a StorEdge Multipack Enclosure

This chapter contains the procedures to install and maintain a Sun StorEdge™ MultiPack enclosure in a Sun™ Cluster environment.

This chapter contains the following procedures.

- “How to Install a Storage Array” on page 11
- “How to Add a Disk Drive” on page 17
- “How to Replace a Disk Drive” on page 19
- “How to Remove a Disk Drive” on page 22
- “How to Add a Storage Array” on page 22
- “How to Replace a Storage Array” on page 28
- “How to Remove a Storage Array” on page 30

For conceptual information on multihost disks, see your Sun Cluster concepts documentation.

Installing Storage Arrays

This section contains the procedure for an initial installation of a storage array in a new cluster configuration. To add a storage array to an operating cluster, use the procedure, “How to Add a Storage Array” on page 22 .

▼ How to Install a Storage Array

Use this procedure to install and configure the storage array before installing the Solaris operating environment and Sun Cluster software on your cluster nodes. Perform this procedure in conjunction with the procedures in your Sun Cluster software installation documentation and your server hardware manual.

This procedure uses an updated method for setting the `scsi-initiator-id`. For this storage array, the method published in earlier documentation is still applicable. However, if your cluster configuration uses a Sun StorEdge PCI Dual Ultra3 SCSI host adapter to connect to any other shared storage, you will need to update your `nvrnrc` script and set the `scsi-initiator-id` by using this procedure.

Multihost storage in clusters uses the multi-initiator capability of the Small Computer System Interface (SCSI) specification. For conceptual information on multi-initiator capability, see your Sun Cluster concepts documentation.



Caution – SCSI-reservations failures have been observed with storage arrays that contain a particular models of Quantum disk drives: SUN4.2G VK4550J and SUN2.1GB VK2275J. Do not use this particular models of Quantum disk drives for clustering with storage arrays. If you must use this model of disk drive, you must set the `scsi-initiator-id` of Node 1 to 6. If you use a six-slot storage array, you must set the storage array for the 9-through-14 SCSI target address range. For more information, see the *Sun StorEdge MultiPack Storage Guide*.

The single-ended SCSI specifications specify that bus lengths are based on speed and number of devices. The bus lengths in the following table outline a typical implementation of the single-ended SCSI specifications for Sun hardware. For details, see your host adapter and storage documentation.

TABLE 1-1 Typical Single-Ended, Wide SCSI Bus Lengths

Number of Devices ¹	Maximum Length
<i>UltraSCSI</i> (40MB/s: FAST-20, Wide)	
up to 4	3 meters
up to 8	1.5 meters
<i>FastSCSI</i> (20MB/s: FAST-10, Wide)	
up to 16	6 meters

¹ Devices include both targets and initiators.

If you exceed these specifications, you might experience SCSI errors. The host adapter or the driver might recover from these errors by retrying the request. If this action does not succeed, the host adapter or the driver might recover by renegotiating to a less demanding mode of operation. In some cases, the host adapter or the driver might not be able to recover from these errors, and I/O might fail. You experience delays in I/O if the host adapter or the driver needs to perform this recovery.

If your configuration uses UltraSCSI and requires the 6-meter bus length, use the host adapter driver's `scsi-options` property to limit the speed negotiation to FastSCSI operation. Use the following `/kernel/drv/glm.conf` file as an example to set the `scsi-options` property.

```
name="glm" parent="/pci@1f,4000" unit-address="2" scsi-options=0x3f8;
```

This example uses specific hardware. Change this example to accommodate the hardware in your configuration. In this example, the `scsi-options` property sets the following support.

- Disconnect/reconnect
- Synchronous transfer
- Parity
- Fast SCSI
- Wide SCSI

For more information, see your `isp(7D)`, `glm(7D)` or other host adapter driver man page and documentation on <http://sunsolve.sun.com>.

1. Install the host adapters in the nodes to be connected to the storage array.

For the procedure on how to install a host adapter, see the documentation that shipped with your host adapter and node hardware.

2. Connect the cables to the storage array, as shown in Figure 1-1.

Make sure that the *entire* SCSI bus length to each storage array conforms to the distances in Typical Single-Ended, Wide SCSI Bus Lengths table. This measurement includes the cables to both nodes, as well as the bus length that is internal to each storage array, node, and host adapter. Refer to the documentation that shipped with the storage array for other restrictions about SCSI operation.




FIGURE 1-1 Installing a Mirrored Pair Configuration: An Example

3. **Connect the AC power cord for each storage array of the mirrored pair to a different power source.**

4. **Ensure that each device in the SCSI chain has a unique SCSI address.**

To avoid SCSI-chain conflicts, this procedure instructs you to reserve SCSI address 7 for one host adapter in the SCSI chain and change the other host adapter's global `scsi-initiator-id` to an available SCSI address. Then the procedure instructs you to change the `scsi-initiator-id` for local devices back to 7.

Note – A slot in the storage array might not be in use. However, do not set the `scsi-initiator-id` to a SCSI address for that disk slot. This precaution minimizes future complications if you install additional disk drives.

- a. **If necessary, power on the storage devices.**

For the procedure on powering on a storage device, see the service manual that shipped with your storage device.

- b. **If necessary, power on a node, but do not allow it to boot. If necessary, halt the system to continue with OBP Monitor tasks.**

- c. **Set the `scsi-initiator-id` for one node to 6.**

```
{1} ok setenv scsi-initiator-id 6
scsi-initiator-id = 6
```

- d. **Find the paths to the host adapters that connect to the local disk drives.**

```
{0} ok show-disks
```

Use this information to change the SCSI addresses in the `nvrarc` script. Do not include the `/sd` directories in the device paths.

- e. **Edit the `nvrarc` script to set the `scsi-initiator-id` for the local devices on the first node to 7.**

For a full list of commands, see the *OpenBoot 2.x Command Reference Manual*.



Caution – Insert exactly one space after the first double quote and before `scsi-initiator-id`.

```
{0} ok nvedit
0: probe-all
1: cd /pci@1f,4000/scsi@2
2: 7 " scsi-initiator-id" integer-property
3: device-end
4: cd /pci@1f,4000/scsi@3
5: 7 " scsi-initiator-id" integer-property
6: device-end
```

```

7: install-console
8: banner[Control C]
{0} ok

```

f. Store the changes.

The changes you make through the `nvedit` command are recorded on a temporary copy of the `nvrarc` 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
{1} ok

```

- To discard the changes, type:

```

{0} ok nvquit
{1} ok

```

g. Verify the contents of the `nvrarc` script that you created, as shown in the following example.

If the contents of the `nvrarc` script are incorrect, use the `nvedit` command to make corrections.

```

{0} ok printenv nvrarc
nvrarc =
        probe-all
        cd /pci@1f,4000/scsi@2
        7 "scsi-initiator-id" integer-property
        device-end
        cd /pci@1f,4000/scsi@3
        7 "scsi-initiator-id" integer-property
        device-end
        install-console
        banner
{1} ok

```

h. Instruct the OpenBoot PROM (OBP) Monitor to use the `nvrarc` script, as shown in the following example.

```

{0} ok setenv use-nvrarc? true
use-nvrarc? = true
{1} ok

```

5. Verify that the `scsi-initiator-id` is set correctly on the second node.

- a. If necessary, power on the second node, but do not allow it to boot. If necessary, halt the system to continue with OBP Monitor tasks.
- b. Verify that the `scsi-initiator-id` for each host adapter on the second node is set to 7

Use the `show-disks` command to find the paths to the host adapters that are 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@6,4000/pci@3/scsi@5
{0} ok .properties
scsi-initiator-id      00000007
...
```

6. Continue with the Solaris operating environment, Sun Cluster software, and volume management software installation tasks.

For software installation procedures, see your Sun Cluster software installation documentation.

Maintaining Storage Arrays

This section contains the procedures to maintain a storage array. The following table lists these procedures.

TABLE 1-2 Task Map:Maintaining a Storage Array

Task	Information
Add a disk drive.	"How to Add a Disk Drive" on page 17
Replace a disk drive.	"How to Replace a Disk Drive" on page 19
Remove a disk drive.	"How to Remove a Disk Drive" on page 22
Add a storage array.	"How to Add a Storage Array" on page 22
Replace a storage array.	"How to Replace a Storage Array" on page 28
Remove a storage array.	"How to Remove a Storage Array" on page 30
Replace a SCSI cable.	"How to Replace a SCSI Cable" on page 31
Replace a node's host adapter	"How to Replace a Host Adapter" on page 32

▼ How to Add a Disk Drive

Use this procedure to add a disk drive to a running cluster. Perform the steps in this procedure in conjunction with the procedures in your Sun Cluster software installation documentation and your server hardware manual. “Example—Adding a Disk Drive” on page 18 shows how to apply this procedure.

For conceptual information on quorums, quorum devices, global devices, and device IDs, see your Sun Cluster concepts documentation.



Caution – SCSI-reservations failures have been observed when clustering storage arrays that contain a particular model of Quantum disk drive: SUN4.2G VK4550J. Do not use this particular model of Quantum disk drive for clustering with storage arrays. If you must use this model of disk drive, you must set the `scsi-initiator-id` of the *first node* to 6. If you use a six-slot storage array, you must set the storage array for the 9-through-14 SCSI target address range. For more information, see the *Sun StorEdge MultiPack Storage Guide*.

1. **Locate an empty disk slot in the storage array for the disk drive you are adding.**

Identify the empty slots either by using one of the following methods.

- Observe the disk drive LEDs on the front of the storage array.
- Remove the side-cover of the unit.

The target address IDs that correspond to the slots appear on the middle partition of the drive bay.

2. **Install the disk drive.**

For instructions, see the documentation that shipped with your storage array.

3. **On all nodes that are attached to the storage array, 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 necessary, 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 that is 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.

```
# sddidadm -l
```

Note – As “Example—Adding a Disk Drive” on page 18 shows, DID 35, which is assigned to the new disk drive, might not be in sequential order in the storage array.

8. Perform volume management administration to add the new disk drive to the configuration.

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

Example—Adding a Disk Drive

The following example shows how to apply the procedure on how to add a storage array disk drive.

```
# sddidadm -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 cl
# 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
# sddidadm -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/rdisk/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 your Sun Cluster system administration documentation for the procedure on how to add a quorum device.

▼ How to Replace a Disk Drive

Use this procedure to replace a disk drive in a storage array . Perform the steps in this procedure in conjunction with the procedures in your Sun Cluster software installation documentation 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 your Sun Cluster concepts documentation.



Caution – SCSI-reservations failures have been observed when clustering storage arrays that contain a particular model of Quantum disk drive: SUN4.2G VK4550J. Do not use this particular model of Quantum disk drive for clustering with storage arrays. If you must use this model of disk drive, you must set the `scsi-initiator-id` of the *first node* to 6. If you use a six-slot storage array, you must also set the storage array for the 9-through-14 SCSI target address range. For more information, see the *Sun StorEdge MultiPack Storage Guide*.

1. Identify the failed disk drive.

```
# scdidadm -o diskid -l cNtXdY
```

2. Determine if the disk drive you want to replace is a quorum device.

```
# scstat -q
```

- If the disk is not a quorum device, skip this step.
- If the disk drive you want to replace is a quorum device, add a new quorum device on a different storage device and remove the old quorum device.

For the procedure on how to add and remove a quorum device, see your Sun Cluster system administration documentation.

3. If possible, back up the metadvice or volume.

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

4. If you are using Solstice DiskSuite/Solaris Volume Manager as your volume manager, save the disk partitioning for use when you partition the new disk drive.

```
# 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.

5. Replace the failed disk.

- a. Determine which node owns the device group.

```
# scstat -D
```

- b. If you are using VERITAS Volume Manager, remove the LUNs from the VERITAS Volume Manager control on a node that does not have ownership of the device group.

```
# vxdisk offline cNtXdY
# vxdisk rm cNtXdY
```

- c. On a node that does not have ownership of the device group, suspend activity on the SCSI bus.

```
# cfgadm -x replace_device cX::disk/cXtYdZ
```

When prompted, enter **y** to suspend activity on the SCSI bus.

- d. After SCSI bus activity stops, replace the disk and enter **y** at the prompt.

After replacing the disk, warning messages might display. These messages can be ignored.

- e. Switch the primary for the device group.

```
# scswitch -z -D disk-device-group -h node
```

- f. If you are using VERITAS Volume Manager, remove the LUNs from the VERITAS Volume Manager control on a node that does not have ownership of the device group.

```
# vxdisk offline cNtXdY
# vxdisk rm cNtXdY
```

- g. Suspend activity on the device group's original primary.

```
# cfgadm -x replace_device cX::disk/cXtYdZ
```

When prompted, enter **y** to suspend activity on the SCSI bus.

- h. When prompted to proceed with the disk replacement, enter **y** to proceed.

Warning messages might display. These messages can be ignored.

6. On all nodes that are attached to the storage device, 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 require at least five minutes to complete.

```
# devfsadm
```

7. Label the new disk drive by using the `format` command.
8. If you are using Solstice DiskSuite/Solaris Volume Manager as your volume manager, from any node that is connected to the storage device, partition the new disk drive by using the partitioning you saved.

If you are using VERITAS Volume Manager, skip this step.

```
# fmthard -s filename /dev/rdsk/cNtXdYsZ
```

9. From all nodes that are connected to the disk drive, update the DID database and driver.

```
# sddidadm -R deviceID
```

Note – After running `sddidadm -R` on the first node, each subsequent node that you run the command on might display a warning. This warning can be ignored.

10. From any node, confirm that the failed disk drive has been replaced by comparing the new physical DID to the physical DID that was identified at the beginning of this procedure.

If the new physical DID is different, you successfully replaced the failed disk drive with a new disk drive.

```
# sddidadm -o diskid -l cNtXdY
```

11. Perform volume management administration to add the disk drive back to its diskset or disk group.

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

12. If you want this new disk drive to be a quorum device, add the quorum device.

For procedures on how to add quorum devices, see your Sun Cluster system administration documentation.

▼ How to Remove a Disk Drive

Use this procedure to remove a disk drive from a storage array. Perform the steps in this procedure in your Sun Cluster software installation documentation and your server hardware manual.

For conceptual information on quorum, quorum devices, global devices, and device IDs, see your Sun Cluster concepts documentation.

1. Is the disk drive that you are removing a quorum device?

```
# scstat -q
```

- If no, proceed to Step 2.
- If yes, choose and configure another device to be the new quorum device. Remove the old quorum device.

For procedures on how to add and remove quorum devices, see your Sun Cluster system administration documentation.

2. Perform volume management administration to remove the disk drive from the configuration.

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

3. Identify the disk drive that needs to be removed. Identify 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 information on the procedure on how to remove 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 Storage Array

Use this procedure to add a storage array to a running cluster. Perform the steps in this procedure in conjunction with the procedures in your Sun Cluster software installation documentation and your server hardware manual.

This procedure uses an updated method for setting the `scsi-initiator-id`. For this storage array, the method published in earlier documentation is still applicable. However, if your cluster configuration uses a Sun StorEdge PCI Dual Ultra3 SCSI host adapter to connect to any other shared storage, you will need to update your `nvrarc` script and set the `scsi-initiator-id` by using this procedure.

For conceptual information on multi-initiator SCSI and device IDs, see the Sun Cluster concepts documentation.



Caution – Quorum failures have been observed when clustering storage arrays that contain a particular model of Quantum disk drive: SUN4.2G VK4550J. Do not use this particular model of Quantum disk drive for clustering with storage arrays. If you must use this model of disk drive, you must set the `scsi-initiator-id` of the Node 1 to 6. If you use a six-slot StorEdge Multipack, you must set the storage array for the 9-through-14 SCSI target address range. For more information, see the *Sun StorEdge MultiPack Storage Guide*.

The single-ended SCSI specifications specify that bus lengths are based on speed and number of devices. The bus lengths in the following table outline a typical implementation of the single-ended SCSI specifications for Sun hardware. For details, see your host adapter and storage documentation.

TABLE 1-3 Typical Single-Ended, Wide SCSI Bus Lengths

Number of Devices ¹	Maximum Length
<i>UltraSCSI</i> (40MB/s: FAST-20, Wide)	
up to 4	3 meters
up to 8	1.5 meters
<i>FastSCSI</i> (20MB/s: FAST-10, Wide)	
up to 16	6 meters

¹ Devices include both targets and initiators.

If you exceed these specifications, you might experience SCSI errors. The host adapter or the driver might recover from these errors by retrying the request. If this action does not succeed, the host adapter or the driver might recover by renegotiating to a less demanding mode of operation. In some cases, the host adapter or the driver might not be able to recover from these errors, and I/O might fail. You experience delays in I/O if the host adapter or the driver needs to perform this recovery.

If your configuration uses UltraSCSI and requires the 6-meter bus length, use the host adapter driver's `scsi-options` property to limit the speed negotiation to FastSCSI operation. Use the following `/kernel/drv/glm.conf` file as an example to set the `scsi-options` property.

```
name="glm" parent="/pci@1f,4000" unit-address="2" scsi-options=0x3f8;
```

This example uses specific hardware. Change this example to accommodate the hardware in your configuration. In this example, the `scsi-options` property sets the following support.

- Disconnect/reconnect
- Synchronous transfer
- Parity
- Fast SCSI
- Wide SCSI

For more information, see your `isp(7D)`, `glm(7D)` or other host adapter driver man page and documentation on <http://sunsolve.sun.com>.

1. Shut down and power off Node 1.

For the procedure on how to shut down and power off a node, see your Sun Cluster system administration documentation.

2. Install the host adapters in Node 1.

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

3. Connect the single-ended SCSI cable between the node and the storage arrays, as shown in Figure 1–2.

Make sure that the *entire* SCSI bus length to each storage array conforms to the distances in the Typical Single-Ended, Wide SCSI Bus Lengths table. This measurement includes the cables to both nodes, as well as the bus length that is internal to each storage array, node, and host adapter. Refer to the documentation that shipped with the storage array for other restrictions about SCSI operation.




FIGURE 1–2 Adding a Mirrored Pair Configuration: Part I

4. Install a single-ended terminator on the SCSI IN port of the *second* storage array, as shown in Figure 1–2. This terminator is temporary.
5. Connect each storage array of the mirrored pair to different power sources.
6. Ensure that each device in the SCSI chain has a unique SCSI address.

To avoid SCSI-chain conflicts, this procedure instructs you to reserve SCSI address 7 for one host adapter in the SCSI chain and change the other host adapter's global `scsi-initiator-id` to an available SCSI address. Then the procedure instructs you to change the `scsi-initiator-id` for local devices back to 7.

Note – A slot in the storage array might not be in use. However, do not set the `scsi-initiator-id` to a SCSI address for that disk slot. This precaution minimizes future complications if you install additional disk drives.

- a. If necessary, power on the storage devices.

For the procedure on powering on a storage device, see the service manual that shipped with your storage device.

- b. If necessary, power on a node, but do not allow it to boot. If necessary, halt the system to continue with OBP Monitor tasks.

- c. Set the `scsi-initiator-id` for one node to 6.

```
{1} ok setenv scsi-initiator-id 6
scsi-initiator-id = 6
```

- d. Find the paths to the host adapters that connect to the local disk drives.

```
{0} ok show-disks
```

Use this information to change the SCSI addresses in the `nvrwrc` script. Do not include the `/sd` directories in the device paths.

- e. Edit the `nvrwrc` script to set the `scsi-initiator-id` for the local devices on the first node to 7.

For a full list of commands, see the *OpenBoot 2.x Command Reference Manual*.



Caution – Insert exactly one space after the first double quote and before `scsi-initiator-id`.

```
{0} ok nvedit
0: probe-all
1: cd /pci@1f,4000/scsi@2
2: 7 " scsi-initiator-id" integer-property
3: device-end
4: cd /pci@1f,4000/scsi@3
```

```

5: 7 " scsi-initiator-id" integer-property
6: device-end
7: install-console
8: banner[Control C]
{0} ok

```

f. Store the changes.

The changes you make through the `nvedit` command are recorded on a temporary copy of the `nvrsrc` 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
{1} ok

```

- To discard the changes, type:

```

{0} ok nvquit
{1} ok

```

g. Verify the contents of the `nvrsrc` script that you created, as shown in the following example.

If the contents of the `nvrsrc` script are incorrect, use the `nvedit` command to make corrections.

```

{0} ok printenv nvrsrc
nvrsrc =
        probe-all
        cd /pci@1f,4000/scsi@2
        7 " scsi-initiator-id" integer-property
        device-end
        cd /pci@1f,4000/scsi@3
        7 " scsi-initiator-id" integer-property
        device-end
        install-console
        banner
{1} ok

```

h. Instruct the OpenBoot PROM (OBP) Monitor to use the `nvrsrc` script, as shown in the following example.

```

{0} ok setenv use-nvrsrc? true
use-nvrsrc? = true
{1} ok

```

7. Perform a reconfiguration boot on Node 1 to create the new Solaris device files and links.

8. On all nodes, verify that the DIDs have been assigned to the disk drives in the storage array.

```
# scdidadm -l
```

9. Shut down and power off the *second* node.

For the procedure on how to shut down and power off a node, see your Sun Cluster system administration documentation.

10. Install the host adapters in the second node.

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

11. Remove the SCSI terminator that you installed in Step 4.

12. Connect the storage arrays to the host adapters by using single-ended SCSI cables, as show in Figure 1-3.




FIGURE 1-3 Adding a Mirrored Pair COnfiguration: Part II

13. Verify that the `scsi-initiator-id` is set correctly on the second node.

- a. If necessary, power on the second node, but do not allow it to boot. If necessary, halt the system to continue with OBP Monitor tasks.
- b. Verify that the `scsi-initiator-id` for each host adapter on the second node is set to 7

Use the `show-disks` command to find the paths to the host adapters that are 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@6,4000/pci@3/scsi@5
{0} ok .properties
scsi-initiator-id      00000007
...
```

14. Perform a reconfiguration boot on the second node to create the new Solaris device files and links.

15. On all nodes, verify that the DIDs have been assigned to the disk drives in the storage array.

```
# scdidadm -l
```

16. Perform volume management administration to add the disk drives in the storage array to the volume management configuration.

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

▼ How to Replace a Storage Array

Use this procedure to replace a storage array in a running cluster. This procedure assumes that you are retaining the disk drives in the storage array that you are replacing. This procedure also assumes that you are retaining the references to these disk drives.

If you want to replace your disk drives, see “How to Replace a Disk Drive” on page 19.

1. **If possible, back up the metadevices or volumes that reside in the disk array.**

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

3. **Disconnect the SCSI cables from the storage array. Disconnect the SCSI cable on the SCSI OUT connector, and then disconnect the cable on the SCSI IN connector as shown in Figure 1-4.**




FIGURE 1-4 Replacing a Storage Array

4. **Power off, and disconnect the storage array from the AC power source.**
For more information, see the documentation that shipped with your storage array and the labels inside the lid of the storage array.
5. **Connect the new storage array to an AC power source.**
Refer to the documentation that shipped with the storage array and the labels inside the lid of the storage array.
6. **Connect the SCSI cables to the new disk array. Reverse the order in which you disconnected the cables as shown in Figure 1-4.**
7. **One disk drive at a time, remove the disk drives from the old disk array. Insert the disk drives into the same slots in the new disk array.**
8. **Power on the storage array.**
9. **On all nodes that are attached to the storage array, run the `devfsadm(1M)` command.**

```
# devfsadm
```
10. **One at a time, shut down and reboot the nodes that are connected to the storage array.**
For the procedure on how to shut down and power off a node, see your Sun Cluster system administration documentation.
11. **Perform volume management administration to add the storage array to the configuration.**
For more information, see your Solstice DiskSuite/Solaris Volume Manager or VERITAS Volume Manager documentation.

▼ How to Remove a Storage Array

Use this procedure to remove a storage array from a running cluster. This procedure assumes that you want to remove the references to the disk drives in the storage array.

1. **Perform volume management administration to remove the storage array from the configuration.**

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

2. **Disconnect the SCSI cables from the storage array. Disconnect the SCSI cables in the order that is shown in Figure 1–5.**




FIGURE 1–5 Removing a Storage Array

3. **Power off and disconnect the storage array from the AC power source.**

For more information, see the documentation that shipped with the storage array and the labels inside the lid of the storage array.

4. **Remove the storage array.**

For the procedure on how to remove an storage array, see the *Sun StorEdge MultiPack Storage Guide*.

5. **Identify the disk drives that you need to remove from the cluster.**

```
# cftadm -al
```

6. **On all nodes, remove references to the disk drives that were in the storage array. This storage array is the storage array that you removed in Step 4.**

```
# cftadm -c unconfigure cN::dsk/cNtXdY
# devfsadm -C
# scdidadm -C
```

7. If you plan to remove a host adapter that has an entry in the `nvrामrc` script, delete the references to the host adapters in the `nvrामrc` script.

Note – If there are no other parallel SCSI devices connected to the nodes, you can delete the contents of the `nvrामrc` script and, at the OpenBoot PROM, set `setenv use-nvrामrc? false`.

8. If necessary, remove any unused host adapters from the nodes.

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

▼ How to Replace a SCSI Cable

Use this procedure to replace a failed SCSI cable in a running cluster.

1. Is one of the disks in the storage array a quorum device? This array is the array that is attached to the cable you are replacing.

```
# scstat -q
```

- If no, skip to Step 3
- If yes, relocate that quorum device to another suitable storage array

For procedures on how to add and remove quorum devices, see your Sun Cluster system administration documentation.

2. Detach the submirror(s) on the storage array.

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

3. Replace the SCSI cable.

4. Reattach the submirror(s) on the storage array.

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

5. (Optional) Did you relocate the quorum device in Step 1?

- If no, this step completes the procedure
- If yes, relocate the quorum device function back to this storage array

For procedures on how to add and remove quorum devices, see your Sun Cluster system administration documentation.

▼ How to Replace a Host Adapter

Use this procedure to replace a node's host adapter in a running cluster. This procedure defines Node A as the node with the failed host adapter you are replacing.

1. Determine the resource groups and device groups running on Node A.

Record this information because you use this information in Step 13 of this procedure to return resource groups and device groups to Node A.

```
# scstat
```

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

```
# scswitch -S -h nodename
```

3. Is one of the disks in the storage array a quorum device? This array is the array that is attached to the host adapter you are replacing.

```
# scstat -q
```

- If no, proceed to Step 4

- If yes, relocate that quorum device to another suitable storage array

For procedures on adding and removing quorum devices, see your Sun Cluster system administration documentation.

4. Detach the submirror(s) on the storage array.

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

5. Is Node A enabled with the Solaris dynamic reconfiguration (DR) feature?

- If yes, skip to Step 8
- If no, proceed to Step 6

For more information, see your Sun Cluster system administration documentation.

6. Shut down Node A.

For the full procedure about how to shut down and power off a node, see your Sun Cluster system administration documentation.

7. Power off Node A.

8. Replace the failed host adapter.

For the procedure about how to remove and add host adapters, see the documentation that shipped with your nodes.

9. Is Node A enabled with the Solaris dynamic reconfiguration (DR) feature?

- If yes, skip to Step 12
- If no, proceed to Step 10

For more information, see your Sun Cluster system administration documentation.

10. Power on Node A.

11. Boot Node A into cluster mode.

For more information about how to boot nodes, see your Sun Cluster system administration documentation.

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

See your Sun Cluster release notes documentation for information about how to access Sun's EarlyNotifier web pages, which list information about patches or firmware levels that are available for download. For the procedure about how to apply any host adapter firmware patch, see the firmware patch README file.

13. (Optional) Return the resource groups and device groups you identified in Step 1 to Node A.

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

For more information, see your Sun Cluster system administration documentation.

14. Reattach the submirror(s) on the storage array.

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

15. (Optional) Did you relocate the quorum device in Step 1?

- If no, this step completes the procedure
- If yes, relocate the quorum device function back to this storage array

For procedures on adding and removing quorum devices, see your Sun Cluster system administration documentation.

Index

A

adding
 See also installing
 disk drives, 17
 storage arrays, 22
arrays, *See* storage arrays

D

disk drives
 adding, 17
 removing, 22
 replacing, 19

E

enclosures, *See* storage arrays

H

help, 8
host adapters, replacing, 32

I

installing
 See also adding
 storage arrays, 11

R

removing
 disk drives, 22
 storage arrays, 30
replacing
 disk drives, 19
 host adapters, 32
 SCSI cable, 31
 storage arrays, 28

S

SCSI cable, replacing, 31
storage arrays
 adding, 22
 installing, 11
 removing, 30
 replacing, 28

T

technical support, 8

