

# **Sun Cluster 3.1 - 3.2 with StorageTek RAID Arrays Manual for Solaris OS**



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

Part No: 819-7306-12  
January 2009, Revision 12

Copyright 2009 Sun Microsystems, Inc. 4150 Network Circle, Santa Clara, CA 95054 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 U.S. patents or pending patent applications in the U.S. and in other countries.

U.S. Government Rights – Commercial software. Government users are subject to the Sun Microsystems, Inc. standard license agreement and applicable provisions of the FAR and its supplements.

This distribution may include materials developed by third parties.

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, the Solaris logo, the Java Coffee Cup logo, docs.sun.com, Sun StorEdge, StorageTek, Java, and Solaris are trademarks or registered trademarks of Sun Microsystems, Inc. or its subsidiaries 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.

Products covered by and information contained in this publication are controlled by U.S. Export Control laws and may be subject to the export or import laws in other countries. Nuclear, missile, chemical or biological weapons or nuclear maritime end uses or end users, whether direct or indirect, are strictly prohibited. Export or reexport to countries subject to U.S. embargo or to entities identified on U.S. export exclusion lists, including, but not limited to, the denied persons and specially designated nationals lists is strictly prohibited.

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 2009 Sun Microsystems, Inc. 4150 Network Circle, Santa Clara, CA 95054 U.S.A. Tous droits réservés.

Sun Microsystems, Inc. détient les droits de propriété intellectuelle relatifs à la technologie incorporée dans le produit qui est décrit dans ce document. En particulier, et ce sans limitation, ces droits de propriété intellectuelle peuvent inclure un ou plusieurs brevets américains ou des applications de brevet en attente aux Etats-Unis et dans d'autres pays.

Cette distribution peut comprendre des composants développés par des tierces personnes.

Certaines composants de ce produit peuvent être dérivées du logiciel Berkeley BSD, licenciés par l'Université de Californie. UNIX est une marque déposée aux Etats-Unis et dans d'autres pays; elle est licenciée exclusivement par X/Open Company, Ltd.

Sun, Sun Microsystems, le logo Sun, le logo Solaris, le logo Java Coffee Cup, docs.sun.com, Sun StorEdge, StorageTek, Java et Solaris sont des marques de fabrique ou des marques déposées de Sun Microsystems, Inc., ou ses filiales, 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.

Les produits qui font l'objet de cette publication et les informations qu'il contient sont régis par la législation américaine en matière de contrôle des exportations et peuvent être soumis au droit d'autres pays dans le domaine des exportations et importations. Les utilisations finales, ou utilisateurs finaux, pour des armes nucléaires, des missiles, des armes chimiques ou biologiques ou pour le nucléaire maritime, directement ou indirectement, sont strictement interdites. Les exportations ou réexportations vers des pays sous embargo des Etats-Unis, ou vers des entités figurant sur les listes d'exclusion d'exportation américaines, y compris, mais de manière non exclusive, la liste de personnes qui font l'objet d'un ordre de ne pas participer, d'une façon directe ou indirecte, aux exportations des produits ou des services qui sont régis par la législation américaine en matière de contrôle des exportations et la liste de ressortissants spécifiquement désignés, sont rigoureusement interdites.

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 CONTREFACON.

# Contents

---

<b>Preface .....</b>	5
<b>1    Restrictions and Requirements .....</b>	11
Requirements .....	11
Restrictions .....	11
<b>2    Installing and Configuring a StorageTek Array .....</b>	13
Installing Storage Arrays .....	13
Storage Array Cabling Configurations .....	13
How to Install Storage Arrays in a New Cluster .....	16
▼ How to Add Storage Arrays to an Existing Cluster .....	17
Configuring Storage Arrays .....	18
▼ How to Create a Logical Volume .....	19
▼ How to Remove a Logical Volume .....	21
<b>3    Maintaining a StorageTek Array .....</b>	25
FRUs That <i>Do Not</i> Require Sun Cluster Maintenance Procedures .....	25
Maintaining Storage Arrays .....	26
▼ How to Upgrade Storage Array Firmware .....	26
▼ How to Remove a Storage Array .....	27
<b>Index .....</b>	29



# Preface

---

The *Sun Cluster 3.1 - 3.2 with StorageTek RAID Arrays Manual for Solaris OS* provides procedures that are specific to StorageTek™ RAID arrays that are placed in a Sun™ Cluster environment.

Use this manual with any version of Sun Cluster 3.1 or 3.2 software on SPARC® based clusters and on any supported version of Sun Cluster 3.1 or 3.2 software on x86 based clusters. Unless otherwise noted, procedures are the same for all supported Sun Cluster 3.1 or 3.2 versions. See the “[Revision History](#)” on page 6 for a list of changes to this manual.

---

**Note** – This Sun Cluster release supports systems that use the SPARC and x86 families of processor architectures: UltraSPARC, SPARC64, AMD64, and Intel 64. In this document, x86 refers to the larger family of 64-bit x86 compatible products. Information in this document pertains to all platforms unless otherwise specified.

---

This book assumes that you are performing one or more of the following tasks:

- You want to replace an array component to *prevent* a failure.
- You want to replace an array component because you have an *existing* failure.
- You want to add (to an *established* cluster) or install (to a *new* cluster) a storage array.

## 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 the following chapters:

- [Chapter 1, “Restrictions and Requirements,”](#) lists limitations of your use of StorageTek storage arrays in a Sun Cluster environment.
- [Chapter 2, “Installing and Configuring a StorageTek Array,”](#) discusses how to install StorageTek storage arrays and how to configure logical units on them.
- [Chapter 3, “Maintaining a StorageTek Array,”](#) describes how to maintain StorageTek storage arrays 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 3.1 - 3.2 with StorageTek RAID Arrays Manual for Solaris OS

Revision Date	New Information
May 2008	General edits to make the guide generic to all types of RAID arrays.
March 2008	Replaced outdated information about SunSolve with information about <a href="#">Sun Connection Update Manager</a> .

# Related Documentation

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

The following Sun Cluster books support the Sun Cluster 3.1 and 3.2 releases. If you are maintaining a different version of Sun Cluster software, refer to the appropriate documentation. All Sun Cluster documentation is available at <http://docs.sun.com>. Documentation that is not available at <http://docs.sun.com> is listed with the appropriate URL.

Refer to your storage array's documentation on [docs.sun.com](http://docs.sun.com) for detailed information on each storage array. For the StorageTek array, refer to the online product documentation

**TABLE P-2** Sun Cluster Documentation

Documentation
<a href="#">Solaris Cluster 3.2</a>
<a href="#">Sun Cluster 3.1</a>

## Using UNIX Commands

This document contains information about commands that are used to install, configure, or upgrade a Sun Cluster configuration. This document might not contain complete information about 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™ Operating System (Solaris OS)
- Other software documentation that you received with your system
- Solaris Operating System 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 10)
- The release number of Sun Cluster (for example, Sun Cluster 3.2)

Use the following commands to gather information about 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

Command	Function
<code>/usr/cluster/bin/clnode show-rev</code>	Displays Sun Cluster release and package version information
<code>/usr/cluster/bin/scinstall -pv</code>	

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

## Documentation, Support, and Training

The Sun web site provides information about the following additional resources:

- Documentation (<http://www.sun.com/documentation/>)
- Support (<http://www.sun.com/support/>)
- Training (<http://www.sun.com/training/>)

## Sun Welcomes Your Comments

Sun is interested in improving its documentation and welcomes your comments and suggestions. To share your comments, go to <http://docs.sun.com> and click Feedback.

## Typographic Conventions

The following table describes the typographic conventions that are used in this book.

TABLE P-3 Typographic Conventions

Typeface	Meaning	Example
AaBbCc123	The names of commands, files, and directories, and onscreen computer output	Edit your <code>.login</code> file. Use <code>ls -a</code> to list all files. <code>machine_name% you have mail.</code>
AaBbCc123	What you type, contrasted with onscreen computer output	<code>machine_name% su</code> <code>Password:</code>
aabbcc123	Placeholder: replace with a real name or value	The command to remove a file is <code>rm filename</code> .

**TABLE P-3** Typographic Conventions *(Continued)*

Typeface	Meaning	Example
<i>AaBbCc123</i>	Book titles, new terms, and terms to be emphasized	<p>Read Chapter 6 in the <i>User's Guide</i>.</p> <p>A <i>cache</i> is a copy that is stored locally.</p> <p>Do <i>not</i> save the file.</p> <p><b>Note:</b> Some emphasized items appear bold online.</p>

## Shell Prompts in Command Examples

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

**TABLE P-4** Shell Prompts

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



## Restrictions and Requirements

---

This chapter includes only restrictions and requirements that have a direct impact on the procedures in this book. For general support information, contact your Sun service provider.

### Requirements

If you are replacing a Host Bus Adapter (HBA), you must re-mask the HBA's World Wide Name (WWN) numbers to the respective Logical Unit Numbers (LUNs) on the array.

### Restrictions

When using arrays that support storage-based replication (for example a StorageTek 9900 array), do not configure a replicated volume as a quorum device. Locate any quorum devices on an unreplicated volume. See “[Using Storage-Based Data Replication](#)” in *Sun Cluster System Administration Guide for Solaris OS* for more information on storage-based replication.



## Installing and Configuring a StorageTek Array

---

This chapter contains the procedures about how to install and configure StorageTek RAID arrays. These procedures are specific to a Sun™ Cluster environment.

This chapter contains the following main topics:

- “[Installing Storage Arrays](#)” on page 13
- “[Configuring Storage Arrays](#)” on page 18

For detailed information about storage array architecture, features, configuration utilities, and installation, see “[Related Documentation](#)” on page 6.

### Installing Storage Arrays

This section contains the procedures listed in [Table 2–1](#).

**TABLE 2–1** Task Map: Installing Storage Arrays

Task	Information
Install a storage array in a new cluster, before the OS and Sun Cluster software are installed.	<a href="#">“How to Install Storage Arrays in a New Cluster” on page 16</a>
Add a storage array to an existing cluster.	<a href="#">“How to Add Storage Arrays to an Existing Cluster” on page 17</a>

### Storage Array Cabling Configurations

You can install your storage array in several different configurations; see [Figure 2–1](#) through [Figure 2–4](#) for examples.

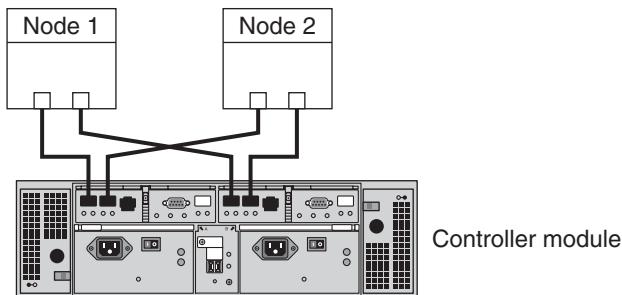


FIGURE 2–1 StorageTek Array Direct-Connect Configuration

The StorageTek 6140 array houses two controllers; each controller has four host ports. The cabling approach is the same as shown in [Figure 2–1](#), but it can support up to four nodes in a direct-attach configuration.

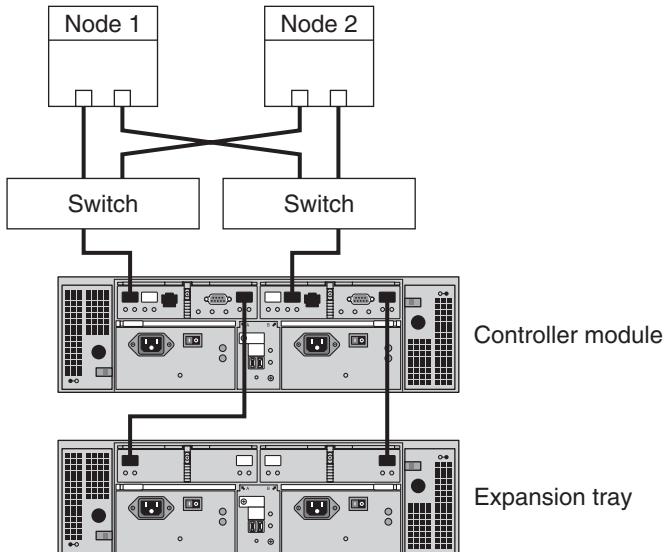
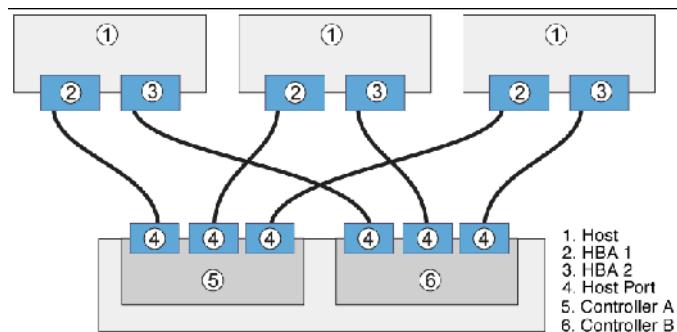


FIGURE 2–2 StorageTek Array Switched Configuration

[Figure 2–2](#) shows a switched configuration for a two-node storage array.



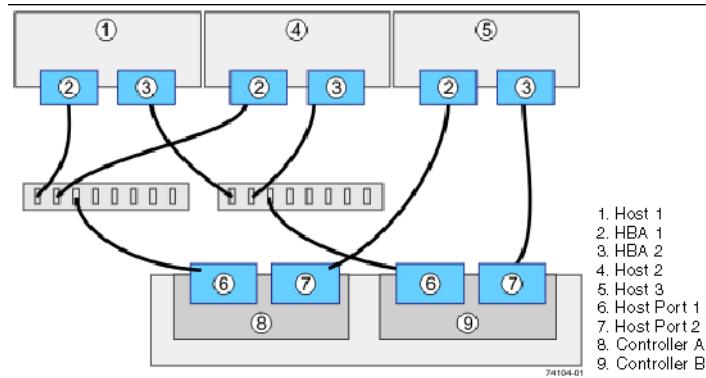
**FIGURE 2–3** Direct Connections from Three Data Hosts with Dual HBAs

You can connect one or more hosts to a storage array. Figure 2–3 shows an example of a direct host connection from each data host with dual HBAs.

---

**Note** – For maximum hardware redundancy, you should install a minimum of two HBAs in each host and distribute I/O paths between these HBAs. A single, dual-port HBA can provide both data paths to the storage array but does not ensure redundancy if the HBA fails.

---



**FIGURE 2–4** Mixed Topology—Three Hosts Connected Through a Switch or Connected Directly

Figure 2–4 shows that three hosts can be connected directly or through a switch.

# How to Install Storage Arrays in a New Cluster

Use this procedure to install a storage array in a new cluster. To add a storage array to an existing cluster, use the procedure in “How to Replace a Host Adapter” in *Sun Cluster 3.1 - 3.2 With Sun StorEdge A3500FC System Manual for Solaris OS*.

This procedure relies on the following assumptions:

- You have *not* installed the Solaris Operating System.
- You have *not* installed the Sun Cluster software.
- You have enough host adapters to connect the nodes and the storage array.

## ▼ Install and Cable the Hardware

### 1 Unpack, place, and level the storage array.

For instructions, see the StorageTek online documentation.

### 2 If necessary, install the Fibre Channel (FC) switch for the storage array (if the switch is not already installed).

For the procedure about how to install an FC switch, see the documentation that shipped with your FC switch hardware.

### 3 Connect the nodes to the storage array.

- SAN Configuration — Connect the FC switches to the storage array
- Direct-Attached Configuration — Connect each node directly to the storage array
- SAS Direct-Attached Configuration
- iSCSI Direct-Attached Configuration
- iSCSI Switched Configuration

For instructions, see your storage array documentation and the “Related Documentation” on page 6 section.

### 4 Hook up the cards for the storage array.

For instructions, see your storage array documentation.

### 5 Power on the storage array and the nodes.

For instructions, see your storage array documentation.

**6 Configure the storage array, if needed.**

For instructions, see “[Configuring Storage Arrays](#)” on page 18 and consult your storage array documentation.

▼ **Install the Solaris Operating System and Configure Multipathing**

**1 On all nodes, install the Solaris operating system and apply the required Solaris patches for Sun Cluster software and storage array support.**

For the procedure about how to install the Solaris operating environment, see “[How to Install Solaris Software](#)” in *Sun Cluster Software Installation Guide for Solaris OS*.

**2 Install any required patches or software for Solaris I/O multipathing software support to nodes and enable multipathing.**

For the procedure about how to install the Solaris I/O multipathing software, see “[How to Install Sun Multipathing Software](#)” in *Sun Cluster Software Installation Guide for Solaris OS*.

**See Also**

- To create a logical volume, see “[How to Create a Logical Volume](#)” on page 19.
- To continue with Sun Cluster software installation tasks, see your Sun Cluster software installation documentation.

▼ **How to Add Storage Arrays to an Existing Cluster**

Use this procedure to add a new storage array to a *running* cluster. To install a new storage array in a Sun Cluster configuration that is not running (the nodes are in noncluster mode), use the procedure in “[How to Install Storage Arrays in a New Cluster](#)” on page 16.

**Before You Begin**

This procedure relies on the following assumptions:

- (*Veritas Volume Manager Only*) You have a version of Veritas Volume Manager that includes Array Support Library (ASL).
- You have enough host adapters to connect the nodes and the storage array.  
If you need to install host adapters, see “[How to Replace a Host Adapter](#)” in *Sun Cluster 3.1 - 3.2 With Sun StorEdge A3500FC System Manual for Solaris OS*. When this procedure asks you to replace the failed host adapter, install the new host adapter instead.
- All cluster nodes have joined the cluster.  
If you need to add a node to your cluster, see your Sun Cluster system administration documentation. Ensure that you install the required Solaris patches for storage array support.

**1 Unpack, place, and level the storage array.**

For instructions, see the StorageTek online documentation.

**2 If necessary, install the Fibre Channel (FC) switch for the storage array (if the switch is not already installed).**

For the procedure about how to install an FC switch, see the documentation that shipped with your FC switch hardware.

**3 Connect the nodes to the storage array.**

- **SAN Configuration — Connect the FC switches to the storage array**
- **Direct-Attached Configuration — Connect each node directly to the storage array**
- **SAS Direct-Attached Configuration**
- **iSCSI Direct-Attached Configuration**
- **iSCSI Switched Configuration**

For instructions, see your storage array documentation and the “[Related Documentation](#)” on [page 6](#) section.

**4 Hook up the cards for the storage array.**

For instructions, see your storage array documentation.

**5 Power on the storage array and the nodes.**

For instructions, see your storage array documentation.

**6 Configure the storage array, if needed.**

For instructions, see “[Configuring Storage Arrays](#)” on [page 18](#) and consult your storage array documentation.

**See Also** □ To create a logical volume, see “[How to Create a Logical Volume](#)” on [page 19](#).

## Configuring Storage Arrays

This section contains the procedures to configure a storage array in a running cluster. [Table 2–2](#) lists these procedures.

**TABLE 2-2** Task Map: Configuring a Storage Array

Task	Information
Create a logical volume	<a href="#">“How to Create a Logical Volume” on page 19</a>
Remove a logical volume	<a href="#">“How to Remove a Logical Volume” on page 21</a>

The following is a list of administrative tasks that do not require cluster-specific procedures. See the storage array's documentation [“Related Documentation” on page 6](#) for the following procedures.

- Creating a storage pool
- Removing a storage pool
- Creating a volume group
- Removing a volume group
- Creating an initiator group
- Adding an initiator group
- Removing an initiator group

## ▼ How to Create a Logical Volume

Use this procedure to create a logical volume from unassigned storage capacity.

---

**Note** – Sun storage documentation uses the following terms:

- Logical volume
- Logical device
- Logical unit number (LUN)

This manual uses *logical volume* to refer to all such logical constructs.

---

**Before You Begin** This procedure relies on the following prerequisites and assumptions.

- All nodes are booted in cluster mode and attached to the storage device.
- The storage device is installed and configured. If you are using multipathing, the storage device is configured as described in the installation procedure.
- If you are using Solaris I/O multipathing (MPxIO) for the Solaris 10 OS, previously called Sun StorEdge Traffic Manager in the Solaris 9 OS, verify that it is installed and configured and the path to the storage device is functioning. To configure the Traffic Manager for the Solaris 9 OS, see the *Sun StorEdge Traffic Manager Installation and Configuration Guide*. To configure multipathing for the Solaris 10 OS, see the *Solaris Fibre Channel Storage Configuration and Multipathing Support Guide*.

This procedure provides the long forms of the Sun Cluster commands. Most commands also have short forms. Except for the forms of the command names, the commands are identical. For a list of the commands and their short forms, see [Appendix A, “Sun Cluster Object-Oriented Commands,” in \*Sun Cluster 3.1 - 3.2 Hardware Administration Manual for Solaris OS\*.](#)

- 1 Become superuser or assume a role that provides `solaris.cluster.modify` role-based access control (RBAC) authorization.**
- 2 Follow the instructions in your storage device's documentation to create and map the logical volume. For a URL to this storage documentation, see “[Related Documentation](#)” on page 6.**
  - Completely set up the logical volume. When you are finished, the volume must be created, mapped, mounted, and initialized.
  - If necessary, partition the volume.
  - To allow multiple clusters and nonclustered nodes to access the storage device, create initiator groups by using LUN masking.
- 3 If you are not using multipathing, skip to [Step 5](#).**
- 4 If you are using multipathing, and if any devices that are associated with the volume you created are at an unconfigured state, configure the multipathing paths on each node that is connected to the storage device.**

To determine whether any devices that are associated with the volume you created are at an unconfigured state, use the following command.

```
# cfgadm -al | grep disk
```

---

**Note –** To configure the Solaris I/O multipathing paths on each node that is connected to the storage device, use the following command.

```
# cfgadm -o force_update -c configure controllerinstance
```

---

To configure the Traffic Manager for the Solaris 9 OS, see the *Sun StorEdge Traffic Manager Installation and Configuration Guide*. To configure multipathing for the Solaris 10 OS, see the *Solaris Fibre Channel Storage Configuration and Multipathing Support Guide*.

- 5 On one node that is connected to the storage device, use the `format` command to label the new logical volume.**
- 6 From any node in the cluster, update the global device namespace.**
  - **If you are using Sun Cluster 3.2, use the following command:**  
`# cldevice populate`

- If you are using Sun Cluster 3.1, use the following command:

```
# scgdevs
```

---

**Note** – You might have a volume management daemon such as vold running on your node, and have a DVD drive connected to the node. Under these conditions, a device busy error might be returned even if no disk is inserted in the drive. This error is expected behavior. You can safely ignore this error message.

---

- 7 To manage this volume with volume management software, use Solaris Volume Manager or Veritas Volume Manager commands to update the list of devices on all nodes that are attached to the new volume that you created.

For more information, see your Solaris Volume Manager or Veritas Volume Manager documentation.

**See Also**

- To configure a logical volume as a quorum device, see Chapter 6, “Administering Quorum,” in *Sun Cluster System Administration Guide for Solaris OS*.
- To create a new resource or configure a running resource to use the new logical volume, see Chapter 2, “Administering Data Service Resources,” in *Sun Cluster Data Services Planning and Administration Guide for Solaris OS*.

## ▼ How to Remove a Logical Volume

Use this procedure to remove a logical volume. This procedure defines Node A as the node with which you begin working.

---

**Note** – Sun storage documentation uses the following terms:

- Logical volume
- Logical device
- Logical unit number (LUN)

This manual uses *logical volume* to refer to all such logical constructs.

---

**Before You Begin**

This procedure relies on the following prerequisites and assumptions.

- All nodes are booted in cluster mode and attached to the storage device.
- The logical volume and the path between the nodes and the storage device are both operational.

This procedure provides the long forms of the Sun Cluster commands. Most commands also have short forms. Except for the forms of the command names, the commands are identical. For

a list of the commands and their short forms, see [Appendix A, “Sun Cluster Object-Oriented Commands,” in \*Sun Cluster 3.1 - 3.2 Hardware Administration Manual for Solaris OS\*](#).

- 1 Become superuser or assume a role that provides `solaris.cluster.read` and `solaris.cluster.modify` RBAC authorization.**

- 2 Identify the logical volume that you are removing.**

Refer to your Solaris Volume Manager or Veritas Volume Manager documentation for more information.

- 3 (Optional) Migrate all data off the logical volume that you are removing. Alternatively, back up that data.**

- 4 If the LUN that you are removing is configured as a quorum device, choose and configure another device as the quorum device. Then remove the old quorum device.**

To determine whether the LUN is configured as a quorum device, use one of the following commands.

- **If you are using Sun Cluster 3.2, use the following command:**

```
# clquorum show
```

- **If you are using Sun Cluster 3.1, use the following command:**

```
# scstat -q
```

For procedures about how to add and remove quorum devices, see [Chapter 6, “Administering Quorum,” in \*Sun Cluster System Administration Guide for Solaris OS\*](#).

- 5 If you are using volume management software, use that software to update the list of devices on all nodes that are attached to the logical volume that you are removing.**

For instructions about how to update the list of devices, see your Solaris Volume Manager or Veritas Volume Manager documentation.

- 6 If you are using volume management software, run the appropriate Solaris Volume Manager or Veritas Volume Manager commands to remove the logical volume from any diskset or disk group.**

For more information, see your Solaris Volume Manager or Veritas Volume Manager documentation.

---

**Note** – Volumes that were managed by Veritas Volume Manager must be completely removed from Veritas Volume Manager control before you can delete them from the Sun Cluster environment. After you delete the volume from any disk group, use the following commands on both nodes to remove the volume from Veritas Volume Manager control.

```
# vxdisk offline Accessname  
# vxdisk rm Accessname  
  
Accessname      Disk access name
```

---

**7 If you are using multipathing, unconfigure the volume in Solaris I/O multipathing.**

```
# cfgadm -o force_update -c unconfigure Logical_Volume
```

**8 Access the storage device and remove the logical volume.**

To remove the volume, see your storage documentation. For a list of storage documentation, see “[Related Documentation](#)” on page 6.

**9 Determine the resource groups and device groups that are running on all nodes.**

Record this information because you use it in [Step 14](#) and [Step 15](#) of this procedure to return resource groups and device groups to these nodes.

- **If you are using Sun Cluster 3.2, use the following commands:**

```
# clresourcegroup status +  
# cldevicegroup status +
```

- **If you are using Sun Cluster 3.1, use the following command:**

```
# scstat
```

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

- **If you are using Sun Cluster 3.2, use the following command:**

```
# clnode evacuate nodename
```

- **If you are using Sun Cluster 3.1, use the following command:**

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

**11 Shut down and reboot Node A.**

To shut down and boot a node, see [Chapter 3, “Shutting Down and Booting a Cluster,” in \*Sun Cluster System Administration Guide for Solaris OS\*.](#)

- 12 On Node A, remove the paths to the logical volume that you removed. Remove obsolete device IDs.
  - If you are using Sun Cluster 3.2, use the following command:

```
# devfsadm -C  
# cldevice clear
```
  - If you are using Sun Cluster 3.1, use the following command:

```
# devfsadm -C  
# scdidadm -C
```
- 13 For each additional node that is connected to the shared storage that hosted the logical volume, repeat Step 9 to Step 12.
- 14 (Optional) Restore the device groups to the original node.

Do the following for each device group that you want to return to the original node.

  - If you are using Sun Cluster 3.2, use the following command:

```
# cldevicegroup switch -n nodename devicegroup1[ devicegroup2 ...]  
-n nodename           The node to which you are restoring device groups.  
devicegroup1[ devicegroup2 ...]   The device group or groups that you are restoring to the node.
```
  - If you are using Sun Cluster 3.1, use the following command:

```
# scswitch -z -D devicegroup -h nodename
```
- 15 (Optional) Restore the resource groups to the original node.

Do the following for each resource group that you want to return to the original node.

  - If you are using Sun Cluster 3.2, use the following command:

```
# clresourcegroup switch -n nodename resourcegroup1[ resourcegroup2 ...]  
nodename               For failover resource groups, the node to which the groups are returned. For scalable resource groups, the node list to which the groups are returned.  
resourcegroup1[ resourcegroup2 ...]   The resource group or groups that you are returning to the node or nodes.
```
  - If you are using Sun Cluster 3.1, use the following command:

```
# scswitch -z -g resourcegroup -h nodename
```

## Maintaining a StorageTek Array

---

This chapter contains the procedures about how to maintain a StorageTek array. These procedures are specific to a Sun™ Cluster environment.

This chapter contains the following procedures:

- “[How to Upgrade Storage Array Firmware](#)” on page 26
- “[How to Remove a Storage Array](#)” on page 27

For detailed information about storage array architecture, features, and configuration utilities, see the StorageTek documentation listed in “[Related Documentation](#)” on page 6.

## FRUs That Do Not Require Sun Cluster Maintenance Procedures

In general, the following is a list of administrative tasks that require no cluster-specific procedures. See the base-product documentation for these procedures. Refer to your documentation for components not listed below.

Depending on your configuration type and the state of your cluster, a few of the following FRUs might require cluster-specific steps. Some FRUs include the DSP and the storage array.

- Adding a disk drive
- Replacing a storage array's chassis
- Replacing an Ethernet cable
- Replacing a power supply
- Replacing the power cable on the storage array
- Replacing a power and cooling unit (PCU)
- Replacing a controller

# Maintaining Storage Arrays

This section contains the procedures about how to maintain a storage system in a running cluster. [Table 3–1](#) lists these procedures.

**TABLE 3–1** Task Map: Maintaining a Storage System

Task	Information
Remove a storage array	<a href="#">“How to Remove a Storage Array” on page 27</a>
Upgrade storage array firmware	<a href="#">“How to Upgrade Storage Array Firmware” on page 26</a>
Replace a node-to-switch component	<a href="#">“How to Replace a Node-to-Switch Component in a Cluster Without Multipathing” in <i>Sun Cluster 3.1 - 3.2 With Sun StorEdge or StorageTek 9900 Series Storage Device Manual for Solaris OS</i></a>
Replace a node's host adapter	<a href="#">“How to Replace a Host Adapter” in <i>Sun Cluster 3.1 - 3.2 With Sun StorEdge A3500FC System Manual for Solaris OS</i></a>
Replace a disk drive	<a href="#">“How to Replace a Failed Disk Drive in a Running Cluster” in <i>Sun Cluster 3.1 - 3.2 With Sun StorEdge A3500FC System Manual for Solaris OS</i></a>
Add a node to the storage array	Sun Cluster system administration documentation
Remove a node from the storage array	Sun Cluster system administration documentation

---

**Note** – Most storage arrays are maintained through common array management software. For example, see [Sun StorageTek Common Array Manager Software](#).

---

## ▼ How to Upgrade Storage Array Firmware

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

---

**Note** – When you upgrade firmware on a storage device or on an enclosure, redefine the stripe size of a LUN, or perform other LUN operations, a device ID might change unexpectedly. When you perform a check of the device ID configuration by running the `cldevice check` or `scdidadm -c` command, the following error message appears on your console if the device ID changed unexpectedly.

```
device id for nodename:/dev/rdsckXtYdZsN does not match physical  
device's id for ddecimalnumber, device may have been replaced.
```

To fix device IDs that report this error, run the `cldevice repair` or `scdidadm -R` command for each affected device.

---

- 1 **Stop all I/O to the storage arrays you are upgrading.**
- 2 **Apply the controller, disk drive, and loop-card firmware patches by using the arrays' GUI tools.**  
For specific instructions, see your storage array's documentation.
- 3 **Confirm that all storage arrays that you upgraded are visible to all nodes.**  
`# luxadm probe`
- 4 **Restart all I/O to the storage arrays.**  
You stopped I/O to these storage arrays in Step 1.

## ▼ How to Remove a Storage Array

Use this procedure to permanently remove a storage array from a running cluster.



---

**Caution** – During this procedure, you lose access to the data that resides on the storage array that you are removing. Back up the data before you proceed.

---

- 1 **If you are using Sun Cluster Geographic Edition, you might need to back up all database tables, data services, and volumes that are associated with each partner group that is affected.**
- 2 **Remove references to the volumes that reside on the storage array that you are removing. For instructions, see “[How to Remove a Volume From a Device Group \(VERITAS Volume Manager\)](#)” in *Sun Cluster System Administration Guide for Solaris OS*.**  
For more information, see your Solaris Volume Manager or Veritas Volume Manager documentation.
- 3 **If you have a Fibre Channel array, disconnect the cables that connected Node N to the FC switches in your storage array.**

- 4 On all nodes, remove the obsolete Solaris links and device IDs.
  - If you are using Sun Cluster 3.2, use the following command:

```
# devfsadm -C  
# cldevice clear
```
  - If you are using Sun Cluster 3.1, use the following command:

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

- 5 Repeat [Step 3](#) and [Step 4](#) for each node that is connected to the storage array.

# Index

---

## A

adding  
  *See also* installing  
disk drives, 25  
initiator groups, 19  
storage arrays, 17-18  
arrays, *See* storage arrays

## C

cables, replacing, 25  
chassis, replacing, 25  
controllers, replacing, 25  
cooling units, replacing, 25  
creating  
  initiator groups, 19  
  logical volumes, 19-21  
  storage pools, 19  
  volume groups, 19

## D

deleting logical volumes, 21-24  
disk drives, adding, 25

## E

Ethernet cables, replacing, 25

## F

firmware, upgrade storage array firmware, 26-27  
FRUs, 25

## H

help, 7-8

## I

initiator groups, 19  
installing  
  *See also* adding  
storage arrays, 16-17

## L

logical devices, *See* logical volumes  
logical unit numbers, *See* logical volumes  
logical volumes  
  creating, 19-21  
  removing, 21-24  
LUN masking, 11  
LUNs, *See* logical volumes

## M

midplane, replacing, 25  
modifying, initiator groups, 19

## N

nodes, adding and removing, 26

## T

technical support, 7-8

## P

pools, 19

power cables, replacing, 25

power supplies, replacing, 25

power units, replacing, 25

## U

upgrading, storage array firmware, 26-27

## V

volume groups, 19

## R

removing

initiator groups, 19

logical volumes, 21-24

nodes, 26

storage arrays, 27-28

storage pools, 19

volume groups, 19

replacing

chassis, 25

controllers, 25

cooling units, 25

Ethernet cables, 25

midplane, 25

power cables, 25

power supplies, 25

power units, 25

## S

storage array firmware, upgrading, 26-27

storage arrays

adding, 17-18

installing, 16-17

removing, 27-28

storage-based replication, 11

storage pools, 19

systems, *See* storage arrays