Sun Cluster 3.0-3.1 With Sun StorEdge T3 or T3+ Array Manual for Solaris OS

SPARC Platform Edition

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

Part No: 817–0176–13
January 2006, Revision A

060118@13215

Contents

Preface

The *Sun Cluster 3.0-3.1 With Sun StorEdge T3 or T3+ Array Manual for Solaris OS* provides procedures specific to Sun StorEdge™ T3 or T3+ arrays in single-controller and partner-group configurations when placed in a Sun™ Cluster environment.

Use this manual with any version of Sun Cluster 3.0 and 3.1 software. Unless otherwise noted, procedures are the same for all Sun Cluster 3.0 and 3.1 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 2 chapters.

Chapter 1 discusses how to install and configure the StorEdge T3 and T3+ storage array.

Chapter 2 discusses how to maintain the StorEdge T3 and T3+ storage array.

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 3.0-3.1 With Sun StorEdge T3 or T3+ Array Manual for Solaris OS

Revision Date	New Information
January 2006	Added information about single, dual-port HBA configurations.

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, 2004, and 2005 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.

TABLE P–2 Hardware Documentation

Title	Part Number
Sun StorEdge Traffic Manager Installation and Configuration Guide	816-1420
Available on `http://www.sun.com/products-n-solutions/hardware/docs`	
Sun StorEdge T3 and T3+ Array Configuration Guide	816-4771
Available on `http://www.sun.com/products-n-solutions/hardware/docs`	

TABLE P–2 Hardware Documentation *(Continued)*

Title	Part Number
Sun StorEdge T3 and T3+ Array Installation, Operation, and Service Manual	816-0773
Available on `http://www.sun.com/products-n-solutions/hardware/docs`	
Sun StorEdge T3 and T3+ Array Release Notes	816-3652
Available on `http://www.sun.com/products-n-solutions/hardware/docs`	
Sun StorEdge T3 Array Controller Upgrade Manual	816-0780
Available on `http://www.sun.com/products-n-solutions/hardware/docs`	

TABLE P–3 Sun Cluster Documentation

Application	Title
Concepts	*Sun Cluster Concepts Guide for Solaris OS*
Overview	*Sun Cluster Overview for Solaris OS*
Hardware administration	*Sun Cluster 3.0-3.1 Hardware Administration Manual for Solaris OS*
	Individual hardware administration guides
Software installation	*Sun Cluster Software Installation Guide for Solaris OS*
Data service administration	*Sun Cluster Data Services Planning and Administration Guide for Solaris OS*
	Individual data service guides
Data service development	*Sun Cluster Data Services Developer's Guide for Solaris OS*
System administration	*Sun Cluster System Administration Guide for Solaris OS*
Error messages	*Sun Cluster Error Messages Guide for Solaris OS*
Command and function reference	*Sun Cluster Reference Manual for Solaris OS*
Release Notes	*Sun Cluster 3.1 8/05 Release Notes for Solaris OS*
	Sun Cluster 3.0-3.1 Release Notes Supplement

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 9)
- The release number of Sun Cluster (for example, Sun Cluster 3.1 8/05)

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

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

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

Typographic Conventions

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

TABLE P–4 Typographic Conventions

Typeface	Meaning	Example
`AaBbCc123`	The names of commands, files, and directories, and onscreen computer output	Edit your `.login` file. Use `ls -a` to list all files. `machine_name% you have mail.`
`AaBbCc123`	What you type, contrasted with onscreen computer output	`machine_name%` **`su`** `Password:`
aabbcc123	Placeholder: replace with a real name or value	The command to remove a file is `rm` *filename*.
AaBbCc123	Book titles, new terms, and terms to be emphasized	Read Chapter 6 in the *User's Guide*. A *cache* is a copy that is stored locally. Do *not* save the file. **Note:** Some emphasized items appear bold online.

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–5 Shell Prompts

Shell	Prompt
C shell	`machine_name%`
C shell for superuser	`machine_name#`
Bourne shell and Korn shell	`$`

TABLE P–5 Shell Prompts *(Continued)*

Shell	Prompt
Bourne shell and Korn shell for superuser	#

Product Training

Sun Microsystems offers training in many Sun technologies through a variety of instructor-led courses and self-paced courses. For information about the training courses that Sun offers and to enroll in a class, visit Sun Training at `http://training.sun.com/`.

CHAPTER 1

Installing and Configuring a Sun StorEdge T3 or T3+ Array

This chapter contains the procedures about how to install and configure Sun™ StorEdge™ T3 and Sun StorEdge T3+ arrays in a single-controller (noninterconnected) or partner-group (interconnected) configuration in a Sun™ Cluster environment.

This chapter contains the following topics.

For information about how to use a storage array in a storage area network (SAN), see "SAN Solutions in a Sun Cluster Environment" in *Sun Cluster 3.0-3.1 Hardware Administration Manual for Solaris OS*.

Installing Sun StorEdge T3 and T3+ Arrays

This section contains the procedures about how to install storage arrays in a cluster. The following table lists these procedures.

TABLE 1–1 Task Map: Installing a Storage Array

Task	Information
Installing arrays in a new cluster, using a single-controller configuration	"How to Install a Storage Array in a New Cluster Using a Single-Controller Configuration" on page 12

TABLE 1–1 Task Map: Installing a Storage Array *(Continued)*

Task	Information
Installing arrays in a new cluster, using a partner-group configuration	"How to Install a Storage Array in a New Cluster, Using a Partner-Group Configuration" on page 16
Adding arrays to an existing cluster, using a single-controller configuration.	"How to Add a Storage Array to an Existing Cluster, Using a Single-Controller Configuration" on page 20
Adding arrays to an existing cluster, using a partner-group configuration.	"How to Add a Storage Array to an Existing Cluster, Using a Partner-Group Configuration" on page 29
Upgrading a T3 storage array to a T3+ array.	"How to Upgrade a StorEdge T3 Controller to a StorEdge T3+ Controller" on page 65
Migrating a single-controller array configuration to a partner-group configuration.	"How to Migrate From a Single-Controller Configuration to a Partner-Group Configuration" on page 66

▼ How to Install a Storage Array in a New Cluster Using a Single-Controller Configuration

Use this procedure to install and configure the first storage array in a new cluster, using a single-controller configuration. Perform the steps in this procedure in conjunction with the procedures in the Sun Cluster software installation documentation and your server hardware manual.

The following procedures contain instructions for other array-installation situations:

- "How to Install a Storage Array in a New Cluster, Using a Partner-Group Configuration" on page 16
- "How to Add a Storage Array to an Existing Cluster, Using a Single-Controller Configuration" on page 20
- "How to Add a Storage Array to an Existing Cluster, Using a Partner-Group Configuration" on page 29

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 `scdidadm -c` command, the following error message appears on your console if the device ID changed unexpectedly.

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

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

Steps

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

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

2. **Install the FC hubs/switches.**

   For the procedure about how to install FC hubs/switches, see the documentation that shipped with your FC hub/switch hardware.

   Note – If you are using two FC switches and Sun SAN software to create a storage area network (SAN), see "SAN Solutions in a Sun Cluster Environment" in *Sun Cluster 3.0-3.1 Hardware Administration Manual for Solaris OS* for more information.

3. **Set up a Reverse Address Resolution Protocol (RARP) server on the network on which the new storage arrays are to reside.**

   This RARP server enables you to assign an IP address to the new storage array by using each storage array's unique MAC address.

   For the procedure about how to set up a RARP server, see the *Sun StorEdge T3 and T3+ Array Installation, Operation, and Service Manual*.

4. **Are you adding a StorEdge T3+ array?**

   - **If yes, proceed to Step 5.**
   - **If no, install the media interface adapters (MIAs) in the storage array that you are installing, as shown in Figure 1–1.**

     For the procedure about how to install a media interface adapter (MIA), see the *Sun StorEdge T3 and T3+ Array Configuration Guide*.

5. **If necessary, install gigabit interface converters (GBICs) or Small Form-Factor Pluggables (SFPs) in the FC hubs/switches, as shown in Figure 1–1.**

   The GBICs or SFPs let you connect the FC hubs/switches to the storage array that you are installing. For the procedure about how to install an FC hub/switch GBIC or an SFP, see the documentation that shipped with your FC hub/switch hardware.

6. **Install fiber-optic cables between the FC hubs/switches and the storage array, as shown in Figure 1–1.**

   For the procedure about how to install a fiber-optic cable, see the *Sun StorEdge T3 and T3+ Array Configuration Guide*.

7. **Install fiber-optic cables between the FC hubs/switches and the nodes, as shown in Figure 1–1.**

8. **Install the Ethernet cables between the storage array and the Local Area Network (LAN), as shown in Figure 1–1.**

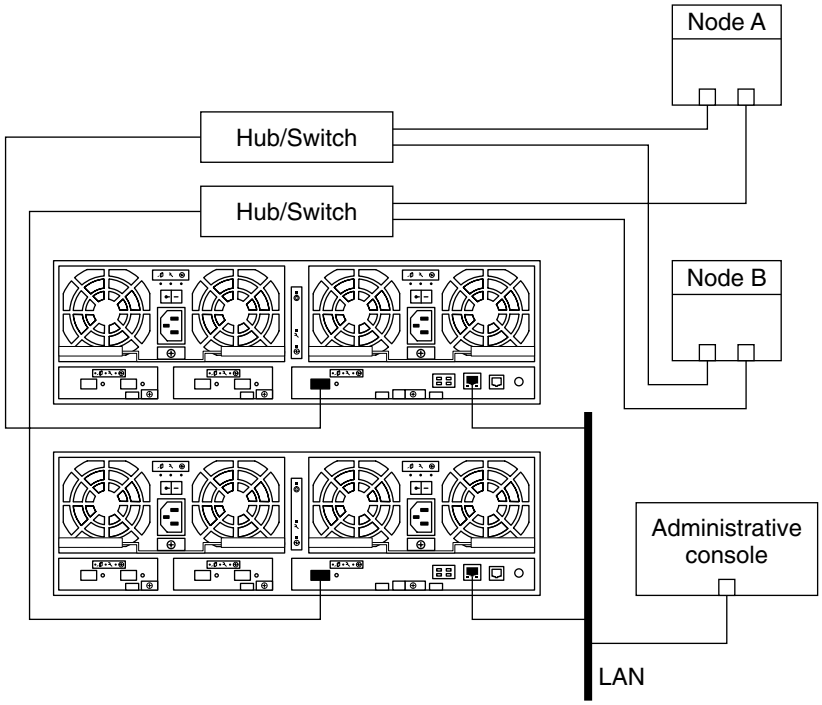

FIGURE 1–1 Installing a Single-Controller Configuration

Note – Figure 1–1 shows how to cable two storage arrays to enable data sharing and host-based mirroring. This configuration prevents a single-point of failure.

---

9. **Install power cords to each storage array that you are installing.**

10. **Power on the storage array and confirm that all components are powered on and functional.**

Note – The storage array might require a few minutes to boot.

For the procedure about how to power on a storage array, see the *Sun StorEdge T3 and T3+ Array Installation, Operation, and Service Manual*.

11. **(Optional) Configure the storage array with logical volumes.**

For the procedure about how to configure the storage array with logical volumes, see the *Sun StorEdge T3 and T3+ Array Administrator's Guide*.

12. **Access each storage array that you are adding. Install the required controller firmware for the storage array.**

PatchPro is a patch-management tool that eases the selection and download of patches required for installation or maintenance of Sun Cluster software. PatchPro provides an Interactive Mode tool especially for Sun Cluster. The Interactive Tool makes the installation of patches easier. PatchPro's Expert Mode tool helps you to maintain your configuration with the latest set of patches. Expert Mode is especially useful for obtaining all of the latest patches, not just the high availability and security patches.

To access the PatchPro tool for Sun Cluster software, go to http://www.sun.com/PatchPro/, click Sun Cluster, then choose either Interactive Mode or Expert Mode. Follow the instructions in the PatchPro tool to describe your cluster configuration and download the patches.

For third-party firmware patches, see the SunSolve^{SM} Online site at `http://sunsolve.ebay.sun.com`.

13. **Ensure that this new storage array has a unique target address.**

For the procedure on how to verify and assign a target address, see the *Sun StorEdge T3 and T3+ Array Configuration Guide*.

14. **Reset the storage array.**

For the procedure about how to reboot or reset a storage array, see the *Sun StorEdge T3 and T3+ Array Installation, Operation, and Service Manual*.

15. **Install to the nodes the Solaris operating environment. Apply any required Solaris patches for Sun Cluster software and storage array support.**

For the procedure about how to install the Solaris operating environment, see your Sun Cluster software installation documentation. For the location of required Solaris patches and installation instructions for Sun Cluster software support, see your Sun Cluster release notes documentation. For a list of required Solaris patches

for storage array support, see the *Sun StorEdge T3 and T3+ Array Release Notes*.

See Also To continue with Sun Cluster software installation tasks, see your Sun Cluster software installation documentation.

▼ How to Install a Storage Array in a New Cluster, Using a Partner-Group Configuration

Use this procedure to install and configure the first storage array partner groups in a new cluster. Perform the steps in this procedure in conjunction with the procedures in the Sun Cluster software installation documentation and your server hardware manual.

Make certain that you are using the correct procedure. This procedure contains instructions about how to install a partner group into a new cluster, before the cluster is operational. The following procedures contain instructions for other array-installation situations:

- "How to Add a Storage Array to an Existing Cluster, Using a Partner-Group Configuration" on page 29
- "How to Install a Storage Array in a New Cluster Using a Single-Controller Configuration" on page 12
- "How to Add a Storage Array to an Existing Cluster, Using a Single-Controller Configuration" on page 20

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 `scdidadm -c` command, the following error message appears on your console if the device ID changed unexpectedly.

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

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

Steps

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

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

2. **Install the Fibre Channel (FC) switches.**

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

Note – If you are using two FC switches and Sun SAN software to create a storage area network (SAN), see "SAN Considerations" on page 69 for more information.

3. **Are you installing Sun StorEdge T3+ arrays?**

   - **If no, proceed to Step 4**
   - **If yes, install the media interface adapters (MIAs) in the Sun StorEdge T3 arrays that you are installing, as shown in Figure 1–2.**

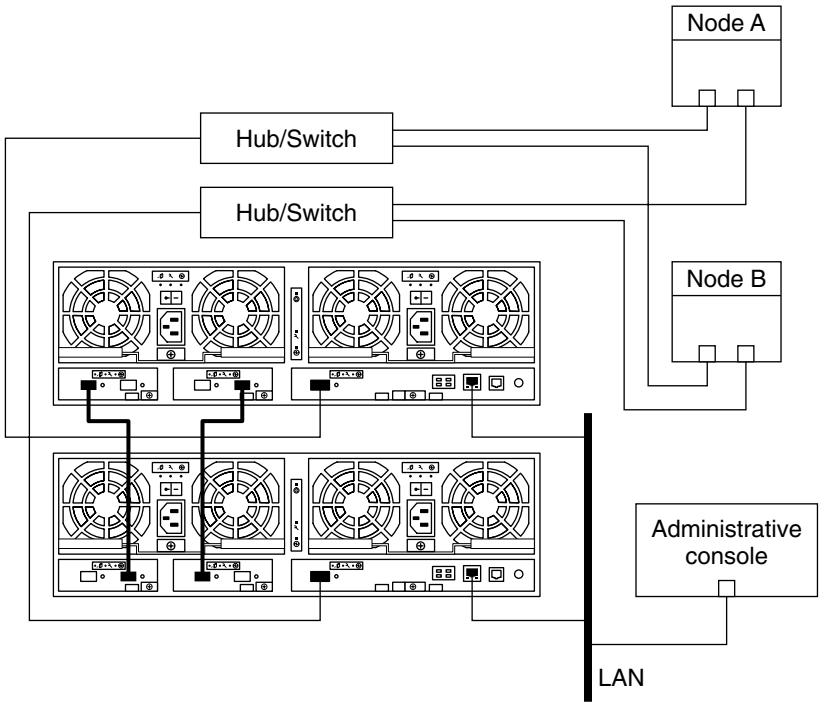

FIGURE 1–2 Installing a Partner-Group Configuration

For the procedure about how to install a media interface adapter (MIA), see the *Sun StorEdge T3 and T3+ Array Configuration Guide*.

4. **If necessary, install GBICs or SFPs in the FC switches, as shown in Figure 1–2.**

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

5. **Set up a Reverse Address Resolution Protocol (RARP) server on the network on which the new storage arrays are to reside.**

   This RARP server enables you to assign an IP address to the new storage arrays. Assign an IP address by using the storage array's unique MAC address. For the procedure about how to set up a RARP server, see the *Sun StorEdge T3 and T3+ Array Installation, Operation, and Service Manual*.

6. **Cable the storage arrays, as shown in Figure 1–2.**

   a. **Connect the storage arrays to the FC switches by using fiber-optic cables.**

   b. **Connect the Ethernet cables from each storage array to the LAN.**

   c. **Connect interconnect cables between the two storage arrays of each partner group.**

   d. **Connect power cords to each storage array.**

   For the procedure about how to install fiber-optic, Ethernet, and interconnect cables, see the *Sun StorEdge T3 and T3+ Array Installation, Operation, and Service Manual*.

7. **Power on the storage arrays. Verify that all components are powered on and functional.**

   For the procedure about how to power on the storage arrays and verify the hardware configuration, see the *Sun StorEdge T3 and T3+ Array Installation, Operation, and Service Manual*.

8. **Administer the storage arrays' network settings.**

   Use the `telnet` command to access the *master* controller unit and administer the storage arrays. For the procedure on how to administer the storage array network addresses and settings, see the *Sun StorEdge T3 and T3+ Array Installation, Operation, and Service Manual*.

   The master controller unit is the storage array that has the interconnect cables attached to the second port of each interconnect card, when viewed from the rear of the storage arrays. For example, Figure 1–2 shows the master controller unit of the partner group as the lower storage array. In this diagram, the interconnect cables are connected to the second port of each interconnect card on the master controller unit.

9. **Install any required storage array controller firmware.**

   For partner-group configurations, use the `telnet` command to access the *master* controller unit. Install the required controller firmware.

PatchPro is a patch-management tool that eases the selection and download of patches required for installation or maintenance of Sun Cluster software. PatchPro provides an Interactive Mode tool especially for Sun Cluster. The Interactive Tool makes the installation of patches easier. PatchPro's Expert Mode tool helps you to maintain your configuration with the latest set of patches. Expert Mode is especially useful for obtaining all of the latest patches, not just the high availability and security patches.

To access the PatchPro tool for Sun Cluster software, go to http://www.sun.com/PatchPro/, click Sun Cluster, then choose either Interactive Mode or Expert Mode. Follow the instructions in the PatchPro tool to describe your cluster configuration and download the patches.

For third-party firmware patches, see the SunSolve[SM] Online site at `http://sunsolve.ebay.sun.com`.

10. **Ensure that each storage array has a unique target address.**

    For the procedure about how to verify and assign a target address to a storage array, see the *Sun StorEdge T3 and T3+ Array Configuration Guide*.

11. **Ensure that the `cache` and `mirror` settings for each storage array are set to `auto`.**

    For more information, see the *Sun StorEdge T3 and T3+ Array Administrator's Guide*.

12. **Ensure that the `mp_support` parameter for each storage array is set to `mpxio`.**

    For more information, see the *Sun StorEdge T3 and T3+ Array Administrator's Guide*.

13. **Ensure that both storage array controllers are online.**

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

14. **(Optional) Configure the storage arrays with the desired logical volumes.**

    For the procedure about how to create and initialize a logical volume, see the *Sun StorEdge T3 and T3+ Array Administrator's Guide*. For the procedure about how to mount a logical volume, see the *Sun StorEdge T3 and T3+ Array Installation, Operation, and Service Manual*.

15. **Reset the storage arrays.**

    For the procedure about how to reboot or reset a storage array, see the *Sun StorEdge T3 and T3+ Array Installation, Operation, and Service Manual*.

16. **On all nodes, install the Solaris operating environment. 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*

17. **Install any required patches or software for Sun StorEdge Traffic Manager software support to nodes.**

    For the procedure about how to install the Sun StorEdge Traffic Manager software, see "How to Install Sun Multipathing Software" in *Sun Cluster Software Installation Guide for Solaris OS*.

18. **On all nodes, update the `/devices` and `/dev` entries.**

    ```
    # devfsadm -C
    ```

19. **On all nodes, confirm that all storage arrays that you installed are visible.**

    ```
    # luxadm display
    ```

See Also To continue with Sun Cluster software installation tasks, see your Sun Cluster software installation documentation.

▼ How to Add a Storage Array to an Existing Cluster, Using a Single-Controller Configuration

This procedure contains instructions about how to add a new storage array to a running cluster in a single-controller configuration. The following procedures contain instructions for other array-installation situations:

- "How to Add a Storage Array to an Existing Cluster, Using a Partner-Group Configuration" on page 29
- "How to Install a Storage Array in a New Cluster Using a Single-Controller Configuration" on page 12
- "How to Install a Storage Array in a New Cluster, Using a Partner-Group Configuration" on page 16

This procedure defines `Node A` as the node with which you begin working. `Node B` is another node in the cluster.

Steps

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

   This RARP server enables you to assign an IP address to the new storage array by using the storage array's unique MAC address.

   For the procedure about how to set up a RARP server, see the *Sun StorEdge T3 and T3+ Array Installation, Operation, and Service Manual*.

2. **Are you adding a StorEdge T3+ array?**

   - **If yes, proceed to Step 3.**
   - **If no, skip to Step 4.**

3. **Install the media interface adapter (MIA) in storage array that you are adding, as shown in Figure 1–3.**

   For the procedure about how to install a media interface adapter (MIA), see the *Sun StorEdge T3 and T3+ Array Configuration Guide*.

4. **If necessary, install gigabit interface converters (GBICs) or Small Form-Factor Pluggables (SFPs) in the FC hub/switch, as shown in Figure 1–3.**

   The GBICs or SFPs enables you to connect the FC hubs/switches to the storage array that you are adding.

   For the procedure about how to install an FC hub/switch GBIC or an SFP, see the documentation that shipped with your FC hub/switch hardware.

   Note – If you are using two FC switches and Sun SAN software to create a storage area network (SAN), see "SAN Solutions in a Sun Cluster Environment" in *Sun Cluster 3.0-3.1 Hardware Administration Manual for Solaris OS* for more information.

   ---

5. **Install the Ethernet cable between the storage array and the Local Area Network (LAN), as shown in Figure 1–3.**

6. **Power on the storage array array.**

   Note – The storage array might require a few minutes to boot.

   For the procedure about how to power on a storage array, see the *Sun StorEdge T3 and T3+ Array Installation, Operation, and Service Manual*.

7. **Access the storage array that you are adding. If necessary, install the required controller firmware for the storage array.**

   PatchPro is a patch-management tool that eases the selection and download of patches required for installation or maintenance of Sun Cluster software. PatchPro provides an Interactive Mode tool especially for Sun Cluster. The Interactive Tool makes the installation of patches easier. PatchPro's Expert Mode tool helps you to maintain your configuration with the latest set of patches. Expert Mode is especially useful for obtaining all of the latest patches, not just the high availability and security patches.

   To access the PatchPro tool for Sun Cluster software, go to http://www.sun.com/PatchPro/, click Sun Cluster, then choose either Interactive Mode or Expert Mode. Follow the instructions in the PatchPro tool to describe your cluster configuration and download the patches.

   For third-party firmware patches, see the SunSolve^{SM} Online site at `http://sunsolve.ebay.sun.com`.

8. **Does this new storage array have a unique target address?**

   - **If yes, proceed to Step 9.**
   - **If no, change the target address for this new storage array.**

   For the procedure about how to verify and assign a target address, see the *Sun StorEdge T3 and T3+ Array Configuration Guide*.

9. **Install a fiber-optic cable between the FC hub/switch and the storage array, as shown in Figure 1–3.**

   For the procedure about how to install a fiber-optic cable, see the *Sun StorEdge T3 and T3+ Array Configuration Guide*.

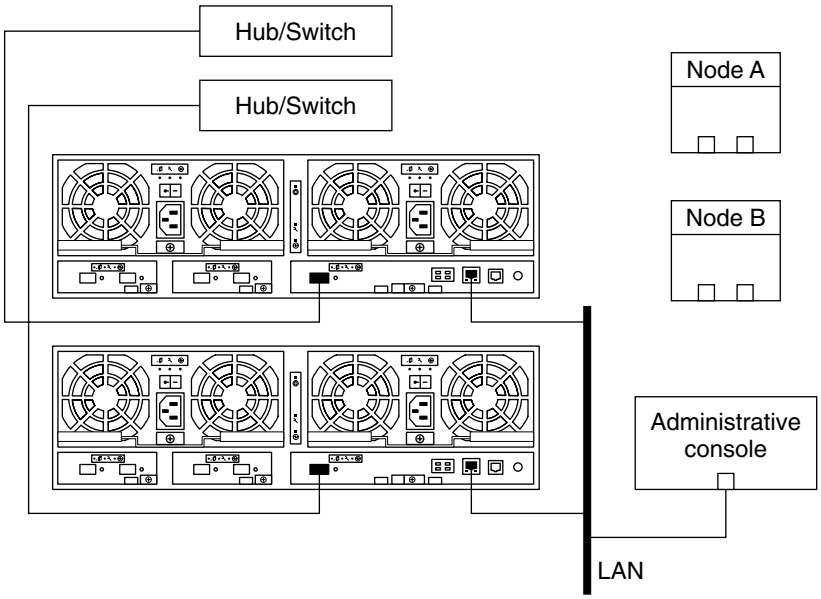

FIGURE 1–3 Adding a Single-Controller Configuration: Part I

Note – Figure 1–3 shows how to cable two storage arrays to enable data sharing and host-based mirroring. This configuration prevents a single-point of failure.

---

10. **Configure the new storage array.**

    For the procedure about how to create a logical volume, see the *Sun StorEdge T3 and T3+ Array Administrator's Guide*.

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

    Record this information because you use this information in Step 41 of this procedure to return resource groups and device groups to these nodes.

    ```
    # scstat
    ```

12. **Move all resource groups and device groups off `Node A`.**

    ```
    # scswitch -S -h from-node
    ```

13. **Do you need to install a host adapter in `Node A`?**

    - **If yes, proceed to Step 14.**
    - **If no, skip to Step 36.**

14. **Is the host adapter that you are installing the first FC host adapter on `Node A`?**

    - **If no, skip to Step 16.**
    - **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 socal Device Driver
    system    SUNWluxlx    Sun Enterprise Network Array socal Device Driver
                                    (64-bit)
    system    SUNWluxop    Sun Enterprise Network Array firmware and utilities
    ```

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

    - **If yes, proceed to Step 16.**
    - **If no, install the required support packages that are missing.**

    The storage array packages are located in the `Product` directory of the Solaris CD-ROM. Add any necessary packages.

16. **Shut down `Node A`.**

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

17. **Power off `Node A`.**

18. **Install the host adapter in `Node A`.**

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

19. **If necessary, power on and boot `Node A` into noncluster mode.**

    For the procedure about how to boot a node in noncluster mode, see your Sun Cluster system administration documentation.

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

    PatchPro is a patch-management tool that eases the selection and download of patches required for installation or maintenance of Sun Cluster software. PatchPro provides an Interactive Mode tool especially for Sun Cluster. The Interactive Tool makes the installation of patches easier. PatchPro's Expert Mode tool helps you to maintain your configuration with the latest set of patches. Expert Mode is especially useful for obtaining all of the latest patches, not just the high availability and security patches.

    To access the PatchPro tool for Sun Cluster software, go to http://www.sun.com/PatchPro/, click Sun Cluster, then choose either Interactive Mode or Expert Mode. Follow the instructions in the PatchPro tool to describe your cluster configuration and download the patches.

    For third-party firmware patches, see the SunSolve^{SM} Online site at `http://sunsolve.ebay.sun.com`.

21. **Connect a fiber-optic cable between the FC hub/switch and `Node A`, as shown in Figure 1–4.**

    For the procedure about how to install an FC host adapter GBIC or an SFP, see your host adapter documentation. For the procedure about how to install a fiber-optic cable, see the *Sun StorEdge T3 and T3+ Array Configuration Guide*.

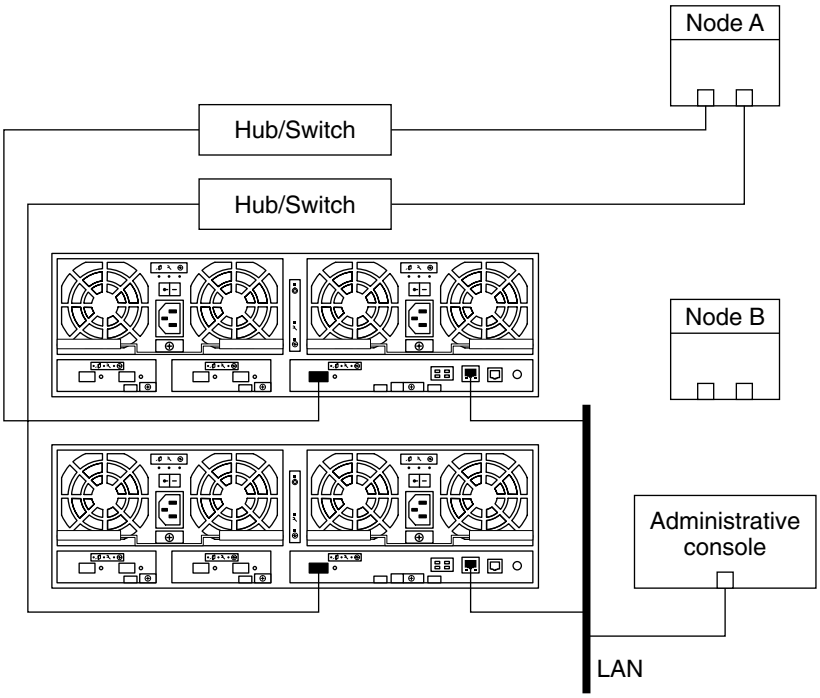

FIGURE 1–4 Adding a Single-Controller Configuration: Part II

22. **If necessary, install the required Solaris patches for array support on `Node A`.**

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

23. **Shut down `Node A`.**

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

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

    ```
    # boot -r
    ```

25. **Label the new logical volume.**

    For the procedure about how to label a logical volume, see the *Sun StorEdge T3 and T3+ Array Administrator's Guide*.

26. **(Optional) On `Node A`, verify that the device IDs (DIDs) are assigned to the new LUNs.**

    ```
    # 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 35.**

28. Is the host adapter that you are installing the first FC 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 socal Device Driver
system    SUNWluxlx    Sun Enterprise Network Array socal 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 the required support packages that are missing.**

The storage array packages are located in the `Product` directory of the Solaris CD-ROM. Add any necessary packages.

30. Shut down `Node B`.

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

31. Power off `Node B`.

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

32. Install the host adapter in `Node B`.

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

33. If necessary, power on and boot `Node B` into noncluster mode.

For the procedure about how to boot a node in noncluster mode, see your Sun Cluster system administration documentation.

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

PatchPro is a patch-management tool that eases the selection and download of patches required for installation or maintenance of Sun Cluster software. PatchPro provides an Interactive Mode tool especially for Sun Cluster. The Interactive Tool makes the installation of patches easier. PatchPro's Expert Mode tool helps you to maintain your configuration with the latest set of patches. Expert Mode is especially useful for obtaining all of the latest patches, not just the high availability and security patches.

To access the PatchPro tool for Sun Cluster software, go to http://www.sun.com/PatchPro/, click Sun Cluster, then choose either Interactive Mode or Expert Mode. Follow the instructions in the PatchPro tool to describe your cluster configuration and download the patches.

For third-party firmware patches, see the SunSolve^{SM} Online site at `http://sunsolve.ebay.sun.com`.

35. **If necessary, install a GBIC or an SFP, as shown in Figure 1–5.**

    For the procedure about how to install an FC hub/switch GBIC or an SFP, see the documentation that shipped with your FC hub/switch hardware.

Note – If you are using two FC switches and Sun SAN software to create a storage area network (SAN), see "SAN Solutions in a Sun Cluster Environment" in *Sun Cluster 3.0-3.1 Hardware Administration Manual for Solaris OS* for more information.

---

36. **If necessary, connect a fiber-optic cable between the FC hub/switch and `Node B`, as shown in Figure 1–5.**

    For the procedure about how to install an FC host adapter GBIC or an SFP, see your host adapter documentation. For the procedure about how to install a fiber-optic cable, see the *Sun StorEdge T3 and T3+ Array Configuration Guide*.

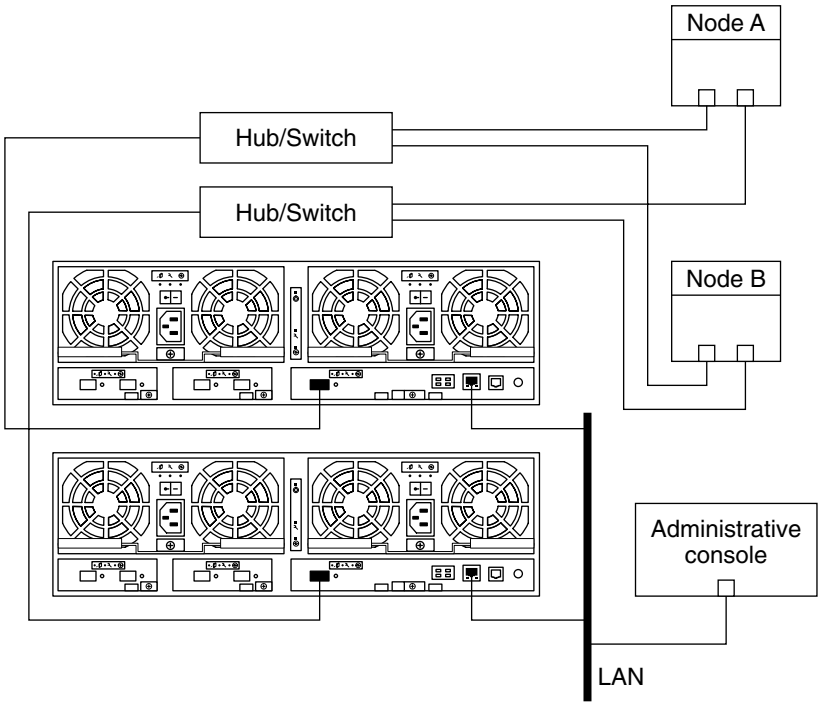

FIGURE 1–5 Adding a Single-Controller Configuration: Part III

37. **If necessary, install the required Solaris patches for storage array support on `Node B`.**

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

38. **Shut down `Node B`.**

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

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

    ```
    # boot -r
    ```

40. **(Optional) On `Node B`, verify that the device IDs (DIDs) are assigned to the new LUNs.**

    ```
    # scdidadm -l
    ```

41. **Return the resource groups and device groups that you identified in Step 11 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 your Sun Cluster system administration documentation.

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

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

▼ How to Add a Storage Array to an Existing Cluster, Using a Partner-Group Configuration

This procedure contains instructions about how to add new storage array partner groups to a running cluster. The following procedures contain instructions for other array-installation situations:

- "How to Install a Storage Array in a New Cluster, Using a Partner-Group Configuration" on page 16
- "How to Install a Storage Array in a New Cluster Using a Single-Controller Configuration" on page 12
- "How to Add a Storage Array to an Existing Cluster, Using a Single-Controller Configuration" on page 20

This procedure defines `Node A` as the node that you begin working with. `Node B` is the second node.

Steps

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

   Note – Assign an IP address to the *master* controller unit only. The master controller unit is the storage array that has the interconnect cables attached to the to the second port of each interconnect card, as shown in Figure 1–6.

   This RARP server enables you to assign an IP address to the new storage arrays. Assign an IP address by using the storage array's unique MAC address. For the procedure about how to set up a RARP server, see the *Sun StorEdge T3 and T3+ Array Installation, Operation, and Service Manual*.

2. **Install the Ethernet cable between the storage arrays and the local area network (LAN), as shown in Figure 1–6.**

3. **If not already installed, install interconnect cables between the two storage arrays of each partner group, as shown in Figure 1–6.**

   For the procedure about how to install interconnect cables, see the *Sun StorEdge T3 and T3+ Array Installation, Operation, and Service Manual*.

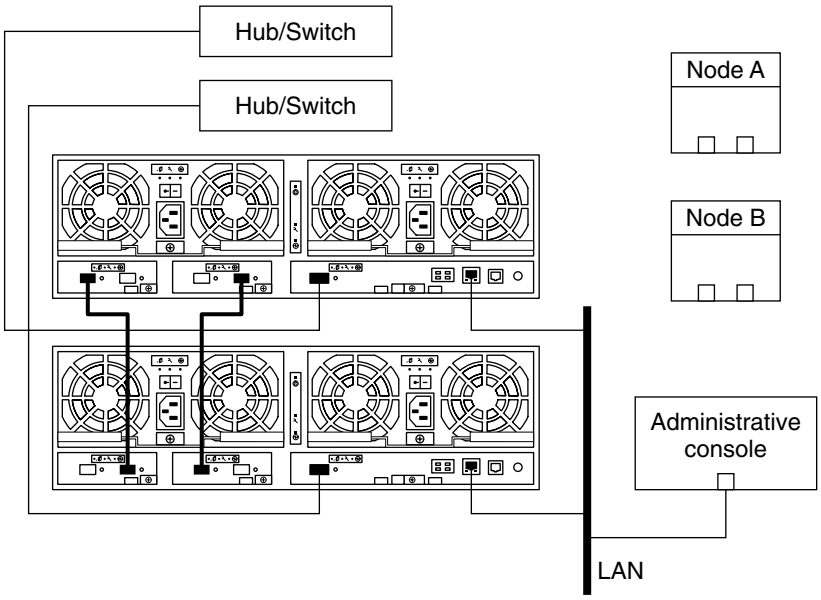

FIGURE 1–6 Adding a Partner-Group Configuration: Part I

4. **Power on the storage arrays.**

Note – The storage arrays might require several minutes to boot.

For the procedure about how to power on storage arrays, see the *Sun StorEdge T3 and T3+ Array Installation, Operation, and Service Manual*.

5. **Administer the storage arrays' network addresses and settings.**

   Use the `telnet` command to access the *master* controller unit and administer the storage arrays.

   For the procedure about how to administer the network address and the settings of a storage array, see the *Sun StorEdge T3 and T3+ Array Installation, Operation, and Service Manual*.

6. **Install any required storage array controller firmware upgrades.**

   For partner-group configurations, use the `telnet` command to the *master* controller unit. If necessary, install the required controller firmware for the storage array.

   For the required revision number of the storage array controller firmware, see the *Sun StorEdge T3 Disk Tray Release Notes*.

7. **Ensure that each storage array has a unique target address.**

   For the procedure about how to verify and assign a target address to a storage array, see the *Sun StorEdge T3 and T3+ Array Configuration Guide*.

8. **Ensure that the `cache` and `mirror` settings for each storage array are set to `auto`.**

   For more information, see the *Sun StorEdge T3 and T3+ Array Administrator's Guide*.

9. **Ensure that the `mp_support` parameter for each storage array is set to `mpxio`.**

   For more information, see the *Sun StorEdge T3 and T3+ Array Administrator's Guide*.

10. **Configure the new storage arrays with the desired logical volumes.**

    For the procedure about how to create and initialize a logical volume, see the *Sun StorEdge T3 and T3+ Array Administrator's Guide*. For the procedure about how to mount a logical volume, see the *Sun StorEdge T3 and T3+ Array Installation, Operation, and Service Manual*.

11. **Reset the storage arrays.**

    For the procedure about how to reboot or reset a storage array, see the *Sun StorEdge T3 and T3+ Array Installation, Operation, and Service Manual*.

12. **Are you adding Sun StorEdge T3+ arrays?**

    - If no, proceed to Step 13.
    - If yes, install the media interface adapter (MIA) in the Sun StorEdge T3 arrays that you are adding, as shown in Figure 1–6.

      For the procedure about how to install an MIA, see the *Sun StorEdge T3 and T3+ Array Configuration Guide*.

13. **If necessary, install GBICs or SFPs in the FC switches, as shown in Figure 1–6.**

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

14. **Install a fiber-optic cable between each FC switch and both new storage arrays of the partner group, as shown in Figure 1–6.**

    For the procedure about how to install a fiber-optic cable, see the *Sun StorEdge T3 and T3+ Array Installation, Operation, and Service Manual*.

Note – If you are using two FC switches and Sun SAN software to create a storage area network (SAN), see "SAN Considerations" on page 69 for more information.

---

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

    Record this information because you use this information in Step 48 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 from-node
    ```

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

    - **If no, skip to Step 39.**
    - **If yes, proceed to Step 18.**

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

    - **If no, skip to Step 20.**
    - **If yes, 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 socal Device Driver
      system     SUNWluxlx     Sun Enterprise Network Array socal 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 yes, proceed to Step 20.**
    - **If no, install the required support packages that are missing.**

    The support packages are located in the `Product` directory of the Solaris CD-ROM. Add any missing packages.

20. **Shut down and power off `Node A`.**

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

21. **Install the host adapters in `Node A`.**

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

22. **Power on and boot `Node A` into noncluster mode.**

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

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

    PatchPro is a patch-management tool that eases the selection and download of patches required for installation or maintenance of Sun Cluster software. PatchPro provides an Interactive Mode tool especially for Sun Cluster. The Interactive Tool

makes the installation of patches easier. PatchPro's Expert Mode tool helps you to maintain your configuration with the latest set of patches. Expert Mode is especially useful for obtaining all of the latest patches, not just the high availability and security patches.

To access the PatchPro tool for Sun Cluster software, go to http://www.sun.com/PatchPro/, click Sun Cluster, then choose either Interactive Mode or Expert Mode. Follow the instructions in the PatchPro tool to describe your cluster configuration and download the patches.

For third-party firmware patches, see the SunSolve^{SM} Online site at `http://sunsolve.ebay.sun.com`.

24. **Connect fiber-optic cables between `Node A` and the FC switches, as shown in Figure 1–7.**

    For the procedure about how to install a fiber-optic cable, see the *Sun StorEdge T3 and T3+ Array Installation, Operation, and Service Manual*.

Note – If you are using two FC switches and Sun SAN software to create a storage area network (SAN), see "SAN Considerations" on page 69 for more information.

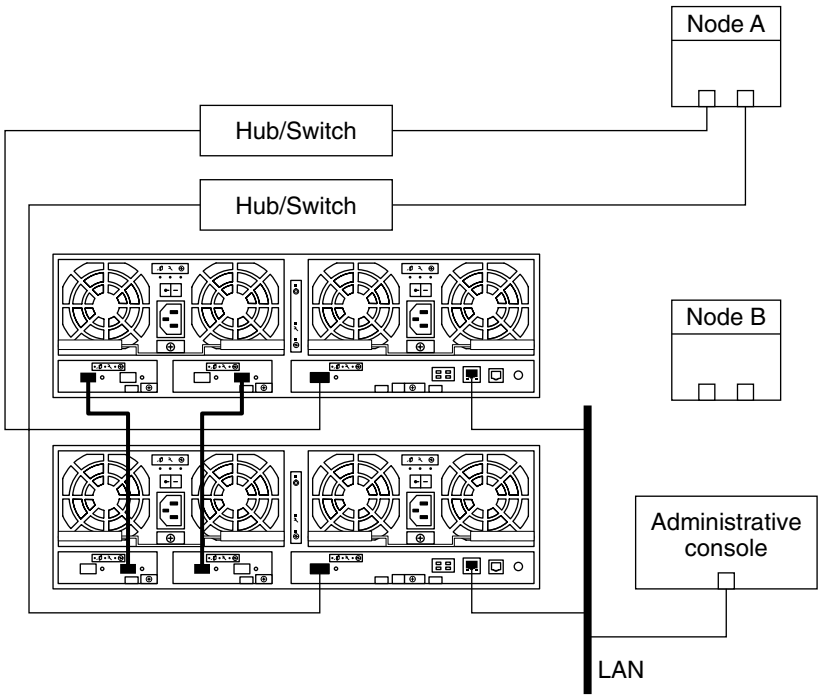

FIGURE 1–7 Adding a Partner-Group Configuration: Part II

25. If necessary, install the required Solaris patches for storage array support on `Node A`.

PatchPro is a patch-management tool that eases the selection and download of patches required for installation or maintenance of Sun Cluster software. PatchPro provides an Interactive Mode tool especially for Sun Cluster. The Interactive Tool makes the installation of patches easier. PatchPro's Expert Mode tool helps you to maintain your configuration with the latest set of patches. Expert Mode is especially useful for obtaining all of the latest patches, not just the high availability and security patches.

To access the PatchPro tool for Sun Cluster software, go to http://www.sun.com/PatchPro/, click Sun Cluster, then choose either Interactive Mode or Expert Mode. Follow the instructions in the PatchPro tool to describe your cluster configuration and download the patches.

For third-party firmware patches, see the SunSolve^{SM} Online site at `http://sunsolve.ebay.sun.com`.

26. **Shut down `Node A`.**

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

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

```
# boot -r
```

28. **On `Node A`, update the `/devices` and `/dev` entries.**

```
# devfsadm -C
```

29. **On `Node A`, update the paths to the DID instances.**

```
# scdidadm -C
```

30. **Label the new storage array logical volume.**

    For the procedure about how to label a logical volume, see the *Sun StorEdge T3 and T3+ Array Administrator's Guide.*

31. **(Optional) On `Node A`, verify that the device IDs (DIDs) are assigned to the new LUNs.**

```
# scdidadm -l
```

32. **Do you need to install host adapters in `Node B`?**

    - **If no, skip to Step 39.**
    - **If yes, proceed to Step 33.**

33. **Is the host adapter that you are installing the first host adapter on `Node B`**

    - **If no, skip to Step 35.**
    - **If yes, determine whether the required support packages for host adapters 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 socal Device Driver
system    SUNWluxlx    Sun Enterprise Network Array socal Device Driver
                                (64-bit)
system    SUNWluxop    Sun Enterprise Network Array firmware and utilities
system    SUNWluxox    Sun Enterprise Network Array libraries (64-bit)
```

34. **Are the required support packages already installed?**

    - **If yes, skip to Step 35.**
    - **If no, install the missing support packages.**

The support packages are located in the `Product` directory of the Solaris CD-ROM. Add any missing packages.

35. **Shut down and power off `Node B`.**

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

36. **Install the host adapters in `Node B`.**

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

37. **Power on and boot `Node B`.**

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

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

    PatchPro is a patch-management tool that eases the selection and download of patches required for installation or maintenance of Sun Cluster software. PatchPro provides an Interactive Mode tool especially for Sun Cluster. The Interactive Tool makes the installation of patches easier. PatchPro's Expert Mode tool helps you to maintain your configuration with the latest set of patches. Expert Mode is especially useful for obtaining all of the latest patches, not just the high availability and security patches.

    To access the PatchPro tool for Sun Cluster software, go to http://www.sun.com/PatchPro/, click Sun Cluster, then choose either Interactive Mode or Expert Mode. Follow the instructions in the PatchPro tool to describe your cluster configuration and download the patches.

    For third-party firmware patches, see the SunSolve^{SM} Online site at `http://sunsolve.ebay.sun.com`.

39. **If necessary, install GBICs or SFPs to the FC switches, as shown in Figure 1–8.**

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

40. **Connect fiber-optic cables between the FC switches and `Node B`, as shown in Figure 1–8.**

    For the procedure about how to install fiber-optic cables, see the *Sun StorEdge T3 and T3+ Array Installation, Operation, and Service Manual*.

Note – If you are using two FC switches and Sun SAN software to create a storage area network (SAN), see "SAN Considerations" on page 69 for more information.

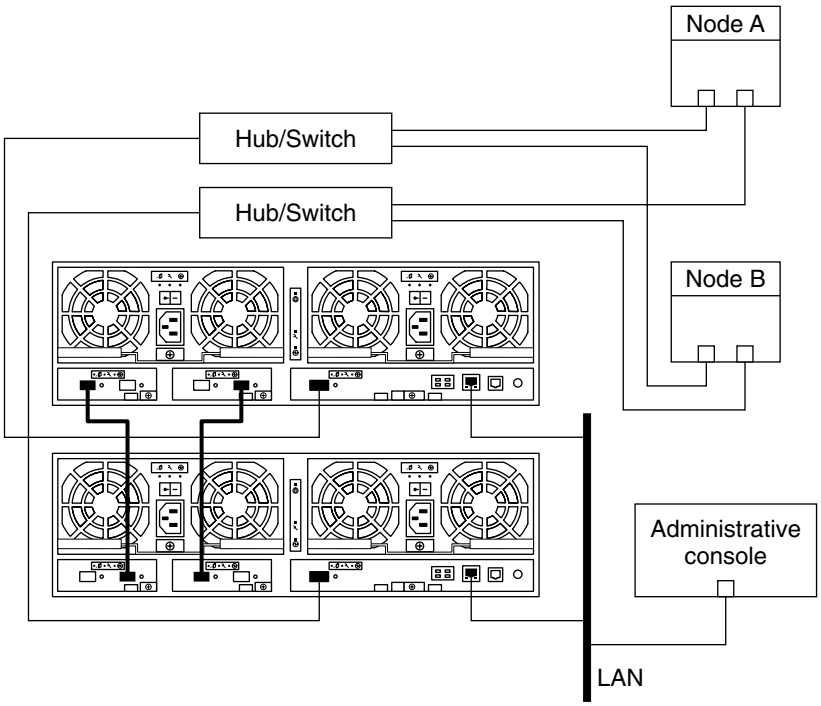

FIGURE 1–8 Adding a Partner-Group Configuration: Part III

41. If necessary, install the required Solaris patches for storage array support on `Node B`.

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

42. Shut down `Node B`.

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

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

```
# boot -r
```

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

```
# devfsadm -C
```

45. On `Node B`, update the paths to the DID instances.

```
# scdidadm -C
```

46. **(Optional) On `Node B`, verify that the DIDs are assigned to the new LUNs.**

    ```
    # scdidadm -l
    ```

47. **On one node that is attached to the new storage arrays, reset the SCSI reservation state.**

    ```
    # scdidadm -R n
    ```

    Where *n* is the DID instance of a storage array LUN that you are adding to the cluster.

    Note – Repeat this command on the same node for each storage array LUN that you are adding to the cluster.

    ---

48. **Return the resource groups and device groups that 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 your Sun Cluster system administration documentation.

49. **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.

Next Steps The best way to enable multipathing for a cluster is to install the multipathing software and enable multipathing before installing the Sun Cluster software and establishing the cluster. For this procedure, see "How to Install Sun Multipathing Software" in *Sun Cluster Software Installation Guide for Solaris OS*. If you need to add multipathing software to an established cluster, see "How to Install Sun Multipathing Software" in *Sun Cluster Software Installation Guide for Solaris OS* and follow the troubleshooting steps to clean up the device IDs.

Configuring Sun StorEdge T3 and T3+ Arrays

This section contains the procedures about how to configure a storage array in a running cluster. The following table lists these procedures.

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 `scdidadm -c` command, the following error message appears on your console if the device ID changed unexpectedly.

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

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

TABLE 1–2 Task Map: Configuring a Storage Array

Task	Information
Create a logical volume.	"How to Create a Logical Volume" on page 39
Remove a logical volume.	"How to Remove a Logical Volume" on page 41

▼ 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 devices is configured as described in the installation procedure.
- If you are using Sun StorEdge Traffic Manager, it is installed and configured and the path between Traffic Manager and the storage device is functioning.

Steps

1. **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.

2. **Are you using multipathing?**

   - **If yes, proceed to Step 3.**
   - **If no, skip to Step 4.**

3. **Are any devices that are associated with the volume you created at an `unconfigured` state?**

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

   - **If no, proceed to Step 4.**
   - **If yes, configure the Traffic Manager paths on each node that is connected to the storage device.**

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

     For the procedure about how to configure Traffic Manager paths, see the *Sun StorEdge Traffic Manager Installation and Configuration Guide*.

4. **On one node that is connected to the storage device, use the `format` command to label the new logical volume.**

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

   ```
   # scgdevs
   ```

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

   ---

6. **To manage this volume with volume management software, use the appropriate Solstice DiskSuite/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 Solstice DiskSuite/Solaris Volume Manager or VERITAS Volume Manager documentation.

See Also

- To configure a logical volume as a quorum device, see Chapter 5, "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.

Steps

1. **Identify the logical volume that you are removing.**

   Refer to your Solstice DiskSuite/Solaris Volume Manager or VERITAS Volume Manager documentation for more information.

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

3. **Check if the logical volume that you are removing is a quorum device.**

   ```
   # scstat -q
   ```

   If yes, choose and configure another device as the quorum device. Then remove the old quorum device.

   For procedures about how to add and remove quorum devices, see Chapter 5, "Administering Quorum," in *Sun Cluster System Administration Guide for Solaris OS*.

4. **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 Solstice DiskSuite/Solaris Volume Manager or VERITAS Volume Manager documentation.

5. **If you are using volume management software, run the appropriate Solstice DiskSuite/Solaris Volume Manager or VERITAS Volume Manager commands to remove the logical volume from any diskset or disk group.**

   For more information, see your Solstice DiskSuite/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

---

6. **If you are using multipathing, unconfigure the volume in Sun StorEdge Traffic Manager.**

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

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

   For the procedure about how to remove the volume, see your storage documentation. For a list of storage documentation, see "Related Documentation" on page 6.

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

   Record this information because you use it in Step 13 of this procedure to return resource groups and device groups to these nodes.

   ```
   # scstat
   ```

9. **Move all resource groups and device groups off `Node A`.**

   ```
   # scswitch -s -h from-node
   ```

10. **Shut down and reboot `Node A`.**

    For the procedure about how to shut down and power off a node, see Chapter 3, "Shutting Down and Booting a Cluster," in *Sun Cluster System Administration Guide for Solaris OS*.

11. **On `Node A`, remove the paths to the logical volume that you removed. Remove obsolete device IDs.**

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

12. **For each additional node that is connected to the shared storage that hosted the logical volume, repeat Step 8 to Step 11.**

13. **(Optional) Return the resource groups and device groups that you identified in Step 8 to all cluster nodes.**

CHAPTER 2

Maintaining and Upgrading a Sun StorEdge T3 or T3+ Array

This chapter contains the procedures about how to maintain Sun StorEdge™ T3 and Sun StorEdge T3+ arrays in a single-controller (noninterconnected) or partner-group (interconnected) configuration in a Sun™ Cluster environment.

This book contains the following topics.

For information about how to use a storage array in a storage area network (SAN), see "SAN Solutions in a Sun Cluster Environment" in *Sun Cluster 3.0-3.1 Hardware Administration Manual for Solaris OS*.

Maintaining Sun StorEdge T3 or T3+ Storage Array Components

This section contains the procedures about how to maintain a storage array. The following table lists these procedures. This section does not include a procedure about how to add a disk drive and a procedure about how to remove a disk drive, because a storage array 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 storage array is designed so that an orderly shutdown occurs. This 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. If you do not replace the FRU within this time, the storage array, and all attached storage array, shut down and power off.

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 `scdidadm -c` command, the following error message appears on your console if the device ID changed unexpectedly.

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

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

TABLE 2–1 Task Map: Maintaining a Storage Array

Task	Information
Replace a disk drive.	"How to Replace a Disk Drive" on page 48
Replace a host-to-hub/switch fiber-optic cable.	"Replacing a Node-to-Switch Component" on page 49
Replace an FC host adapter GBIC or an SFP.	"Replacing a Node-to-Switch Component" on page 49
Replace an FC hub/switch GBIC or an SFP that connects an FC hub/switch to a host.	"Replacing a Node-to-Switch Component" on page 49
Replace a hub/switch-to-storage array fiber-optic cable.	"How to Replace a Hub/Switch or Hub/Switch-to-Array Component in a Single-Controller Configuration" on page 51
Replace an FC hub/switch GBIC or an SFP that connects the FC hub/switch to a storage array.	"How to Replace a Hub/Switch or Hub/Switch-to-Array Component in a Single-Controller Configuration" on page 51

TABLE 2–1 Task Map: Maintaining a Storage Array *(Continued)*

Task	Information
Replace an FC hub/switch.	"How to Replace a Hub/Switch or Hub/Switch-to-Array Component in a Single-Controller Configuration" on page 51
Replace an FC switch. This procedure applies to SAN-configured clusters only.	"How to Replace a Hub/Switch or Hub/Switch-to-Array Component in a Single-Controller Configuration" on page 51
Replace an FC hub/switch power cord.	"How to Replace a Hub/Switch or Hub/Switch-to-Array Component in a Single-Controller Configuration" on page 51
Replace a media interface adapter (MIA) on a storage array. This procedure does not apply to storage arrays.	"How to Replace a Hub/Switch or Hub/Switch-to-Array Component in a Single-Controller Configuration" on page 51
Replace a storage array controller.	"How to Replace a Storage Array Controller" on page 54
Replace a chassis.	"How to Replace a Chassis" on page 54
Replace a host adapter.	"How to Replace a Host Adapter" on page 55
Remove a storage array.	"How to Remove a Storage Array in a Single-Controller Configuration" on page 57
Remove a storage array.	"How to Remove a Storage Array in a Single-Controller Configuration" on page 57
Remove a partner group.	"How to Remove a Partner Group" on page 59
Add a node to the storage device.	Sun Cluster system administration documentation
Remove a node from the storage device.	Sun Cluster system administration documentation

Sun StorEdge T3 and T3+ Array FRUs

The following is a list of administrative tasks that require no cluster-specific procedures. See the *Sun StorEdge T3 and T3+ Array Installation, Operation, and Service Manual* for the following procedures.

- Replacing an Ethernet cable.
- Replacing a power cable on the storage array.
- Replacing a power and cooling unit (PCU).
- Replacing a unit interconnect card (UIC).
- Adding a disk drive.
- Removing a disk drive.

Upgrading a Sun StorEdge T3 array controller to a Sun StorEdge T3+ array controller requires no cluster-specific procedures. See the *Sun StorEdge T3 Array Controller Upgrade Manual* for the following procedures.

▼ How to Replace a Disk Drive

Use this procedure to replace a failed disk drive in a storage array in a running cluster.

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.

Steps

1. **If the failed disk drive affect the storage array logical volume's availability, If yes, use volume manager commands to detach the submirror or plex.**

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

2. **Check if the logical volume (in Step 1) is a quorum device.**

   ```
   # scstat -q
   ```

   If yes, choose and configure another device to be the new quorum device. Remove the old quorum device. For procedures about how to add and remove quorum devices, see your Sun Cluster system administration documentation.

3. **Replace the failed disk drive.**

   For instructions, refer to the *Sun StorEdge T3 and T3+ Array Installation, Operation, and Service Manual*.

4. **(Optional) If the new disk drive is part of a logical volume that you want to be a quorum device, add the quorum device.**

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

5. **Did you detach a submirror or plex in Step 1?**

   - If no, you are finished with this procedure.
   - If yes, use volume manager commands to reattach the submirror or plex.

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

Replacing a Node-to-Switch Component

Use this procedure to replace a node-to-switch component that has failed or that you suspect might be contributing to a problem.

Note – Node-to-switch components that are covered by this procedure include the following components:

- Node-to-switch fiber-optic cables
- Gigabit interface converters (GBICs) or small form-factor pluggables (SFPs) on an FC switch
- FC switches

For the procedure about how to replace a host adapter, see "How to Replace a Host Adapter" on page 55.

This procedure defines `Node A` as the node that is connected to the node-to-switch component that you are replacing. This procedure assumes that, except for the component you are replacing, your cluster is operational.

Ensure that you are following the appropriate instructions:

- If your cluster uses multipathing, see "How to Replace a Node-to-Switch Component in a Cluster That Uses Multipathing" on page 50.
- If your cluster does *not* use multipathing, see "How to Replace a Node-to-Switch Component in a Cluster Without Multipathing" on page 50.

▼ How to Replace a Node-to-Switch Component in a Cluster That Uses Multipathing

Steps

1. **Is your configuration active-passive?**

   If yes, and the active path is the path that needs a component replaced, make that path passive.

2. **Replace the component.**

   Refer to your hardware documentation for any component-specific instructions.

3. **(Optional) If your configuration is active-passive and you changed your configuration in Step 1, switch your original data path back to active.**

▼ How to Replace a Node-to-Switch Component in a Cluster Without Multipathing

Steps

1. **Check if the physical data path failed.**

   If no, proceed to Step 2.

   If yes:

   a. **Replace the component.**

      Refer to your hardware documentation for any component-specific instructions.

   b. **Fix the volume manager error that was caused by the failed data path.**

   c. **(Optional) If necessary, return resource groups and device groups to this node.**

   You have completed this procedure.

2. **Determine the resource groups and device groups that are running on `Node A`.**

   ```
   # scstat
   ```

3. **Move all resource groups and device groups to another node.**

   ```
   # scswitch -s -h from-node
   ```

4. **Replace the node-to-switch component.**

   Refer to your hardware documentation for any component-specific instructions.

5. **(Optional) If necessary, return the resource groups and device groups that you identified in Step 2 to `Node A`.**

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

▼ How to Replace a Hub/Switch or Hub/Switch-to-Array Component in a Single-Controller Configuration

Use this procedure to replace a hub/switch or the following hub/switch-to-array components for an array in a single-controller configuration. To replace these components for an array in a partner-group configuration, see "How to Replace an FC Switch or Storage Array-to-Switch Component in a Partner-Group Configuration" on page 52. You can use storage array in single-controller configuration with FC switches when you create a SAN.

- Fiber-optic cable that connects a hub/switch to a storage array
- FC hub/switch GBIC or an SFP that connects a hub/switch to a storage array
- FC hub/switch
- FC hub/switch power cord
- Media interface adapter (MIA) on a StorEdge T3 array

  This component does not apply to StorEdge T3+ arrays.

Steps

1. **Detach the submirrors on the storage array. This array is connected to the hub/switch-to-array fiber-optic cable that you are replacing. Detach the submirrors to stop all I/O activity to this storage array.**

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

2. **Replace the hub/switch or hub/switch-to-array component.**

   - For the procedure about how to replace a fiber-optic cable, see the *Sun StorEdge T3 and T3+ Array Configuration Guide*.
   - For the procedure about how to replace an FC hub/switch GBIC or an SFP, an FC hub/switch, or an FC hub/switch power cord, see the documentation that shipped with your FC hub/switch hardware.
   - For the procedure about how to replace an MIA, see the *Sun StorEdge T3 and T3+ Array Installation, Operation, and Service Manual*.
   - If you are replacing FC switches in a SAN, follow the hardware installation and SAN configuration instructions in the documentation that shipped with your switch hardware.

Note – If you are replacing an FC switch and you intend to save the switch configuration for restoration to the replacement switch, do not connect the cables to the replacement switch until *after* you recall the Fabric configuration to the replacement switch. For more information about how to save and recall switch configurations, see the documentation that shipped with your switch hardware.

Note – Before you replace an FC switch, be sure that the `probe_timeout` parameter of your data service software is set to more than 90 seconds. If you increase the value of the `probe_timeout` parameter to more than 90 seconds, you avoid unnecessary resource group restarts. Resource group restarts occur when one of the FC switches is powered off.

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

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

▼ How to Replace an FC Switch or Storage Array-to-Switch Component in a Partner-Group Configuration

Use this procedure to replace components for a storage array in a partner-group configuration in a running cluster. To replace components for arrays in a single-controller configuration, see "How to Replace a Hub/Switch or Hub/Switch-to-Array Component in a Single-Controller Configuration" on page 51. Use this procedure to replace the following storage array-to-switch components.

- Fiber-optic cable that connects an FC switch to a storage array.
- GBIC or an SFP on an FC switch, connecting to a storage array.
- FC switch.
- Media Interface Adapter (MIA) on a Sun StorEdge T3 storage array. This component does not apply to Sun StorEdge T3+ storage arrays.
- Interconnect cables between two interconnected storage arrays of a partner group.

Steps

1. **Access the storage array. This storage array is connected to the FC switch or component that you are replacing.**

2. **View the controller status for the two storage arrays in the partner group.**

   For more information, see the *Sun StorEdge T3 and T3+ Array Administrator's Guide*.

3. **Are you replacing an FC switch or component that is attached to a storage array controller that is `ONLINE` or `DISABLED`?**

   - **If the controller is already `DISABLED`, skip to Step 5.**
   - **If the controller is `ONLINE`, disable the controller.**

   For more information, see the *Sun StorEdge T3 and T3+ Array Administrator's Guide*.

4. **Verify that the controller's state has been changed to `DISABLED`.**

5. **Replace the component by using the following references.**

   - **For the procedure about how to replace a fiber-optic cable between a storage array and an FC switch, see the documentation that shipped with your switch hardware.**
   - **For the procedure about how to replace a GBIC or an SFP on an FC switch, see the documentation that shipped with your FC switch hardware.**
   - **For the procedure about how to replace an FC switch, see the documentation that shipped with your switch hardware.**

     Note – If you are replacing an FC switch and you intend to save the switch IP configuration for restoration to the replacement switch, wait to connect the cables to the replacement switch. Connect the cables to the replacement switch *after* you recall the Fabric configuration to the replacement switch. For more information about how to save and recall switch configurations, see the documentation that shipped with your switch hardware.

     Note – Before you replace an FC switch, be sure that the `probe_timeout` parameter of your data service software is set to more than 90 seconds. If you increase the value of the `probe_timeout` parameter to more than 90 seconds, you avoid unnecessary resource group restarts. Resource group restarts occur when one of the FC switches is powered off.

   - **For the procedure about how to replace an MIA, see the** *Sun StorEdge T3 and T3+ Array Installation, Operation, and Service Manual***.**
   - **For the procedure about how to replace interconnect cables, see the** *Sun StorEdge T3 and T3+ Array Installation, Operation, and Service Manual***.**

6. **If necessary, access the storage array. This storage array is the storage array in the partner group that is still online.**

7. **Enable the storage array that you disabled in Step 3.**

   For more information, see the *Sun StorEdge T3 and T3+ Array Administrator's Guide*.

8. **Verify that the controller's state has been changed to `ONLINE`.**

▼ How to Replace a Storage Array Controller

Steps

1. **Detach the submirrors on the storage array. This array is connected to the controller that you are replacing. Detach the submirrors to stop all I/O activity to this storage array.**

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

2. **Replace the controller.**

   For the procedure about how to replace a storage array controller, see the *Sun StorEdge T3 and T3+ Array Installation, Operation, and Service Manual*.

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

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

▼ How to Replace a Chassis

Use this procedure to replace a storage array chassis. This procedure assumes that you are retaining all FRUs other than the chassis and the midplane. To replace the chassis, you must replace both the chassis and the midplane because these components are manufactured as one part.

Caution – You must be a Sun service provider to perform this procedure. If you need to replace a chassis, contact your Sun service provider.

Steps

1. **Detach the submirrors on the storage array. This array is connected to the chassis that you are replacing. Detach the submirrors to stop all I/O activity to this storage array.**

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

2. **Replace the chassis and the midplane.**

   For the procedure about how to replace a storage array chassis, see the *Sun StorEdge T3 and T3+ Array Field Service Manual*.

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

Caution – The world wide names (WWNs) change as a result of this procedure. You must reconfigure your volume manager software to recognize the new WWNs.

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

▼ How to Replace a Host Adapter

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

Before You Begin

This procedure relies on the following prerequisites and assumptions.

- Except for the failed host adapter, your cluster is operational and all nodes are powered on.
- Your nodes are not configured with dynamic reconfiguration functionality.

  If your nodes are configured for dynamic reconfiguration **and** you are using two entirely separate hardware paths to your shared data, see the *Sun Cluster Hardware Administration Manual for Solaris OS* and skip steps that instruct you to shut down the cluster.

  You cannot replace a single, dual-port HBA that has quorum configured on that storage path by using DR. Follow all steps in the procedure. For the details on the risks and limitations of this configuration, see "Configuring Cluster Nodes With a Single, Dual-Port HBA" in *Sun Cluster 3.0-3.1 Hardware Administration Manual for Solaris OS*.

  Exceptions to this restriction include three-node or larger cluster configurations where no storage device has a quorum device configured.

Steps

1. **Determine the resource groups and device groups that are running on `Node A`.**

   Record this information because you use this information in Step 9 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. **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.

4. **Power off `Node A`.**

5. **Replace the failed host adapter.**

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

6. **Do you need to upgrade the node's host adapter firmware?**

   - **If yes, boot `Node A` into noncluster mode. Proceed to Step 7.**

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

   - **If no, proceed to Step 8.**

7. **Upgrade the host adapter firmware on `Node A`.**

   PatchPro is a patch-management tool that eases the selection and download of patches required for installation or maintenance of Sun Cluster software. PatchPro provides an Interactive Mode tool especially for Sun Cluster. The Interactive Tool makes the installation of patches easier. PatchPro's Expert Mode tool helps you to maintain your configuration with the latest set of patches. Expert Mode is especially useful for obtaining all of the latest patches, not just the high availability and security patches.

   To access the PatchPro tool for Sun Cluster software, go to http://www.sun.com/PatchPro/, click Sun Cluster, then choose either Interactive Mode or Expert Mode. Follow the instructions in the PatchPro tool to describe your cluster configuration and download the patches.

   For third-party firmware patches, see the SunSolve^{SM} Online site at `http://sunsolve.ebay.sun.com`.

8. **Boot `Node A` into cluster mode.**

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

9. **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.

▼ How to Remove a Storage Array in a Single-Controller Configuration

Use this procedure to remove a storage array and its submirrors from a running cluster. This procedure provides the flexibility to remove the host adapters from the nodes for the storage array that you are removing. To remove a partner group from the cluster, see "How to Remove a Partner Group" on page 59.

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 `scdidadm -c` command, the following error message appears on your console if the device ID changed unexpectedly.

`device id for` *`nodename`*`:/dev/rdsk/cXtYdZsN does not match physical device's id for d`*`decimalnumber`*`, device may have been replaced.`

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

This procedure defines `Node A` as the node with which you begin working. `Node B` is another node in the cluster.

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

Steps

1. **Back up all database tables, data services, and volumes that are associated with the storage array. This storage array is the storage array you are removing.**

2. **Detach the submirrors from the storage array that you are removing. Detach the submirrors to stop all I/O activity to the storage array.**

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

3. **Remove the references to the LUN(s) from any diskset or disk group.**

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

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

   ```
   # scstat
   ```

5. **Shut down `Node A`.**

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

6. **Is the storage array that you are removing the last storage array that is connected to `Node A`?**

   - **If yes, disconnect the fiber-optic cable between `Node A` and the FC hub/switch that is connected to this storage array. Afterward, disconnect the fiber-optic cable between the FC hub/switch and this storage array.**
   - **If no, proceed to Step 7.**

   For the procedure about how to remove a fiber-optic cable, see the *Sun StorEdge T3 and T3+ Array Configuration Guide*.

   Note – If you are use your storage array in a SAN-configured cluster, you must keep two FC switches configured in parallel. This configuration maintains cluster availability. See "SAN Solutions in a Sun Cluster Environment" in *Sun Cluster 3.0-3.1 Hardware Administration Manual for Solaris OS* for more information.

   ---

7. **Do you want to remove the host adapter from `Node A`?**

   - **If yes, power off `Node A`.**
   - **If no, skip to Step 10.**

8. **Remove the host adapter from `Node A`.**

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

9. **Without enabling the node to boot, power on `Node A`.**

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

10. **Boot `Node A` into cluster mode.**

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

11. **Shut down `Node B`.**

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

12. **Is the storage array that you are removing the last storage array that is connected to the FC hub/switch.**

    - **If yes, disconnect the fiber-optic cable that connects this FC hub/switch and `Node B`.**

- **If no, proceed to Step 13.**

For the procedure about how to remove a fiber-optic cable, see the *Sun StorEdge T3 and T3+ Array Configuration Guide*.

Note – If you are use your storage array in a SAN-configured cluster, you must keep two FC switches configured in parallel. This configuration maintains cluster availability. See "SAN Solutions in a Sun Cluster Environment" in *Sun Cluster 3.0-3.1 Hardware Administration Manual for Solaris OS* for more information.

---

13. **Do you want to remove the host adapter from `Node B`?**

    - **If yes, power off `Node B`.**
    - **If no, skip to Step 16.**

14. **Remove the host adapter from`Node B`.**

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

15. **Without enabling the node to boot, power on `Node B`**

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

16. **Boot `Node B` into cluster mode.**

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

17. **On all nodes, update the `/devices` and `/dev` entries.**

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

18. **Return the resource groups and device groups that 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 Remove a Partner Group

Use this procedure to permanently remove storage array partner groups and their submirrors from a running cluster. To remove a storage array in single-controller configuration, see "How to Remove a Storage Array in a Single-Controller Configuration" on page 57.

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 `scdidadm -c` command, the following error message appears on your console if the device ID changed unexpectedly.

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

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

This procedure defines `Node A` as the cluster node that you begin working with. `Node B` is another node in the cluster.

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

Steps

1. **If necessary, back up all database tables, data services, and volumes associated with each partner group that you are removing.**

2. **If necessary, detach the submirrors from each storage array or partner group that you are removing. Detach the submirrors to stop all I/O activity to the storage array or partner group.**

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

3. **Remove references to each LUN. Each LUN belongs to the storage array 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 that are running on all nodes.**

   Record this information because you use this information in Step 19 of this procedure to return resource groups and device groups to these nodes.

   ```
   # scstat
   ```

5. **Shut down `Node A`.**

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

6. **Disconnect from both storage arrays the fiber-optic cables connecting to the FC switches, then the Ethernet cable(s).**

7. **Is any storage array that you are removing the last storage array connected to an FC switch that is on `Node A`?**

   - **If no, skip to Step 11.**
   - **If yes, disconnect the fiber-optic cable between `Node A` and the FC switch that was connected to this storage array.**

     For the procedure about how to remove a fiber-optic cable, see the *Sun StorEdge T3 and T3+ Array Configuration Guide*.

Note – If you are using your storage arrays in a SAN-configured cluster, you must keep two FC switches configured in parallel. This configuration maintains cluster availability. See "SAN Considerations" on page 69 for more information.

---

8. **Do you want to remove the host adapters from `Node A`?**

   - **If no, skip to Step 11.**
   - **If yes, power off `Node A`.**

9. **Remove the host adapters from `Node A`.**

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

10. **Without enabling the node to boot, power on `Node A`**

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

11. **Boot `Node A` into cluster mode.**

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

12. **Shut down `Node B`.**

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

13. **Is any storage array that you are removing the last storage array connected to an FC switch that is on `Node B`?**

    - **If no, proceed to Step 14.**
    - **If yes, disconnect the fiber-optic cable that connects this FC switch to `Node B`.**

      For the procedure about how to remove a fiber-optic cable, see the *Sun StorEdge T3 and T3+ Array Configuration Guide*.

Note – If you are using your storage arrays in a SAN-configured cluster, you must keep two FC switches configured in parallel. This configuration maintains cluster availability. See "SAN Considerations" on page 69 for more information.

14. **Do you want to remove the host adapters from `Node B`**

    - **If no, proceed to Step 17.**
    - **If yes, power off `Node B`.**

15. **Remove the host adapters from `Node B`.**

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

16. **Without enabling the node to boot, power on `Node B`**

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

17. **Boot `Node B` into cluster mode.**

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

18. **On all nodes, update the `/devices` and `/dev` entries.**

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

19. **Return the resource groups and device groups that you identified in Step 4 to all nodes.**

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

Upgrading Sun StorEdge T3 or T3+ Storage Arrays

This section contains the procedures about how to upgrade storage arrays. The following table lists these procedures.

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 `scdidadm -c` command, the following error message appears on your console if the device ID changed unexpectedly.

`device id for` *nodename*`:/dev/rdsk/c`*X*`t`*Y*`d`*Z*`s`*N* `does not match physical device's id for d`*decimalnumber*`, device may have been replaced.`

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

TABLE 2–2 Task Map: Maintaining a Storage Array

Task	Information
Upgrade storage array firmware.	"How to Upgrade Storage Array Firmware (No Submirrors)" on page 63
Upgrade a StorEdge T3 array controller to a StorEdge T3+ array controller.	"How to Upgrade a StorEdge T3 Controller to a StorEdge T3+ Controller" on page 65
Migrate to a partner group	

▼ How to Upgrade Storage Array Firmware (No Submirrors)

Use this procedure to upgrade storage array firmware in a running cluster, when your arrays are not configured to support submirrors. To upgrade firmware when you are using submirrors, see "How to Upgrade Storage Array Firmware When Using Mirroring" on page 64. Firmware includes controller firmware, unit interconnect card (UIC) firmware, and disk drive firmware.

Caution – Perform this procedure on one storage array at a time. This procedure requires that you reset the storage arrays that you are upgrading. If you reset more than one storage array at a time, your cluster loses access to data.

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

Steps

1. **On one node that is attached to the storage array, detach the submirrors. This storage array is the storage array that you are upgrading.**

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

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

   PatchPro is a patch-management tool that eases the selection and download of patches required for installation or maintenance of Sun Cluster software. PatchPro provides an Interactive Mode tool especially for Sun Cluster. The Interactive Tool makes the installation of patches easier. PatchPro's Expert Mode tool helps you to maintain your configuration with the latest set of patches. Expert Mode is especially useful for obtaining all of the latest patches, not just the high availability and security patches.

   To access the PatchPro tool for Sun Cluster software, go to http://www.sun.com/PatchPro/, click Sun Cluster, then choose either Interactive Mode or Expert Mode. Follow the instructions in the PatchPro tool to describe your cluster configuration and download the patches.

   For third-party firmware patches, see the SunSolve^{SM} Online site at `http://sunsolve.ebay.sun.com`.

3. **Reset the storage array, if you have not already done so.**

   For the procedure about how to reboot a storage array, see the *Sun StorEdge T3 and T3+ Array Installation, Operation, and Service Manual*.

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

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

▼ How to Upgrade Storage Array Firmware When Using Mirroring

Use this procedure to upgrade out-of-date controller firmware, disk drive firmware, or unit interconnect card (UIC) firmware. This procedure assumes that your cluster is operational. This procedures defines `Node A` as the node on which you are upgrading firmware. `Node B` is another node in the cluster.

Caution – Perform this procedure on one storage array at a time. This procedure requires that you reset the storage arrays that you are upgrading. If you reset more than one storage array at a time, your cluster loses access to data.

Steps

1. **On the node that currently owns the disk group or diskset to which the mirror belongs, detach the storage array logical volume. This storage array is the storage**

array on which you are upgrading firmware.

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

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

   For the list of required storage array patches, see the *Sun StorEdge T3 Disk Tray Release Notes*. For the procedure about how to apply firmware patches, see the firmware patch `README` file. For the procedure about how to verify the firmware level, see the *Sun StorEdge T3 Disk Tray Release Notes*.

3. **Disable the storage array controller that is attached to `Node B`. Disable the controller so that all logical volumes are managed by the remaining controller.**

   For more information, see the *Sun StorEdge T3 and T3+ Array Administrator's Guide*.

4. **On one node that is connected to the partner group, verify that the storage array controllers are visible to the node.**

   ```
   # format
   ```

5. **Enable the storage array controller that you disabled in Step 3.**

6. **Reattach the mirrors that you detached in Step 1 to resynchronize the mirrors.**

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

▼ How to Upgrade a StorEdge T3 Controller to a StorEdge T3+ Controller

Use the following procedure to upgrade a StorEdge T3 storage array controller to a StorEdge T3+ storage array controller in a running cluster.

Caution – Perform this procedure on one storage array at a time. This procedure requires you to reset the storage arrays that you are upgrading. If you reset more than one storage array at a time, your cluster loses access to data.

Steps

1. **On one node that is attached to the StorEdge T3 storage array in which you are upgrading the controller, detach that storage array's submirrors.**

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

2. **Upgrade the StorEdge T3 storage array controller to a StorEdge T3+ storage array controller.**

   For instructions, see the *Sun StorEdge T3 Array Controller Upgrade Manual* .

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

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

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

Use this procedure to migrate your storage arrays from a single-controller (noninterconnected) configuration to a partner-group (interconnected) configuration. This procedure assumes that the two storage arrays in the partner-group configuration are correctly isolated from each other on separate FC switches. Do not disconnect the cables from the FC switches or the nodes.

Caution – You must be a Sun service provider to perform this procedure. If you need to migrate from a single-controller configuration to a partner-group configuration, contact your Sun service provider.

Steps

1. **Back up all data on the storage arrays before you remove the storage arrays from the Sun Cluster configuration.**

2. **Remove the noninterconnected storage arrays to be in your partner group from the cluster configuration.**

   Follow the procedure in "How to Remove a Partner Group" on page 59.

3. **Connect and configure the single storage arrays to form a partner group.**

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

4. **Ensure that each storage array has a unique target address.**

   For the procedure about how to verify and assign a target address to a storage array, see the *Sun StorEdge T3 and T3+ Array Configuration Guide*.

5. **Ensure that the `cache` and `mirror` settings for each storage array are set to `auto`.**

6. **Ensure that the `mp_support` parameter for each storage array is set to `mpxio`.**

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

   PatchPro is a patch-management tool that eases the selection and download of patches required for installation or maintenance of Sun Cluster software. PatchPro provides an Interactive Mode tool especially for Sun Cluster. The Interactive Tool makes the installation of patches easier. PatchPro's Expert Mode tool helps you to maintain your configuration with the latest set of patches. Expert Mode is especially useful for obtaining all of the latest patches, not just the high availability and security patches.

To access the PatchPro tool for Sun Cluster software, go to http://www.sun.com/PatchPro/, click Sun Cluster, then choose either Interactive Mode or Expert Mode. Follow the instructions in the PatchPro tool to describe your cluster configuration and download the patches.

For third-party firmware patches, see the SunSolve^{SM} Online site at `http://sunsolve.ebay.sun.com`.

8. **If necessary, install the required Solaris patches for storage array support on `Node A`.**

   PatchPro is a patch-management tool that eases the selection and download of patches required for installation or maintenance of Sun Cluster software. PatchPro provides an Interactive Mode tool especially for Sun Cluster. The Interactive Tool makes the installation of patches easier. PatchPro's Expert Mode tool helps you to maintain your configuration with the latest set of patches. Expert Mode is especially useful for obtaining all of the latest patches, not just the high availability and security patches.

   To access the PatchPro tool for Sun Cluster software, go to http://www.sun.com/PatchPro/, click Sun Cluster, then choose either Interactive Mode or Expert Mode. Follow the instructions in the PatchPro tool to describe your cluster configuration and download the patches.

   For third-party firmware patches, see the SunSolve^{SM} Online site at `http://sunsolve.ebay.sun.com`.

9. **Shut down `Node A`.**

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

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

    ```
    # boot -r
    ```

11. **On `Node A`, update the `/devices` and `/dev` entries.**

    ```
    # devfsadm -C
    ```

12. **On `Node A`, update the paths to the DID instances.**

    ```
    # scdidadm -C
    ```

13. **Label the new storage array logical volume.**

    For the procedure about how to label a logical volume, see the *Sun StorEdge T3 and T3+ Array Administrator's Guide*.

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

    PatchPro is a patch-management tool that eases the selection and download of patches required for installation or maintenance of Sun Cluster software. PatchPro provides an Interactive Mode tool especially for Sun Cluster. The Interactive Tool

makes the installation of patches easier. PatchPro's Expert Mode tool helps you to maintain your configuration with the latest set of patches. Expert Mode is especially useful for obtaining all of the latest patches, not just the high availability and security patches.

To access the PatchPro tool for Sun Cluster software, go to http://www.sun.com/PatchPro/, click Sun Cluster, then choose either Interactive Mode or Expert Mode. Follow the instructions in the PatchPro tool to describe your cluster configuration and download the patches.

For third-party firmware patches, see the SunSolve^{SM} Online site at `http://sunsolve.ebay.sun.com`.

15. **If necessary, install the required Solaris patches for storage array support on `Node B`.**

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

16. **Shut down `Node B`.**

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

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

    ```
    # boot -r
    ```

18. **On `Node B`, update the `/devices` and `/dev` entries.**

    ```
    # devfsadm -C
    ```

19. **On `Node B`, update the paths to the DID instances.**

    ```
    # scdidadm -C
    ```

20. **(Optional) On `Node B`, verify that the DIDs are assigned to the new LUNs.**

    ```
    # scdidadm -l
    ```

21. **On one node that is attached to the new storage arrays, reset the SCSI reservation state.**

    ```
    # scdidadm -R n
    ```

    Where *n* is the DID instance of a storage array LUN that you are adding to the cluster.

    Note – Repeat this command on the same node for each storage array LUN that you are adding to the cluster.

    ---

22. **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.

Next Steps The best way to enable multipathing for a cluster is to install the multipathing software and enable multipathing before installing the Sun Cluster software and establishing the cluster. For this procedure, see "How to Install Sun Multipathing Software" in *Sun Cluster Software Installation Guide for Solaris OS*. If you need to add multipathing software to an established cluster, see "How to Install Sun Multipathing Software" in *Sun Cluster Software Installation Guide for Solaris OS* and follow the troubleshooting steps to clean up the device IDs.

SAN Considerations

This section contains information about how to use storage array in a SAN. This information is specific to a SAN in a Sun Cluster environment. Use the cluster-specific procedures in this chapter to install and maintain a storage array in your cluster.

For instructions about how to create and maintain a SAN, see the documentation that shipped with your switch hardware. For information on switch ports, zoning, and required software and firmware, also see the documentation that shipped with your switch hardware.

SAN hardware includes the following components.

- FC switches
- FC host adapters
- Storage devices
- Enclosures

SAN software includes the following components.

- Drivers that are packaged with the operating system
- Firmware for the switches
- Management tools for the switches and storage devices
- Volume management software
- Other administration tools

SAN Clustering Considerations

If you are replacing an FC switch and you intend to save the switch IP configuration for restoration to the replacement switch, wait to connect the cables to the replacement switch. Connect the cables to the replacement switch *after* you recall the Fabric configuration to the replacement switch. For more information about how to save and recall switch configurations, see the documentation that shipped with your switch hardware.

Index

M

N

P

R

S

T

U