Sun StorEdge™ 5310 NAS Troubleshooting Guide

Sun Microsystems, Inc.
www.sun.com

Part No. 817-7513-11
August 2004, Revision A

Submit comments about this document at: `http://www.sun.com/hwdocs/feedback`

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

This document and the product to which it pertains are distributed under licenses restricting their use, copying, distribution, and decompilation. No part of the product or of this document may be reproduced in any form by any means without prior written authorization of Sun and its licensors, if any.

Third-party software, including font technology, is copyrighted and licensed from Sun suppliers.

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

Sun, Sun Microsystems, the Sun logo, AnswerBook2, docs.sun.com, Sun StorEdge, Java, and Solaris are trademarks or registered trademarks of Sun Microsystems, Inc. in the U.S. and in other countries.

Mozilla, Netscape, and Netscape Navigator are trademarks or registered trademarks of Netscape Communications Corporation in the United States and other countries.

All SPARC trademarks are used under license and are trademarks or registered trademarks of SPARC International, Inc. in the U.S. and in other countries. Products bearing SPARC trademarks are based upon an architecture developed by Sun Microsystems, Inc.

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

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

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

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

Ce produit ou document est protégé par un copyright et distribué avec des licences qui en restreignent l'utilisation, la copie, la distribution, et la décompilation. Aucune partie de ce produit ou document ne peut être reproduite sous aucune forme, par quelque moyen que ce soit, sans l'autorisation préalable et écrite de Sun et de ses bailleurs de licence, s'il y ena.

Le logiciel détenu par des tiers, et qui comprend la technologie relative aux polices de caractères, est protégé par un copyright et licencié par des fournisseurs de Sun.

Des parties de ce produit pourront être dérivées des systèmes Berkeley BSD licenciés par l'Université de Californie. UNIX est une marque déposée aux Etats-Unis et dans d'autres pays et licenciée exclusivement par X/Open Company, Ltd.

Sun, Sun Microsystems, le logo Sun, AnswerBook2, docs.sun.com, Sun StorEdge, Java, et Solaris sont des marques de fabrique ou des marques déposées de Sun Microsystems, Inc. aux Etats-Unis et dans d'autres pays.

Mozilla, Netscape, et Netscape Navigator sont des marques de Netscape Communications Corporation aux Etats-Unis et dans d'autres pays.

Toutes les marques SPARC sont utilisées sous licence et sont des marques de fabrique ou des marques déposées de SPARC International, Inc. aux Etats-Unis et dans d'autres pays. Les produits portant les marques SPARC sont basés sur une architecture développée par Sun Microsystems, Inc.

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

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

Please
Recycle

Adobe PostScript™

Contents

Tables

Figures

Preface

This Troubleshooting Guide provides information on how to identify, isolate, and fix problems with the Sun StorEdge^{TM} 5310 NAS. It also explains how to remove and replace certain key server components.

Topics in this chapter include:

Who Should Use This Book

The intended audience for this book is Sun field service personnel who are responsible for maintaining Sun StorEdge 5310 NAS.

How This Manual is Organized

This manual contains the following chapters:

- Chapter 1, “Troubleshooting Overview” on page 1-1
- Chapter 2, “NAS Head” on page 2-1
- Chapter 3, “Storage Arrays” on page 3-1
- Chapter 4, “StorEdge File Replicator” on page 4-1
- Chapter 5, “Clustering” on page 5-1
- Chapter 6, “Checkpoints/Snapshots” on page 6-1
- Chapter 7, “FRU/CRU Replacement Procedures” on page 7-1

Typographic Conventions

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

TABLE P-1 Typographic Conventions

Typeface or Symbol	Meaning	Example

courier font	Names of commands;
	Names of files;
	On-screen computer output;
		Use `ls -a` to list all files.
		Edit your `.login` file.
		`machine_name% You have mail`.

italics	Book titles, new words;	Read Chapter 6 in the *User’s Guide*;

	Terms to be emphasized;
	Variables that you replace with a real value;
		These are called *class* options;
		You must be *root* to do this;
		To delete a file, type *rm filename*.

boldface courier font	What you type	**`machine_name% su`**

Related Documentation

These documents contain information related to the tasks described in this book:

Sun StorEdge 5310 NAS Quick Reference Manual
Sun StorEdge 5310 NAS Hardware Installation, Configuration, and User Guide
Sun StorEdge 5310 NAS Software Installation, Configuration, and User Guide
Sun StorEdge 5310 NAS Setup Poster

Ordering Sun Documents

The SunDocsSM program provides more than 250 manuals from Sun Microsystems, Inc. If you are in the United States, Canada, Europe or Japan, you can purchase documentation sets or individual manuals by using this program.

For a list of documents and how to order them, see the catalog section of the SunExpress™ Internet site at http://store.sun.com.

Accessing Sun Documentation Online

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

Shell Prompts in Command Examples

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

TABLE P-2 Shell Prompt

Shell	Prompt
Bourne shell and Korn shell prompt	machine name$
Bourne shell and Korn shell superuser prompt	machine name#

Sun Welcomes Your Comments

Sun is interested in improving its documentation and welcomes your comments and suggestions. You can email your comments to Sun at:

`docfeedback@sun.com`

Please include the part number (8xx-xxxx-xx) of your document in the subject line of your email.

CHAPTER 1

Troubleshooting Overview

This chapter provides an overview of diagnostic functions and tools needed for troubleshooting the Sun StorEdge 5310 NAS.

This chapter contains the following sections:

1.1 How to Use This Manual

Before going deep into this manual, check the following to ensure that common problems have been resolved.

- Are both of the power cords plugged in?
- Are green LEDs displaying on the power sources? If no, check the power source.
- Does the LCD Display panel show the system name and CPU% on it? If no, check the power source.
- Can you ping the system? If no, check the network cables and IP address on the LCD Display. If you are still having problems, check with your system administrator.
- If the user can't access shares, are the shares set up on the system? Check the shares section to make sure that the shares are set up with the proper name.
- Is an NFS client having permissions issues on a CIFS file? Vice versa? Check the FAQ for file permission issues to resolve.

1.2 Important Notices and Information on the Sun StorEdge 5310 NAS

Caution – Do not plug a USB keyboard into the front USB connector. This will cause the system to crash.

Caution – Do Not power on the Sun StorEdge 5310 NAS, until two minutes after the JBOD has been powered up, to ensure that the disk drives have finished spinning up.

Caution – /dvol/etc folder contains config information and needs to be backed up to ensure that all configuration information is available upon a failure. Back up the /dvol/etc folder to an existing LUN on the Sun StorEdge 5310 NAS.

Note – `/dvol/etc` folder contains config information and needs to be backed up to ensure that all configuration information is available upon a failure. It is recommended to back the `/dvol/etc` folder up to an existing LUN on the Sun StorEdge 5310 NAS.

Note – You must enable FTP from the CLI using the `load ftpd` command. Currently, enabling the FTP from the web interface does not work.

Note – When configuring the Sun StorEdge 5310 NAS through a firewall, ensure that the correct ports are not blocked. Refer to “StorEdge Web Admin does not work properly through a firewall.” on page 2-80 for more details.

Note – There is a line of tape that must be removed to be able to remove the fan tray.

1.3 Troubleshooting Tools

1.3.0.1 Storage Automated Diagnostic Environment (StorAde)

If you have the Storage Automated Diagnostic Environment installed in the host, check the internal status of the array with this tool. See the documentation for this tool for further information.

All that you need to use the Storage Automated Diagnostic Environment is web browser access to the host where it is installed.

1.3.0.2 Command Line Interface (CLI)

The CLI can be accessed through the MENU system or by using Telnet. This is a useful sections for troubleshooting many types of issues. The CLI is also where you load tools like FTP. See the Diagnostic Tools and Procedures section for details.

1.3.0.3 Log Error Messages

Both the Sun StorEdge 5310 NAS and attached hosts create log message files or error messages of system conditions and events. These log files are the most useful *immediate* tools for troubleshooting.

1.3.0.4 Sun StorEdge 5310 NAS Generated Messages

A `syslog` daemon in the array writes system error message logs to a location determined by the site system administrator. Consult with the site system administrator to obtain access to this log.

1.3.0.5 Client Generated Messages

CIFS clients will get messages on the monitor when they have attached shares on the Sun StorEdge 5310 NAS. These messages will be useful in determining issues that arise.

NFS clients will have messages generated in its /var/adm/messages file.

A variety of software logging tools monitor the various branches of the storage network. When an error is detected, the error's severity level is categorized and classified. Errors are reported or logged according to severity level.

1.3.0.6 Log Message Severity Levels

- Emergency—Specifies emergency messages. These messages are not distributed to all users. Emergency priority messages are logged into a separate file for reviewing.
- Alert—Specifies important messages that require immediate attention. These messages are distributed to all users.
- Critical—Specifies critical messages not classified as errors, such as hardware problems. Critical and higher-priority messages are sent to the system console.
- Error—Specifies any messages that represent error conditions, such as an unsuccessful disk write.
- Warning—Specifies any messages for abnormal, but recoverable, conditions.
- Notice—Specifies important informational messages. Messages without a priority designation are mapped into this priority message.
- Information—Specifies informational messages. These messages are useful in analyzing the system.
- Debug—Specifies debugging messages.

1.4 Troubleshooting Procedures

1.4.0.1 High-Level Troubleshooting Tasks

This section lists the high-level steps you can take to isolate and troubleshoot problems in the array. It offers a methodical approach, and lists the tools and resources available at each step.

1. **Discover the error by checking one or more of the following messages or files:**
   - Storage Automated Diagnostic Environment alerts or email messages, if available
   - "event log" from the Sun StorEdge 5310 NAS
   - `/var/adm/messages` file at the host system
   - CIFS clients messages
2. **Determine the extent of the problem by using one or more of the following methods:**

- Review the Storage Automated Diagnostic Environment topology view
- Using the Storage Automated Diagnostic Environment revision checking
- functionality, determine whether the package or patch is installed

3. **Check the status of a Sun StorEdge 5310 NAS by using one or more of the following methods:**
   - Review the status of the light-emitting diodes (LED) on the array
   - Run the commands that check and display the configuration
   - Manually open a telnet session to the array and check the system status
   - Review the Storage Automated Diagnostic Environment device monitoring reports, if available
4. **Test and isolate field-replaceable units (FRUs) using the following tools:**
   - Storage Automated Diagnostic Environment diagnostic tests, if available (these tests might require a loopback cable for isolation)
   - Use the Troubleshooting Guide procedures documentation to help isolate FRU failures

---

Note – These tests isolate the problem to a FRU that must be replaced. Follow the instructions in the *Sun StorEdge 5310 NAS Troubleshooting Guide* for proper FRU replacement procedures.

---

5. **Replace the failed FRU.**
6. **Verify the fix using the following tools:**
   - Storage Automated Diagnostic Environment GUI Topology View and Diagnostic Tests, if available
   - `/var/adm/messages` on the data host
   - CIFS client Access
   - Array LEDs
   - `syslog` file

1.4.0.2 Initial Troubleshooting Guidelines

To begin a problem analysis, check one or more of the following information sources for troubleshooting and perform one or more of the following checks:

- The LED's can help you quickly identify if a problem is occurring. See the Hardware Troubleshooting section to help isolate the failed component.

- Sun StorEdge 5310 NAS messages, found in the `syslog` file, indicating a problem. See Error Messages section for more information about array generated messages.
- Host-generated message, found in the `/var/adm/messages` file, CIFS clients may have errors on their monitor or in the event log.

1.5 Troubleshooting Flow Charts

Use the flow charts below to diagnose problems.

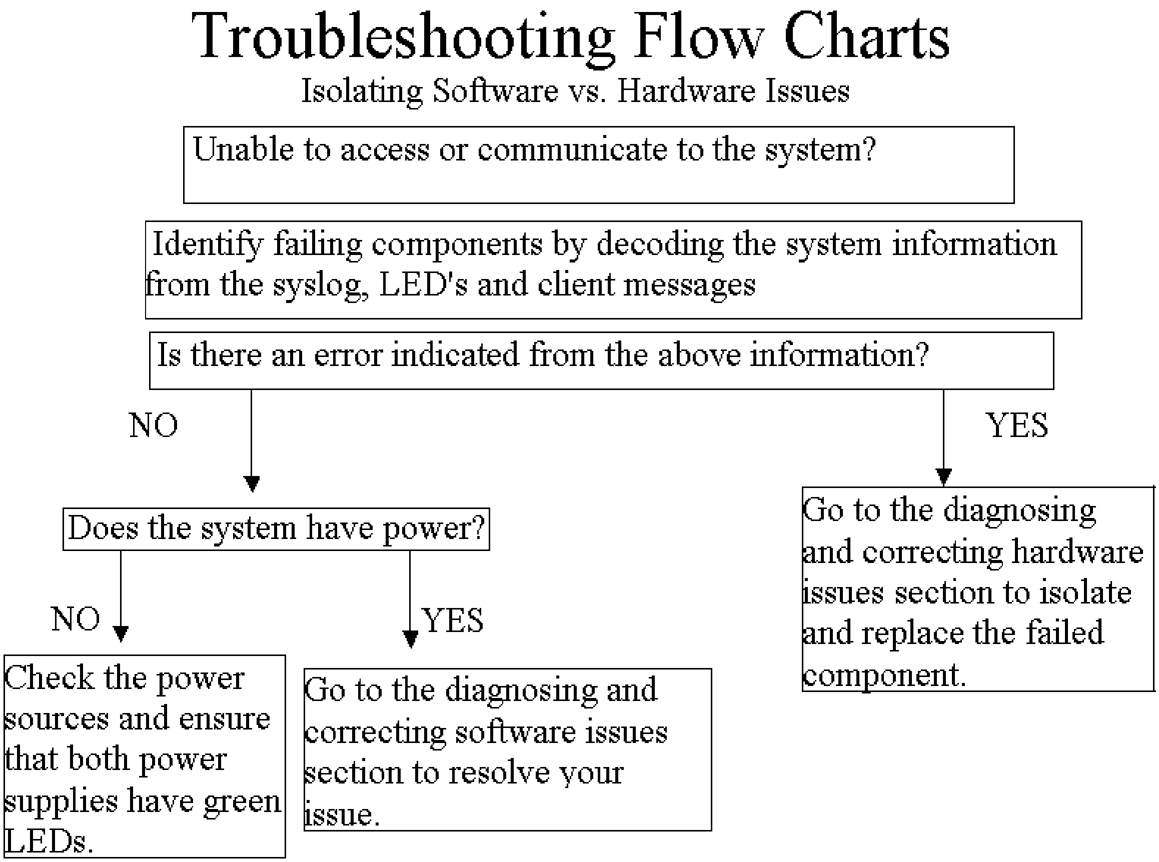

Troubleshooting Flow Charts

Isolating Software vs. Hardware Issues

Unable to access or communicate to the system?

Identify failing components by decoding the system information from the syslog, LED's and client messages

Is there an error indicated from the above information?

NO

Does the system have power?

NO

Check the power sources and ensure that both power supplies have green LEDs.

YES

Go to the diagnosing and correcting software issues section to resolve your issue.

YES

Go to the diagnosing and correcting hardware issues section to isolate and replace the failed component.

Follow the steps below to diagnose hardware problems.

Troubleshooting Flow Charts

Hardware Flow Chart

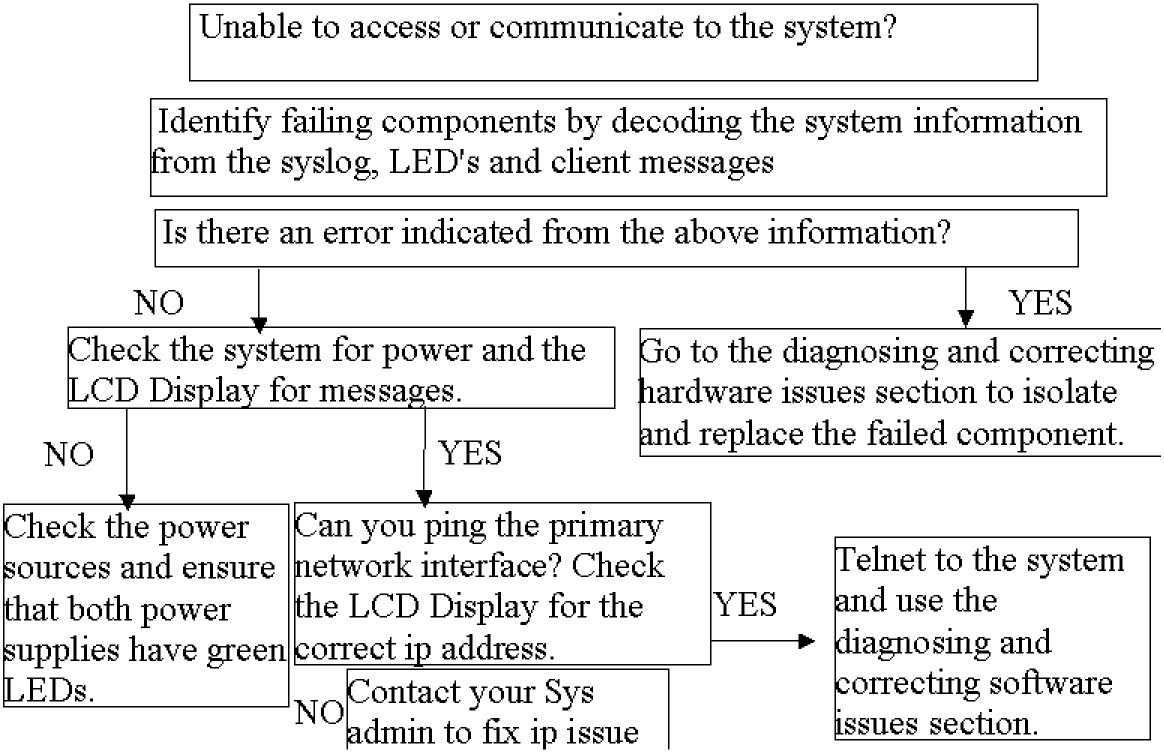
Unable to access or communicate to the system?
Identify failing components by decoding the system information from the syslog, LED's and client messages
Is there an error indicated from the above information?
NO
YES
Check the system for power and the LCD Display for messages.
Go to the diagnosing and correcting hardware issues section to isolate and replace the failed component.
NO
YES
Check the power sources and ensure that both power supplies have green LEDs.
Can you ping the primary network interface? Check the LCD Display for the correct ip address.
YES
Telnet to the system and use the diagnosing and correcting software issues section.
NO
Contact your Sys admin to fix ip issue

Follow the steps below to diagnose software problems.

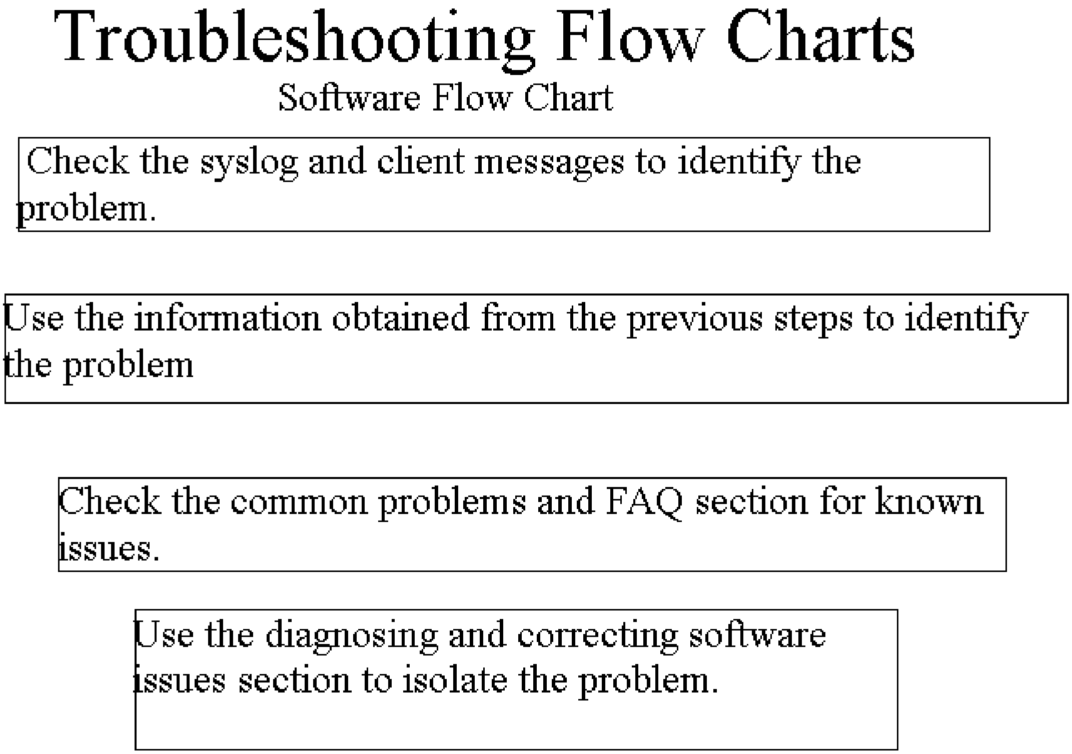

Troubleshooting Flow Charts

Software Flow Chart

Check the syslog and client messages to identify the problem.

Use the information obtained from the previous steps to identify the problem

Check the common problems and FAQ section for known issues.

Use the diagnosing and correcting software issues section to isolate the problem.

1.6 Diagnostic Information Sources

1.6.1 StorEdge Diagnostic Email

The diagnostic email includes information about the StorEdge system configuration, disk subsystem, file system, network configuration, SMB shares, backup/restore information, /etc information, system log, environment data and administrator information. The diagnostics are a primary tool for checking configuration and troubleshooting.

Before you can send email diagnostics from the StorEdge, SMTP (email) must be configured. Please see the FAQ, “How do I set up SMTP (email)?”

To collect diagnostics, proceed as follows:

1. **Access the StorEdge via Telnet or serial console.**
2. **Press enter at the [menu] prompt and enter the administrator password.**
3. **Press the spacebar until “Diagnostics” is displayed under “Extensions” at the lower right.**
4. **Select the letter corresponding to “Diagnostics”.**
5. **Wait a few seconds while the StorEdge builds the diagnostic.**
6. **Select option “2”, Send Email**
7. **Select option “1”, Edit problem description**
8. **Enter a precise description of the problem**
9. **Press [Enter]**
10. **Select option “8”, Send Email**

    Diagnostic is sent

If an email server is not configured or not available, it is also possible to save the diagnostics to a file on the StorEdge. To do this, proceed as above to access the “Diagnostics” menu.

1. **Select option “1”, Save File.**
2. **Select option “1”, Edit path**
3. **Enter a valid path name in the path box. Format is /<volumename>/<directory>/<new filename>.**
4. **Press [Enter]**
5. **Select option “2”, save diagnostics file**

   System will respond with diagnostic saved
6. **Access the volume that you saved the file to with SMB or NFS.**
7. **Copy the file to a local workstation**

Important – Saving the diagnostics file locally will not include the necessary attachments. When escalating an issue with diagnostics, you must also include the contents of the /etc directory, and the contents of /cvol/log.

This functionality is also available through the StorEdge Web Admin. To access these settings, log in, and click the envelope icon on the top taskbar. All of the options described above are available.

1.6.2 Data Collection for Escalations

1.6.2.1 Collecting Information from the Sun StorEdge 5310 NAS

The following are important considerations for data collection. Data collection is critical in cases that require escalation. We should always collect as much data as needed to resolve the worst-case scenario, in order to be able to resolve all scenarios. The worst-case scenario in this case, is that the issue has never before been seen, and we'll need to recreate the problem in the lab. To do this, we'll need to know about the client systems, the workload, the network, and so on.

1.6.2.2 Accurately quantify the problem

First, the problem must be quantified. We have identified a negative behavior of some type. We must precisely identify the scope of the problem and all possible details in order to resolve the issue. For example, if the StorEdge has a performance issue, we must exactly measure the performance, identify which problems exhibit the problem, and determine under what circumstances the problem occurs.

1.6.2.3 Collect general data

The first part of the data collection is to collect information that will be useful in every case. Much of this is contained in the StorEdge system diagnostics. From the diagnostics, we can see the StorEdge OS version, internal settings, recent log activity, and more. It is very important to generate the diagnostics during or immediately after the manifestation of the problem. Otherwise, the log and statistics will not show any data on the failure. Always collect a diagnostic email when escalating issues.

You should also collect any error messages generated by this problem, and any steps already taken in the attempt to resolve the problem, and the results obtained.

1.6.2.4 Collect specific data

Based on the above data, additional information may be required. This document will help you to tailor this data collection. Here are some examples:

- Version(s) of software on client system(s)
- Version(s) of software on server system(s)
- Network topology
- Steps and/or sequence of events leading to the failure
- What was the user doing or attempting to do when the failure occurred?
- Problem symptom (error codes, failed operation, crash)
- Syslog data
- Network traces
- Diagnostic email

1.6.2.5 Check remote access capabilities

In some cases, it is useful for one of your escalation resources to directly access the system. This can be a way to greatly simplify advanced data collection. Please note that this step is not always necessary or useful, but it can be a very valuable tool at times. When you know that advanced investigation will be required, it’s always wise to ask if remote access via TCP/IP or dial-up is available.

1.6.2.6 Data Collection for Specific Issues

Software compatibility issues

Some applications do not function properly when StorEdge is used in place of a server running a native operating system. Most, but not all, of these issues can be resolved with data collection and troubleshooting. It may be necessary to upgrade the application, the client operating system, or the StorEdge operating system. Keep in mind that the problem may lie in any of these, or a combination of all three.

The first step is to do research. Check to see if a newer version of the application or the StorEdge operating system is available. Check the release notes to see if the compatibility issue is addressed. If either version is far out of date, perform an upgrade to see if the problem is resolved. Another useful step is to try to operation on a other available network clients.

To escalate the issue, begin data collection by generating a system diagnostic with all attachments. If there is a specific symptom which can be identified, generate the system diagnostic as close as possible after this time, so that any effects can be observed in the logs and statistics.

The procedure for this can be found later in this document under Diagnostic Procedures. Next, it is necessary to collect as much data as possible on the client and application. At a minimum, the following information is required:

- Client Operating System version, including any service packs or minor revisions

- Software version, including any service packs, options or minor revisions
- Client configuration information– mount options, NIC configuration, platform, etc.
- Network information – topology, switch and router information, path from client to StorEdge
- Server information – Detailed information on any application or authentication servers, including all of the above details.
- An exact set of steps to reproduce the problem. This should be very detailed, including every menu selection and text entry
- Details on any symptoms experienced by the client

The goal of this data set is to allow someone in a remote location to reproduce and resolve the issue without impacting the customer.

The next step is to verify the problem and collect network traces. If possible, copy the data residing on the StorEdge to another server temporarily. Verify that it works as expected. If it still exhibits the same symptom, the issue likely resides with the application.

Use a network capture utility to capture the network traffic generated by the failure condition between the client, the StorEdge and any other server involved in the issue. Define traffic filters so that only this traffic is captured.

Next, repeat the network capture, using the server which the application runs successfully on. This will allow engineering to make a direct comparison of a successful operation and an unsuccessful operation.

StorEdge has a built-in network monitoring tool. Details on the operation of this tool can be found in the Diagnostic Procedures section of this document. However, in this case it would be best to use a network analysis tool on the client. The main reason for this is that the StorEdge tool will not be able to capture the data when an alternate server is used for comparison.

1.6.2.7 Security Issues

When troubleshooting security problems, it is useful to experiment. Try other workstations, other operating systems and different user accounts, including a root or a Domain Admin account. These are very useful in locating the source of the problem.

When escalating a security issue collect the following data:

Cacls

For issues with access to a file or directory, collect the output of the cacls command. This command is available from the CLI. At the CLI, enter “cacls <full pathname>”. The full pathname should begin with the volume name, as in this example: “cacls /vol1/testfile.txt”.

Cacls output contains the following information:

First, the basic mode information and UID/GID of the owner is displayed. Here is an example:

```
drwxrw----        34      22      /vol1/data
```

In this case, we can see that the item is a directory, with 750 permissions: Read/write/execute (7) for the owner (UID 34), Read/write for members of the owner’s group (GID 22), and no permissions (0) for everyone else.

Listed next are Creation time, FS Creation time, and FS mtime. These are timestamps associated with the file and the filesystem, generally only useful for troubleshooting timestamp issues.

Next is the Windows security descriptor. In its simplest form, it will read “No security descriptor”. This means that no Windows security is present, and that Windows will simulate security based on the above NFS permissions.

If a Windows security descriptor is present, the following information is displayed:

- Security Descriptor:The type of security descriptor. This can be disregarded.
- Owner:The user name or SID of the owner.
- Primary Group: The group name or SID of the group owner.
- Discretionary Access Control List (DACL):A list of users who have access to the file, by SID.

A SID is a number that uniquely identifies a user or group. The data to the right of the final dash identifies the user within the domain; the rest of the number indicates domain and type of account information. This user information is known as the RID (relative ID). The RID is the number used for user mapping. It can be cross-referenced with the StorEdge user or group mapping data determine the user/group name and NFS UID/GID.

User access token

For issues with the access of a particular user, it may be useful to capture the access token. The access token identifies an SMB user along with other details such as domain and group memberships. See the instructions under /proc filesystem. This item is particularly useful when the issue involves group membership. Note that this data is only useful for SMB issues.

Proc filesystem

The /proc filesystem is a virtual filesystem used to collect system data. The location of some of the more useful data is listed below. To collect the data, copy the file, or use the "cat" CLI command to dump it to the screen while logging the terminal session.

`/proc/cifs/DOMAIN.USER.6789ABCD…`

These are user access tokens. They may be useful in troubleshooting SMB issues.

These file names begin with the domain name, then the username, then some hexadecimal digits. The hexadecimal digits are a representation of the IP address, which can be used to discern between multiple logins for a user. If you do not see the user token that you need, it may be necessary to log the user off for thirty seconds, and then back on in order to capture the token.

`/proc/cifs/pdc`

The currently connected domain, domain controller, and the IP address of the domain controller.

`/proc/cifs/ntdomain`

A list of all trusted domains, their related SIDs, and the local machine and local domain SIDs.

Network trace

A network trace can be very valuable towards diagnosing problems that involve network communication. Set the trace to filter traffic between StorEdge, the client, and any authentication server. In this case, it is usually best to use the StorEdge built-in packet capture utility.

1.6.2.8 StorEdge network capture utility

StorEdge includes a built-in network monitoring tool. This allows you to capture packets from the network and save them to a file. This can be a valuable troubleshooting tool.

To configure network monitoring, it must first be loaded at the StorEdge CLI.

1. **To access the StorEdge CLI, connect to the StorEdge via Telnet or serial console, and type "admin" at the [menu] prompt and enter the administrator password.**
2. **At the CLI, enter "load netm". Then type "menu" to configure capture and capture packets.**

3. **Press the spacebar until "Packet Capture" is displayed under "Extensions" at the lower right.**

4. **Select the letter corresponding to "Packet Capture".**

5. **Select option "1", Edit Fields.**

   The available options are as follows:

   - Capture FileWhere to save the capture file. </volumename/directory/filename>
   - Frame Size (B)Size in bytes of each frame to capture. The default is normally used.
   - IP Packet Filter"No" captures all traffic, "Yes" allow you to filter what is received.

     A filter allows you to select which IP address or addresses you will capture traffic from. You can also filter on a particular TCP or UDP port.
   - Dump EnableSelect "Yes" to allow StorEdge to save the capture in the event of a problem.

6. **After configuring these options, select option "7", "Start Capture"**

7. **Reproduce the network event you wish to capture.**

8. **Select option "7", "Stop Capture".**

9. **Access the file via NFS or SMB and copy the file as needed.**

Client and Server data

Collect all possible information on the client system having the issue and any authentication or application servers involved in the issue. This information should include operating system version, patch level and platform.

Duplication instructions

If possible, provide a step-by-step procedure to recreate this problem. Include every setting and every configuration detail.

1.6.2.9 TCP/IP Connectivity problems

A good tool to investigate network connectivity problems is the netstat command. This command is available from the StorEdge CLI. Simply type “netstat” at the CLI and a list of all network interfaces and routes is displayed, along with several useful statistics. Two tables are displayed, as follows:

TABLE 1-1 List of Adapters

Name	Mtu	Netmask	Address	Ipackets	Ierr	Opackets	Oerr	Coll
lo0	1536	255.0.0.0	127.0.0.1	77	0	77	0	0
fxp1	1500	255.255.255.0	10.10.35.2	269947	0	97815	0	0
fxp2	1500	--no-address--	0	0	0	0	0	0

The first table is a list of adapters and statistics for each.

TABLE 1-2 Routing Table

Netmask	Destination	Gateway	Interf	Flags	Refs	Use
l0.0.0.0	l0.0.0.0	64.60.56.1	fxp1ug	5	70796	
255.255.255.0		64.60.56.0	10.10.35.2	fxp1	uc	00
255.255.255.255		127.0.0.1	127.0.0.1	lo0	uh	077

The second table is the routing table. The adapter “lo0” is the loopback device and does not represent a physical adapter. The route “0.0.0.0” is the default gateway. The following should be checked in this display:

- Check for typos in IP addresses and netmasks.
- Check “Ierr”, “Oerr”, and Coll”. These are all packet errors. They may indicate a bad NIC or cable, connected to the StorEdge or elsewhere, or possibly, in the case of the “coll” statistic, an incorrect speed and duplex setting.
- Check Ipackets and Opackets for the appropriate network adapter. These are packets received and sent by each adapter. A disconnected or bad cable will result in no Ipackets for the connected interface. No Opackets may indicate that there is no route defined which uses this interface.
- Check for modified gateways. A “d” or “m” in the flags column indicates a dynamically added or dynamically modified route. If an important route is modified, it may no longer be able to send packets to the desired destination.

These are the result of an ICMP message from another router or firewall, typically due to mis-configuration of that device. It is also possible to configure StorEdge to ignore ICMP requests to change the default gateway.

- Check the "Use" statistic in the routing table. This statistic indicates how many times a route has been used. If you have defined a route for a specific purpose, such as mirroring, and this counter is not incrementing, then the route was most likely not defined correctly.

Also, check the basics. Try another client on the same subnet, try another cable, try another switch port for both client and StorEdge.

To escalate TCP/IP connectivity issues collect a network trace from the StorEdge, using the internal utility, and also from the client or server attempting to connect to the StorEdge if possible. Also include details on the client system, especially network configuration information and operating system version. A network diagram which includes IP addresses and information on switches and router hardware on the network is also very helpful.

1.6.2.10 Performance Issues

The following is a general list of barriers to peak performance:

Network Configuration:

- Verify speed and duplex negotiation.
- Verify that port aggregation is configured when multiple NICs are connected to a subnet.
- Ensure that Jumbo frames are not configured.
- Ensure that Spanning Tree Protocol is not configured.
- Ensure that all configured NIS, DNS, SMTP servers and etc. are reachable and resolvable. (Note: always configure by IP rather than name where possible)

Configuration:

Checkpoints: Checkpoints can be overused and have a drastic effect of performance of the system. Verify that customers understand the use of checkpoints and how the retention policy can play a significant role in system performance.

If using ADS, improper configuration of dynamic DNS configuration can adversely affect performance.

Other processes / High CPU Utilization

When performance is low, one possible reason is that the system is busy with other processes. One way to check this is to observe the CPU utilization. This is best viewed from the activity monitor screen in the telnet interface. The CPU utilization can be found in the lower right corner, listed as a percentage.

The rest of the activity monitor screen may also be helpful, as it may give an indication of the source of the demand on resources. The display is arranged in four columns. The left most column lists each volume, and for each volume, the current disk space in use as a percentage of the volume and I/O requests. Note that a volume utilization of over 75% can cause a significant slowdown. The second column shows the load on each resource, such as CPU, memory or network adapters. These numbers do not correspond to any defined performance parameters, so they are only useful for relative comparison to another point in time. The third and fourth columns list clients currently connected to the StorEdge, and how many network I/O requests are coming from each.

Having determined that the slow server response corresponds with high CPU utilization, the next step is to collect system diagnostic while the CPU utilization is high (usually 90% or higher). The diagnostics provide a per-process breakdown of CPU and memory utilization, along with all associated log messages and configuration.

It is also possible to acquire this per-process utilization breakdown at the CLI with the “status” command. This can be useful when the CPU utilization spikes are very brief in duration, rendering them difficult to capture via a diagnostic. In this case, you would log the telnet or terminal session, and run the status command several times in succession while a performance problem is occurring. System diagnostics should also be captured to supplement this information.

Command Line performance utilities

StorEdge provides several built-in utilities designed to measure performance. These are best used to isolate a problem. For example, using `aratewrite` to write directly to the RAID set may help to determine whether a write performance problem is on a particular volume, or even the network.

Usage for these utilities is as follows:

- ratewrite: write contents of a file, report performance. The file creation does not use network connection. This can determine if issue is disk or network related.

  usage: ratewrite FILENAME [+OFFSET] TOTALKB [BLOCKSIZE]

  example:
  ```
  support > ratewrite /vol1/testfile 1000000 4096 1024000000 bytes
  (976.5M) in 36.844 seconds  26.50MB/sec
  ```

- rateread: read contents of a file, report performance. The file read does not use network connection. This can determine if issue is disk or network related and also if problem is in reading or writing data.

  usage: rateread FILENAME [+OFFSET] TOTALKB [BLOCKSIZE]

  example:
  ```
  support > rateread /vol1/testfile 8192 1024000000 bytes (976.5M) in
  0.877 seconds  1.086GB/sec
  ```
- ratecopy: copy a file, test the performance of a file copy from source to target. Uses network connection and can be used to determine if any network issues exsist.

  usage: ratecopy SOURCE_FILENAME DEST_FILENAME [BLOCKSIZE]

  example:
  ```
  support > ratecopy /vol1/testfile vol1/testout 1024000000 bytes
  (976.5M) in 25.116 seconds  38.88MB/sec
  ```
- aratewritewrite a file direct asynchronously. Test performance VS ratewrite.

  usage: aratewrite FILENAME [+OFFSET] TOTALKB [BLOCKSIZE] [MB_PER_COMMIT]

  example:
  ```
  support > aratewrite /vol1/testfile 1000000 4096 Writing 976MB in
  4KB blocksize with 0MB per commit. 1024000000 bytes (976.5M) in
  14.982 seconds  65.18MB/sec
  ```

1.6.3 Log Error Messages

1.6.4 SYSLOG

The syslog is an important tool for troubleshooting. It provides a place to begin isolating system issues. There are many levels of warnings that can be used to notify you via email that there is a problem.

1.6.4.1 Understanding Sun StorEdge 5310 NAS Log Messages

The Sun StorEdge 5310 NAS provides an Event Management subsystem that monitors the chassis and reports event information to:
- The system log, which is only in memory
- A syslog server
- SNMP Traps (SNMP v1 and v2)
- A local file on one of the created volumes
- Email notification

Components of an Event Message

Time/date   Severity

```
05/23/04 05:55:30 C sysmon[63]: Disk drive at enclosure 1 row 0 column 2
failed.
```

Facility FID   Message body

- Time/date- Time and Date of the event
- Severity- Severity can be one of those listed below
- Facility- The system module that reported the message
- FID- The kernel ID of the Facility
- Message Body- The contents of the message

Severity Level Definitions (highest to lowest)

- Emergency—Specifies emergency messages. These messages are not distributed to all users. Emergency priority messages are logged into a separate file for reviewing.
- Alert—Specifies important messages that require immediate attention. These messages are distributed to all users.
- Critical—Specifies critical messages not classified as errors, such as hardware problems. Critical and higher-priority messages are sent to the system console.
- Error—Specifies any messages that represent error conditions, such as an unsuccessful disk write.
- Warning—Specifies any messages for abnormal, but recoverable, conditions.
- Notice—Specifies important informational messages. Messages without a priority designation are mapped into this priority message.
- Information—Specifies informational messages. These messages are useful in analyzing the system.
- Debug—Specifies debugging messages.

1.6.5 Error Codes from the Sun StorEdge 5310 NAS LCD Display and syslog

This section details the specific error messages sent through e-mail, SNMP notification, the LCD panel, and the system log to notify the administrator in the event of a system error. *SysMon*, the monitoring thread in the Sun StorEdge 5310

NAS, monitors the status of RAID devices, UPSs, file systems, head units, enclosure subsystems, and environmental variables. Monitoring and error messages vary depending on model and configuration.

In the tables in this section, table columns with no entries have been deleted.

About SysMon Error Notification

SysMon, the monitoring thread in the Sun StorEdge 5310 NAS, captures events generated as a result of subsystem errors. It then takes the appropriate action of sending an e-mail, notifying the SNMP server, displaying the error on the LCD panel, writing an error message to the system log, or some combination of these actions. E-mail notification and the system log include the time of the event.

Sun StorEdge 5310 NAS Error Messages

The following sections show error messages for the Sun StorEdge 5310 NAS UPS, file system usage, and the PEMS.

UPS Subsystem Errors

Refer to Table A-3 for descriptions of UPS error conditions.

TABLE A-3 UPS Error Messages

Event	E-Mail Subject: Text	SNMP Trap	LCD Panel	Log
Power Failure	**AC Power Failure:** AC power failure. System is running on UPS battery. Action: Restore system power. Severity = Error	EnvUpsOn Battery	U20 on battery	UPS: AC power failure. System is running on UPS battery.
Power Restored	**AC power restored:** AC power restored. System is running on AC power. Severity = Notice	EnvUpsOff Battery	U21 power restored	UPS: AC power restored.
Low Battery	**UPS battery low:** UPS battery is low. The system will shut down if AC power is not restored soon. Action: Restore AC power as soon as possible. Severity = Critical	EnvUpsLow Battery	U22 low battery	UPS: Low battery condition.
Normal Battery	**UPS battery recharged:** The UPS battery has been recharged. Severity = Notice	EnvUps Normal Battery	U22 battery normal	UPS: Battery recharged to normal condition.
Replace Battery	**Replace UPS Battery:** The UPS battery is faulty. Action: Replace the battery. Severity = Notice	EnvUps Replace Battery	U23 battery fault	UPS: Battery requires replacement.
UPS Alarms - Ambient temperature or humidity outside acceptable thresholds	**UPS abnormal temperature/humidity:** Abnormal temperature/humidity detected in the system. Action: 1. Check UPS unit installation, OR 2. Contact technical support. Severity = Error	EnvUps Abnormal	U24 abnormal ambient	UPS: Abnormal temperature and/or humidity detected.

TABLE A-3 UPS Error Messages

Event	E-Mail Subject: Text	SNMP Trap	LCD Panel	Log
Write-back cache is disabled.	**Controller Cache Disabled:** Either AC power or UPS is not charged completely. Action: 1 - If AC power has failed, restore system power. 2 - If after a long time UPS is not charged completely, check UPS. Severity = Warning		Cache Disabled	write-back cache for ctrl x disabled
Write-back cache is enabled.	**Controller Cache Enabled:** System AC power and UPS are reliable again. Write-back cache is enabled. Severity = Notice		Cache Enabled	write-back cache for ctlr n enabled
The UPS is shutting down.	**UPS shutdown:** The system is being shut down because there is no AC power and the UPS battery is depleted. Severity = Critical			!UPS: Shutting down
UPS Failure	**UPS failure:** Communication with the UPS unit has failed. Action: 1. Check the serial cable connecting the UPS unit to one of the CPU enclosures, OR 2. Check the UPS unit and replace if necessary. Severity = Critical	EnvUpsFail	U25 UPS failure	UPS: Communication failure.

File System Errors

File system error messages occur when the file system usage exceeds a defined usage threshold. The default usage threshold is 95%.

TABLE A-4 File System Errors

Event	E-Mail Subject: Text	SNMP Trap	LCD Panel	Log
File System Full	**File system full:** File system <name> is xx% full. Action: 1. Delete any unused or temporary files, OR 2. Extend the partition by using an unused partition, OR 3. Add additional disk drives and extend the partition after creating a new partition. (Severity=Error)	PartitionFull	F40 FileSystemName full	File system <name> usage capacity is xx%.

PEMS Events

Sun StorEdge 5310 NAS employs the PEMS board to monitor environmental systems and to send messages regarding fan, power supply, and temperature anomalies.

Note – Device locations are shown in the *Sun StorEdge 5310 NAS Hardware Installation, Configuration, and User Guide* included in your documentation CD.

Table A-5 describes the PEMS error messages for the Sun StorEdge 5310 NAS.

TABLE A-5 PEMS Error Messages

Event	E-Mail Subject: Text	SNMP Trap	LCD Panel	Log
CPU Fan Error	**Fan Failure:** The CPU fan has failed. Fan speed = xx RPM. Action: The system will shut down in 10 seconds to protect the CPU from damage. You should replace the CPU fan before turning the system back on. Severity = Critical	envFanFail trap	P11 CPU fan failed	The CPU fan has failed! Better shut down.

TABLE A-5 PEMS Error Messages

Event	E-Mail Subject: Text	SNMP Trap	LCD Panel	Log
Fan Error	**Fan Failure:** Blower fan xx has failed. Fan speed = xx RPM. Action: The fan must be replaced as soon as possible. If the temperature begins to rise, the situation could become critical. Severity = Error	envFanFail trap	P11 Fan xx failed	Blower fan xx has failed!
Power Supply Module Failure	**Power supply failure:** The power supply unit xx has failed. Action: The power supply unit must be replaced as soon as possible. Severity = Error	envPowerFail trap	P12 Power xx failed	Power supply unit xx has failed.
Power Supply Module Temperature	**Power supply temperature critical:** The power supply unit xx is overheating. Action: Replace the power supply to avoid any permanent damage. Severity = Critical	envPowerTemp Critical trap	P22 Power xx overheated	Power supply unit xx is overheating.
Temperature Error	**Temperature critical:** Temperature in the system is critical. It is xxx Degrees Celsius. Action: 1. Check for any fan failures, OR 2. Check for blockage of the ventilation, OR 3. Move the system to a cooler place. Severity = Error	envTemperatue Error trap	P51 Temp error	The temperature is critical.
Primary Power Cord Failure	**Power cord failure:** The primary power cord has failed or been disconnected. Action: 1. Check the power cord connections at both ends, OR 2. Replace the power cord. Severity = Error	envPrimary PowerFail trap	P31 Fail PWR cord 1	The primary power cord has failed.
Secondary Power Cord Failure	**Power cord failure:** The secondary power cord has failed or been disconnected. Action: 1. Check the power cord connections at both ends, OR 2. Replace the power cord. Severity = Error	envSecondary PowerFail trap	P32 Fail PWR cord 2	The secondary power cord has failed.

1.7 Maintenance Precautions

The sections that follow provide subassembly-level removal and installation guidelines. After completing all necessary removal and replacement procedures, verify that all components are working properly.

1.7.0.1 Tools Required

To service the Sun StorEdge 5310 NAS, you need:

- Phillips screw driver
- Flat head screw driver

1.7.0.2 Electrostatic Discharge Information

Static electricity can cause damage to static-sensitive devices and/or microcircuitry. For this reason, it is important that proper packaging and grounding techniques be observed. To further ensure the prevention of electrostatic damage, observe these procedures:

- Transport products in static-safe containers.
- Cover work stations with approved static-dissipating material.
- Wear a wrist strap, and always be properly grounded when touching static-sensitive equipment/parts.
- Use only properly grounded tools and equipment.
- Avoid touching pins, lead or circuitry.

Note – The following section can be ignored if you are swapping out a fan, power supply or hard drive.

1.7.0.3 Preparation Procedures

Complete the following steps before you begin the removal/installation procedures:

1. **Shut the system down properly according to your operating system's instructions.**

2. **Turn the Sun StorEdge 5310 NAS off.**

3. **Disconnect the power cord from the power source, then from the Sun StorEdge 5310 NAS server.**

4. **Shut down the storage enclosure and remove its power cords.**

5. **Disconnect the all other external peripheral devices from the Sun StorEdge 5310 NAS server if applicable.**
6. **Disconnect all optical fibre and network interface cables from the Sun StorEdge 5310 NAS server and the storage enclosure.**
7. **Remove the Sun StorEdge 5310 NAS and the storage enclosure from the rack.**

1.8 Static Electricity Precautions

1.8.0.1 Grounding Procedure

You must maintain reliable grounding of this equipment. The Sun StorEdge 5310 NAS system (including head and optional Expansion Unit) must be connected to a dedicated 20A receptacle.

1.8.0.2 Static Electricity

The Sun StorEdge 5310 NAS server and Expansion Unit contain several components sensitive to static-electrical discharge. Surges of static electricity (caused by shuffling your feet across a floor and touching a metallic surface, for example) can cause damage to electrical components.

Static electricity can cause damage to static-sensitive devices and/or microcircuitry. For this reason, it is important that proper packaging and grounding techniques be observed. To further ensure the prevention of electrostatic damage, observe these procedures:

- Transport products in static-safe containers.
- Cover work stations with approved static-dissipating material.
- Wear a wrist strap, and always be properly grounded when touching static-sensitive equipment/parts.
- Use only properly grounded tools and equipment.
- Avoid touching pins, leads, or circuitry.

To avoid damaging Sun StorEdge 5310 NAS and Expansion Unit internal components with static electricity, follow these instructions before performing any installation procedures.

1. **Make sure both of the Sun StorEdge 5310 NAS (and optional Expansion Unit) AC power cables are plugged in, and that the unit is turned off.**

2. **Wear a wrist strap, and always be properly grounded when touching static-sensitive equipment/parts.**

   If a wrist strap is not available, touch any unpainted metal surface on the Sun StorEdge 5310 NAS (and optional Sun StorEdge 5310 NAS Expansion Unit) back panel to dissipate static electricity. Repeat this procedure several times during installation.

3. **Avoid touching exposed circuitry, and handle components by their edges only.**

Caution – Do not power on the Sun StorEdge 5310 NAS nor Sun StorEdge 5310 NAS Expansion Unit units until after you have connected to the Network.

The AC source must be electrically isolated by double or reinforced insulation from any hazardous AC or DC source. The AC source must be capable of providing up to 500 W of continuous power per feed pair.

Mains AC Power Disconnect—You are responsible for installing an AC power disconnect for the entire rack unit. This power source disconnect must be readily accessible, and it must be labeled as controlling power to the entire unit, not just to the server(s).

CHAPTER 2

NAS Head

This chapter addresses frequently asked questions for the Sun StorEdge 5310 NAS. The chapter contains these sections:

- "Hardware" on page 2-1
- "OS Operations" on page 2-36
- "Updating the OS on the Sun StorEdge 5310 NAS" on page 2-40
- "Sun StorEdge 5310 NAS Firmware" on page 2-40
- "Common Problems Encountered on the Sun StorEdge 5310 NAS" on page 2-42
- "Frequently Asked Questions" on page 2-92

2.1 Hardware

2.2 Contacting Technical Support

For technical support, call the phone numbers listed below, according to your location.

United States1-800-USA-4SUN (1-800-872-4786)

UK	Tel: +44 870-600-3222
France	Tel: +33 1 34 03 5080
Germany	Tel: +49 1805 20 2241
Italy	Tel: +39 02 92595228, Toll Free 800 605228

Spain	Tel: +011 3491 767 6000

See the following link for US, Europe, South America, Africa, and APAC local country telephone numbers:

`http://www.sun.com/service/contacting/solution.html`

For general support and documentation on the servers, see the following link:

`http://www.sun.com/supporttraining/`

2.2.1 Problems With Initial System Startup

Problems that occur at initial system startup are usually caused by incorrect installation or configuration. Hardware failure is a less frequent cause.

2.2.1.1 Checklist

- Are all cables correctly connected and secured?
- Is the power cord properly inserted and fully seated?
- Are there any Baseboard Management Controller (BMC) beep codes? You may have to listen carefully two or three times to hear them. See "POST Error Beep Codes" on page 2-27 for beep code details.
- Is the BMC running? Try pressing the ID button on the front panel. If the blue ID LED fails to illuminate, the BMC is not responding.
- Are the cables going to the front panel board installed and seated properly (check the front panel cable, the USB cable, and the 100-pin flex cable).
- Are the processors fully seated in their sockets on the server board?
- Are all add-in PCI boards fully seated in their slots on the server board?
- Are all jumper and switch settings on add-in boards and peripheral devices correct? To check these settings, refer to the manufacturer's documentation that comes with them. If applicable, ensure that there are no conflicts—for example, two add-in boards sharing the same interrupt.
- Are all DIMMs installed correctly?
- Are all peripheral devices installed correctly?
- If the system has a hard disk drive, is it properly formatted or configured?
- Are all device drivers properly installed?
- Are the configuration settings made in BIOS Setup correct?
- Did you press the system power on/off switch on the front panel to turn the server on (power on light should be lit)?
- Is the system power cord properly connected to the system and plugged into a NEMA 5-15R outlet for 100-120 V or a NEMA 6-15R outlet for 200-240V?
- Is AC power available at the wall outlet?
- Are there any POST LEDs illuminated? If so check "Power-On Self-Test (POST)" on page 2-7.

- Are there any POST beep codes? If so check “POST Error Beep Codes” on page 2-27.

2.2.2 Resetting the Server

Quite often, a problem can be solved merely be resetting the server or shutting it down and powering it back up. You may restart or shut down the Sun StorEdge 5310 NAS using software or hardware.

2.2.2.1 Shutdown Commands for Software Menu

To shutdown the system using the menu:

1. **Use the Web Administrator or Telnet to the Sun StorEdge 5310 NAS to shutdown the server.**
2. **Via Web Admin, go to Managing the System and choose Shutdown Server.**
3. **Via Telnet go to the main menu.**
4. **Press 0 for Server Shutdown.**

   This screen will give you the option of reboot or halt.
5. **Choose one of the options and the server will shut down.**

Note – There could be a few second delay before the server shuts down.

2.2.2.2 Shutdown Commands for Hardware LCD Display

To shutdown the system using the LCD display:

1. **Press the Select button on the LCD panel to access menus.**
2. **The LCD panel displays options A and B. Press the Down Arrow to select option “B. Shutdown Server” then press the Select button.**
3. **Press Select to select the “A. Power Off” option.**
4. **Press the Down Arrow to change “No” to “Yes”.**
5. **Press Select to confirm and begin shutting down.**

2.2.3 Preparing the System for Diagnostic Testing

Caution – Turn off devices before disconnecting cables. Before disconnecting any peripheral cables from the system, turn off the system and any external peripheral devices. Failure to do so can cause permanent damage to the system and/or the peripheral devices.

1. **Turn off the system and all external peripheral devices. Disconnect all of them from the system, except the keyboard and video monitor.**
2. **Make sure the system power cord is plugged into a properly grounded AC outlet.**
3. **Make sure your video display monitor and keyboard are correctly connected to the system. Turn on the video monitor. Set its brightness and contrast controls to at least two thirds of their maximum ranges (see the documentation supplied with your video display monitor).**
4. **Turn on the system. If the power LED does not light, see "Power LED Does Not Light" on page 2-8.**
5. **If errors are encountered, power off the system, remove all add-in cards, and turn the power back on.**

2.2.3.1 Specific Problems and Corrective Actions

This section provides possible solutions for the specific problems listed in Table 2-1.

TABLE 2-1 Index to Problems

Problems	Reference
"Problems Starting Up"	page 2-5
"Power LED Does Not Light"	page 2-8
"System Cooling Fans Do Not Rotate Properly"	page 2-8
"Cannot Connect to a Server"	page 2-9
"Problems with Network"	page 2-9

Try the solutions in the order given.

Problems Starting Up

If the server does not start up properly, use the information in this section to diagnose problems.

Server Does Not Power On

If the server does not power on, check the following:

- Does the main server board have power? Open the chassis lid and check the 5V Standby LED on the baseboard to see if it is illuminated. If your server is plugged in, this LED should be green. See Figure 2-5, "Fault and Status LEDs on the Server Board," on page 2-21 for the location of this LED.
- Check the power cord connection. The Sun StorEdge 5310 NAS allows the use of two power supplies, and the system will not power on if one power cord is used and it is plugged into the wrong power connector.
- Remove all add-in cards and see if the server boots using just the on-board components. If the server boots successfully, add the cards back in one at a time with a reboot after each addition to see if you can isolate a suspect card.
- Remove and reseat the memory modules. Ensure that you have properly populated the memory modules. On the main board, memory is populated in pairs. See "Memory" on page 7-6 for memory module installation and placement. Refer to the silkscreen on the main board for proper memory module placement. Try using memory modules from a known, compatible, server.
- Check the internal cable connections to ensure that they are properly connected.
- Remove the processor(s) and reseat as a last resort.

Caution – Removing and replacing the processors is not recommended and should only be done as a last resort. This is a procedure that should be attempted by Sun qualified service personnel.

Front Panel is Unresponsive and Video is Disabled

If the front panel is unresponsive to any pushbuttons you press, and video is disabled, it could be that the front panel is locked. By default, front panel locking is disabled; however, it is possible to enable front panel locking through the BIOS setup. To do this, an administrative password must be set using Security > Set Admin Password.

When the password is set, the front panel, mouse, and keyboard are locked after a timeout expires. The video is also blanked. The purpose of this is to prevent unauthorized access to a server by someone who plugs in a keyboard and video monitor. Access is regained simply by using the keyboard to type the password.

Note – A corded PS/2 keyboard (not a wireless one) must be plugged into the keyboard/mouse connector at the back of the server. When the front panel is locked, the lights on the keyboard flash, but the server is still fully functional.

Server Beeps at Power On or When Booting

The server indicates problems with "beep codes" during Power-On Self Test (POST) in the event there is no displayed video. A complete list of beep codes is given in "POST Error Beep Codes" on page 2-27.

Note – The RAID card also will beep when a disk drive has failed. Check the system log to help isolate the problem.

The following beep codes identify system events during POST in case video fails to display.

TABLE 2-2 Bootup Beep Codes

Beeps	Reason
1	One short beep before boot (normal, not an error)
1-2	Search for option ROMs. One long beep and two short beeps on checksum failure.
1-2-2-3	BIOS ROM checksum
1-3-1-1	Test DRAM refresh
1-3-1-3	Test 8742 keyboard controller
1-3-3-1	Auto size DRAM. System BIOS stops execution here if the BIOS does not detect any usable memory DIMMs.
1-3-4-1	Base RAM failure. BIOS stops execution here if entire memory is bad.
2-1-2-3	Check ROM copyright notice.
2-2-3-1	Test for unexpected interrupts.
1-5-1-1	FRB failure (processor failure)
1-5-2-2	No processors installed
1-5-2-3	Processor configuration error (for example, mismatched VIDs).
1-5-2-4	Front-side bus select configuration error (for example, mismatched BSELs)

TABLE 2-2 Bootup Beep Codes

Beeps	Reason
1-5-4-2	Power fault
1-5-4-3	Chipset control failure
1-5-4-4	Power control failure

Server Starts Booting Automatically at Power On

The server board saves the last known power state in the event of a power failure. If you remove power before powering down the system using the power switch on the front panel, your system might automatically attempt to restore itself back to the state it was in after you restore power.

You can configure how you would like your server system to react when power is restored in the BIOS set-up (Security menu). You can have the server remain off or return to the last known power state.

- Please keep in mind that unplugging the system or flipping a switch on the power strip both remove power.
- Follow the correct power removal sequence (make sure the system has shut down before removing the power cord).

Power-On Self-Test (POST)

Each time you turn on the system, the BIOS begins execution of POST. POST discovers, configures, and tests the processors, memory, keyboard, and most installed peripheral devices. The time needed to test memory depends on the amount of memory installed. POST is stored in flash memory.

To execute and monitor POST:

1. **Turn on your video monitor and system. After a few seconds, POST begins to run and displays a splash screen.**

2. **While the splash screen is displayed:**

   - Press <F2> to enter the BIOS Setup

   OR

   - Press <Esc> to view POST diagnostic messages and change the boot device priority for this boot only.

   OR

   - If the Service Partition is installed, press <F4> to run the System Setup Utility

3. **If you do not press <F2> or <Esc> or <F4> and do NOT have a device with an operating system loaded, the boot process continues and the system beeps once. The following message is displayed:**

   ```
   Operating System not found
   ```

4. **At this time, pressing any key causes the system to attempt a reboot. The system searches all removable devices in the order defined by the boot priority.**

During POST, the server BIOS presents screen messages to indicate error conditions. POST also provides beep codes to give you audible clues regarding the performance and operation of the server when there is no video display that can present error messages. In addition, a set of four bi-color diagnostic LEDs is located on the back edge of the server main board. These LEDs are active during POST and indicate the state of the server. Each of the four LEDs can have one of four states: Off, Green, Red, or Amber. See "Power-On Self-Test (POST)" on page 2-7 for a complete description of the screen messages, beep codes, and diagnostic LEDs.

Verifying Proper Operation of Key System LEDs

As POST determines the system configuration, it tests for the presence of each mass storage device installed in the system. As each device is checked, its activity light should turn on briefly. Check to see if the disk drive activity light for each drive turns on briefly.

2.2.3.2 Power LED Does Not Light

Check the following:

- Is the system operating normally? If so, the power LED is probably defective or the cable from the front panel to the server board is loose.
- Are there other problems with the system? If so, check the items listed under "System Cooling Fans Do Not Rotate Properly" on page 2-8.

If all items are correct and problems persist, contact your service representative or authorized dealer for help.

2.2.3.3 System Cooling Fans Do Not Rotate Properly

If the system cooling fans are not operating properly, system components could be damaged.

Check the following:

- Is AC power available at the wall outlet?
- Is the system power cord properly connected to the system and the wall outlet?

- Did you press the power button?
- Is the power on light illuminated?
- Have any of the fan motors stopped (use the server management subsystem to check the fan status)?
- Are the fan power connectors properly connected to the server board?
- Is the cable from the front panel board connected to the server board?
- Are the power supply cables properly connected to the server board?
- Are there any shorted wires caused by pinched cables or power connector plugs forced into power connector sockets the wrong way?

If the switches and connections are correct and AC power is available at the wall outlet, contact your service representative or authorized dealer for help.

2.2.3.4 Cannot Connect to a Server

Check the following:

- Make sure the network cable is securely attached to the connector at the system back panel. If the cable is attached but the problem persists, try a different cable.
- Make sure the hub port is configured for the same duplex mode as the network controller.
- If you are directly connecting two servers (no hub), you will need a crossover cable (see your hub documentation for more information on crossover cables).
- Check the network controller LEDs that are visible through an opening at the system back panel.

2.2.3.5 Problems with Network

If diagnostics pass, but the connection fails:

- Make sure the network cable is securely attached.

The Activity LED does not light:

- Make sure the network hub has power.

If the controller stopped working when an add-in adapter was installed:

- Make sure the cable is connected to the port from the onboard network controller.
- Try reseating the add in adapter.

If the add-in adapter stopped working without apparent cause:

- Try reseating the adapter first; then try a different slot if necessary.

2.2.3.6 Other Problems

If the preceding information does not fix the problem with your server, try the following:

- Check for proper processor installation. Systems with a single processor must have the CPU installed in CPU socket 1. If two processors are installed, the processors must be of the same speed and voltage (and within one stepping). Do not attempt to over clock the processors or other components on this system. Over clocking is generally not possible and may damage components and void the warranty of your server board and your boxed or tray processor.
- Memory must be of the approved type and be properly seated.
- Verify that all chassis and power supply fans are properly installed and functioning.
- Approved heat sinks must be properly installed on the processors. Do not attempt to run the processors without a heat sink for even a few moments.

2.3 Troubleshooting the Server Using Built-In Tools

This chapter explains how to detect and isolate faulty components within the Sun StorEdge 5310 NAS. The chapter contains these sections:

- "CIFS/SMB/Domain Issues" on page 2-92
- "LEDs and Pushbuttons" on page 2-11
- "Power-On Self Test (POST)" on page 2-24
- "Contacting Technical Support" on page 2-1

2.4 Diagnosing System Errors

Use the following tools to help you isolate server problems:

- "LEDs" on page 2-11
- "Beep Codes" on page 2-11
- "POST Screen Messages" on page 2-11

2.4.1 LEDs

You can use the diagnostic LED indications to isolate faults. See "LEDs and Pushbuttons" on page 2-11.

2.4.2 Beep Codes

A built-in server speaker indicates failures with audible beeps. See "POST Error Beep Codes" on page 2-27.

2.4.3 POST Screen Messages

For many failures, the BIOS sends error codes and message to the screen. See "POST Screen Messages" on page 2-24

2.5 LEDs and Pushbuttons

Note – This section addresses LEDs and Pushbuttons on the Sun StorEdge 5310 NAS. The LEDs on the Sun StorEdge 5210 Expansion Unit are different.

This section describes the LEDs and pushbuttons on the Sun StorEdge 5310 NAS.

TABLE 2-3 Server LEDs

LED Name	Function	Location	Color	Status
ID	Helps identify the server from the front or rear	One LED on front panel and one at rear corner	Blue	On = ID
System status	Visible fault indicator	One LED on front panel and one at rear corner	Green or amber	Off = POST in progress or system stop Green steady on = no fault Green blinking = degraded Amber steady = critical or non-recoverable state Amber blinking = non-critical state
Disk activity	Indicates hard disk activity	Front panel and main board left side	Green	Blinking = HDD activity

TABLE 2-3 Server LEDs

LED Name	Function	Location	Color	Status
Memory DIMM fault (1 - 6)	Identifies failing DIMM module	At the front of each DIMM location on main board	Amber	On = fault
POST LEDs (1 - 4)	Displays boot 80 POST codes	Left rear of main board	Each LED can be off, green, red, or amber	See "POST Progress Code LED Indicators" on page 2-30 for POST code LED details.
Fan fault (1 - 4)	Identifies Sun StorEdge 5310 NAS fan failure	On Sun StorEdge 5310 NAS fan module board	Amber	On = fault
CPU 1 and 2 fault	Identify CPU failure	Back corner of processor socket on main board	Amber	On = fault
5V standby	Identify 5V standby power on state	Front left on main board	Green	Green = 5V standby power on

Main power LED	Identifies power state of the server	Front panel	Green	Off = power is off
				On = power is on

2.5.1 Front Panel LEDs and Pushbuttons

The front panel contains the pushbuttons and LEDs shown in Figure 2-1. Note that the illustration has the bezel removed.

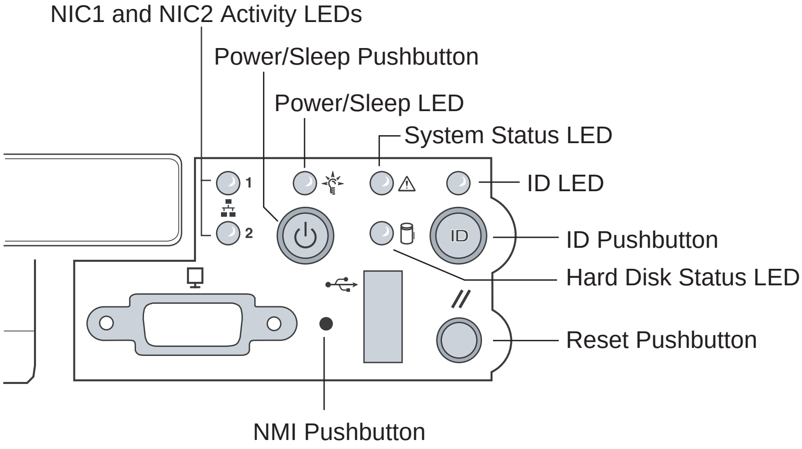

FIGURE 2-1 Front Panel Pushbuttons and LEDs

2.5.1.1 Front Panel LEDs

The front panel LEDs are summarized in Table 2-4.

TABLE 2-4 Front Panel LEDs

LED	Color	Function
Power	Green	This LED is controlled by software. It turns steady when the server is powered up and is off when the system is off or in sleep mode.
NIC1 and NIC2	Green	These LEDs are on when a good network link has been established. They blink green to reflect network data activity.

TABLE 2-4 Front Panel LEDs

LED	Color	Function
System Status/Fault	Green/ Amber	This LED can assume different states (green, amber, steady, blinking) to indicate critical, non-critical, or degraded server operation. Steady green: Indicates the system is operating normally Blinking green: Indicates the system is operating in a degraded condition. Blinking amber: Indicates the system is in a non-critical condition. Steady amber: Indicates the system is in a critical or non-recoverable condition. Off: Indicates POST/system stop. See "Front-Panel System Status LED" on page 2-18 for more details regarding this LED.
Hard Disk Drive Activity	Green	The Drive Activity LED on the front panel is used to indicate drive activity from the onboard SCSI controller. The server Main Board also provides a header, giving access to this LED for add-in IDE or SCSI controllers. Blinking green (random): Hard disk activity Steady amber: Hard disk fault Off: No disk activity nor fault condition (or power is off).
System ID	Blue	The blue System Identification LED is used to help identify a system for servicing when it is installed within a high density rack or cabinet that is populated with several other similar systems. The System ID LED is illuminated when the system ID button, located on the front panel, is pressed. If activated by the front panel pushbutton, the LED remains on until the pushbutton is depressed again. The LED also illuminates when the server receives a remote System Identify command from a remote management console. In this case, the LED turns off after a timeout period. The timeout period is configurable, with a default of 15 seconds. An additional blue System ID LED on the Main Board is visible through the rear panel. It mirrors the operation of the front panel LED.

2.5.1.2 Front Panel Pushbuttons

The front panel pushbuttons are summarized in Table 2-5.

TABLE 2-5 Front Panel Pushbuttons

Switch	Function
Power/Sleep	This pushbutton is used to toggle the system power on and off. This button is also used as a sleep button for operating systems that follow the ACPI specification. Linux, for example, configures the power button to the instant off mode. There is no ACPI support for the Solaris OS.
Reset	Depressing this pushbutton reboots and initializes the system.
NMI	Pushing this recessed pushbutton causes a non-maskable interrupt to occur. **Note: NMI is not currently supported.**
System ID	This pushbutton toggles the state of the front panel ID LED and the server Main Board ID LED. The Main Board ID LED is visible through the rear of the chassis and allows you to locate a particular server from behind a rack of servers.

2.5.2 Rear Panel LEDs

The rear panel contains the LEDs shown in Figure 2-2.

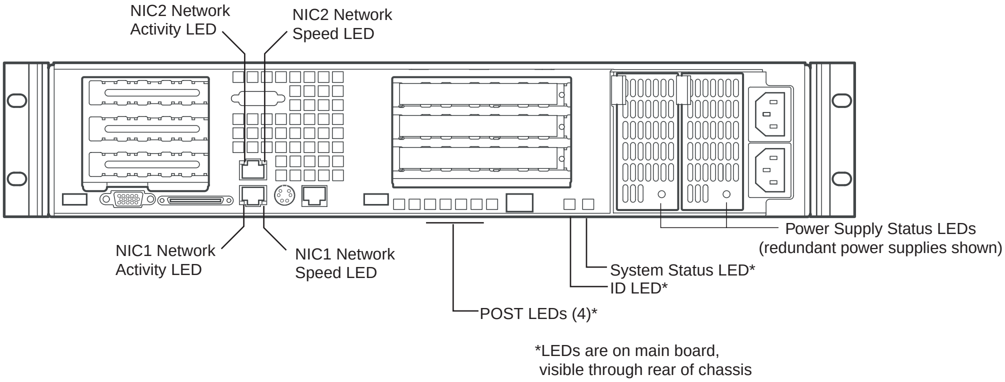
NIC2 Network Activity LED
NIC2 Network Speed LED
NIC1 Network Activity LED
NIC1 Network Speed LED
POST LEDs (4)*
System Status LED*
ID LED*
Power Supply Status LEDs (redundant power supplies shown)
*LEDs are on main board, visible through rear of chassis

FIGURE 2-2 Rear Panel LEDs

TABLE 2-6 Rear Panel LEDs

LED	Color	Function
Network Connection/ Network Activity	Green	This LED is on the left side of each NIC connector. Green = valid network connection. Blinking = transmit or receive activity.

Network Speed	Amber/Green	This LED is on the right side of the NIC connector.
		Off = 10 Mbps operation.
		Green = 100 Mbps operation.
		Amber = 1000 Mbps operation.

POST LEDs (four)	Multicolor (Red/Green/Amber)	To help diagnose power-on self test (POST) failures, a set of four bi-color diagnostic LEDs is located on the back edge of the server Main Board. These LEDs are visible through holes in the rear panel. Each of the four LEDs can have one of four states: Off, Green, Red, or Amber. For detailed information on these LEDs, see “POST Progress Code LED Indicators” on page 2-30.

TABLE 2-6 Rear Panel LEDs

LED	Color	Function
System ID	Blue	This LED is located on the Main Board and is visible through holes in the rear panel. It can provide a mechanism for identifying one system out of a group of identical systems. This can be particularly useful if the server is used in a rack-mount chassis in a high-density, multiple-system application. The LED is activated by depressing the front panel System ID pushbutton or if the server receives a remote System Identify command from a remote management console. If activated by the front panel pushbutton, the LED remains on until the pushbutton is depressed again. When the LED illuminates due to a remote System Identify command, the LED turns off after a timeout period. An additional blue System ID LED is located on the front panel that mirrors the operation of the rear Main Board LED.
System Status/Fault	Green/Amber	This LED reflects the state of the System Status LED on the front panel.
Power Supply	Green/Amber	This is a bi-color LED that can be on, off, green, amber, or blinking, or combination thereof. See “Rear Panel Power Supply Status LED” on page 2-20 for more detailed information.

2.5.3 Front-Panel System Status LED

The front-panel system status LED is located as shown in Figure 2-3.

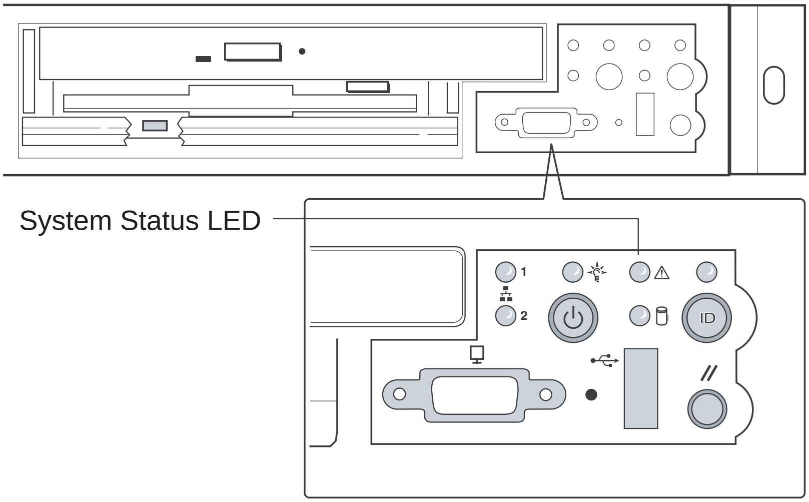

FIGURE 2-3 Location of Front-Panel System Status LED

The front-panel system status LED has the states indicated in Table 2-7.

TABLE 2-7 System Status LED States

System Status LED State	System Condition
CONTINUOUS GREEN	Indicates the system is operating normally.
BLINKING GREEN	Indicates the system is operating in a degraded condition.
BLINKING AMBER	Indicates the system is in a non-critical condition.
CONTINUOUS AMBER	Indicates the system is in a critical or non-recoverable condition.
OFF	Indicates POST/system stop.

Critical Condition

A critical condition or non-recoverable threshold crossing is indicated with a continuous amber status LED and is associated with the following events:

- Temperature, voltage, or fan critical threshold crossing.

- Power subsystem failure. The Baseboard[1] Management Controller (BMC) asserts this failure whenever it detects a power control fault (for example, the BMC detects that the system power is remaining on even though the BMC has deasserted the signal to turn off power to the system).
- The system is unable to power up due to incorrectly installed processor(s), or processor incompatibility.
- A satellite controller such as the HSC, or another IMPI-capable device, such as an add-in server management PCI card, sends a critical or non-recoverable state, via the Set Fault Indication command to the BMC.
- Critical Event Logging errors, including System Memory Uncorrectable ECC error and Fatal/Uncorrectable Bus errors, such as PCI SERR and PERR.

Non-Critical Condition

A non-critical condition is indicated with a blinking amber status LED and signifies that at least one of the following conditions is present:

- Temperature, voltage, or fan non-critical threshold crossing.
- Chassis intrusion.
- Satellite controller sends a non-critical state, via the Set Fault Indication command, to the BMC.
- A Set Fault Indication command from the system BIOS. The BIOS may use the Set Fault Indication command to indicate additional, non-critical status such as system memory or CPU configuration changes.

Degraded Condition

A degraded condition is indicated with a blinking green status LED and signifies that at least one of the following conditions is present:

- Non-redundant power supply operation. This only applies when the BMC is configured for a redundant power subsystem. The power unit configuration is configured via OEM SDR records.
- A processor is disabled by FRB or BIOS.
- BIOS has disabled or mapped out some of the system memory.

This Troubleshooting Guide gives information on how to isolate the server component responsible for any of the critical, non-critical, or degraded conditions listed above.

[1] Baseboard refers to the server Main Board.

2.5.4 Rear Panel Power Supply Status LED

The rear-panel power supply status LEDs are located as shown in Figure 2-4.

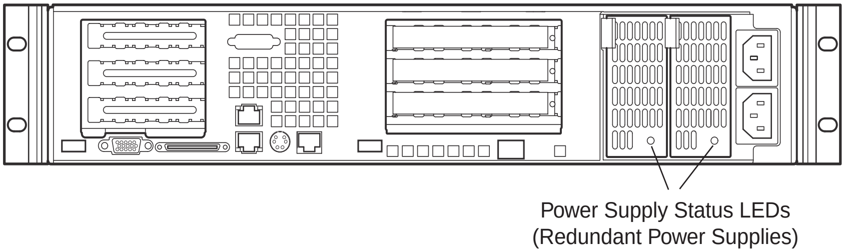

FIGURE 2-4 Location of Rear-Panel Power Supply Status LEDs

The rear-panel power supply status LED has the states indicated in Table 2-8.

TABLE 2-8 Power Supply Status LED States

Power Supply LED State	Power Supply Condition
OFF	No AC power present to power supply
BLINKING GREEN	AC power present, but only the standby outputs are on
GREEN	Power supply DC outputs are on and OK
BLINKING AMBER	PSAlert# signal asserted, power supply on
AMBER	Power supply shutdown due to over current, over temperature, over voltage, or undervoltage
AMBER or OFF	Power supply failed and AC fuse open or other critical failure

Note – If redundant power supplies are used in the Sun StorEdge 5310 NAS, the power supply LEDs have the following meaning:

Both LEDs off = no power to power supplies or both power supplies bad.
Both LEDs blinking green = power supplies receiving AC power, but server is off.
Both LEDs solid green = server is fully powered on and power supplies are good.
One LED solid green and one LED amber = AC power missing from one of the power supplies.

2.5.5 Server Main Board Fault LEDs

There are several fault and status LEDs built into the server board (see Figure 2-5). Some of these LEDs are visible only when the chassis cover is removed. The LEDs are explained in this section.

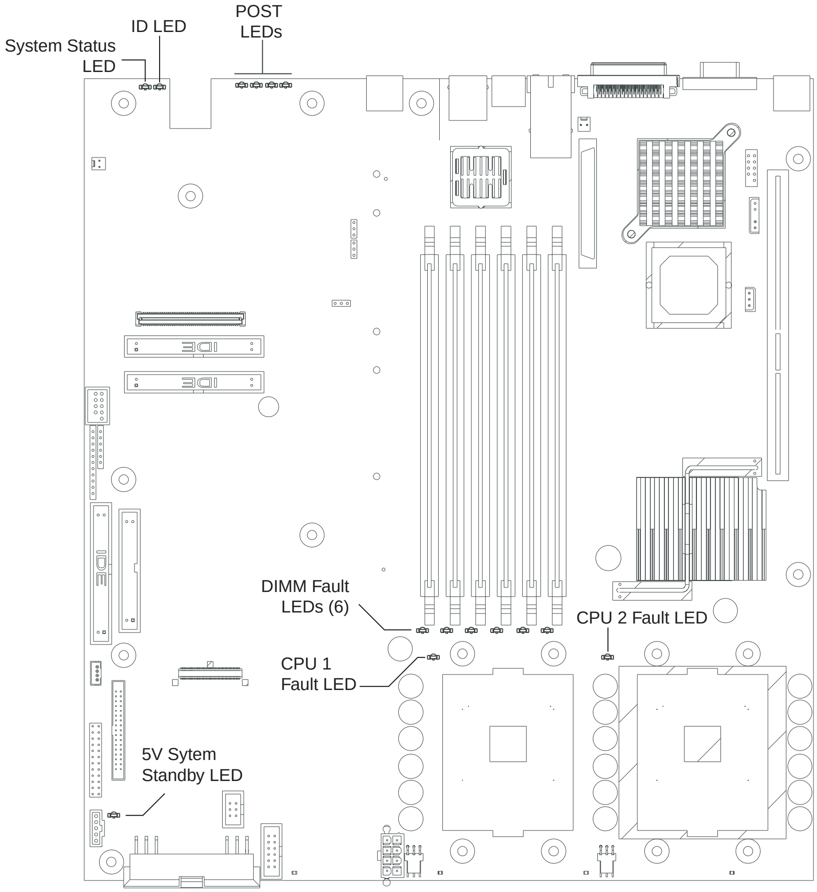

FIGURE 2-5 Fault and Status LEDs on the Server Board

The fault LEDs are summarized below.

- **POST LEDs:** To help diagnose POST failures, a set of four bi-color diagnostic LEDs is located on the back edge of the baseboard. Each of the four LEDs can have one of four states:

  Off, Green, Red, or Amber. During the POST process, each light sequence represents a specific Port-80 POST code. If a system should hang during POST, the diagnostic LEDs present the last test executed before the hang. When reading the lights, the LEDs should be observed from the back of the system. The most significant bit (MSB) is the first LED on the left, and the least significant bit (LSB) is the last LED on the right.

  See "POST Progress Code LED Indicators" on page 2-30 for details regarding the POST LED display.

- **CPU Fault LEDs:** A fault indicator LED is located next to each of the processor sockets. If the server Baseboard Management Controller (BMC) detects a fault in any processor, the corresponding LED illuminates.
- **Memory Fault LEDs:** A fault indicator LED is located next to each of the DIMM sockets. If the BMC detects a fault in a given DIMM, the corresponding LED illuminates.

  One LED for each DIMM is illuminated if that DIMM has an uncorrectable or multi-bit memory error. The LEDs maintain the same state across power switch, power down, or loss of AC power.

- **Fan Fault LEDs:** Depending on the server model, the fan header may include a fan fault LED. If the BMC detects a fan fault, the LED illuminates. If the fan fault LED is lit, the entire fan module must be replaced.
- **System Status LED:** Indicates functional status of the server board. Glows green when all systems are operating normally. Glows amber when one or more systems are in a fault status. This LED mirrors the function of the system status LED on the front panel.

  See Table 2-7 on page 2-18 for a description of the LED states.

- **+5V Standby LED**. This green LED is on when the server is plugged into AC power, whether or not the server is actually powered on. AC power is applied to the system as soon as the AC cord is plugged into the power supply.
- **System ID LED**. This blue LED can be illuminated to identify the server when it is part of a large stack of servers. See "System ID LEDs" on page 2-23 for details.

2.5.6 System ID LEDs

A pair of blue LEDs, one at the rear of the server, and one on the front panel, can be used to easily identify the server when it is part of a large stack of servers. A single blue LED located at the back edge of the server board next to the backup battery is visible through the rear panel. The two LEDs mirror each other and can be illuminated by the Baseboard Management Controller (BMC) either by pressing a button on the chassis front panel or through server-management software. When the button is pressed on the front panel, both LEDs illuminate and stay illuminated until the button is pressed again. If the LED is illuminated through a remote System Identify command, the LED turns off after a timeout period. See Figure 2-5 on page 2-21 for the location of the rear Main Board LED. The front panel ID LED and the ID activation button are shown in Figure 2-6.

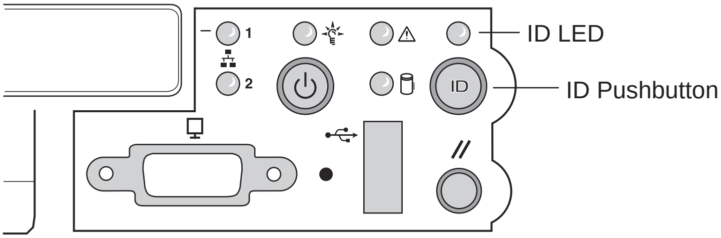

FIGURE 2-6 Location of Front-Panel ID Pushbutton and LED

2.6 Power-On Self Test (POST)

The BIOS indicates the current testing phase during POST by writing a hex code to the Enhanced Diagnostic LEDs, located on the rear of the server main board and visible through the back of the chassis.

If errors are encountered, error messages or codes will either be displayed to the video screen, or if an error has occurred prior to video initialization, errors will be reported through a series of audible beep codes. POST errors are logged in to the System Event Log (SEL).

During the power-on self test (POST), the server may indicate a system fault by:

- Displaying error codes and messages at the display screen
- Beeping the speaker in a coded sequence
- Illuminating the POST LEDs, visible from the rear panel, in a coded fashion

2.6.1 POST Screen Messages

During POST, if an error is detected, the BIOS displays an error code and message to the screen. The tables in this section describe the standard and extended POST error codes and their associated messages. The BIOS prompts the user to press a key in case of serious errors. Some of the error messages are preceded by the string "Error" to highlight the fact that the system may be malfunctioning. All POST errors and warnings are logged in the System Event Log (SEL) unless it is full.

Note – All POST errors are logged to the SEL, which is capable of holding approximately 3200 entries. After the SEL is full, no further errors are logged. The SEL can be cleared using the SSU or the BIOS setup. The SEL is automatically cleared after running the PCT.

Table 2-9 and Table 2-10 contain the POST error messages and error codes.

TABLE 2-9 Standard POST Error Messages and Codes

Error Code	Error Message	Pause On Boot
100	Timer Channel 2 error	Yes
101	Master Interrupt Controller	Yes
102	Slave Interrupt Controller	Yes
103	CMOS battery failure	Yes

TABLE 2-9 Standard POST Error Messages and Codes *(Continued)*

Error Code	Error Message	Pause On Boot
104	CMOS options not set	Yes
105	CMOS checksum failure	Yes
106	CMOS display error	Yes
107	Insert key pressed	Yes
108	Keyboard locked message	Yes
109	Keyboard stuck key	Yes
10A	Keyboard interface error	Yes
10B	System memory size error	Yes
10E	External cache failure	Yes
113	Hard disk 0 error	Yes
114	Hard disk 1 error	Yes
115	Hard disk 2 error	Yes
116	Hard disk 3 error	Yes
11B	Date/time not set	Yes
11E	Cache memory bad	Yes
120	CMOS clear	Yes
121	Password clear	Yes
140	PCI error	Yes
141	PCI memory allocation error	Yes
142	PCI IO allocation error	Yes
143	PCI IRQ allocation error	Yes
144	Shadow of PCI ROM failed	Yes
145	PCI ROM not found	Yes
146	Insufficient memory to shadow PCI ROM	Yes

TABLE 2-10 Extended POST Error Messages and Codes

Error Code	Error Message	Pause On Boot
8100	Processor 1 failed BIST	No
8101	Processor 2 failed BIST	No
8110	Processor 1 internal error (IERR)	No
8111	Processor 2 internal error (IERR)	No
8120	Processor 1 thermal trip error	No
8121	Processor 2 thermal trip error	No
8130	Processor 1 disabled	No
8131	Processor 2 disabled	No
8140	Processor 1 failed FRB-3 timer	No
8141	Processor 2 failed FRB-3 timer	No
8150	Processor 1 failed initialization on last boot.	No
8151	Processor 2 failed initialization on last boot.	No
8160	Processor 01: unable to apply BIOS update	Yes
8161	Processor 02: unable to apply BIOS update	Yes
8170	Processor P1 :L2 cache failed	Yes
8171	Processor P2 :L2 cache failed	Yes
8180	BIOS does not support current stepping for Processor P1	Yes
8181	BIOS does not support current stepping for Processor P2	Yes
8190	Watchdog timer failed on last boot	No
8191	4:1 core to bus ratio: processor cache disabled	Yes
8192	L2 Cache size mismatch	Yes
8193	CPUID, processor stepping are different	Yes
8194	CPUID, processor family are different	Yes
8195	Front side bus speed mismatch: System halted	Yes, Halt
8196	Processor models are different	Yes
8197	CPU speed mismatch	Yes
8198	Failed to load processor microcode	Yes
8300	Baseboard Management Controller (BMC) failed to function	Yes
8301	Front panel controller failed to function	Yes

TABLE 2-10 Extended POST Error Messages and Codes *(Continued)*

Error Code	Error Message	Pause On Boot
8305	Hotswap controller failed to function	Yes
8420	Intelligent System Monitoring chassis opened	Yes
84F1	Intelligent System Monitoring forced shutdown	Yes
84F2	Server Management Interface failed	Yes
84F3	BMC in update mode	Yes
84F4	Sensor Data Record (SDR) empty	Yes
84FF	System event log full	No
8500	Bad or missing memory in slot 3A	Yes
8501	Bad or missing memory in slot 2A	Yes
8502	Bad or missing memory in slot 1A	Yes
8504	Bad or missing memory in slot 3B	Yes
8505	Bad or missing memory in slot 2B	Yes
8506	Bad or missing memory in slot 1B	Yes
8601	All memory marked as fail: forcing minimum back online	Yes

2.6.2 POST Error Beep Codes

The tables in this section list the POST error beep codes. Prior to system video initialization, the BIOS and BMC use these beep codes to notify users of error conditions.

TABLE 2-11 BMC-Generated POST Beep Codes

Beep Code[1]	Description
1	One short beep before boot (normal, not an error)
1-2	Search for option ROMs. One long beep and two short beeps on checksum failure.
1-2-2-3	BIOS ROM checksum
1-3-1-1	Test DRAM refresh
1-3-1-3	Test 8742 keyboard controller
1-3-3-1	Auto size DRAM. System BIOS stops execution here if the BIOS does not detect any usable memory DIMMs.
1-3-4-1	Base RAM failure. BIOS stops execution here if entire memory is bad.

TABLE 2-11 BMC-Generated POST Beep Codes

Beep Code[1]	Description
2-1-2-3	Check ROM copyright notice.
2-2-3-1	Test for unexpected interrupts.
1-5-1-1	FRB failure (processor failure)
1-5-2-2	No processors installed or processor socket 1 is empty
1-5-2-3	Processor configuration error (for example, mismatched VIDs)
1-5-2-4	Front-side bus select configuration error (for example, mismatched BSELs)
1-5-4-2	Power fault: DC power unexpectedly lost (for example, power good from the power supply was deasserted)
1-5-4-3	Chipset control failure
1-5-4-4	Power control failure (for example, power good from the power supply did not respond to power request)

[1] The code indicates the beep sequence; for example, 1-5-1-1 means a single beep, then a pause, then 5 beeps in a row, then a pause, then a single beep, then a pause, and then finally a single beep.

TABLE 2-12 BIOS-Generated Boot Block POST Beep Codes

Beep Code	Error Message	Description
1	Refresh timer failure	The memory refresh circuitry on the motherboard is faulty.
2	Parity error	Parity can not be reset
3	Base memory failure	Base memory test failure. See Table 2-13 on page 2-29 for additional error details.
4	System timer	System timer is not operational
5	Processor failure	Processor failure detected
6	Keyboard controller Gate A20 failure	The keyboard controller may be bad. The BIOS cannot switch to protected mode.
7	Processor exception interrupt error	The CPU generated an exception interrupt.
8	Display memory read/write error	The system video adapter is either missing or its memory is faulty. This is not a fatal error.
9	ROM checksum error	System BIOS ROM checksum error
10	Shutdown register error	Shutdown CMOS register read/write error detected
11	Invalid BIOS	General BIOS ROM error

TABLE 2-13 Memory 3-Beep and LED POST Error Codes

Beep Code	Debug Port 80h Error Indicator	Diagnostic LED Decoder (G = green, R = red, A = amber)				Meaning
		MSB			LSB	
3	00h	Off	Off	Off	Off	No memory was found in the system
3	01h	Off	Off	Off	G	Memory mixed type detected
3	02h	Off	Off	G	Off	EDO is not supported
3	03h	Off	Off	G	G	First row memory test failure
3	04h	Off	G	Off	Off	Mismatched DIMMs in a row
3	05h	Off	G	Off	G	Base memory test failure
3	06h	Off	G	G	Off	Failure on decompressing post module
3	07h	Off	G	G	G	Generic memory error
	08h	G	Off	Off	Off	
	09h	G	Off	Off	G	
	0Ah	G	Off	G	Off	
	0Bh	G	Off	G	G	
	0Ch	G	G	Off	Off	
	0Dh	G	G	Off	G	
3	0Eh	G	G	G	Off	SMBUS protocol error
3	0Fh	G	G	G	G	Generic memory error

2.6.2.1 BIOS Recovery Beep Codes

In rare cases, when the system BIOS has been corrupted, a BIOS recovery process must be followed to restore system operability. During recovery mode, the video controller is not initialized. One high-pitched beep announces the start of the recovery process. The entire process takes two to four minutes. A successful update ends with two high-pitched beeps. In the event of a failure, two short beeps are generated and a flash code sequence of 0E9h, 0EAh, 0EBh, 0ECh, and 0EFh appears at the Port 80 diagnostic LEDs (see Table 2-14 on page 2-30).

TABLE 2-14 BIOS Recovery Beep Codes

Beep Code	Error Message	Port 80h LED Indicators	Description
1	Recovery started		Start recovery process.
Series of long low-pitched single beeps	Recovery failed	EEh	Unable to process valid BIOS recovery images. BIOS already passed control to OS and flash utility.
Two long high pitched beeps	Recovery complete	EFh	BIOS recovery succeeded, ready for powerdown, reboot.

2.6.3 POST Progress Code LED Indicators

To help diagnose POST failures, a set of four bi-color diagnostic LEDs is located on the back edge of the server main board. Each of the four LEDs can have one of four states: Off, Green, Red, or Amber.

The LED diagnostics feature consists of a hardware decoder and four dual color LEDs. During boot block POST and post boot block POST, the LEDs display all normal Port 80 codes representing the progress of the BIOS POST. Each POST code is represented by a combination of colors from the four LEDs. The LEDs are in pairs of green and red. The POST codes are broken into two nibbles, an upper and a lower nibble. Each bit in the upper nibble is represented by a red LED and each bit in the lower nibble is represented by a green LED. If both bits are set in the upper and lower nibble then both red and green LEDs are lit, resulting in an amber color. Likewise, if both bits are clear, the red and green LEDs are off.

Figure 2-7 shows examples of how the POST LEDs are coded.

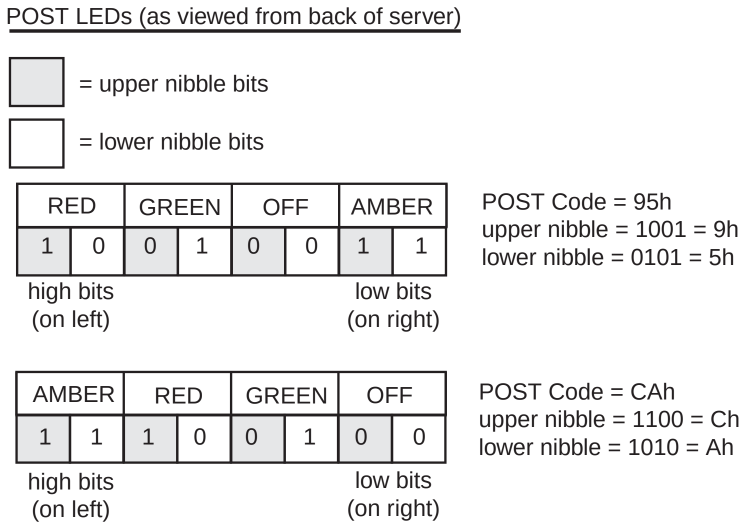

FIGURE 2-7 Examples of POST LED Coding

During the POST process, each light sequence represents a specific Port-80 POST code. If a system should hang during POST, the diagnostic LEDs present the last test executed before the hang. When you read the LEDs, observe them from the back of the system. The most significant bit (MSB) is the leftmost LED, and the least significant bit (LSB) is the rightmost LED.

Note – When comparing a diagnostic LED color sequence from the server Main Board to those listed in the diagnostic LED decoder in the following tables, the LEDs on the Main Board should be referenced when viewed by looking into the system from the back. Reading the LEDs from left to right, the most-significant bit is located on the left.

TABLE 2-15 Boot Block POST Progress LED Code Table (Port 80h Codes)

POST Code	Diagnostic LED Decoder (G = green, R = red, A = amber)				Description
	MSB			LSB	
	MSB			LSB	
10h	Off	Off	Off	R	The NMI is disabled. Start power-on delay. Initialization code checksum verified.
11h	Off	Off	Off	A	Initialize the DMA controller, perform the keyboard controller BAT test, start memory refresh, and enter 4 GB flat mode.

TABLE 2-15 Boot Block POST Progress LED Code Table (Port 80h Codes) *(Continued)*

POST	Diagnostic LED Decoder				
Code	(G = green, R = red, A = amber)				Description
12h	Off	Off	G	R	Get start of initialization code and check BIOS header.
13h	Off	Off	G	A	Memory sizing.
14h	Off	G	Off	R	Test base 512K of memory. Return to real mode. Execute any OEM patches and set up the stack.
15h	Off	G	Off	A	Pass control to the uncompressed code in shadow RAM. The initialization code is copied to segment 0 and control will be transferred to segment 0.
16h	Off	G	G	R	Control is in segment 0. Verify the system BIOS checksum. If the system BIOS checksum is bad, go to checkpoint code E0h; otherwise, going to checkpoint code D7h.
17h	Off	G	G	A	Pass control to the interface module.
18h	G	Off	Off	R	Decompression of the main system BIOS failed.
19h	G	Off	Off	A	Build the BIOS stack. Disable USB controller. Disable cache.
1Ah	G	Off	G	R	Uncompress the POST code module. Pass control to the POST code module.
1Bh	A	R	Off	R	Decompress the main system BIOS runtime code.
1Ch	A	R	Off	A	Pass control to the main system BIOS in shadow RAM.
E0h	R	R	R	Off	Start of recovery BIOS. Initialize interrupt vectors, system timer, DMA controller, and interrupt controller.
E8h	A	R	R	Off	Initialize extra module if present.
EEh	A	A	A	Off	Jump to boot sector.

TABLE 2-16 POST Progress LED Code Table (Port 80h Codes)

POST	Diagnostic LED Decoder				
Code	(G = green, R = red, A = amber)				Description
	MSB			LSB	
20h	Off	Off	R	Off	Uncompress various BIOS modules.
22h	Off	Off	A	Off	Verify password checksum.
24h	Off	G	R	Off	Verify CMOS checksum.
26h	Off	G	A	Off	Read microcode updates from BIOS ROM.
28h	G	Off	R	Off	Initializing the processors. Set up processor registers. Select least featured processor as the BSP.

TABLE 2-16 POST Progress LED Code Table (Port 80h Codes) *(Continued)*

POST Code	Diagnostic LED Decoder (G = green, R = red, A = amber)				Description
2Ah	G	Off	A	Off	Go to Big Real mode.
2Ch	G	G	R	Off	Decompress INT13 module.
2Eh	G	G	A	Off	Keyboard controller test: the keyboard controller input buffer is free. Next, the BAT command will be issued to the keyboard controller.
30h	Off	Off	R	R	Swap keyboard and mouse ports, if needed.
32h	Off	Off	A	R	Write command byte 8042: the initialization after the keyboard controller BAT command test is done. The keyboard command byte will be written next.
34h	Off	G	R	R	Keyboard Init: the keyboard controller command byte is written. Next, the pin 23 and 24 blocking and unblocking commands will be issued.
36h	Off	G	A	R	Disable and initialize the 8259 programmable interrupt controller.
38h	G	Off	R	R	Detect configuration mode, such as CMOS clear.
3Ah	G	Off	A	R	Chipset initialization before CMOS initialization.
3Ch	G	G	R	R	Init system timer: the 8254 timer test is over. Starting the legacy memory refresh test next.
3Eh	G	G	A	R	Check refresh toggle: the memory refresh line is toggling. Checking the 15 second on/off time next.
40h	Off	R	Off	Off	Calculate CPU speed.
42h	Off	R	G	Off	Init interrupt vectors: interrupt vector initialization is done.
44h	Off	A	Off	Off	Enable USB controller in chipset.
46h	Off	A	G	Off	Initialize SMM handler. Initialize USB emulation.
48h	G	R	Off	Off	Validate NVRAM areas. Restore from backup if corrupted.
4Ah	G	R	G	Off	Load defaults in CMOS RAM if bad checksum or CMOS clear jumper is detected.
4Ch	G	A	Off	Off	Validate date and time in RTC.
4Eh	G	A	G	Off	Determine number of microcode patches present.
50h	Off	R	Off	R	Load microcode to all CPUs.
52h	Off	R	G	R	Scan SMBIOS GPNV areas.
54h	Off	A	Off	R	Early extended memory tests.
56h	Off	A	G	R	Disable DMA.
58h	G	R	Off	R	Disable video controller.

TABLE 2-16 POST Progress LED Code Table (Port 80h Codes) *(Continued)*

POST Code	Diagnostic LED Decoder (G = green, R = red, A = amber)				Description
5Ah	G	R	G	R	8254 timer test on channel 2.
5Ch	G	A	Off	R	Enable 8042. Enable timer and keyboard IRQs. Set video mode initialization before setting the video mode is complete. Configuring the monochrome mode and color mode settings next.
5Eh	G	A	G	R	Initialize PCI devices and motherboard devices. Pass control to video BIOS. Start serial console redirection.
60h	Off	R	R	Off	Initialize memory test parameters.
62h	Off	R	A	Off	Initialize AMI display manager module. Initialize support code for headless system if no video controller is detected.
64h	Off	A	R	Off	Start USB controllers in chipset.
66h	Off	A	A	Off	Set up video parameters in BIOS data area.
68h	G	R	R	Off	Activate ADM: the display mode is set. Displaying the power-on message next.
6Ah	G	R	A	Off	Initialize language module. Display splash logo.
6Ch	G	A	R	Off	Display sign on message, BIOS ID, and processor information.
6Eh	G	A	A	Off	Detect USB devices.
70h	Off	R	R	R	Reset IDE Controllers.
72h	Off	R	A	R	Displaying bus initialization error messages.
74h	Off	A	R	R	Display setup message: the new cursor position has been read and saved. Displaying the hit setup message next.
76h	Off	A	A	R	Ensure timer keyboard interrupts are on.
78h	G	R	R	R	Extended background memory test start.
7Ah	G	R	A	R	Disable parity and NMI reporting.
7Ch	G	A	R	R	Test 8237 DMA controller: the DMA page register test passed. Performing the DMA controller 1 base register test next.
7Eh	G	A	A	R	Initialize 8237 DMA controller: the DMA controller 2 base register test passed. Programming DMA controllers 1 and 2 next.
80h	R	Off	Off	Off	Enable mouse and keyboard: the keyboard test has started. Clearing the output buffer and checking for stuck keys. Issuing the keyboard reset command next
82h	R	Off	G	Off	Keyboard interface test: A keyboard reset error or stuck key was found. Issuing the keyboard controller interface test command next.

TABLE 2-16 POST Progress LED Code Table (Port 80h Codes) *(Continued)*

POST Code	Diagnostic LED Decoder (G = green, R = red, A = amber)				Description
84h	R	G	Off	Off	Check stuck key enable keyboard: the keyboard controller interface test is complete. Writing the command byte and initializing the circular buffer next.
86h	R	G	G	Off	Disable parity NMI: the command byte was written and global data initialization has completed. Checking for a locked key next.
88h	A	Off	Off	Off	Display USB devices.
8Ah	A	Off	G	Off	Verify RAM size: Checking for a memory size mismatch with CMOS RAM data next.
8Ch	A	G	Off	Off	Lock out PS/2 keyboard/mouse if unattended start is enabled.
8Eh	A	G	G	Off	Initialize boot devices: the adapter ROM had control and has now returned control to the BIOS POST. Performing any required processing after the option ROM returned control.
90h	R	Off	Off	R	Display IDE mass storage devices.
92h	R	Off	G	R	Display USB mass storage devices.
94h	R	G	Off	R	Report the first set of POST errors to Error Manager.
96h	R	G	G	R	Boot password check: the password was checked. Performing any required programming before Setup next.
98h	A	Off	Off	R	Float processor initialize: performing any required initialization before the coprocessor test next.
9Ah	A	Off	G	R	Enable Interrupts 0, 1, 2: checking the extended keyboard, keyboard ID, and NUM Lock key next. Issuing the keyboard ID command next.
9Ch	A	G	Off	R	Initialize FDD devices. Report second set of POST errors to error messager.
9Eh	A	G	G	R	Extended background memory test end.
A0h	R	Off	R	Off	Prepare and run setup: Error manager displays and logs POST errors. Waits for user input for certain errors. Execute setup.
A2h	R	Off	A	Off	Set base expansion memory size.
A4h	R	G	R	Off	Program chipset setup options, build ACPI Tables, and build INT15h E820h table.
A6h	R	G	A	Off	Set display mode.
A8h	A	Off	R	Off	Build SMBIOS table and MP tables.
AAh	A	Off	A	Off	Clear video screen.

TABLE 2-16 POST Progress LED Code Table (Port 80h Codes) *(Continued)*

POST Code	Diagnostic LED Decoder (G = green, R = red, A = amber)				Description
ACh	A	G	R	Off	Prepare USB controllers for operating system.
AEh	A	G	A	Off	One beep to indicate end of POST. No beep if silent boot is enabled.
000h	Off	Off	Off	Off	POST completed. Passing control to INT 19h boot loader next.

2.7 OS Operations

2.7.1 Filesystem Check (fsck) Procedure

The first step in filesystem repair is to ensure that you have a complete, tested backup. The filesystem check carries some risk. Directories, files and filenames may be lost. A tested backup means that the data has been restored from tape, and checked for validity.

After the backup, the next step is to schedule the file system check. The volume that you are running the filesystem check against will be unavailable for the duration of the process. In addition, if this is the volume containing the /etc directory, all other volumes will be offline for the duration of the process. In any case, there will be a heavy load on the filesystem that will affect all clients. It is difficult to determine how long the process will take, as there are several variables which cause this time to vary, such as system specifications, size of volume, workload, and how many errors are found. The check should be run as soon as possible, as the filesystem problems can potentially worsen when writing to a damaged volume.

As a general rule, allow five hours for each run, more if a large number of errors are expected. Also note that if any errors are found, multiple runs are always required. Because of the time involved, consideration should be given to recreating the volume and restoring from a backup. This decision should be made based on the severity of the problem. A read-only filesystem check may be helpful in making this determination, but this may add several hours to the process.

Next, run the fsck procedure. This is done at the StorEdge CLI. It is strongly recommended to log the output of the filesystem check session for escalation purposes. Therefore, you should access the CLI with a client that is capable of logging, such as a LAN connected client or a serial console. Using a dial-up or WAN connected client is not recommended, as this can extend the run time of the procedure.

At the CLI, enter “fsck <volumename>”. You will then be prompted whether repairs should be made if errors are found. Generally, the answer should be “y” for “yes”. The other potentially useful option is “n” for “no”. This will run a check against the volume without writing the repairs. As noted above, this can be used to make decisions about running the filesystem check.

If errors are reported by the filesystem check, the filesystem check must be repeated until there are no errors. This may require several runs of the filesystem check. In this case, the following message is displayed:

```
sfs2ck vol1: no errors
```

It is also possible, but very rare, that the above message will never be seen. This can occur in extreme cases where the filesystem check is unable to completely repair a volume. In these cases, the volume should be deleted and restored from tape.

Another rare possibility is that the filesystem check can fail and either hang or reboot. In this case, proceed according to the instructions under the heading “System hang or reboot during normal operation” above, and escalate the issue immediately.

If repairs were made by the filesystem check, file and directory names are sometimes lost. These files are issued a name that begins with “Node”, followed by numbers related to the inode location in the filesystem. This number is generally not useful, other than to ensure a unique filename. These files and directories retain their original contents, to the extent possible. Manual inspection of these files is required to determine the original file type and filesystem location.

2.7.2 StorEdge Network Capture Utility

Sun StorEdge 5310 NAS includes a built-in network monitoring tool. This allows you to capture packets from the network and save them to a file. This can be a valuable troubleshooting tool.

To configure network monitoring, it must first be loaded at the StorEdge CLI.

To access the StorEdge CLI:

1. **Connect to the StorEdge via Telnet or serial console, and type `admin` at the [menu] prompt and enter the administrator password.**
2. **At the CLI, enter `load netm`.**
3. **Then type `menu` to configure capture and capture packets.**
4. **Press the spacebar until “Packet Capture” is displayed under “Extensions” at the lower right.**
5. **Select the letter corresponding to “Packet Capture”.**

6. **Select option "1", Edit Fields.**

   The available options are as follows:
   - Capture File—Where to save the capture file, in the format /volumename/directory/filename
   - Frame Size (B)—Size in bytes of each frame to capture. The default is normally used.
   - IP Packet Filter—"No" captures all traffic, "Yes" allow you to filter what is received.

     A filter allows you to select which IP address or addresses you will capture traffic from. You can also filter on a particular TCP or UDP port.
   - Dump Enable—Select "Yes" to allow StorEdge to save the capture in the event of a problem.

7. **After configuring these options, select option "7", "Start Capture"**

8. **Reproduce the network event you wish to capture.**

9. **Select option "7", "Stop Capture".**

10. **Access the file via NFS or SMB and copy the file as needed.**

2.7.3 Upgrades

2.7.4 Cacls - Access Control List

For issues with access to a file or directory, collect the output of the cacls command. This command is available from the CLI. At the CLI, enter "cacls <full pathname>". The full pathname should begin with the volume name, as in this example: "cacls /vol1/testfile.txt".

Cacls output contains the following information:

First, the basic mode information and UID/GID of the owner is displayed. Here is an example:

```
drwxrw----       34      22      /vol1/data
```

In this case, we can see that the item is a directory, with 750 permissions: Read/write/execute (7) for the owner (UID 34), Read/write for members of the owner's group (GID 22), and no permissions (0) for everyone else.

Listed next are Creation time, FS Creation time, and FS mtime. These are timestamps associated with the file and the filesystem, generally only useful for troubleshooting timestamp issues.

Next is the Windows security descriptor. In its simplest form, it will read “No security descriptor”. This means that no Windows security is present, and that Windows will simulate security based on the above NFS permissions.

If a Windows security descriptor is present, the following information is displayed:

- Security Descriptor:The type of security descriptor. This can be disregarded.
- Owner:The user name or SID of the owner.
- Primary Group: The group name or SID of the group owner.
- Discretionary Access Control List (DACL):A list of users who have access to the file, by SID.

A SID is a number that uniquely identifies a user or group. The data to the right of the final dash identifies the user within the domain; the rest of the number indicates domain and type of account information. This user information is known as the RID (relative ID). The RID is the number used for user mapping. It can be cross-referenced with the StorEdge user or group mapping data determine the user/group name and NFS UID/GID.

2.7.5 Proc filesystem

The /proc filesystem is a virtual filesystem used to collect system data. The location of some of the more useful data is listed below. To collect the data, copy the file, or use the “cat” CLI command to dump it to the screen while logging the terminal session.

```
/proc/cifs/DOMAIN.USER.6789ABCD…
```

These are user access tokens. They may be useful in troubleshooting SMB issues.

These file names begin with the domain name, then the username, then some hexadecimal digits. The hexadecimal digits are a representation of the IP address, which can be used to discern between multiple logins for a user. If you do not see the user token that you need, it may be necessary to log the user off for thirty seconds, and then back on in order to capture the token.

```
/proc/cifs/pdc
```

The currently connected domain, domain controller, and the IP address of the domain controller.

```
/proc/cifs/ntdomain
```

A list of all trusted domains, their related SIDs, and the local machine and local domain SIDs.

2.7.6 FTP Server

To use the built in ftp server, you need to load the ftp daemon from the command line. The command is as follows:

```
load ftpd <CR>
```

This will allow you to ftp files to and from the Sun StorEdge 5310 NAS.

2.8 Updating the OS on the Sun StorEdge 5310 NAS

This section provides information on firmware and BIOS upgrades

This section contains the following topics:

- "Sun StorEdge 5310 NAS Firmware" on page 2-40

2.9 Sun StorEdge 5310 NAS Firmware

2.9.1 Operating System

2.9.1.1 Upgrading the StorEdge Operating System

The StorEdge software can be updated either via the StorEdge Web Admin or via copying the file directly. Before performing the upgrade, you must have downloaded the software and extracted the upgrade file. This file should have the extension ".img". This file should be stored locally on the client from which the software upgrade will be done. The operating system upgrade requires a system reboot which should be done immediately after copying new OS to the system.

Important – After the reboot, the system may take as long as five minutes to complete the software upgrade and return to service. There is no visual indication that this process is taking place. The StorEdge LCD displays "…booting…" during this process. If it is necessary to check the status of the upgrade, connect a display to the StorEdge.

To update the operating system via web interface:

1. **To use the Web Admin, connect with a Web browser to http://<hostname or IP address of your StorEdge>.**
2. **Click "Grant" or "Yes" to accept any Java software authorization windows and you will reach the login screen.**
3. **Type the administrator password to access the administration interface.**
4. **Navigate to System Operations/ Update Software.**

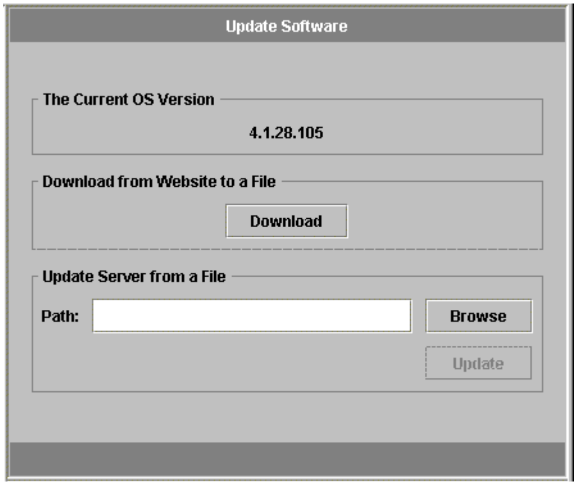

Update Software

The Current OS Version

4.1.28.105

FIGURE 2-8 The Update Software Panel

5. **Click on the Browse button and navigate to the directory containing the new OS image file.**
6. **Click to select the OS image file**
7. **Click the "Open" button. The file name and path are now displayed in Update Software window**
8. **Click on the Update button**
9. **Wait for the OS image file to upload to the StorEdge.**

10. **When the update process is complete, click Yes to reboot, or No to continue without rebooting. The update does not take effect until the system is rebooted.**

    To update the operating system via file copy:

1. **Access the StorEdge via SMB or NFS.**
2. **Via SMB, access the share c$.**

   You must be a member of the local Administrators group to access this share. Via NFS, mount to /cvol. By default, this is only possible from a trusted host.
3. **In either case, copy the operating system image to the root of /cvol.**
4. **Next, reboot the StorEdge via one of the administration interfaces.**

   The operating system upgrade will take place before the system comes up.

2.10 Common Problems Encountered on the Sun StorEdge 5310 NAS

This chapter describes common problems with the Sun StorEdge 5310 NAS.

It includes the following sections:

2.11 CIFS/SMB/Domain

Changes to Windows group membership do not take effect. Changes to user mapping do not take effect.

Windows clients use a device called an access token to assign user data and group membership. This token is assigned when the client connects to the StorEdge. Any changes to this token are not implemented until the next time the user connects.

To cause any changes to take effect immediately, ensure that the user closes all sessions with the StorEdge.

The easiest way to do this is to log the user out of all connected workstations. It is necessary that the user remain disconnected for approximately 30 seconds because the token is cached for a short time.

To ensure that all users' tokens are updated after making large-scale changes, reboot StorEdge. This action ensures that all sessions are disconnected.

Windows clients cannot connect by NetBIOS name. StorEdge not present in browse list / Network Neighborhood.

A master browser is a server that is configured to manage CIFS/SMB browse lists and respond to client requests for them. Windows server operating systems are configured to do this by default.

StorEdge is configured not to act as a master browser. This is done to dedicate all StorEdge resources to file sharing.

For the browsing to function correctly, each subnet or physical network segment must have a master browser. Therefore, if you wish to make the StorEdge available via browse lists, it should be located on the same segment and subnet as a Windows Server.

Note that configuring a WINS server improves the performance of browsing, and in some cases may compensate for the lack of a master browser on some segments. If possible, a WINS server should always be configured.

To determine which master browser, if any, StorEdge has located, generate a system diagnostic.

1. **To access this functionality, access the StorEdge via Telnet.**

2. **Press enter at the [menu] prompt and enter the administrator password.**

3. **Press the spacebar until "Diagnostics" is displayed under "Extensions" at the lower right.**

4. **Select the letter corresponding to "Diagnostics."**

   "Please wait…" is then displayed in the upper left.

   After a short time, the system diagnostics is displayed.

5. **Scroll through the diagnostics with the [spacebar] and [b] keys.**

6. **Under the heading "NETBIOS Cache" look for an entry with a <1D> tag.**

   <1D> is a segment master browser.

7. **Verify that this <ID> entry matches your domain name and IP subnet.**

8. **If no browser is found, either move a server to the subnet that the StorEdge is on, or move StorEdge to a subnet with Windows servers.**

Cannot join Windows Domain.

To authenticate users from a Windows Domain, StorEdge must locate a Domain Controller, authenticate, and then add a computer account to the domain.

Users from the domain are not able to establish a connection to the StorEdge until this entire process has succeeded.

The first step towards resolving this issue is data collection. The two primary sources of data are the system log and the StorEdge NetBIOS cache. Note that this data collection must take place as soon as possible after the failed attempt to join the domain.

To check the system log, proceed as follows:

1. **Access the StorEdge via Telnet.**

2. **Press enter at the [menu] prompt and enter the administrator password.**

3. **Select option "2", Show Log.**

   The fourteen most recent syslog messages are displayed.

4. **Look for messages related to the attempt to join the domain.**

   The first message typically contains the words "join domain".

5. **If no messages are found, select option "1", Show Entire Log.**

6. **Page through the log with the space bar, scrolling to the approximate time and date that you made the most recent attempt to join the domain.**

7. **Look again for the messages related to joining the domain.**

8. **If no applicable messages are found, repeat the attempt to join the domain, and check the log again.**

   The system log is also available through the StorEdge Web Admin.

   To access it, log in, and navigate to: Notification and Monitoring/View System Log.

   You can scroll through the log, or save it as a file.

   To check the NetBIOS cache, proceed as follows:

1. **Access the StorEdge via Telnet.**
2. **Press enter at the [menu] prompt and enter the administrator password.**
3. **Press the spacebar until "Diagnostics" is displayed under "Extensions" at the lower right.**
4. **Select the letter corresponding to "Diagnostics".**
5. **Wait a few seconds while the StorEdge builds the diagnostic.**
6. **When the diagnostic is ready, you can page through it here, with [space] and [b], or you can email it or save it to a file.**
7. **In either case, search through the file for the heading "NETBIOS Cache". Note each of the NetBIOS tags.**

   Each NetBIOS tag is displayed in the form: Hostname<##>, or Domain<##>, with one or more IP addresses associated with it. <##> is a number expressing a particular NetBIOS service being advertised.

   The tags you should be concerned with are as follows:

   - Hostname<00>: Local workstation service for hostname.
   - Hostname<20>: Local server service for hostname.
   - Domain<00>: Indicates inclusion in the domain or workgroup for the included IP address.

Note – Does not necessarily indicate domain membership.

   - Domain<1D>: Segment master browser(s) for the listed domain. This server provides browsing services for this domain only on this IP subnet.
   - Domain<1C>: Domain Controller for listed domain. Either a Primary (PDC) or Backup (BDC).
   - Domain<1B>: Primary Domain Controller for listed domain. By definition, the browse master for its own subnet, and the collector of all data from other browse servers.

Using these two information sources, you can begin to diagnose the problem. The following are the most common possible problems along with their indicating symptoms.

Wrong password / insufficient permissions: This is usually indicated by a logon failure or access denied message in the system log. The user account that is entered into the StorEdge Domain configuration screen must have the correct password, and must have the authority to create computer accounts. Typically, a user account that is a member of the Domain Admins global group is used.

No master browser on the subnet: CIFS/SMB relies on a hierarchical system of browser servers. Each IP subnet and network segment must have at least one such server, known as a "master browser" in order for systems on that subnet to locate network resources. StorEdge does not provide master browser services.

The first indication of this is the log message "No Master Browsers found for <domain>". Check the NetBIOS cache for the <1D> or <1B> tag with an IP matching your subnet. Double check the domain name used against the one in the NetBIOS tag of the master browser. It may be necessary to move the StorEdge to the same subnet and segment as a master browser. All Windows server operating systems provide master browser services by default. Installing StorEdge on the same subnet as a Domain Controller is the best practice when possible.

If the problem persists after ensuring that StorEdge has a local master browser, check the solutions below under "Multiple subnets connected to StorEdge".

Other browsing problems: The log message, "Join domain [local]: locate failed" indicates that a Domain Controller could not be found. Note that this message also appears in conjunction with the above "No Master Browser found" message. When that message is present, the above solutions should be followed first.

Start by looking at the NetBIOS cache. Look for <1B> or <1C> domain controller tags. If you see any of these, ensure that the domain name matches the one configured on StorEdge. If you see a <1D> segment master browser, but no <1B> or <1C> tags, check the NetBIOS cache on the master browser system. This is done with the "nbtstat –c" at the Windows command prompt. The output is essentially the same as the StorEdge NetBIOS cache display. If no domain controllers are present in the master browser's NetBIOS cache, then there is a network browsing issue that needs to be addressed.

One possible solution is to add a WINS server. WINS helps to speed Windows browsing and compensates for browsing problems. WINS can compensate for browser deficiencies, and should be used whenever possible. In order for WINS to function properly, the browsers, domain controller and StorEdge must be configured to use the server for lookups. Another possible solution is to move the StorEdge to the same subnet as the Domain Controller, though this does not address the larger browsing problem.

If these solutions have been attempted to no avail, also see the following solution.

Multiple subnets connected to StorEdge: Care must be taken when StorEdge is connected to multiple subnets, particularly when the subnets are disjoint, i.e. not connected to one another. A common example of this is a direct connection to a backup or database server.

The problem created by the disjoint subnets is that StorEdge registers each of its IP addresses via NetBIOS broadcast and/or WINS. The Domain Controller may select one of the addresses on a disjoint subnet, and fail to communicate with StorEdge, resulting in a failure to join the domain. The solution to this is to prevent NetBIOS registration of those addresses not connected to the main network.

In the case of the backup or database server, the solution is easy. The StorEdge "independent" NIC role was created expressly for this purpose. To configure the NIC role, proceed as follows:

1. **Access the StorEdge via Telnet.**
2. **Press enter at the [menu] prompt and enter the administrator password.**
3. **Select option "A", Host Name & Network.**
4. **Select option "1", Edit fields.**
5. **Navigate through the fields with [Tab] or [Enter] until the "Role" field of the desired NIC is highlighted.**
6. **Select option "3", Independent.**
7. **Select option "7", Save Changes.**

It is also possible to disable the NetBIOS registration without changing the NIC role. This can be done if you have a problem after attempting the above, or if you have a requirement to leave the role as primary. To make this configuration change, proceed as follows:

1. **Connect to the StorEdge via Telnet, and type "admin" at the [menu] prompt and enter the administrator password.**
2. **At the CLI, enter "load smbtools", and then "smbwins exclude addr= 192.168.243.1".**

   This action prevents these IP addresses from being registered via NetBIOS. However, the master browsers and WINS servers do not immediately remove these addresses. To accomplish this, proceed as follows:

3. **Remove the entry for StorEdge from any WINS server databases.**
4. **Locate the master browser for the local subnet and any local Domain Controllers.**
5. **Enter "nbtstat –R" at the Windows CLI on each of these systems.**

6. **Reboot the StorEdge.**

   Note that the above changes do not take effect until after the reboot.

   This action removes the undesired entries in almost every case. The only case where the entries may persist is in a multiple server WINS environment using replication. In this case, consult the provider of the WINS server operating system for removal instructions.

   Anonymous connections restricted by Domain Controller: In this case, the master browser and domain controller are both located, but the system log shows a number of RPC errors related to security, along with the name and IP address of the Domain Controller to which it is attempting to authenticate.

   Windows 2000 and later operating systems can be configured to refuse anonymous connections, otherwise known as null sessions. Typically, this is done for security reasons. Restricting anonymous connections is not recommended unless all clients and servers in the domain are running Windows 2000 or newer. StorEdge and other non-Windows servers require a change to this policy.

   This setting is accessed via the registry editor on the Windows domain controller. Using the Registry Editor, navigate to the key: "HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Control\LSA". Locate the value RestrictAnonymous. If it is set to "2", modify it to "0" or "1". A setting of "0"

   The Domain Controller must be rebooted for this change to take effect.

   Connected to a DC across a WAN link: In rare cases, it is possible that StorEdge will join a domain using a distant Domain Controller across a slow link. The symptoms in this case will vary. You could see timeouts, authentication failures due to a firewall, or even success with poor performance. The primary indication will be log messages indicating any of the above problems, and referring to communications with a Domain Controller on a faraway subnet.

   To resolve this issue, first check the NetBIOS cache as directed above to ensure that the local domain controllers are present. If not, proceed as above to correct any difficulty locating them. After verifying the presence of one or more nearby Domain Controllers (<1B> or <1C> NetBIOS tags), proceed as follows to force StorEdge to use a particular Domain Controller:

1. **Connect to the StorEdge via Telnet, and type "admin" at the [menu] prompt and enter the administrator password.**

2. **At the CLI, enter "set smb.pdc <IP address>", replacing <IP address> with the IP address of one of the above domain controllers. In spite of the variable name, it is acceptable to use either a PDC <1B>, or a BDC <1C>.**

3. **After setting the variable, retry the attempt to join the domain. Check the system log to ensure success.**

Assuming that the difficulty connecting to the Domain Controller is temporary, and related to network load, it should not be necessary to save this variable with the savevars command. Doing so will limit the ability of StorEdge to find an alternate Domain Controller in the case that this one fails.

Cannot connect or authenticate to Windows 2003 Domain Controller.

By default Windows 2003 is configured to require signed digital communications from clients. This is also known as SMB packet signing. StorEdge does not support packet signing. Therefore, Windows 2003 must be configured to negotiate packet signing rather than assuming that it is present.

1. **To configure this, you must access the Local Security Policy Editor on the Windows 2003 Server.**
2. **Next, navigate to Security Settings/Local Policies/Security Options.**
3. **Scroll down to "Microsoft network server: Digitally sign network communications (always)"**
4. **Double click the entry and click the "Disabled" button.**
5. **Click "OK".**

Changing this setting does not restrict the Windows 2003 server from using packet signing with those clients that support it.

Lost Connection with Windows Domain.

In some conditions, it is possible for StorEdge to lose connection to the Domain Controller. In this case, Windows users will be denied access to the StorEdge, and they will be prompted for a password.

Possible reasons for this include modification of administrative user password, network problems or failure of PDC.

The solution to each of these is the same. It is necessary to re-enter the user and password information in the domain setup screen. This is done as follows:

1. **Access the StorEdge via Telnet.**
2. **Press enter at the [menu] prompt and enter the administrator password.**
3. **Press the spacebar until "CIFS/SMB Configuration" is displayed under "Extensions" at the lower right.**
4. **Select the letter corresponding to "CIFS/SMB Configuration".**

5. **Select the letter corresponding to "Domain Configuration".**
6. **Use the [Enter] or [Tab] key to navigate to the User name field.**
7. **Enter a user name for the listed domain with the rights to add a computer account.**
8. **Press [Enter] to move to the "Password" field.**
9. **Enter the password for this user.**
10. **Select option "7", Save Changes.**

If the attempt to join the domain is unsuccessful, proceed according to the instructions in the Troubleshooting Guide: "Cannot join Windows Domain".

This functionality is also available through the StorEdge Web Admin. To use the Web Admin, connect with a Web browser to http://<hostname or IP address of your StorEdge>. Click "Grant" or "Yes" to accept any Java software authorization windows and you will reach the login screen. Type the administrator password to access the administration interface.

Navigate to Windows Configuration/Configure Domains and Workgroups. Enter a user name for the listed domain with the rights to add a computer account and the associated password.

If the attempt to join the domain is unsuccessful, proceed according to the instructions in the Troubleshooting Guide: "Cannot join Windows Domain".

CIFS/SMB share changed to hidden is still visible on network. Renamed share, but old name is still displayed in browse list.

Share lists are sometimes cached by the client's network redirector. This problem will clear itself within a short time, 30 minutes at the most.

Cannot set share security, all shares inherit the security of the directory object.

The StorEdge security implementation allows only for securing files and directories. The effective security of a CIFS/SMB share is always the security of the directory to which it points.

StorEdge has same files in 2 different shares.

This is caused by creating multiple share names that point to the same directory or volume. Shares always point to a directory. Root level shares will always contain all files on the volume, regardless of how many shares are created to this volume. View shares as pointers, with the understanding that many of these pointers may exist to a single location.

User maps are incorrect.
User maps are not automatically created.

The requirements for successful user mapping are to import all NFS users to the StorEdge and define a mapping rule. These requirements must be met before any CIFS/SMB users have connected. If CIFS/SMB users connect before both of these are in place, the user will be mapped to a StorEdge-generated UID.

Once the mapping has been created, it will not be overwritten by subsequent connections with updated credential info.

Windows users are not mapped to the expected NFS group.
Mappings are not created for most Windows groups.

Although Windows users can maintain membership in many groups, the StorEdge user and group mapping functionality only recognizes the Primary Group. By default, all Windows Users are assigned the primary group, “Domain Users”. The only exception to this is if they are a member of the “Domain Admins” group at the time the user account is created, in which case this group is assigned as the primary group.

In order for group mapping between CIFS/SMB and NFS to be effective, primary group assignments must be made selectively. It may be necessary to create some groups. Primary group assignment is done is Windows User Manager for Domains, usually from a Domain Controller. See your Windows documentation for details on how to configure this setting. It is important to use only Windows “Global Groups” for this purpose. Windows Local Groups are intended to be used only locally, on the Domain Controllers themselves.

After making modifications to Windows users’ primary groups, groups with no mappings will be mapped to NFS groups according to the mapping policy. The mapping will be automatically created as soon as a CIFS/SMB user connects with a primary group which is not in the StorEdge group.map file. Before this happens, make sure that group information is imported to StorEdge, either manually or via NFS, and the desired mapping policy is in place.

Another way to resolve this, for users with primary group assignments in the passwd file, is to use the “Map to Primary Group” policy.

Can’t copy greater than 4G file from Windows to StorEdge.

This problem may be seen on Windows 2000 and prior versions. If running Windows 2000, it can be fixed applying the latest service pack. If running an older version, there is no fix available, though you may be able to work around the problem with the Windows backup utility or a similar third party solution.

Can’t map drives via CIFS/SMB.

In order to map a drive or connect to a share, you must have read access to the directory to which the share points. If StorEdge is in domain mode, you must also be logged in to the domain. File and directory security can be checked at the StorEdge CLI.

1. **To access the StorEdge CLI, connect to the StorEdge via Telnet, and type “admin” at the [menu] prompt and enter the administrator password.**
2. **At the CLI, enter “cacls <path>”. The path must include the volume name. If the path includes spaces, enclose the argument in double quotes, as in cacls “/vol1/my directory/my file”.**

   Cacls output contains the following information:

   First, the basic mode information and UID/GID of the owner is displayed. Here is an example:

   ```
   drwxrw----        34      22      /vol1/data
   ```

   In this case, we can see that the item is a directory, with 750 permissions: Read/write/execute (7) for the owner (UID 34), Read/write for members of the owner’s group (GID 22), and no permissions (0) for everyone else.

   Listed next are Creation time, FS Creation time, and FS mtime. These are timestamps associated with the file and the filesystem, generally only useful for troubleshooting timestamp issues.

   Next is the Windows security descriptor. In its simplest form, it will read “No security descriptor”. This means that no Windows security is present, and that Windows will simulate security based on the above NFS permissions.

   If a Windows security descriptor is present, the following information is displayed:

   - Security Descriptor:The type of security descriptor. This can be disregarded.
   - Owner:The user name or SID of the owner.

- Primary Group: The group name or SID of the group owner.
- Discretionary Access Control List (DACL):A list of users who have access to the file, by SID.

A SID is a number that uniquely identifies a user or group. The data to the right of the final dash identifies the user within the domain; the rest of the number indicates domain and type of account information. This user information is known as the RID (relative ID). The RID is the number used for user mapping. It can be cross-referenced with the StorEdge user or group mapping data determine the user/group name and NFS UID/GID.

From there, it is simply a matter of assigning appropriate rights to the user attempting to access the directory. Set security as desired using a Windows Domain Admin account.

Can’t set Windows security at the root of a volume or at the base of a share.

Windows security is set by right clicking on an object, and then selecting the security tab. If you wish to do this for the root of a volume, first map a drive to the share, then right click on the mapped drive within “My Computer”. You will then be able to access the security tab as normal.

Cannot see the security tab from Windows clients.

Current versions of Windows do not display the security tab unless you have the right to view or change security.

File and directory security can be checked at the StorEdge CLI.

1. **To access the StorEdge CLI, connect to the StorEdge via Telnet, and type “admin” at the [menu] prompt and enter the administrator password.**
2. **At the CLI, enter “cacls <path>”. The path must include the volume name. If the path includes spaces, enclose the argument in double quotes, as in cacls “/vol1/my directory/my file”.**

   Cacls output contains the following information:

   First, the basic mode information and UID/GID of the owner is displayed. Here is an example:

```
drwxrw----        34      22      /vol1/data
```

   In this case, we can see that the item is a directory, with 750 permissions: Read/write/execute (7) for the owner (UID 34), Read/write for members of the owner’s group (GID 22), and no permissions (0) for everyone else.

Listed next are Creation time, FS Creation time, and FS mtime. These are timestamps associated with the file and the filesystem, generally only useful for troubleshooting timestamp issues.

Next is the Windows security descriptor. In its simplest form, it will read "No security descriptor". This means that no Windows security is present, and that Windows will simulate security based on the above NFS permissions.

If a Windows security descriptor is present, the following information is displayed:

- Security Descriptor:The type of security descriptor. This can be disregarded.
- Owner:The user name or SID of the owner.
- Primary Group: The group name or SID of the group owner.
- Discretionary Access Control List (DACL):A list of users who have access to the file, by SID.

A SID is a number that uniquely identifies a user or group. The data to the right of the final dash identifies the user within the domain; the rest of the number indicates domain and type of account information. This user information is known as the RID (relative ID). The RID is the number used for user mapping. It can be cross-referenced with the StorEdge user or group mapping data determine the user/group name and NFS UID/GID.

From there, it is simply a matter of assigning appropriate rights to the user attempting to access the directory. Set security as desired using a Windows Domain Admin account.

Windows anti-virus, backup or file management software runs endlessly, following symbolic links.

By default, StorEdge follows symbolic links in Windows. Windows cannot differentiate between links and standard files. Therefore, if a symbolic link points to a location in the filesystem above its own location, Windows applications can get stuck in a loop following these links.

To correct this behavior, you can either manually exclude such links from the scan or backup, or you can set a variable to disable the following of symbolic links from CIFS/SMB clients. The variable affects all volumes and all CIFS/SMB clients.

This functionality is only available at the StorEdge CLI (command line interface).

1. **To access the StorEdge CLI, connect to the StorEdge via Telnet, and type "admin" at the [menu] prompt and enter the administrator password.**
2. **At the CLI, enter `set smb.dir_symlink.disable yes`**

3. **After setting any variables on the StorEdge, i.e. anytime the "set" command is used, the command `savevars` must be entered at the command line in order for the settings to persist though future server reboots.**

CIFS/SMB share created to /cvol is not visible or accessible.

StorEdge does not allow the export of /cvol by default. The /cvol volume exists on compact flash memory which is very space limited and contains sensitive operating system files. This volume should only be accessed under while following documented procedures or on the direct advice of technical support.

Also, the administrative share "c$" is created for administrator only access to /cvol.

Still, if it is necessary to create a user share, this is possible by modifying a variable at the CLI.

1. **To access the StorEdge CLI, connect to the StorEdge via Telnet, and type "admin" at the [menu] prompt and enter the administrator password.**
2. **At the CLI, enter "set system.export.bootdir yes". This will allow access to shares which point to /cvol.**

Having problems with Outlook .pst files stored on the StorEdge.

Microsoft recommends against using .pst files for anything except local (on the workstation), temporary mail storage. The recommended solution is either offline storage files (.ost) or Windows Terminal Server. For additional information on these solutions, and the reasoning behind them, please see document #297019 in the Microsoft Knowledge Base.

Cannot access administrative shares in Workgroup mode. Cannot create/remove shares via rmtshare in Workgroup mode.

These operations are normally allowed only to Domain Admin users. Since Workgroup mode does not use access tokens, these actions are disabled for security reasons. The following instructions explain how to enable these features in Workgroup mode. Please note that this is extremely insecure, and not recommended for any environment that requires CIFS/SMB security. It allows unrestricted access to all data on the StorEdge for those who are aware of the hidden shares.

This functionality is only available at the StorEdge CLI (command line interface).

1. **To access the StorEdge CLI, connect to the StorEdge via Telnet, and type "admin" at the [menu] prompt and enter the administrator password.**

2. **At the CLI, enter "set srvsvc.netshare.enable yes". After setting any variables on the StorEdge, i.e. anytime the "set" command is used, the command "savevars" must be entered at the command line in order for the settings to persist though future server reboots.**

CIFS/SMB disconnects from MS SQL Server.
CIFS/SMB disconnects from MS Access.

Though we do not provide support for either of these environments, we have discovered a setting to improve operations. This setting is potentially useful for any client/server type CIFS/SMB application accesses the StorEdge on behalf of many clients. The setting prevents StorEdge from caching client information to improve handling of consecutive requests from multiple clients coming from a single system (the server).

1. **To access this setting, access the StorEdge via Telnet.**
2. **Press enter at the [menu] prompt and enter the administrator password.**
3. **Press the spacebar until "CIFS/SMB Configuration" is displayed under "Extensions" at the lower right.**
4. **Select the letter corresponding to "CIFS/SMB Configuration".**
5. **Select A, "Domain Configuration."**
6. **Press the [Enter] or [Tab] key to move through the fields to the "Keep Alive" field.**
7. **Enter the value "0" to turn off SMB keep alive.**
8. **Press [Enter] or [Tab] to navigate to the end of the form.**
9. **Select option "7", Save Changes.**

How can I check/modify which Domain Controller StorEdge is using for authentication?

This functionality is only available at the StorEdge CLI (command line interface).

1. **To access the StorEdge CLI, connect to the StorEdge via Telnet, and type "admin" at the [menu] prompt and enter the administrator password.**
2. **At the CLI, enter "cat /proc/cifs/pdc".**

   The response will be in the form:

   Domain:

Server:

Ipaddr:

3. **To force the StorEdge to a preferred domain controller, set the smb.pdc variable to the IP address of your preferred DC and (re)join the domain. From the command line interface type in**

   `set smb.pdc 192.168.200.136` (IP address of domain controller)

   `savevars`

   `menu`

4. **Press the space bar until "SMB/CIFS Setup" option is displayed in the extension section in the lower right.**
5. **Select the letter of that option**
6. **Enter "1" to edit**
7. **Enter in domain information**
8. **Enter "7" to save**
9. **Reboot system**

Users from trusted domains cannot access StorEdge

The indication of this issue is that clients from the local domain can access the StorEdge, but clients from a trusted domain cannot. You will typically see a log message similar to the following:

```
mlrpc[0x0F]: error: NO_TRUST_SAM_ACCOUNT (0xC000018B)
```

This message indicates that StorEdge is successfully communicating with the local Domain Controller, and that this Domain Controller is denying access to the remote domain. Other Windows servers may allow access due to longer caching of user token information, but the fact remains that this message is coming from the domain controller. The solution is to access the Domain Controllers to reestablish the trust relationship according to Windows documentation. Also, verify that the remote Domain Controller is available, as a temporary outage may show similar symptoms.

Connection to SAMBA domain controller fails.

Although support for Samba to act as a primary domain controller has recently been announced (http://www.samba.org), the current implementation has several problems, as outlined below. Use of StorEdge with a Samba domain controller is not recommended at this time.

The Samba PDC implementation is an ASCII-only implementation, i.e. it does not support Unicode, which impacts foreign language support. All Windows domain controllers support Unicode.

Samba supports a limited subset of the full domain controller interface. If StorEdge is used with a Samba domain controller, NOT_SUPPORTED status messages may appear in the StorEdge log. These status messages originate from the Samba server.

The Samba Net Logon support is limited and problematic. It requires a very specific sequence of commands and logon problems have been observed using both Windows and StorEdge that do not appear when using Windows domain controllers.

The StorEdge currently does not work with SAMBA PDC implementations. By default, the StorEdge is set not to attempt communication with a SAMBA PDC. If future releases of SAMBA properly support a Windows-style PDC implementation then the StorEdge default must be changed to be able to utilize this new implementation.

This setting can only be modified at the StorEdge CLI. To access the StorEdge CLI, proceed as follows:

1. **Connect to the StorEdge via Telnet, and type "admin" at the [menu] prompt and enter the administrator password.**
2. **At the CLI, enter "set smb.samba.pdc yes".**
3. **After setting any variables on the StorEdge, i.e. anytime the "set" command is used, the command "savevars" must be entered at the command line in order for the settings to persist though future server reboots.**

Dial-up CIFS/SMB clients cannot connect to StorEdge.

There are internal and third party solutions that allow CIFS/SMB users to connect to networks remotely via dial-up. Some users have reported problems connecting from these dial-up clients.

In the cases we investigated, the problem was caused by the clients using non-standard ports for CIFS/SMB access. StorEdge only supports NetBIOS/SMB traffic on the standard ports of 137 through 139. The specific port assignments are as follows:

137/tcp	NETBIOS Name Service
137/udp	NETBIOS Name Service
138/tcp	NETBIOS Datagram Service
138/udp	NETBIOS Datagram Service
139/tcp	NETBIOS Session Service
139/udp	NETBIOS Session Service

Windows local groups cannot be added to Access Control List.

Windows Local Groups cannot be used to assign security on remote systems. Local groups are not stored in the Domain SAM database. They exist in the database of individual computers, for use on that computer only. An exception to this is Windows Domain Controllers, which share a set of local groups. However, these are shared only with other domain controllers. Global groups should be used to make security assignments to StorEdge.

StorEdge has its own set of Local Groups. These groups are provided for Windows compatibility purposes. They allow a limited set of permissions, and they cannot be used for security assignments to individual files and folders.

Note – Windows Domain Local Groups are also not supported.

Log Message: Share database corrupt.

This message indicates a problem with the database that contains the CIFS/SMB shares. Typically this can be fixed by repairing the database. The database consists of two files per volume: share.db and a share.inx.

This functionality is only available at the StorEdge CLI (command line interface).

1. **To access the StorEdge CLI, connect to the StorEdge via Telnet, and type “admin” at the [menu] prompt and enter the administrator password.**

2. **At the CLI, enter “load dbck”.**

3. **Then, “dbck /<volumename>/share”.**

   The following will display

```
*** Checking /vol1/share

Should repairs be needed, do you want them made?
```

4. **First, select option “N”, “No”. This will allow the database to be checked read-only.**
5. **If errors are reported, run dbck again as above, and select “Y”, “Yes” to perform the repairs.**

Can’t create new DTQ

A DTQ cannot be defined for existing regular directories. A DTQ must be created using StorEdge administration interfaces (telnet, GUI or command-line), which automatically creates a new directory as part of the DTQ setup.

The maximum number of DTQs per volume is 255. If creation of the 256th DTQ on a volume is attempted, the DTQ will not be created, and an error will appear in the system log.

Nested DTQs are supported. This means that a DTQ can be created as a sub-directory of an existing DTQ. The sub directory DTQ must be created from the command line, using the dtq create command.

If an existing DTQ is moved under another DTQ, the moved DTQ will be converted to a regular directory object and its quota record will be removed. The containing DTQ will take account of the moved objects.

Hard links can be created only inside the same DTQ. Creating a hard link between different DTQs is not allowed. There is no limitation for soft (symbolic) links.

Files moved from DTQ to another directory get new time stamps and copy more slowly.

When moving objects into a DTQ, out of a DTQ or between different DTQs, but still within a single file system volume, the operation is treated as if the object had been moved from one volume to another. The StorEdge will perform a copy & delete operation rather than a rename operation. This takes much more time.

2.12 NFS Issues

NFS root user doesn't have appropriate access.

StorEdge implements a feature known as "root squash". When a user connects as root (UID 0) from an NFS client, StorEdge causes the UID to be mapped to UID 60001, the "nobody" account. In order for an NFS client to have root access to StorEdge, you must create a trusted hosts entry, or explicitly define root access for a particular export.

To test whether you have root access: create a file with the touch command, then use ls –ln filename to view the ownership. If the owner is UID 60001, then correct as above.

NFS root user can't change ownership.
NFS root user can't change security.

This typically occurs when a file or directory has been created or modified by a Windows client. Windows uses complex security descriptors, known as ACLs or Access Control Lists. These cannot always be accurately represented using NFS security attributes. Therefore, to prevent NFS users from circumventing these security descriptors, modification of security or ownership is not permitted on files with ACLs.

It is possible to remove ACL information for a single file, or for an entire volume. Removing the ACL information allows the security and ownership for these objects to be edited from NFS clients once again, with appropriate permissions and ownership.

This functionality is only available at the StorEdge CLI (command line interface).

1. **To access the StorEdge CLI, connect to the StorEdge via Telnet, and type "admin" at the [menu] prompt and enter the administrator password.**
2. **At the CLI, enter "chsmb <filename>" or "chsmb <volumename>".**

   For <filename>, use a full path, including volume. A directory is acceptable for <filename>, but chsmb can not be run recursively. A warning is displayed only when a <volumename> argument is used.

   It is also possible to modify this behavior as a system policy. Generally, this is not recommended in environments where Windows security is important.

Modifying the system policy is also done at the CLI. Access the CLI as above, and enter “set acl.overwrite.allowed”. After setting any variables on the StorEdge, i.e. anytime the “set” command is used, the command "savevars" must be entered at the command line in order for the settings to persist though future server reboots. This particular variable setting will not take effect until the next reboot of the StorEdge.

Trusted host does not have root access.

If a host on the trusted list cannot access StorEdge as root, there is likely a problem with the /etc/approve file.

To check for the presence of a host on the trusted list, proceed as follows:

1. **Access the StorEdge via Telnet.**
2. **Press enter at the [menu] prompt and enter the administrator password.**
3. **Select option “F”, Hosts. Check for the presence of the desired hosts.**
4. **Before checking the approve file, confirm that you have correctly diagnosed the problem by creating a file with the touch command.**
5. **After creating the file, use ls –ln <filename> to view the ownership. If the owner is UID 60001, then correct as follows.**
6. **To determine the location of the active approve file, access the StorEdge CLI (command line interface).**
7. **To access the StorEdge CLI, connect to the StorEdge via Telnet, and type “admin” at the [menu] prompt and enter the administrator password.**
8. **At the CLI, enter “show file.approve”. This will return the location of the active approve file.**
9. **Next, enter “cat”, followed by the output of the show command, e.g. “cat /sysvol/etc/approve”.**

   The first line that is not a comment line (comment lines begin with the “#” character) should read as follows:

   ```
   files/@trustedaccess=rw uid0=0
   ```

   If this line is missing, or if the uid0=0 is not present, correct as follows:

1. **From the CLI type “load unixtools”**
2. **Next, type “cp <active approve file and path> <active approve file and path>.bak”. This backs up the current file with a .bak extension.**
3. **Next, type “cp /cvol/nf0/approve <active approve file and path>”. This overwrites the current approve file with a default copy.**

4. **Finally, type “approve update”. This causes the new file to become active.**

   Now, you should be able to mount StorEdge from a trusted host via NFS.

Local NIS files are no longer updating.

The first step is to check the system log. This will tell us if there is a problem connecting to the NIS server.

1. **To do so, access the StorEdge via Telnet.**
2. **Type “admin” at the [menu] prompt and enter the administrator password.**
3. **At the CLI, enter “menu”.**
4. **Select option “2”, Show Log. This will display the fourteen most recent syslog messages.**
5. **Select option “1”, Show Entire Log. Browse through the log one screen at a time, looking for any messages relating to the NIS server. You might find errors that a server is unreachable via TCP, or if the NIS server is defined by name, there could be a problem communicating with the DNS server.**

   If no such messages are found, it is likely that the NIS monitor thread needs to be reinitialized. To perform this operation, press [Esc] until you exit to the CLI, then enter “niscfg –k” at the CLI. After this, type “menu” to return to the menu, and check the log once again as above to verify that NIS files have been updated.

Windows created files are root owned when viewed via NFS. (In Windows Domain mode)

This occurs when a Windows user is a member of the Domain Administrators group. The Domain Admins group of the local domain is always mapped to root user and group (UID and GID 0). If it is necessary for a member of Domain Admins to share files with a specific UNIX$^{(R)}$ user account, you must change the user’s primary group to a group other than Domain Admins. This is done by editing the Windows Domain user account.

It is also necessary to change the ownership of the file to the user rather than the Domain Admins group. Do this by right clicking on the object within Windows, and selecting Properties/Security. Please note that you must have the appropriate rights to view or set security, and that you can perform these operations recursively with a checkbox.

Windows created files are root owned when viewed via NFS. (In Workgroup mode)

Workgroup mode assigns ownership per share, based on the UID and GID settings configured when the share was defined. By default, this is set to UID and GID 0, leaving the files root owned. The best way to manage ownership of files in Workgroup mode is to have each user access StorEdge via a unique share, and define the UID/GID settings accordingly.

International NFS filenames are garbled or cannot be read from Windows. (Or vice versa.)

Windows uses Unicode UTF-8 for directory and filename storage. In order to read extended characters cross-platform, you must also use a UTF-8 codepage on the NFS clients.

By default, StorEdge assumes that all filename and directory name data received from NFS clients is ASCII text. If your NFS clients are using UTF-8 encoding, StorEdge needs to be configured to accept UTF-8 data. Note that this setting is only important if you intend to share data with Windows clients.

It is imperative that this configuration setting is made before any filenames with extended characters are written from NFS clients. Otherwise, filenames written prior to the change may become completely inaccessible. If a change is necessary on a system that already contains this type of data, any such data should be moved from the StorEdge, and re-migrated after configuring the system correctly. Please note that tape backup is not acceptable for this purpose. The data must be moved to a system that is using the same codepage as the NFS client that wrote the data.

The NFS UTF-8 setting is only available at the StorEdge CLI (command line interface).

1. **To access the StorEdge CLI, connect to the StorEdge via Telnet, and type “admin” at the [menu] prompt and enter the administrator password.**

2. **At the CLI, enter “set nfs.utf8 yes”. This tells the StorEdge that the NFS clients send data in UTF-8 format. After setting any variables on the StorEdge, i.e. anytime the “set” command is used, the command "savevars" must be entered at the command line in order for the settings to persist though future server reboots.**

In addition, you must configure the language codepage that StorEdge should use. The best way to do this is with the StorEdge Web Admin.

1. **To use the Web Admin, connect with a Web browser to http://<hostname or IP address of your StorEdge>.**

2. **Click “Grant” or “Yes” to accept any Java software authorization windows and you will reach the login screen.**

3. **Type the administrator password to access the administration interface.**
4. **Navigate to System Operations/Assign Language. Select the desired language, and click the "Apply" button.**

GID for new NFS objects is incorrect. StorEdge doesn't recognize the set GID bit.

The StorEdge software supports three ways of setting the group ID of new files and directories. The default is to inherit GID from the parent directory in all cases. This behavior is configurable only at the Command Line Interface (CLI).

1. **To access the StorEdge CLI, connect to the StorEdge via Telnet, and type "admin" at the [menu] prompt and enter the administrator password.**
2. **At the CLI, enter "fsctl gidmode <type>". <Type> is one of the following:**
   - bsdInherit the group id from the parent directory
   - sysvApply the creating user's primary group id
   - sgidSet the group id according to the set GID bit (02000) in the parent directory mode.

For sgid, when the set GID (S_ISGID) bit is set, both the group ID and the set GID bit will be inherited from the parent directory. Otherwise, the creating user's group ID is applied.

Chown by root fails if it would put target user over hard quota.

This behavior is by design. StorEdge does not allow users to exceed their hard quota. The solution is to modify the quota either temporarily or permanent.

NFS User can't access his own files created with CIFS/SMB account.

This is an indication that there is a problem with user or group mapping. From an NFS mount, use ls –ln to determine the current ownership of the file. If the owner's UID or GID is not what is expected, proceed according to the instructions in the Troubleshooting Guide: "Windows users are not mapped to the expected NFS group." or "User maps are incorrect."

2.13 Network Issues

When is it necessary to add a TCP/IP route?

By default, StorEdge creates a route for each connected subnet. StorEdge also allows for the configuration of a default gateway. The local routes are used to send packets to the attached subnets, and packets to all other IP addresses are sent via the default gateway. This configuration works for the vast majority of networks.

A manually configured TCP/IP route is required when packets for a particular network or host must travel through a particular network or gateway. This is usually necessary only in cases where there are complex (multiple subnet), disjoint (not connected to one another) networks attached to more than one interface on the StorEdge. A manual route can also be used for performance, where it is known that a particular network connection, other than the default gateway, is a faster route to a particular network location. An example of this is when mirroring is used, and a private connection is desired to a mirror system on another subnet.

How do I manually add a TCP/IP route?

1. **To access this functionality, access the StorEdge via Telnet.**
2. **Press enter at the [menu] prompt and enter the administrator password.**
3. **Select option A, “Host name and Network”.**
4. **Select option 2, “Manage Routes”.**
5. **Select option 1, “Add route”**
6. **Select option 1, “Edit”**
7. **Select either Host, Network, Host Gateway or Network Gateway.**

   A host route defines a route to a particular host; likewise a network route defines a route to a network. The gateway designation specifies that this route defines an external router or gateway that should already be reachable via other defined or default routes. When the Host or Network route is used without the gateway designation, any gateway argument will be ignored, and a local NIC will be used as the gateway for this route.

Note – The list of routes displayed in this menu are only the user-defined static routes. The

---

8. **Select option 7, "Save Changes"**
9. **Press [Esc] to return to the menu, or proceed as above to define another route.**

2.13.1 NIC speed and duplex negotiation issues.

StorEdge is reporting Ethernet transmit and receive errors on a switched network.

By default, the StorEdge Ethernet driver is set to auto-negotiate speed and duplex. This works well the great majority of the time, but occasionally there is a problem in the negotiation between the switch and NIC. The primary indications of this problem are extremely slow transfer speed and an increase in packet errors, usually collisions, which can be found in the port statistics on either the StorEdge or the switch.

Note – These problems are not isolated to StorEdge. We have investigated performance issues where other workstations on the network were the source of performance problems for exactly this reason.

The source of this problem is that the speed and duplex were not successfully negotiated. When a NIC or switch port is initialized, assuming that it is configured to auto-negotiate, the NIC and the port will negotiate the highest available speed and duplex.

One possible cause for this is that either StorEdge or the switch port is not set to auto-negotiate. Forcing a particular setting, such as 1000Mb/full duplex, causes the device not to negotiate. Therefore, if only one side is set to auto-negotiate, the negotiation will fail.

It is also possible that the negotiation can fail even though enabled on both sides. First, you should try another cable, and another switch port, as negotiation problems can be caused by hardware. If negotiation is still unsuccessful, the final recourse is to force both StorEdge and the connected switch port to the desired speed and duplex.

The following commands are used to check and/or force speed and duplex for the StorEdge Ethernet driver.

This functionality is only available at the StorEdge CLI (command line interface).

1. **To access the StorEdge CLI, connect to the StorEdge via Telnet, and type "admin" at the [menu] prompt and enter the administrator password.**

2. **To view the current negotiated rate: Enter "em show all". The name of each NIC and current speed, duplex and link status is displayed.**

3. **To force a particular rate (and disable auto-negotiation): enter "em set <NICname> duplex=<duplex> speed=<speed>". Replace <NICname> with the name of the NIC found in the above show command, <duplex> with either "full" or "half", and <speed with either "10", "100" or "1000". Be sure to force the same settings on the switch.**

   For example, to set the second NIC card to 100baseT full duplex:

   ```
   PROMPT> em show all

   emc1: 100Mb/s FULL-DUPLEX, Link is UP

   emc2: 10Mb/s HALF-DUPLEX, Link is UP

   PROMPT> em set emc2 duplex=full speed=100
   ```

   Here, we checked the speed and duplex, found NIC emc2 to be running at 10Mb/s half duplex, then forced NIC emc2 to 100Mb/s full duplex.

Default gateway changes without user input.

Under certain circumstances, the default gateway on StorEdge can be modified via ICMP. Typically, this is an indication of a network configuration problem. However, StorEdge can be configured to prevent this situation.

This functionality is only available at the StorEdge CLI (command line interface).

1. **To access the StorEdge CLI, connect to the StorEdge via Telnet, and type "admin" at the [menu] prompt and enter the administrator password.**

2. **To disable ICMP requests to change the default gateway, enter "set default.gateway.redirect no" at the CLI. Press the [Enter] key.**

   After setting any variables on the StorEdge, i.e. anytime the "set" command is used, the command "savevars" must be entered at the command line in order for the settings to persist though future server reboots.

How can I disable RIP routing?

RIP, or Routing Information Protocol, is a method for exchanging routing table information among routers. As this can generate a large amount of traffic in some cases, and the information may not be needed by StorEdge, it may be desirable in some cases to disable RIP on StorEdge.

This functionality is only available at the StorEdge CLI (command line interface).

1. **To access the StorEdge CLI, connect to the StorEdge via Telnet, and type "admin" at the [menu] prompt and enter the administrator password.**
2. **To disable RIP, enter "set routed.active no" at the CLI. Press the [Enter] key.**

   After setting any variables on the StorEdge, i.e. anytime the "set" command is used, the command "savevars" must be entered at the command line in order for the settings to persist though future server reboots.

StorEdge is only reachable from systems on the local subnet.

This is a clear indication of a routing problem. The clients are able to reach StorEdge, but StorEdge cannot successfully reply, because it has no route back to the client.

Possible causes are as follows: incorrect default gateway setting, the default gateway has been overwritten via ICMP or a manually configured TCP/IP route is needed.

To check the default gateway currently in use, proceed as follows:

1. **To access the StorEdge CLI, connect to the StorEdge via Telnet, and type "admin" at the [menu] prompt and enter the administrator password.**
2. **Enter "netstat" at the CLI.**

   The resulting display shows two sections, the local NIC configuration, and then the routing table. Look for a route in the routing table with the destination 0.0.0.0. It is normally the first route. Check the IP address in the "gateway" column. This is the current default gateway.
3. **Next, enter "show inet.gateway".**
4. **Compare this IP address to the actual gateway found above. If they do not match, the gateway has been overwritten.**

   If they do match, check to make sure that this is the correct gateway for your network. If this information is not readily available, check the gateway setting of other systems on the same subnet of other systems who are successfully communicating across subnets.

   The other possibility is that a manually configured route is needed.

Multiple NICs are installed, but all outbound traffic is being sent through a single NIC.

This is the result of configuring more than one network interface on a single subnet. The reason that all the network traffic travels through a single interface is the fact that TCP/IP can only define a single route to each subnet, and this route can only use one network interface.

This unnecessarily limits network bandwidth. The easiest solution is to link the cards at a lower level via port aggregation.

2.14 File System Issues

Note – A full backup should be done before performing the following procedures.

File system inaccessible (mount failure)

Under certain circumstances, volumes may fail to mount. This will typically manifest as an "access denied" message returned to users attempting to access the data on the affected volume.

The first place to look for information on this is the StorEdge system log. The boot log is the beginning of the system log for each boot.

Review the bootlog for messages about the volume mount. The beginning of the mount process can be identified by the "sfs2" process identifying disk segments, as in this example:

```
sfs2: /vol1 - id=3F142D97, extent 1 of 1, version 0
```

A successful mount is indicated by the "<volumename> is complete" message, as in this example:

```
/vol1 is complete
```

If the current boot log indicates a successful mount, then the problem lies elsewhere, most likely in the area of Windows Domain security.

The most recent boot log is named "bootlog". The next most recent is named "bootlog.1", the naming convention continues on to the oldest file "bootlog.9". Beginning with /cvol/bootlog, search through the files for the most recent successful mount of the volume. Then, check the following bootlog to locate the first occurrence of the mount problem.

The most common reason for the mount failure is a power loss or crash during the mount process. This causes a "mount failure" flag to be set on the volume. When StorEdge encounters this flag upon boot, it does not mount the volume, and logs the following message:

```
/vol1 not mounted, previous mount did not complete.
```

Before making this diagnosis, it is very important to ensure that you are checking the bootlog containing the first unsuccessful mount attempt. Regardless of what the original problem is, the second attempt and all subsequent unsuccessful attempts to mount the volume will always log the “previous mount did not complete” message.

If you are 100% certain that the problem is an interrupted mount attempt, you should be able to correct the problem by entering “mount –f <volumename>” at the CLI. Check the system log to ensure that the mount is successful, and that there are no errors.

The above issue is the most common reason for mount failure, and a relatively minor problem. However, please be aware that other mount problems can be much more severe. If the mount failure occurs for any reason other than the above, or if the above solution does not work, the issue must be escalated. Do not use the “mount –f” command in an attempt to resolve a mount problem other than the one listed above.

Information required for escalation: A diagnostic email with all attachments should be sufficient for escalation. Verify that the boot logs are included with the diagnostics, and that they contain the mount attempt. Also, if possible, include information on the circumstances surrounding the mount problem and any attempts to correct it. Find out whether remote access is available to the site, and if appropriate, provide the necessary details to accomplish this.

Can’t write to file system

This is caused by a problem found at mount time. The problem can vary in degree, and the causes for this vary widely. This problem manifests as a complete inability to write for all users, including root.

The following are the messages that you will see in this case

```
/vol1 mounted read-only due to errors, run check
```

```
/vol1 is complete
```

The message “run check” refers to the filesystem check, or “fsck”. The filesystem check is a time consuming process, which is potentially destructive to data. This being the case, data should now be collected for escalation. Involving the engineering group at this time may make it possible to circumvent the lengthy volume repair process.

A diagnostic email, with all attachments, is required to escalate this type of issue. The primary source of information for this case is the bootlog. Also, collect as much information as possible about the circumstances surrounding the failure, e.g. did the system lose power, what symptoms were seen by the clients, was any hardware or

RAID configuration changed. If syslogd logging was enabled, these results should be included as well. Find out whether remote access is available to the site, and if appropriate, provide the necessary details to accomplish this.

After reviewing the case, engineering may make specific recommendations and modifications, or they may recommend that you proceed with the filesystem repair. For instructions on how to complete a filesystem repair, see the heading "Filesystem check procedure" under Diagnostic Procedures at the end of this document.

File system related error messages

The following are examples of messages that indicate a problem with the filesystem:

```
Broken directory, run check
```

```
Corrupted/run check
```

```
Cleaner error
```

These messages are generated when unsuccessful attempts are made to access particular files or directories. The message may occur only when accessing a particular directory; or it may fill the system log entirely. You can get a general idea of the degree of the problem by the frequency of the error messages.

This problem manifests as an inability to access particular files and folders, which often causes a hang condition or a timeout.

The message "run check" refers to the filesystem check, or "fsck". The filesystem check is a time consuming process, which is potentially destructive to data. Generally speaking, unless you have a very good idea about the source of the problem, the issue should be escalated. It is almost a certainty that a filesystem check will be needed in this case. However, it is important to work toward locating the source of the problem to ensure that it will not recur.

A diagnostic email, with all attachments, is required to escalate this type of issue. The primary source of information for this case is the system log. Also, collect as much information as possible about the circumstances surrounding the failure, e.g. did the system lose power, what symptoms were seen by the clients, was any hardware or RAID configuration changed. These particular log messages always correspond to an attempt to access a particular directory or file, this may provide a clue to the source of the problem.

If syslogd logging was enabled, these results should be included as well. Find out whether remote access is available to the site, and if appropriate, provide the necessary details to accomplish this.

After reviewing the case, engineering may make specific recommendations and modifications, or they may recommend that you proceed with the filesystem repair. For instructions on how to complete a filesystem repair, see “Filesystem check procedure” under Diagnostic Procedures at the end of this document.

Reoccurrence of filesystem related error messages / mount problems after repair

If you have run a filesystem check until no errors were reported, or recreated a volume, this should permanently resolve the filesystem errors. If the errors return, the source of the problem remains. The most likely source is a hardware problem. A good first step is to replace the system board memory and the RAID controller, or failing that, the entire system. Once the source of the problem has been resolved, it will be necessary to proceed according to the “Filesystem check procedure” under Diagnostic Procedures at the end of this document.

Checkpoint database problems reported in system log Can’t delete checkpoints

The indication of a checkpoint database problem is either a hard error (e.g. cannot write) in the system log when attempting to delete a checkpoint, or an error message which specifically states “error in checkpoint database”. As the checkpoint filesystem is read-only, and treated as a separate filesystem in many ways, this problem must be addressed at the filesystem level. Specifically, via the chkpntabort command and a file system check.

It is generally recommended that this issue be escalated for assistance in accurately identifying the problem, and also to locate the source of the problem. The messages can vary considerably from the above; and similar checkpoint related messages could lead one down the wrong path toward applying an unnecessarily severe solution.

A diagnostic email, with all attachments, is required to escalate this type of issue. The primary source of information for this case is the system log. The diagnostic should be captured as close as possible to the time the messages occur, so that they may be seen in context in the system log. Also, collect as much information as possible about the circumstances surrounding the failure, e.g. when did the messages first appear, what was happening at the time, symptoms reported by users.

Typically in this case, it is necessary to abort checkpoints on the volume. This is done from the CLI. After verifying the diagnosis with engineering, access the CLI and enter “chkpntabort <volumename>”. StorEdge will prompt for confirmation. Answering “y”, “yes” to the prompt will result in the immediate deletion all checkpoints. A file system check is required as soon as possible after aborting checkpoints. It may be necessary to disable the “Use checkpoints for backup” option in the StorEdge Volume Configuration screen in order to perform the prerequisite backup.

2.15 Drive Failure Messages

Note – Check the WebAdmin and system log to ensure the drive rebuild is completed before performing the following procedures.

The light on one of the hard drives is red.

Check to ensure the drive rebuild has completed in the syslog.

This is an indication of a failed drive. Before performing any work on hard drive and the drive subsystem, verify that an accurate and full system backup is available. If a backup cannot be confirmed, then backup system immediately.

LCD panel may display message R11 Drive failure

Controller alarm may be beeping

Replace the drive immediately.

Drive failure in log or an email received stating drive has failed or is about to fail.

Before performing any work on hard drive and the drive subsystem, verify that an accurate and full system backup is available. If a backup cannot be confirmed, then backup system immediately.

LCD panel may display message R11 Drive failure

A red light may or may not exist on this type of error. Verify location via the log or email message. The message will state what slot the drive that is failing is located at (drive failure slot 9).

Controller alarm may be beeping

Replace the drive immediately.

A drive has failed how do I replace it?

Failed drives are usually evident by the following:

Red light is on the drive.

Log will display “Failed drive at slot #”.

Diagnostic email will list drive as failed.

LUN will be reported as degraded in the log and in the diagnostic email.

Controller alarm may be beeping.

Some or all of these symptoms may exist. The process is the same regardless of how the failed drive is reporting.

This functionality is only available through the StorEdge Web Admin.

1. **To access these settings, log in, and navigate to RAID/Manage RAID.**

   The screen may take several minutes to display, as the StorEdge must scan the RAID subsystem for configuration data.

   This screen displays the drives, RAID sets and LUN configuration.

   Functional drives have a green LED displayed beside them. The failed drive will have a red LED. Replace the failed drive.

   After a few minutes, the display will change. The replaced drive will display a yellow LED, with status “new drive added not configured”.

   RMV LUN, ADD LUN, RMV HS and Add HS buttons will be grayed out.

2. **Select the Rebuild button and the system will start a LUN rebuild.**

   The time for completion is dependent on drive size, system load and LUN size. The average is approximately 2-3 hours to complete a rebuild.

   Alarm will continue to beep until rebuild is complete

   When complete, the LED displayed beside the drive will change to green, with a status of “online”.

Log message: LUN critical.
Email alert: LUN critical.

This is an indication that the system has a failure in the disk subsystem, and that there is no spare drive. The term “Critical” is used because another drive failure would result in failure of the LUN and all the data contained there.

Verify the following:

Are there any red lights on any hard drive in StorEdge?

Are there any messages in the log other than “LUN Critical”? (Drive failed etc.)

Controller alarm may be beeping.

Follow instructions on drive replacement. Once the drive has been replaced, system will rebuild the drive and the LUN will go from critical to online.

Alarm will continue to beep until rebuild is complete. The alarm can be silenced from the RAID page of the GUI or from the menu.

2.16 File and Volume Operations

StorEdge doesn’t allow creation of volumes larger than 256 GB.

The StorEdge volume creation screens allow you to create volumes up to 256GB. In order to create larger volumes, segments (also limited to 256GB) must be created and joined to the volume.

Free space not immediately available after delete.

By default, the StorEdge will delete files as a background process and frees blocks as this completes. This is done to provide better performance to foreground processes. If many files have been deleted and the system is very busy, it may take some time. This functionality is provided by the .attic$ directory feature.

The .attic$ directory can be found at the root of each volume. In rare cases, on very busy filesystems, the .attic$ directory can be filled faster than it can process deletes. This leads to a lack of free space and slow performance. In these cases, disabling the .attic$ directory is recommended.

This is configurable only at the Command Line Interface (CLI).

1. **To access the StorEdge CLI, connect to the StorEdge via Telnet, and type “admin” at the [menu] prompt and enter the administrator password.**

2. **At the CLI, enter “fsctl attic disable <volumename>”.**

   This must be done for each volume that requires the change. It is not necessary to delete the directory, but it is permissible to delete the files within the directory to reclaim the disk space.

After deleting files, volume free space remains the same.

The most likely cause of this is the checkpoint feature. Checkpoints store deleted and changed data for a defined period of time so that customers can retrieve deleted files and prior versions for data security. This means that the data is not removed from disk until the checkpoint is expired, a maximum of two weeks. If you are deleting data to free disk space, it will be necessary to remove or disable checkpoints.

Can’t delete a file.

If a file cannot be deleted, the first thing to check is security. Make sure that the deleting user has the appropriate rights. Files can also be deleted from the command line. For details on this procedure.

In some cases, a “file not found” message will be generated when attempting to delete a file, even from the command line. Usually, this is due to an invalid character in the filename that cannot be parsed. To delete such a file, use the “\” character. Use any of the command line utilities, and proceed the illegal character with “\”. An example of an illegal character would be a colon, a comma or a double quote.

An exception to this is the “$” character in a filename. In this case, the character must be preceded with another “$” in order for it to be accessible via command line operations.

There is one more item to be aware of when referencing filenames at the CLI. If the filename contains one or more spaces, the entire argument must be enclosed in quotes.

Example: del “/vol1/sales/my big file name”

Can’t delete files from /etc directory.

This is by design as many of these files are required for proper system operations. There may some times that files must be deleted, especially when configuring host files and user files manually. To enable deletion of files from the etc directory the immutable bet must be set.

From the command line type in “cleari /voname/etc/”

Now files can be deleted. Caution must be exercised it is not recommended to delete files from the etc.

Log message “mbtowc[0xXX]: invalid first byte”, or “: invalid sequence”.

This message is generated when StorEdge receives a filename or network name with a character that is unreadable. ASCII (plain text) characters are expressed with a byte value of 0x7F or below. Values above this range are expected to be Unicode encoded, per the UTF-8 specification. The encoding requires multiple bytes per character. In order for the character to be valid, the first byte must be in the range 0xC0 through 0xFD. All other bytes in the sequence must be in the range 0x80 to 0xBF. When the character is out of these ranges, the invalid first byte or invalid sequence message is generated.

The most likely source of this is a client sending non-ASCII names that are not Unicode encoded. It can also be an indication that NFS clients are not properly set up for Unicode.

2.17 Administration Interfaces

Can't run GUI, Java certificate expired message.

The Java certificate for the StorEdge Web Admin is valid for a fixed period of one year. To stop this message, please contact Technical Support for an operating system upgrade.

Can't run StorEdge Web Admin or some screens incorrect.

The following are the known issues that may interfere with the operation of StorEdge Web Admin.

Verify your browser version. StorEdge Web Admin requires Internet Explorer 5.5 or newer, or Netscape 4.77 or newer, with the exception that Netscape 6.0x is not supported.

Check for a proxy server. If your site uses a proxy server, disable proxy for local IP addresses or disable proxy altogether.

Clear your browser cache. Delete all files in your browser cache. Check your browser documentation for instructions on this procedure. This is usually a problem after a StorEdge operating system upgrade, as outdated pages may be saved in cache.

Check whether workstation is on same subnet as StorEdge.

Verify Java client version. Version 1.3.1.1 or newer of the Java client is required. If no Java Client is installed on the client connecting to the StorEdge, you will normally be prompted for installation. If not, the client can be found at http://www.java.com. If you receive error messages related to the Java Client (or JRE), uninstall and reinstall the client.

If problems continue, one of the most helpful troubleshooting steps is to try another workstation or two. These problems are often client-specific, and this step often leads to the resolution.

StorEdge Web Admin missing ICONS or GUI does not respond

The GUI interface requires at least version 1.3.1.1 of the Java plug-in to be installed on the administrator client machine and should automatically download the plug-in when needed.

We have seen that sporadic GUI behavior, missing icons or repeated messages requiring a re-certification is often an indicator of a damaged Java plug-in.

1. **From Internet Explorer select "Tools" then "Internet options".**
2. **Select "Settings" then "View objects".**
3. **Look for Java runtime version.**
4. **In the status area, verify that it is not damaged.**
5. **If it says Damaged or if it has uninstalled even after the install has completed then remove the Java client from workstation and reinstall.**

Web GUI login failure.

First, try clearing your browser cache. For Internet Explorer, this can be found under Tools/Internet options.

Next, try resetting the password via Telnet. Proceed as follows:

1. **Access the StorEdge via Telnet.**
2. **When the menu prompt appears enter "menu" press the enter key.**
3. **Enter the administrator password for the StorEdge. Press the enter key.**
4. **Select "K", "Admin access".**
5. **Enter password information. (It can stay the same no need to change password.)**
6. **Enter "7", "Save changes".**

Web GUI session aborted while performing administration.

For integrity reasons, only one user is permitted in the StorEdge Web Admin at a given time. A second connection to the Web Admin will terminate the first. This implementation is necessary to allow recovery from client hangs.

How do I reset the StorEdge administrator password?

The management administrator password can only be reset via a direct connection to the StorEdge. This will require that a keyboard and display be connected to the rear of the StorEdge.

1. **Access the StorEdge via keyboard or serial console.**
2. **Type [Enter] at the [menu] prompt and enter the administrator password.**
3. **Select option "K", "Admin access" in the Access control Section**
4. **Enter "Y" to enable admin password**
5. **Enter the new password**
6. **Enter the new password again for verification**
7. **Enter "7", "Save changes".**

StorEdge Web Admin does not work properly through a firewall.

In order for the StorEdge Web Admin to be used across a firewall, the following ports must be opened:

If using http, port 80/TCP and UDP.

If using https, port 443/TCP and UDP.

For the portmapper service, port 111/TCP and UDP

One additional port is required for Java communications; determine the needed port as follows:

The rpcinfo program can be used on a UNIX client to determine the port required by the java package. If only Windows clients are available, a third party rpcinfo utility will be required.

Run rpcinfo and search for program number 805898577. The program number is the left most column.

Example:

```
PROMPT> rpcinfo -p <StorEdge-hostname>

programversprotoport

1000002tcp111portmapper

1000002udp111portmapper

...

8058985771udp693webadmin
```

In the above example, UDP port 693 would have to be opened. The port is always in the range 600 to 1023 but may vary based on system parameters.

The keyboard arrows keys do not respond properly when using telnet to the StorEdge

The arrow keys generally do not work within StorEdge Telnet menus, and often cause an immediate exit from the current screen.

Use only the following keys within the Telnet menus:

[Backspace] or [Del]Deletes the previous character.

[Ctrl]+UDeletes the entire current field.

[Esc]Exits the menu with no change.

[Enter] or [Tab]Enters data and proceeds to the next field.

If no data entered, proceeds to the next field with no change.

Having problems after software upgrade, how do I return to the previous version?

1. **Access the StorEdge via Telnet.**
2. **Press [Enter] at the [menu] prompt and enter the administrator password.**
3. **Enter "0", "Shutdown" in the operations Section.**
4. **Enter "P", "Boot previous version".**
5. **At the verification screen, enter "Y", "Yes" to continue or [Esc] to cancel.**

   System reboots and loads previous version of software

   This functionality is also available from the StorEdge Web Admin.

6. **Navigate to System Operations/Shut down the server.**
7. **Select the "Reboot previous version" radio button.**
8. **Click Apply.**

After software upgrade and reboot, system appears to hang, LCD displays "…booting".

The software upgrade process may take as long as five minutes. No indication of progress is available. This is standard, the system will boot normally in a few minutes. If the process takes longer than five minutes, connect a VGA display to StorEdge to check status.

Software Upgrade is not working. The file transfer says it's complete and then I reboot, but the software does not go to the new version. After the file transfer has completed, telnet to the Sun StorEdge 5310 NAS and type the following commands.

```
Load unixtools
```

```
ll /cvol/*
```

Check the image name and make sure that the file is named .img and should look something like this;

```
nf420b149.img
```

Some versions of WINZIP will change the file to a .zip extension and the Sun StorEdge 5310 NAS will not use it for the upgrade.

2.18 StorEdge Features and Utilities

StorEdge fails to create scheduled checkpoints
Previously created checkpoints are missing.

The checkpoint feature requires some disk space to operate. In addition, StorEdge (and other servers) perform much better with filesystems that have some free space to work with. As a rough guideline, you should consider 70% the maximum file volume utilization. It may be necessary to consider adding storage earlier than this in a performance sensitive environment.

Once a StorEdge volume reaches 90% disk space utilization, StorEdge will cease to create scheduled checkpoints. Once the volume reaches 95% disk space utilization, StorEdge will delete checkpoints, beginning with the oldest.

Another possible reason for checkpoint creation failure is that the checkpoint limit for a particular volume has been reached. A log message similar to the following will be recorded when StorEdge cannot create a checkpoint:

```
7/30 22:00 I ndmpd[327]: ENOSPC No space left on device
chkpbkup prepare vol1 (create)
```

This message can indicate that the disk space threshold (90%) has been exceeded, or that the limit of 16 checkpoints for a volume has been reached.

To check disk space utilization on the StorEdge, proceed as follows:

1. **Connect to the StorEdge via Telnet, and type "admin" at the [menu] prompt and enter the administrator password.**
2. **At the CLI, enter "df".**

StorEdge fails to send diagnostic email

SMTP (email) configuration allows StorEdge to send diagnostics and urgent notifications directly to your mail server.

First, check the system log. Most SMTP problems can be identified precisely by checking the system log.

1. **To do so, access the StorEdge via Telnet.**
2. **Type "admin" at the [menu] prompt and enter the administrator password.**
3. **At the CLI, enter "menu".**
4. **Select option "2", "Show Log".**

   The fourteen most recent syslog messages are displayed.
5. **If email related messages are not found, select option "1", "Show Entire Log".**
6. **Browse through the log one screen at a time, looking for any messages relating to email issues.**
7. **If no messages are found, attempt to send a diagnostic again.**

   All attempts to send email, successful or not, generate a log message.

The following are some of the most common email related messages and their causes. Note that the messages may vary slightly, depending on the mail server.

```
smtp: Could not send mail, err=-1
```

This means that the email could not be sent. It is always accompanied by a more specific error.

```
Error in tcp_open for servername
```

This means that StorEdge could not open a TCP connection to the configured mail server. Possible reasons are name resolution, incorrect IP, IP unreachable due to network problem. To correct this issue, enter the mail server by IP address, and make sure the IP address is correct and reachable. The ping command from the StorEdge CLI may be helpful in this case, both by name and IP address.

```
Unknown code <501 Syntax error, parameters in command "MAIL
From: <servername@>" unrecognized or missing.
```

The key here is the “from” address ending in the “@” character. The SMTP server usually refuses email from StorEdge unless a DNS domain is configured, dependent on mail server configuration. This is true whether you are using DNS or not.

```
Unknown code <501 This MTA is configured NOT to relay message
from [servername] to [domain.com].
```

This message indicates that the mail server is configured not to relay messages from other SMTP senders (i.e. StorEdge) to outside domains. The solution to this is to remove all email addresses not in the local DNS domain. The email messages may then be forwarded as needed.

2.19 Hardware Warning Messages

Hardware replacement procedures can be located in Chapter 8, “FRU Replacement Procedures” in this book.

Log message: Controller write-back cache is disabled.
Log message: System on battery backup.

The power source has become unstable, and StorEdge now writes data to disk before confirming writes to clients, in case of complete power loss.

Specifically, this is an indication the RAID cache has been disabled. This message is caused by either an AC power failure or a discharged UPS. If there is an obvious indication of AC power loss that could last for some time, client systems should be disconnected and shut down.

If AC power has failed, the condition corrects itself after AC power is restored.

If AC power has been restored for 30 minutes, and the UPS is not charged completely, check the UPS battery.

Log message: Low Battery.

After StorEdge receives the Low Battery notification from the UPS, this message is logged and the shutdown process is initiated in order to protect customer data.

Log message: Controller write-back cache is enabled.

System AC power and UPS have returned to a reliable state. Write-back cache is enabled.

Log Message: Blower Fan has failed.

This is an indication that one of the StorEdge cooling fans has failed.

Verify that one or more of the following indicators are present:

The unit system status on the front of the StorEdge is a solid red.

The system status light on the back of the StorEdge is a solid red.

The log should have a message “Blower fan # has failed”.

Front panel may have message in display “P11 Fan # has failed”.

The # is number of the fan 1, 2, 3, 4 all four fans are on the fan enclosure assembly located inside the StorEdge unit.

The failed fan # should be lit on the fan assembly.

Contact Technical Support to replace the failed fan.

UPS Messages

If a UPS is connected to the Sun StorEdge 5310 NAS when booted, you will see the following message in the log:

```
07/08/04 10:48:34 I sysmon[51]: UPS: Smart mode set.
```

If the power to the Sun StorEdge 5310 NAS fails and you have a UPS connected, you will see the following message in the log:

```
07/08/04 11:00:30 E sysmon[51]: UPS: AC power failure. System
is running on UPS battery.
```

And then the following message.

```
07/08/04 11:00:38 I sysmon[51]: Ctlr0: write-back cache
disabled
```

If the power is returned to the UPS, you will see the following message:

```
07/08/04 11:01:09 I sysmon[51]: UPS: AC power restored
```

Front LCD panel message: "P21 Power 1 Failed".

This is an indication that one of the two power supplies has failed.

Verify that one or more of the following indicators are present:

The system status light on the front of the StorEdge is flashing green.

The system status light on the back of the StorEdge is flashing green.

The log should have a message "power supply unit 1 has failed"

Front panel may have message in display "P21 power 1 has failed"

The StorEdge has 2 power supplies located in the back of the unit. The failed power supply may have a solid red LED.

If both power supply lights are green, see hardware reference guide for unit replacement locations.

Replace the failed power supply.

The light on one of the power supplies is amber.

This is an indication of a failed power cord. Verify the following:

Verify that one or more of the following indicators are present:

The system status light on the front of the StorEdge is solid amber.

The system status light on the back of the StorEdge is solid amber.

Verify that the cord is plugged into an AC source.

Reseat the power cord in the power supply.

Replace the power cord if reseating does not fix the problem.

Replace the power supply if power cord replacement does not work.

What do the Status LED Indicators on front panel indicate?

LED Status indicators at the front panel signal current activities taking place in the system..

TABLE 2-17 Status LED Indicators

Power LED	A continuous green LED indicates the system is powered on. No light indicates the system is off.
Built-in NIC 1 LED	A green LED indicates network activity via the built in NIC port 1.
Built-in NIC 2 LED	A green LED indicates network activity via the built in NIC port 2.
Hard Drive Status LED	➧ A random blinking green LED indicates hard drive activities. ➧ A continuous amber light indicates a hard drive fault. ➧ No light indicates no activities or faults.
System Status LED	➧ A continuous green LED indicates the system is in normal operation. ➧ A blinking green LED indicates the system is operating in a degraded mode. ➧ A continuous amber LED indicates the system is in a critical or nonrecoverable condition. ➧ A blinking amber LED indicates the system is in a non-critical condition. ➧ No light indicates the system is halted assuming the power LED is green.
Flashdisk Activity	➧ A random blinking green LED indicates flashdisk activity. ➧ A continuous amber LED indicates the system is in a critical or nonrecoverable condition.
System ID LED	➧ A continuous blue LED indicates the ID button is depressed. This light is intended to help identify this chassis among several. It can also help to illuminate the rear of the chassis for service. ➧ No light indicates the ID button is not depressed.

2.20 Backup Issues

Tape library not recognized.

Make sure the tape drive is on the list of supported tape units.

SCSI ID of tape library should be higher than the tape drive. Set ID of library to 0, ID of tape drive to 5.

Does the SCSI card recognize the drive on system boot up?

See if card can talk to tape drive. On boot get into SCSI card BIOS.

Run scan utility.

If no device is found check cables, termination.

Try another tape drive.

Network backup fails due to .attic$ directory.

The .attic$ directory is a StorEdge system directory at the root of each volume. Some third party backup software has trouble with this directory. The simplest solution to this is to configure the software to ignore this directory.

We have encountered some backup software that is not capable of ignoring the presence of this directory to the extent necessary to successfully backup the StorEdge. In this case we recommend disabling the .attic$ directory.

This is configurable only at the Command Line Interface (CLI).

1. **To access the StorEdge CLI, connect to the StorEdge via Telnet, and type "admin" at the [menu] prompt and enter the administrator password.**
2. **At the CLI, enter "fsctl attic disable <volumename>".**

This must be done for each volume that is to be used with this backup software. It is not necessary to delete the directory, but it is permissible to delete the files within the directory.

NDMP backup fails: access denied message.

NDMP software must authenticate to the StorEdge in order to backup files and directories. Each NDMP software solution has a place to configure a username and password for a device. For StorEdge, the username is “administrator”, and should be accompanied by the console password.

NDMP: Can’t browse backup history.

Certain NDMP backup utilities are not able to browse backup history without a configuration change on the StorEdge. To make this configuration change, it is necessary to access the StorEdge CLI.

This behavior is configurable only at the Command Line Interface (CLI).

1. **To access the StorEdge CLI, connect to the StorEdge via Telnet, and type “admin” at the [menu] prompt and enter the administrator password.**
2. **At the CLI, enter “ndmp set dump.pathnode=yes”.**
3. **Then “ndmp set tar.pathnode=yes”.**
4. **Then “ndmp save”.**

NDMP incremental/differential backups back up all files.

Some Windows applications, such as virus scanning software, update timestamps on all files that they scan. NDMP software uses this same timestamp to identify whether the file has been backed up since the last full backup. As a result, NDMP sees all these files as modified, and effectively performs a full backup instead of the desired incremental or differential.

To avoid this problem, it is possible to configure StorEdge to ignore this timestamp when performing NDMP backup.

To access the StorEdge CLI, connect to the StorEdge via Telnet, and type “admin” at the [menu] prompt and enter the administrator password.

From the command line enter:

```
ndmp set ignore.ctime=yes

ndmp save
```

2.21 Direct Attached Tape Libraries

Tape library not recognized.

Check the following settings:

- Make sure the tape drive is on the list of supported tape units.
- Set SCSI ID of tape library to 0, ID of tape drive to 5.
- Does the SCSI card recognize the drive on system boot up?
- See if card can talk to tape drive. On boot get into SCSI card bios.
- Run scan utility.
- If no device is found check cables, termination.
- Try another tape drive.

Network backup fails due to .attic$ directory.

The .attic$ directory is a StorEdge system directory at the root of each volume. Some third party backup software has difficulty processing this directory. The simplest solution to this is to configure the software to ignore this directory.

We have encountered some backup software that is not capable of ignoring the presence of this directory to the extent necessary to successfully backup the StorEdge. In this case we recommend disabling the .attic$ directory. To do so, access the Command Line Interface (CLI) as follows:

1. **Connect to the StorEdge via Telnet, serial console, or keyboard console.**
2. **Type ?admin? at the [menu] prompt and enter the administrator password.**
3. **At the CLI, enter ?fsctl attic disable <volumename>?**

   This must be done for each volume that is to be used with this backup software. It is not necessary to delete the directory, but it is permissible to delete the files within the directory.

NDMP backup fails: access denied message.

NDMP software must authenticate to the StorEdge in order to backup files and directories. Each NDMP software solution has a place to configure a username and password for a device. For StorEdge, the username is ?administrator?, and should be accompanied by the console password.

Local backup or restore fails with ?PNReduce error? in log.

This message indicates that StorEdge could not read the pathname (PN) provided for backup or restore. The local backup and restore utilities require a full path, they are case sensitive, and they do not allow the use of wildcards (?*? or ???). Specifying a directory causes it to be backed up in its entirety.

NDMP: Can?t browse backup history.

Certain NDMP backup utilities are not able to browse backup history without a configuration change on the StorEdge. To make this configuration change, access the StorEdge CLI, as follows:

1. **Connect to the StorEdge via Telnet, serial console, or keyboard console.**
2. **Type "admin" at the [menu] prompt and enter the administrator password.**
3. **At the CLI, enter "ndmp set dump.pathnode=yes"**
4. **Type "ndmp set tar.pathnode=yes"**
5. **Type "ndmp save"**

NDMP incremental/differential backups back up all files.

Some Windows applications, such as virus scanning software, update timestamps on all files that they scan. NDMP software uses this same timestamp to identify whether the file has been backed up since the last full backup. As a result, NDMP sees all these files as modified, and effectively performs a full backup instead of the desired incremental or differential.

To avoid this problem, it is possible to configure StorEdge to ignore this timestamp when performing NDMP backup. To do so, proceed as follows:

1. **Connect to the StorEdge via Telnet, serial console, or keyboard console.**
2. **Type "admin" at the [menu] prompt and enter the administrator password.**
3. **Enter "ndmp set ignore.ctime=yes"**
4. **Enter "ndmp save"**

2.22 Frequently Asked Questions

This section addresses frequently asked questions for the Sun StorEdge 5310 NAS. The section contains these topics:

- "CIFS/SMB/Domain Issues" on page 2-92
- "NIS/NIS+ Issues" on page 2-104
- "TCP/IP and Network Configuration" on page 2-106
- "Quota Configuration" on page 2-109
- "Checkpoint Configuration" on page 2-115
- "Volume Creation and Expansion" on page 2-120
- "Reserved Filesystems and Directories" on page 2-123
- "NFS Issues" on page 2-124
- "Administration Interfaces and Utilities" on page 2-128
- "Backup and Migration Issues" on page 2-142
- "Macintosh Connectivity" on page 2-146
- "Miscellaneous Log Messages" on page 2-147
- "Direct Attached Tape Libraries" on page 2-148
- "StorEdge File Replicator" on page 2-149

2.23 CIFS/SMB/Domain Issues

How do I configure local SMB groups?

A group is a named collection of users. In a Windows domain environment there are two types of groups: global groups and local groups. A global group is visible to any computer participating in a domain. A local group is defined on an individual computer and may contain user accounts created on the local computer as well as user accounts and global groups from the domain to which the local computer belongs or trusts.

To access StorEdge local group configuration, proceed as follows:

1. **Access the StorEdge via Telnet or serial console.**
2. **Press [Enter] at the [menu] prompt and enter the administrator password.**
3. **Press the spacebar until "CIFS/SMB Configuration" is displayed under "Extensions" at the lower right.**
4. **Select the letter corresponding to "CIFS/SMB Configuration".**

5. **Select the letter corresponding to "Local Groups".**

6. **This takes you to a menu where you will see a list of all currently configured groups. By default, in Domain mode, the "Administrators", "Power Users" and "Backup Operators" local groups exist.**

   To add a group, press "8", Add a Group, from this screen. To edit group settings, or to delete a group, press the letter to the left of the group name. Please note that the "Administrators" group is always understood to contain the Domain Admins group of the configured Windows domain, though it is not listed. This behavior is expected by Windows Domains.

   Inside the edit menu, there are five options listed at the bottom of the screen, as follows:

   - "1", "Fields": Selecting this option allows you to edit the name of user-defined groups, or the comment field for any group. The names of default groups cannot be changed.
   - "2", "Members": Selecting this option allows you to add members to the group. Inside the submenu, select option "8", "Add" to bring up a text box. For a user in the configured domain, simply enter the username. For a user in a trusted domain, enter the name in the format <domainname\username>. In order to successfully add the account, the account must already exist, and any required Windows Domain trust relationship must already be in place.
   - "3", "Privileges": Selecting this option allows you to configure privileges for the selected group. In this submenu, there is a list of security privileges, as follows: Take ownership of files or other objects, Back up files and directories, Restore files and directories. These correspond to standard Windows privileges. Press the letter corresponding to the privilege to enable or disable it.
   - "8", "Delete Group": This deletes the local group. This option is only valid for user-defined groups.
   - "0", "Cancel": This returns you to the list of groups.

7. **After configuring desired options, select option "7", "Save Changes".**

   This functionality is also available through the StorEdge Web Admin.

1. **To use the Web Admin, connect with a Web browser to http://<hostname or IP address of your StorEdge>.**

2. **Click "Grant" or "Yes" to accept any Java software authorization windows and you will reach the login screen.**

3. **Type the administrator password to access the administration interface.**

4. **Navigate to Windows Configuration/Configure Groups. All of the options explained above are available in this menu.**

Note – Workgroup mode refers not only to the lack of domain membership, but the use of share-level security.

For more information on this topic, refer to the *Sun StorEdge 5310 NAS Software Installation, Configuration, and User Guide*.

How do I share files with SMB users?
How do I create SMB shares?

To share files via SMB, shares must be created. A share allows access to a particular location in the directory tree. To access this functionality, access the StorEdge via Telnet or serial console.

1. **Press [Enter] at the [menu] prompt and enter the administrator password.**
2. **Press the spacebar until "CIFS/SMB Configuration" is displayed under "Extensions" at the lower right.**
3. **Select the letter corresponding to "CIFS/SMB Configuration".**
4. **Select the letter corresponding to "Shares".**

This will bring up a list of all existing shares, ten per page. There is a default share created for each volume, which is only accessible to members of the Domain Admins group from the configured Windows Domain. If there are more than ten, select option "1" and option "3" to move through the pages. To add a share, select option "8". To edit share settings press the letter corresponding to an existing share.

The "add" and "edit" options bring up the SMB/CIFS share setup menu. There, you will see a list of options for each share as follows:

- Share name—Name of share
- Directory—Full directory path shared, including volume name
- Comment—Optional comment field, displayed in browse list
- ADS Container—For Active Directory only, name of container to publish shares to.
- Macintosh Extensions
- Desktop DB—Only used when Mac SMB clients are connected.
- Password Protection—In Workgroup mode, this enables or disables share-level security.
- Access Password—In Workgroup mode, shares are secured by password only.
- Read/write—In Workgroup mode, this password allows read/write access.
- Read-only—In Workgroup mode, this password allows read/only access.
- User ID—In Workgroup mode, files written via this share are owned by this UID.
- Group ID—In Workgroup mode, files written via this share are owned by this UID.

- Umask—In Workgroup mode, these NFS permission bits will be cleared.
- (when creating new files.)

Workgroup mode settings are ignored when Windows Domain Security is enabled. Workgroup mode on the Sun StorEdge 5310 NAS also implies use of what Microsoft calls "Share-level Security."

This functionality is also available through the StorEdge Web Admin.

1. **To use the Web Admin, connect with a Web browser to http://<hostname or IP address of your StorEdge>.**
2. **Click "Grant" or "Yes" to accept any Java software authorization windows and you will reach the login screen.**
3. **Type the administrator password to access the administration interface.**
4. **Navigate to Windows Configuration/Configure Name Service. All of the options described above are available.**

Note – It is recommended that you avoid sharing user data at the root of a volume. Best practice is to create subdirectory structures and share these. This eases security administration, and removes the need to hide or secure system folders.

How do I create hidden shares?

Create an SMB share with a name ending in "$". The share will be accessible by name, but will not appear in network browse lists.

What are the default SMB shares?
How do I use SMB administrative shares?

For each volume, a default share is created for the root directory. These shares are known as "administrative shares".

This behavior is expected by Windows Domain security. The share name is a single letter, followed by "$". The first volume, /cvol, is associated with the share c$, and the first user volume created is associated with the share e$. The "$" character at the end of the name causes these shares to be hidden from the network browse list.

Only members of the local Administrators group have access to these shares. Please note that the Domain Admins global group from the configured Windows Domain is always a member of this local group.

What does the umask setting do?

The umask setting, allows the permissions of new files and directories to be specified on a per-share basis, which is consistent with the per-share UID and GID specification.

A umask is a file creation mask. It defines the permission bits to turn off when creating a file. Bits that are set in the umask are cleared in the mode of a newly created file. The umask is defined in octal because octal numbers comprise three bits, which maps easily to the UNIX file permission representation, for example; the UNIX permissions rwxr-xr-- can be represented as 754.

The umask is applied on a per share basis using standard UNIX rules, with the exception of the DOS read-only attribute. If the DOS read-only attribute is set in the file creation request, all write bits will be removed from the permission mode after the umask has been applied.

How do I set up user and group credential mapping? How do I share files between NFS and SMB users?

In Windows Domain Security Mode, this is accomplished via user and group mapping. Every time a Windows user accesses the StorEdge for the first time, a new user mapping is created. Similarly, a new group mapping is created the first time each user from a particular Windows primary group logs in.

These mappings permanently associate the Windows user's ID with a particular NFS UID. This allows users with both NFS and SMB accounts to access their own data from either type of client, and to share data with heterogeneous workgroups. The mapping rules determine how the NFS UID or GID for a particular Windows user or group is obtained.

It is strongly recommended that you define a mapping rule and import NFS accounts to StorEdge prior to the migration of data. This will minimize the amount of manual configuration required.

The primary tool to accomplish this is the selection of user and group mapping rules. This can be accomplished via the Web Admin, or via the CLI. Each method is detailed below.

1. **To use the Web Admin, connect with a Web browser to http://<hostname or IP address of your StorEdge>.**
2. **Click "Grant" or "Yes" to accept any Java software authorization windows and you will reach the login screen.**
3. **Type the administrator password to access the administration interface.**

4. **Navigate to Windows Configuration/Manage SMB CIFS Mapping/Configure User Mapping. There you will see radio buttons for each of three user mapping options and each of three group mapping options. The user mapping options are as follows:**

   - No mapping: This is the default setting. When a new user connects, a new UID is generated by StorEdge. This UID will be one larger than the largest current UID found on the StorEdge. Any desired mapping of SMB users to NFS users must be done manually.
   - Map by User Name: This setting specifies that the Windows user's name is looked up via the configured passwd lookup service. If the lookup is successful, the NFS UID is taken from the matching entry. If the lookup fails, a new UID is generated as with the "no mapping" rule.
   - Map by Full Name: This setting specifies that the NT users full name is looked up via the configured passwd lookup service. If the lookup is successful, the NFS UID is taken from the matching entry. If the lookup fails, a new UID is generated as with the "no mapping" rule.

   The group mapping options are as follows:

   - No mapping: This is the default setting. When a new user connects, a new GID is generated by StorEdge. This GID will be one larger than the largest current GID found on the StorEdge. Any desired mapping of SMB groups to NFS groups must be done manually.
   - Map by Group Name: This setting specifies that the NT group name is looked up via the configured group lookup service. If the lookup is successful, the NFS GID is taken from the matching entry. If the lookup fails, a new GID is generated as with the "no mapping" rule.
   - Map to Primary Group: This setting specifies that the NT group name is looked up via the configured passwd lookup service, in the primary group field. If the lookup is successful, the NFS GID is taken from the matching entry. If the lookup fails, a new GID is generated as with the "no mapping" rule.

5. **To set this up with the StorEdge CLI, connect to the StorEdge via Telnet or serial console.**

6. **Type "admin" at the [menu] prompt and enter the administrator password.**

7. **At the CLI, enter "show map*", this will return the current mapping rules. The syntax to set the user mapping rule is as follows, "set <variable> <mapping rule>". Replace <variable> with either smb.map.users or smb.map.groups, depending on which mapping rule you wish to set. The options for <mapping rule> all match the descriptions above, as follows:**

   Valid options for smb.map.users:

   - MAP_NONE—No Mapping
   - MAP_USERNAME—Map by User Name

- MAP_FULLNAME—Map by Full Name

Valid options for smb.map.groups:

- MAP_NONE—No Mapping
- MAP_GROUPNAME—Map by Group Name
- MAP_UNIXGID—Map to Primary Group

Example: **`set smb.map.users MAP_USERNAME`** will define the mapping rule for users to Map by User Name.

Note – All variable names and values are case sensitive. After setting any variables on the StorEdge, i.e. anytime the "set" command is used, the command "savevars" must be entered at the command line in order for the settings to persist though future server reboots.

How do I modify existing user and group credential mappings?

User and group mappings are stored in the configuration files users.map and group.map. We provide a menu interface to edit these mappings. This will be necessary in cases where the NFS user account name does not match the SMB user account name; and in cases where mapping was not configured prior to migration of users and data.

To access this functionality, proceed as follows:.

1. **Access the StorEdge via Telnet or serial console.**
2. **Press [Enter] at the [menu] prompt and enter the administrator password.**
3. **Press the spacebar until "CIFS/SMB Configuration" is displayed under "Extensions" at the lower right.**
4. **Select the letter corresponding to "CIFS/SMB Configuration".**
5. **Select the letter corresponding to "User Mapping". (Or "Group Mapping" for groups)**

   This will bring up a list of all existing user maps, ten per page. If there are more than ten, select option "1" and option "3" to move through the pages. For each user, there is a Windows username, domain and RID on the left side of the screen. The RID is roughly equivalent to the NFS UID or GID. RID information is stored in a database on the Windows Domain Controllers. Note that changing a user's RID in the StorEdge administration interface is not possible. Modifying the value collected from the Domain Controller will simply invalidate the mapping.

On the right side of the screen, you will see the NFS username, which may or may not have been automatically generated based on the defined mapping rule.

Option "7" refreshes the list of mappings, adding any new users.

Option "8" will allow you to manually add a mapping, but this is rarely used, as the RID information is relatively difficult to retrieve from the Domain Controllers. Since the RID is retrieved automatically from each user that logs in, it's easier to edit the mappings after they've been collected.

6. **To edit the mappings, select the letter corresponding to the user information. Then select option "1", "Edit fields". This will display user mapping options as follows:**
   - NT User (or group)
     - Account: This is the Windows user (or group) account in <domain\username> (or group name) format.
     - RID: This is the Windows RID as described above. Usually this is left as is, as this can only be truly changed from the Windows Domain controller.
   - UNIX User (or group)
     - Account: This is the NFS user account name. Changes here will only be saved locally.
     - ID: This is the NFS UID (or GID). This is where the changes are made to effect mapping.

   You can also select option "8" to delete the mapping. This is useful if you have recently changed the mapping rule, as deleted mappings will be remapped according to the current mapping rule the next time this user connect to the StorEdge.

   This functionality is also available through the StorEdge Web Admin.

1. **To use the Web Admin, connect with a Web browser to http://<hostname or IP address of your StorEdge>.**
2. **Click "Grant" or "Yes" to accept any Java software authorization windows and you will reach the login screen.**
3. **Type the administrator password to access the administration interface.**
4. **Navigate to Windows Configuration/Manage SMB CIFS Mapping/Configure maps. All of the options described above are available. Double click user entries to edit them.**

   It is not recommended to edit the map files directly, as the StorEdge must be rebooted immediately in order for changes to take effect.

How do I set up the SMB Autohome directory feature?

Autohome shares are temporary shares that are created when a user logs on to the system and removed when the user logs off. The autohome path defines the base directory path for the shares. For example, if a user's home directory is /usr/home/john, then the autohome path should be set to /usr/home. The temporary share will be named john. It is assumed that the user's home directory name is the same as the user's logon name.

To set up Autohome, proceed as follows:

1. **Access the StorEdge via Telnet or serial console.**
2. **Press [Enter] at the [menu] prompt and enter the administrator password.**
3. **Press the spacebar until "CIFS/SMB Configuration" is displayed under "Extensions" at the lower right.**
4. **Select the letter corresponding to "CIFS/SMB" Configuration".**
5. **Select the letter corresponding to "Autohome Setup".**
6. **Select option "1", "Edit fields".**
7. **Select "Y", "Yes" to enable Autohome support.**
8. **Define the path, including volume name, where user home directories are located.**
9. **Define the ADS container that these shares should be published to. (This assumes that ADS is already configured)**
10. **After configuring desired options, select option "7", "Save Changes".**

This functionality is also available through the StorEdge Web Admin. To use the Web Admin, connect with a Web browser to http://<hostname or IP address of your StorEdge>. Click "Grant" or "Yes" to accept any Java software authorization windows and you will reach the login screen. Type the administrator password to access the administration interface.

Navigate to Windows Configuration/Configure Autohome. Therein you will find all the options listed above.

How do I configure StorEdge to authenticate to a Windows Domain?

StorEdge is capable of providing pass-through authentication to existing Windows domains. To configure this support, proceed as follows:

1. **Access the StorEdge via Telnet or serial console.**

2. **Press [Enter] at the [menu] prompt and enter the administrator password.**

3. **Press the spacebar until "CIFS/SMB Configuration" is displayed under "Extensions" at the lower right.**

4. **Select the letter corresponding to "CIFS/SMB" Configuration".**

5. **Select the letter corresponding to "Domain Configuration".**

   Therein, you will see a list of options as follows:

   - Domain—Name of Windows domain.
   - Scope—SMB scope, this is typically left blank.
   - Description—This is displayed in the network browse list. (Optional)
   - Primary/Secondary WINS—IP address of WINS server(s).
   - Keep Alive—Time in seconds before disconnecting idle connections.
   - Security Mode—Select "2", NT Domain (Auto UID)
   - Username—A user account with the rights to add a computer to the above domain.
   - Password—The password for this user account.

6. **After configuring desired options, select option "7", "Save Changes".**

   This functionality is also available through the StorEdge Web Admin.

1. **To use the Web Admin, connect with a Web browser to http://<hostname or IP address of your StorEdge>.**

2. **Click "Grant" or "Yes" to accept any Java software authorization windows and you will reach the login screen.**

3. **Type the administrator password to access the administration interface.**

4. **Navigate to Windows Configuration/Configure Domains and Workgroups.**

   All of the options described above are available, with the exception of "Keep Alive" and "Scope". WINS configuration can be found separately, under Windows Configuration/Set Up WINS.

How do I set up Active Directory? How do I set up dynamic DNS?

Active Directory is the Windows 2000 directory service that provides centralized access to domain resources such as users, groups and shared data. As StorEdge acquires users and groups from other sources, StorEdge ADS support essentially consists of making shares available via ADS.

ADS relies on the Internet Domain Name System (DNS) to provide name resolution services. The DNS provided with ADS supports the ability for clients to dynamically update their entries in the DNS database; this is known as dynamic DNS.

To configure ADS, proceed as follows:

1. **Access the StorEdge via Telnet or serial console.**
2. **Press [Enter] at the [menu] prompt and enter the administrator password.**
3. **Press the spacebar until "ADS setup" is displayed under "Extensions" at the lower right.**
4. **Select the letter corresponding to ADS setup.**

   Therein, you will see a list of options for both Active Directory Setup, as follows:

   - Enable—Enable ADS
   - ADS Domain—The name of the Windows Domain
   - User—An Windows user account name with rights to update ADS
   - Password—Password for this account
   - User Container—The ADS container that is the location of the above user account.
   - (LDAP distinguished name, without domain, e.g. ou=users)
   - ADS Site—Enter the local ADS site, if different from ADS domain. Usually left blank.
   - Kerberos Realm—Name of Kerberos realm for secure ADS and DNS, usually the ADS domain name.
   - KDC Server—Hostname for Key Distribution Server, usually a domain controller. (This field can usually be left blank, as it can normally be resolved by DNS)

The ADS configuration must be in place before configuring Dynamic DNS.

To configure Dynamic DNS, proceed as follows:

1. **Return to the main menu by pressing the [Esc] key.**
2. **Select option "H", "DNS & Syslogd".**
3. **Select option "1", "Edit Fields".**
4. **Use [Enter] or [Tab] to navigate through the fields.**
5. **Ensure that standard DNS is set up, with a domain name and server(s) configured.**
6. **Select option "Y", "Yes" to enable Dynamic DNS.**
7. **Enter a username and password with sufficient rights to perform secure DNS updates.**
8. **After configuring desired options, select option "7", "Save Changes".**

After you have successfully configured these settings, you will be able to publish shares to ADS using the SMB/CIFS shares menu. Please refer to the FAQ "How do I create SMB shares?" for details on this procedure.

This functionality is also available through the StorEdge Web Admin. This functionality is also available through the StorEdge Web Admin.

1. **To use the Web Admin, connect with a Web browser to http://<hostname or IP address of your StorEdge>.**
2. **Click "Grant" or "Yes" to accept any Java software authorization windows and you will reach the login screen.**
3. **Type the administrator password to access the administration interface.**
4. **Navigate to Windows Configuration/Configure Domains and Workgroups. All of the above described settings can be configured therein.**

What are the limitations of Workgroup mode?

Workgroup mode on theSun StorEdge 5310 NAS also implies use of what Microsoft calls "Share-level Security." In this mode, user tokens are not used, and ACL data cannot be written or read. Resources are accessed as a particular UID/GID combination assigned to each share. Shares are secured by password only, rather than by a username/password combination.

Only NFS style permissions are possible. This mode is intended for only for use with a small number of clients with very low security requirements, such as temporary file transfer, or a small number of servers in a physically secure NFS environment. If there is any requirement for security, or storage of individual user data, Windows domain mode is strongly recommended.

Does StorEdge support Domain Local Groups?

No. StorEdge does not allow files and folders to be secured with local group accounts, including Domain Local Groups. Only Domain Global groups are supported.

Can StorEdge serve as a domain controller in a Windows Domain?

No. The primary purpose of StorEdge is to provide dedicated file service. We provide pass-through authentication to existing domain controllers and Active Directory servers.

Does StorEdge support DFS?

DFS (distributed file system) is a hierarchical file system that allows files to be stored across multiple servers and managed as a single group.

StorEdge can serve as a DFS target. This means that DFS referrals can redirect clients to StorEdge, but StorEdge does not provide referrals and cannot be configured as a root replica.

What support is there for Windows Server Manager?

StorEdge currently supports the following operations in Server Manager:

- View a list of shared resources
- Add a new share
- Delete a share
- View a list of services running on StorEdge
- View a list of connected users

Note – Dialogs that offer the supported operations may also offer other operations but support is limited to the options listed above.

2.24 NIS/NIS+ Issues

How do I set up NIS or NIS+?

NIS and NIS+ provide information to a simple networked database with information about users and hosts. NIS works by copying files from the NIS server to the StorEdge, NIS+ works by performing remote lookups to these files stored on a NIS+ server, and adds several other features.

1. **To access this functionality, access the StorEdge via Telnet or serial console.**

2. **Press [Enter] at the [menu] prompt and enter the administrator password.**

3. **Select option “I”, “NIS & NIS+”. A list of options for both NIS and NIS+ is displayed, as follows:**

   - Network Information Services (NIS)
     - Enable—Enables or disable NIS.
     - NIS Domain—Defines the NIS domain.

  - Broadcast—Enables or disables broadcast search for NIS servers.
  - Server—IP address of NIS server.
  - Files Hosts Users Groups Netgroups—Select which files should be imported from NIS with “Y”.
  - Check Rate minutes—How often to check the NIS server for changes.
- Network Information Services (NIS+)
  - Enable—Enables or disable NIS+.
  - NIS+ Domain—Defines the NIS+ domain.
  - Broadcast—Enables or disables broadcast search for NIS+ servers.
  - Home Domain Server—IP address of primary NIS+ server.
  - Secure RPC Password—Password for NIS+ server
  - Search Path (optional)—List of other NIS domains to search. Separate with colons.

4. **After configuring desired options, select option “7”, Save Changes.**

   This functionality is also available through the StorEdge Web Admin.

1. **To use the Web Admin, connect with a Web browser to http://<hostname or IP address of your StorEdge>.**
2. **Click “Grant” or “Yes” to accept any Java software authorization windows and you will reach the login screen.**
3. **Type the administrator password to access the administration interface.**
4. **Navigate to Unix Configuration/Configure NIS or Unix Configuration/Configure NIS+. Most, but not all of the configuration fields defined above are available.**

How do I configure the NS lookup order?

1. **To access this functionality, access the StorEdge via Telnet or serial console.**
2. **Press [Enter] at the [menu] prompt and enter the administrator password.**

   The lookup order tells StorEdge which set of hosts/groups/hostgrps files to use. The options are LOCAL, NIS, DNS (for hosts only) and NIS+. Each source will be searched in the order you select, and the search will stop when the object is found. Therefore, if you wish for your local hosts file to be used for lookup, it should be first on the list. Otherwise StorEdge will look to first resolve the host via the defined network resources. Netgroups are also supported for lookup, via NIS and NIS+ only.

   This functionality is also available through the StorEdge Web Admin.

1. **To use the Web Admin, connect with a Web browser to http://<hostname or IP address of your StorEdge>.**

2. **Click “Grant” or “Yes” to accept any Java software authorization windows and you will reach the login screen.**

3. **Type the administrator password to access the administration interface.**

4. **Navigate to Unix Configuration/Configure Name Service. Therein you will find tabs for users, hosts, hostgrps and netgroups.**

5. **For each tab, move desired lookup services from the “Services Not Selected” to the “Services Selected” column by highlighting them and clicking the right arrow. Remove services which are not desired by highlighting them and clicking the left arrow. Edit the lookup order by selecting a service in the “Services Selected” column and clicking the up or down arrows. The services will be queried from top to bottom. You must click the apply button in order for any changes to take effect.**

2.25 TCP/IP and Network Configuration

What is port aggregation?

Port aggregation gives you the flexibility to scale your network I/O in port aggregation or to provide NIC port redundancy in high availability.

Port Aggregation is also known as “channel bonding” or "trunking." This type of bonding lets you scale network I/O by joining adjacent NIC ports. It forms a single network channel of high bandwidth from two or more channels of lower bandwidth. You must have a minimum of two available NIC ports for port bonding, and they must be of the same interface type (e.g., Fast Ethernet with Fast Ethernet).

StorEdge’s port aggregation uses Cisco’s Fast EtherChannel architecture. The switch must support EtherChannel bonding, all NICs must be connected to the same switch, and the ports must be specifically configured for EtherChannel. Please refer to the switch user manual for details on how to set up EtherChannel bonding.

High Availability port bonding provides NIC port redundancy or failover. More than one NIC port is bonded to a primary port as backup ports. If the primary port fails, the StorEdge switches over to the backup port that is first on the list of "high availability" bonded ports. If that port also fails, the port next on the list is used and so on.

Any type of switch may be used for High Availability port bonding. Each NIC can be connected to a separate switch, and the switch hardware need not be similar. The only requirement is that all switches used for the HA bond are connected to the same subnet.

How do I set up port aggregation?

1. **Access the StorEdge via Telnet or serial console.**
2. **Press [Enter] at the [menu] prompt and enter the administrator password.**
3. **Select option "A", "Host Name & Network" under the configuration section.**
4. **Select option "3", "Manage bond".**
5. **Select option "1", "Create".**
6. **In the type field enter a "0" for PA Port aggregation or a "1" for HA High Availability.**
7. **Fill in the IP, Netmask and Broadcast parameters for your Network.**
8. **Select a slave NIC (the NIC card in the system that you want to bond to).**
9. **Answer "Yes" to "Create bond with the above configuration?".**

How do I disable port aggregation?

If you wish to disable port aggregation, or if an aborted attempt was made to install port aggregation, it should be disabled on the system. This will prevent future NIC card configuration issues from occurring.

Disable port aggregation at the StorEdge CLI (command line interface).

To disable port aggregation, proceed as follows:

1. **To access the StorEdge CLI, connect to the StorEdge via Telnet or serial console.**
2. **Type "admin" at the [menu] prompt and enter the administrator password.**
3. **At the CLI, enter "set bonding.enable no". Please note that all variable names and values are case sensitive.**
4. **After setting any variables on the StorEdge, i.e. anytime the "set" command is used, the command "savevars" must be entered at the command line in order for the settings to persist though future server reboots.**

What is IP aliasing?

IP Aliasing is a networking feature that allows you to assign multiple IP addresses to a single NIC port. This is useful when StorEdge is replacing multiple servers. All of the IP aliases for the selected NIC port must be on the same physical network and share the same netmask and broadcast address as the first, or primary IP address specified for the selected adapter. Up to nine alias IP addresses can be added to the primary IP address of each NIC port. Therefore, a single network interface card with two ports could provide up to 20 usable IP addresses.

Important – Alias IP addresses can only be added to NIC ports that are assigned a primary role. There are three possible role options for the StorEdge NICs. They are: primary, independent, and mirror. The primary NIC role is not to be confused with the primary IP address. The primary NIC role is an assignment indicating how the adapter will function in a system. The primary IP address is the first address assigned to a selected adapter.

How do I set up IP aliasing?

1. **Access the StorEdge via Telnet or serial console.**
2. **Press [Enter] at the [menu] prompt and enter the administrator password.**
3. **Select option "A", "Host Name & Network".**
4. **Find the page containing the desired network interface with the spacebar.**
5. **Select option "1", "Edit fields".**
6. **Use [Tab] or [Enter] to navigate to the IP Alias field for the desired NIC.**
7. **Select option "1", "Setup".**
8. **Enter an Alias IP for the NIC card.**
9. **Continue to enter alias IP addresses. If no more aliases are required, simply press the [Enter] key.**
10. **System will return to setup screen.**
11. **Navigate to the end of the menu with [Tab] or [Enter].**
12. **Select option "7", "Save changes".**

How do I configure Jumbo Frames support?

Currently this is not supported by the StorEdge software.

Can I set more than one default gateway?

No. The default gateway is the gateway used when a TCP/IP client needs to send data to a network to which it does not have a specific route. After checking the destination network against the routing table and finding no match, the data is sent to the default gateway. There is no provision for TCP/IP to choose between default gateways.

Some operating systems allow the administrator to configure a second default gateway to be used in the case of failure of the primary default gateway. StorEdge does not currently support this feature.

What will happen if I configure multiple network adapters on the same subnet?

In this case, all outbound traffic will be sent via one network interface. This unnecessarily limits network bandwidth. The reason that all the network traffic travels through a single interface is the fact that TCP/IP can only define a single route to each subnet, and this route can only use one network interface. The solution is to link the cards at a lower level via port aggregation.

2.26 Quota Configuration

How do I configure user and group quotas? How do I view current user quotas and disk usage?

Quotas determine how much disk space is available to a user or group, and/or how many files a user or group can write to a particular volume. The quota allocation is enumerated according to file ownership. Changing file ownership of files will change the quota availability.

It is important to note that group quotas apply to an entire group, rather than to each member. For example, if the users group has a 2GB quota, members will be able to own a total of 2GB, regardless of how much space is allocated to any particular user.

The primary interface for quota administration is the StorEdge Web Admin. To use the Web Admin, connect with a Web browser to http://<hostname or IP address of your StorEdge>. Click "Grant" or "Yes" to accept any Java software authorization windows and you will reach the login screen. Type the administrator password to access the administration interface.

First, navigate to File Volume Operations/Edit Properties. There you will find a checkbox to enable quotas. You must highlight each volume to enable or disable quotas for that volume.

Next, navigate to File Volume Operations/Manage Quotas/Configure Group and User Quotas. This brings up a screen which displays all current user quotas, and current disk space and file allocation for each user. Radio buttons at the top of the screen allow you to switch between group and user quota information.

The first line in the display is an entry for root. The purpose of this is to display statistics on files owned by the root user. It is not possible to set a quota for the root user.

The second line in the display is the default quota. Newly added users are automatically limited by the default quota restrictions. Setting the default quota to Unlimited effectively disables this feature. Also, setting the default quota to "Default" has the same effect as setting it to "Unlimited"

The columns in the display are defined as follows:

ID: The NFS UID (GID in the case of group quotas).

Name: The NFS username (group name in the case of group quotas). An entry of "unknown" in this field or the next indicates that files are owned by a UID which does not have an associated username on StorEdge. Typically this is the result of a change in mapping rules or deleted users.

Windows Name: The Windows DOMAIN/username.

KB Used: Storage space in KB currently allocated to files owned by this user or group.

Hard KB Limits: The hard limit is the absolute limit in KB of the total size of the data owned by a particular user or group. When the quota is exceeded, the user will no longer be able to write to this volume.

Soft KB Limits: The soft limit can be exceeded for a period of seven days. If the quota is exceeded for longer than seven days, the user will no longer be able to write to this volume. This allows the temporary allocation of extra storage space for users or groups. If both hard and soft quotas are used, the soft quota must be smaller than the hard quota.

Files Used: Number of files currently owned by this user or group.

Hard Limits/Soft Limits: Same as KB Limits above, but applies to the number of files rather than their size.

To define a quota, locate the desired user on the list, and double click the user entry. This will pop up a window which will allow you to define hard and soft KB limits, as well as hard and soft file limits. You can choose Default, No Limit, or Custom via radio buttons. Default applies the quotas defined for the default user, if any. No Limit allows unlimited storage, and Custom allows you to define a quota in KB, MB or GB. Click “apply” to set the newly defined quota, or click cancel to close the window with no changes.

If the user does not appear in the list, click the add button. Select the UNIX or Windows radio button, and select a user from the list. You can also type in a username, but the account must already exist on the Sun StorEdge 5310 NAS. Having selected the user, proceed as above to define a quota.

How do I configure user and group quotas at the CLI?

All quota operations are also available at the CLI, as follows:

There are four commands available which control quotas.

quotaon /volumename: This enables quotas for a particular volume.

quotaoff /volumename: This disables quotas for a particular volume

quota: This command is used to set quotas. The official syntax is as follows:

quota [-g|-u] /VOLUME|/* NAME|ID|NAME/ID [bh=KB] [bs=KB] [fh=N] [fs=N]

-g and -u are used to specify a user or group name or GID/UID for display or setting

/VOLUME is mandatory, as all quotas are volume specific

A "block" is one KB, bh= defines a hard quota in blocks, bs= defines a soft quota in blocks, fh= and fs= define hard and soft quotas for number of files.

A soft quota may be exceeded temporarily, for a period of seven days. After the seven days, writes from this user will be denied due to quota.

A hard quota cannot be exceeded and will return an error upon the attempt. The most common implementation is to use hard quotas only.

It is permissible to use a user/group name instead of UID/GID, but the mapping must be present in /etc/passwd.nis. UID/GID will always work.

example: "quota -u 100 /vol1 bs=1000000 bh-2000000" sets the soft quota for UID 100 to 1GB and the hard quota to 2GB.

repquota /volumename: This lists all quotas for a given volume. The output is formatted to ten columns, as follows:

Column 1: UID/GID or user group name.

Column 2: By default, this appears to be a two line dash. However, if the listed user has exceeded either soft or hard quota, the respective dash turns to a "+"

Column 3: Current disk usage in blocks for this user.

Column 4: Current soft block quota for this user.

Column 5: Current hard block quota for this user.

Column 6: If currently in excess of soft block quota, time remaining in seven day grace period. Field is blank if user is within soft quota.

Column 7: Current disk usage in number of files for this user.

Column 8: Current soft files quota for this user.

Column 9: Current hard files quota for this user.

Column 10: If currently in excess of soft files quota, time remaining in seven day grace period. Field is blank if user is within soft quota.

There is also a totals line at the end of the quota list, showing how many blocks and files exist on the specified volume.

The quotas are assigned by UID and GID. Every file on the Sun StorEdge 5310 NAS has a GID and UID. In the case of Windows users and groups, the UID/GID is assigned by user mapping.

Please note that the user mapping is only functional in NT domain mode. In secure share mode, the UID and GID are set directly on the share in the Sun StorEdge 5310 NAS UI.

One more important note, check the current file or block usage (via the repquota command) before setting quotas. A common mistake is to not realize that the user or group already owns many files, resulting in a quota that is exceeded as soon as it is defined.

How do I set up Directory Tree Quotas (DTQs)?

Directory Tree Quotas are assigned on sub-directory trees within a volume to limit the amount of space and/or the number of files created under each sub-directory. There is no relationship between directory tree quotas and user/group quotas.

Note – The DTQ creation interface only allows the creation of a DTQ on a new directory. The primary interface for configuring DTQs is the StorEdge Web Admin.

1. **To use the Web Admin, connect with a Web browser to http://<hostname or IP address of your StorEdge>.**
2. **Click "Grant" or "Yes" to accept any Java software authorization windows and you will reach the login screen.**
3. **Type the administrator password to access the administration interface.**
4. **Navigate to File Volume Operations/Manage Quotas/Configure Directory Tree Quotas. There you will find a list of existing DTQs**
5. **To add a DTQ, click the add button. This pops up a window.**
6. **In the pop-up window, select a volume from the pulldown menu.**
7. **Type a unique name for this DTQ.**
8. **Type or browse to the intended parent directory for the DTQ.**
9. **Type a new, unique directory name. A directory with this name will be created at the specified path.**
10. **Define disk space and/or file limits for this directory.**
11. **Click the apply button to set the new DTQ, click cancel to go back to the main screen.**
12. **To edit an existing DTQ, double click the entry.**
13. **In the pop-up window, all the above fields are accessible, with the exception of the directory name, which can no longer be changed.**

    This functionality is also available at the StorEdge CLI (command line interface).

1. **To access the StorEdge CLI, connect to the StorEdge via Telnet or serial console.**
2. **Type "admin" at the [menu] prompt and enter the administrator password.**
3. **At the CLI, enter one of the following commands:**
    - **`dtq create volume=<volume-name> name=<dtq-name> path=<dtq-path> [flimit=N] [slimit=MB]`**

      This creates a new DTQ. "flimit" is limit in number of files, "slimit" is the size limit in MB.
    - **`dtq remove volume=<volume-name> <name=dtq-name>`**

      This removes an existing DTQ. The directory is left intact, but the limits are removed.

- **`dtq rename volume=volume-name from=dtq-name to=dtq-name`**

  This changes the name of the DTQ. Note that this does not change the name of the directory.

- **`dtq set volume=volume-name name=dtq-name [flimit=N] [slimit= MB]`**

  This modifies file or size limits for an existing DTQ.

- **`dtq status [volume=volume-name] [name=dtq-name] [file= file_path]`**

  This shows detail on existing DTQ. “dtq status <volumename>” will return status of all DTQs on the volume.

- **`dtq help`**

  This displays help and syntax for the DTQ command.

How do I create a DTQ for an existing directory?

It is strongly recommended that DTQs be defined at the time of directory creation.

It is also possible to set a DTQ on an existing directory. This method is much faster, but it prevents write access to the entire volume during the quota calculation process.

This functionality is only available at the StorEdge CLI (command line interface).

1. **To access the StorEdge CLI, connect to the StorEdge via Telnet or serial console.**
2. **Type “admin” at the [menu] prompt and enter the administrator password.**
3. **At the CLI, enter “dtq add <path>”. The path must include the volume name.**

How do I disable quotas?

1. **To use the Web Admin, connect with a Web browser to http://<hostname or IP address of your StorEdge>.**
2. **Click “Grant” or “Yes” to accept any Java software authorization windows and you will reach the login screen.**
3. **Type the administrator password to access the administration interface.**
4. **Navigate to File Volume Operations/Edit Properties. Select the desired volume and clear the “Enable Quotas” checkbox.**

Important – This will delete all previously defined quotas. If quotas are enabled in the future, all quotas must be redefined.

Note – Deleting directories also deletes the DTQ set for that directory.

This functionality is also available from the StorEdge CLI.

1. **To access the StorEdge CLI, connect to the StorEdge via Telnet or serial console.**
2. **Type `admin` at the [menu] prompt and enter the administrator password.**
3. **At the CLI, enter `quotaoff /<volumename>`**

What's the difference between a hard quota and a soft quota?

A hard quota cannot be exceeded and will return an error upon any attempt to write.

A soft quota may be exceeded temporarily, for a period of seven days. After the seven days, writes from this user will no longer be allowed.

How does the default quota work?

The default quota is applied to all new users. It is a means to define a per volume quota limit for all current and future users. By default, this quota is unlimited. For specific instructions on how to set the default quota, please refer to the FAQ, "How do I configure user and group quotas?"

2.27 Checkpoint Configuration

What are checkpoints?

A checkpoint, otherwise known as a consistency spot (or c-spot), is a virtual read-only copy of a primary file volume. While the file volume remains in read/write operation, all data existing at the time the checkpoint was created remains available. Checkpoints are used to retrieve mistakenly modified or deleted files, and to stabilize backups. It is important to note that a checkpoint is a virtual, or imaginary, copy of the file volume. It is not an online backup. If the file volume is lost, so are all the checkpoints. A large amount of space and system memory is required for checkpoints. The more checkpoints there are on a system, the more it will affect the system's performance.

Checkpoints are generated from changed blocks. If system usage consists of large amount of file changes then the system will base the next checkpoint off of the changed blocks. It is recommended that when using checkpoints the volume should be allocated at no more than 70% this would allow for system operation to maintain maximum performance as well as give space to the system OS. At 90% disk space utilization, StorEdge will stop creating scheduled checkpoints. At 95% StorEdge will automatically delete checkpoints to free space for system operations.

Checkpoints can be created in two ways: automatic and manual. If the user selects the automatic checkpoints, checkpoints are created and removed based on the scheduling that user specifies for the checkpoints. This scheduling is enforced by a checkpoint manager thread. On the other hand, checkpoint manager does not control manually created checkpoints (though users can create manual checkpoints that will be removed automatically by using the same naming convention that system uses for automatic checkpoints).

How do I set up checkpoints?

To utilize this StorEdge feature, you must first go to the file volume operations menu and enable checkpoints. To do so, proceed as follows:

1. **Access the StorEdge via Telnet or serial console.**
2. **Press [Enter] at the [menu] prompt and enter the administrator password.**
3. **Enter "D", "Disks & Volumes" in the configuration section to set up checkpoints.**
4. **Enter the letter corresponding to the system disk that contains volume that requires checkpoints.**
5. **Select the number corresponding to the volume that requires checkpoints.**
6. **Select option "6", "Checkpoints".**
7. **Select option "1", "Edit fields".**
8. **Use [Tab] or [Enter] to navigate through fields.**
9. **Select option "Y", "Yes" to enable checkpoints on the selected volume**
10. **Pseudo volume checkpoint name is forced to "Yes"**
11. **Select option "Y", "Yes" to make the checkpoint volume visible.**

Checkpoints can be set up automatically or manually. If an automatic schedule is selected then day and time information must be configured for these checkpoints to occur. Checkpoints have a negative effect on system performance. This effect increases as checkpoints are added. Use them judiciously.

To create automatic checkpoints, select "Yes" to "Automatic" and fill in each of the fields listed below:

- Description – Enter a short description of the checkpoint. This is a mandatory field.
- Days – Select the days on which you want the checkpoint to be created. Enter a "Y" for each day that is scheduled. The space bar will leave current data in tact and skip to next day. [Enter] key will skip to the next field.
- AM Hours – Select the AM hours at which you would like the checkpoint to be created. Enter a "Y" for hour that checkpoint will be taken. The space bar will leave current data in tact and skip to next hour. [Enter] key will skip to the next field.
- PM Hours – Select the PM hours at which you would like the checkpoint to be created. Enter a "Y" for hour that checkpoint will be taken. The space bar will leave current data intact and skip to next hour. [Enter] key will skip to the next field.
- Keep Days + Hours – Enter the number of days and hours the checkpoint will be retained. The Days box contains all integer values between 0 and 14, while the Hours box contains all integer values between 0 and 23. This is a mandatory field.

12. **Select option "7", "Save Changes" or press [Esc] to cancel**

    This functionality is also available through the StorEdge Web Admin.

1. **To use the Web Admin, connect with a Web browser to http://<hostname or IP address of your StorEdge>.**
2. **Click "Grant" or "Yes" to accept any Java software authorization windows and you will reach the login screen.**
3. **Type the administrator password to access the administration interface.**
4. **To set up scheduled checkpoints, navigate to File Volume Operations/Configure Checkpoints/Schedule Checkpoints. All options listed above are available.**

    To manage checkpoints, navigate to File Volume Operations/Configure Checkpoints/Manage Checkpoints. Buttons for Create, Remove and Rename are present, along with a list of current checkpoints. Remove and Rename are very straight forward.

    To create a checkpoint, click create, and then select a volume from the pulldown menu. Define a name for the checkpoint, or a time at which it will be deleted. These two options are mutually exclusive.

How do I manage checkpoints from the command line?

To view all checkpoints on a volume:

```
chkpntls VOLUME_NAME
```

A list of all checkpoints on volume is displayed.

To create a single checkpoint, enter the following command from the CLI:

```
chkpntmk VOLUME-NAME CHECKPOINT-NAME
```

Volume-name is the volume that will be checkpointed.

Checkpoint_name is name of checkpoint that is to be created.

Example: `chkpntmk vol1 mycp`

To delete a single checkpoint, enter the following command from the CLI:

```
chkpntrm VOLUME-NAME CHECKPOINT-NAME
```

Volume-name is the volume that contains the checkpoint.

Checkpoint_name is name of checkpoint that is to be deleted.

Example: `chkpntrm vol1 20020117-200000,14d2h`

To rename a single checkpoint, enter the following command from the CLI:

```
chkpntmv  VOLUME-NAME  OLD-CHECKPOINT-NAME NEW-CHECKPOINT-NAME
```

Renaming a scheduled automatic checkpoint and not using the same naming convention that the system uses for automatic checkpoints causes it not be removed automatically by the checkpoint manager. In effect, it now becomes a manual checkpoint.

For additional information on this topic, please refer to the document, “Checkpoint White Paper”.

How do I disable checkpoints?

1. **Access the StorEdge via Telnet or serial console.**
2. **Press [Enter] at the [menu] prompt and enter the administrator password.**
3. **Select option “D”, “Disks & Volumes” in the configuration Section to set up Checkpoints**
4. **Enter letter of system disk that contains volume that requires checkpoints to be disabled**

5. **Enter the number of the volume that contains the checkpoints no longer required.**
6. **Select option "6", "Checkpoints" to configure checkpoints on the selected volume**
7. **Select option "1", "Edit fields".**
8. **Select option "N", "No" to disable checkpoints on the selected volume**
9. **Select option "7", "Save Changes".**

   All existing checkpoints will be deleted and the volume will no longer create new checkpoints.

How do I change the name of automatically created checkpoints?

This feature is not currently supported by the StorEdge operating system.

How do I recover deleted files from a checkpoint?

If an end-user loses a file from the live file system, it can be recovered if the file was on the last checkpoint. There are three methods to use in order to recover checkpointed data.

The simplest way to recover the data is to navigate to hidden directory ".chkpnt" below each directory.

1. **From a Windows client, open a window to the directory which contained the data that you wish to recover.**
2. **In the path Window, append ".chkpnt" to the path. Therein, you will see all checkpoints on the system. If the checkpoint is a scheduled checkpoint, you will be able to tell when it was created from the filename.**
3. **Locate the desired file within the desired checkpoint.**

   Once found, you can "copy and paste" or "drag and drop" the file to the desired location. Please note that "cut and paste" will fail, due to the fact that the checkpoint filesystem is read-only.

   From an NFS client, use the cp command or similar. As above, navigate to the .chkpnt directory, and then locate and copy the desired data.

2.28 Volume Creation and Expansion

How do I create a volume?

The first step is to scan for new disks. The Scan for New Disks option on the Create File Volume panel allows you to scan for new disks that may have been recently added to the system.

To scan for new disks:

1. **Access the StorEdge via Telnet or serial console.**
2. **Press [Enter] at the [menu] prompt and enter the administrator password.**
3. **Select option "D", "Disks and Volumes".**
4. **Select option "9", "scan for new disk".**

A primary file volume is limited to 256GB.

To create a Volume:

1. **Enter the letter corresponding to the system drive where the volume will be created. Space is displayed in GB.**
2. **Select the option number corresponding to "Create" (this will vary depending on how many volumes are already created)**
3. **Select option "1", "sfs2" for type of partition you would like to create. Sfs2 is a primary volume.**
4. **Type a unique name for the new volume in the Name box. Up to 12 characters can be used; however the first character must be a letter. Valid characters include alphanumeric (a - z, A - Z, 0 - 9) and "_" (underscore) characters. The volume name is case sensitive. It is recommended that you always use lower case names, as this will make the addition of extensions and shares much easier.**
5. **Enter the size of extension in MB.**
6. **Select option "7", "Proceed with create".**
7. **StorEdge will format the volume.**
8. **When creation is complete hit the [Esc] key to return to the menu.**

This functionality is also available through the StorEdge Web Admin.

To access these settings, log in, and navigate to File Volume Operations/Create File Volumes. All of the options described above are available.

How do I extend the size of an existing volume?

A primary file volume is limited to 256GB; however, its size can be extended by attaching segments to it. Up to 63 segments can be attached to a single primary file volume.

The first step is to create the segment:

1. **Access the StorEdge via Telnet or serial console.**
2. **Press [Enter] at the [menu] prompt and enter the administrator password.**
3. **Select option "D", "Disks and Volumes".**
4. **Enter the letter corresponding to the system drive where the extension volume will be created. Space is displayed in GB.**
5. **Enter the number for "Create" (this will vary depending on how many volumes have previously been created).**
6. **Select option "2", "sfs2ext" for the type of partition you would like to create. Sfs2ext is an extension volume.**
7. **Type a unique name for the new segment in the "Name" box. Up to 12 characters may be used; however the first character must be a letter. Valid characters include alphanumeric (a - z, A - Z, 0 - 9) and "_" (underscore) characters. This name is case sensitive. It is recommended to use lower case names, as this will make configuration of extensions and shares much easier.**
8. **Enter the size of extension in MB.**
9. **Select option "7", "Proceed with create".**
10. **StorEdge will format the extension volume.**
11. **When complete hit the [Esc] key to proceed.**
12. **Enter the number of the volume to which you want to attach segment.**
13. **Select option "5", "Segments".**
14. **Select option "1", "Add an extension segment".**
15. **Select option "A", "Select extension".**
16. **Read the warning:**

You are about to add this new extension segment to this file volume. Doing so will increase the free space of the volume. This will only take a moment during which the file volume will remain in operation.

ONCE THE EXTENSION IS ATTACHED TO THE FILE VOLUME, IT CANNOT BE DETACHED. THIS IS AN IRREVERSIBLE OPERATION. BE SURE!

17. **Select option "1", "Edit choice" to return to the previous screen, or option "7", "Proceed" to continue with the attachment.**

    This functionality is also available through the StorEdge Web Admin.

    To access these settings, log in, and navigate to File Volume Operations/Create File Volumes. All of the options described above are available.

How do I rename an existing volume?

1. **Access the StorEdge via Telnet or serial console.**
2. **Press [Enter] at the [menu] prompt and enter the administrator password.**
3. **Select option "D", "Disks and Volumes".**
4. **Enter the letter corresponding to the system drive that contains the volume to be renamed.**
5. **Enter the number corresponding to the volume that is to be renamed.**
6. **Select option "3", "Rename".**
7. **Type a unique name for the new volume name in the Name box. Up to 12 characters can be used; however the first character must be a letter. Valid characters include alphanumeric (a - z, A - Z, 0 - 9) and "_" (underscore) characters. This is case sensitive it s recommended to always use lower case as this will make addition of extensions and shares much easier.**
8. **Select option "7", "Proceed with rename".**

    This functionality is also available through the StorEdge Web Admin.

    To access these settings, log in, and navigate to File Volume Operations/Edit Properties. Simply change the name and click "Apply".

How do I delete an existing volume?

1. **Access the StorEdge via Telnet or serial console.**
2. **Press [Enter] at the [menu] prompt and enter the administrator password.**

3. **Select option “D”, “Disks and Volumes”.**

4. **Enter number of volume that is to be deleted.**

5. **Select option “8”, “Delete”.**

6. **As a sanity check, the system will prompt for the volume name. This is case sensitive and it must be typed in exactly as it was entered when the volume was created.**

7. **Select option “7”, “Proceed with delete”.**

8. **When complete, press the [Esc] key to return to the menu.**

   This functionality is also available through the StorEdge Web Admin.

   To access these settings, log in, and navigate to File Volume Operations/Delete File Volumes. Simply highlight the volume, click “Apply”, and answer “Yes” to the confirmation box.

2.29 Reserved Filesystems and Directories

What is the function of the .attic$ directory?

The .attic$ directory is present by default at the root of all user created volumes. Deleted files are temporarily stored here while being processed. Copying user data here will result in that data being deleted.

What is the function of the /etc directory?

The /etc directory stores configuration and log files for the StorEdge. It is strongly recommended that you do not write data to /cvol or modify data on /cvol. This volume should only be accessed when following specific instructions from official documentation or technical support.

What is the function of the /proc volume?

This is the procfs pseudo filesystem. It contains technical detail on running processes and configuration parameters. This information is purely for diagnostic purposes, and therefore isn’t needed for normal operation. It is not possible to copy user data here.

What is the function of the /cvol volume?

The /cvol is a DOS file system volume that is located on the flash memory which the StorEdge boots to. It contains the operating system, the most recent previous version of the operating system, a base version of the operating system, and configuration and log information. It is strongly recommended that you do not write data to /cvol or modify data on /cvol. This volume should only be accessed when following specific instructions from official documentation or technical support.

What is the function of the /dvol volume?

The /dvol is a sfs2 file system volume that is located on the flash memory which the StorEdge boots to. It contains the /etc directory, which is used to store StorEdge configuration data. It is strongly recommended that you do not write data to /dvol or modify data on /dvol. This volume should only be accessed when following specific instructions from official documentation or technical support.

2.30 NFS Issues

How do I share files with NFS users?

1. **To begin sharing files via NFS, access the StorEdge via Telnet or serial console.**
2. **Press [Enter] at the [menu] prompt and enter the administrator password.**
3. **Select option "L", "Volume Access". Therein, you will see a list of all your volumes, including checkpoint and special volumes. On the right side, you will see the current access allowed.**
4. **To change the access, select the letter corresponding to the volume. Three options will appear at the bottom of the screen, as follows:**
5. **Choose the general access for /volumename**

   1. Read/write   2. Read only   3. None

6. **Select the desired access. You will then be prompted to save changes "7", or cancel "0".**

Note – This screen provides access to the host group @general, which by default includes everyone who can reach the StorEdge. To provide NFS access in a more limited and secure way, see the following FAQs:

How do I authorize a trusted host?
How do I authorize an entire subnet as trusted hosts?
How do I manage NFS exports via the StorEdge Web Admin?
How do I manage NFS exports via the configuration files?

How do I authorize a trusted host?

Authorizing trusted hosts allows unrestricted access to all files and folders on the StorEdge via NFS for particular IP addresses.

1. **To access this functionality, access the StorEdge via Telnet or serial console.**
2. **Press [Enter] at the [menu] prompt and enter the administrator password.**
3. **Select option "F", "Hosts".**

   A list of the currently configured hosts is displayed. If the host you wish to add to the trusted list is not present, type the hostname, then the [Enter] key to add host, then "7" to save changes. You will then be prompted for the IP address of the host. If you are using NIS or NIS+ to resolve hosts, then disregard this section and move on to the next step.
4. **Select option "M", "Trusted Hosts".**

   Type the hostname you wish to add to the trusted host list. Select "7" to add it to the list. This setting will take effect for all subsequent NFS mounts. Currently connected users will need to disconnect, and then remount.

   This functionality is also available through the StorEdge Web Admin.

1. **To use the Web Admin, connect with a Web browser to http://<hostname or IP address of your StorEdge>.**
2. **Click "Grant" or "Yes" to accept any Java software authorization windows and you will reach the login screen.**
3. **Type the administrator password to access the administration interface.**
4. **Navigate to Unix Configuration/Configure NFS/Configure Exports.**
5. **Click the add button to add a host, double click an existing host to edit the address or change whether the host is trusted or not.**
6. **Check or clear the checkbox in the pop-up window to add or remove them from the trusted host list.**

How do I authorize an entire subnet as trusted hosts?

This can only be done by directly editing the configuration file, /dvol/etc/hostgrps.

1. **Access this file via NFS or SMB, and open it with an editor.**
2. **Edit the line which begins with the word "trusted". Entries in hostgrps are plain text, separated by spaces.**

   To allow trusted access to the entire class B subnet 192.168.0.0, you would add the entry "192.168.*" The finished hostgrps file should look something like this:

   ```
   general *

   trusted host1 host2 192.168*
   ```

How do I manage NFS exports via the StorEdge Web Admin?

1. **To use the Web Admin, connect with a Web browser to http://<hostname or IP address of your StorEdge>.**
2. **Click "Grant" or "Yes" to accept any Java software authorization windows and you will reach the login screen.**
3. **Type the administrator password to access the administration interface.**
4. **Manage NFS exports by navigating to Unix Configuration/Configure NFS/Configure Exports.**

   Exports are created or deleted with the add/remove buttons. The add button brings up a pop-up window with several options, as follows:

   - Volume: Pull down menu with a complete list of volumes. Select the volume for which you wish to create an export.
   - Path: This defines the directory path relative to the root directory of the volume selected above. This parameter can be omitted, which will result in an export created for the root of the volume.
   - Full Path: This is a display only field, showing the full export path, combining the volume and directory path.
   - Access: Select a radio button to define read/write, read only, or no access for this export. The "no access" export is used to restrict access to a group which would otherwise have access to the export.
   - Hosts: Select the host or group of hosts to which this directory path will be exported. Pull down menus are available for existing hosts, hostgrps and netgroups entries. There is also a text box, in which you can type an IP address or resolvable hostname for a previously unknown host.
5. **After filling these fields, click the apply button to save the new export.**

Selecting an export and clicking the remove button will remove the export.

Selecting an export and clicking the edit button, or double clicking an existing export will bring up the edit screen. Only the “access” field may be changed in this screen. Other changes must be made by deleting the export and recreating it, or by editing the configuration files manually. Also, please note that it is not currently possible to change order of exports in the Web Admin.

For additional information on this topic, and also for an additional level of configuration detail, please refer to the FAQ, “How do I manage NFS exports via the configuration files?”

Important – It is recommended that you avoid sharing user data at the root of a volume. The best practice is to create subdirectory structures and share these. This eases security administration, and removes the need to hide or secure system folders.

How do I manage NFS exports via the configuration files?

Editing the configuration files directly affords the greatest level of control over NFS exports. The primary files which are used for this are /dvol/etc/hostgrps and /dvol/etc/approve. Access these file via NFS or SMB, and open them with an editor.

Important – You must enter the command “approve update” at the CLI after the editing of files is complete. Otherwise, all changes to the approve files will be lost the next time a change is made via one of the administration interfaces.

Entries in the hostgrps file are plain text, separated by spaces. The name of the host group is always the first entry on each line, followed by one or more hostnames or IP addresses.

Entries in the approve file are also plain text, separated by spaces. Comment lines are preceded by the “#” character. The active lines each define an NFS export. The syntax for the active lines is as follows:

```
<Object type> <path> <hosts or groups> <security>
```

The object type will always be “files”.

The path is the full directory path to the data exported, including volume name.

Hosts or groups will typically be a hostgrps entry, but can also be a netgroup entry or host specification.

Security is the security setting for this export.

Here's an example:

```
files          /              @trusted          access=rw  uid0=0
```

This entry gives access to all files (/), to the hostgrps entry trusted. (The "@" symbol defines it as a hostgrps entry) with read/write (rw) security. The last entry "uid0=0" is a special entry which disables root squash for an export. Root squash is the default, and it causes all mounts done as UID 0, to be done as UID 60001, the "nobody" account. The above entry creates a trusted host, as it gives root access to the root directory.

2.31 Administration Interfaces and Utilities

How do I access the StorEdge administration interfaces?

StorEdge offers several administration interfaces, as follows:

- StorEdge Web Admin: To use the Web Admin, connect with a Web browser to http://<hostname or IP address of your StorEdge>. Click "Grant" or "Yes" to accept any Java software authorization windows and you will reach the login screen. Type the administrator password to access the administration interface.
- Telnet: Connect with your telnet client to the IP or hostname of the StorEdge Server. Press [Enter] at the [menu] prompt and enter the administrator password.

What is the default administrator password?

How do I configure an administrator password?

By default, there is no password on the StorEdge. To set one, proceed as follows:

1. **Access the StorEdge via Telnet or serial console.**
2. **Press [Enter] at the [menu] prompt.**
3. **Select option "K", "Admin Access".**
4. **Select "1", "Edit Fields".**
5. **Select "Y", "Yes" to enable password protection.**
6. **Type a new admin password in the text box.**
7. **Type it again for verification purposes.**

8. **Select option "7", "Save Changes".**

   This functionality is also available through the StorEdge Web Admin.

1. **To use the Web Admin, connect with a Web browser to http://<hostname or IP address of your StorEdge>.**
2. **Click "Grant" or "Yes" to accept any Java software authorization windows and you will reach the login screen.**
3. **Type the administrator password to access the administration interface.**
4. **Navigate to System Operations/Set Administrator Password. In this menu, there is no enable/disable field. The presence of a password enables password protection.**

How do I collect diagnostics from StorEdge?

The diagnostic email includes information about the StorEdge system configuration, disk subsystem, file system, network configuration, SMB shares, backup/restore information, /etc information, system log, environment data and administrator information. The diagnostics are a primary tool for checking configuration and troubleshooting.

To collect diagnostics, proceed as follows:

1. **Access the StorEdge via Telnet or serial console.**
2. **Press [Enter] at the [menu] prompt and enter the administrator password.**
3. **Press the spacebar until "Diagnostics" is displayed under "Extensions" at the lower right.**
4. **Select the letter corresponding to "Diagnostics".**
5. **Wait a few seconds while the StorEdge builds the diagnostic.**
6. **Select option "2", "Send Email".**
7. **Select option "1", "Edit problem description".**
8. **Enter a precise description of the problem.**
9. **Select option "8", "Send Email".**

   Diagnostic is sent

   If an email server is not configured or not available, it is also possible to save the diagnostics to a file on the StorEdge. To do this, use Steps 1-5 above to access the "Diagnostics" menu.

1. **Select option "1", "Save File".**

2. **Select option “1”, “Edit path”.**

3. **Enter a valid path name in the path box. Format is /<volumename>/<directory>/<new filename>.**

4. **Select option “2”, “Save diagnostics file”**

5. **StorEdge will respond “Diagnostic saved”.**

6. **Access the volume that you saved the file to via SMB or NFS.**

7. **Copy the file to a local workstation.**

   This functionality is also available through the StorEdge Web Admin.

1. **To use the Web Admin, connect with a Web browser to http://<hostname or IP address of your StorEdge>.**

2. **Click “Grant” or “Yes” to accept any Java software authorization windows and you will reach the login screen.**

3. **Type the administrator password to access the administration interface.**

4. **Click the envelope icon on the top taskbar. All of the options described above are available.**

How do I set up SMTP (email)?

SMTP (email) configuration allows StorEdge to send email directly to your mail server. You can use this functionality to collect diagnostic information from the StorEdge or to send warning messages for critical system events via email.

Email from StorEdge usually will not be accepted unless a DNS domain is configured. This is true whether or not you are using DNS. To check or configure DNS settings, proceed as follows:

1. **Access the StorEdge via Telnet or serial console.**

2. **Press [Enter] at the [menu] prompt and enter the administrator password.**

3. **Select option “H”, “DNS & SYSLOGD” in the configuration Section to set up DNS.**

4. **Check to see if there is an entry in the DNS domain field. If the entry is present, proceed to the next step.**

5. **Otherwise, select option “1”, “Edit fields”.**

6. **Press [Tab] or [Enter] to move through fields.**

7. **In the DNS Domain field, enter the Domain name for the StorEdge.**

8. **Select option “7”, “Save Changes”. This will return you to the menu.**

9. **Access the StorEdge via Telnet or serial console.**

10. **Press [Enter] at the [menu] prompt and enter the administrator password.**

11. **Press the space bar until the “Email Configuration” option is displayed under “Extensions” at the lower right.**

12. **Select the letter corresponding to “Email Configuration”.**

13. **Select option “1”, “Edit fields”.**

14. **Press [Tab] or [Enter] to navigate through the fields.**

15. **Enter the IP address of your SMTP (email) server.**

16. **If your mail server requires that email be sent from a domain other than the DNS domain, enter it here.**

17. **Enter up to 4 recipient email addresses.**

18. **Check the desired message options for each recipient: “Notifications”, “Diagnostics”, or both.**

19. **Set Notifications to either “errors and warnings” or “none”.**

20. **Select option “7”, “Save Changes”.**

    This functionality is also available through the StorEdge Web Admin.

1. **To use the Web Admin, connect with a Web browser to http://<hostname or IP address of your StorEdge>.**

2. **Click “Grant” or “Yes” to accept any Java software authorization windows and you will reach the login screen.**

3. **Type the administrator password to access the administration interface.**

4. **Navigate to Monitoring and Notification/Email Notification. All of the options**

How do I set up local or remote logging?

The Set up Logging panel enables the System Message Logger and its designated server. This is only applicable if your system includes a syslogd (pronounced “syslog dee”) server on the network that can receive the StorEdge system log. StorEdge can also save the log files locally. Configuring one or both of these options is required to save logs through system reboots.

1. **To access this functionality, access the StorEdge via Telnet or serial console.**

2. **Press [Enter] at the [menu] prompt and enter the administrator password.**

3. **Select option "H", "DNS & SYSLOGD" in the configuration Section to set up remote logging**

4. **Select option "1", "Edit fields".**

5. **Use [Tab] or [Enter] to navigate through the fields.**

6. **Select option "Y", "Yes", to enable SYSLOGD.**

7. **Enter the IP address of the SYSLOGD server that will receive the StorEdge system log (if applicable).**

8. **Select the appropriate Facility. The facility indicates the application or system component generating the messages. All messages sent to the syslogd server will have this facility value.**

   The possible facility values in the Set up Remote Logging panel include:

   - Kern – Messages generated by the kernel. These cannot be generated by any user processes.
   - User – Messages generated by random user processes. This is the default facility identifier if none is specified.
   - Mail – The mail system.
   - Daemon – System or network daemons.
   - Auth – Authorization systems, such as login.
   - Syslog – Messages generated internally by syslogd.
   - Lpr – The line printer spooling system.
   - News – Reserved for the USENET network news system.
   - Uucp – Reserved for the UUCP system, which does not currently use syslog.
   - Local0 - Local7 – Reserved for local use.

9. **Next are the settings for local logging, these are independent of the remote log settings.**

   - Local Log - Enable/disable local logging
   - Local File - Filename for local logging, must include full path and filename
   - Archives - These two settings define how much log data to save, archives is the number
   - Size (KB) - of files to save, KB is the size of each.

10. **Select option "7", "Save Changes"**

    This functionality is also available through the StorEdge Web Admin.

1. **To use the Web Admin, connect with a Web browser to http://<hostname or IP address of your StorEdge>.**
2. **Click “Grant” or “Yes” to accept any Java software authorization windows and you will reach the login screen.**
3. **Type the administrator password to access the administration interface.**
4. **Navigate to Monitoring and Notification/View System Events/Set Up Logging. All of the options described above are available.**

How do I set up NTP or RDATE?

The StorEdge can be configured to synchronize its time with either Network Time Protocol (NTP) protocol or an RDATE server. NTP is an Internet protocol used to synchronize the clocks of computers to a reference time source, such as a radio, satellite receiver or modem. Typical NTP configurations use multiple redundant servers and diverse network paths to achieve high accuracy and reliability.

The RDATE time protocol provides a site-independent date and time. It is a protocol that can retrieve the time from another machine on your network. RDATE servers are commonly present on UNIX systems, and allow you to synchronize StorEdge server time with RDATE server time.

1. **To set up NTP, proceed as follows:**
2. **Access the StorEdge via Telnet or serial console.**
3. **Press [Enter] at the [menu] prompt and enter the administrator password.**
4. **Press the spacebar until “NTP Configuration” is displayed under “Extensions” at the lower right.**
5. **Select the letter corresponding to “NTP Configuration”.**
   - NTP Enable - Enable/disable NTP
   - Server 1 enable - Enable the NTP server configured below
   - Server - IP address
   - Authentication - Select none or symmetric key
   - Key ID - Enter symmetric key if enabled above
   - Server 2 enable - Enable the NTP server configured below
   - Server - IP address
   - Authentication - Select none or symmetric key
   - Key ID - Enter symmetric key if enabled above
   - Min Polling Interval - Minimum polling rate for NTP messages

- Max Polling Interval - Maximum polling rate for NTP messages

---

Note – Above two fields are in seconds, raised to the power of two. For example, an entry of 4 sets the interval to 16 seconds. Valid range is 4 to 17.

- Broadcast Client Enabled - Allows StorEdge to respond to NTP broadcasts.
- Require Server authentication - Allows NTP communication only with authentication.

To set up RDATE, proceed as follows:

1. **Press [Esc] to return to the menu.**
2. **Press the spacebar until the "RDATE time update" option is displayed in the extension section in the lower right.**
3. **Select the letter corresponding to "RDATE time update".**
4. **Select option "1", "Edit Fields".**
5. **Use [Tab] and [Enter] to navigate through fields.**

   Settings are as follows:

   - Enable - Enables/disables RDATE service
   - Server - IP address of RDATE server
   - Tolerance - Maximum amount in seconds to modify time via RDATE
6. **After configuring desired options, select option "7", Save Changes.**

   To use RDATE, NTP service must be disabled first. When an RDATE server is set, the server is consulted once a day at 23:45 for the current time. If the reply time is within the Delta Tolerance (+/- seconds), this servers time is updated. The time is also updated when changes are saved in this screen.

How can I capture network traffic to and from the StorEdge?

StorEdge includes a built-in network monitoring tool. This allows you to capture packets from the network and save them to a file. This can be a valuable troubleshooting tool.

To configure network monitoring, it must first be loaded at the StorEdge CLI (command line interface).

1. **To access the StorEdge CLI, connect to the StorEdge via Telnet or serial console.**
2. **Type "admin" at the [menu] prompt and enter the administrator password.**

3. **At the CLI, enter “load netm”. Then type “menu” to configure capture and capture packets.**

4. **Press the spacebar until “Packet Capture” is displayed under “Extensions” at the lower right.**

5. **Select the letter corresponding to “Packet Capture”.**

6. **Select option “1”, Edit Fields.**

   The available options are as follows:

   - Capture File - Where to save the capture file. </volumename/directory/filename>
   - Frame Size (B) - Size in bytes of each frame to capture. The default is normally used.
   - IP Packet Filter - “No” captures all traffic, “Yes” allows you to filter what is received. A filter allows you to select which IP address or addresses you will capture traffic from. You can also filter on a particular TCP or UDP port.
   - Dump Enable - Select “Yes” to allow StorEdge to save the capture in the event of a problem.

7. **After configuring these options, select option “7”, “Start Capturing”**

8. **Reproduce the network event you wish to capture.**

9. **Select option “7”, “Stop Capture”.**

10. **Access the file via NFS or SMB and copy the file as needed.**

How do I access command history at the CLI?

When typing commands at the CLI, it is sometimes desirable to access previously typed commands. StorEdge has several command history options available. Note that the CLI does not provide Unicode support. The features described here are intended for use with ASCII characters only.

1. **To access the StorEdge CLI, connect to the StorEdge via Telnet or serial console.**

2. **Type “admin” at the [menu] prompt and enter the administrator password.**

   The history list contains a maximum of 32 entries and will be saved to /cvol/log/history whenever the exit or quit commands are issued. The contents of /cvol/log/history will be read on boot. The shell built-in history command can be used to display a numbered list of previously executed commands. Previous commands can be executed by entering “!nn” at the CLI, where nn is the entry number. The history command is not displayed in the help list.

The command line can be edited using the following key bindings. There is no overwrite mode. Characters will always be inserted at the cursor position. [Ctrl] + t can be used to display the key bindings at any time. The current command line will be redisplayed following the key list.

From the command line:

The StorEdge supports standard keyboard functionality for command history. The cursor movement keys can be used to select the following:

- Previous commands (using up arrow).
- Scroll forward through history (using down arrow).
- Previous character (using left arrow).
- Forward one character (using right arrow).

Full command list:

- [Ctrl] + a (beginning-of-line) - Move to the start of the line
- [Ctrl] + b (backward-char) - Move backward one character
- [Ctrl] + c (abort-key) - Cancel current operation
- [Ctrl] + d (delete) - Delete the character under the cursor
- [Ctrl] + e (end-of-line) - Move to the end of the line
- [Ctrl] + f (forward-char) - Move forward one character
- [Ctrl] + g (abort-key) - Cancel current operation
- [Ctrl] + h (rubout) - Delete the character behind the cursor
- [Ctrl] + i (Tab) - File completion
- [Ctrl] + j (Enter) - Execute command
- [Ctrl] + k (kill-end-of-line) - Kill to the end of the line
- [Ctrl] + l (refresh-line) - Redisplay the line
- [Ctrl] + n (next-line) - Move down to the next line
- [Ctrl] + p (prev-line) - Move up to the previous line
- [Ctrl] + r (backward-word) - Move backward one word
- [Ctrl] + t (help) - Display CLI key help
- [Ctrl] + u (kill-line) - Kill entire line
- [Ctrl] + w (forward-word) - Move forward one word

The Tab key ([Ctrl] + i) performs file name completion. If there is a single match the path will be updated. If the matching object is a directory a / will be appended, otherwise a space will be appended. If there are multiple potential matches a list will be displayed and the command line will be redisplayed.

Examples:

- To repeat the last command, enter [Ctrl] + p.
- To back up and edit the current line, use [Ctrl] + b.
- To delete the character under the cursor, enter [Ctrl] + d.

How do I delete files from the StorEdge administration utilities?

The operating system has some CLI commands available to perform advanced system administration. Caution must be exercised, as these commands can change data paths and structures. Under certain circumstances, a mistyped command can result in downtime or data loss.

These tools must be loaded from the Command Line Interface (CLI).

1. **To access the StorEdge CLI, connect to the StorEdge via Telnet or serial console.**
2. **Type "admin" at the [menu] prompt and enter the administrator password.**
3. **From the CLI, enter "load unixtools".**
4. **To delete files from the CLI using del**

   The del command deletes one or more specified files. It does not delete directories or files with the immutable bit set. The command requires the file's full path and filename. If no arguments are entered, the usage description is displayed. The del command does not interpret any wildcard characters.

   Specified files will be deleted regardless of the file's security or the permission. You will not be prompted for confirmation even if the file(s) is read only.

   The del command also accepts filenames containing possibly illegal characters, such as the double quote (") character, to be deleted. To remove a filename containing such characters, prepend the character in question with the backslash "\" character. All commands are entered from the command line.

   Examples:

   - Delete a file with an illegal character: SE5310> del /path/file\"test
   - Delete multiple files: SE5310> del /path/file1 /path/file2 /path/file3

   To delete files from the CLI using rm

   If the recursive flag is specified, the file hierarchy rooted at DIRECTORY is removed. The rm command removes symbolic links, not the files referenced by the links.

   Examples:

   Remove the file 'file1' from the directory '/vol1/dir1'.

   ```
   SE5310 > rm /vol1/dir1/file1
   ```

Remove the directory '/vol1/dir1' if it is empty.

```
SE5310 > rm /vol1/dir1
```

Remove the file hierarchy rooted at '/vol1/dir1' displaying each file as it is removed.

`SE5310 > rm -r -v /vol1/dir1` This removes all files and the directory.

Note – All paths must be absolute paths from the root directory.

To delete directories from the CLI using rmdir

The rmdir utility removes the directory entry specified by each directory argument, provided it is empty. Arguments are processed in the order given. In order to remove both a parent directory and a subdirectory of that parent, the subdirectory must be specified first so the parent directory is empty when rmdir tries to remove it.

Example:

```
rmdir /vol1/d1 /vol1/d2 /vol1/d3 /vol1
```

Note – All paths must be absolute paths from the root directory.

How do I recursively delete directories from the StorEdge administration utilities?

StorEdge provides a utility to delete entire directory trees, including their contents. This is a very powerful tool, but one that comes with some risk. This command immediately and permanently deletes files and directories, or even the entire contents of a volume. This tool should be used very carefully, and the entry should be carefully checked before entering. No messages asking to verify will be presented.

In order to prevent accidental or unauthorized deletions, this command must be manually loaded before use. This functionality is only available at the StorEdge CLI (command line interface).

1. **To access the StorEdge CLI, connect to the StorEdge via Telnet or serial console.**
2. **Type “admin” at the [menu] prompt and enter the administrator password.**
3. **At the CLI, enter “load rdel.nsm”.**
4. **Enter “rdel <pathname>”. The pathname must be a full path including volume name.**

Currently, the rdel utility cannot be unloaded from memory, and therefore will not be removed from memory until the next reboot.

How do I set up an FTP server on StorEdge?

StorEdge has a built-in FTP server. Before using it, you must load it via the CLI (command line interface).

1. **To access the StorEdge CLI, connect to the StorEdge via Telnet or serial console.**
2. **Enter "admin" at the [menu] prompt and enter the administrator password.**
3. **At the CLI, enter "load ftpd" to initialize the FTP service.**
4. **Then type "menu" to access FTP configuration.**
5. **Press the spacebar until "FTP configuration" is displayed under "Extensions" at the lower right.**
6. **Select the letter corresponding to FTP configuration.**

   The following options are available:

   - Enable FTP - enable or disable the FTP server.
   - Allow guest access - allow anonymous users
   - Allow user access - allow non-anonymous users
   - Allow admin access - allow access to admin user (root access)
   - Enable logging - log access to FTP server

7. **Select option "1", "Edit fields". Move through the options with the [enter] or [tab] keys.**
8. **After selecting desired options, select option "7", "save changes".**

   Individual FTP usernames and passwords can be configured under menu option "E", Users.

9. **Type a new username to add, type an existing username to edit.**

How do I configure the FTP service to load automatically?

To configure the FTP service to load automatically on every boot, proceed as follows:

1. **Create a text file named inetload.ncf. Note that the name must be all lower case, and the file must be plain text. The file should contain only the following two lines:**

```
# Load the FTP service
```

```
ftpd
```

2. **Next, this file must be copied to the StorEdge /dvol/etc directory. Access this directory via NFS or SMB and copy the file.**

   On future system reboots, the inetload service will read and act on the file automatically at boot time.

How do I enable server-to-server FTP copying on the StorEdge?

The FXP protocol for FTP allows for server-to-server file transfers. The environment variables “ftp.fxp.<user/admin/guest>” must be set to yes to enable FXP for the appropriate user class.

This functionality is only available at the StorEdge CLI (command line interface).

1. **To access the StorEdge CLI, connect to the StorEdge via Telnet or serial console.**
2. **Type “admin” at the [menu] prompt and enter the administrator password.**
3. **To enable FXP for admin users, at the CLI, enter “set ftp.fxp.admin yes”.**
4. **To enable FXP for standard users, at the CLI, enter “set ftp.fxp.user yes”.**
5. **To enable FXP for anonymous users, at the CLI, enter “set ftp.fxp.guest yes”.**

Note – All variable names and values are case sensitive. After setting any variables on the StorEdge, i.e. anytime the “set” command is used, the command "savevars" must be entered at the command line in order for the settings to persist though future server reboots.

How do I configure and secure rsh access?

Rsh (remote shell) is typically used to run StorEdge CLI commands remotely, often as part of a script operation. Some possible applications are checking disk space, creating checkpoints, or graphing CPU utilization.

Rsh is enabled by default on StorEdge servers. Also by default, rsh commands can only be run after responding to a prompt for a password. This assumes that you have an administrator password set on the StorEdge. If there is no password defined, access to rsh is unrestricted, as would be the case with all of the administrative interfaces.

In order to allow rsh access without a password, StorEdge system variables are used.

Users allowed explicit access by one of the following environment variables will not be prompted for a password. Variables are set at the StorEdge CLI (command line interface).

1. **To access the StorEdge CLI, connect to the StorEdge via Telnet or serial console.**
2. **Type "admin" at the [menu] prompt, and enter the administrator password.**

   The syntax is as follows:

   `set rshd.allow.<cmd>.<user> yes`

   The <user> parameter is optional. If it is not used, it allows rsh execution of the specified command for all users.

   For example:

   `set rshd.allow.help yes` allows everyone rsh access to the help command.

   When the <user> parameter is used, it allows rsh access to the command only for the specified user. The user must be specified in the format "user@host". The host argument can be an IP address, a simple hostname in the local domain, or a fully qualified hostname.

   For example:

   `set rshd.allow.df.root@websys.procom.com yes` gives the root user at websys.procom.com rsh access to the df command.

Note – All variable names and values are case sensitive. After setting any variables on the StorEdge, i.e. anytime the "set" command is used, the command "savevars" must be entered at the command line in order for the settings to persist though future server reboots. Any host names or user names used must be resolvable via the hosts and passwd files, or via the DNS service, dependent on your settings for lookup order.

How do I access a shell prompt from the StorEdge CLI?

A true shell prompt is not available. However, you can gain access to the chown, du, ll, find, mkdir, rmdir, cp and rm commands by entering "load unixtools" at the StorEdge CLI (command line interface).

1. **To access the StorEdge CLI, connect to the StorEdge via Telnet or serial console.**
2. **Type "admin" at the [menu] prompt and enter the administrator password.**

Help is available for each of these commands at the prompt by typing "help <command>". Additionally, man pages are available for mkdir, rmdir, cp and rm. Access these by entering "man <command>".

How do I enable or disable ftp, tftp, rlogin, rsh, telnet, ssh, smb or the Web Admin?

By default, all of these services are enabled, with the exception of tftp. Enabling or disabling of these services is done via the netserv command at the CLI (command line interface).

1. **To access the StorEdge CLI, connect to the StorEdge via Telnet or serial console.**
2. **Type `admin` at the [menu] prompt and enter the administrator password.**
3. **Type `netserv` alone to retrieve a list of currently enabled services.**
4. **Type `netserv enable <service name>` to enable a particular service.**
5. **Type `netserv disable <service name>` to disable a particular service.**
6. **Type `netserv status <service name>` to check whether a particular service is running.**

For additional information on this topic, please see the man page. This can be accessed by typing `man netserv` at the StorEdge CLI.

2.32 Backup and Migration Issues

What steps do I need to take before migrating data to StorEdge?

The following steps are recommended before migrating data to StorEdge.

Migrate user database information: If you use NIS or NIS+, this is as simple as defining server information. If you manually maintain your passwd, hosts, and groups files, these files must be copied to the /etc directory on the StorEdge.

To determine the location of this file, access the StorEdge CLI (command line interface).

1. **To access the StorEdge CLI, connect to the StorEdge via Telnet or serial console.**
2. **Type "admin" at the [menu] prompt and enter the administrator password.**

3. **At the CLI, enter “show file.hosts”. This will return the location of the active hosts file. You can safely assume that the active passwd and group file are located in the same directory.**

4. **Next, run the following commands from the CLI: “cleari /<volumename>/etc/hosts”, “cleari /<volumename>/etc/passwd”, and “cleari/<volumename>/etc/group”. Press the [Enter] key after each of these.**

5. **Next, copy the updated version of these files to the /etc directory located above via NFS or SMB.**

   Determine NFS Unicode and language settings: Several international character sets use double-byte character encoding to define extended characters. It is important to know whether these clients use Unicode encoding or not. This setting is configured on the client systems. Windows systems use Unicode code pages exclusively, so in a mixed environment, it’s best to configure the NFS client for Unicode. This primarily affects filenames, but also network names and few other items.

   The StorEdge assumes that NFS clients send this data in ASCII format, and converts this data to UTF-8 before storing it. If the NFS clients use Unicode, this behavior needs to be changed, which can only be done at the StorEdge CLI (command line interface).

1. **To access the StorEdge CLI, connect to the StorEdge via Telnet or serial console.**

2. **Type “admin” at the [menu] prompt and enter the administrator password.**

3. **At the CLI, enter “set nfs.utf8 yes”.**

   This tells the StorEdge that the information is already in Unicode UTF-8 format. Please note that all variable names and values are case sensitive. After setting any variables on the StorEdge, i.e. anytime the “set” command is used, the command "savevars" must be entered at the command line in order for the settings to persist though future server reboots.

   Also, you should configure the language codepage that StorEdge should use. The best way to do this is with the StorEdge Web Admin.

   This functionality is also available through the StorEdge Web Admin.

1. **To use the Web Admin, connect with a Web browser to http://<hostname or IP address of your StorEdge>.**

2. **Click “Grant” or “Yes” to accept any Java software authorization windows and you will reach the login screen.**

3. **Type the administrator password to access the administration interface.**

4. **Navigate to System Operations/Assign Language. Select the desired language, and click the “Apply” button.**

Define Directory Tree Quotas: It is best to define DTQs before migrating data to the StorEdge. There are difficulties associated with setting DTQs on existing data which can be completely avoided by planning ahead.

Set an administrator password: Administrator access allows many powerful options, including deletion of volumes and override of security settings. Define a secure password to protect your data.

What the best way to migrate data to or from StorEdge?

First, please see the preparatory steps in the FAQ, "What steps do I need to take before migrating data to StorEdge".

If you are migrating data from an SMB environment, and the data contains ACL information, you need to use a utility which is aware of the ACL information, and is configured to copy it. Previous versions of Windows used a Resource Kit utility called scopy. Later versions of Windows 2000 and XP support the /X and /O flags for the built-in xcopy command. which cause ACL and ownership information to be copied with the data.

If you are migrating data from an NFS environment, you primarily need to be concerned with file ownership and mode security. Be sure to configure your copy mechanism to preserve owner and mode information, or alternatively, use chmod and chown to set security immediately after migration. This information is used to grant access and calculate quota information, so make sure that root doesn't own all the files. Also, please note that StorEdge does not support any NFS ACL information at this time.

How do I set up NDMP to backup the StorEdge?

The Network Data Management Protocol (NDMP) is an open protocol for network-based backup. NDMP architecture allows network attached storage vendors to ship NDMP-compliant servers that can be used with any NDMP-compliant backup administration application.

To backup using NDMP, you must first enable checkpoints, and configure them for backup. To do so, proceed as follows:

1. **Access the StorEdge via Telnet or serial console.**

2. **Press [Enter] at the [menu] prompt and enter the administrator password.**

3. **Select option "D", "Disks & Volumes" in the configuration Section to set up Checkpoints.**

4. **Enter the letter corresponding to the system disk that contains volume that requires checkpoints.**

5. **Enter the number corresponding to the volume that requires checkpoints.**
6. **Select option "6", "Checkpoints".**
7. **Select option "1", "Edit fields".**
8. **Use [Tab] or [Enter] to navigate through fields.**
9. **Select option "Y", "Yes" to enable checkpoints on the selected volume.**
10. **Select option "Y", "Yes" to use the checkpoint for backups. Backing up from a checkpoint avoids all issues related to backing up open files.**
11. **Select option "7", "Save Changes". This will take you to the main menu.**

To enable a NDMP device to backup the StorEdge:

1. **Press the space bar until "NDMP Setup" option is displayed in the extension section in the lower right.**
2. **Select the letter corresponding to "NDMP Setup".**
3. **Select the NIC that will be used to transfer data to the tape drive server. If you have no preference, select a NIC that is on same network as the gateway.**

How do I find the names of NDMP devices connected to the StorEdge?

This functionality is only available at the StorEdge CLI (command line interface).

1. **To access the StorEdge CLI, connect to the StorEdge via Telnet or serial console.**
2. **Type "admin" at the [menu] prompt and enter the administrator password.**
3. **At the CLI, enter "ndmp devices".**

How do I view the NDMP backup log?

1. **Access the StorEdge via Telnet or serial console.**
2. **Press [Enter] at the [menu] prompt and enter the administrator password.**

   The log is stored on the /etc directory of the first volume created on the StorEdge. If the first volume is vol1 then enter the following from the command prompt to view the log file:

   ```
   cat /vol1/etc/backup/ndmp.log
   ```

How do I move files from one disk to another with NDMP?

Currently this function is not supported in the StorEdge software.

What tape libraries and drives have been tested for compatibility with the StorEdge?

NEED TO ADD THIS INFORmation

2.33 Macintosh Connectivity

How do I share files with Mac Users?

StorEdge does not support AppleTalk networking. Macintosh clients require an SMB or NFS client in order to connect to the StorEdge.

Apple OS-X has built-in NFS client software. Third party NFS clients are also available. Using an NFS client is the recommended solution.

We have done very limited testing with the DAVE client from Thursby. This software enables browsing and file access from Mac clients via SMB.

What does the desktop DB option in the shares menu do?

The Desktop DB Calls option allows the StorEdge to access and set Macintosh desktop database information. It speeds up Macintosh client file access and allows non-Macintosh clients to access Macintosh files on the StorEdge.

2.34 Miscellaneous Log Messages

Why can’t I see system log information prior to most recent reboot?

By default, the system log is stored only in memory. Therefore, it is lost upon reboot. StorEdge offers the option to save syslog data locally, or to send it to a syslogd server.

For additional information on this topic, please refer to the FAQ “How do I set up local or remote logging?”

System log message: “nfsd: error 'prog unavail' x.x.x.x (100227) nfs_acl”

This message is generated when an NFS client tries to write a POSIX style ACL. This is not supported by StorEdge at this time.

System log message: “ARP information overwritten”

This message is generated when StorEdge recognizes a new Ethernet (MAC) address associated with an IP address. As each NIC has only one Ethernet address, this is either an indication that the client is using port aggregation, or that there is a duplicate IP address on the network.

System log message: “statmon error: no statmon, remote x.x.x.x”

This message indicates that there was no reply from client x.x.x.x to an NFS status check. Usually due to a client reboot or lost connection.

System log message: “dac_get_dev_info: GetSCSIDevice failed FFFFFFF0”

This message indicates a SCSI device error. It is usually an indication of heavy system load, but not necessarily an indication of a problem.

2.35 Direct Attached Tape Libraries

The Sun StorEdge 5310 NAS supports specific SUN branded Tape Libraries. For the updated list refer to the WWWW for the Sun StorEdge 5310 NAS.

The following Sun Tape Libraries and Tape drives are supported:

TABLE 2-18 Supported Tape Libraries and Tape Drives

Tape Libraries	Tape Drives
L8	Ultrium LTO1
L25	Ultrium LTO2
L100	SDLT 320
L180	

2.35.1 SCSI ID Settings

When installing the Tape Devices, ensure that the SCSI IDs for the devices are set as follows:

Tape Libraries are generally set for SCSI ID of 0.

Tape devices need to be set higher than 0 to be recognized by the Sun StorEdge 5310 NAS.

The SCSI IDs are set on the libraries main LCD panel.

SCSI ID 7 is reserved for the on board Adaptec card. There are also settings for the on board Adaptec card that can be changed during boot up by typing control-a when prompted.

The settings for the Adaptec card are all factory defaults, except the Enable Disconnect settings. Set the Enable Disconnect setting to NO (the default is YES).

In cases where the settings need to be changed in the field, you must connect a monitor and keyboard to the Sun StorEdge 5310 NAS to change the settings.

2.36 StorEdge File Replicator

This section provides the following information:

- How does File Replicator work?
- The applications for File Replicator
- How do I set up File Replicator?

How Does File Replicator work?

Replicating allows you to duplicate any or all of the file volumes of one StorEdge server onto another StorEdge server. The source server is referred to as the active server and the target server is referred to as the mirror server.

In the event that the active server fails, the replicating file volumes on the mirror server can become available to network users within minutes. Once a mirror has been broken on the active server, the replicating file volume can be promoted, or made available for users, on the mirror server.

The replicating method used in the StorEdge is an asynchronous transaction-oriented mirror. Replicating is accomplished through the use of a large mirror buffer to queue file system transactions for transfer to the mirror system. Note that there is a performance cost associated with replicating, as writes to the master server must be done synchronously. Because the mirror is transaction-oriented, the integrity of the mirror file system is guaranteed, even during network interruptions or system outages.

What are the applications for File Replicator?

File Replicator can be used to help address the following data management challenges facing IT professionals today:

- Disaster Recovery
- Backup
- Data Distribution

Disaster Recovery

Without a reliance on slow tape media, File Replicator eliminates the need for lengthy tape restores. File Replicator enhances recovery time in case of a complete loss of data, as businesses can now access mission-critical data from an online backup on a mirror StorEdge.

Backup

A File Replicator target volume may be dedicated for backing up source volumes. File Replicator enhances operations by moving backup I/O to the remote volume. This shadow processing capability reduces CPU load on the production StorEdge, streamlining operations.

Data Distribution

For businesses with remote locations, File Replicator simplifies data distribution. StorEdges running File Replicator can be used to push data from a central location, such as a corporate headquarters, to a satellite office without relying on FTP or other passive file transfer methods.

How do I set up File Replicator?

Before You Begin Replicating

Before you begin, make sure you have the following:

Two StorEdge servers are required for replicating. The StorEdge servers may be of any model, and they don't need to be the same model.

The mirror server must contain an equal or larger amount of storage space than the file volumes to be replicating. This space must be raw unassigned and unformatted. The system files should reside on a volume which is not replicating. For this purpose, you should create a small system volume.

A dedicated network connection is required between the active and mirror StorEdge servers. The servers may be directly connected using a cross-over cable, or connected via a switch or router. If you are connecting the servers to a router, be sure to configure the static route setting to ensure that the replicating data is directed through a private route. If you are connecting the servers to a switch, create a VLAN for each server to isolate network traffic. This means that the mirror interface on each system must not be on the same subnet as the other interfaces in the system.

Both servers must have the same version of the operating system installed.

Both systems must have a File Replicator License installed.

To setup the NIC Card (both systems), do the following:

1. **Access the StorEdge via Telnet, serial console, or keyboard console.**
2. **Press [Enter] at the [menu] prompt and enter the administrator password.**
3. **Select option "A", "Host Name & Network."**

4. **Select option "1", "Edit fields."**

5. **Navigate through the fields with [Tab] or [Enter] until the "Role" field of the NIC that will be used for mirror is highlighted.**

6. **Select option "4", "Mirror" to change the role to mirror.**

To create Host File (both systems), do the following:

1. **From the main menu, enter "F", "Hosts".**

2. **Create a new host entry for the mirror interface selected above. For each system, choose a name similar to the hostname, such as, "host-M." For these host entries, use the IP address assigned to the NIC with the mirror role.**

3. **Select option "7", "Save Changes".**

To activate Mirror License (both systems), do the following:

1. **From the main menu, select option "4", "Licenses" to input a license key.**

2. **Select option "A", "replicating" under features.**

3. **Input the replicating key.**

   It must be exact and is case sensitive

   The StorEdge operating system verifies that the key is correct and provides the expiration date.

4. **Select option "7", "Save changes."**

After all configurations steps have been completed on both systems start the mirror. To start the mirror (Source System), do the following:

1. **Access the StorEdge via Telnet, serial console, or keyboard console.**

2. **Press [Enter] at the [menu] prompt and enter the administrator password.**

3. **Press the spacebar until "Mirrors" is displayed under "Extensions" at the lower right.**

4. **Select the letter corresponding to "Mirrors".**

5. **Select option "8", "Add Mirror".**

6. **Enter letter associated with the volume to be replicating.**

7. **Enter target host name in the host name field.**

8. **The IP address should auto fill, verify that it is correct.**

9. **Enter the desired size of the Mirror Buffer.**

   The mirror buffer stores file system write transactions while they are being transferred to the mirror server. The size of the mirror buffer depends on a variety of factors, but must be at least 100 MB. You may want to create a mirror buffer that is approximately 10% of the size of the file volume you are replicating. The size you choose should depend on how much information is being written to the file volume rather than the size of the file volume. The file volume free space on the active server will be reduced by the allocation size of the mirror buffer. Our general recommendation is 10% of the volume to be replicating.

10. **Select option "7", "Save Changes" to proceed.**

11. **To verify status, return to the Mirrors submenu, and view details by selecting the "Mirror" option, usually "A" or "B."**

    This display shows status complete and details of current operations.

2.37 StorEdge File Replicator Issues

Since File Replicator operates at the disk block level, the mirror system is an exact replica of the master system. However, since replicating operations are not strictly real time, the mirror system may lag the master by a time delta dependent on the speed and quality of the network. While this network lag may prevent the mirror system from being an exact copy of the master at any given point, the integrity of the mirror system is guaranteed at all times. Only complete file system transactions are replicating. In the course of creating a duplicate volume, a mirror goes through three main phases: creation, replication, or sync, and sequencing. File Replicator is a fault-tolerant technology. In all of the three main phases, the mirror handles errors with the intent of self-recovery as much as possible. When errors are encountered that are too severe for the mirror to handle on its own, it enters an ERROR state. In this state, user intervention is required to remedy the error and restart the mirror.

A list of all mirror states and their definitions can be found in the "Error Codes" section at the end of this document. The mirror states are very useful for troubleshooting.

Configuration

File Replicator requires that the NIC card to be used for File Replicator is designate as a mirror role, and that the network connection for this card be private. Also check that a host entry exists in both systems to provide a path to mirror and master.

Mirror Promoted on host

Once a volume has been promoted, the mirror cannot continue. Once a volume has been promoted it can not be a mirror again; however it can function as a master.

Waiting on host, link is down

This is typically a connection problem between the two system. Check the cables and the network connectivity. The master system will continue to send changed data to the buffer. Once the buffer is filled the mirror will crack.

Error Initializing mirror buffer

This error occurs when the master system does not have enough free space to create a mirror buffer. Try creating a smaller buffer. Minimum buffer size is 100mb.

Mirror stops when creating buffer

On the master system Sync status "displays syncing to host 0%" and system log displays "updating mirror and volume information" but nothing is happening.

1. **Go to the mirror system**
2. **Access the StorEdge via Telnet, serial console, or keyboard console.**
3. **When the menu prompt appears enter "admin" and then press the Enter key.**
4. **Enter the administrator password for the StorEdge and then press the Enter key.**
5. **From the command line enter menu and select option 2 system log.**

   The system log displays the following error message:

   Invalid password from host- The password provided on the master system is incorrect.
6. **Go back to master system and break the mirror.**
7. **Add a new mirror and enter the correct administrator password for the mirror.**

   Whenever a mirror problem exsists it is best to view both systems and their respective logs.

Not enough space to create mirror buffer

To check the amount of space available for a mirror buffer, proceed as follows:

From the CLI, enter "fsctl frags <volumename>."

A screen similar to the following is displayed:

```
/vol1: Analyzing volume.  Press ESC to abort ...

EXT        FREE PAGES

0          1244306

RANGEFRAGSLOW/HIGHPAGESSIZE(MB)

1-851/7190

9-3100/000

32-10242039/6312204

1024216399/1226668   12430674855

=========================================================

TOTAL2712443064860
```

The key data on this page are the third and fourth entries in the SIZE column. Adding these two numbers returns the largest mirror buffer that can be created on this volume. This is the total size of the fragments larger than 32 pages.

CHAPTER 3

Storage Arrays

This chapter instructs you on how to solve specific Storage Array problems with the Sun StorEdge 5310 NAS. It contains the following sections:

- "Array Overview" on page 3-1
- "Using the Array" on page 3-8
- "Troubleshooting and Recovery" on page 3-22
- "Relocating a Command Module" on page 3-31
- "Raid Storage Manager (RSM)" on page 3-44

3.1 Fibre Channel FC

3.1.1 Array Overview

This chapter describes the array command module its components.

3.1.1.1 Module Description

The array command module is a compact unit that provides high-capacity disk storage for Fibre Channel environments. This module supports two controllers and is available in both a deskside and a rackmount model.

The command module supports 14 drives within its enclosure, plus seven expansion drive modules, containing 14 drives each for a total of 112 drives.

Figure 3-1 on page 3-2 and Figure 3-2 on page 3-3 show front and back views of the command module and its components.

The front of the command module consists of the following components:

- Front bezel - Molded frame containing drive and global indicator lights and a mute button for the optional audible alarm feature
- Drives - Fourteen removable disk drives

The back of the command module contains the following components:

- Fans - Two removable fan housings, containing two fans each
- Controller - Two removable controllers
- Power supplies - Two removable power supplies

3.1.1.2 Controllers

The command module supports two controllers. Each controller attaches to hosts through fibre optic interface cables, and to drive modules through fibre optic or copper cables. Each controller also supports an Ethernet connection to a host for out-of-band management.

Figure 3-1 shows a controller. Each controller slides into the back of the module and has cache memory, which requires a rechargeable battery.

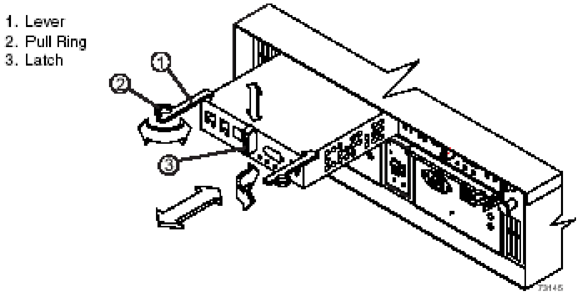

FIGURE 3-1 Controller

3.1.1.3 Controller Battery

Each controller contains a sealed, rechargeable 4-volt, lead acid battery. The battery provides backup power to the cache memory for up to three days in the event of a power loss. The service life of the battery is two years. Replace the battery every two years using the procedure described in “Replacing a Controller Battery” on page 7-29.

The battery performs a self-test at startup and every 25 hours thereafter. If needed, the battery will begin recharging at that time. Data caching starts after the battery completes its startup tests.

The Battery Charging/Charged light flashes during the startup self-test and when the battery is charging. It turns on and does not flash when the battery is fully charged, and turns off if the battery fails.

Figure 3-2 shows the controller labels. Each controller has a media access control (MAC) address label, located on the top or the front of the controller, and a battery label, located on top of the controller, which lists the battery installation and expiration dates.

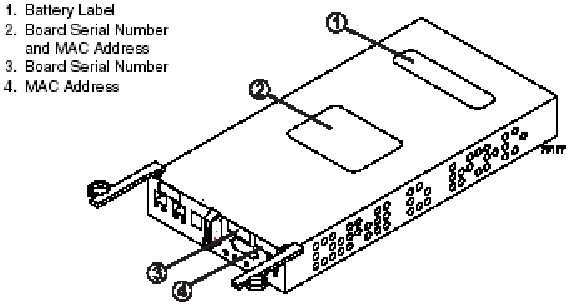

FIGURE 3-2 Label Locations on the Controller

3.1.1.4 Controller Memory

Note – IMPORTANT For specific information on the controller memory, refer to the Product Release Notes.

Each controller has 1 GB of memory for processor memory and data cache. The processor memory is used to store application data, while the data is in cache memory.

If caching has been enabled and data is in the cache, the Cache Active light on the controller turns on but does not flash. If caching is enabled and the Cache Active light never turns on during I/O activity, a cache memory failure or a battery failure has occurred.

Figure 3-3 shows the Cache Active light and the Battery Charging/Charged light.

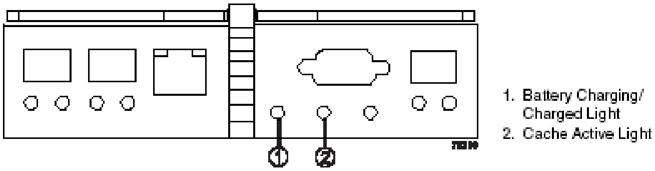

FIGURE 3-3 Battery Charging/Charged and Cache Active Lights

3.1.1.5 Drives

The command module supports up to 14 removable fibre Channel drives internally, plus up to seven expansion drive modules containing 14 drives each for a total of 112 drives.

Figure 3-4 shows the command module drive and its lights. Note that the drives in your command module may differ slightly in appearance from those shown. The variation will not affect the function of the drives.

Figure 3-5 array shows the physical locations of the drives, which are numbered 1 through 14, from left to right. When a drive is installed, the drive/tray slot designation is set automatically.

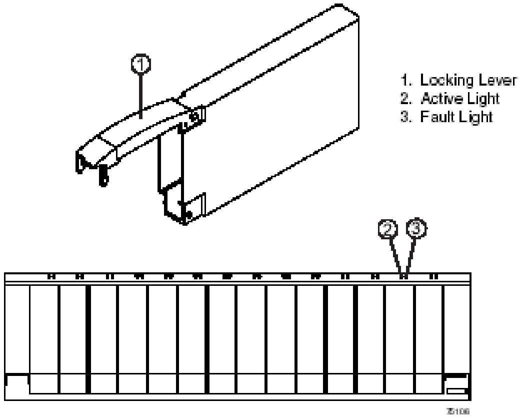

FIGURE 3-4 Drives and Lights

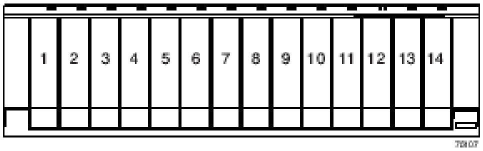

FIGURE 3-5 Drive Numbering – Rackmount Module

3.1.1.6 Fans

Each module has two removable fan housings. Each fan housing contains two fans. The fans provide redundant cooling, which means that if one of the fans in either fan housing fails, the remaining fans will continue to provide sufficient cooling to operate the command module.

Figure 3-6 shows a set of fans in a fan housing. The fans circulate air inside the drive module by pulling air in through the vents on the front and pushing the air out the vents on the back of each fan housing.

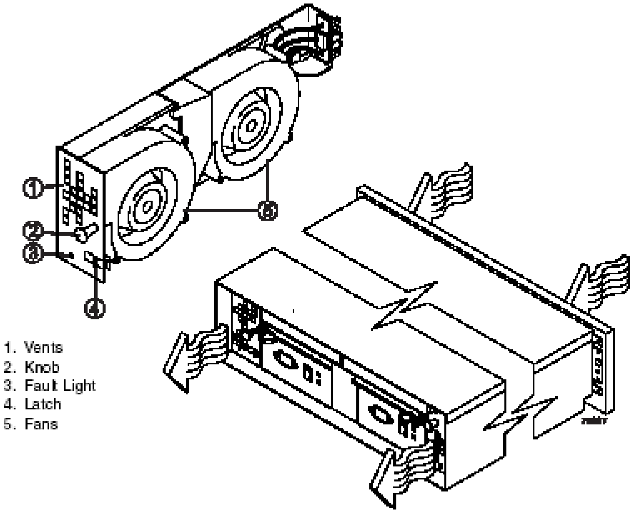

FIGURE 3-6 Fans and Airflow

3.1.1.7 Power Supplies

Each module contains two removable power supplies. The power supplies provide power to the internal components by converting incoming AC voltage to DC voltage. If one of the power supplies is turned off or malfunctions, the other power supply can maintain electrical power to the command module.

Figure 3-7 shows the power supplies, which are interchangeable by reversing the locking levers.

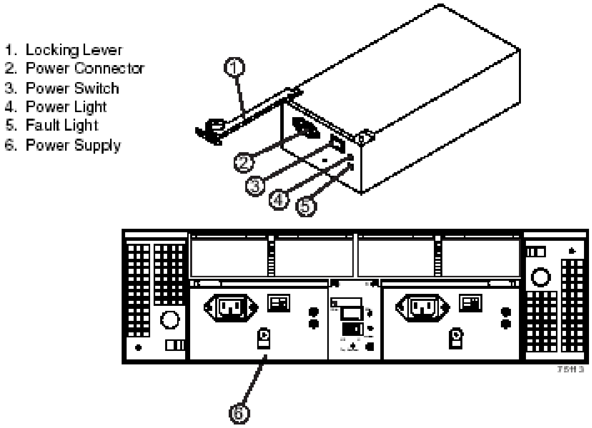

FIGURE 3-7 Power Supplies

3.1.1.8 SFP Transceivers and Cables

Each module supports either copper (for drive extension only) or fibre optic interface cables. If fibre optic cables are used, a small form-factor pluggable (SFP) transceiver must be installed in each interface connector on the controller where a fibre optic cable is to be installed.

Figure 3-8 shows an SFP transceiver and a fibre optic cable. Note that the SFP transceiver shown may look different from those shipped with your unit. The differences will not affect transceiver performance.

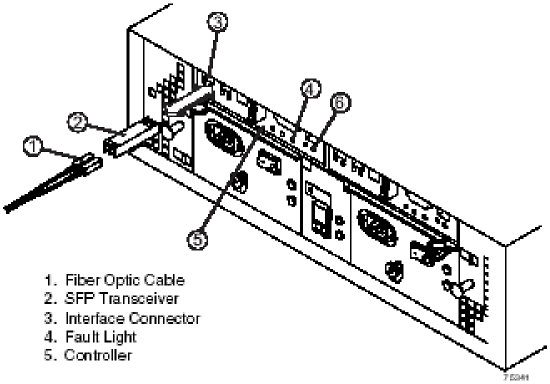

FIGURE 3-8 SFP Transceiver and fibre Optic Cable

3.1.1.9 Tray ID Switch

Note – IMPORTANT Each module in the storage array must have a unique tray ID.

The Tray ID switch is located between the power supplies. The Tray ID switch lets you assign each module a unique tray ID, which is required for proper operation of the storage array. The settings for each digit (X10 and X1) in the Tray ID range from 0 through 7.Recommended unique ID numbers range from 01 through 77.

Figure 3-9 shows the tray ID switch. The switch is covered with a metal cover plate secured with a screw. The screw and metal plate must be removed to view or change the Tray ID switch setting.

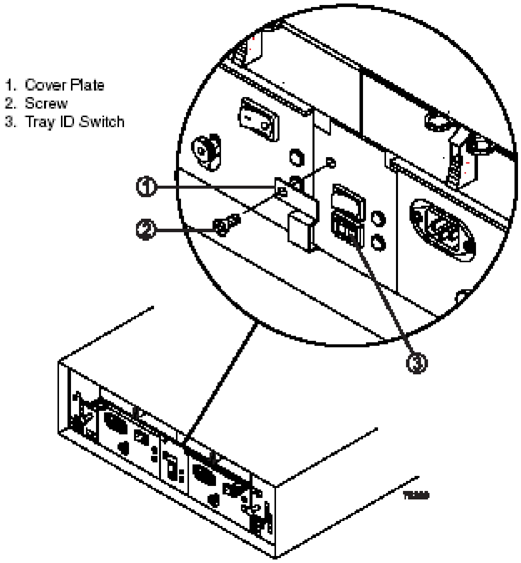

FIGURE 3-9 Tray ID Switch

3.1.2 Using the Array

This chapter provides general operating procedures for the array command module.

3.1.2.1 Removing and Replacing the Back Cover

Caution – Potential damage to cables. To prevent degraded performance or damaged cables, do not bend or pinch the cables between the module and the backfire.

Back covers are available only on deskside modules. If you have a rackmount module, you must open the hinged door or remove the access panel of the rackmount cabinet.Removing the back cover lets you see the indicator lights and access the cables and module components.To remove and replace the back cover, use the following procedure.

Figure 3-10 illustrates the procedure.

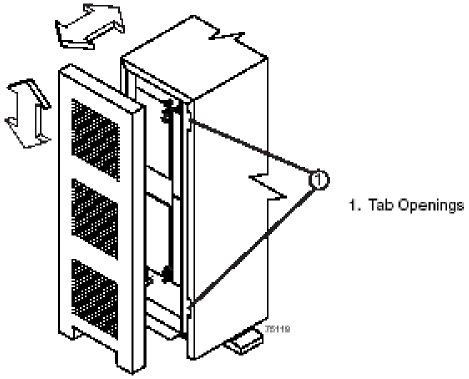

FIGURE 3-10 Removing and Replacing a Deskside Module Back Cover

1. **Remove the back cover.**

   Push the back cover up from the bottom, and pull the cover away from the module.

2. **Replace the back cover.**

   a. **Hold the back cover next to the back of the command module, and carefully route all cables through the opening at the bottom of the cover.If the opening is too small for all cabling, route some cables through the gap between the bottom of the module and the floor.**

   b. **Align the tabs on the cover with the tab openings in the module, and then push the cover onto the module to snap it into place.**

End Of Procedure

3.1.2.2 Turning On the Power

Use the following procedure to turn on power to one or more modules in a storage array at the initial startup or after a planned shutdown. Figure 3-11 on page 3-10 shows the locations of the power supply switches.

Caution – Potential damage to drives and data loss. Turning off the power then turning on the power without waiting for the disk drives to spin down can damage the drives and may cause data loss. Always let at least 30 seconds elapse from when you turn off the power until you turn it on again.

Note – Turn off both power switches on all modules in the configuration before connecting power cords or turning on the main circuit breakers.

1. **Remove the back cover, if needed.**
2. **Are the main circuit breakers in the cabinet turned on?**
   - Yes - Turn off both power switches on all modules that you intend to connect to the power.
   - No - Turn off both power switches on all modules in the cabinet.
3. **Connect the power cords to the power supplies on each module.**

Note – To ensure that the controllers acknowledge each attached drive module, turn on power to the drive modules before turning on power to the command module to ensure that the controllers acknowledge each attached drive module.

4. **Turn on both power switches on the back of each drive module.**

   The drives will not spin up until they receive a Start Unit command from the controller. While the drive modules power up, the green and amber lights on the front and the back of the modules will flash intermittently. Depending on your configuration, the drive modules can take from 20 seconds to several minutes to power up.

5. **Turn on both power switches on the back of the command module.**

   A command module can take up to 30 seconds to power up and up to 15 minutes to complete its controller battery self-test. During this time, the green and amber lights on the front and the back of the respective module flash intermittently.

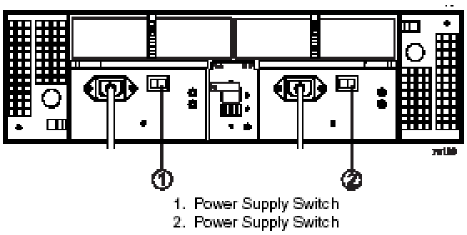

FIGURE 3-11 Power Supply Switches

6. **Check the status of each module and it components.**

a. **Note the status of the lights on the front and the back of each module.**

   A green light indicates a normal status; an amber light indicates a hardware fault.

b. **Open the Array Management Window for the storage array.**

c. **To view the status of its components, select the appropriate component button for each module in the Physical View of the Array Management Window. The status for each component will be either Optimal or Needs Attention.**

7. **Does each module display green lights only, and is the status Optimal on each module component?**

- Yes - Go to step 9.
- No - Go to step 8.

Note – If a fault requires you to power off an attached module, you may need to cycle the power on all remaining modules in the storage array. Contact technical support before powering off any attached modules.

---

8. **8 Diagnose and correct the fault.**

a. **To run the Recovery Guru, select the Recovery Guru toolbar button in the Array Management Window.**

b. **Complete the recovery procedure.**

   If the Recovery Guru directs you to replace a failed component, use the individual lights on the modules to locate the specific failed component. For troubleshooting procedures, refer to "Troubleshooting and Recovery" on page 3-22.

c. **When the recovery procedure is completed, ensure that the problem has been corrected. To re-run the Recovery Guru, select Recheck in the Recovery Guru.**

d. **If the problem persists, contact technical support.**

9. **Replace the back cover, if needed.**

End Of Procedure

3.1.2.3 Turning Off Power for a Planned Shutdown

Storage array modules are designed to run continuously, 24 hours a day. After you power up a module, it should remain on unless you need to perform an upgrade or service procedure that requires powering down.

Use the following procedure to turn on power to one or more modules in a storage array for a planned shutdown. Figure 3-11 on page 3-10 shows the location of the power supply switches.

1. **Stop I/O activity to all modules.**

2. **Remove the front cover from the command module, if applicable.**

3. **Determine the status of each module and its components.**

   a. **Note the status of the lights on the front and the back of each module.**

   A green light indicates a normal status; an amber light indicates a hardware fault.

   b. **Open the Array Management Window for the storage array.**

   c. **To view the status of its components, select the appropriate component button for each module in the Physical View of the Array Management Window.**

   The status for each component will be either Optimal or Needs Attention.

4. **Does each module display green lights only and is the status Optimal on each module component?**

   - Yes - Go to step 6.
   - No - Go to step 5.

Note – If the fault requires you to power off an attached drive module, you may need to cycle the power on all remaining modules in the storage array. Before you power off any attached module, contact technical support.

---

5. **Diagnose and correct the fault.**

   a. **To run the Recovery Guru, select the Recovery Guru toolbar button in the Array Management Window.**

   b. **Complete the recovery procedure.**

   If the Recovery Guru directs you to replace a failed component, use the individual lights on the modules to locate the specific failed component. For troubleshooting procedures, refer to "Troubleshooting and Recovery" on page 3-22.

   c. **When the recovery procedure is completed, select Recheck in the Recovery Guru to re-run the Recovery Guru and to ensure that the problem has been corrected.**

   d. **If the problem persists, contact technical support.**

Caution – Potential data corruption. An abrupt power loss to any module in a storage array can cause data corruption, especially if the power loss occurs when data is being written to a disk or if the power loss occurs when write back caching is enabled and data is being downloaded to cache memory. Before turning off power to the modules, always wait until the Cache Active light turns off and all drive Active lights stop flashing.

---

6. **Check the lights on the front and the back of each drive module, and verify that all drive Active lights are on but not flashing.**

   If one or more drive Active lights are flashing, then data is being written to or from the disks. Wait for all drive Active lights to stop flashing, and then go to step 7.

7. **Check the Cache Active light, then choose one of the following steps, based on the status of the light.**

   Figure 3-12 and Table 3-1shows the locations of the status lights.

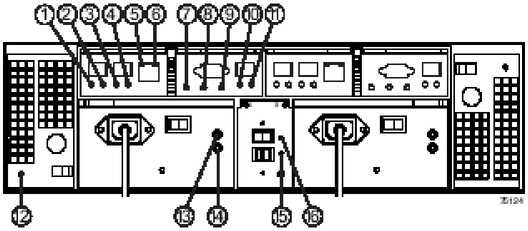

FIGURE 3-12 Lights on the Back of a Command Module

TABLE 3-1 Lights on the Back of a Command Module

Location	Component Light	Color	Normal Status	Problem Status	Procedure
			Controller		
1	Host Connector 1 Link Indicator	Green	On	Off	"Replacing a Drive" on page 7-36
2	Host Connector 1 Speed Indicator	Green	On - 2 Gb/s data rate Off - 1 Gb/s data rate	Not Applicable	
3	Host Connector 2 Link Indicator	Green	On	Off	
4	Host Connector 2 Speed Indicator	Green	On - 2 Gb/s data rate Off - 1 Gb/s data rate	Not Applicable	
5	Ethernet link indicator	Green	On - Connection active Off - Connection inactive	Not Applicable	Not Applicable
6	Ethernet 100Base TX indicator	Green	On - 100BaseTX connection Off - 10BaseT (or inactive)	Not Applicable	
7	Battery Charging/ Charged	Green	On -battery charged Flashing - battery charging	Off	"Replacing a Controller Battery" on page 7-29

TABLE 3-1 Lights on the Back of a Command Module

Location	Component Light	Color	Normal Status	Problem Status	Procedure
8	Cache Active	Green	On	Off (if cache enabled)	"Replacing a Power Supply" on page 7-41
9	Fault	Amber	Off	On	
10	Drive Link	Green	On	Off	
11	Expansion Port Bypass	Amber	Off	On	
			Fan		
12	Fan Fault	Amber	Off	On	"Replacing a Fan" on page 7-39
			Power Supply		
13	Power	Green	On	Off	"Replacing a Power Supply" on page 7-41
14	Fault	Amber	Off	On	
			Link Rate		
15	2 Gb/s Link Rate	Green	On	Not Applicable	Not Applicable
16	Tray ID Conflict	Amber	Off	On - Tray IDs incorrect	"Setting the Tray ID Switch" on page 3-29

Cache Active light is off - The cache contains no data. Go to step 8.

Cache Active light is illuminated - Write-back caching is enabled and data is in the cache. Wait for the data to clear from the cache memory and for the Fast Write Cache or Cache Active light to turn off, and then go to step 8. For information on cache memory protection and settings, refer to the Array Management Window online help.

Caution – Potential damage to drives and data loss. Turning the power off and on without waiting for the disk drives to spin down can damage the drives and may cause data loss. Always let at least 30 seconds elapse from when you turn off the power until you turn it on again.

8. **Turn off both power switches on the back of the command module.**

9. **Turn off both power switches on the back of each drive module**

   End Of Procedure

3.1.2.4 Turning Off Power for an Unplanned Shutdown

Storage array modules are designed to run continuously, 24 hours a day. Certain situations, however, may require you to shut down all storage array modules quickly. These situations might include a power failure or emergency because of a fire, a flood, extreme weather conditions, some other hazardous circumstance, or a power supply shutdown caused byoverheating.Use the following procedure to turn off power to all modules in a storage array for an unplanned shutdown. Figure 3-11 on page 3-10 shows the locations of the power supply switches.

Caution – Potential damage to drives and data loss. Turning the power off and on without waiting for the disk drives to spin down can damage the drives and may cause data loss. Always let at least 30 seconds elapse from when you turn off the power until you turn it on again.

---

1. **Stop all I/O activity to the command module and attached drives.**
2. **Remove the front cover from the command module, if applicable.**

Caution – Potential data corruption. Turning off the power when an amber fault light is illuminated can cause data corruption. To prevent data corruption, always check for faults and correct all problems before turning off the power.

---

3. **Determine the status of each module and its components.**

   a. **Check the lights on the front and the back of each module.**

   A green light indicates a normal status; an amber light indicates a hardware fault.

   b. **Open the Array Management Window for the storage array.**

   c. **To view the status of its components, select the appropriate component button for each module in the Physical View of the Array Management Window.**

   The status for each component will be either Optimal or Needs Attention.

4. **Does each module display green lights only, and is the status Optimal on each module component?**

   Yes - Go to step 6.

   No - Go to step 5.

5. **Diagnose and correct the fault.**

a. **To run the Recovery Guru, select the Recovery Guru toolbar button in the Array Management Window.**

b. **Complete the recovery procedure.**

If the Recovery Guru directs you to replace a failed component, use the individual lights on the modules to locate the failed component. For troubleshooting procedures, refer to "Troubleshooting and Recovery" on page 3-22.

Note – If a fault requires you to power off an attached module, you may need to cycle the power on all remaining modules in the storage array. Before you power off any attached modules, contact technical support.

---

c. **When the recovery procedure is completed, select Recheck in the Recovery Guru to re-run the Recovery Guru and to ensure that the problem has been corrected.**

d. **If the problem persists, contact technical support.**

Caution – Potential data corruption. An abrupt power loss to any module in a storage array can cause data corruption, especially if the power loss occurs when data is being written to a disk or if the power loss occurs when write back caching is enabled and data is being downloaded to cache memory. Before turning off power to the modules, always wait until the Cache Active light turns off and all drive Active lights stop flashing.

---

6. **Check the lights on the front of each attached drive module, and verify that all drive Active lights are on but not flashing.**

   If one or more drive Active lights are flashing, then data is being written to or from the disks. Wait for all drive Active lights to stop flashing, and then go to step 7.

7. **Check the Cache Active light, and then choose one of the following steps, based on the status of the light.**

   Cache Active light is off - The cache contains no data. Go to step 8.

   Cache Active light is illuminated - Write-back caching is enabled, and data is in the cache. Wait for the data to clear from the cache memory and for the Fast WriteCache or Cache Active light to turn off, and then go to step 8. For information on cache memory protection and settings, refer to the Array Management Window online help.

8. **Check the lights on the back of the command module, and then choose one of the following steps, based on the status of the lights.**

   The Host Link and Host Speed lights as well as the Power light are illuminated; all others are off - The module Link or 100BT light might be on if the command module is using an Ethernet connection. Go to step 9.

   One or more amber lights are on - Do not continue with the power off procedure until you have corrected the fault. To diagnose the problem, go to "Troubleshooting the Module" on page 3-22.

9. **Turn off the main circuit breaker in the cabinet.**

10. **Turn off all power switches on all modules affected by the unplanned shutdown.**

11. **Unplug both power cables from each module.**

12. **Replace the front cover on the command module, if applicable.**

13. **After the emergency situation has passed, perform the power recovery procedure in "Restoring Power After an Unplanned Shutdown" on page 3-18.**

    End Of Procedure

3.1.2.5 Restoring Power After an Unplanned Shutdown

Caution – WARNING! Risk of severe electrical shock. Never turn on the power to any equipment if there is evidence of fire, water, or structural damage. Doing so may cause severe electrical shock. After the emergency situation has passed or power is restored to the building, always check all the equipment for physical damage first.

Use the following procedure to restore power to all modules in a storage array. Figure 3-11 on page 3-10 shows the locations of the power switches.

1. **Remove the back cover, if needed.**

2. **Is there evidence of damage to any components or cables?**

   Yes - Do not continue with this procedure if you find any evidence of damage.Call the factory or appropriate service organization for assistance. Depending on the current service agreements, you may need to return the equipment to the factory or local service center for repair.

   No - Go to step 3.

Caution – Potential data loss or corruption. Ensure that all module power switches are turned off before resetting the circuit breakers. Failure to do so can cause data loss or corruption.

---

3. **Verify that both power switches on all modules in the cabinet are turned off.**

4. **Are the main circuit breakers in the cabinet turned off?**

   Yes -Turn on the main circuit breakers in the cabinet.

   No - Reset the main circuit breakers in the cabinet.

5. **Connect the power cables to both power supplies in each module.**

Note – IMPORTANT To ensure that the controllers acknowledge each attached drive module, it is recommended that you turn on power to the drive modules before turning on power to the command module.

---

6. **Turn on both power switches on the back of each drive module.**

   The drives will not spin up until they receive a Start Unit command from the controller. During this time, the green and amber lights on the front and the back of the modules will flash intermittently. Depending on your configuration, the drive modules can take from 20 seconds to several minutes to power up.

7. **Turn on both power switches on the back of each command module.**

   An command module can take up to 30 seconds to power up and up to 15 minutes to complete its controller battery self-test. During this time, the green and amber lights on the front and the back of the command module flash intermittently.

8. **Check the status of each module in the storage array and its components.**

   a. **Note the status of the lights on the front of each module.**

   b. **Open the Array Management Window for the storage array.**

   c. **To view the status of its components, select the appropriate component button for each module in the Physical View of the Array Management Window.**

      The status for each component will be either Optimal or Needs Attention.

9. **Does each module display green lights only, and is the status Optimal on each module component?**

   Yes - Go to step 11.

   No - Go to step 10.

Note – IMPORTANT If a fault requires you to power off an attached module, you may need to cycle the power on all remaining modules in the storage array. Contact technical support before powering off any attached modules.

---

10. **Diagnose and correct the fault.**

a. **To run the Recovery Guru, select the Recovery Guru toolbar button in the Array Management Window.**

b. **Complete the recovery procedure.**

If the Recovery Guru directs you to replace a failed component, use the individual lights on the modules to locate the specific failed component. For troubleshooting procedures, refer to "Troubleshooting and Recovery" on page 3-22.

c. **When the procedure is completed, select Recheck in the Recovery Guru to re-run the Recovery Guru and to ensure that the problem has been corrected.**

d. **If the problem persists, contact technical support.**

11. **Replace the cover on the command module, if applicable.**

End Of Procedure

3.1.2.6 Responding to the Optional Audible Alarm

The command module may have an optional audible alarm. Modules with the optional alarm have a Mute button on the front bezel, below the Power and Global Fault lights. Figure 3-13 on page 3-20 shows the locations of the indicator lights and the mute button.

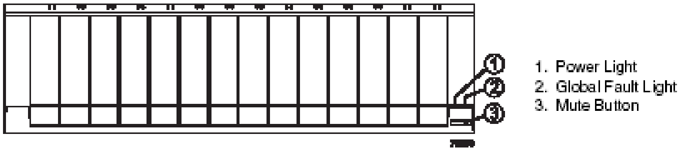

FIGURE 3-13 Alarm Mute Button

When one of the following conditions occurs, the alarm sounds and the Global Fault light illuminates:

- Hardware malfunction in a command module - Malfunctions include: overheating conditions, missing fans, or component failures (failed drives, environmental services monitors [ESMs] or controllers, power supplies, or fans).
- Transmission failures - Transmission failures include: I/O transmission problems with the Small Form-factor Pluggable (SFP) transceivers and the interface cables.

The alarm does not affect the operation of the indicator lights or the ability of the module to report errors to the host.

Use the following procedure to turn off the alarm and to identify the problem that caused the alarm to sound.

1. **Locate the module with the alarm sounding and the amber Global Fault light illuminated.**

2. **Press the Mute button to turn off the alarm. If another fault occurs, the alarm will sound again.**

3. **Remove the back cover, if needed.**

4. **Determine the status of each module and its components.**

   a. **Note the status of the lights on the front and the back of each module.**

      A green light indicates a normal status; an amber light indicates a hardware fault.

   b. **Open the Array Management Window for the storage array.**

   c. **To view the status of its components, select the appropriate component button for each module in the Physical View of the Array Management Window.**

      The status for each component will be either Optimal or Needs Attention.

5. **Does each module display green lights only, and is the status Optimal on each module component?**

   - Yes - Go to step 7.
   - No - Go to step 6.

Note – IMPORTANT If a fault requires you to power off an attached module, you may need to cycle the power on all other modules in the storage array. Contact technical support before powering off any attached modules.

---

6. **Diagnose and correct the fault.**

   a. **To run the Recovery Guru, select the Recovery Guru toolbar button in the Array Management Window.**

   b. **Complete the recovery procedure.**

      If the Recovery Guru directs you to replace a failed component, use the individual lights on the modules to locate the specific failed component. For troubleshooting procedures, refer to “Troubleshooting and Recovery” on page 3-22.

   c. **When the recovery procedure is completed, select Recheck in the Recovery Guru to re-run the Recovery Guru and to ensure that the problem has been corrected.**

   d. **If the problem persists, contact technical support.**

7. **7 Replace the cover, if needed.**

End Of Procedure

3.2 Troubleshooting and Recovery

This chapter provides procedures for diagnosing and correcting problems with the array command module.

3.2.1 Troubleshooting the Module

The storage management software provides the best way to monitor the modules, diagnose problems, and to recover from hardware failures. You should run the storage management software continuously and check the status of the storage array frequently. Use the following procedure to check the status and identify problems with the command module.

1. **Open the Array Management Window for this storage array.**

2. **Select the component button for each module in this array. View the status of all components.**

   The status for each component will be either Optimal or Needs Attention.

3. **Do any components have a Needs Attention status?**

   - Yes - Go to step 4.
   - No - All components are Optimal. Go to step 5.

Note – IMPORTANT If a fault requires you to power off an attached module, you may need to cycle the power on all remaining modules in the storage array. Contact technical support before powering off any attached modules.

---

4. **To correct the problem, select the Recovery Guru toolbar button. Perform the procedure in the Recovery Guru to correct the problem.**

   The Recovery Guru may direct you to replace the failed component. If so, go to step 7.

5. **Check the lights on the front of the module. A green light indicates a normal status; an amber light indicates a hardware fault.**

6. **Are any amber lights on?**

   - Yes - Go to step 7.

- No - You are finished with this procedure. If you are still experiencing a problem with this storage array, go to step 10.

7. **Remove the cover.**

8. **If needed, turn off the alarm.**

9. **Check all of the lights on the front and the back of each module. Figure 3-14 and Figure 3-15 show the locations of indicator lights. Table 3-2 and Table 3-3 refer you to the appropriate procedures for various fault indicators. Go to the page indicated in the tables, and perform the procedure for replacing any failed component as needed.**

10. **If you are still experiencing a problem with this storage array, create, save, and print a storage array profile, and then call technical support for assistance.**

The storage array profile may be helpful when troubleshooting.

End of Procedure

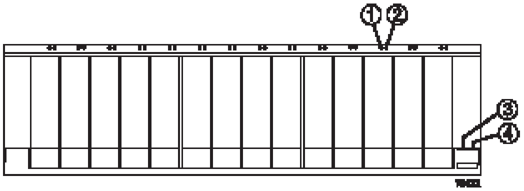

FIGURE 3-14 Lights on the Front of a Command Module

TABLE 3-2 Lights on the Front of a Command Module

Location	Component Light	Color	Normal Status	Problem Status	Procedure
Drives					
1	Drive Active	Green	On, not flashing - no data is being processed Flashing - data is being processed	Off	“Replacing a Drive” on page 7-36
2	Drive Fault	Amber	Off Flashing - drive, volume, or storage array locate function On, not flashing	On, not flashing	

TABLE 3-2 Lights on the Front of a Command Module

Location	Component Light	Color	Normal Status	Problem Status	Procedure
			Global		
3	Global Power	Green	On	Off	"Recovering from an Overheated Power Supply" on page 3-26
4	Global Fault	Amber	Off	On	Recovery Guru procedure

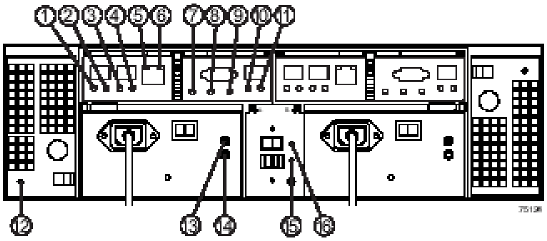

FIGURE 3-15 Lights on the Back of a Command Module

TABLE 3-3 Lights on the Back of a Command Module

Location	Component Light	Color	Normal Status	Problem Status	Procedure
			Controller		

TABLE 3-3 Lights on the Back of a Command Module

Location	Component Light	Color	Normal Status	Problem Status	Procedure
1	Host Connector 1 Link Indicator	Green	On	Off	“Replacing a Drive” on page 7-36
2	Host Connector 1 Speed Indicator	Green	On - 2 Gb/s data rate Off - 1 Gb/s data rate	Not Applicable	
3	Host Connector 2 Link Indicator	Green	On	Off	
4	Host Connector 2 Speed Indicator	Green	On - 2 Gb/s data rate Off - 1 Gb/s data rate	Not Applicable	
5	Ethernet link indicator	Green	On - Connection active Off - Connection inactive	Not Applicable	Not Applicable
6	Ethernet 100Base TX indicator	Green	On - 100BaseTX connection Off - 10BaseT (or inactive)	Not Applicable	
7	Battery Charging/ Charged	Green	On -battery charged Flashing - battery charging	Off	“Replacing a Controller Battery” on page 7-29
8	Cache Active	Green	On	Off (if cache enabled)	“Replacing a Power Supply” on page 7-41
9	Fault	Amber	Off	On	
10	Drive Link	Green	On	Off	
11	Expansion Port Bypass	Amber	Off	On	
			Fan		

TABLE 3-3 Lights on the Back of a Command Module

Location	Component Light	Color	Normal Status	Problem Status	Procedure
12	Fan Fault	Amber	Off	On	"Replacing a Fan" on page 7-39
			Power Supply		
13	Power	Green	On	Off	"Replacing a Power Supply" on page 7-41
14	Fault	Amber	Off	On	
			Link Rate		
15	2 Gb/s Link Rate	Green	On	Not Applicable	Not Applicable
16	Tray ID Conflict	Amber	Off	On - Tray IDs incorrect	"Setting the Tray ID Switch" on page 3-29

3.2.2 Recovering from an Overheated Power Supply

All modules contain two power supplies, each containing a built-in temperature sensor designed to prevent the power supplies from overheating. Under normal operating conditions, with an ambient air temperature range of 5º C to 40º C (40º F to 104º F), the fans in the module maintain a proper operating temperature inside the module.

If the internal temperature exceeds 70° C (158° F), the power supplies shut down automatically. If both power supplies shut down because of overheating, the module has no power and all indicator lights are off.

The following factors can cause the power supplies to overheat:

- An unusually high room temperature
- A fan failures
- Defective circuitry in the power supply
- A blocked air vent
- A failure in another device in the command module cabinet

If a fan failure has caused the overheating, the fan Fault lights and the power supply Fault lights turn on.

If the module temperature exceeds 45° C (113° F), the storage management software displays a Needs Attention icon in the Array Management Window.

If event monitoring is enabled, and if event notification is configured, the software also issues one or both of the following critical problem notifications:

- If one power supply shuts down, the storage management software will display a Needs Attention status in the Array Management Window.
- If both power supplies shut down, the module will shut down, and the storage management software will display a Not Responding status in the Array Management Window.

Use the following procedure to resume normal operation after a power supply shutdown.

Figure 3-16 on page 3-28 shows the locations of the power supply switches.

Caution – Risk of damage from overheating. Power supplies automatically shut down when the air temperature inside the cabinet reaches 70° C (158° F). If the power supplies have shut down, immediately remove all cabinet panels to help cool the cabinet air temperature and prevent damage to the modules.

1. **Remove the cover.**

2. **If needed, turn off the alarm.**

3. **Did you use the procedure "Troubleshooting the Module" on page 3-22 to identify an overheating problem?**
   - Yes - Go to step 4.
   - No - Perform the procedure "Troubleshooting the Module" on page 3-22 to verify that the power supplies have shut down because of an overheating problem. When finished, go to step 4.

4. **Stop I/O activity to the module and all attached modules.**

5. **To alleviate the overheating problem, perform any or all of the following cooling measures:**

   a. **Immediately removing all panels from the cabinet.**

   b. **Using external fans to cool the area.**

   c. **Shut down the power to the modules in the cabinet, using the procedure described in "Turning Off Power for a Planned Shutdown" on page 3-11.**

6. **Wait for the air temperature in and around the module to cool.**

   After the temperature inside the power supplies cools to below 70° C (158° F), the module is capable of a power-up recovery without operator intervention. After the air has cooled, the power supplies should turn on automatically. If the power supplies restart automatically, the controllers reset and return to normal operation.

7. **Did the power supplies restart automatically?**

- Yes - Go to step 9.
- No - Go to step 8.

8. **Turn on both power switches on the back of each drive module connected to the command module.**

   The drive modules will not spin up until they receive a Start Unit command from the controller. While the drive modules power up, the lights on the fronts and backs of the modules will flash intermittently. Depending on your configuration, it can take between 20 seconds and several minutes for the drive modules to power up.

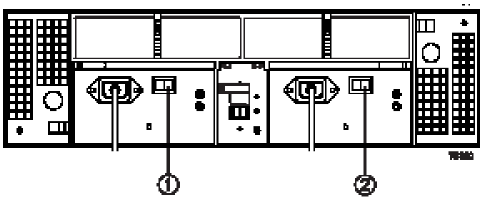

1. Power Supply Switch
2. Power Supply Switch

FIGURE 3-16 Power Supply Switches

9. **Turn on both power switches on the back of the command module.**

   An command module can take 30 seconds to power up and up to 15 minutes for the battery self-test to complete. During this time, the lights on the front and the back of the modules flash intermittently.

10. **Check the status of each module and its components.**

    a. **Note the status of the indicator lights on the front and the back of each module. A green light indicates a normal status; an amber light indicates a hardware fault.**

    b. **Open the Array Management Window for the storage array.**

    c. **Select the appropriate component button for each module in the Physical View of the Array Management Window to view the status of its components.**

    The status for each component will be either Optimal or Needs Attention.

11. **Does each module display green lights only and is the status Optimal on each module component?**

    - Yes - Go to step 13.
    - No - Go to step 12.

12. **Diagnose and correct the fault.**

a. **To run the Recovery Guru, select the Recovery Guru toolbar button in the Array Management Window.**

b. **Complete the recovery procedure.**

If the Recovery Guru directs you to replace a failed component, use the individual lights on the modules to locate the specific failed component.

Figure 3-14 on page 3-23 and Figure 3-15 on page 3-24 show the locations of indicator lights. Table 3-2 and Table 3-3 refer you to the appropriate procedures for various fault indicators. Go to the page indicated in the tables, and perform the procedure for replacing any failed component as needed.

c. **When the procedure is completed, select Recheck in the Recovery Guru to re-run the Recovery Guru and ensure that the problem has been corrected.**

d. **If the problem persists, contact technical support.**

13. **Replace the cover, if needed.**

End Of Procedure

3.2.3 Setting the Tray ID Switch

Use the following procedure to set the Tray ID switch if a Tray ID conflict is indicated.

Figure 3-17 shows the location of the Tray ID switch.

Caution – Electrostatic discharge damage to sensitive components. To prevent electrostatic discharge damage to the module, use proper antistatic protection when handling the module components.

Note – IMPORTANT Each module must have its own unique Tray ID.

1. **Put on antistatic protection.**
2. **Locate the Tray ID switch on the back of the module between the power supplies.**
3. **Remove the screw and the switch cover.**
4. **Set the tray ID switch to the desired switch setting.**

   The settings for each Tray ID digit (X10 and X1) range from 0 through 7.

   Recommended unique IDs range from 01 through 77.
5. **Replace the switch cover and tighten the screw to secure it into place.**

6. **If applicable, repeat for all other modules in the storage array.**

   End Of Procedure

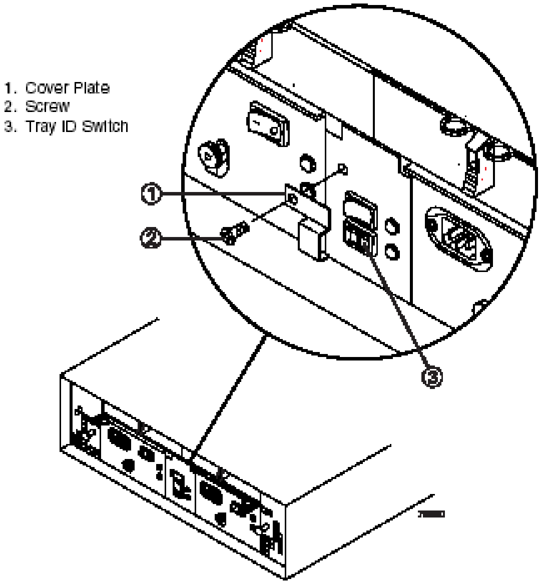

FIGURE 3-17 Setting the Tray ID Switch

3.2.4 Verifying the Link Rate Setting

Use the following procedure to verify the Link Rate setting if a link rate problem is indicated. Figure 3-18 shows the location of the Link Rate switch.

Caution – Electrostatic discharge damage to sensitive components. To prevent electrostatic discharge damage to the module, use proper antistatic protection when handling the module components.

---

1. **Put on antistatic protection.**
2. **Remove the screw and the switch cover from the Link Rate switch.**
3. **Verify that the Link Rate Switch is set to 2 Gb/s.**

   The Link Rate switch is not active. This switch is pre-set to 2 Gb/s at the factory, and defaults to 2 Gb/s.

4. **Replace the switch cover and tighten the screw to secure it into place.**

5. **If applicable, repeat for all other modules in the storage array.**

   End of Procedure

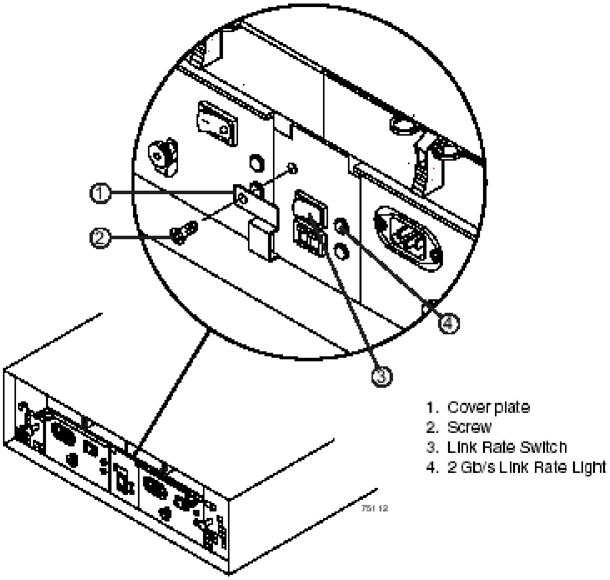

FIGURE 3-18 Verifying the Link Rate Setting

3.3 Relocating a Command Module

This chapter provides procedures for upgrading an E2600 command module for greater storage capacity and guidelines for relocating an E2600 command module.

3.3.1 Upgrade Requirements

You can upgrade an command module for greater storage capacity, but doing so requires careful planning in order to prevent data loss or command module failure. Consider the following cautions before starting any upgrade procedure.

Caution – Potential data loss or data corruption. Never insert drives into a command module without first confirming the drive firmware level. Inserting a drive with the incorrect firmware level may cause data loss or data corruption. For information on supported drive firmware levels, contact technical support.

Caution – Potential command module failure. Use of nonsupported drives in the command modules can cause the command module to fail.

Caution – Potential data loss or data corruption. In configurations with mixed command modules, command modules, or drive modules, all modules must be operating at the same speed. Refer to the Product Release Notes for any model-specific restrictions.

3.3.1.1 Upgrade Methods

You can upgrade the command module either by adding new drives or replacing existing drives. This section describes each method.

Adding New Drives

You can increase the capacity of an command module by installing the additional drives into empty slots in the module. Additional drives can be installed while the command module is in operation. To use this upgrade method, refer to "Adding New Drives to Empty Slots" on page 3-33.

Replace Existing Drives with Greater Capacity Drives

Caution – Potential data loss. Using the wrong drive upgrade procedure can cause data loss. If you are upgrading drives containing RAID 0 volumes, you must use the procedure for replacing all of the drives at once. If you are upgrading a drive containing RAID 1, 3, or 5 volumes, you may use either upgrade procedure.

When replacing existing drives with greater capacity drives, two methods may be used: replacing all the drives at the same time or replacing one drive at a time. The method you choose will depend on the RAID level you are using on the storage array, the amount of time you can afford to keep the command module offline, and the method that most closely matches the upgrade procedure recommended in the storage management software and this guide.

Replace All Drives at the Same Time

If you are upgrading drives containing RAID 0 volumes, you must use this method. This method requires you to back up the command module and turn off the power to the storage array before replacing the drives. After replacing all the drives, you must reconfigure the command module and restore the data from backup media. This is the safest way to exchange drives without losing data. However, this method may take considerable time to complete because of the backup, reconfiguration, and restoration procedures. Also, other users will be unable to use the command module until you finish the procedure. To use this upgrade method, refer to “Replacing All Drives at the Same Time” on page 3-36.

Replace One Drive at a Time

This method works only on drives containing redundant volumes, RAID 1, 3, or 5. If you are upgrading drives containing RAID 0 volumes, you must not use this method. This procedure lets you replace the drives while the command module is in operation, eliminating the necessity to shut down the command module. You manually fail each drive, replace it, and wait for the system to restore data to the new drive before installing the next drive. After installing all of the new drives, you configure them to create additional drive space.

Depending on your configuration, the reconfiguration procedure for this method may require considerable time to complete. Furthermore, you can lose data if the storage array reconfiguration or drive restoration fails. For this reason, you should back up all data on the command module before using this upgrade method. This will safeguard your data if the reconfiguration or restoration fails, or if the new drive malfunctions. To use this upgrade method, refer to “Replacing One Drive at a Time” on page 3-39.

3.3.2 Adding New Drives to Empty Slots

Use the following procedure to install additional drives into empty slots in the command module Figure 3-19 on page 3-35 illustrates inserting and removing a drive. You can install additional drives while the command module is in operation.

Caution – Potential volume group failure. The command module supports a maximum of eight drive modules per loop (112 drives maximum). Do not install new drives into the empty drive slots in the command module enclosure if the module is already at the maximum configuration. Doing so will exceed the fibre channel protocol limit and cause volume groups to fail.

Caution – Electrostatic discharge damage to sensitive components. To prevent electrostatic discharge damage to the module, use proper antistatic protection when handling the module components.

Caution – Potential damage to drive components. Drives can be damaged by bumping them into other objects or surfaces. When removing or installing a drive into the drive module, place your hand under the drive to support its weight. Place drives on an antistatic, cushioned surface.

1. **Put on antistatic protection.**
2. **Unpack the new drives that you intend to install.**

   Set the new drives on a dry, level surface, away from magnetic fields. Save all packing materials in case you need to return the drives.
3. **Locate the blank drives in the command module.**

   Check the Drive Active light on the front of the command module. On an active drive, a green light will be on or flashing. On a blank drive, the light will be off.

Note – IMPORTANT When replacing a blank drive with a new drive, replace the first available blank drive from the left, looking at the front of the module. Functional drives occupy the slots to the left; blank drives occupy the slots to the right.

Caution – Risk of potential data loss. Removing the wrong drive can cause data loss. Remove only blank drives. If you accidentally remove an active drive, wait at least 30 seconds, and then reinstall it. For further recovery procedures, refer to your storage management software.

4. **Lift the locking lever on the blank drive and remove the drive from the slot.**
5. **Slide the new drive all the way into the empty slot and close the drive lever.**

   As the drive spins up, the Fault lights may flash intermittently. A flashing Active light indicates that data is being restored to the new drive.

6. **Repeat step 4 and step 5 to install each new drive.**

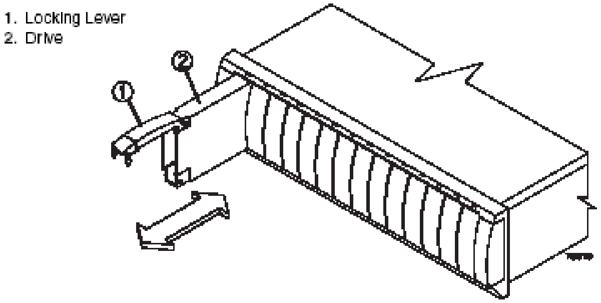

FIGURE 3-19 Removing and Installing a Drive

7. **Based on the status of the Active and Fault lights, choose one of the following steps:**
   - Active lights are on while Fault lights are off - Go to step 9.
   - Active lights are off while Fault lights are off - The drive may be installed incorrectly. Remove the drive, wait 30 seconds, and then reinstall it. Go to step 8.
   - Fault lights are on - The new drive may be defective. Replace it with another new drive, and then go to step 8.

8. **Did this correct the problem?**
   - Yes - Go to step 9.
   - No - Select the Recovery Guru toolbar button in the Array Management Window and complete the recovery procedure. If the problem persists, contact technical support.

9. **Check the module status using the storage management software.**

   Select the appropriate component button for each module in the Physical View of the Array Management Window to view the status of its components. The status for each component will be either Optimal or Needs Attention.

10. **Does any module component have a Needs Attention status?**
    - Yes - Go to step 11.
    - No - Go to step 12.

11. **To run the Recovery Guru, select the Recovery Guru toolbar button in the Array Management Window.**

a. **Complete the recovery procedure.**

If the Recovery Guru directs you to replace a failed component, use the individual lights on the modules to locate the specific failed component. For troubleshooting procedures, refer to "Troubleshooting and Recovery" on page 3-22.

b. **Select Recheck in the Recovery Guru to re-run the Recovery Guru and to ensure that the problem has been corrected.**

c. **If the problem persists, contact technical support.**

12. **Configure the new drives using the storage management software.**

13. **Create, print, and save a new storage array profile.**

End Of Procedure

3.3.3 Replacing All Drives at the Same Time

Use the following procedure to replace all drives at the same time. Figure 3-20 on page 3-38 shows the locations of the power switches. Figure 3-21 on page 3-38 illustrates inserting and removing a drive.

Caution – Potential data loss. Using the wrong drive upgrade procedure can cause data loss. If you are upgrading drives containing RAID 0 volumes, you must use this procedure for replacing all drives at once. If you are upgrading drives containing RAID 1, 3, or 5 volumes, you may use either upgrade procedure.

Caution – Electrostatic discharge damage to sensitive components. To prevent electrostatic discharge damage to the module, use proper antistatic protection when handling the module components.

Caution – Potential damage to drive components. Drives can be damaged by bumping them into other objects or surfaces. When removing or installing a drive into the drive module, place your hand under the drive to support its weight. Place drives on an antistatic, cushioned surface.

1. **Put on antistatic protection.**

2. **Unpack the new drives.**

Set the new drives on a dry, level surface, away from magnetic fields. Save all packing materials in case you need to return the drives.

3. **Read all information provided in "Replace Existing Drives with Greater Capacity Drives" on page 3-32, particularly the paragraphs explaining the differences between the two possible upgrade procedures.**

4. **Compare the SANtricity Storage Manager Product Release Notes with this procedure to determine if you need to modify this procedure, based on more recent information.**

5. **Determine the status of each module and its components.**

   Note the status of the indicator lights on the front and the back of each module. A green light indicates a normal status; an amber light indicates a hardware fault.

   d. **Open the Array Management Window for the storage array.**

   e. **To view the status of all of its components, select the appropriate component button for each module in the Physical View of the Array Management Window.**

   The status for each component will be either Optimal or Needs Attention.

6. **Does the module display green lights only, and is the status Optimal on each module component?**

   - Yes- Go to step 8.
   - No - Go to step 7.

7. **Diagnose and correct the fault.**

   a. **Complete the recovery procedure.**

   If the Recovery Guru directs you to replace a failed component, use the individual lights on the modules to locate the specific failed component. For troubleshooting procedures, refer to "Troubleshooting and Recovery" on page 3-22.

   b. **Select Recheck in the Recovery Guru to re-run the Recovery Guru and ensure that the problem has been corrected.**

   c. **If the problem persists, contact technical support.**

Caution – All data on the drives will be lost when you replace the drives using this method. You must perform a complete back up of the drives and use the backup media to restore the data to the new drives.

8. **Perform a complete backup of the drives you are replacing. You will use the backup media to restore data to the drives later in this procedure.**

9. **Stop all I/O activity to the command module.**

10. **Verify that none of the Active lights above the drives are flashing.**

11. **Turn off both power switches on the back of the command module.**

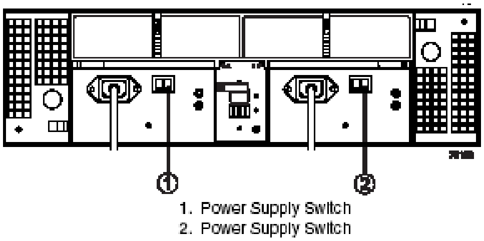

FIGURE 3-20 Power Supply Switches

Note – IMPORTANT If you accidentally remove an active drive, wait at least 30 seconds and then reinstall it. For recovery procedures, refer to your storage management software.

---

12. **Lift the locking lever on the drive and remove it from the slot.**
13. **Slide the new drive all the way into the empty slot and close the drive lever.**
14. **Repeat step 12 and step 13 for each drive you are replacing.**

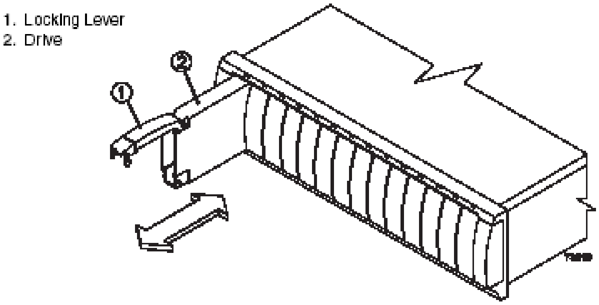

FIGURE 3-21 Removing and Installing a Drive

15. **After installing all of the new drives, turn on both power switches on the command module.**

    As the drives spin up, the Fault lights may flash intermittently. If the Active light begins to flash, data is being restored to the new drive.

16. **Choose one of the following steps, based on the status of the Active and Fault lights:**

- Active lights are on while Fault lights are off - Go to step 18.
- Active lights are off while Fault lights are off - The drive may be installed incorrectly. Remove the drive, wait 30 seconds, and then reinstall it. Go to step 17.
- Fault lights are on - The new drive may be defective. Replace it with another new drive, and then go to step 17.

17. **Did this correct the problem?**

    - Yes - Go to step 18.
    - No - Select the Recovery Guru toolbar button in the Array Management Window and complete the recovery procedure. If the problem persists, contact technical support.

18. **Check the status of the command module using the storage management software.**

    Select the appropriate component button for each module in the Physical View of the Array Management Window to view the status of all its components. The status for each component will be either Optimal or Needs Attention.

19. **Does any module component have a Needs Attention status?**

    - Yes - Go to step 20.
    - No - Go to step 21.

20. **To run the Recovery Guru, select the Recovery Guru toolbar button in the Array Management Window.**

    a. **Complete the recovery procedure.**

    If the Recovery Guru directs you to replace a failed component, use the individual lights on the modules to locate the specific failed component. For troubleshooting procedures, refer to “Troubleshooting and Recovery” on page 3-22.

    b. **Select Recheck in the Recovery Guru to run the Recovery Guru and to ensure that the problem has been corrected.**

    c. **If the problem persists, contact technical support.**

21. **Configure the new drives using the storage management software.**

22. **Create, print, and save a new storage array profile.**

23. **Restore the data to all drives, using the backup media created in step 8.**

End Of Procedure

3.3.4 Replacing One Drive at a Time

Use the following procedure to replace one drive at a time. Figure 3-22 on page 3-42 illustrates inserting and removing a drive.

Caution – Potential data loss. Using the wrong drive upgrade procedure can cause data loss. If you are upgrading drives containing RAID 0 volumes, you must use the procedure for replacing all of the drives at once. If you are upgrading drives containing RAID 1, 3, or 5 volumes, you may use either upgrade procedure.

Caution – Potential data loss. When replacing a drive, make sure the new drive has a storage capacity equal to or greater than the old drive. Using a smaller capacity drive may result in data loss.

Caution – Electrostatic discharge damage to sensitive components. To prevent electrostatic discharge damage to the module, use proper antistatic protection when handling the module components.

Caution – Potential damage to drive components. Drives can be damaged by bumping them into other objects or surfaces. When removing or installing a drive into the drive module, place your hand under the drive to support its weight. Place drives on an antistatic, cushioned surface.

Note – IMPORTANT If you are upgrading drives in a storage array that contains hot spares, drive reconstruction may start on the hot spare before you insert the new drive. The data on the new drive will still be rebuilt, but the process will take longer for each drive. To prevent this delay, you can disable the assigned hot spares while you perform the upgrade procedure. Remember to reassign the hot spares when you are finished upgrading the drives. For information on hot spares, refer to the storage management software.

1. **Put on antistatic protection.**
2. **Unpack the new drives.**

   Set the new drives on a dry, level surface, away from magnetic fields. Save all packing materials in case you need to return the drives.
3. **Read all information provided in "Replace Existing Drives with Greater Capacity Drives" on page 3-32, particularly the paragraphs explaining the differences between the two possible upgrade procedures**
4. **Read the SANtricity Storage Manager Product Release Notes to determine if you need to modify these procedures based on more recent information.**

5. **Determine the status of all modules and their components in the storage array.**

   Note the status of the indicator lights on the front and the back of each module. A green light indicates a normal status; an amber light indicates a hardware fault.

   d. **Select the appropriate component button for each module in the Physical View of the Array Management Window to view the status of all its components.**

   The status for each component will be either Optimal or Needs Attention.

6. **Does each module display green lights only, and is the status Optimal on each module component?**

   - Yes- Go to step 8.
   - No - Go to step 7.

7. **Diagnose and correct the fault.**

   a. **Complete the recovery procedure.**

   If the Recovery Guru directs you to replace a failed component, use the individual lights on the modules to locate the specific failed component. For troubleshooting procedures, refer to "Troubleshooting and Recovery" on page 3-22.

   b. **To ensure that the problem has been corrected, select Recheck in the Recovery Guru to re-run the Recovery Guru.**

   c. **If the problem persists, contact technical support.**

8. **Perform a complete backup of the drives you are replacing.**

9. **Use the storage management software to manually fail the first drive that you are replacing, and then verify the following conditions:**

   - Software shows a Failed status for only one drive (the one that you manually failed)
   - Fault light above the failed drive is illuminated

Caution – Potential data loss. Removing a drive that has not failed can cause data loss. To prevent data loss, remove only drives that have a fault light on or a Failed status in the storage management software.

10. **Lift the locking lever on the failed drive and remove it from the slot.**

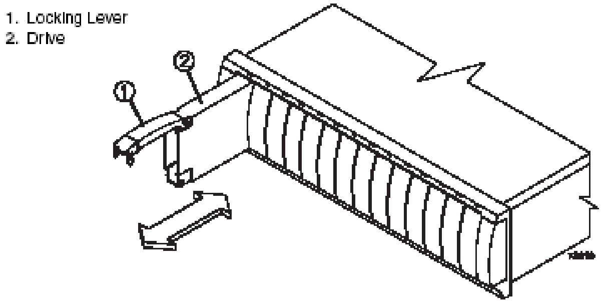

FIGURE 3-22 Removing and Installing a Drive

11. **Slide the new drive all the way into the empty slot and close the locking lever.**

    As the drive spins up, the Fault lights may flash intermittently. The new drive should begin reconstructing automatically after you install it in the drive slot. During reconstruction, the drive's Fault light may come on for a few minutes, and then turn off when the Active light begins flashing. A flashing Active light indicates that data is being restored to the new drive.

Note – IMPORTANT If your storage array has active hot spares, the storage management software may not begin copying data to the new drive until it has been reconstructed on the hot spare. This increases the time required to complete the procedure.

---

12. **Wait for the new drive to spin up, and then choose one of the following steps, based on the status of the Active and Fault lights:**

    - Active lights are on and Fault lights are off - Go to step 14.
    - Active lights are off while Fault lights are off - The drive may be installed incorrectly. Remove the drive, wait 30 seconds, and then reinstall it. Go to step 13.
    - Fault lights are on - The new drive may be defective. Replace it with another new drive, and then go to step 13.

13. **Did this correct the problem?**

    - Yes - Go to step 14.
    - No - Select the Recovery Guru toolbar button in the Array Management Window and complete the recovery procedure. If the problem persists, contact technical support.

14. **Check the status of the command module using the storage management software.**

15. **To view the status of its components, select the appropriate component button for each module in the Physical View of the Array Management Window. The status for each component will be either Optimal or Needs Attention. Do all module components have an Optimal status?**
    - Yes - Go to step 17.
    - No - Go to step 16.

16. **To run the Recovery Guru, select the Recovery Guru toolbar button in the Array Management Window.**

    a. **Complete the recovery procedure.**

    If the Recovery Guru directs you to replace a failed component, use the individual lights on the modules to locate the specific failed component. For troubleshooting procedures, refer to "Troubleshooting and Recovery" on page 3-22.

    b. **When the procedure is completed, select Recheck in the Recovery Guru to re-run the Recovery Guru and to ensure that the problem has been corrected.**

    c. **If the problem persists, contact technical support.**

17. **Configure the new drives using the storage management software.**

18. **Create, save, and print a new storage array profile.**

End Of Procedure

3.3.5 Relocation Considerations

When relocating command module and drives from one storage array to another, use the following guidelines.

Caution – Potential data loss. Moving a storage array or storage array component that is configured as part of a volume group can cause data loss. To prevent data loss, always consult technical support before relocating any configured storage array module.

3.3.5.1 Relocate Drives or Modules

Note – IMPORTANT Always contact technical support before relocating drives or modules.

Drives, drive modules, command modules, or command modules that are part of a volume group configuration should not be moved. If you need to move storage array components, call technical support for detailed procedures. Technical support may direct you to complete several storage array preparation tasks undertaking the relocation. These tasks may include the following:

- Creating, saving, and printing a storage array profile of all storage arrays that will be affected by the relocation of drives, drive modules, command modules, or command modules
- Performing a complete backup of all the data on the drives you intend to move
- Verifying that the volume group and associated volumes on the affected storage array have an Optimal status
- Determining the location and status of any global hot spares associated with the affected storage array

3.3.5.2 Convert an Command Module to a Drive Module

Note – IMPORTANT For information on whether your command module can be converted, contact technical support.

An command module containing one or more controllers can be converted into a drive module containing two environmental services monitors (ESMs). The conversion can be completed without losing the data on the drives.

3.3.6 Raid Storage Manager (RSM)

This section provides an introduction to the Enterprise Management Window and the Array Management Window, and describes the basic software layout. This section also provides an overview of the tools and functions that are performed from each management window, outlines the menus and toolbar buttons available in each management window, and describes the various views and tabs displayed in the Array Management Window.

For additional conceptual information and detailed procedures for the options described in this section, refer to Storage Management Concept, "Introducing the Enterprise Management Window" in the Enterprise Management Window online help, and Learn About Storage Management Software, "Using the Array Management Window" in the Array Management Window online help.

3.3.6.1 Client Software Windows

The client software has two main windows: the Enterprise Management Window (Figure 3-23) and the Array Management Window (Figure 3-24).

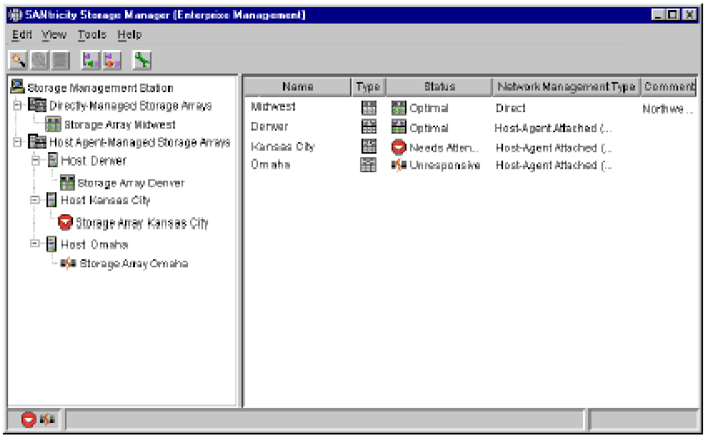
SANtricity Storage Manager (Enterprise Management)
Edit View Tools Help
Storage Management Station
Directly-Managed Storage Arrays
Storage Array Midwest
Host Agent-Managed Storage Arrays
Host Denver
Storage Array Denver
Host Kansas City
Storage Array Kansas City
Host Omaha
Storage Array Omaha

Name	Type	Status	Network Management Type	Comment
Midwest		Optimal	Direct	Northwe ...
Denver		Optimal	Host-Agent Attached (...	
Kansas City		Needs Atten...	Host-Agent Attached (...	
Omaha		Unresponsive	Host-Agent Attached (...	

FIGURE 3-23 Enterprise Management Window

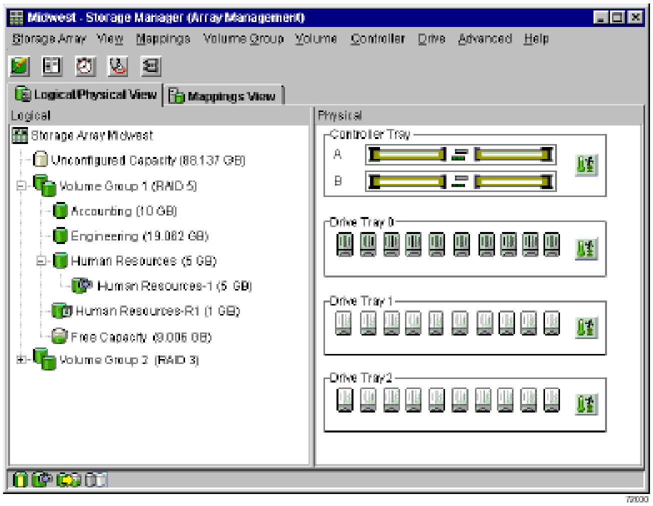

FIGURE 3-24 Array Management Window

3.3.6.2 The Enterprise Management Window

The Enterprise Management Window is the first window to appear when you start the software. It is used to:

- Detect and add the storage arrays you want to manage.
- View the status of all the storage arrays detected or added.
- Execute scripts to perform batch management tasks on a particular storage array using the Script Editor. For example, scripts may be run to create new volumes or download new controller firmware.
- Configure destinations, such as e-mail or Simple Network Management Protocol [SNMP] traps, to receive alert notifications for non-optimal storage arrays.

A local configuration file stores all the information about storage arrays you have added and any e-mail or SNMP destinations you have configured.

The Enterprise Management Window or the event monitor must be running to receive alert notifications of critical events on storage arrays. For more information about the event monitor, refer to "Event Monitor" on page 3-86.

After storage arrays are added, the Enterprise Management Window is used primarily to monitor the storage arrays for a Needs Attention status and for alert notification of critical errors affecting the storage arrays. When you are notified of a non-optimal storage array status in the Enterprise Management Window, starting an Array Management Window for the affected storage array displays detailed information about the storage array condition.

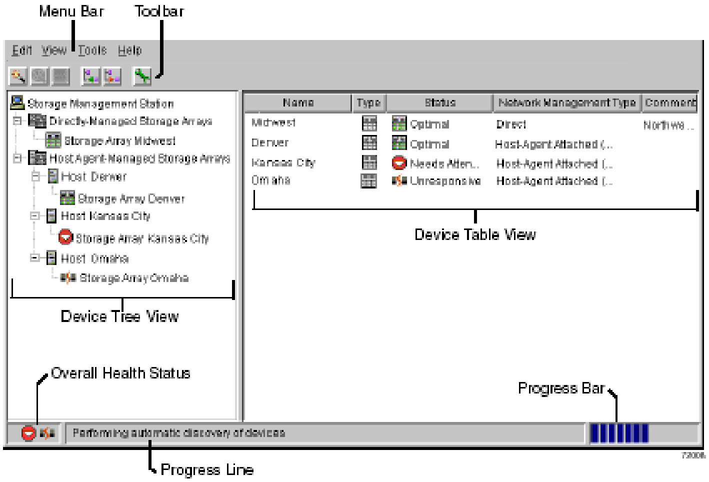

Name	Type	Status	Network Management Type	Comment
Midwest	▦	▦ Optimal	Direct	Northwe...
Denver	▦	▦ Optimal	Host-Agent Attached (...	
Kansas City	▦	● Needs Atten...	Host-Agent Attached (...	
Omaha	▦	Unresponsive	Host-Agent Attached (...	

Performing automatic discovery of devices

FIGURE 3-25 Enterprise Management Window

Device Tree

The Enterprise Management Window Device Tree provides a hierarchical view of all the host-agent and directly managed storage arrays (Figure 3-26). The storage management station node is the root node and sends the storage management commands. When storage arrays are added to the Enterprise Management Window, they are shown in the Device Tree as child nodes of the storage management station node. A storage array can be managed through an Ethernet connection on each controller in the storage array (directly managed) or through a host interface connection to a host with the host-agent installed (host-agent managed).

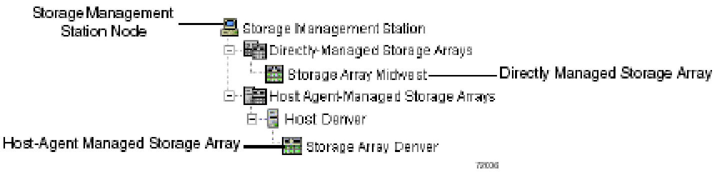
Storage Management Station Node
Storage Management Station
Directly Managed Storage Arrays
Storage Array Midwest
Directly Managed Storage Array
Host Agent-Managed Storage Arrays
Host Denver
Host-Agent Managed Storage Array
Storage Array Denver
72036

FIGURE 3-26 Device Tree Example

There are two ways to add storage arrays to the Enterprise Management Window:

- Automatic Discovery - detects directly managed and host-agent managed storage arrays on the local subnetwork and adds them to the Enterprise Management Window automatically. The Enterprise Management Window detects host-agent managed storage arrays by locating any hosts that provide network management connections to the storage arrays. The hosts appear in the Device Tree with their associated storage arrays.
- Add Device - provides the ability to manually add directly managed and host-agent managed storage arrays to the Enterprise Management Window. For a directly managed storage array, a host name or IP address must be entered for each controller in the storage array. Typically, there are two controllers in a single storage array. For a host-agent managed storage array, a host name or an IP address must be entered for the host that is attached to the storage array.

The first time storage arrays are detected or added to the Enterprise Management Window, they are shown as <unnamed> in the Device Tree and Device Table unless they have been named by another storage management station. Storage Management Station Node Host-Agent Managed Storage Array Directly Managed Storage Array

Device Table

The Device Table lists the name, type of managed device, status, management type (direct network attached for directly managed storage arrays, or host-agent attached for host-agent managed storage arrays), and comments entered for storage arrays (Figure 3-25).

Enterprise Management Window Menus

The Enterprise Management Window menus on the menu bar are described in Table 3-4.

TABLE 3-4 Enterprise Management Window Menus

Menu	Description
Edit	Contains options to add devices or remove devices from the Enterprise Management Window, to configure alert destinations, or to add a comment to the Device Table about a storage array.
View	Provides options to sort the entries in the Device Table by name, status, management type, or comment. Another option shows partially managed devices-a condition occurring when only one controller of a controller pair is defined or can be reached when the storage array is added or detected. In this state, many management operations that require access to both controllers are not available.
Tools	Displays options for automatically detecting devices on the same subnetwork or for rescanning to find storage arrays newly attached to a host. Other items on this menu include options to update the event monitor, open an Array Management Window to manage a selected storage array, open the Script Editor to perform batch management tasks, or load a saved configuration file.
Help	Presents options to display the Enterprise Management Window Help system and to view the software version and copyright information.

Enterprise Management Window Toolbar

The Enterprise Management Window toolbar buttons are described in Table 3-5.

TABLE 3-5 Enterprise Management Window Toolbar Buttons

Toolbar Button	Description
Automatically detect new devices	Activates the Automatic Discovery option that detects hosts and storage arrays on the local subnetwork and adds them to the Enterprise Management Window.
Rescan selected host for new devices	Rescans the highlighted host for any newly attached storage arrays. Before using this option, the new storage arrays must be physically attached to the host and the host-agent software residing on the host must be restarted. Note This option is available only when you select a host in the Device Tree.
Synchronize Event Monitor	Synchronizes the event monitor with any changes made in the Enterprise Management Window, such as adding or removing devices or adding alerts. Note This option is available only if the configurations of the Enterprise Management Window and the event monitor are not synchronized.
Add host/device	Opens the Add Device dialog for manually adding hosts or storage array controllers to the Enterprise Management Window.
Remove host/device	Removes the selected storage array or the selected host and its attached storage arrays from the Enterprise Management Window. Note This option is available only when a storage array in the Device Tree or Device Table is selected, or a host in the Device Tree is selected.
Launch an Array Management Window	Starts an Array Management Window for the selected storage array. Note This option is available only when a storage array in the Device Tree or Device Table is selected. Starts an Array Management Window for the selected.

3.3.6.3 The Array Management Window

The Array Management Window (Figure 3-27) is used to configure and maintain the logical and physical components of a storage array and to view and define volume-to-LUN mappings.

The Array Management Window is specific to an individual storage array; therefore, you can manage only a single storage array within an Array Management Window. However, you can start other Array Management Windows from the Enterprise Management Window to simultaneously manage multiple storage arrays.

The storage management software supports firmware version 5.40 and all firmware versions 4.x and 5.x. For maximum system stability, the recommended minimum is firmware version 4.01.02.30. However, to access all the features of version 8.40, you must upgrade to firmware version 5.40.

Note – IMPORTANT Depending on your version of storage management software, the views, menu options, and functionality may be different to the information presented in this guide. Refer to the documentation supplied with your version of storage management software for information on available functionality.

The features of a particular release of firmware will be accessible when an Array Management Window is launched from the Enterprise Management Window to manage a storage array. For example, you manage two storage arrays using this software; one storage array has firmware version 5.40 and the other has firmware version 4.x. When you open an Array Management Window for a particular storage array, the correct Array Management Window version is used. The storage array with firmware version 5.40 will use version 8.40 of the storage management software, and the storage array with firmware version 4.x will use version 7.x. You can verify the version you are currently using by selecting Help >> About in the Array Management Window.

This bundling of previous versions of the Array Management Window provides the flexibility of upgrading the firmware only on selected storage arrays instead of having to perform an upgrade on all storage arrays at once.

Array Management Window Tabs

The Array Management Window has two tabs: Logical/Physical View and Mappings View (Figure 3-27), which are described in Table 3-6. The tabs display information about the logical components (volume and volume groups), physical components (controllers and drives), topological elements (host groups, hosts, host ports), and volume-to-LUN mappings in a storage array.

Also displayed in the Array Management Window are the toolbar, menu bar, components buttons, premium feature status area, and the storage partition status area.

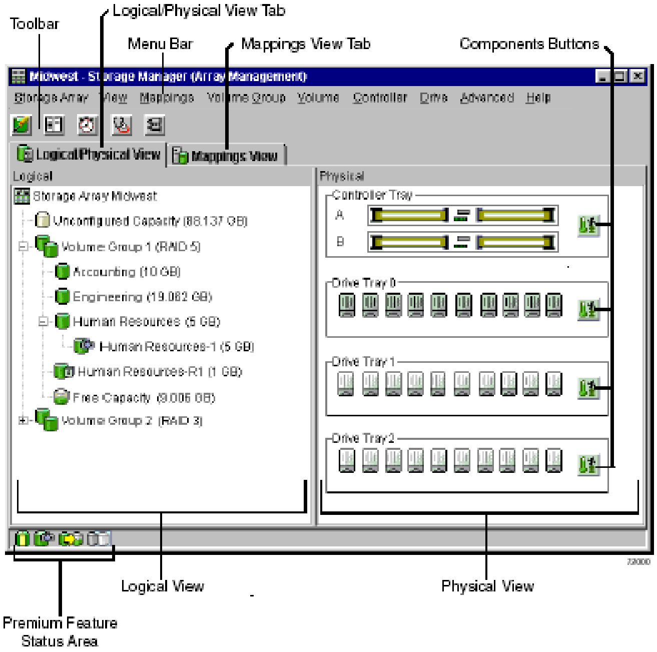

FIGURE 3-27 Array Management Window

TABLE 3-6 Array Management Window Tabs

Tabs	Description
Logical/Physical View	The Array Management Window Logical/Physical View contains two panes: the Logical View and the Physical View. The Logical View (left pane of Figure 3-27 on page 3-51) provides a tree-structured view of logical nodes. This view shows the organization of storage array capacity into volume groups and volumes. The Physical View (right pane of Figure 3-27 on page 3-51) provides a view of the physical devices in a storage array, such as command module and drive module components. Selecting a volume or other entity in the Logical View shows you the associated physical components in the Physical View. There is a Components button in every command module and drive module that, when selected, presents the status of each component and shows the temperature status.
Mappings View	The Mappings View of the Array Management Window contains two panes: the Topology View and the Defined Mappings View. For details, see "Mappings View" on page 3-68. The Topology View provides a tree-structured view of logical nodes related to storage partitions. The Defined Mappings Detail View displays the mappings associated with the selected node in the Topology View.

Array Management Window Menus

The Array Management Window menus are described in Table 3-7. The menus are used to perform storage management operations for a selected storage array or for selected components within a storage array. See Figure 3-27 on page 3-51 for an example of the Array Management Window menus.

TABLE 3-7 Array Management Window Menus (1 of 2)

Menu	Description
Storage Array	Contains options to perform the following storage array management operations: locating functions (locating the storage array or a specific drive channel by flashing indicator lights), configuring the storage array, enabling premium features, starting Recovery Guru, monitoring performance, downloading firmware and NVSRAM files, changing various settings, setting controller clocks, redistributing volumes, running Read Link Status diagnostics, activating or deactivating the remote volume mirroring premium feature, and renaming storage arrays.
View	Allows you to change the display from the Logical/Physical view to the Mappings View, view Associated Components for a selected node, find a particular node in the Logical View or Topology View, locate an associated node in the tree, or access the Event Log or the Storage Array Profile.
Mappings	Permits you to make changes to or retrieve details about mappings associated with a selected node. The options are: Define, Change, Move, Replace Host Port, Show All Host Port Information, Remove, and Rename. Note You must be in the Mappings View to access the options available in this menu.
Volume Group	Presents options to perform the following storage management operations on volume groups: creating or locating volumes, changing Redundant Array of Independent Disks (RAID) level or controller ownership, adding free capacity (drives), defragmenting free capacity, placing controllers online or offline, initializing a volume group, reviving a volume group, checking redundancy, and deleting a volume group. Note These menu options are only available when a volume group is selected.
Volume	Provides options to perform the following storage management operations on volumes: creating volumes, changing ownership or segment size, increasing capacity, initializing, deleting, or renaming a volume, creating a volume copy, viewing volume copies using the Copy Manager, creating or disabling a snapshot volume, creating a remote volume mirror, and viewing volume properties. Note These menu options are only available when a volume is selected.
Controller	Displays options to perform the following storage management operations on controllers: placing a controller online or offline, enabling or disabling data transfer (I/O), changing the controller mode to active, changing the preferred loop ID, modify a controller's IP address, gateway address, or network subnet mask, running diagnostics, and viewing controller properties. Note These menu options are only available when a controller is selected.

TABLE 3-7 Array Management Window Menus (1 of 2)

Menu	Description
Drive	Contains options to perform the following storage management operations on drives: locating a drive, assigning or unassigning a hot spare, failing, reconstructing, reviving or initializing a drive, or viewing drive properties. Note These menu options are only available when a drive is selected.
Advanced	Presents maintenance options which should only be used under the guidance of technical support.
Help	Provides options to display the contents of the Array Management Window Help, view a reference of all Recovery Guru procedures, and to view the software version and copyright information.

Array Management Window Toolbar

The Array Management Window toolbar buttons are used to create new volumes or volume groups, monitor performance, view events, recover from failures, and locate a particular node. Each toolbar button is described in Table 3-8. See Figure 3-27 on page 3-51 for an example of the Array Management Window toolbar.

TABLE 3-8 Array Management Window Toolbar Buttons

Toolbar Button	Description
Create new volumes (and volume groups, if applicable)	Permits you to create volume groups and volumes. Note You must select either a Free Capacity Node or an Unconfigured Capacity Node in the Logical View before this option is available.
View diagnostic event log	Starts the Event Log Viewer which displays a detailed list of events that occur in a storage array.

TABLE 3-8 Array Management Window Toolbar Buttons

Toolbar Button	Description
Monitor performance	Opens the Performance Monitor which provides information about how the storage array is functioning.
Recover from failures	Initiates the Recovery Guru which is used to help troubleshoot storage array problems. Note If the storage array is in a Needs Attention state, the icon on the Recovery Guru toolbar button flashes.
Find node in tree	Starts the Find dialog that allows you to search for a particular node in the Logical/Physical View or Mappings View of the Array Management Window.

3.3.6.4 Protecting Your Data

This section describes storage array configuration options you can use to maximize data availability and to prevent data loss. A storage array includes redundant drives, controllers, power supplies, and fans. This hardware redundancy keeps the storage array working if a component fails. In addition, you can use the storage management software to implement the data protection options in this section. For conceptual information and detailed procedures for the options described in this section, refer to Learn About Data Protection Strategies in the Array Management Window online help.

3.3.6.5 Software Redundancy

The storage management software has three premium features that provide data protection strategies. Remote Volume Mirroring is used to create an online, real-time replication of data between storage arrays over a remote distance, while snapshot volume creation provides a way to more efficiently create a point-in-time image of data either for backup or for application testing. Volume Copy allows you copy data from one volume to another within the same storage array. The following sections provide a brief overview of the premium features used for data protection strategies. For more detailed information, refer to “Configuring Storage Arrays” on page 3-64.

Volume Copy

The Volume Copy premium feature is used to copy data from one volume (the source) to another volume (the target) in a single storage array. The source volume is a standard volume in a volume copy that accepts host I/O requests and stores application. The target volume is a standard volume in a volume copy that maintains a copy of the data from the source volume. A volume copy can be used to back up data, to copy data from volume groups that use smaller capacity drives to volume groups that use larger capacity drives, or to restore snapshot volume data to the base volume.

Remote Volume Mirroring

When you create a remote volume mirror, a mirrored volume pair is created, which consists of a primary volume on a primary storage array and a secondary volume on a secondary storage array. When the mirror relationship is first created, data from the primary volume is copied in its entirety to the secondary volume. The secondary volume maintains a mirror (or copy) of the data from its associated primary volume. In the event of a disaster or catastrophic failure at the primary site, a manual role reversal by the system administrator can be performed to promote the secondary volume to a primary role. Hosts will then be able to access the newly promoted volume and business operations can continue.

Snapshot Volumes

A snapshot volume is a point-in-time image of a volume. It is the logical equivalent of a complete physical copy, but you create it much more quickly and it requires less disk space.

Typically, a snapshot volume is created so that an application, such as a backup application, can access the snapshot volume and read the data while the base volume remains online and user-accessible. When the backup completes, the snapshot volume is no longer needed. You can also create several snapshot volumes of a base volume and write data to the snapshot volumes to perform testing and analysis. For example, before upgrading a database management system, snapshot volumes can be used to test different configurations. The performance data provided by the storage management software can also be used to help decide how to configure the live database system.

3.3.6.6 RAID Levels and Data Redundancy

RAID is an acronym for Redundant Array of Independent Disks. It is a storage solution in which the same data or information about the data (parity) is stored in different places on multiple hard disks. By placing data on multiple disks, I/O operations overlap and performance improves. If a disk drive in a volume group fails, the redundant or parity data can be used to regenerate the user data on replacement disk drives.

RAID relies on a series of configurations, called levels, to determine how user and redundancy data is written and retrieved from the drives. Each level provides different performance and protection features. The storage management software offers four formal RAID level configurations: RAID levels 0, 1, 3, and 5. Table 3-9 describes these configurations.

RAID levels 1, 3, and 5 write redundancy data to the drive media for fault tolerance. The redundancy data might be a copy of the data or an error-correcting code derived from the data. If a drive fails, the redundancy data can be used to quickly reconstruct information.

Only one RAID level can be configured across each volume group. Each volume group stores its own redundancy data. The capacity of the volume group is the aggregate capacity of the member drives, minus the capacity reserved for redundancy data. The amount of capacity needed for redundancy data depends on the RAID level used.

TABLE 3-9 RAID Level Configurations

RAID Level	Short Description	Detailed Description
RAID 0	Non-Redundant, Striping Mode	• Used for high performance needs, but does not provide data redundancy. • Stripes data across all drives in the volume group. • Not recommended for high data availability needs. RAID 0 is better for non-critical data. • A single drive failure causes all associated volumes to fail and data loss can occur.

TABLE 3-9 RAID Level Configurations

RAID Level	Short Description	Detailed Description
RAID 1	Striping/ Mirroring Mode	• Also called RAID 10 or 0+1. • A minimum of two drives is required for RAID 1;one for the user data and one for the mirrored data. • Offers the best availability high performance and the best data availability. Data is written to two duplicate disks simultaneously. If one of the disk drives in a disk-pair fails, the system can instantly switch to the other disk without any loss of data or service. However, only half of the drives in the volume group are available for user data. • Uses disk mirroring to make an exact copy from one drive to another drive. • A single drive failure causes associated volumes become degraded, but the mirror drive allows access to the data. • Can survive multiple drive failures as long as no more than one failure exists per mirrored pair. • A drive-pair failure in a volume group causes all associated volumes to fail and data loss could occur.
RAID 3	High Bandwidth Mode	• Both user data and redundancy data (parity) are striped across the drives. • The equivalent of one drive's capacity is used for redundancy data. • Good for large data transfers in applications such as multimedia or medical imaging that write and read large sequential chunks of data. • A single drive failure in a volume group causes associated volumes to become degraded, but the redundancy data allows access to the data. • Two or more drive failures in a volume group cause all associated volumes to fail and data loss could occur.
RAID 5	High I/O Mode	• Both user data and redundancy data (parity) are striped across the drives. • The equivalent of one drive's capacity is used for redundancy data. • Good for multi-user environments such as database or file system storage, where typical input/output (I/O) size is small and there is a high proportion of read activity. • A single drive failure in a volume group causes associated volumes to become degraded, but the redundancy data allows access to the data. • Two or more drive failures in a volume group causes all associated volumes to fail and data loss could occur.

3.3.6.7 Hardware Redundancy

Data protection strategies provided by the storage system hardware include cache memory, hot spare drives, background media scans, and channel protection.

Controller Cache Memory

Caution – Sometimes write caching is disabled when batteries are low or discharged. If a parameter called Write caching without batteries is enabled on a volume, write caching continues even when batteries in the command module or array module are discharged. If you do not have an uninterruptible power supply (UPS) for power loss protection, do not enable this parameter, because data in the cache will be lost during a power outage if the command module or array module does not have working batteries.

Write caching can increase I/O performance during data transfers. However, it also increases the risk of data loss if a controller (or its memory) fails while unwritten data resides in cache memory. Write cache mirroring protects data during a controller or cache memory failure. When write cache mirroring is enabled, cached data is mirrored across two redundant controllers with the same cache size. The data written to the cache memory of one controller is also written to the cache memory of the alternate controller. Therefore, if one controller fails, the alternate can complete all outstanding write operations.

To prevent data loss or corruption, the controller periodically writes cache data to disk (flushes the cache) when the amount of unwritten data in cache reaches a certain level, called a start percentage, or when data has been in cache for a predetermined amount of time. The controller writes data to disk until the amount of data in cache drops to a stop percentage level. Start and stop percentages can be configured by the user. For example, you can specify that the controller start flushing the cache when it reaches 80% full and stop flushing the cache when it reaches 16% full.

Low start and stop percentages provide for maximum data protection. However, in both cases, this increases the chance that data requested by a read command will not be in the cache, decreasing the cache hit percentage for writes and the I/O request. Choosing low start and stop percentages also increases the number of disk writes necessary to maintain the cache level, increasing system overhead and further decreasing performance.

Data in the controller cache memory is also protected in case of power outages. Command modules and array modules contain batteries that protect the data in cache beekeeping a level of power until the data can be written to the drive media. If a power outage occurs and there is no battery or the battery is damaged, data in the cache that has not been written to the drive media will be lost, even if it is mirrored to the cache memory of both controllers. It is, therefore, important to change the command module and array module batteries at the recommended time intervals. The controllers in the storage array keep track of the age (in days) of the battery. After replacing the battery, the age must be reset so that you will receive an accurate critical alert notification when the battery is nearing expiration and when it has expired.

Hot Spare Drives

A valuable strategy to protect data is to assign available drives in the storage array as hot spare drives. A hot spare is a drive, containing no data, which acts as a standby in the storage array in case a drive fails in a RAID 1, 3, or 5 volume. The hot spare adds another level of redundancy to the storage array. Generally, the drive assigned as a hot spare should have a capacity that is equal to or greater than the capacity of the largest drive on the storage array. If a drive fails in the storage array, the hot spare is automatically substituted for the failed drive without requiring user intervention. If a hot spare is available when a drive fails, the controller uses redundancy data to reconstruct the data onto the hot spare. When you have physically replaced the failed drive, the data from the hot spare is copied back to the replacement drive. This is called copyback.

If you do not have a hot spare, you can still replace a failed disk drive while the storage array is operating. If the drive is part of a RAID 1, 3, or 5 volume group, the controller will use redundancy data to automatically reconstruct the data onto the replacement drive. This is called reconstruction.

Background Media Scan

A media scan is a background process performed by the controllers to provide error detection on the drive media. A media scan detects errors and reports them to the Event Log.

The media scan must be enabled for the entire storage array as well as enabled on each volume.

The media scan runs on all volumes in the storage array for which it has been enabled. The advantage of enabling a media scan is that the process can find media errors before they disrupt normal drive reads and writes. The media scan process scans all volume data to verify that it can be accessed, and if you enable a redundancy check, it also scans the volume redundancy data.

Channel Protection

In a Fibre Channel environment, channel protection is usually present for any volume group candidate because, when the storage array is properly cabled, there are two redundant Fibre Channel Arbitrated Loops for each drive.

3.3.6.8 I/O Data Path Protection

I/O data path protection to redundant controllers in a storage array is accomplished with the Auto-Volume Transfer (AVT) feature and a host multi-path driver.

Note – IMPORTANT Redundant Disk Array Controller (RDAC) must be uninstalled in order for DMP to become the default failover driver.

A multi-path driver is an I/O path failover driver installed on host computers that access the storage array, such as Redundant Disk Array Controller (RDAC). Veritas Volume Manager with Dynamic Multi-Pathing (DMP) is another example of a failover driver. This failover driver requires the Array Support Library (ASL) software to be installed, which provides information to the Volume Manager for setting the path associations for the failover driver.

AVT is a built-in feature of the controller firmware that allows volume-level failover rather than controller-level failover. AVT is disabled by default and will be automatically enabled based on the failover options supported by the host operating system.

AVT or RDAC will transfer volumes to the alternate controller if the preferred controller owner fails. If the volumes are not subsequently transferred back to their preferred controller, a critical event will automatically be generated. An associated alert notification will automatically be sent if you have configured alert destinations for the storage array.

For operating system-specific failover options, refer to the SANtricity Storage Manager Installation Guide.

Multi-Path Driver with AVT Enabled

If AVT is enabled when a volume is created, a controller must be assigned to own the volume (called the preferred controller, or preferred owner). The preferred controller normally receives the I/O requests to the volume. If a problem along the data path (such as a component failure) causes an I/O to fail, the multi-path driver will issue the I/O to the alternate controller.

When AVT is enabled and used in conjunction with a host multi-path driver, it helps ensure an I/O data path is available for the storage array volumes. The AVT feature changes the ownership of the volume receiving the I/O to the alternate controller.

After the I/O data path problem is corrected, the preferred controller will automatically reestablish ownership of the volume as soon as the multi-path driver detects the path is normal again.

Multi-Path Driver with AVT Disabled

When AVT is disabled, the I/O data path will still be protected as long as a multi-path driver is installed on each host connected to the storage array. However, when an I/O request is sent to a specific volume, and a problem occurs along the data path to its preferred controller, all volumes on that controller will be transferred to an alternate controller instead of just the specific volume.

3.3.6.9 Password Protection

Note – IMPORTANT Executing destructive commands on a storage array can cause serious damage, including data loss. Without password protection, all options are available within this storage management software.

Note – IMPORTANT If you forget the password, contact technical support.

The storage management software provides a number of security features to protect data, including generation numbering to prevent replay attacks. Hashing and encryption are employed to guard against client spoofing and snooping.

For added security, you can configure a password for each storage array you manage. Because the password is stored on the storage array, each storage array that you want to be password protected will need a password. A specified password protects any options that the controller firmware deems destructive. These options include any functions that change the state of the storage array such as creation of volumes, modification of cache settings, and so on.

After the password has been set on the storage array, you will be prompted for that password the first time you attempt a destructive operation in the Array Management Window. You will be asked for the password only once during a single management session.

Password Failure Reporting and Lockout

For storage arrays with a password and alert notifications configured, any attempts to access the storage array without the correct password will be reported.

If a password is incorrectly entered, an information major event log (MEL) event is logged, indicating than an invalid password or no password has been entered.

If the password is incorrectly entered 10 times within 10 minutes, both controllers will enter lockout mode. The lockout mode will last for a period of 10 minutes, during which both controllers will deny any attempts to enter a password to access the storage array.

Note – IMPORTANT If the controllers are reset, the password failure counter will be cleared and access to the storage array can be attempted again. If the password is incorrectly entered after 10 attempts within 10 minutes, the controllers will re-enter lockout mode.

A critical MEL event will be logged to the event log, indicating that the controllers have entered lockout mode. After the 10 minute lockout period has elapsed, the controllers will reset the password failure counter and will unlock themselves.

3.3.6.10 Configuring Storage Arrays

This section provides descriptions for volumes and volume groups, Dynamic Volume Expansion (DVE), and premium features such as SANshare Storage Partitioning, snapshot volumes, Remote Volume Mirroring, and Volume Copy. In addition, this section describes the specific functions of the Mappings View in the Array Management Window, an overview of the heterogeneous host setting, and how to manage persistent reservations. For additional conceptual information and detailed procedures for the options described in this section, refer to Learn About Configuring a Storage Array in the Array Management Window online help.

3.3.6.11 Volumes and Volume Groups

When configuring a storage array, appropriate data protection strategies as well as how the total storage capacity will be organized into volumes and shared among hosts must be considered. The storage management software identifies several distinct volumes: standard, snapshot, snapshot repository, primary, secondary, mirror repository, source, and target.

- Standard volume - A logical structure created on a storage array for data storage. A standard volume is created using the Create Volume Wizard. If the premium feature is not enabled for snapshot volumes or Remote Volume Mirroring, then only standard volumes will be created. Standard volumes are also used in conjunction with creating snapshot volumes and remote mirror volumes.
- Snapshot volume - A point-in-time image of a standard volume. A snapshot volume is the logical equivalent of a complete physical copy, but you create it much more quickly and it requires less disk space. The volume from which you are basing the snapshot volume, called the base volume, must be a standard volume in your storage array.
- Snapshot repository volume - A special volume in the storage array created as a resource for a snapshot volume. A snapshot repository volume contains snapshot volume metadata and copy-on-write data for a particular snapshot volume.o Primary volume - A standard volume in a mirror relationship that accepts host I/O and stores application data. When the mirror relationship is first created, data from the primary volume is copied in its entirety to the associated secondary volume.
- Secondary volume - A standard volume in a mirror relationship that maintains a mirror (or copy) of the data from its associated primary volume. The secondary volume remains unavailable to host applications while mirroring is underway. In the event of a disaster or catastrophic failure of the primary site, the system administrator can promote the secondary volume to a primary role.
- Mirror repository volume - A special volume created as a resource for each controller in both the local and remote storage array. The controller stores mirroring information on the mirror repository volume, including information about remote writes that are not yet complete. The controller can use the mirrored information to recover from controller resets and accidental powering-down of storage arrays.
- Source volume - A standard volume that contains the data that will be copied to another volume, which is known as the target volume. A source volume can be either a standard volume, a snapshot volume, the base volume of a snapshot volume, or a primary volume of a mirrored pair.
- Target volume - A standard volume to which the data on the source volume is being copied. When a volume is selected as a target volume, any existing data on the volume will be completely overwritten and the volume will automatically become read-only after the copy operation has completed, to protect it from host write access. After the volume copy completes, you can use the Copy Manager to disable the Read-Only attribute for the target volume.

Volume Groups

A volume group is a set of drives that the controller logically groups together to provide one or more volumes to an application host. When creating a volume from unconfigured capacity, the volume group and the volume are created at the same time. When creating a volume from free capacity, an additional volume is created on an existing volume group (Figure 3-28).

To create a volume group, two parameters must be specified: RAID level and capacity (how large you want the volume group). For the capacity parameter, you can either choose the automatic choices provided by the software or select the manual method to indicate the specific drives to include in the volume group. The automatic method should be used whenever possible, because the software provides the best selections for drive groupings.

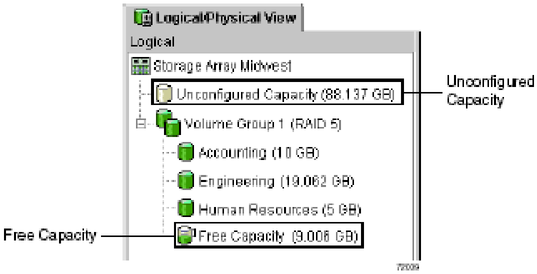

Unconfigured Capacity

Free Capacity

FIGURE 3-28 Unconfigured and Free Capacity Nodes

3.3.6.12 Standard Volumes

Note – IMPORTANT The host operating system may have specific limits on how many volumes the host can access which must be considered when creating volumes for use by a particular host. For operating system restrictions, refer to the SANtricity Storage Manager Product Release Notes that were shipped with the software and your host operating system documentation.

A standard volume is a logical structure created on a storage array for data storage. A volume is defined over a set of drives called a volume group, and has a defined RAID level and capacity. Volumes are created from either unconfigured capacity or free capacity nodes on the storage array. If no volumes on the storage array are configured, the only node available is the unconfigured capacity node (Figure 3-28 on page 3-66).

The Create Volume Wizard is used to create one or more volumes on the storage array. During the volume creation process, the Wizard prompts you to select the capacity to allocate for the volumes and to define basic and optional advanced volume parameters for the volume. Each Wizard screen has context-sensitive help.

Specifying Volume Parameters from Free Capacity

Note – IMPORTANT The free capacity, unconfigured capacity, or unassigned drives selected when starting the Wizard determine the default initial capacity selections. After the Wizard begins, the capacity can be changed by selecting a different free capacity node location for the volume, or by selecting different unassigned drives for the volume group.

The Specify Volume Parameters screen is used to specify the capacity for the volume, the volume name, and whether to use recommended advanced volume settings or customize the advanced volume properties for this volume.

The Specify Advanced Volume Parameters screen is also used to specify the volume I/O characteristics that will apply to the volume based on the needs of your application, or a custom cache read-ahead multiplier and segment size, preferred controller ownership, and a volume-to-LUN mapping parameter. Volumes are tailored to specific application needs by customizing the advanced volume settings.

After the volume creation process is finished, a confirmation dialog is displayed. Use this dialog to restart the Wizard to create another volume.

Specifying Volume Parameters from Unconfigured Capacity

Note – IMPORTANT The free capacity, unconfigured capacity, or unassigned drives selected when starting the Wizard determine the default initial capacity selections. After the Wizard begins, the capacity can be changed by selecting a different free capacity node location for the volume, or by selecting different unassigned drives for the volume group.

The Specify Volume Group Parameters screen is used to specify the RAID level of the volume group to meet your volume data storage and protection requirements, and to select the drives that will comprise the volume group. It also provides a way to specify the capacity for the volume, the volume name, and whether to use recommended advanced volume settings or customize the advanced volume properties for the volume.

The Specify Advanced Volume Parameters screen can further be used to specify the volume I/O characteristics that will apply to the volume based on the needs of your application, including a custom cache read-ahead multiplier and segment size, preferred controller ownership, and a volume-to-LUN mapping parameter.

During the volume creation process, you will be prompted to set the volume-to-LUN mapping preference to specify whether you will be using SANshare Storage Partitioning. There are two settings:

- Automatic - If you are not using SANshare Storage Partitioning, specify this setting. The Automatic setting specifies that a logical unit number (LUN) be automatically assigned to the volume using the next available LUN within the default group. This setting grants volume access to host groups or hosts that have no specific volume-to-LUN mappings (designated by the Default Group node in the Topology View).
- Map later with SANshare Storage Partitioning - If you are using SANshare Storage Partitioning, specify this setting. The Map later setting specifies that a LUN not be assigned to the volume during volume creation. This setting allows definition of a specific volume-to-LUN mapping and creation of storage partitions. After the volume creation process is finished, a confirmation dialog is displayed. Use this dialog to restart the Wizard to create another volume.

3.3.6.13 Mappings View

The Mappings View is used to define the storage topology elements (host groups, hosts, host ports, and so on), to define volume-to-LUN mappings, and to view SANshare Storage

Partitioning and heterogeneous host information. The Mappings View has two views, Topology View and Defined Mappings View, shown in Figure 3-29 on page 3-68 and described in Table 3-10.

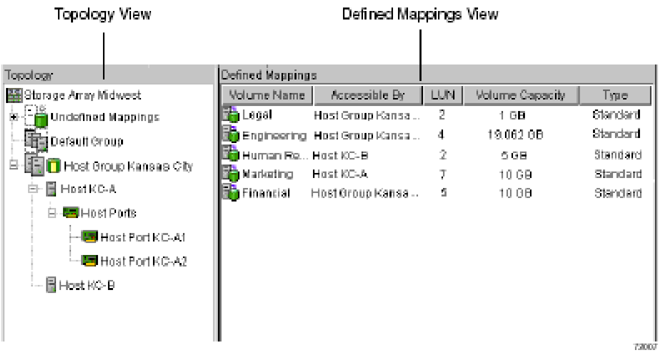

Volume Name	Accessible By	LUN	Volume Capacity	Type
Legal	Host Group Kansa ...	2	1 GB	Standard
Engineering	Host Group Kansa ...	4	19.062 GB	Standard
Human Re...	Host KC-B	2	5 GB	Standard
Marketing	Host KC-A	7	10 GB	Standard
Financial	Host Group Kansa ...	5	10 GB	Standard

FIGURE 3-29 Mappings View Window

TABLE 3-10 Mappings View Tab

View	Description
Topology	Shows defined topological elements (host groups, hosts, and host ports), undefined mappings (volumes that have been created but do not have a defined volume-to-LUN mapping), and the Default Group.
Defined Mappings	Displays the volume-to-LUN mappings in a storage array in table form. Information is displayed about the volumes: topological entities that can access the volume, volume name, volume capacity, and LUN number associated with the volume.

Table 3-11 describes the topological elements displayed in the Mappings View Window in Figure 3-29 on page 3-68.

TABLE 3-11 Volume-to-LUN Terminology

Term	Description
SANshare Storage Partitioning Topology	A collection of nodes (default group, host groups, hosts, and host ports) shown in the Topology View of the Mappings View tab.You must define the various topological elements if you want to define specific volume-to-LUN mappings and storage partitions for host groups or hosts.
Default Group	A node in the Topology View that designates all host groups, hosts, and host ports that:(1) have no specific volume-to-LUN mappings and (2) share access to any volumes that were automatically assigned default LUN mappings by the controller firmware during volume creation.
Host Group	An optional topological element that you define if you want to designate a collection of hosts that will share access to the same volumes. The host group is a logical entity.
Host	A computer that is attached to the storage array and accesses various volumes on the storage array through its host ports (host bus adapters). You can define specific volume-to-LUN mappings to an individual host or assign the host to a host group that shares access to one or more volumes.

TABLE 3-11 Volume-to-LUN Terminology

Term	Description
Host Port	The physical connection that allows a host to gain access to the volumes in the storage array. When the host bus adapter only has one physical connection (host port), the terms host port and host bus adapter are synonymous. Host ports can be automatically detected by the storage management software after the storage array has been connected and powered-up. Therefore, if you want to define specific volume-to-LUN mappings for a particular host or create storage partitions, you must define the host's associated host ports. Initially, all detected host ports belong to the Default Group. Therefore, if during volume creation, you had a LUN automatically assigned to a volume, that volume will be accessible by any of the host ports in the Default Group. If you have the SANshare Storage Partitioning feature enabled, then you should always choose to map the volume later using the options in the Mappings View so that a LUN is not automatically assigned to a volume during volume creation. Use the host bus adapter utility to find the World Wide Name (WWN) of the host port. (This is the host port identifier shown in the Define New Host Port dialog in the Mappings View.) The WWNs for the host ports on a particular host are used to define the host ports and associate them with a particular host using the Define New Host Port dialog. If necessary, refer to your operating system or host bus adapter documentation for more information. If a host port is moved, any volume-to-LUN mappings must be re-mapped. Access to your data will be lost until this is done.

TABLE 3-11 Volume-to-LUN Terminology

Term	Description
Logical Unit Number (LUN)	The number a host uses to access a volume on a storage array. Each host has its own LUN address space. Therefore, the same LUN may be used by different hosts to access different volumes on the storage array. However, a volume can only be mapped to a single LUN. A volume cannot be mapped to more than one host group or host. For example, Figure 3-32 on page 3-76 shows that Host KC-A may access Volume Legal using LUN 2 and Host Group Omaha may access Volume HResources also using LUN 2.
Default Volume-to-LUN mapping	During volume creation, you can specify that you want to have the software assign a LUN automatically to the volume or that you want to map a LUN to the volume later. If you have the SANshare Storage Partitioning feature enabled, than you should always choose to map the volume later using the options in the Mappings View so that a LUN is not automatically assigned to a volume during volume creation. Any volumes that are given automatic (default) volume-to-LUN mappings can be accessed by all host groups or hosts that do not have specific volume-to-LUN mappings. These host groups and hosts are shown as part of the Default Group in the Topology section of the Mappings View.
Specific Volume-to-LUN mapping	A specific volume-to-LUN mapping occurs when you select a defined host group or host in the Topology View and select the SANshare Storage Partitioning Wizard or Define Additional Mapping option to assign a volume a specific LUN (volume-to-LUN mapping). This designates that only the selected host group or host has access to that particular volume through the assigned LUN.You can define one or more specific volume-to-LUN mappings for a host group or host. Note The SANshare Storage Partitioning feature must be enabled to create specific mappings.

Register the Volume with the Operating System

Note – IMPORTANT The hot_add utility is not available for all operating systems. Refer to the SANtricity Storage Manager Installation Guide to verify if the hot_add utility is available for your operating system and how to run this utility.

After creating all volumes and assigning volume-to-LUN mappings, the host-based hot_add utility is used to register the volume with the operating system.

Once volumes have been created and volume-to-LUN mappings have been defined, this utility is run to ensure that the operating system is aware of the newly created volumes.

The host-based SM devices utility (if available for your operating system) is used to associate the physical device name and the volume name. Refer to "SANshare Storage Partitioning" on page 3-72 for more information on assigning volume-to-LUN mappings.

3.3.6.14 SANshare Storage Partitioning

This is a premium feature of the storage management software and must be enabled either by you or your storage vendor. The SANshare Storage Partitioning feature enables hosts with different operating systems (heterogeneous hosts) to share access to a storage array.

A storage partition is a logical entity consisting of one or more storage array volumes that can be shared among hosts. To create a storage partition after the total storage capacity has been configured into volumes, you must define a single host or collection of hosts (or host group) that will access the storage array. Then you will need to define a volume-to-LUN mapping, which will allow you to specify the host group or host that will have access to a particular volume in your storage array.

Storage partitions can be created quickly with the SANshare Storage Partitioning Wizard. The Wizard contains the major steps required to specify which hosts, volumes, and associated logical unit numbers (LUNs) will be included in the partition.

Based on the premium feature key purchased, a maximum of 64 storage partitions can be supported by the storage management software.

Note – IMPORTANT Windows NT, Solaris with RDAC, NetWare 5.1, and HP-UX 11.0 are restricted to 32 volumes per partition for this release.

A maximum of 256 volumes per partition can be defined; this is limited to the total number of volumes on your storage array. The software can further support up to two host ports in each host and up to eight ports in each host group, allowing a four-way cluster of dual-adapter hosts.

SANshare Storage Partitioning Example

In the example shown in Figure 3-30 on page 3-73, four hosts (Omaha A and B, and KC-A and B) are connected to Storage Array Midwest. Three storage partitions have been created, allowing these hosts to share access to the volumes on the storage array.

The first partition is composed of Volume Financial. This volume is accessed by Host KC-B using LUN 5. Even though Host KC-B is part of the logical Host Group Kansas City, Host KC-A cannot access this volume because the volume-to-LUN mapping was created with Host KC-B rather than the Host Group Kansas City.

The second partition consists of Volumes Legal and Engineering. This volume-to-LUN mapping was created using Host Group Kansas City. These volumes are accessed by Hosts KC-A and KC-B in Host Group Kansas City using both LUNs 2 and 4.

The third partition consists of Volumes Marketing and HResources. This volume-to-LUN mapping was created using Host Group Omaha. These volumes are accessed by Hosts Omaha A and Omaha B in Host Group Omaha using both LUNs 7 and 2.

Note – A host accesses the volumes on the storage array through the physical host ports residing on the installed host bus adapters. To ensure redundant paths to each volume, each host must have at least two host ports.

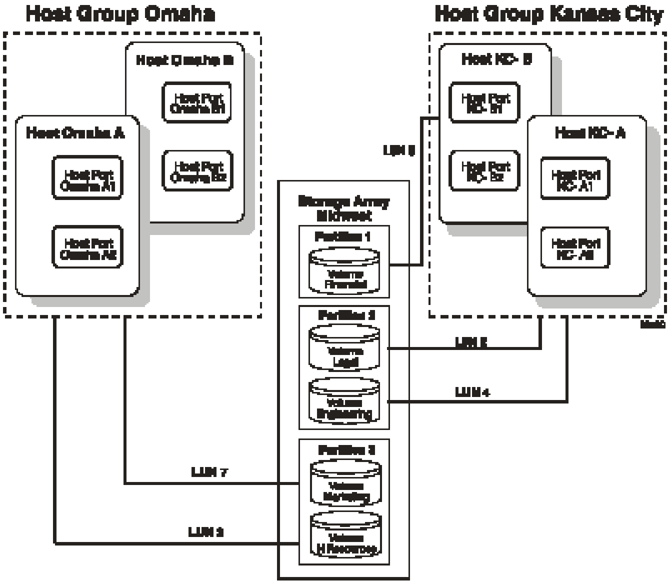

FIGURE 3-30 SANshare Storage Partitioning Example

SANshare Storage Partitioning involves three key steps:

- Create volumes on the storage array. As part of the volume creation, specify one of two volume-to-LUN mapping settings:
  - Automatic - If you are not using SANshare Storage Partitioning, specify this setting. The Automatic setting specifies that a LUN be automatically assigned to the volume using the next available LUN within the Default Group. This setting will grant volume access to host groups or hosts that have no specific volume-to-LUN mappings (designated by the Default Group in the Topology View).
  - Map later with SANshare Storage Partitioning - If you are using SANshare Storage Partitioning, specify this setting. The Map later setting specifies that a LUN not be assigned to the volume during volume creation. This setting allows definition of a specific volume-to-LUN mapping and creation of storage partitions.
- Define the storage partition topology (including host groups, hosts, and host ports) that access the volumes. Storage partition topology is reconfigurable. You can:
  - Move a host port
  - Replace a host port
  - Move a host from one host group into another host group
  - Delete a host group, host, or host port
  - Rename a host group, host, or host port
  - Change a volume-to-LUN mapping
  - Define additional volume-to-LUN mappings
- Grant volume access to defined host groups or hosts by defining volume-to-LUN mappings, using the SANshare Storage Partitioning Wizard. Each host group or host is granted a unique view of partitioned storage. A defined host group or host can either access:
  - Volumes with default volume-to-LUN mappings - The host group or host is part of the Default Group.
  - Volumes to which they have been granted access through a specific volume-to-LUN mapping - The host group or host will be part of a storage partition.

3.3.6.15 Heterogeneous Hosts

The heterogeneous hosts portion of the SANshare Storage Partitioning feature allows hosts running different operating systems to access a single storage array. To specify different operating systems for attached hosts, you must specify the appropriate host type when you define the host ports for each host.

Host types can be completely different operating systems, such as Solaris and Windows NT, or variants of the same operating system, such as Windows NT - clustered and Windows NT - non-clustered. By specifying a host type, you are defining how the controllers in the storage array will work with the particular operating system on the hosts that are connected to it.

Heterogeneous Hosts Example

Note – IMPORTANT Heterogeneous host settings are only available with SANshare Storage Partitioning enabled. In Figure 3-32 on page 3-76, the SANshare Storage Partitioning feature is enabled.

In a heterogeneous environment, you must set each host type to the appropriate operating system during host port definition (Figure 3-31 on page 3-75). By doing this, the firmware on each controller can respond correctly for that host's operating system.

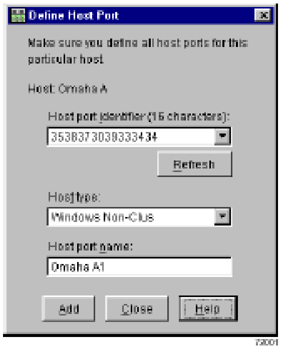

Define Host Port

Make sure you define all host ports for this particular host

Host: Omaha A

Host port identifier (16 characters): 3538373038333434

Host type: Windows Non-Clus

Host port name: Omaha A1

FIGURE 3-31 Host Port Definitions Dialog

In the example shown in Figure 3-32 on page 3-76, four hosts (Omaha A and B, and KC-A and B) are connected to Storage Array Midwest. Three storage partitions have been created, allowing these hosts to share access to the volumes on the storage array.

Hosts Omaha A and Omaha B share access to Volumes Marketing and HResources. Host KC-A has exclusive access to Volumes Legal and Engineering, and Host KC-B has exclusive access to Volume Financial. Because there are four hosts running three different operating systems, the appropriate host types must be defined for each host port to support these heterogeneous hosts.

After you define the host type for each host port, you can display the host port's host type in the Topology View by placing your cursor over the specific host port; a tooltip will display the associated host type.

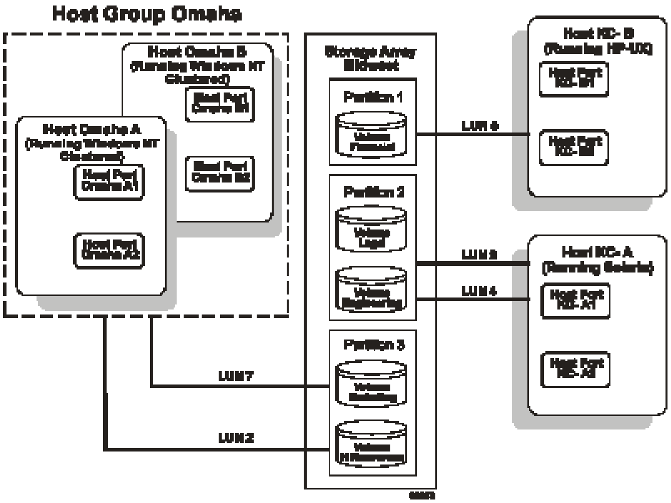

FIGURE 3-32 Heterogeneous Hosts Example

Snapshot Volumes

This is a premium feature of the storage management software and must be enabled either by you or your storage vendor. The Snapshot Volume feature is used to create a logical point-in-time image of another volume.

Typically, you create a snapshot so that an application, for example a backup application, can access the snapshot and read the data while the base volume remains online and user-accessible. When the backup completes, the snapshot volume is no longer needed.

You can also create snapshots of a base volume and write data to the snapshot volumes in order to perform testing and analysis. Before upgrading your database management system, for example, you can use snapshot volumes to test different configurations. Then you can use the performance data provided by the storage management software to help you decide how to configure your live database system.

The maximum number of snapshot volumes allowed is one half of the total volumes supported by your controller model, while the maximum number of snapshot volumes supported by a volume is four.

When a snapshot volume is created, the controller suspends I/O to the base volume for a few seconds while it creates a physical volume, called the snapshot repository volume, to store the snapshot volume metadata and copy-on-write data. Because the only data blocks that are physically stored in the snapshot repository volume are those that have changed since the time the snapshot volume was created, the snapshot volume uses less disk space than a full physical copy.

The storage management software provides a warning message when the snapshot repository volume nears a user-specified threshold (a percentage of its full capacity; the default is 50%). When this condition occurs, the capacity of the snapshot repository volume can be expanded from free capacity on the volume group. If you are out of free capacity on the volume group, unconfigured capacity can be added to the volume group to expand the snapshot repository volume.

The Create Snapshot Volume Wizard is used to create snapshot volumes by defining the snapshot volume's name, the name of its associated snapshot repository volume, and to specify the snapshot repository volume's capacity as a percentage of the base volume's capacity. Either the Simple or Advanced path can be selected through the Create Snapshot Volume Wizard:

- Simple Path - provides a preview of the preconfigured snapshot volume and snapshot repository volume parameters.
- Advanced Path - provides a way to select a free capacity or unconfigured capacity node on which to place the snapshot repository volume, and allows you to change snapshot repository volume parameters. The Advanced Path can be chosen regardless of whether free capacity or unconfigured capacity is selected.

Disabling or Deleting a Snapshot Volume

As long as a snapshot volume is enabled, storage array performance is impacted by the copy-on-write activity to the associated snapshot repository volume. If a snapshot volume is no longer needed, it can be disabled, which will stop the copy-on-write activity.

If the snapshot volume is disabled, you can retain it and its associated snapshot repository volume. When you need to create a different point-in-time image of the same base volume, you can use the re-create option to reuse the disabled snapshot volume and its associated snapshot repository volume. This takes less time than creating a new one.

When you disable a snapshot volume:

- You cannot use the snapshot volume again until you use the re-create option.
- Only the snapshot volume is disabled. All other snapshot volumes remain functional.

If you do not intend to re-create a snapshot volume, you can delete the snapshot volume instead of disabling it. When you delete a snapshot volume, the associated snapshot repository volume is also deleted.

Re-creating a Snapshot Volume

If you have a snapshot volume that you no longer need, instead of deleting it, you can reuse it (and its associated repository volume) to create a different point-in-time image of the same base volume. Re-creating a snapshot volume takes less time than creating a new one.

When you re-create a snapshot volume:

- The snapshot volume must have either an optimal or a disabled state.
- All copy-on-write data previously on the snapshot repository volume is deleted.
- Snapshot volume and snapshot repository volume parameters remain the same as the previously disabled snapshot volume and its associated snapshot repository volume. After the snapshot volume is re-created, you can change parameters on the snapshot repository volume through the appropriate menu options.
- The original names for the snapshot volume and snapshot repository volumes will be retained. You can change these names after the re-create option is completed.

Performing Defragment Operations

Using an operating system-specific defragment utility to perform a defragment operation on a base volume with an associated snapshot repository volume will cause a copy-on-write of every data block in the base volume. This can cause the snapshot repository volume to fill before the defragment operation is completed.

As a result, the snapshot volume will fail or host writes will fail to the base volume, depending on the Snapshot Volume Full condition that was originally set for the snapshot volume.

To prevent this from occurring, ensure that the snapshot repository volumes' capacity is set to at least 105% of the size of the base volume before using a defragment utility. This is the minimum size needed to support a copy-on-write of every data block in the base volume, which will occur as a result of a defragment operation.

Dynamic Volume Expansion (DVE)

Caution – Increasing the capacity of a standard volume is only supported on certain operating systems. If volume capacity is increased on a host operating system that is unsupported, the expanded capacity will be unusable, and you cannot restore the original volume capacity. For information on supported operating systems, refer to the SANtricity Storage Manager Product Release Notes shipped with the storage management software.

Dynamic Volume Expansion (DVE) is a modification operation used to increase the capacity of standard or snapshot repository volumes. The increase in capacity can be achieved by using any free capacity available on the volume group of the standard or snapshot repository volume.

Data will be accessible on volume groups, volumes, and disk drives throughout the entire modification operation.

During the modification operation, the volume having its capacity increased shows a status of Operation in Progress, together with its original capacity and the total capacity being added (Figure 3-33). After the increase in capacity is completed, the volumes expanded capacity is displayed, and the final capacity for the Free Capacity node involved will show a reduction in capacity. If all of the free capacity is used to increase the volumes size, then the Free Capacity node involved will be removed from the Logical View.

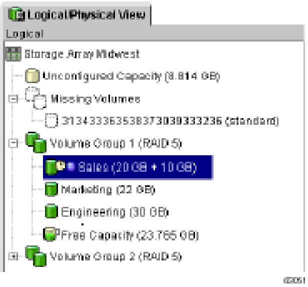

FIGURE 3-33 DVE Modification Operation in Progress

An increase in storage capacity for snapshot repository volumes would be completed if a warning is received that the snapshot repository volume is in danger of becoming full. Increasing the capacity of a snapshot repository volume does not increase the capacity of the associated snapshot volume. The snapshot volume's capacity is always based on the capacity of the base volume at the time the snapshot volume is created.

Remote Volume Mirroring

This is a premium feature of the storage management software and must be enabled either by you or your storage vendor. The Remote Volume Mirroring feature is used for online, real-time replication of data between storage arrays over a remote distance.

Prior to creating a mirror relationship, the Remote Volume Mirroring feature must be enabled and activated on both the primary and secondary storage arrays. The primary volume is the volume that accepts host I/O and stores application data. When you create a remote volume mirror, a mirrored volume pair is created and consists of a primary volume at the primary storage array and a secondary volume at the secondary storage array.

Data from the primary volume is copied in its entirety to the secondary volume. The secondary volume maintains a mirror (or copy) of the data from its associated primary volume. The secondary volume remains unavailable to host applications while mirroring is underway. In the event of a disaster or catastrophic failure of the primary site, the secondary volume can be promoted to a primary role.

For detailed information on this premium feature, refer to the Array Management Window online help or to the SANtricity Storage Manager Remote Volume Mirroring Feature Guide.

Mirror Relationships

A secondary volume must be created on the secondary site if one does not already exist and must be a standard volume of equal or greater capacity than the associated primary volume.

When a secondary volume is available, a mirror relationship can be established in the storage management software by identifying the storage array containing the primary volume and the storage array containing the secondary volume.

Mirror Repository Volumes

A mirror repository volume is a special volume in the storage array created as a resource for the controller owner of the primary volume in a Remote Volume Mirror. The controller stores mirroring information on this volume, including information about remote writes that are not yet complete. The controller can use this information to recover from controller resets and accidental powering-down of storage arrays.

When you activate the Remote Volume Mirroring feature on the storage array, you create two mirror repository volumes, one for each controller in the storage array. An individual mirror repository volume is not needed for each Remote Volume Mirror.

When you create the mirror repository volumes, you specify the location of the volumes. You can either use existing free capacity or you can create a volume group for the volumes from unconfigured capacity and then specify the RAID level.

Data Replication

Data replication between the primary volume and the secondary volume is managed by the controllers and is transparent to host machines and applications. When the controller owner of the primary volume receives a write request from a host, the controller first logs information about the write to a mirror repository volume, then writes the data to the primary volume. The controller then initiates a remote write operation to copy the affected data blocks to the secondary volume at the secondary storage array.

After the host write request has been written to the primary volume and the data has been successfully copied to the secondary volume, the controller removes the log record on the mirror repository volume and sends an I/O completion indication back to the host system.

Volume Copy

This is a premium feature of the storage management software and must be enabled either by you or your storage vendor. The Volume Copy premium feature is used to copy data from one volume (the source) to another volume (the target) in a single storage array.

This feature can be used to back up data, to copy data from volume groups that use smaller capacity drives to volume groups that use larger capacity drives, or to restore snapshot volume data to the base volume. For detailed information on this premium feature, refer to the Array Management Window online help or to the SANtricity Storage Manager Volume Copy Feature Guide.

Source Volume

When you create a volume copy, a copy pair is created and consists of a source volume and a target volume located on the same storage array. The source volume is the volume that accepts host I/O and stores application data, and can be a standard volume, snapshot volume, base volume of a snapshot volume, or a Remote Volume Mirror primary volume.

Note – IMPORTANT If a primary volume is selected as the source volume for a volume copy, you must ensure that the capacity of the target volume is equal to, or greater than the usable capacity of the primary volume. The usable capacity for the primary volume is the minimum of the primary and secondary volume's actual capacities.

When a volume copy is started, data from the source volume is copied in its entirety to the target volume. While the volume copy has a status of In Progress, Pending, or Failed, the source volume is available for read I/O activity. After the volume copy is completed, write requests are allowed to the source volume.

Target Volume

Caution – A volume copy will overwrite all data on the target volume. Ensure that you no longer need the data or have backed up the data on the target volume before starting a volume copy.

A target volume maintains a copy of the data from the source volume, and can be a standard volume, the base volume of a Failed or Disabled snapshot volume, or a Remote Volume Mirror primary volume in an active mirrored pair.

Note – IMPORTANT The target volume capacity must be equal to or greater than the source volume capacity.

When a volume copy is started, data from the source volume is copied in its entirety from the source volume to the target volume.

While the volume copy has a status of In Progress, Pending, or Failed, no read or write requests to the target volume will be allowed. After the volume copy is complete, the target volume automatically becomes read-only to hosts, and write requests to the target volume will not take place. The Read-Only attribute can be changed in the Copy Manager only after the volume copy is completed.

Creating a Volume Copy

The Create Copy Wizard guides you through the process of selecting a source volume from a list of available volumes, selecting a target volume from a list of available volumes, and setting the copy priority for the volume copy. After you have completed the Wizard dialogs, the volume copy starts and data is read from the source volume and then written to the target volume.

Caution – A volume copy will overwrite all data on the target volume and automatically make the target volume read-only to hosts. After the volume copy completes, you can use the Copy Manager to disable the Read-Only attribute for the target volume.

The Copy Manager allows you to monitor the volume copy after it has been created. From the Copy Manager, a volume copy may be re-copied, stopped, or removed, and its attributes, including the copy priority and the target volume Read-Only attribute, can be modified. The status of a volume copy can be viewed in the Copy Manager. Also, if you need to find out what volumes are involved in a volume copy, use the Copy Manager or the Storage Array Profile.

3.3.6.16 Managing Persistent Reservations

Caution – The Persistent Reservations option should be used only under the guidance of a technical support representative.

The Persistent Reservations option enables you to view and clear volume reservations and associated registrations. Persistent reservations are configured and managed through the cluster server software, and prevent other hosts from accessing particular volumes. Unlike other types of reservations, a persistent reservation reserves access across multiple host ports, provides various levels of access control, offers the ability to query the storage array about registered ports and reservations, and optionally, provides for persistence of reservations in the event of a storage system power loss.

The storage management software provides functionality for managing persistent reservations in the Array Management Window (Figure 3-34). The Persistent Reservation option enables you to:

- View registration and reservation information for all volumes in the storage array
- Save detailed information on volume reservations and registrations
- Clear all registrations and reservations for a single volume or for all volumes in the storage array

For detailed procedures, refer to the Array Management Window online help.

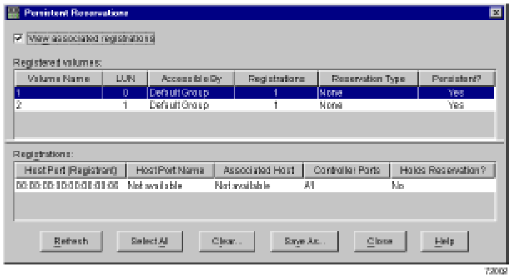

Persistent Reservations

☑ View associated registrations

Registered volumes:

Volume Name	LUN	Accessible By	Registrations	Reservation Type	Persistent?
1	0	Default Group	1	None	Yes
2	1	Default Group	1	None	Yes

Registrations:

Host Port (Registrant)	Host Port Name	Associated Host	Controller Ports	Holds Reservation?
00:00:00:00:00:00:00:00	Not available	Not available	All	No

FIGURE 3-34 Persistent Reservations Dialog

Management of persistent reservations through the script engine and command line interface is also supported. For more information, refer to the Enterprise Management Window online help.

3.3.6.17 Maintaining and Monitoring Storage Arrays

This section describes methods for maintaining storage arrays, including troubleshooting storage array problems, recovering from a storage array problem using the Recovery Guru, and configuring alert notifications using the event monitor. For additional conceptual information and detailed procedures for the options described in this section, refer to Learn About Monitoring Storage Arrays in the Enterprise Management Window online help.

Storage Array Health

Note – IMPORTANT The Enterprise Management Window or the event monitor must be running to receive notification of critical events for the storage arrays. In addition, alert notifications must be configured in the Enterprise Management Window.

The Enterprise Management Window provides a summary of the conditions of all known storage arrays being managed. Appropriate status indicators will be shown in the Device Tree, the Device Table, and in the health summary status area in the lower-left corner of the Enterprise Management Window (Figure 3-35).

The *Device Tree* in this window provides a status for each storage array being managed. A status is shown for each network management connection.

The *Device Table* contains a status column that displays the status of each storage array.

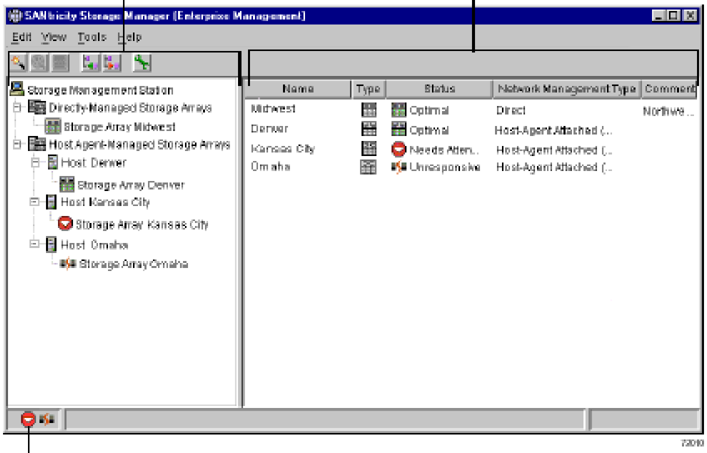

- Storage Management Station
  - Directly-Managed Storage Arrays
    - Storage Array Midwest
  - Host Agent-Managed Storage Arrays
    - Host Denver
      - Storage Array Denver
    - Host Kansas City
      - Storage Array Kansas City
    - Host Omaha
      - Storage Array Omaha

Name	Type	Status	Network Management Type	Comment
Midwest	▦	Optimal	Direct	Northwe...
Denver	▦	Optimal	Host-Agent Attached (...	
Kansas City	▦	Needs Atten...	Host-Agent Attached (...	
Omaha	▦	Unresponsive	Host-Agent Attached (...	

The *Overall Health Status* is a consolidated health status for all of the known storage arrays.

FIGURE 3-35 Monitoring Storage Array Health Using the Enterprise Management Window

Storage Array Status Icons

Table 3-12 provides information about the storage array status icons that display:

- In the Device Tree, Device Table, and Overall Health Status area in the Enterprise Management Window
- As the Root Node in the Logical View Tree in the Array Management Window

TABLE 3-12 Storage Array Status Icon Quick Reference

Status	Description
Optimal	Indicates every component in the storage array is in the desired working condition.
Needs Attention	Specifies a problem on a storage array that requires intervention to correct. To correct the problem, start the Array Management Window for that particular storage array, and then use Recovery Guru to pinpoint the cause of the problem and obtain appropriate procedures.
Fixing	Signifies a Needs Attention condition has been corrected and the storage array is transitioning to an Optimal status; for example, a reconstruction operation is in progress. A Fixing status requires no action unless you want to check on the progress of the operation in the Array Management Window. Note Some recovery actions cause the storage array status to change directly from Needs Attention to Optimal, without an interim status of Fixing. This icon is not displayed in the Overall Health Status area. The Optimal status icon is displayed instead.
Unresponsive	Means the storage management station cannot communicate with the only controller or both controllers in the storage array over its network management connection. Note This icon is not displayed in the Logical View of the Array Management Window. If the Array Management Window is open and the storage array becomes Unresponsive, the last known status icon (Optimal, Needs Attention, or Fixing) is shown.
Contacting Device	Designates that you have started the Enterprise Management Window and the storage management software is establishing contact with the storage array. Note This icon is not displayed in the Logical View of the Array Management Window.

Event Monitor

The event monitor runs continuously in the background monitoring activity on a storage array and checking for critical problems (for example, impending drive failures or failed controllers). If the event monitor detects any critical problems, it can notify a remote system using e-mail and/or simple network management protocol (SNMP) trap messages whenever the Enterprise Management Window is not running.

The event monitor is a separate program bundled with the client software and must be installed with the storage management software. The client/event monitor is installed on a storage management station or host connected to the storage arrays. For continuous monitoring, install the event monitor on a computer that runs 24 hours a day. Even if you choose not to install the event monitor, alert notifications must still be configured on the computer where the client software is installed, because alerts will be sent as long as the Enterprise Management Window is running.

Figure 3-36 shows how the event monitor and the Enterprise Management Window client software send alerts to a remote system. The storage management station contains a file with the name of the storage array being monitored and the address where alerts will be sent. The alerts and errors that occur on the storage array are continuously being monitored by the client software and the event monitor. The event monitor takes over for the client after the client software package is shut down. When an event is detected, a notification is sent to the remote system.

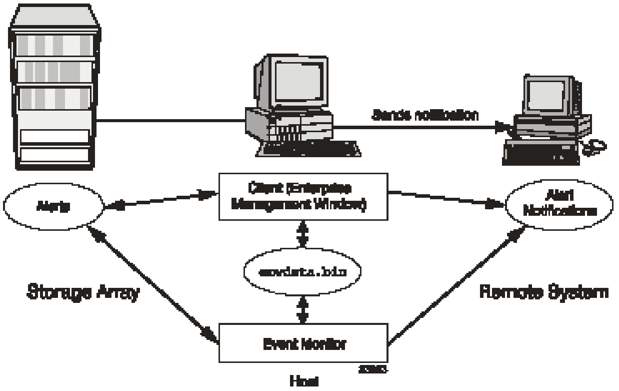

FIGURE 3-36 Event Monitor Configuration

Because the event monitor and the Enterprise Management Window share the information to send alert messages, the Enterprise Management Window contains some visual cues to assist in the event monitor installation and synchronization. The parts of the Enterprise Management Window that are related to event monitoring are shown in Figure 3-37 on page 3-88.

Using the event monitor involves three key steps:

1. **Install the client software. The event monitor is packaged with the client software and installs automatically with the client software. It is recommended that you run the event monitor on one machine that will run continuously. To prevent receipt of duplicate alert notifications of the same critical event on a storage array, disable the event monitor on all but one storage management station.**

You must have administrative permissions to install software on the computer where the event monitor will reside. After the storage management software has been installed, the icon shown in Figure 3-37 on page 3-88 will be present in the lower-left corner of the Enterprise Management Window.

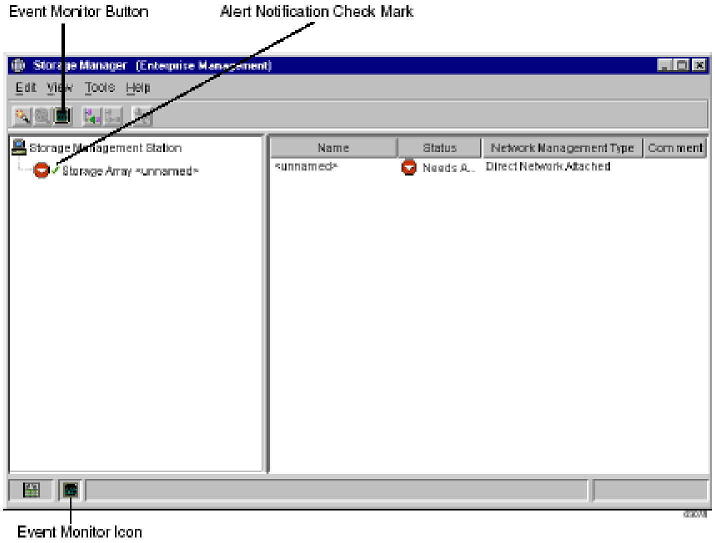
Event Monitor Button
Alert Notification Check Mark
Storage Manager (Enterprise Management)
Edit View Tools Help
Storage Management Station
Storage Array <unnamed>
Event Monitor Icon

Name	Status	Network Management Type	Comment
<unnamed>	Needs A...	Direct Network Attached	

FIGURE 3-37 Event Monitor Example

2. **Set up the alert destinations for the storage arrays you want to monitor from the Enterprise Management Window. A check mark indicates where the alert is set (storage management station, host, or storage array). When a critical problem occurs on the storage array, the event monitor will send a notification to the appropriate alert destinations that were specified.**

3. **Synchronize the Enterprise Management Window and the event monitor. After it has been installed, the event monitor continues to monitor storage arrays and send alerts as long as it continues to run. If you make a configuration change in the Enterprise Management Window, such as adding or removing a storage array or setting additional alert destinations, you should manually synchronize the Enterprise Management Window and the event monitor.**

Alert Notifications

Alert notification settings must be configured to receive e-mail or SNMP notifications if a critical event occurs on a storage array. The notification displays a summary of the critical event and details about the affected storage array, including:

- Name of the affected storage array
- Host IP address (only for a storage array managed through a host-agent)
- Host name/ID (shown as directly managed if the storage array is managed through each controller's Ethernet connection)
- Event error type related to an Event Log entry
- Date and time when the event occurred
- Brief description of the event

Note – IMPORTANT To set up alert notifications using SNMP traps, you must copy and compile a management information base (MIB) file on the designated network management station. Refer to the SANtricity Storage Manager Installation Guide for more information.

There are three key steps involved in configuring alert notifications:

1. **Select a node in the Enterprise Management Window that will display alert notifications for the storage arrays you want to monitor. You can set the alert notifications at any level:**
   - Every storage array being managed
   - Every storage array attached and managed through a particular host
   - An individual storage array
2. **Configure e-mail destinations, if desired. You must provide a mail server name and an e-mail sender address for the e-mail addresses to work.**
3. **Configure SNMP trap destinations, if desired. The SNMP trap destination is the IP address or the host name of a station running an SNMP service, such as a Network Management Station.**

   Customer Support Alert Notifications

Note – IMPORTANT If you do not configure the customer support alert notifications option, thee-mail alert notification will contain only a summary of the critical event. If you do configure this option, all specified e-mail addresses will receive the summary, detailed information about the affected storage array, and the specified contact information.

The Enterprise Management Window contains options to configure the system to send e-mail notifications to a specified customer support group if a critical event occurs on a storage array. After it is configured, the e-mail alert notification includes a summary of the critical event, details about the affected storage array, and customer contact information. Contact technical support for more information about setting up this file.

Configuring customer support alert notifications involves the following:

1. **Create a text file containing the contact information you want to send to the customer support group. For example, include the names and pager numbers of the system administrators.**

2. **Name the file userdata.txt and save it in the home directory (for example, Winnt\ profiles\) on the client machine you are using to manage the storage array (This may be your host machine if you installed the client software on the host.).**

3. **Configure the alert notifications using e-mail or SNMP trap destinations.**

Problem Notification

Note – IMPORTANT The Enterprise Management Window or the event monitor must be running to receive notification of critical events for the storage arrays. In addition, you must have configured the alert notifications in the Enterprise Management Window.

Use Recovery Guru to help troubleshoot storage array problems. Where necessary, use the hardware documentation in conjunction with the recovery steps to replace failed components.

Typically, storage array problems are indicated by:

- A Needs Attention status icon displayed in:
  - The Overall Health Status area, Device Tree View, and Device Table of the Enterprise Management Window.
  - The Array Management Window Logical View.
- The Recovery Guru Optimal toolbar button in the Array Management Window changes from an Optimal to a Needs Attention status and flashes.
- Non-optimal component icons are displayed in the Array Management Window Logical and Physical View.
- Receipt of critical SNMP or e-mail notifications.
- The hardware fault lights display.

In Figure 3-38 on page 3-91, the Array Management Window for Storage Array Engineering indicates a components problem and a Needs Attention status in the Logical/Physical View.

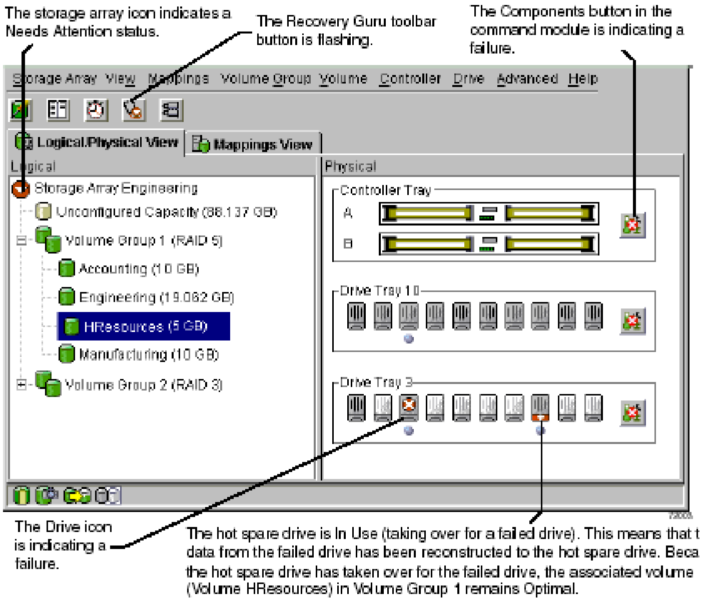

The storage array icon indicates a Needs Attention status.

The Recovery Guru toolbar button is flashing.

The Components button in the command module is indicating a failure.

The Drive icon is indicating a failure.

The hot spare drive is In Use (taking over for a failed drive). This means that t data from the failed drive has been reconstructed to the hot spare drive. Beca the hot spare drive has taken over for the failed drive, the associated volume (Volume HResources) in Volume Group 1 remains Optimal.

FIGURE 3-38 Problem Notification in the Array Management Window

Storage Array Problem Recovery

When you suspect a storage array problem, launch the Recovery Guru. The Recovery Guru is a component of the Array Management Window that will diagnose the problem and provide the appropriate procedure to use for recovery. The Recovery Guru can be displayed by selecting the Recovery Guru toolbar button in the Array Management Window, shown in Figure 3-39 on page 3-92.

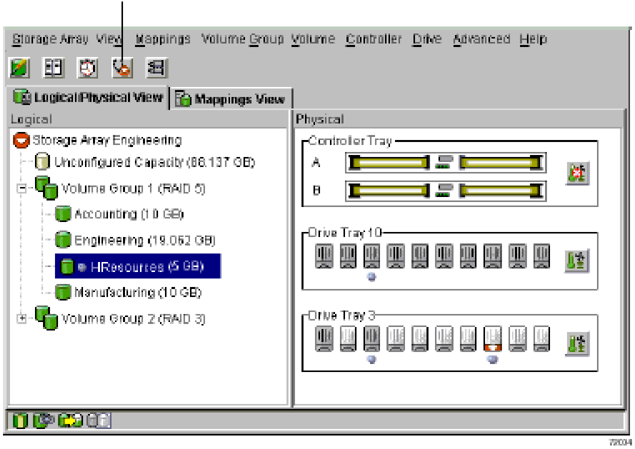
The Recovery Guru Toolbar Button

FIGURE 3-39 Displaying the Recovery Guru Window

Recovery Guru Example

The Recovery Guru window is divided into three views: Summary, Details, and Recovery Procedures. The Summary view presents a list of storage array problems. The Details view displays information about the selected problem in the Summary area. The Recovery Procedure view lists the appropriate steps to follow for the selected problem in the Summary view.

For example, in Figure 3-40 on page 3-93, the Summary area displays two different failures in this storage array; a hot spare in use and a failed battery canister. The Details area shows that in Volume HResources, a hot spare drive in tray 10, slot 6 has replaced a failed drive in tray 3, slot 7. The Recovery Procedure window explains the cause of the selected problem (seen in the Summary view), and describes the appropriate procedures needed to recover from this failure.

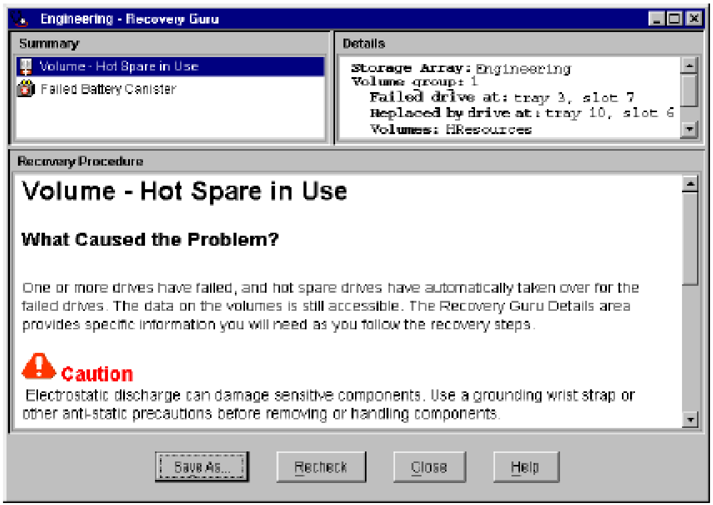

Volume - Hot Spare in Use
Failed Battery Canister

```
Storage Array: Engineering
Volume group: 1
   Failed drive at: tray 3, slot 7
   Replaced by drive at: tray 10, slot 6
   Volumes: HResources
```

Volume - Hot Spare in Use

What Caused the Problem?

One or more drives have failed, and hot spare drives have automatically taken over for the failed drives. The data on the volumes is still accessible. The Recovery Guru Details area provides specific information you will need as you follow the recovery steps.

⚠ Caution

Electrostatic discharge can damage sensitive components. Use a grounding wrist strap or other anti-static precautions before removing or handling components.

FIGURE 3-40 Recovery Guru Window Example

As you follow the recovery procedure to replace the failed drive, the storage array status changes to Fixing, the associated volume (HResources) status changes from failed to Degraded-Copyback in Progress, and the replaced drive status changes to Replaced. The data that was reconstructed to the hot spare drive is now being copied back to the replaced drive. These changes are shown in Figure 3-41 on page 3-94.

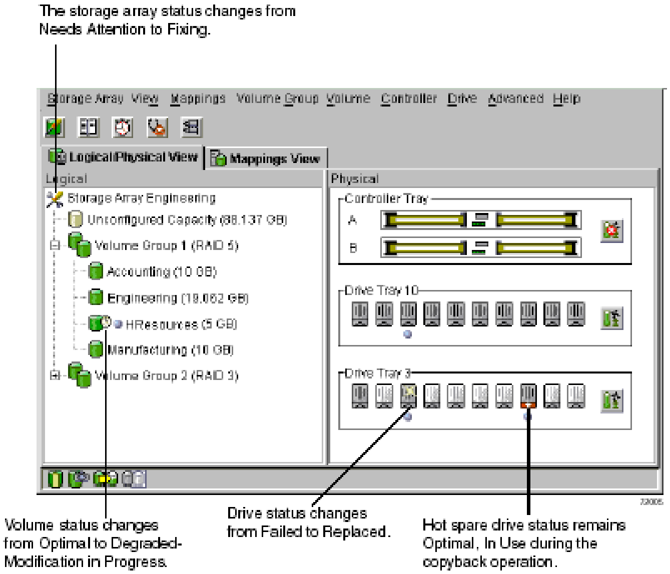
The storage array status changes from Needs Attention to Fixing.
Storage Array View Mappings Volume Group Volume Controller Drive Advanced Help
Logical/Physical View
Mappings View
Logical
Storage Array Engineering
Unconfigured Capacity (88.137 GB)
Volume Group 1 (RAID 5)
Accounting (10 GB)
Engineering (18.062 GB)
HResources (5 GB)
Manufacturing (10 GB)
Volume Group 2 (RAID 3)
Physical
Controller Tray
A
B
Drive Tray 10
Drive Tray 3
72005
Volume status changes from Optimal to Degraded-Modification in Progress.
Drive status changes from Failed to Replaced.
Hot spare drive status remains Optimal, In Use during the copyback operation.

FIGURE 3-41 Status Changes During an Example Recovery Operation

When the copyback operation is finished, the status change to reflect the optimal status of the components, as shown in Figure 3-42 on page 3-95.

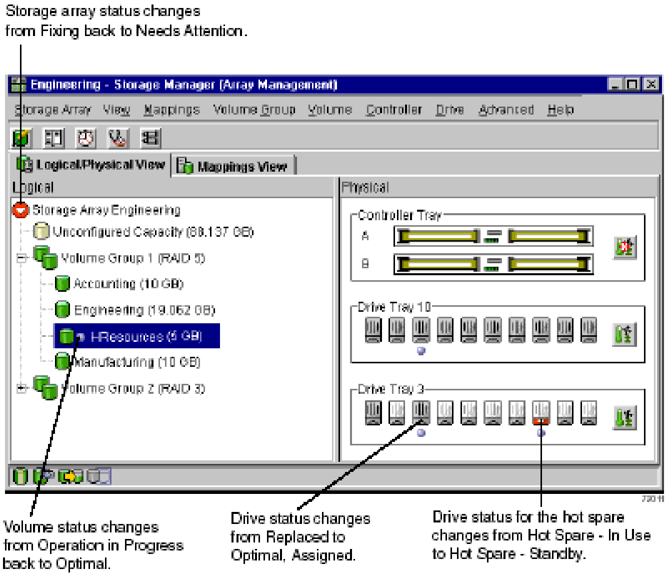

Storage array status changes from Fixing back to Needs Attention.

Volume status changes from Operation in Progress back to Optimal.

Drive status changes from Replaced to Optimal, Assigned.

Drive status for the hot spare changes from Hot Spare - In Use to Hot Spare - Standby.

FIGURE 3-42 Status Changes When The Example Recovery Operation is Completed

After you replace the failed drive in the drive module:

- The storage array status in the Logical View returns to Optimal.
- The storage array status in the Enterprise Management Window changes from Needs Attention to Optimal.
- The Recovery Guru button stops blinking.

Note – For the Recovery Guru button to register Optimal status, the failed battery must be replaced as well.

3.4 Updating Firmware and NVSRAM on the Array

To upgrade the controller firmware and NVSRAM, do the following:

1. **Start the SMclient.**

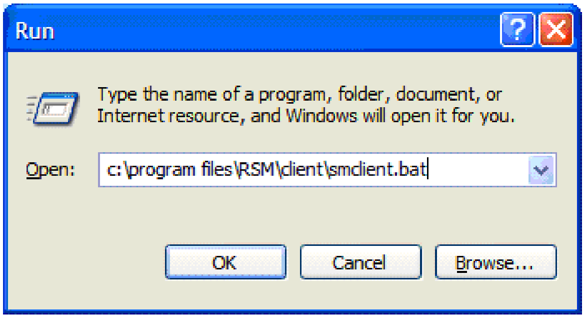

Run

Type the name of a program, folder, document, or Internet resource, and Windows will open it for you.

Open: c:\program files\RSM\client\smclient.bat

A window similar to the one below is displayed after a short delay.

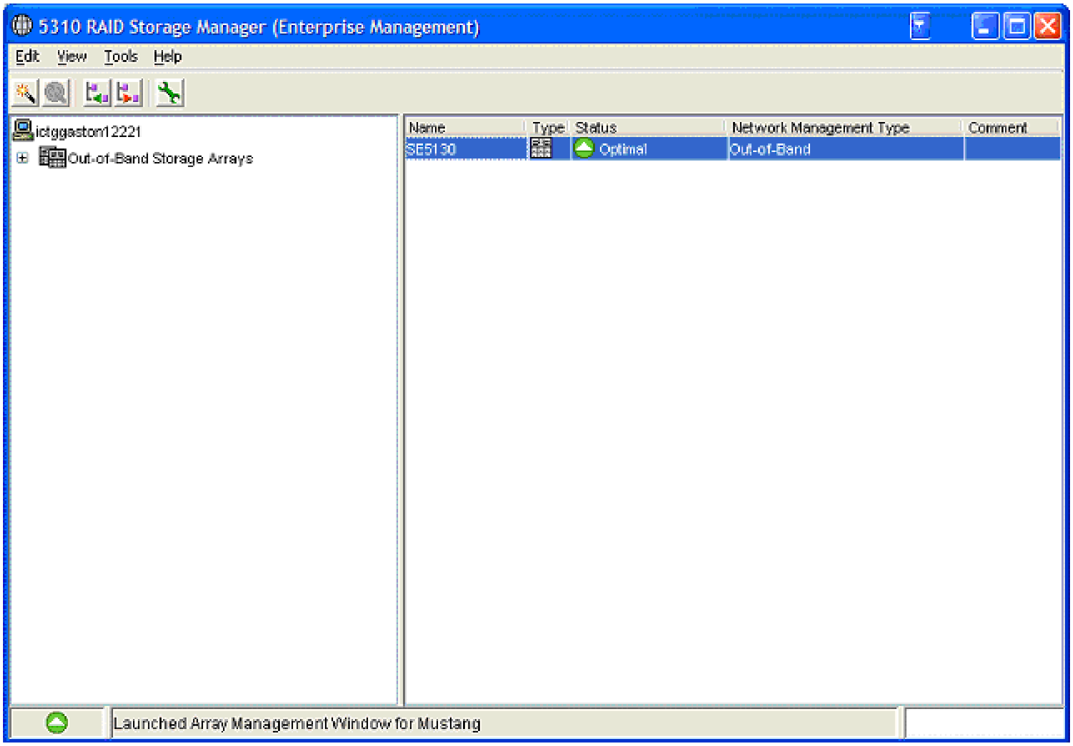

5310 RAID Storage Manager (Enterprise Management)

Name	Type	Status	Network Management Type	Comment
SE5130		Optimal	Out-of-Band	

Launched Array Management Window for Mustang

If the storage array does display in the client window, also known as the Enterprise Management Window (EMW), do the following:

a. **Select Edit>>Add Device and add the storage array by entering the IP address for the A controller and then click Add.**

b. **When the Add Device window returns, enter the IP address for the B controller and click Add.**

c. **When the Add Device window returns, click close.**

2. **Double-click the array name in the right window.**

   The Array Management Window (AMW) is displayed.

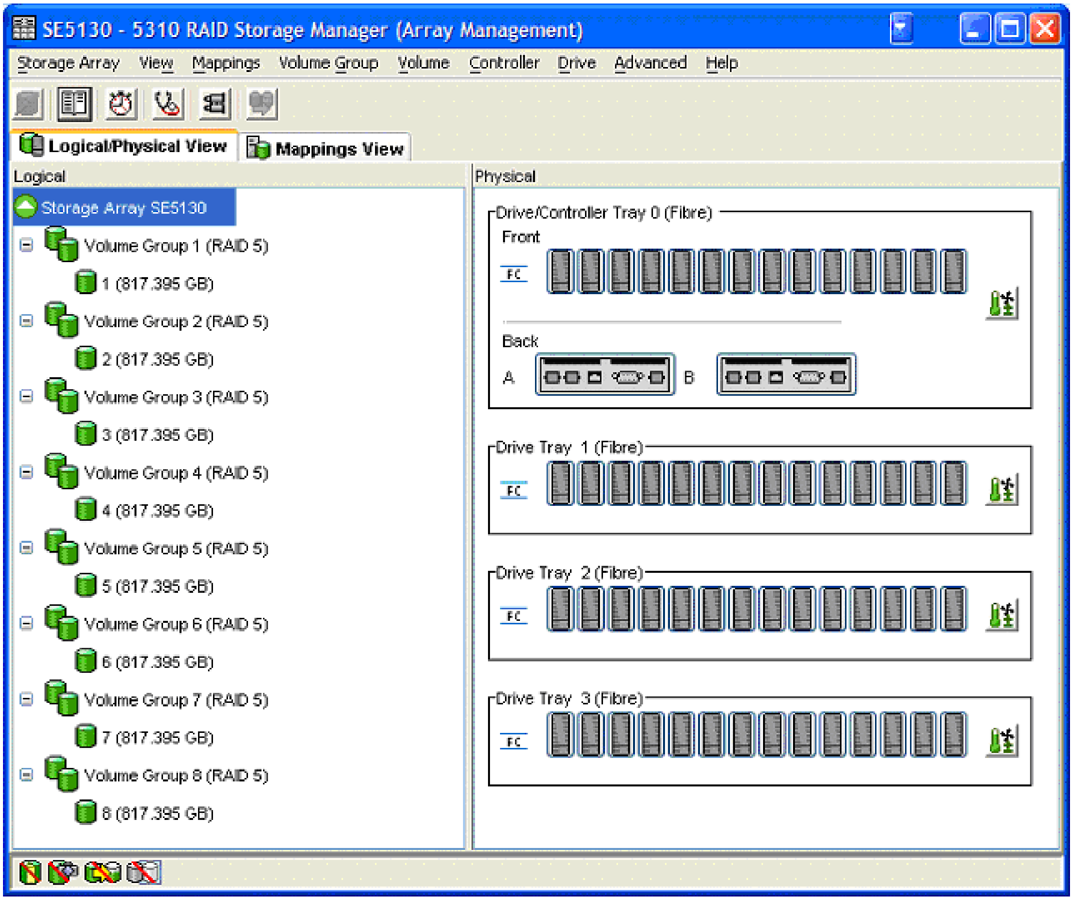

3. **Select Advanced>>Maintenance>>Download>>Controller Firmware.**

   The Download Firmware screen is displayed.

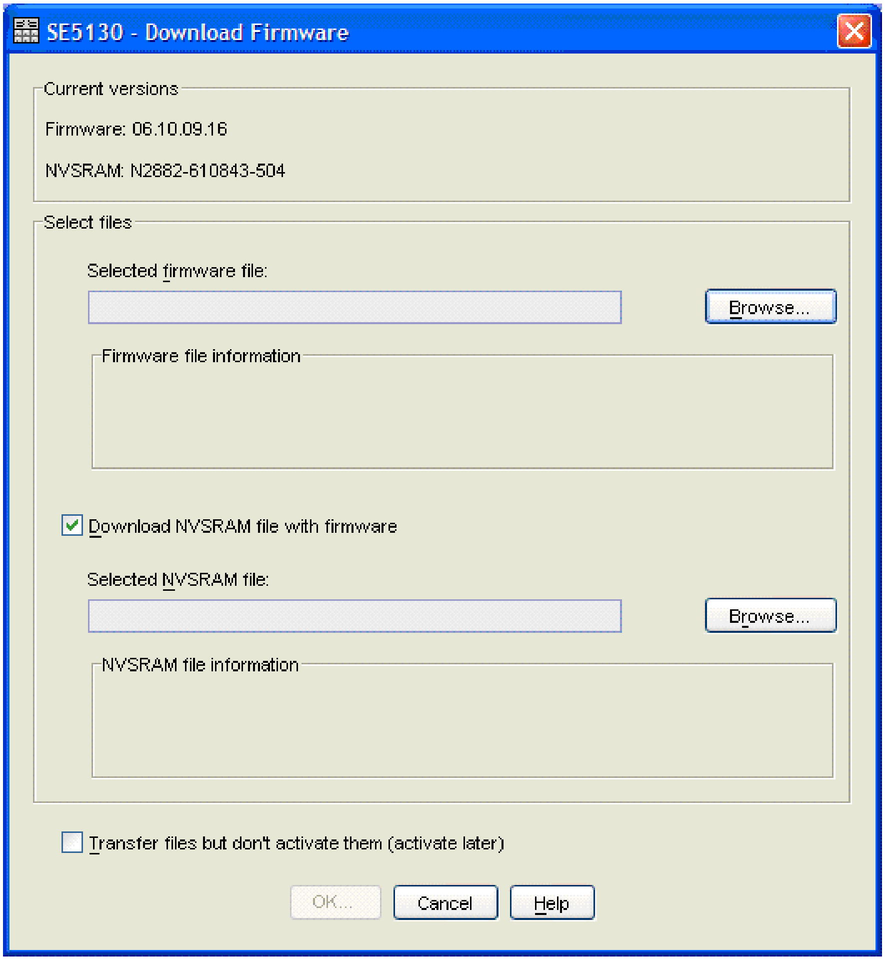

SE5130 - Download Firmware

Current versions

Firmware: 06.10.09.16

NVSRAM: N2882-610843-504

Select files

Selected firmware file:

Firmware file information

☑ Download NVSRAM file with firmware

Selected NVSRAM file:

NVSRAM file information

□ Transfer files but don't activate them (activate later)

4. **Click the browse button to navigate to the directory where the firmware is located.**

   (Example, c:\ temp\rsmfw \) Select SNAP_0610xxxx.dlp.

5. **If upgrading NVSRAM, check the box marked "Download NVSRAM file with firmware" and using the browse button, navigate to the directory where the NVSRAM file is located.**

   (Example, c:\temp\rsmfw\) Select N2882-610843-5xx.dlp.

6. **Make sure the checkbox marked "Transfer files but don't activate them (activate later)" is clear.**

7. **After selecting the firmware and NVSRAM files click OK and then Yes to begin the firmware upgrade.**

   The downloading screen is displayed.

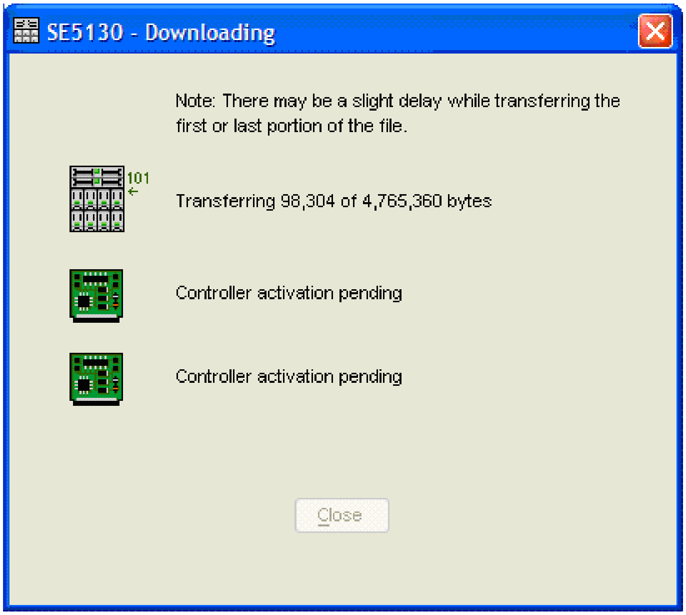

SE5130 - Downloading

Note: There may be a slight delay while transferring the first or last portion of the file.

Transferring 98,304 of 4,765,360 bytes

Controller activation pending

Controller activation pending

3.5 Updating ESM Firmware

To update the ESM (CSM100_E_FC_S) firmware, do the following:

1. **Select Advanced>>Maintenance>>Download>>ESM firmware.**
2. **Select the drive tray(s) to upgrade.**
3. **Browse to where the file is located (Example, c:\temp\rsmfw\)**
4. **Select ESM FW**

5. **Select Start to begin the upgrade.**

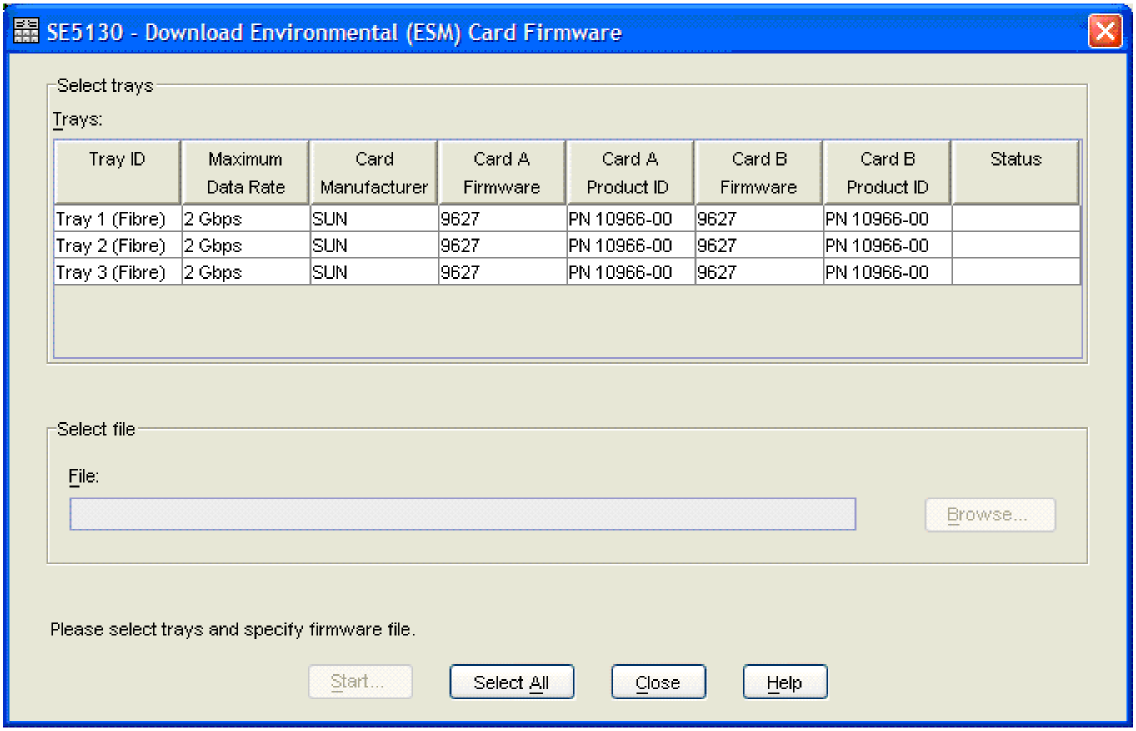

SE5130 - Download Environmental (ESM) Card Firmware

Select trays

Trays:

Tray ID	Maximum Data Rate	Card Manufacturer	Card A Firmware	Card A Product ID	Card B Firmware	Card B Product ID	Status
Tray 1 (Fibre)	2 Gbps	SUN	9627	PN 10966-00	9627	PN 10966-00	
Tray 2 (Fibre)	2 Gbps	SUN	9627	PN 10966-00	9627	PN 10966-00	
Tray 3 (Fibre)	2 Gbps	SUN	9627	PN 10966-00	9627	PN 10966-00	

Select file

File:

Please select trays and specify firmware file.

CHAPTER 4

StorEdge File Replicator

This chapter provides an overview of the StorEdge File Replicator.

4.1 Overview

The Sun StorEdge 5310 NAS data appliances provide fault-tolerant features such as redundant hardware devices, extensive monitoring and notification of both software and hardware components, checkpointing and controller and server Failover. StorEdge File Replicator extends these capabilities to include mirroring. This section provides an overview of the StorEdge File Replicator feature.

A discussion of the different types of mirroring is necessary to lay the groundwork for a discussion of the StorEdge File Replicator mirroring implementation. Mirroring implementations can be loosely categorized into 3 buckets:

- Checkpoint (or Snapshot) Mirroring
- Real-time (or Synchronous) Mirroring
- Pseudo Real-time (or Asynchronous) Mirroring

Throughout this document, some standard terms are used and are defined as follows:

TABLE 4-1 Standard Terms

Term	Definition
Master	The system that is being mirrored or the source system
Mirror	The system that is being used to mirror the Master system, or the target system
Checkpoint / Snapshot	A static image of the file system at a fixed point in time
Client	A network computer that initiates a read or write request
Delta	The filesystem blocks that have changed during a fixed period of time, usually between successive checkpoints
Disaster Recovery (DR)	The act of recovering access to computer systems, networks and data subsequent to a catastrophe, e.g., the loss of a Datacenter
Synchronous Mirroring	Transaction complete is not reflected back to the client until the transaction has been committed to both the Master and Mirror systems
Asynchronous Mirroring	Transaction complete is reflected back to the client when the transaction is committed to the Master system
Mirror	Abstract definition of the system involving a master volume and mirror volume, controlled by the mirror service.
Master System	The Sun StorEdge 5310 NAS system on which the source, or live, volume is located.
Mirror System	The Sun StorEdge 5310 NAS system on which the target, or duplicate, volume is located.
NBD	Network Block Device. A network interface to a remote volume. NBD is used as a transport mechanism used by the mirror service. NBD is also the partition type for the remote mirror volume.
Quality of Service (QOS)	For the purposes of this document, a generic term referring to the quality of service provided to a network user or system over a network link. It should not be confused with QoS, which is a standard for controlling data/packet flow on an IP network.
Transaction Complete	A confirmation from the storage subsystem to the client that a write transaction has been committed to disk.

4.1.1 Real-time Mirroring

Real-time mirroring is the simplest to describe, and the most difficult and expensive to implement. The requirement and guarantee of real-time mirroring is that data is committed in a persistent manner on both the Master and Mirror prior to reflecting transaction complete to the client. If the mirroring is remote, e.g., over a WAN, the expense can be quite great because the user must ensure that the link between the systems is very fast and of exceptional quality, or risk serious reductions in the quality of service locally due to the latencies associated with remote communication. Real-time mirroring systems typically provide for extensive parameter control to enable the user to define policies that manage the mirror link. For instance, the user may want to automatically break the link if serious local QOS issues arise due to telecommunications issues. If the link were broken, the systems would function in pseudo real-time mode until the mirroring system 'caught up' with the mirrored system, at which time the real-time link would be automatically reinstated.

Real-time mirroring is frequently referred to as synchronous mirroring because of the requirement to commit the transaction both locally and remotely prior to reflecting transaction complete to the client.

4.1.2 Pseudo Real-time Mirroring

Pseudo real-time mirroring provides mirroring capabilities approximating those of real-time mirroring, but does not require that transaction complete be received from the mirroring system before reflecting back transaction complete to the client. Pseudo real-time mirroring would typically be implemented where greater than Checkpoint mirroring protection were required, but factors such as economics, infrastructure etc. precluded the need for or feasibility of real-time mirroring.

As with Checkpoint mirroring, pseudo real-time mirroring is an asynchronous operation.

4.1.3 StorEdge File Replicator

The StorEdge File Replicator employs a Pseudo Real-time mirroring approach. A diagram of the lifecycle of a transaction follows:

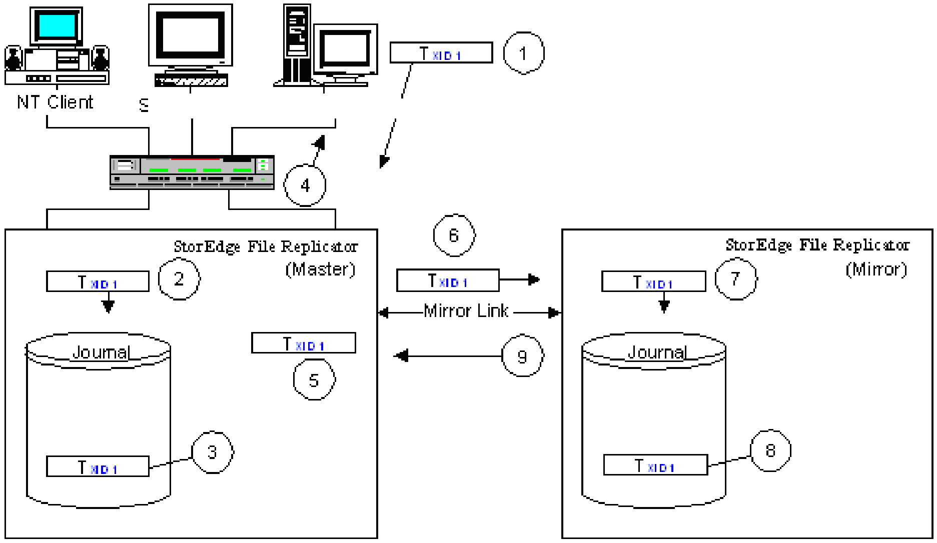

1)A client issues a write for transaction TXID 1

2)Sun StorEdge 5310 NAS receives transaction TXID 1

3)Transaction TXID 1 is committed to the Master system's journal

4)Transaction complete is reflected back to the client for TXID 1

5)Transaction TXID 1 is queued to the Mirror system

6)Transaction TXID 1 is sent over the network to the Mirror system

7)TXID 1 is received by the Mirror system

8)TXID 1 is committed to the Mirror system's journal

9)Transaction complete is reflected to the Master system for TXID 1

FIGURE 4-1 The lifecycle of a transaction in StorEdge File Replicator

Note again that the Sun StorEdge 5310 NAS queues the transaction almost immediately to the Mirror system, the result being that the Mirror stays in very close synchronization with the Master. Sun StorEdge 5310 NAS mirrors on a block level to ensure high performance, but commits data to the mirror on transactional boundaries. This approach guarantees the integrity of the filesystem on the Mirror; at anytime, the Mirror can be promoted and given a network link of suitable quality, the Mirror will be at most a small number of transactions behind the Master. This architecture provides clear advantages over Checkpoint or Snapshot Mirroring, in that the state of the Mirror's data is typically only seconds behind that of the Master's filesystem.

Extensive effort has gone into ensuring not only that the performance of the Mirror does not lag that of the Master, thereby maintaining the Mirror in a high state of synchronization with the Master, but that data integrity - including write ordering - is preserved (Figure 4-2 on page 4-5).

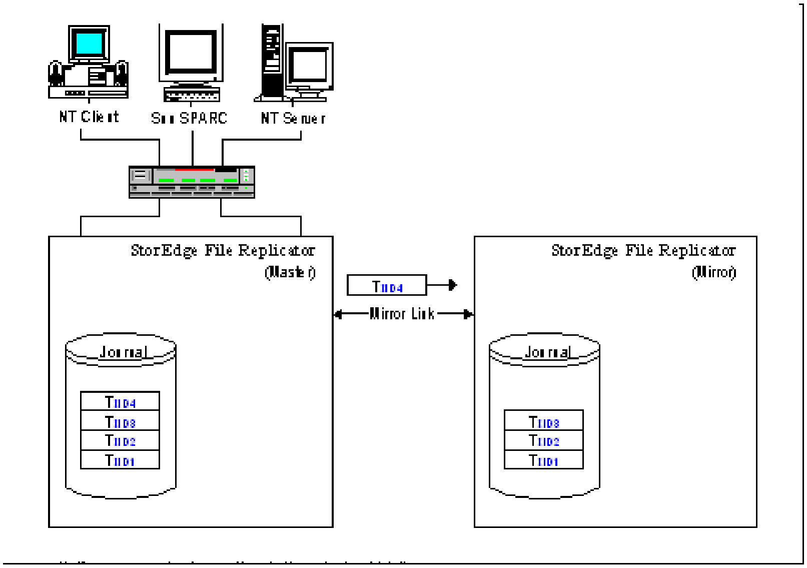

FIGURE 4-2 Write ordering on the Mirror

Figure 4-3 on page 4-6 depicts a scenario where TXID 4 is received prior to TXID 3, either because TXID 3 was lost, during transmission or for some other reason.

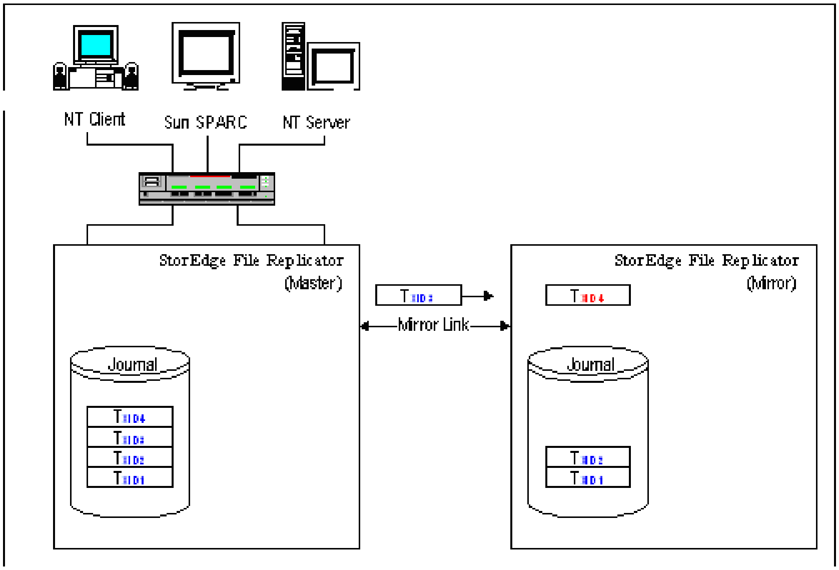

FIGURE 4-3 Lost transaction handling on the Mirror

In the event an out-of-order transaction is received, the Master is notified and it re-sends the missing transaction (Figure 4-3 on page 4-6). The out-of-order transaction - TXID 4 in this example - is not committed until the missing transaction(s) is (are) received and committed (TXID 3). This is crucial as there are many applications (e.g., databases) that require write ordering to be preserved.

Sun StorEdge 5310 NAS provides for handling of network or mirror link issues - e.g., loss of link, poor link quality-of-service etc. - through the use of an expanded journal called the Mirror Log (Figure 4-4 on page 4-7). The Mirror Log tracks write transactions, serving as a repository for those transactions until they can be committed to the Mirror system.

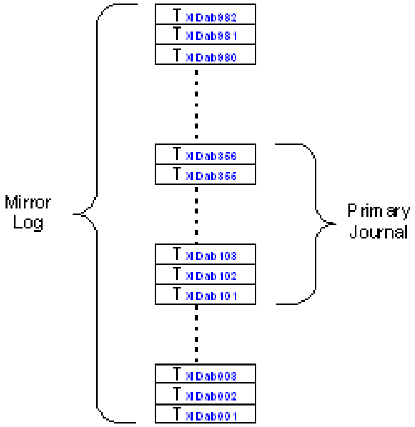

FIGURE 4-4 The Mirror Log and Primary Journal

The Primary Journal floats through the Mirror Log functioning in a similar manner to that in which it does now. Sun StorEdge 5310 NAS maintains hints in the filesystem that enable it to quickly locate its position in both the Primary Journal and Mirror Log if an outage (e.g., a complete power loss, component failure, etc.) is experienced. As with all Sun StorEdge 5310 NAS products, the integrity of the filesystem is guaranteed by the presence and use of the journal, irrespective of the type of outage experienced.

4.1.4 Mirroring Variations

As with many technologies, there are a number of different ways mirroring can be employed. These variations can be employed irrespective of the type of mirroring (Checkpoint/Snapshot, Real-time/Synchronous or Pseudo Real-time/Asynchronous) used.

4.1.4.1 Many-to-One Mirroring

In Many-to-One mirroring, several Master systems are mirrored to a single Mirror target. Many-to-One mirroring is frequently used by customers for DR purposes, to collect exact duplicates of multiple systems at multiple remote locations at a single, central site, where recovery scenarios can be centrally managed. For instance, remote departmental systems in Los Angeles, Houston and San Francisco might be mirrored to a single corporate Datacenter in Kansas City. In the event any one of the remote systems fails or is lost, the data is preserved at the central, corporate site, providing a fallback position for the remote locations and expediting the recovery of the remote site.

Sun StorEdge 5310 NAS supports Many-to-One Mirroring.

4.1.4.2 One-to-Many Mirroring

One-to-Many Mirroring refers to the ability for a Master system to mirror simultaneously to multiple Mirror locations. It may seem a simple variation on the mirroring theme, but it actually introduces a number of complexities to mirroring.

A One-to-Many mirror must be able to cope with different QOS levels on the different connections to remote systems. For example, if a system in Los Angeles is mirroring Many-to-One to systems in San Francisco and Houston and the link between Los Angeles and Houston becomes compromised (increased latencies, packet loss etc.), then the system must decide whether and when to - in the case of Real-time Mirroring - sever the link with the Houston system to preserve the QOS of local users. If the Los Angeles system in the prior example employs Pseudo Real-time Mirroring, then the system must decide how to manage situations where the mirror buffer is overrun for the San Francisco system, but remains intact for the Houston system.

Note that QOS policies must typically be decided on and managed by the user, placing not insignificant burden on them.

One-to-Many Mirroring is frequently confused with Replication and Push Technology.

Sun StorEdge 5310 NAS does not support One-to-Many mirroring.

4.1.4.3 Piggyback or Cascading Mirroring

Piggyback or Cascading Mirroring refers to a more common implementation of One-to-Many Mirroring, whereby a Master system mirrors to a Mirror, which in turn mirrors to another Mirror system (Los Angeles mirrors to Houston, and Houston mirrors to San Francisco).

Sun StorEdge 5310 NAS does not currently support Piggyback or Cascading mirroring.

4.1.4.4 Bi-directional Mirroring

Bi-directional Mirroring refers to the ability for systems at sister locations to mirror to each other. For instance, a system in Los Angeles may be configured to mirror its volumes to a sister system in Houston, which in turn and simultaneously mirrors its volumes to the Los Angeles system. In the event either site experiences a problem, the data is readily available at the sister site. Bi-directional Mirroring is popular because it enables users to deploy systems at partner locations which provide both day-to-day local storage and DR services for their sister location.

Sun StorEdge 5310 NAS supports Bi-directional Mirroring.

4.1.4.5 Replication

Replication is not mirroring, but is closely related to mirroring; mirroring is typically used to effect replication. Replication in the mirroring arena is most frequently associated with data warehousing. In a typical data warehousing replication scenario, data is collected from multiple remote locations, collated at a central site, and subsequently re-distributed to the remote locations. Replication is essentially the initial synchronization of a Master to Mirror sans updates committed to the Master on an ongoing basis.

Sun StorEdge 5310 NAS does not currently support Replication, but is working on this capability.

4.1.4.6 Push Technology

Push Technology is most frequently associated with content delivery. It can be considered a subset of Replication, as it is really selective or object-based Replication. Push Technology would be something like the ability to denote a file or directory as a 'hot' object, with any modification of it resulting in propagation (replication) of the so-denoted object to a set of pre-defined locations or systems.

Sun StorEdge 5310 NAS does not currently support Push Technology, but is working on these capabilities.

4.2 Operational State

StorEdge File Replicator is an active/passive block level journaled mirroring mechanism, similar to RAID 1 disk mirroring except that the devices are connected by a network rather than by a local bus. The network connection media itself is opaque to the mirror service. Connections are made with standard UDP and TCP protocols, so Sun StorEdge 5310 NAS servers can be mirrored across any reachable network.

Since StorEdge File Replicator operates at the disk block level, the mirror system is an exact replica of the master system. However, since mirroring operations are not strictly real time, the mirror system may lag the master by a time delta dependent on the speed and quality of the network. While this network lag may prevent the mirror system from being an exact copy of the master at any given point, the integrity of the mirror system is guaranteed at all times. Only complete file system transactions are mirrored.

In the course of creating a duplicate volume, a mirror goes through three main phases: creation, replication, or sync, and sequencing. StorEdge File Replicator is a fault-tolerant technology. In all of the three main phases, the mirror handles errors with the intent of self-recovery as much as possible. When errors are encountered that are too severe for the mirror to handle on its own, it enters an ERROR state. In this state, user intervention is required to remedy the error and restart the mirror.

4.2.1 Mirror Creation

There are several steps involved in the creation of a mirror. While the mirror is actually only created once, each step is revisited whenever the master system starts up after a reboot. As a general rule, when minor errors are encountered at any of those steps, the mirror enters a RESET state to wait a short time before re-attempting the failed step. More severe errors drive the mirror into the ERROR state, with the mirror status specifying the nature of the error.

When a mirror is started, it enters the NEW state. In this state, a buffer area is allocated within the mirror volume on the master system to store file system transactions while they are transferred to the mirror system. If enough free space cannot be found to accommodate the buffer, the mirror enters the ERROR state. If, as in the case of the reboot of a master system, the mirror buffer already exists, it is simply checked for validity.

Once the mirror buffer is created and/or validated, the mirror enters the INIT state. At this point, the master system attempts to establish a network connection to the mirror system via the NBD. If the mirror system cannot be reached at all, the mirror enters the ERROR state. If the mirror system can be reached, but there are errors encountered in setting up a connection, the mirror enters the RESET state. After a short wait, the mirror re-enters the INIT state and again attempts to establish a connection.

After a connection is set up between the master and mirror systems, the mirror enters the MAKEPARTS state. During this state, a replica of the volume on the master system is created on the mirror system. The volume on the mirror system is constituted with the same structure as that on the master system. File system segments, or extents, are created on the mirror system in numbers and sizes matching those on the master system. In the case of a master reboot, this state consists of simply validating each extent on the mirror system.

The mirror volume is identical to that of the master, except that the partition type is reported as NBD, and all access to the volume outside the mirror service is limited to read-only. The mirror volume will remain unmounted until the entire creation and replication processes have completed.

As in the INIT state, minor errors encountered during the MAKEPARTS state cause the mirror to enter the RESET state, and severe errors drive the mirror into the ERROR state. Some of the more severe errors that might be encountered are a volume on the mirror system with the same name as that on the master system, insufficient disk space on the mirror system, or failure to transfer file system extent information to the mirror system.

When the mirror has been successfully created and/or validated, the mirror enters the READY state. At this point, the master system determines whether to proceed to a replication state or the INSYNC state.

4.2.2 Mirror Replication

Mirror replication is a process that is executed in both the REPLICATE and OUTOFSYNC states; the two states are operationally identical. The REPLICATE state is entered only when the mirror is first starting; either when it has just been created, or when the master system is rebooted while replicating. The OUTOFSYNC state is entered if the mirror cracks while in the sequencing phase. The process of sequencing is explained further below.

During replication, all allocated disk blocks on the volume are copied directly from the master to the mirror system. This process bypasses the file system transaction mechanism normally used in volume read and write operations. Blocks are read from disk and transferred to the mirror system, where they are written straight to their home locations on disk. The mirror buffer is not used by the replication process.

However, the mirror buffer is used by the master file system, during replication, to store file system transactions that need to be transferred to the mirror system. The master volume is live during replication; any changes to the volume must be stored in the buffer until replication completes, at which time those changes are transferred to the mirror system. It is very important that the mirror buffer is large enough to store all transactions while the master is initially syncing itself with the mirror system.

4.2.3 Mirror Sequencing

Once a mirror has been created and fully replicated, it enters the INSYNC state. The mirror volume is finally mounted, read-only, and declared to be of partition type NBD. Then, the mirror begins sequencing. It is at this point that the mirror buffer begins to be actively used. Each file system transaction written to the mirror buffer on the master system is sent, or sequenced, to the mirror system. The master system waits for an acknowledgement from the mirror system that the transaction has been successfully written to its own mirror buffer before discarding it.

The mirror is considered to be "in sync", and will stay in this state, as long as file system changes on the master volume do not outpace the transfer of those changes to the mirror system. In the INSYNC state, the progress percentage represents the percentage of available transaction storage space in the mirror buffer; this percentage may fluctuate depending on disk usage and network latency.

4.2.4 Link Down and Idle Conditions

At any point after the connection between the master and mirror systems has been established in the INIT state, a loss of communication between the two systems will cause the mirror to enter the LINKDOWN state. Data transfer from master to mirror is halted until the connection is re-established.

A mirror can also be manually paused, effecting the same behavior as a LINKDOWN condition. Pausing a mirror puts it into the IDLE state, and is only possible from the INSYNC state. In the IDLE state, transaction sequencing from the master to mirror is suspended.

4.2.5 Cracked and Broken Mirrors

If the rate of file system changes on the master volume exceeds the rate of transfer to the mirror volume, the percentage will drop until it reaches zero. Once there is no more room in the mirror buffer for new transactions, the oldest transactions still pending acknowledgement from the mirror system will start to be overwritten. At this point, the mirror is considered to be "cracked", and will enter the CRACKED state. From the CRACKED state, the mirror reverts to the OUTOFSYNC state, the mirror volume is unmounted, and a full re-sync is started.

It is possible for the mirror to crack both when the mirror is sequencing, in the INSYNC state, and when the mirror is replicating, in the REPLICATE/OUTOFSYNC states. If the mirror buffer is not large enough to store all file system changes made before it finishes syncing to the mirror, it will continue in an endless OUTOFSYNC-CRACKED-OUTOFSYNC cycle.

To stop mirroring on a volume, the user must "break" the mirror. The mirror can be broken from either the master or mirror system. When the mirror is broken, it enters the BREAKING state until the last file system transactions in transit have been acknowledged, and the break request has been communicated to both master and mirror. Data transfer between master and mirror is then stopped, the mirror definition is removed from the mirror service, and the mirror buffer is removed on the master volume. The mirror volume remains mounted as a read-only NBD volume. A broken mirror can be restarted at any time, but must go through the entire creation and replication process from the beginning.

4.2.6 Cannot perform first-time synchronization of mirror system:

File system activity on the master system must be stopped or reduced to a very low level. It is much more filesystem intensive to synchronize the mirror for the first time than it is to maintain it.

4.2.7 Filesystem errors, such as run check, directory broken, etc.:

StorEdge File Replicator is designed to immediately break any active mirrors to a volume with filesystem problems, to avoid replicating these errors. The errors must be corrected via the fsck procedure before reestablishing the mirror.

4.2.8 Error messages, panics or hang condition when enabling mirror:

Ensure that the master and mirror systems are running the same version of the StorEdge operating system.

CHAPTER 5

Clustering

This chapter

5.1 Overview

This section will be updated when the information is available.

CHAPTER 6

Checkpoints/Snapshots

This chapter provides an insight on how checkpoints are created, maintained and deleted.

6.1 Overview

The goal of the checkpoints is to minimize the number of copies when creating a checkpoint. This document discusses what happens to a checkpoint from the time it is created to when it is removed.

Checkpoint lifetime can be divided into three main stages: 1. Creation, 2. Active as a pseudo filesystem, and 3. Deletion. These stages are described in the following sections.

6.1.1 Volumes

In the operating system, every mounted file system is represented by an in-memory data structures called fs_online. The fs_online of a volume is similar to a gate - all accesses to the volume are routed from there. fs_online keeps information about a volume, including the capacity, file-handle of the root directory and status flags. A file-handle is the virtual identifier required for accessing any file system object in system. A file-handle maintains information about the virtual volume, and the corresponding object in that volume. In the case of checkpoint volumes, file-handles also contain information about the checkpoint identifier (or cpid) that contains the actual checkpointed object.

Checkpoints of a volume are accessed through a separate fs_online. This volume corresponds to the virtual checkpoint volume created when checkpoints are made enabled on a volume.

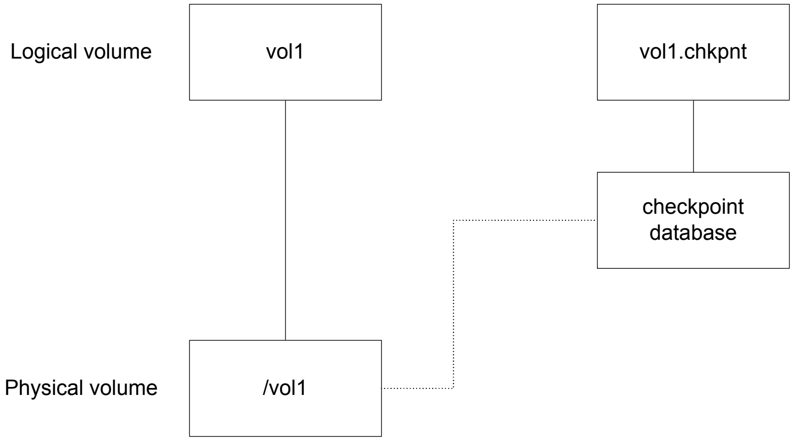

FIGURE 6-1 Physical and Logical Volume Relationship

As shown in Figure 6-1, the existence of checkpoint database distinguishes the checkpoint volume or CFS (Checkpoint File System) from the main volume or LFS (Live File System). The checkpoint database is the data structure that virtually holds different versions of an LFS. It is functioning as a mapping function that maps a virtual block address to its corresponding real address on the live file system.

The checkpoint database is a flat, sparse file with one entry per each block address in the live file system. Each entry is an array of 16 block addresses. When the mapping function tries to resolve a virtual block address, it first locates the corresponding entry in the checkpoint database. In order to do this, it uses the virtual block address as an index to the file. It then uses the cpid (stored in the corresponding file-handle of the object) as index to the entry array to find the proper real block address.

6.1.2 Checkpoint Lifecycle

Checkpoints are created and managed using the fs_chkpntcl( ) function of filesystem which is sfs2_chkpntctl( ) for SFS2 filesystem. The checkpoint management interfaces use this call to create, delete and deactivate checkpoints and checkpointing on sfs2 filesystems.

Checkpoints have three states:

- Active: while checkpoints are active, they can be accessed for most of the read-only file system operations.
- Delete pending: when a checkpoint has expired and automatically removed by the system or users explicitly remove them, they are marked as "delete pending." Later, one of the file system workers (or threads) called checkpoint cleaner will actually remove it. When a checkpoint status is changed from active, it will no longer be accessible and a new checkpoint with the same name can be created.
- Deleting: while the checkpoint cleaner is removing a checkpoint, its state is deleting.

6.1.2.1 Checkpoint Creation

There is a special mode page for checkpointing that is like a directory and contains all of the information about checkpoints on an sfs2 filesystem. The address of this page is in the volume label of the SFS2 filesystem. It is allocated when checkpoints are active on an SFS2 filesystem.

This page contains a table of active checkpoints on filesystem, properties of checkpoints (whether they are visible or not…) and also an array of pointers which are the first of a three-level indirection to pages containing the mappings for the LFS blocks. (The mappings are described in the next section.) There is also a stack of active checkpoints. Each checkpoint that is created will be pushed on top of the stack. When creating a checkpoint, the sfs2cp_dirop( ) function is called by the sfs2_chkpntctl( ) function. This function creates a checkpoint in the checkpoint pseudo-directory.

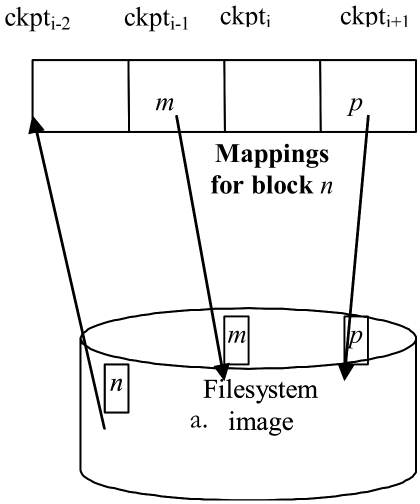

FIGURE 6-2 The Copy-On-Write Mechanism for Checkpoints

6.1.2.2 Active Checkpoint

Each CFS has a mapping function for all of the blocks on LFS. This mapping function returns a value for each page of LFS. Currently, there can be no more than 16 active checkpoints on a filesystem. This means the checkpoint filesystem can hold at most 16 mappings for each block. Also not all of the blocks in the LFS are mapped. Block allocation table and journaling area are not mapped and checkpointed.

Active checkpoints have a stack-like structure and each checkpoint uses the mapping of itself. If there is no mapping, it will use the mapping for the next checkpoint created after it (this selection of proper mapping is handled by the mapping function). If none of the entries have a mapping, this implies the file system block has never been modified since the creation of oldest checkpoint - in this case the block address will be mapped to itself.

When a new checkpoint is created, a new directory entry is added to the list of checkpoints in checkpoint control page. Also, an available slot in the checkpoint stack is assigned to this newly-created checkpoint. After doing this, the new checkpoint becomes the owner of the corresponding entry in mapping entries for all blocks of the file system. The entire operation is performed in a single transaction and is instantaneous. As soon as a new checkpoint is created, the checkpoint entry becomes available and can be used for all read-only file system operations. There is no partial state visible from other file system users.

Figure 6-2 shows part of array of mappings for block n. In this example, ckpti+1 is the most recent checkpoint and thus, the last entry in the mappings and ckpti-2 is the oldest checkpoint (in this example).

Assume a filesystem operation on checkpoint ckpti-2, and block n is accessed. The mapping function first checks the mapping for ckpti-2 and given it is empty, it moves forward and checks the mapping for ckpti-1. It finds m and will use block m instead of block n.

When accessing block n from ckpti, because there is no mapping, the mapping function moves forward and will use the mapping for ckpti+1, which is block p. Note that when searching for a mapping entry, the system will always move forward from current checkpoint toward more recently created checkpoints. The mapping entries of checkpoints older than the current checkpoint are never used.

Now let's see how the mappings are created. While checkpointing is active for a volume, all of the blocks of LFS are shared by checkpoint of its corresponding CFS until they are changed. As show in Figure 6-3, while block n is not modified, all the checkpoints of the filesystem use the block n itself with no mapping. If a block in the live file system is modified, then sfs2cp_notice_update( ) is called. This function first checks to see if there is an existing mapping for this block. If there is a mapping for the block, it will do nothing. If no mappings can be found and the change is not allocation, then it duplicates the block and puts the address of the new block into the mapping for the most recently created checkpoint. If the change is for block allocation, the system does not duplicate the block; it just puts a special value (SFS2CP_ALLO_MARKER) into the mapping. This is because if the block has just been allocated, there can't be any object on the checkpoints using this block (and no need to duplicate this block).

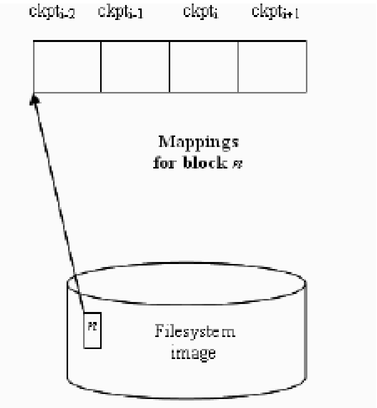

FIGURE 6-3 Mappings for Block n Before Modification

In Figure 6-4, there are two active checkpoints ckpti-2 and ckpti-1. If block n is going to be modified, a new block m will be allocated and the old content of n is copied to it. The address for block m will be inserted into the list of mappings for the block n in the latest checkpoint, which is ckpti-1. From this point, any request for block n in ckpti-2 and ckpti-1 will be redirected to the block m.

From this, we can tell from a mapping table when system blocks are modified relative to active checkpoints.

Referring to Figure 6-2, ckpti-2 is created and then without modifying block n - ckpti-1 is created. Therefore ckpti-2 and ckpti-1 represent the same content for block n. Before the creation of ckpti, block n had been modified. A copy of the before change content is made and put into the mapping for the most recent checkpoint which at that time was ckpti-1. Because both ckpti-1 and ckpti-2 expect to see the same content for block n, the mapping from n to m is shared by both of them.

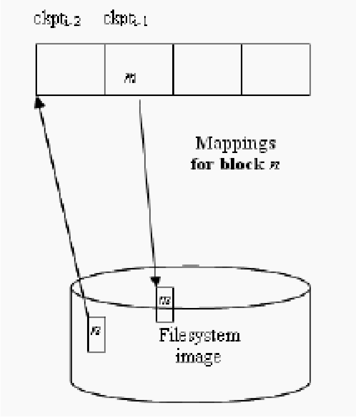

FIGURE 6-4 Mappings for Block n After Modification

In Figure 6-2, Ckpti is created and without modifying block n, ckpti+1 is created and then block n is modified again. This causes the creation of block p and mapping from n to p for checkpoint ckpti and ckpti+1.

When deleting a block in LFS, the system first checks if there is a mapping for the block in the most recent checkpoint or not (considering SFS2CP_ALLO_MARKER a valid mapping). If there is a mapping, it frees the block. Otherwise, it puts the address of block into the mapping for it, creating a one to one mapping for that block. This is called page stealing of CFS from LFS.

6.1.2.3 Translation of File System Objects in Checkpoints

As previously discussed, because the checkpointing mechanism is applied to filesystem blocks rather than filesystem objects, there is no special consideration for the type of object that is checkpointed.

Hardlinks and symlinks in checkpoints will continue to have the same semantics they had in the live filesystem at the time the checkpoint is created. To keep the hardlinks unchanged, the directory entries referring to the same inode and the inode itself should have the unchanged values. Also, checkpoints in all the disk blocks continue to have their old values regardless of whether they are meta-data or data block - all the hardlinks remain unchanged. Symlinks are merely data files that are (probably) references to other objects - so the same applies to them.

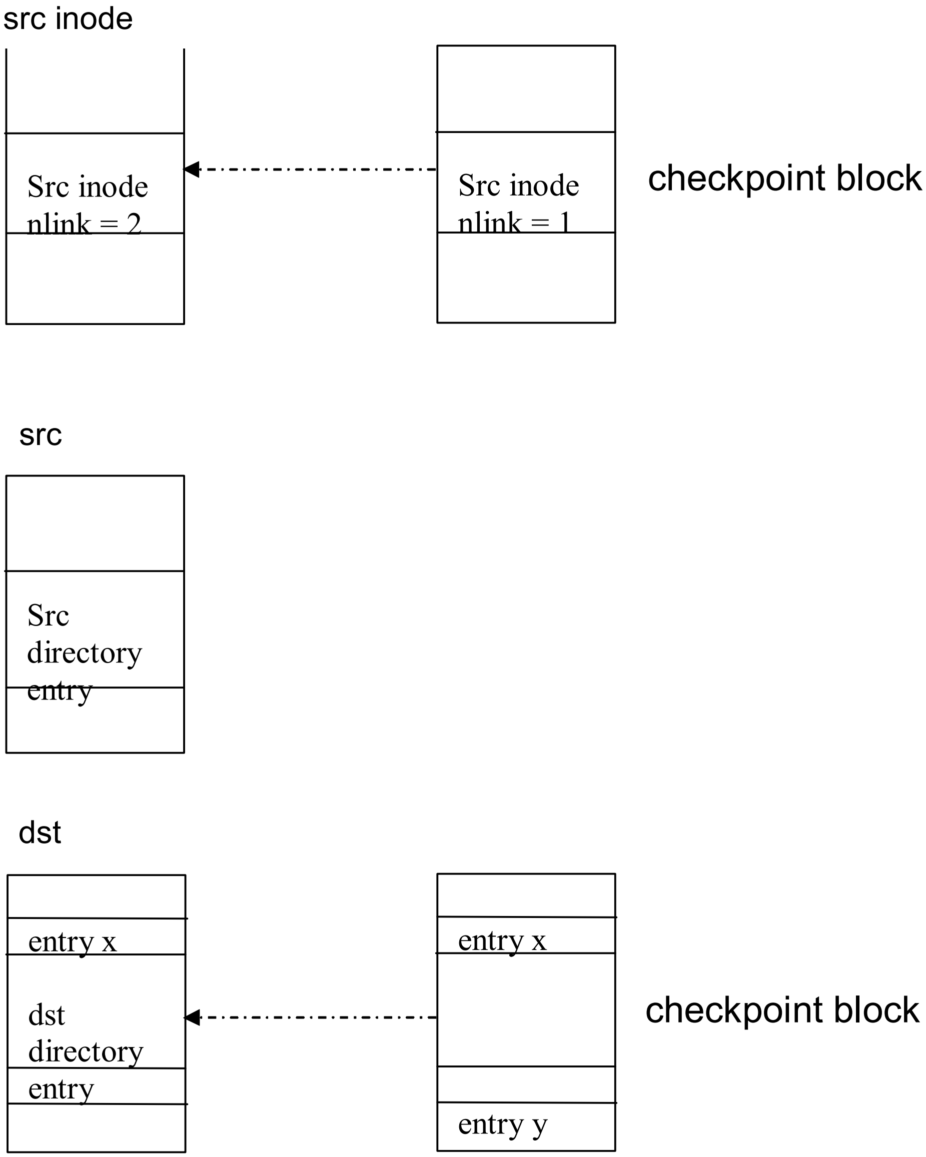

FIGURE 6-5 Creating a hardlink when a volume is checkpointed and has active checkpoints

Another example is when a hard-link is created in a directory named dst to a file residing in a directory called src. First, a new directory is inserted in dst. However, since volume is checkpointed, a copy of corresponding disk block is made first, and then the disk block is modified. When accessing the checkpoint version (and thus accessing the copied block), the old directory content (without the new entry) is seen. Next, the target inode is modified so the link count reflects the new directory entry (pointing to it by incrementing it). This will cause the corresponding disk block to inode data be modified and the copy containing old content will be assigned to checkpoint. Again, for someone accessing the checkpoints, the inode seems to be unchanged and everything remains perfectly consistent both in checkpoint and live file systems.

6.1.2.4 Checkpoint Deletion

When a checkpoint is to be deleted, a flag is set in the entry for that checkpoint in the checkpoint control page. Later the cleaner thread for checkpoint filesystem catches the flag and scans the entire mapping for each block in filesystem and if there is a mapping entry for the checkpoint and if the previous checkpoint has no mapping for the block, then the entry is moved to the entry for the previous checkpoint. Otherwise if the mapping is not SFS2CP_ALLO_MARKER, the block is freed. Then all of the entries for checkpoint after the one that is to be deleted are moved to the left.

For example in Figure 6-2, when deleting ckpti-1, m is moved to the entry for ckpti-2 and all of the entries after ckpti-1 are moved to the left (this operation is done logically in the checkpoint stack, so the actual entries are not moved, only block ownership will be delegated). So ckpti replaces ckpti-1. After changing the mappings, checkpoint stack is also updated similar to what has been done for checkpoint mapping entries. So after deleting ckpti-1, the mapping entries for block n will be as in Figure 6-6.

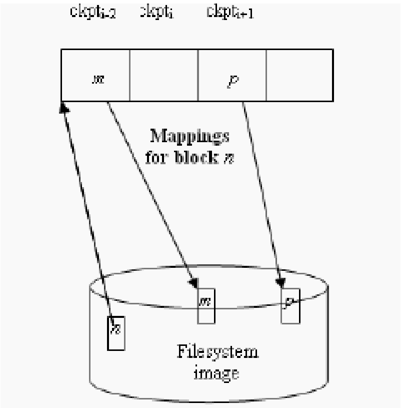

FIGURE 6-6 Mappings for Block n After Deleting ckpti-1

Another example assumes checkpoint ckpti+1 is to be removed. Because ckpti has no mapping for itself, p is copied to the entry for ckpti. Because there is no entry after ckpti, nothing else need be done. The result is depicted in Figure 6-7.

Yet, another example where deleting ckpti-2 is exactly like deleting ckpti-1 with the difference that in Figure 6-6, ckpti-2 is replaced with ckpti-1.

Checkpoint are deleted using sfs2cp_dirop( ) function. This function just marks the checkpoint for deletion and the cleaner thread does rest of the work (described above). When a checkpoint is deleted, it will be removed from stack of active checkpoints in the mode page.

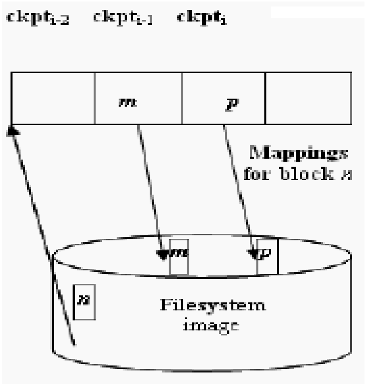

FIGURE 6-7 After Deleting ckpti+1.

6.1.2.5 Checkpoint Scheduling

Checkpoints can be created in two ways: automatic and manual. If the user selects the automatic checkpoints, checkpoints are created and removed based on the scheduling that user specifies for the checkpoints. This scheduling is enforced by a checkpoint manager thread. On the other hand, checkpoint manager does not control manually created checkpoints (users can create manual checkpoints that will be removed automatically by using the same naming convention that system uses for automatic checkpoints).

Also there is a feature (currently disabled) that allows the root user of an NFS client to create manual checkpoints. If the root user creates a directory in the root directory of a CFS, it will be interpreted as a request for creation of a new checkpoint and a checkpoint with specified name will be created.

6.1.2.6 Checkpoint management commands

In order to manage checkpoints, the following commands are provided in StorEdge Command Line interface (CLI):

chkpnt : A general command for changing different checkpointing options

chkpntls: Shows the status of checkpoints along with a list of exiting checkpoints

chkpntmk: Creates a new checkpoint

chkpntmv: Renames an existing checkpoint

chkpntoff: Disables checkpoints on a volume. This will remove all the checkpoints and frees up all the allocated resources

chkpnton: Enables checkpoint on a volume. Does not create any checkpoints

chkpntrm: Removes a checkpoint.

6.1.2.7 Local Directory Checkpoint Access

Checkpoints are read-only point-in-time images of a volume. They can either create manually or schedule to be create and remove by the system without user intervention. Checkpoints are accessible in two ways:

- Via the standard mechanism of mapping or mounting the virtual volume
- Via the local directory checkpoint access feature

The local directory checkpoint access feature allows the users access any version of a directory by simply using normal client commands. This white paper describes the access of checkpoints using the local directory checkpoint access feature and provides usage examples.

Accessing Checkpointed Data

Access to the checkpointed versions of directories and files is achieved through the provision of a hidden, virtual directory - named .chkpnt - within each live directory. Changing the current directory to the virtual .chkpnt directory enables users to access checkpointed or prior versions of filesystem objects. Note the .chkpnt directories are hidden to prevent problems with applications that search through directory hierarchies, e.g., backup or virus scanning applications. Users can access objects in the .chkpnt directories by explicitly navigating to them.

For example:

"cd /live_directory/.chkpnt"

Will navigate the user to the checkpointed version of the directory live_directory, where they'll be able to view prior versions of objects from live_directory. Users can also reference objects explicitly.

For example, the command:

cp /live_directory/.chkpnt/cp1/old_file1.txt /live_directory

Would copy the cp1 version of old_file1 back to live_directory.

Note: .chkpnt is the default name for the hidden directories containing the checkpointed data and can be changed. See Compatibility Issues below.

If a directory is removed and a user wishes to access a checkpointed version of that directory, they can do so by navigating first to the parent directory (of the removed directory) and then traversing the directory hierarchy until they reach the desired version of the removed directory. In Figure 6-8, the directory d3 has been removed but the user can access a checkpointed version of d3 through the .chkpnt directory in its parent directory, d2.

In the example inFigure 6-8, cp1 is the name of a checkpoint that contains the desired version of the d2 directory.

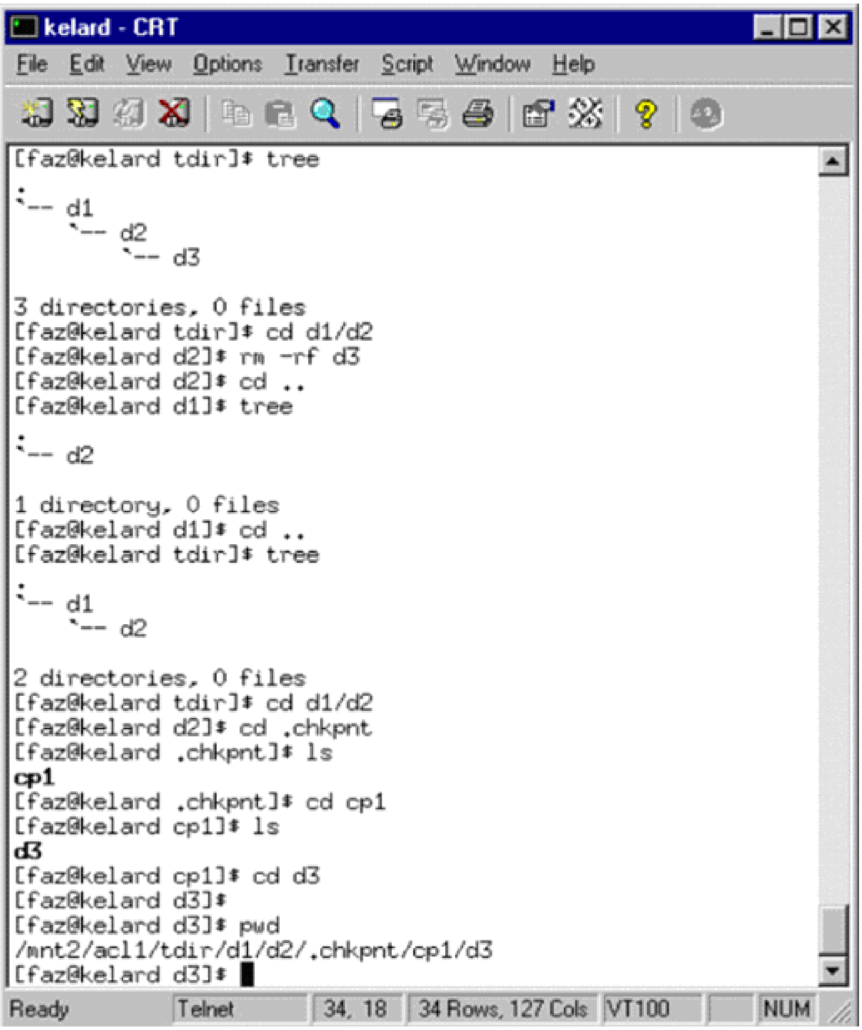

```
[faz@kelard tdir]$ tree
.
`-- d1
    `-- d2
        `-- d3

3 directories, 0 files
[faz@kelard tdir]$ cd d1/d2
[faz@kelard d2]$ rm -rf d3
[faz@kelard d2]$ cd ..
[faz@kelard d1]$ tree
.
`-- d2

1 directory, 0 files
[faz@kelard d1]$ cd ..
[faz@kelard tdir]$ tree
.
`-- d1
    `-- d2

2 directories, 0 files
[faz@kelard tdir]$ cd d1/d2
[faz@kelard d2]$ cd .chkpnt
[faz@kelard .chkpnt]$ ls
cp1
[faz@kelard .chkpnt]$ cd cp1
[faz@kelard cp1]$ ls
d3
[faz@kelard cp1]$ cd d3
[faz@kelard d3]$
[faz@kelard d3]$ pwd
/mnt2/acl1/tdir/d1/d2/.chkpnt/cp1/d3
[faz@kelard d3]$
```

FIGURE 6-8 Accessing .chkpnt in UNIX

In order to access the ".chkpnt" directory, UNIX clients use standard file system commands (as shown in Figure 6-8). The operation from a Windows client is similarly simple; Windows clients can either provide the complete explicit path name or use the "Go to" option in Windows Explorer. This option can be found in Tools menu (see Figure 6-9 and Figure 6-10).

Once the user has navigated to the .chkpnt directory at any point in the directory hierarchy, there are no .chkpnt subdirectories, i.e., there are no nested .chkpnt directories. Users can then access checkpointed versions of the entire volume by navigating to the .chkpnt directory at the root of the volume.

Compatibility Issues

In order to avoid name conflicts between user applications and existing files, each .chkpnt entry possesses a number of special characteristics. As previously noted, it will not show up in directory listings. Applications not aware of the .chkpnt directories will not be able to "see" the directories - thereby preventing potential problems with applications that traverse directory hierarchies like backup and virus scanning programs.

Second, if users already have existing files or directories named .chkpnt, these objects will retain their existing meaning. Users with these pre-existing objects have one of two options:

a. they can rename the existing objects

or

b. the StorEdge OS provides a mechanism as part of the chkpnt command structure to enable users so affected to change the access name for the virtual checkpoint directories.

The syntax for the command is:

```
chkpnt vname {volname}
```

```
shows the current hidden directory name for {volname}
```

```
chkpnt vname {volname} [new-name]
```

```
sets the hidden directory name to new-name
```

Once the virtual checkpoint directory name is set with the chkpnt vname command, StorEdge will disallow the creation of objects with the same name as that specified by the user for the virtual checkpoint directories.

Accessing .chkpnt Using Windows Explorer

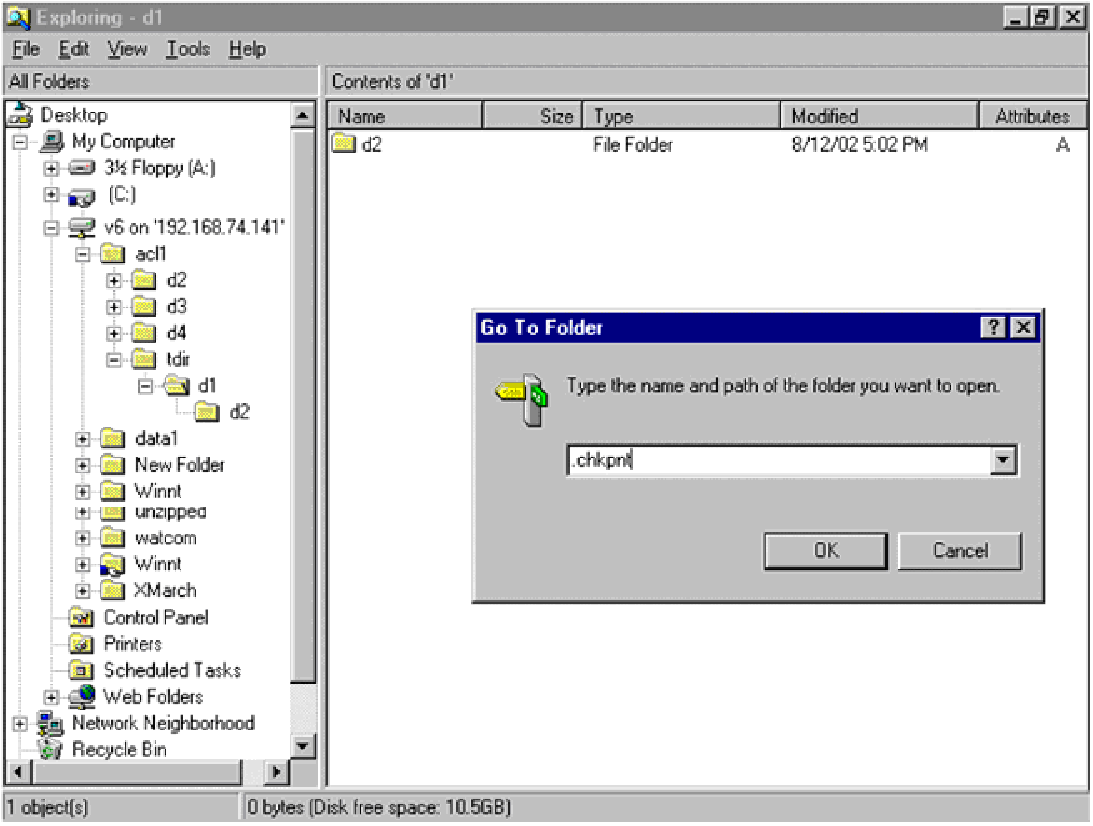

Contents of 'd1'

Name	Size	Type	Modified	Attributes
d2		File Folder	8/12/02 5:02 PM	A

Type the name and path of the folder you want to open.

.chkpnt

1 object(s)

0 bytes (Disk free space: 10.5GB)

FIGURE 6-9 Accessing ".chkpnt" in Windows Explorer

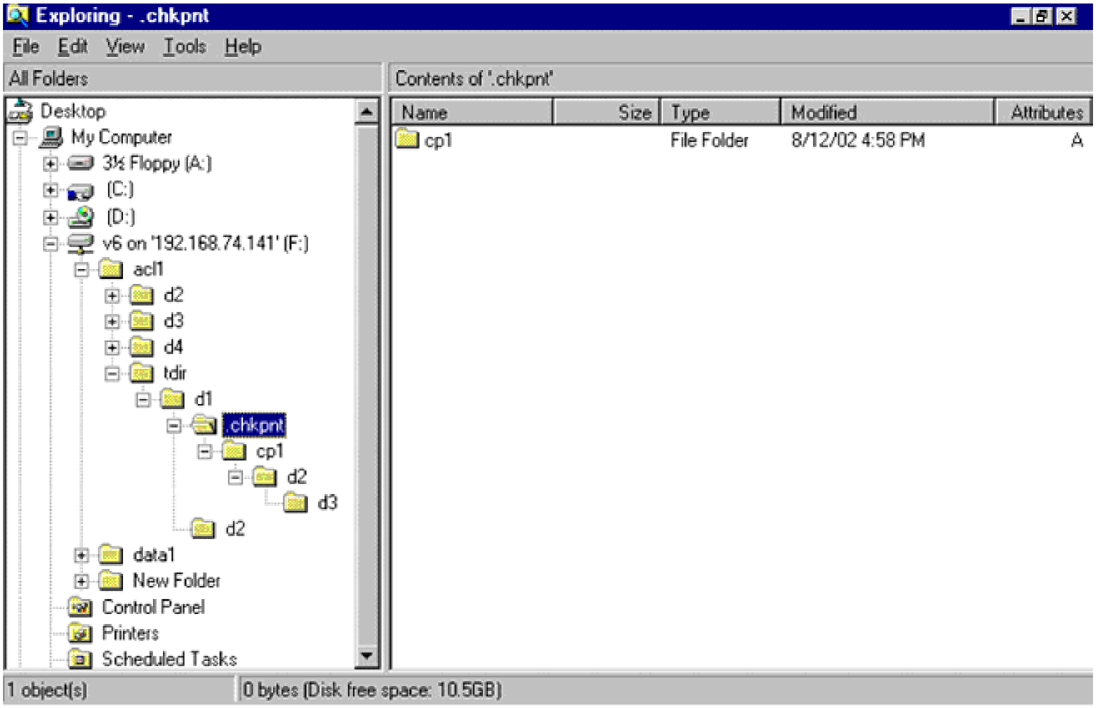

Contents of '.chkpnt'

Name	Size	Type	Modified	Attributes
cp1		File Folder	8/12/02 4:58 PM	A

1 object(s)

0 bytes (Disk free space: 10.5GB)

FIGURE 6-10 Viewing ".chkpnt" in Windows Explorer

6.1.3 Object Checkpoint Restore

Checkpoints are read-only point-in-time images of a volume. They can be created manually or scheduled to be created (and subsequently removed) by the system without user intervention. By definition, objects within the checkpointed volume cannot be modified; this is the very nature of checkpointing and desirable behavior.

Should a user wish to modify a checkpointed version of an object, they must first return an instance of it to the live filesystem. Previously, the only way to accomplish this was to employ standard client-based copy mechanisms, e.g., drag-and-drop on Windows Explorer, the cp command on Unix systems. This method is inefficient for a number of reasons:

- The data is copied twice - from StorEdge to the client and from the client back to StorEdge.
- Precious network bandwidth is used.
- Client resources (CPU, network, memory) are used to effect the copy operation and not available for other purposes.

- Overhead is imposed at each layer of the operation
- The copy operation is a block-for-block operation because the client system is not cognizant of the structure of the StorEdge filesystem, and therefore of the underlying relationship between the blocks in the live and checkpointed versions of the filesystem.

The new StorEdge internal cp command was engineered to confine the operation to the filer, resulting in far greater efficiency and speed.

6.1.4 StorEdge cp Command

The StorEdge Command Line Interface (CLI) cp command effects the copy operation within the filer. As an ancillary benefit, the command is cognizant of the filesystem architecture and it can copy only those blocks that are not referenced in the live filesystem, using StorEdge's copy-on-write implementation.

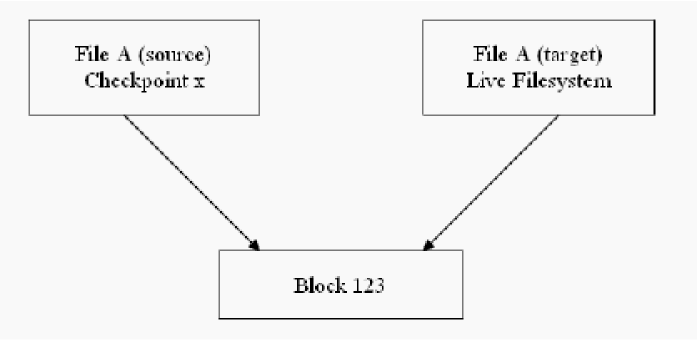

FIGURE 6-11 Sharing Blocks Between Live and Checkpoint File Systems

For example, in Figure 6-11 after a cp restore operation, two versions of a File A - the Live Filesystem version and Checkpoint x's version - share the same block: 123. In this case when Block 123 is referenced by either version of File A, the same block will be accessed. Should a client modify Block 123 of File A in the live filesystem, StorEdge's checkpointing mechanism will ensure that the original Block 123 is preserved in checkpointed versions of the filesystem.

The syntax of the StorEdge cp command is:

```
cp [-c] source destination
```

This command simply copies the source to destination. The source should be a regular file. If the destination is a regular file, it will be overwritten by source. Otherwise, if it is a directory, the source will be copied to that directory. When using the "-c" option, the checkpointed file source will be restored to the destination. If the destination is omitted then system will try to restore to the original (non-checkpointed) path.

For example:

```
cp -c /v6.chkpnt/ckpt1/docs/sample.doc
```

will restore to:

```
/v6/docs/sample.doc
```

The cp command is also available as part of the chkpnt command line operation. The syntax is identical except that the command is prefaced with chkpnt, and the "-c" option is not required:

```
chkpnt cp source [destination]
```

The "-c" option is not required when cp is part of the chkpnt command because the chkpnt prefix conveys the operation's context: to restore a checkpointed object to the live filesystem.

While a checkpoint file restore operation is underway, the file is locked for the duration of the operation and clients can perform no other action against that file. Any attempt to access the file during the restore operation will result in an error. For example, when trying to access the file from a UNIX client, users may see the following:

```
$ cat sample.doc > /dev/null

cat: sample.doc: Input/Output error
```

Similarly, a Windows NT user may see one of the errors shown in Figure 6-12 or Figure 6-13 when attempting to access a file when a checkpoint restore operation is underway against that same file.

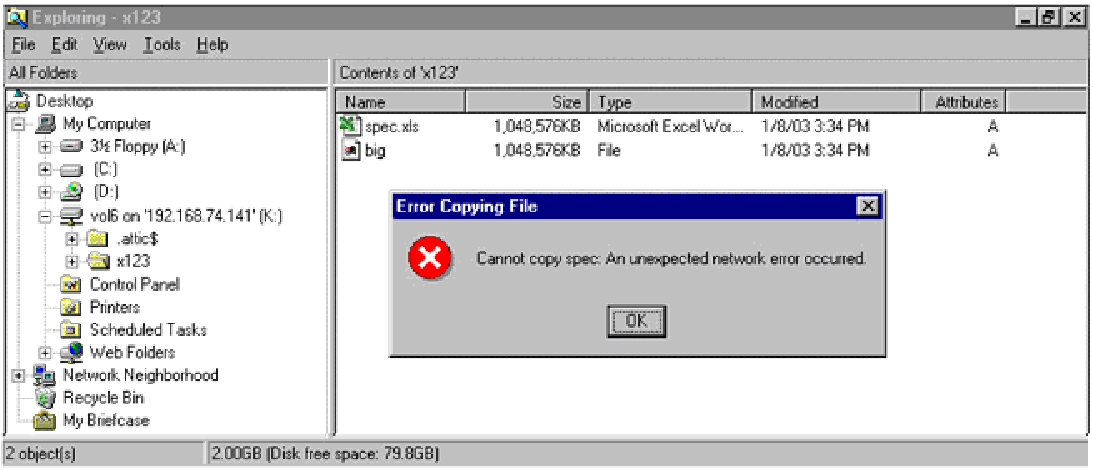

Name	Size	Type	Modified	Attributes
spec.xls	1,048,576KB	Microsoft Excel Wor...	1/8/03 3:34 PM	A
big	1,048,576KB	File	1/8/03 3:34 PM	A

Error Copying File

Cannot copy spec: An unexpected network error occurred.

2 object(s) | 2.00GB (Disk free space: 79.8GB)

FIGURE 6-12 Windows File Copy Error Message During a Checkpoint Restore Operation

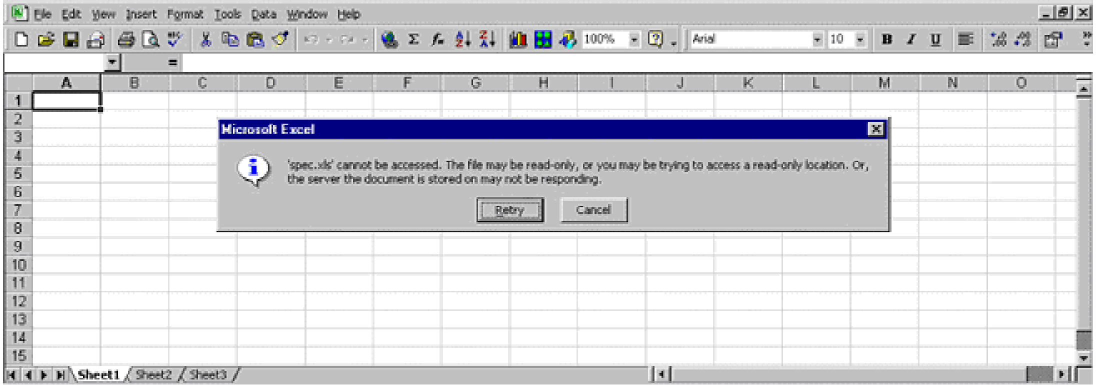
File Edit View Insert Format Tools Data Window Help
100%
Arial
10
A B C D E F G H I J K L M N O
1 2 3 4 5 6 7 8 9 10 11 12 13 14 15
Sheet1 Sheet2 Sheet3
Retry
Cancel

Microsoft Excel

'spec.xls' cannot be accessed. The file may be read-only, or you may be trying to access a read-only location. Or, the server the document is stored on may not be responding.

FIGURE 6-13 Windows Excel Open Error Message During a Checkpoint Restore Operation

The CLI copy (cp) command is a general command and can be used to copy files on StorEdge whether or not the intention is to restore a checkpointed version of an object to the live filesystem. For instance, the cp command can be used to copy a file from one StorEdge volume to another volume on the same filer. This behavior will occur only when the user requests restoration of an object from a checkpoint to the corresponding live filesystem, either by specifying the -c option on the cp command or by using the chkpnt cp command. Otherwise, a general copy operation is performed.

Finally, should a checkpoint restore operation be interrupted for any reason, e.g., a power failure - the operation will be resumed automatically on restart. Other checkpoint-related operations are also suspended for the duration of the checkpoint restore.

CHAPTER **7**

FRU/CRU Replacement Procedures

This chapter describes how to replace components in the Sun StorEdge 5310 NAS after they have been set up. It contains the following sections:

- "Tools and Supplies Needed" on page 7-1
- "Determining a Faulty Component" on page 7-2
- "Safety: Before You Remove the Cover" on page 7-2
- "Removing and Replacing the Cover" on page 7-2
- "Field Replaceable Unit (FRU) Procedures" on page 7-4

Note – The procedures in this chapter for servicing field replaceable faulty components are for the attention of qualified service engineers only. If a Field Replaceable Unit (FRU) needs replacement, contact your local Sun Sales representative who will put you in contact with the Sun Enterprise Service branch for your area. You can arrange to return the system to Sun for repair under the terms of your warranty. Or, if under a Sun Service agreement, the FRU will be replaced by a Sun Service engineer.

Note – When working on a server, you may want to turn on the blue System ID LED to identify the server that is being worked on. See "LEDs" on page 2-11 for instructions on how to turn on this LED.

7.1 Tools and Supplies Needed

All that is needed is an antistatic wrist strap (recommended).

7.2 Determining a Faulty Component

To determine and isolate a faulty component, refer to "Troubleshooting the Server Using Built-In Tools" on page 2-10." This section can help you isolate a faulty component using the following methods:

- Fault and Status LEDs (see "Front Panel LEDs" on page 2-13)
- POST LEDs, beep codes, and displayed error messages (see "Diagnosing System Errors" on page 2-10)

7.3 Safety: Before You Remove the Cover

Before removing the system cover to work inside the server, observe these safety guidelines:

1. **Turn off all peripheral devices connected to the system.**
2. **Turn off the system by pressing the power button on the front of the system. Then unplug the AC power cord from the system or wall outlet.**
3. **Label and disconnect all peripheral cables and all telecommunication lines connected to I/O connectors or ports on the back of the system.**
4. **Before handling components, attach a wrist strap to a chassis ground of the system (any unpainted metal surface).**

7.4 Removing and Replacing the Cover

Many of the equipment replacement procedures require that you remove the chassis cover. Before you remove the cover, observe the safety instructions in the section titled "Safety: Before You Remove the Cover" on page 7-2

To remove the cover, follow these steps:

1. **While pressing the blue latch button (A) with your left thumb, push down on the top cover and slide it back using the heel of your right hand on the blue pad (see Figure 7-1).**

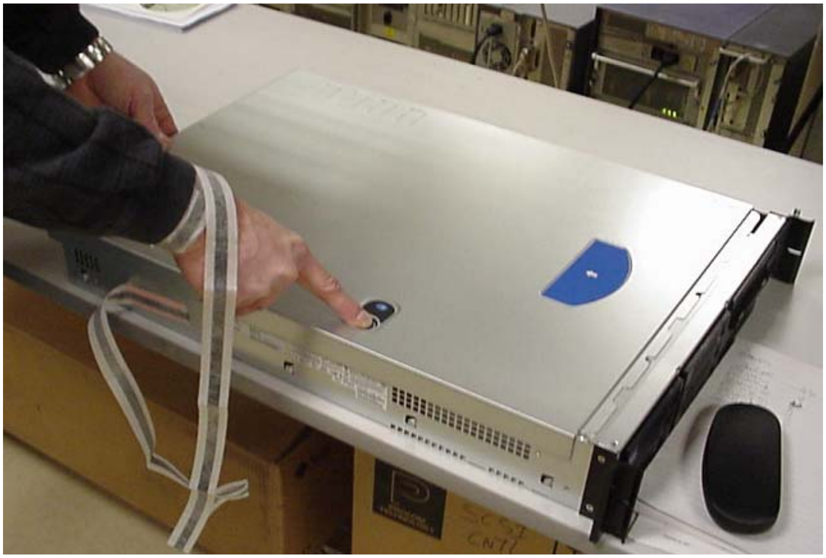

FIGURE 7-1 Removing the Cover

2. **Set the cover aside and away from the immediate work area.**

Note – A non-skid surface or a stop behind the chassis may be needed if attempting to remove the top cover on a flat surface. Sliding the server chassis on a wooden surface may mar the surface (there are no rubber feet on the bottom of the chassis).

7.5 Field Replaceable Unit (FRU) Procedures

This section explains how to replace the FRUs in the Sun StorEdge 5310 NAS:

- Opening the Front Bezel
- Memory
- Power Supply Unit
- Hard Disk Drives
- Fan Module
- High Profile Riser PCI Card
- Gigabit Ethernet Card
- TBBU Transportable Battery Backup Unit
- LCD Display Module
- Flash Disk Module
- System FRU - Super FRU
- Sun StorEdge 5310 NAS Expansion Unit base chassis
- Environmental Monitor Unit (EMU)
- IO Module
- Power Supply/Fan Module
- SCSI Cables (1 or 1.5 feet)

7.6 NAS Head FRU Replacement Procedures

7.6.1 Opening the Front Bezel

To access the system controls and peripherals when a front bezel is installed, grasp the bezel at the finger hole on the left side and gently pull it towards you, unhinging it at the right, until it unsnaps from the chassis. Replace the bezel using the reverse process (see Figure 7-2).

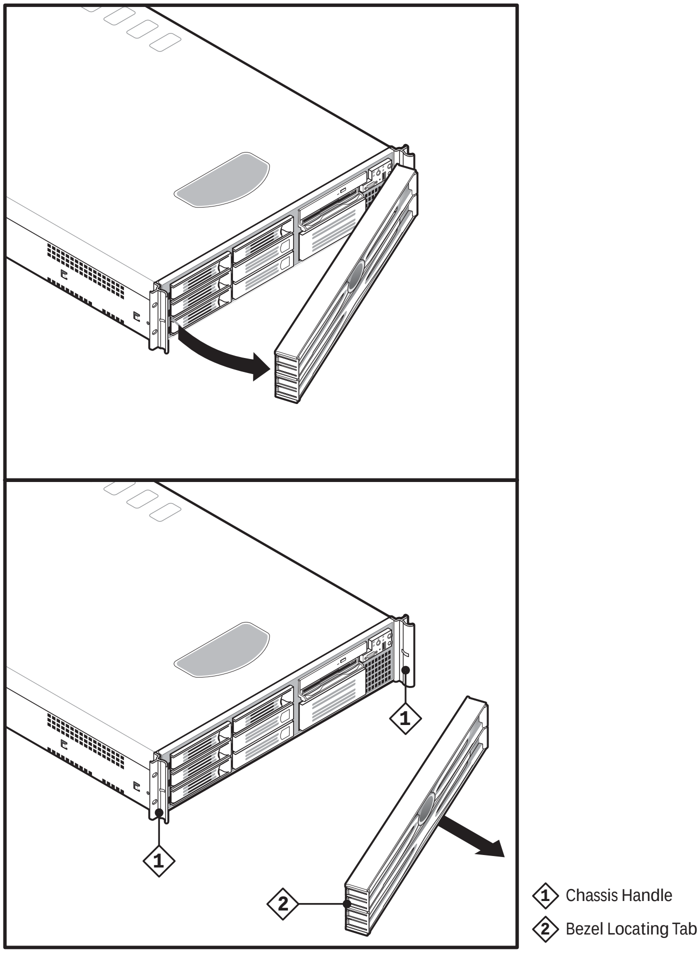

FIGURE 7-2 Sun StorEdge 5310 NAS Bezel Replacement

The EU front bezel uses a key system to lock the bezel. To remove the front bezel:

1. **Unlock both sides of the bezel.**
2. **Gently pull the bezel forward. The bottom section of the bezel is connected via hinges.**

3. **Rotate the bezel towards the front.**
4. **Remove the bezel by pressing the hinges in and pulling it out.**

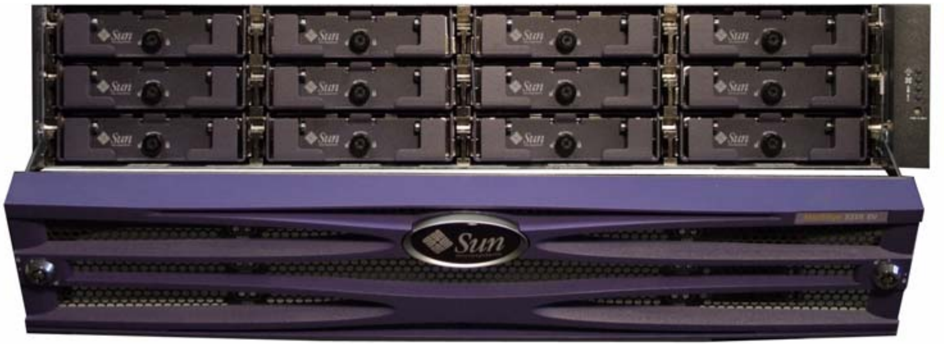

FIGURE 7-3 Sun StorEdge 5310 NAS Expansion Unit

7.6.2 Memory

Caution – Before touching or replacing any component inside the server, disconnect all external cables and follow the instructions in "Safety: Before You Remove the Cover" on page 7-2 and "Removing and Replacing the Cover" on page 7-2. Always place the server on a grounded ESD pad and wear a properly grounded antistatic wrist strap.

The main board supports DDR-266 compliant registered ECC DIMMs operating at 266 MHz.

Each system contains four registered one gig DIMMS.

ECC single-bit errors are corrected and multiple-bit errors are detected.

- Single bit error correction: If a single bit error is detected, the ECC logic generates a new "recovered" 64 bit QWord with a pattern that corresponds to the originally received 8 bit ECC parity code. The corrected data is returned to the requestor (the processor or PCI master).
- Multiple-bit error detection: Additional errors within the same QWord constitute a multiple-bit error, which is unrecoverable. When a multiple-bit memory error is detected, a non-maskable interrupt (NMI) is issued that instructs the system to shut down to avoid data corruption. Multiple-bit errors are very rare.

Note – NMI is not currently supported.

- Memory scrubbing: Error correction is performed on data being read from memory. The correction is then passed to the requestor and at the same time the error is “scrubbed” (corrected) in main memory. Memory scrubbing prevents the accumulation of single-bit errors in main memory that would then become unrecoverable multiple-bit errors.
- X4 single device data correction (x4 SDDC): When x4 memory is installed, the ECC function can detect and correct a four-bit error caused by a single failed memory chip and the system continues to function, though system performance will be affected. When x8 memory is installed, the ECC function will detect an eight-bit error caused by a single failed memory chip but will not be able to correct the error. In this situation a fatal error will be issued.

Caution – Use of unqualified DIMM modules may damage the server and may void the warranty.

7.6.3 Power Supply Unit

Caution – Before touching or replacing any component inside the server, disconnect all external cables and follow the instructions in “Safety: Before You Remove the Cover” on page 7-2 and “Removing and Replacing the Cover” on page 7-2. Always place the server on a grounded ESD pad and wear a properly grounded antistatic wrist strap.

Warning – Power supply may be hot to the touch.

7.6.3.1 Sun StorEdge 5310 NAS Power Supply

To replace a power supply:

1. **Squeeze the module handle to depress the latch (Figure 7-4, panel 1).**
2. **Rotate the handle down while pulling the module toward you (Figure 7-4, panel 2).**
3. **As you pull the module out, support the module with your free hand.**
4. **Insert a new power supply module in the bay.**
5. **Grip the module handle, rotate it down, and push the module into the bay.**

6. **When the module is nearly all of the way in, the handle will rotate up. At this time, push firmly on the front of the handle to lock the latch.**

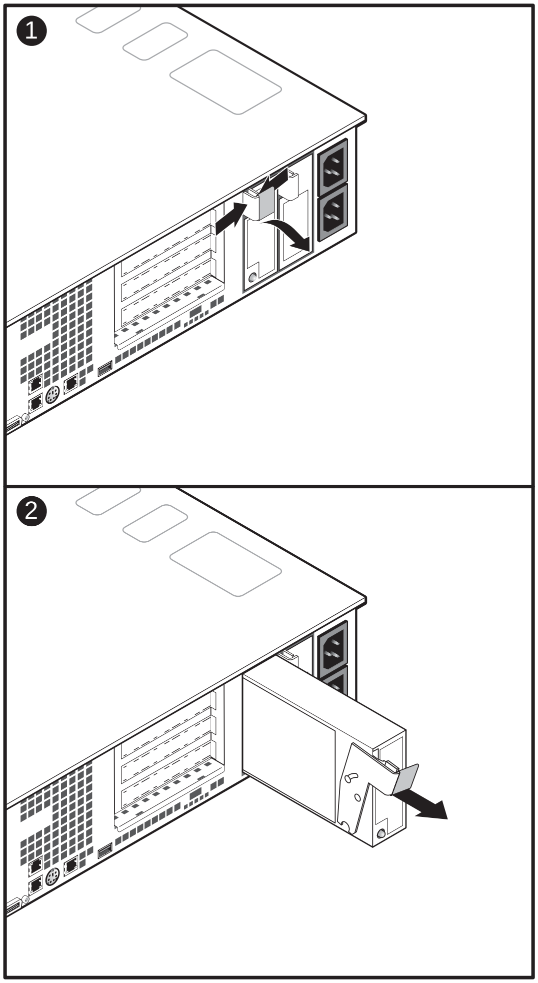

FIGURE 7-4 Replacing the Power Supply

7.6.4 Fan Module

Caution – Before touching or replacing any component inside the Sun StorEdge 5310 NAS, disconnect all external cables and follow the instructions in "Safety: Before You Remove the Cover" on page 7-2 and "Removing and Replacing the Cover" on page 7-2. Always place the server on a grounded ESD pad and wear a properly grounded antistatic wrist strap.

7.6.4.1 Sun StorEdge 5310 NAS Fan Module Removal

The fans in the Sun StorEdge 5310 NAS are individually replaceable. To replace an individual fan, first remove the fan module according to the instructions below while referring to Figure 7-5.

1. **Remove the full-height PCI riser board.**
2. **Unthread the memory fan power cable from the retaining hooks on the plastic processor air duct.**
3. **Push the air duct slightly toward the back of the chassis, then lift it by its front edge and remove it from the chassis.**
4. **Remove the flex circuit cable retention clip.**
5. **Disconnect the flex circuit cable from the backplane.**
6. **Unthread and remove the USB cable from the clips on top of the fan module.**
7. **Unplug the fan cables from the server board system fan connectors.**
8. **At the end of the fan module closest to the chassis centerline, push on the tab to release it from the chassis**
9. **While pushing on the tab, lift up on the module to clear the retention stub.**
10. **Slide the module towards the power supply until it comes free.**
11. **Lift the fan module out of the chassis.**

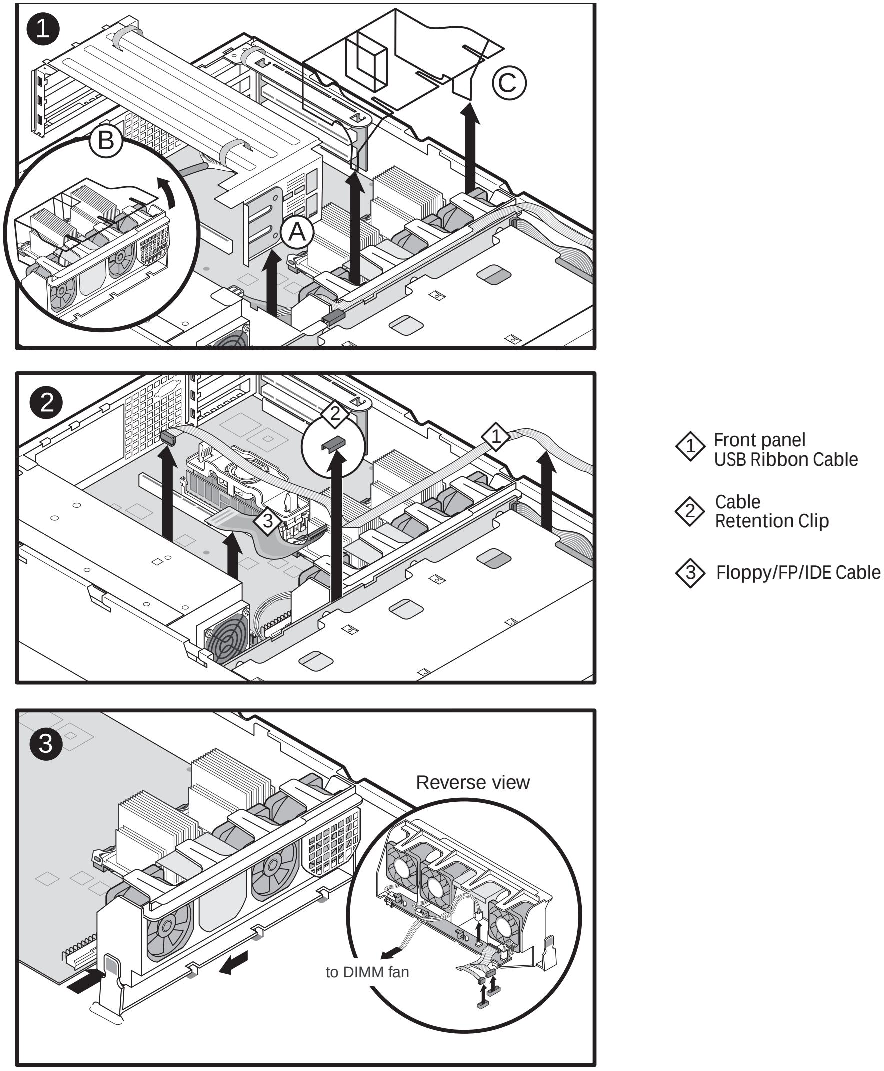

FIGURE 7-5 Removing the Fan Module

7.6.4.2 Sun StorEdge 5310 NAS Fan Module Replacement

Replacing the fan module is essentially the reverse of the procedure described in "Sun StorEdge 5310 NAS Fan Module Removal" on page 7-9.

1. **Note the raised tabs on the chassis floor and the corresponding notches in the bottom of the fan module.**
2. **Lower the fan module until it is just above the chassis floor.**
3. **Align the notches in the fan module with the raised tabs on the chassis and lower the fan module onto the floor.**
4. **While pressing down on the fan module, slide it to the right until the latch snaps into place.**
5. **Plug the fan cables into the server board system fan connectors.**
6. **Make sure the USB cable is routed along the top of the fan module.**
7. **Connect the flex circuit cable to the backplane.**
8. **Install the flex circuit cable retention clip.**
9. **Install the full-height PCI riser board.**
10. **Replace the plastic processor air duct.**
11. **Thread the memory fan power cable through the retaining hooks on the plastic processor air duct.**
12. **Replace the chassis cover.**

7.6.5 High Profile Riser PCI Cards

Note – Add-in cards must be replaced while the riser board is removed from the chassis.

The server supports 3V only and Universal PCI cards. It does not support 5V only cards.

Caution – Before touching or replacing any component inside the Sun StorEdge 5310 NAS, disconnect all external cables and follow the instructions in "Safety: Before You Remove the Cover" on page 7-2 and "Removing and Replacing the Cover" on page 7-2. Always place the server on a grounded ESD pad and wear a properly grounded antistatic wrist strap.

To replace a PCI card, follow these steps:

1. **Before removing the cover to work inside the system, observe the previously stated safety guidelines.**
2. **Remove the chassis cover.**
3. **Insert your finger in the plastic loop on the PCI riser assembly.**
4. **Pull straight up and remove the riser assembly from the chassis.**
5. **Open the retainer clip on the riser card retention bracket.**
6. **Pull the PCI card out of the riser board slot.**
7. **Install the new PCI add-in card on the riser assembly.**
8. **Insert the riser assembly connector in the server board slot while aligning the tabs on the rear retention bracket with the holes in the chassis.**
9. **Firmly press the riser assembly straight down until it is seated in the server board slot.**
10. **Replace the chassis cover if you have no additional work to do inside the chassis.**

7.6.6 Gigabit Ethernet Card

Note – Add-in cards must be replaced while the riser board is removed from the chassis.

The server supports 3V only and Universal PCI cards. It does not support 5V only cards. Alternately, a 10/100 Ethernetcard may be used for Cluster HeartBeat . The procedure to replace the 10/100 Ethernet is the same as the fibre Ethernet card.

Caution – Before touching or replacing any component inside the Sun StorEdge 5310 NAS, disconnect all external cables and follow the instructions in "Safety: Before You Remove the Cover" on page 7-2 and "Removing and Replacing the Cover" on page 7-2. Always place the server on a grounded ESD pad and wear a properly grounded antistatic wrist strap.

To replace a PCI card, follow these steps:

1. **Before removing the cover to work inside the system, observe the previously stated safety guidelines.**
2. **Remove the chassis cover.**
3. **Disconnect the network cable from the interface on the network card.**
4. **Insert your finger in the plastic loop on the PCI riser assembly.**
5. **Pull straight up and remove the riser assembly from the chassis.**
6. **Open the retainer clip on the riser card retention bracket.**
7. **Pull the PCI card out of the riser board slot.**
8. **Install the new PCI add-in card on the riser assembly.**
9. **Insert the riser assembly connector in the server board slot while aligning the tabs on the rear retention bracket with the holes in the chassis.**
10. **Firmly press the riser assembly straight down until it is seated in the server board slot.**

11. **Replace the chassis cover if you have no additional work to do inside the chassis.**

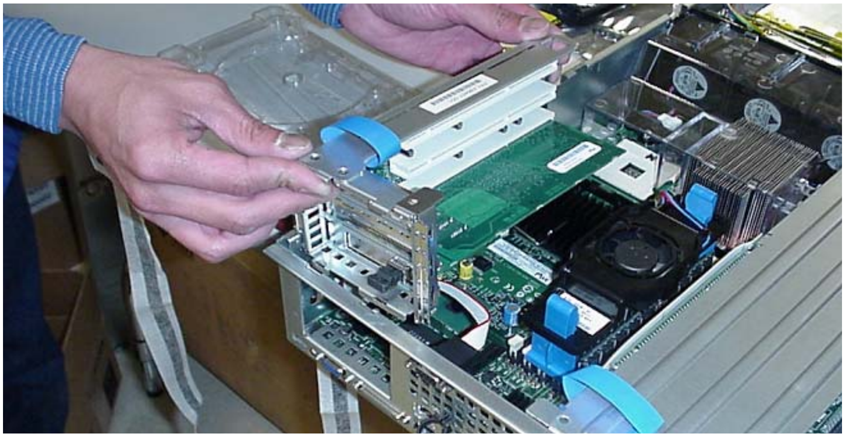

FIGURE 7-6 The Gigabit Ethernet Card in the Low Profile Riser Slot

7.6.7 Low Profile Riser PCI Cards

Note – Add-in cards must be replaced while the riser board is removed from the chassis.

The server supports 3V only and Universal PCI cards. It does not support 5V only cards.

Caution – Before touching or replacing any component inside the Sun StorEdge 5310 NAS, disconnect all external cables and follow the instructions in "Safety: Before You Remove the Cover" on page 7-2 and "Removing and Replacing the Cover" on page 7-2. Always place the server on a grounded ESD pad and wear a properly grounded antistatic wrist strap.

To replace a PCI card, follow these steps:

1. **Before removing the cover to work inside the system, observe the previously stated safety guidelines.**
2. **Remove the chassis cover.**
3. **Insert your finger in the plastic loop on the PCI riser assembly.**
4. **Pull straight up and remove the riser assembly from the chassis.**
5. **Open the retainer clip on the riser card retention bracket.**
6. **Pull the PCI card out of the riser board slot.**
7. **Install the new PCI add-in card on the riser assembly.**
8. **Insert the riser assembly connector in the server board slot while aligning the tabs on the rear retention bracket with the holes in the chassis.**
9. **Firmly press the riser assembly straight down until it is seated in the server board slot.**
10. **Replace the chassis cover if you have no additional work to do inside the chassis.**

7.6.8 Qlogic HBA Removal and Replacement

Note – Add-in cards must be replaced while the riser board is removed from the chassis. The server supports 3V only and Universal PCI cards. It does not support 5V only cards.

Caution – Before touching or replacing any component inside the Sun StorEdge 5310 NAS, disconnect all external cables and follow the instructions in "“Safety: Before You Remove the Cover” on page 7-2 and “Removing and Replacing the Cover” on page 7-2. Always place the server on a grounded ESD pad and wear a properly grounded antistatic wrist strap.

To replace a HBA card, follow these steps:

1. **Before removing the cover to work inside the system, observe the previously stated safety guidelines. Remove the chassis cover.**
2. **Insert your finger in the plastic loop on the Low Profile PCI riser assembly.**
3. **Pull straight up and remove the riser assembly from the chassis.**
4. **Open the retainer clip on the riser card retention bracket.**
5. **Remove the HBA from the Low Profile Riser**
6. **Install the new HBA card on the riser assembly.**
7. **Insert the riser assembly connector in the server board slot while aligning the tabs on the rear retention bracket with the holes in the chassis.**
8. **Firmly press the riser assembly straight down until it is seated in the server board slot.**
9. **Replace the chassis cover if you have no additional work to do inside the chassis.**

7.6.9 LCD Display Module

Note – The LCD Display must be replaced while the cover is removed from the chassis.

Caution – Before touching or replacing any component inside the Sun StorEdge 5310 NAS, disconnect all external cables and follow the instructions in "Safety: Before You Remove the Cover" on page 7-2 and "Removing and Replacing the Cover" on page 7-2. Always place the server on a grounded ESD pad and wear a properly grounded antistatic wrist strap.

To replace the LCD Display, follow these steps:

1. **Before removing the cover to work inside the system, observe the previously stated safety guidelines.**
2. **Remove the chassis cover.**
3. **Remove the ribbon cable from the SCSI connector on the RAID card. Make a note of the connector that the ribbon cable is on.**
4. **Remove the filler plate above the LCD Display.**
5. **Cut the two (2) tie wraps that secure the USB cable to the Flash Drive cables. The tie wraps are located near the Flash Disk.**
6. **Remove the tape holding the USB cable in place.**
7. **Use a long, thin screwdriver to unlatch the LCD bay. Place the screwdriver in the bottom row fourth (4th) hole over from the end of the three (3) rows of holes on the right side of the chassis. Press the Blue tab, visible from the top and side, and gently pull the LCD Display out while feeding the slack to the LCD Display.**
8. **Unplug the USB cable from the LCD Display**
9. **Remove the slider bay bracket from the LCD Display and install it on the new LCD Display.**
10. **Plug the USB cable into the LCD Display.**
11. **Slide the LCD Display back into the bay until you hear a click, carefully pulling the slack from the USB cable.**
12. **Tape the USB cable back in place.**
13. **Add two tie wraps on the USB cable.**
14. **Replace the filler plate above the LCD bay.**

15. **Replace the chassis cover if you have no additional work to do inside the chassis.**

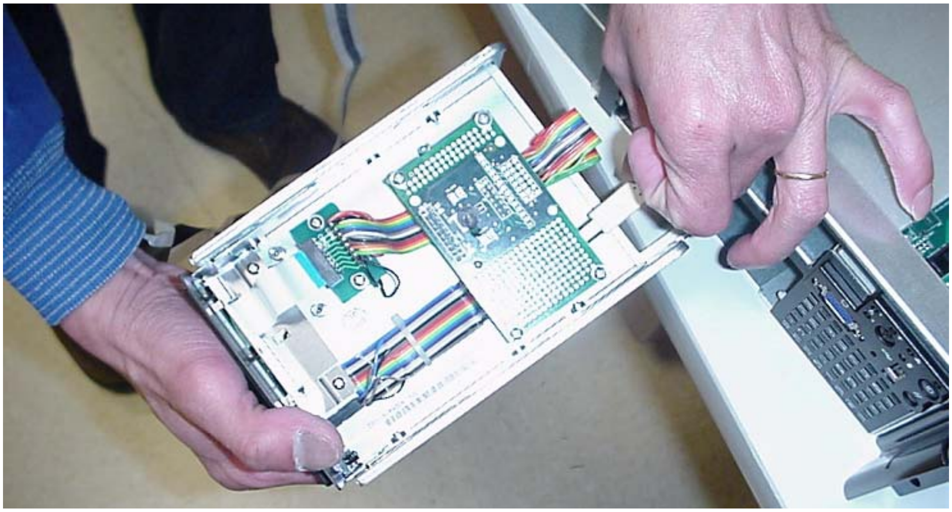

FIGURE 7-7 Connecting the LCD Display

7.6.10 Flash Disk Module

7.6.10.1 Backup of /dvol/etc

Assuming that the flash disk and the /etc directory are still accessible and in usable condition, the /dvol/etc directory should be backed up. This backup saves some configuration steps.

1. **Telnet to the StorEdge and access the CLI.**
2. **Type load unixtools**
3. **Type cp –r –v /dvol/etc <backup path>, replacing <backup path> with the full path, including volume name, to the desired directory location for the configuration files backup. (The destination directory must already exist, and should be empty.)**
4. **Check the output of the cp command to ensure that all files were successfully copied.**

7.6.10.2 Replacing the Flash Disk

Note – The Flash Disk must be replaced while the cover is removed from the chassis.

Caution – Before touching or replacing any component inside the Sun StorEdge 5310 NAS, disconnect all external cables and follow the instructions in "Safety: Before You Remove the Cover" on page 7-2 and "Removing and Replacing the Cover" on page 7-2. Always place the server on a grounded ESD pad and wear a properly grounded antistatic wrist strap.

To replace the Flash disk, follow these steps:

1. **Before removing the cover to work inside the system, observe the previously stated safety guidelines.**
2. **Remove the chassis cover.**
3. **Locate the Flash disk on the mother board, next to the ribbon cable and power supply housing unit.**
4. **Disconnect the Flash disk cable from the housing.**
5. **Remove the Flash disk from its holder.**
6. **Install the new Flash disk.**
7. **Install the cable from the Flash disk and plug it into the connector.**

8. **Replace the chassis cover if you have no additional work to do inside the chassis.**

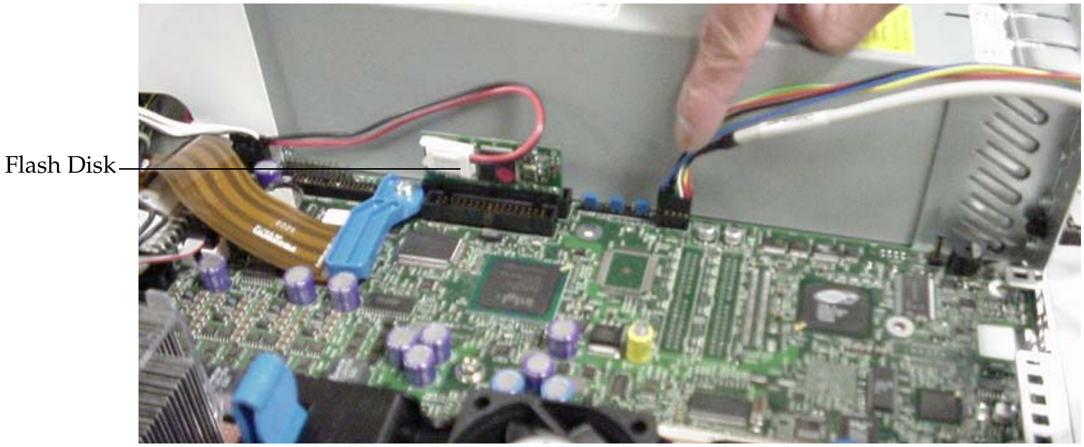

FIGURE 7-8 The Flash Disk

Note – After completing the flash disk replacement, you must recover the configuration information to bring the system back online.

7.6.10.3 Upgrade and Configuration Recovery

In this step, you will restore the system configuration and upgrade from the base operating system to a full version of the OS. The base OS contains only a limited subset of the StorEdge functionality, and cannot access the RAID volumes. Please note that this upgrade procedure is only valid when upgrading from the Base OS. The standard OS upgrade procedure is described elsewhere.

1. **Configure the server's IP address. This may be done automatically at startup via DHCP or it can be configured manually via the LCD display or the console Host Name & Network screen. Please see the software manual or the FAQ for IP setup instructions.**

2. **Telnet to the StorEdge at the newly assigned IP address. Verify that the system is running the Base OS by entering the version command on the CLI. Confirm that the output is similar to the following:**

```
StorEdge Model INSTALL S/N 0 Version 4.02 M38 (Build 154)
```

3. **Connect to the upgrade page on the server. The upgrade page is used to send the latest OS image to the server. The URL for the upgrade page is (please note the dot before BUILT-IN):**

http://<server-IP-address>/.BUILT-IN/upgrade/

4. **Ensure that the model field reads “NOMODEL”.**

5. **Enter 0 (zero) for the serial number field. (A blank serial number may NOT be used, and may cause this procedure to fail.)**

6. **Download the full version of the operating system to the local workstation. Contact Technical Support to obtain this file.**

7. **Click the browse button and navigate to the OS upgrade file on the workstation. Click on Install to copy the image to the server. The file transfer may take several minutes.**

   When the file download has finished, you should see the following message in the browser (your filename may vary):

   ```
   Received NF402B154.IMG for installation.
   You must now reboot to continue installation.
   ```

8. **Reboot the server to complete the upgrade process.**

   If the system has at least one existing volume, StorEdge will automatically restore the basic configuration information after the reboot, such as IP address, DNS and NT domain information. StorEdge will reboot once again if the information is found and successfully imported.

   If no volumes are configured, this setup information must be entered manually, according to the instructions in the Setup Poster.

Note – The system variables are backed up every four hours, at 4, 8 and 12am, and at 4, 8 and 12pm. Therefore, this procedure may not be able to recover configuration information for a newly installed system. In this case, you must also enter the information manually.

If the configuration information in the /dvol/etc was backed up to tape or the RAID volume, restore it to the same location. To restore from tape, simply restore the entire contents of the directory using the backup software. To restore from the RAID volume, use the console copy as follows:

1. **Telnet to the StorEdge and access the CLI (Please see the FAQ or software manual for instructions on accessing the CLI.)**

2. **Type load unixtools**

3. **Type cp –r –v <backup path> /dvol/etc, replacing <backup path> with the full path, including volume name, to the directory containing the backed up files.**

4. **Check the output of the cp command to ensure that all files were successfully copied.**

5. **Type reboot at the CLI to reboot the system and ensure that the new settings are in effect.**

   If a backup of /dvol/etc is not available, the following must be reconfigured manually: user maps, local group members and permissions, ssh keys, exports and related security, and local users/groups/hostgrps entries. Please see the software manual or the FAQ for setup instructions for these items.

7.6.11 System FRU (Super FRU)

Caution – The procedure below is for the attention of qualified service engineers only. Before touching or replacing any component inside the server, disconnect all external cables and follow the instructions in “Safety: Before You Remove the Cover” on page 7-2 and “Removing and Replacing the Cover” on page 7-2. Always place the server on a grounded ESD pad and wear a properly grounded antistatic wrist strap.

A System FRU is the Main Board with SCSI backplane, power supply, front-panel board, fan module and all cables. The System FRU contains no CPU(s), HDDs, or DIMMs. The field engineer transfers the customer's CPU(s), HDDs, and DIMMs to the new assembly.

7.6.11.1 Super FRU Installation

Before removing the cover to work inside the system, observe the safety guidelines previously given.

Please note that Super FRU is the entire system without the following components:

- CPU
- Power Supply
- System Fan
- Memory Fan
- Flash Disk
- LCD Module
- Optical GigE
- RAID Card
- TBBU
- HDD
- DIMMS

The System FRU is supplied with glue to install the Flash Disk. The items listed above must be installed into the Super FRU to restore the system.

7.7 Array FRU replacement Procedures

This chapter provides procedures for replacing failed components in an array command module. Before using the procedures in this chapter, perform the appropriate troubleshooting steps described in Chapter 3, "Troubleshooting and Recovery" and in the Recovery Guru.

The replacement procedures described in this chapter can be performed as hot swap procedures. Hot swap refers to the ability to remove and replace a component of the command module while it is processing I/O activity.

In a fully-configured command module, each component is redundant except for the midplane circuit board. If any redundant component fails and the storage array has been configured for data redundancy, the component can be replaced without powering off the array module and without interrupting data processing.

7.7.1 Replacing a Controller

Use the following procedure to replace a controller in a command module or command module.

Note – IMPORTANT To provide full functionality in dual controller configurations, the two controllers should have the same memory capacity. Although two controllers of different memories can be paired in a command module, the mismatch will cause some functions to be disabled (e.g., the cache mirroring function).

Note – IMPORTANT On dual controller command modules, a controller can be replaced without interrupting data transfer from the host. On a single controller command module, data transfer from the host must be terminated before the controller can be replaced.

7.7.1.1 Tools and Equipment

- Antistatic protection
- Phillips-head screwdriver
- Replacement controller

7.7.1.2 Procedure

1. **If needed, use the storage management software to create, save, and print a new storage array profile.**

2. **Did the Recovery Guru direct you to replace a failed controller?**
   - Yes - Go to step 3.
   - No - Run the Recovery Guru to identify the failed component. Go to step 3.

3. **Remove the back cover.**

4. **If needed, turn off the alarm.**

Caution – Electrostatic discharge damage to sensitive components. To prevent electrostatic discharge damage to the module, use proper antistatic protection when handling the module components.

---

5. **Put on antistatic protection.**

6. **Unpack the new controller.**

   Set the new controller on a dry, level surface. Save all packing materials in case you need to return the controller.

7. **Check the Fault lights to locate the failed controller. On the failed controller, the amber fault light will be illuminated. Figure 7-9 shows the location of the fault light.**

Caution – Potential data loss or degraded performance. To prevent data loss or damage to a cable, do not twist, fold, pinch, or step on a fibre optic cable, and do not bend a cable tighter than a 2-inch radius.

---

8. **Disconnect the SFP transceivers and all attached interface cables from the failed controller. Label all cables such that you can reconnect them correctly to the new controller.**

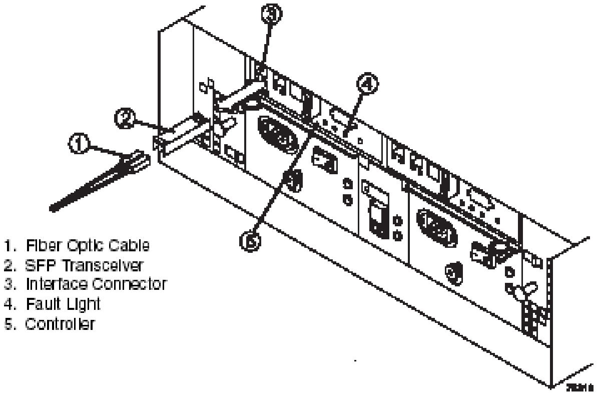

FIGURE 7-9 Removing an SFP Transceiver and fibre Optic Cable

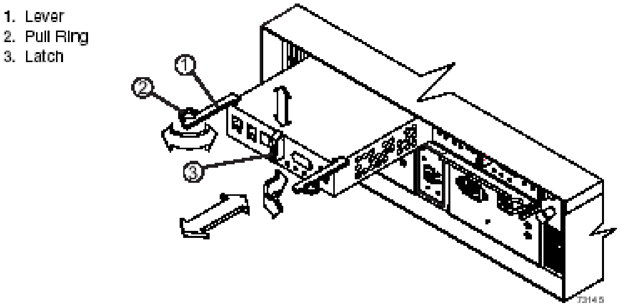

FIGURE 7-10 Removing and Replacing a Controller

9. **Remove the failed controller. Figure 7-10 illustrates the following steps:**

   a. **Push down on the latch.**

   b. **Open the levers.**

   c. **Removed the controller.**

10. **Does the replacement controller already have a battery installed?**

- Yes - Go to step 14.
- No - If the battery in the failed controller is still viable, you have the option of using that battery in the replacement controller. Go to step 11.

11. **Are you using the battery from the old controller?**

    - Yes - Go to step 12.
    - No - Unpack the new battery.

    Set the new battery on a dry, level surface. Save all packing materials in case you need to return the battery, and then go to step 13.

12. **Remove the battery from the failed controller.**

    a. **Turn the controller upside down, remove the screws securing the controller cover, and remove the cover. Figure 7-11 shows the cover and screws.**

    b. **Remove the single screw securing the battery bracket, slide the bracket sideways to clear the lugs, and lift the bracket up. Figure 7-12 on page 7-27 shows the bracket in relation to the controller.**

    c. **Disconnect the battery harness from its controller board connector.**

    d. **Remove the battery from the controller.**

    You may need to hold the controller close above a flat surface and let the battery fall out. Do not let the battery pull on the battery harness.

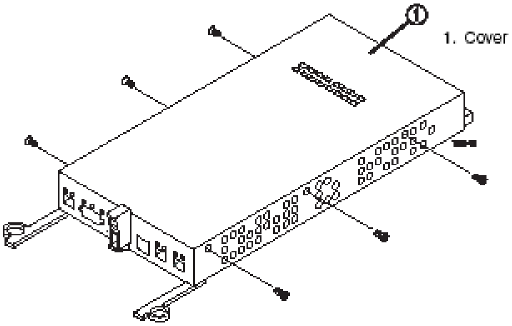

FIGURE 7-11 Removing the Controller Cover (Upside Down View)

13. **Install the replacement battery in the new controller.**

    a. **Connect the battery harness on the replacement battery to the connector on the controller board in the new controller. Figure 7-12 illustrates this connection.**

    b. **Position the battery inside the new controller.**

c. **Replace the controller cover and secure the screws. Figure 7-11 on page 7-26 shows these screws.**

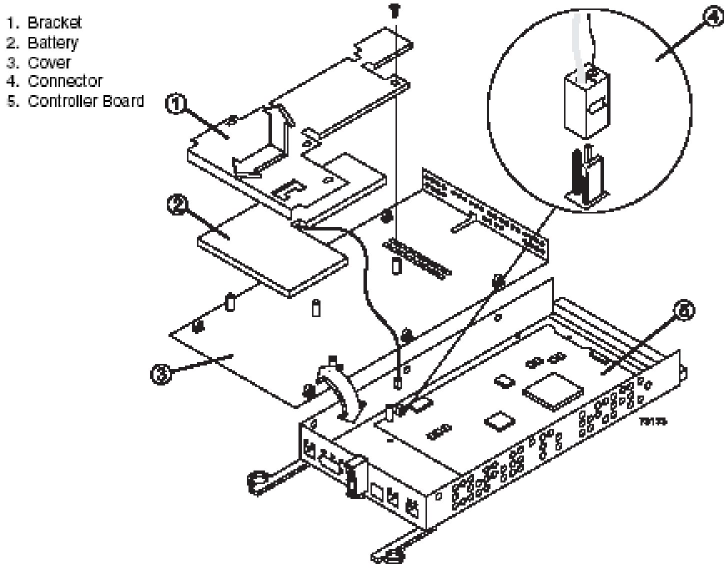

FIGURE 7-12 Replacing the Controller Battery

14. **Update the following information on the controller labels. Figure 7-13 on page 7-28 shows the location of the labels.**
    - Date of Installation - Enter today's date.
    - Replacement Date - Enter the date two years from now.
    - MAC Address - Record the MAC address for the new controller. You will need this information in step 17.

15. **Slide the replacement controller all the way into the empty slot and lock the levers into place. Figure 7-10 on page 7-25 illustrates removing and inserting a controller.**

16. **Reconnect the SFP transceivers and attach the host interface cables and drive interface cables to their original locations. Figure 7-9 on page 7-25 illustrates connecting an SFP transceiver.**

17. **Change the bootstrap protocol (BOOTP) server configuration to the MAC address you recorded in step 14. For detailed information on the configuration procedure, refer to your specific operating system administrator's guide.**

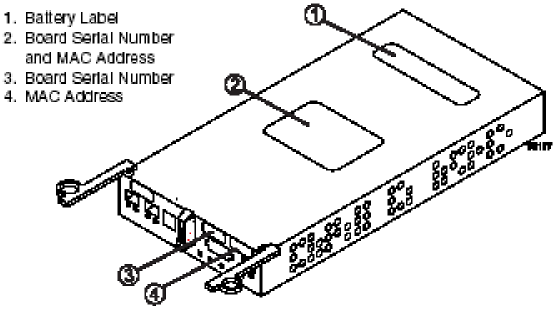

FIGURE 7-13 Label Locations for the Controller

18. **Wait approximately 60 seconds for the storage management software to recognize the new controller, and then go to step 19.**

Caution – Potential data loss. If the battery age is set incorrectly, you may not be notified to change the battery at the correct time, and data loss could occur. Contact technical support if the battery age is mistakenly reset.

Note – IMPORTANT During the Recovery Guru procedure, you will be asked to reset the battery age to zero for the battery and controller that you just installed. If you used the battery from the failed controller, do not reset the age. If you installed the new battery, be sure that you reset the age for the battery in the replacement controller. The software may display a message indicating that the battery in the new controller has failed or is nearing its expiration date. This message will disappear after you reset the battery age and the battery is fully charged.

19. **Complete any remaining Recovery Guru procedures for battery replacement, if needed.**

20. **Based on the status of the Host Link, Drive Link, and Fault lights, choose one of the following steps. Figure 7-14 shows the locations of these lights.**
    - All Link lights are off or the Fault light is illuminated - Verify that the controller has been installed correctly. Reinstall the controller, and then go to step 21.
    - All Link lights are on and the Fault light is off - Go to step 22.

21. **Is the problem corrected?**
    - Yes - Go to step 22.
    - No - Contact technical support.

22. **Remove the antistatic protection.**

23. **If needed, replace the back cover.**

24. **Use the Array Management Window to check the status of each module.**

25. **Do any components have a Needs Attention status?**
    - Yes - Select the Recovery Guru toolbar button in the Array Management Window and complete the recovery procedure. If the problem persists, contact technical support.
    - No - Go to step 26.

26. **Create, save, and print a new storage array profile.**

    End Of Procedure

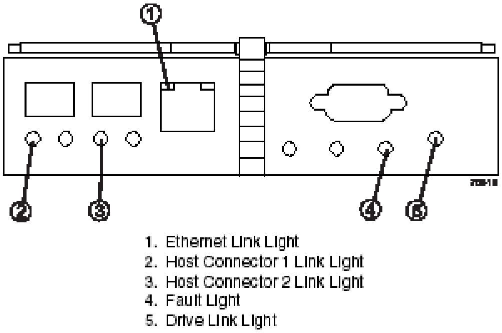
1
2
3
4
5
1. Ethernet Link Light
2. Host Connector 1 Link Light
3. Host Connector 2 Link Light
4. Fault Light
5. Drive Link Light

FIGURE 7-14 Controller Host Link, Drive Link, and Fault Lights

7.7.2 Replacing a Controller Battery

Use the following procedure to replace a controller battery and command module or command module.

7.7.2.1 Tools and Equipment

- Antistatic protection
- Phillips-head screwdriver
- Replacement controller battery

7.7.2.2 Procedure

1. **Use the storage management software to create, save, and print a new storage array profile.**

2. **Did the Recovery Guru direct you to replace a failed controller battery?**
   - Yes - Go to step 3.
   - No - Run the Recovery Guru to identify the failed component. Go to step 3.

3. **Remove the back cover.**

4. **If needed, mute the alarm.**

Caution – Electrostatic discharge damage to sensitive components. To prevent electrostatic discharge damage to the module, use proper antistatic protection when handling the module components.

---

5. **Put on antistatic protection.**

6. **Unpack the new battery.**

   Set the new battery on a dry, level surface. Save all packing materials in case you need to return the battery.

7. **Check the Fault lights to locate the failed controller battery. If a fault is detected, the amber Fault light will be illuminated. Figure 7-15 on page 7-31 shows the locations of these lights.**

Caution – Potential data loss or degraded performance. To prevent data loss or damage to a cable, do not twist, fold, pinch, or step on a fibre optic cable, and do not bend a cable tighter than a 2-inch radius.

---

8. Disconnect the SFP transceivers and all attached interface cables from the failed controller. Label all cables such that you can reconnect them correctly to the new controller. Figure 7-15 illustrates disconnecting a cable.

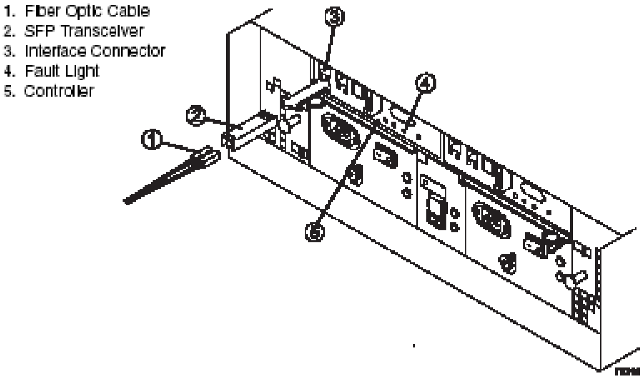

FIGURE 7-15 Removing the SFP Transceiver and fibre Optic Cable

9. Remove the failed controller. Figure 7-16 illustrates the following steps:

   a. Push down on the latch.

   b. Open the levers.

   c. Remove the controller.

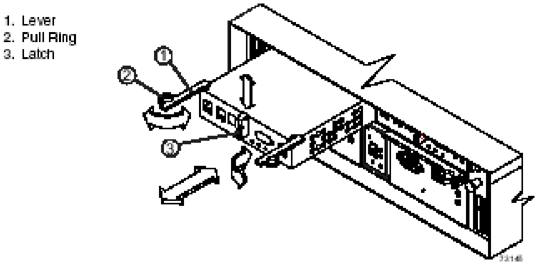

FIGURE 7-16 Removing and Replacing a Controller

10. Complete the following steps to replace the controller battery:

    a. Turn the controller upside down, remove the screws securing the controller cover, and remove the cover. Figure 7-17 shows the cover and screws.

b. **Remove the single screw securing the battery bracket, slide the bracket sideways to clear the lugs, and lift the bracket up. Figure 7-18 show the bracket in relation to the controller.**

c. **Disconnect the battery harness from its controller board connector.**

d. **Remove the battery from the controller.**

   You may need to hold the controller close above a flat surface and let the battery fall out. Do not let the battery pull on the battery harness.

e. **Connect the battery harness on the replacement battery to the connector on the controller board in the new controller. Figure 7-18 illustrates this connection.**

f. **Position the battery inside the new controller.**

g. **Replace the controller cover and secure the screws. Figure 7-17 shows these screws.**

11. **Update the following information on the new controller labels and attach the labels to the replacement controller. Figure 7-19 shows the locations of the labels.**

    - Date of Installation - Enter today's date.

- Replacement Date - If a new battery is used, enter the date two years from now. If the battery from the old controller is used, copy the date from the battery label on the old controller.

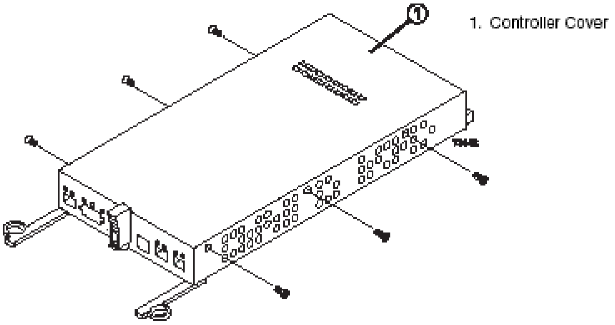

FIGURE 7-17 Removing the Controller Cover (Upside Down View)

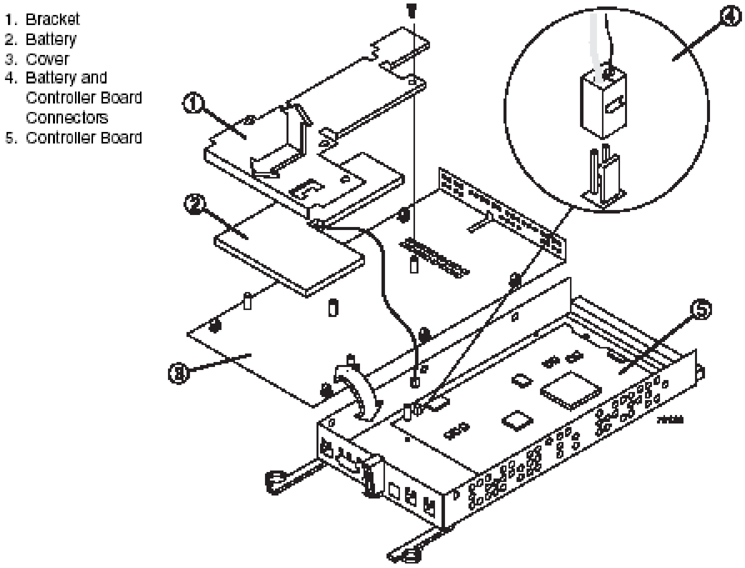

FIGURE 7-18 Removing and Installing the Controller Battery

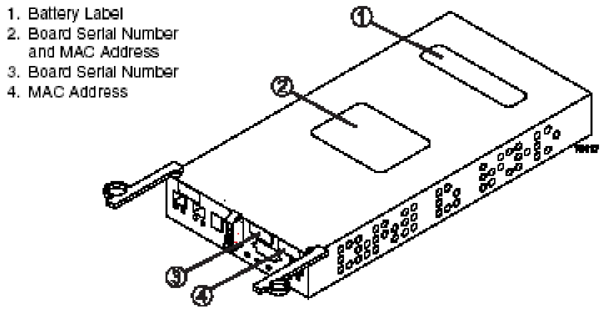

FIGURE 7-19 Label Locations on the Controller

12. **Slide the replacement controller all the way into the empty slot and lock the levers into place. Figure 7-16 on page 7-31 illustrates installing a controller.**

13. **Reinstall the SFP transceivers to their original connectors, and attach the host and drive interface cables to their respective SFP transceivers. Figure 7-15 on page 7-31 illustrates installing an SFP transceiver and cable.**

14. **Wait approximately 60 seconds for the storage management software to recognize the new controller.**

Caution – Potential data loss. If the battery age is set incorrectly, you may not be notified to change the battery at the correct time, and data loss could occur. If the battery age is mistakenly reset, contact technical support.

Note – IMPORTANT Reset the age of the new battery to zero when asked to do so during the Recovery Guru procedure. The software may display a message indicating that the battery in the new controller has failed or is nearing its expiration date. This message will disappear after you reset the battery age and the battery is fully charged.

---

15. **Complete any remaining Recovery Guru procedures for the controller replacement.**

16. **Choose one of the following steps, based on the status of the Host Link, Drive Link, and Fault lights. Figure 7-20 on page 7-36 shows the locations of these lights.**

    - All Link lights are off or the Fault light is illuminated - Verify that the controller has been installed correctly. Reinstall the controller, and then go to step17.
    - All Link lights are on and the Fault light is off - Go to step 18.

17. **Is the problem corrected?**
    - Yes - Go to step 18.
    - No - Contact technical support.

18. **Remove the antistatic protection and replace the back cover, if needed.**

19. **Check the status of all modules.**

20. **Do any components have a Needs Attention status?**
    - Yes - Select the Recovery Guru toolbar button in the Array Management Window and complete the recovery procedure. If the problem persists, contact technical support.
    - No - Go to step 21.

21. **Create, save, and print a new storage array profile.**

⚠ **Caution –** WARNING Potentially hazardous material. The battery canister contains sealed lead acid batteries that may be considered hazardous material. Use proper facilities to recycle the used battery. You must handle this unit in accordance to all applicable local and federal regulations.

⚠ **Caution –** WARNING Potentially hazardous material. If the used battery is physically damaged and is leaking electrolyte gel, DO NOT ship it to a recycling center. Doing so exposes you and others to potentially hazardous material. You must dispose of damaged batteries according to local regulations, which may include procedures for handling batteries as a hazardous waste.

---

22. **Dispose of the used battery according to local and federal regulations, which may include hazardous material handling procedures.**

23. **After 24 hours, check the host link, drive link, fault, and battery lights to ensure the battery is working properly. Figure 7-20 shows the locations of these lights. If the battery has a fault, use the storage management software to check the command module status and obtain the recovery procedure.**

End Of Procedure

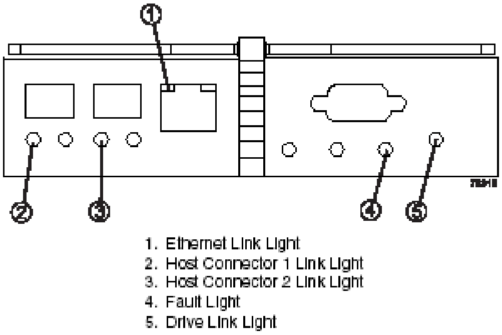
1
2
3
4
5
1. Ethernet Link Light
2. Host Connector 1 Link Light
3. Host Connector 2 Link Light
4. Fault Light
5. Drive Link Light

FIGURE 7-20 Drive Link, Host Link, Battery, and Fault Lights

7.7.3 Replacing a Drive

Use the following procedure to replace a drive in a command module. Figure 7-21 illustrates inserting and removing a drive.

7.7.3.1 Tools and Equipment

- Antistatic protection
- Replacement drive

7.7.3.2 Procedure

Caution – Potential data loss or data corruption. Never insert drives into a drive module without first confirming the drive firmware level. Inserting a drive with the incorrect firmware level may cause data loss or data corruption. For information on supported drive firmware levels, contact technical support.

Caution – Mixed configurations speed requirements. In configurations involving various models of command modules, command modules, or drive modules, all modules must be operating at the same speed. Refer to the Product Release Notes for any model-specific restrictions.

Caution – Risk of data loss and permanent damage. Magnetic fields will destroy all data on a disk drive and cause irreparable damage to its circuitry. To prevent data loss and damage to disk drives, always keep drive modules and drives away from magnetic devices.

1. **Use the storage management software to create, save, and print a new storage array profile.**
2. **Did the Recovery Guru direct you to replace a failed drive?**
   - Yes - Go to step 3.
   - No - Run the Recovery Guru to identify the failed component,

   and then go to step 3.
3. **Remove the back cover**
4. **4 If needed, mute the alarm.**

Caution – Electrostatic discharge damage to sensitive components. To prevent electrostatic discharge damage to the module, use proper antistatic protection when handling the module components.

5. **5 Put on antistatic protection.**

Caution – Potential damage to drives. Bumping disk drives against another surface can damage the drive mechanism or connectors. To prevent damage when removing or installing a drive, always place your hand under the drive to support its weight

6 Unpack the new drive.

Set the new drive on a dry, level surface, away from magnetic fields. Save all packing materials in case you need to return the drive.

Potential data loss. Removing a drive that has not failed can cause data loss. To prevent data loss, remove only a failed drive that has a Fault (amber) light on or a Failed status in the storage management software.

Caution – Potential data loss. Removing a drive that has not failed can cause data loss. To prevent data loss, remove only a failed drive that has a Fault (amber) light on or a Failed status in the storage management software.

6. **Check the Fault lights on the front of the module. If a fault is detected, the amber Fault light will be on.**

Note – IMPORTANT If you remove an active drive accidentally, wait 30 seconds and then reinstall it. Refer to your storage management software for the recovery procedure.

7. **Lift the locking lever and remove the failed drive.**

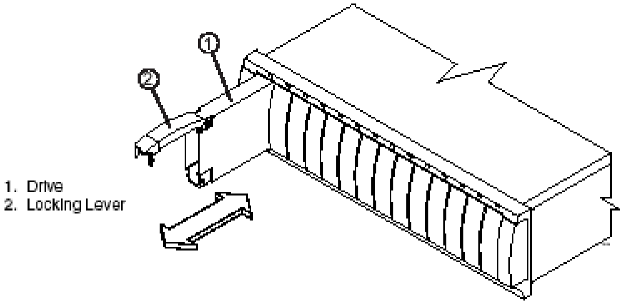

FIGURE 7-21 Replacing a Drive

8. **Wait 30 seconds for the storage management software to recognize that the drive has been removed.**

9. **Slide the new drive all the way into the empty slot and close the drive lever.**

   As the drive spins up, the Fault light may flash intermittently. A flashing Active light indicates that data is being restored to the new drive.

Note – IMPORTANT Depending on your storage array configuration, the storage array may automatically reconstruct data to the new drive. If the storage array uses hot spares, it may have to complete reconstruction on the hot spare before it copies the data to the replaced drive. This increases the time required to complete this procedure.

10. **Choose one of the following steps, based on the status of the Active and Fault lights:**
    - Active light is off - The drive may be installed incorrectly. Remove the drive, wait 30 seconds, and then reinstall it. When finished, go to step 12.
    - Fault light is illuminated - The new drive may be defective. Replace it with another new drive, and then go to step 12.
    - Active lights are on and Fault lights are off - Go to step 13. 12 Is the problem corrected?
    - Yes - Go to step 13.
    - No - Contact technical support.
11. **Remove the antistatic protection.**
12. **Bring the new drive online using the storage management software.**
13. **Complete any remaining Recovery Guru procedures, if needed.**
14. **Check the status of all modules in the storage array.**
15. **Do any components have a Needs Attention status?**
    - Yes - Select the Recovery Guru toolbar button in the Array Management Window and complete the recovery procedure. If the problem persists, contact technical support.
    - No - Go to step 18.
16. **Create, save, and print a new storage array profile.**

    End Of Procedure

7.7.4 Replacing a Fan

Use the following procedure to replace a fan in a command module. Figure 7-22 illustrates inserting and removing a fan.

7.7.4.1 Tools and Equipment

- Antistatic protection
- Replacement fan

7.7.4.2 Procedure

1. **If needed, use the storage management software to create, save, and print a new storage array profile.**
2. **Did the Recovery Guru direct you to replace a failed fan?**

- Yes - Go to step 3.
- No - Run the Recovery Guru to identify the failed component. Go to step 3.

3. **Remove the back cover.**

4. **If needed, mute the alarm.**

Caution – Electrostatic discharge damage to sensitive components. To prevent electrostatic discharge damage to the module, use proper antistatic protection when handling the module components.

---

5. **Put on antistatic protection.**

6. **Unpack the new fan.**

   Set the new fan on a dry, level surface. Save all packing materials in case you need to return the fan.

7. **If the storage management software instructs you to do so, turn off both power switches on the module. Otherwise, leave the power on.**

8. **Check the Fault lights to locate the failed fan. If a fan has a fault, its light will be illuminated and amber.**

9. **Slide the latch left or right (up or down on a deskside model) on the failed fan to unlock the component, and then pull on the knob to remove the failed fan.**

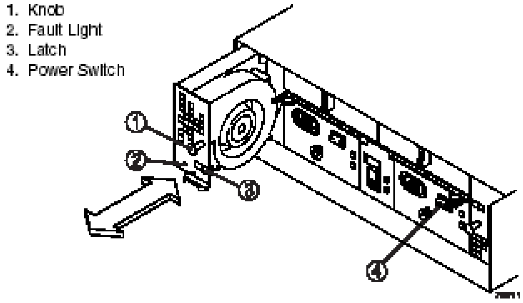

FIGURE 7-22 Replacing a Fan

10. **Slide the new fan all the way into the empty slot, until it snaps into place.**

11. **If you turned off the power in step 7, turn it on again.**

12. **Based on the status of the fan Fault light, choose one of the following steps:**

- Fault light is illuminated - The fan may be installed incorrectly. Reinstall the fan and then go to step 13.
- Fault light is off - Go to step 14.

13. **Is the problem corrected?**
    - Yes - Go to step 14.
    - No - Contact technical support.

14. **Remove the antistatic protection, and replace the back cover, if needed.**

15. **Complete the remaining Recovery Guru procedures, if needed.**

16. **Check the status of all the modules in the storage array.**

17. **Do any components have a Needs Attention status?**
    - Yes - Select the Recovery Guru toolbar button in the Array Management Window and complete the recovery procedure. If the problem persists, contact technical support.
    - No - Go to step 18.

18. **Create, save, and print a new storage array profile.**

    End Of Procedure

7.7.5 Replacing a Power Supply

Use the following procedure to replace a power supply in a command module. Figure 7-23 illustrates inserting and removing a power supply.

7.7.5.1 Tools and Equipment

- Antistatic protection
- Replacement power supply

7.7.5.2 Procedure

⚠ **Caution –** Electrostatic discharge damage to sensitive components. To prevent electrostatic discharge damage to the module, use proper antistatic protection when handling the module components.

1. **If needed, use the storage management software to create, save, and print a new storage array profile.**

2. **Did the Recovery Guru direct you to replace a failed power supply?**

- Yes - Go to step 3.
- No - Run the Recovery Guru to identify the failed component, and then go to step 3.

3. **Mute the alarm, and remove the back cover, if needed.**

4. **Put on antistatic protection.**

5. **Unpack the new power supply.**

   Set the new power supply on a dry, level surface near the command module. Save all packing materials in case you need to return it.

6. **Turn off the power switch on the new power supply.**

7. **Check the Fault lights to locate the failed power supply. If a fault is detected, the light will be illuminated and amber.**

Caution – WARNING Risk of electrical shock. Never remove or install a power supply that has its power cord plugged in and its power switch turned on. Doing so may expose you to the risk of electrical shock. Always turn off the power switch and unplug the power cord before removing or installing a power supply.

---

8. **Turn off the power switch, and unplug the power cord from the failed power supply.**

9. **Squeeze the pull ring on the failed power supply to release the lever. Open the lever and remove the power supply.**

10. **Verify that the lever on the new power supply opens in the same direction as the one you just removed. If it does not, move the lever to the other pivot post.**

11. **Slide the new power supply into the empty slot and close the lever.**

12. **Plug in the power cord and turn on the power.**

13. **Check the Power and Fault light on the new power supply.**

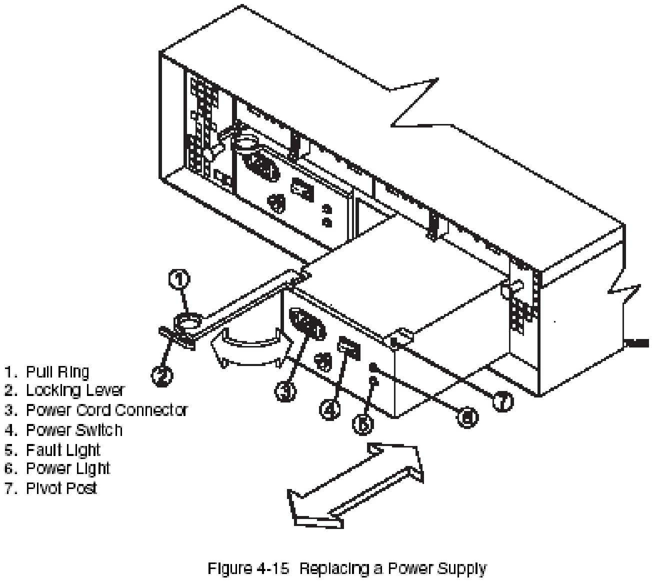

FIGURE 7-23 Replacing a Power Supply

14. **Choose one of the following steps, based on the status of the Power and Fault lights:**

    - Power light is off or Fault light is illuminated - The power supply may be

    installed incorrectly. Reinstall the power supply, and then go to step 15.

    - Power light is illuminated and Fault light is off - Go to step 16.

15. **Is the problem corrected?**

    - Yes - Go to step 16.
    - No - Contact technical support.

16. **Remove the antistatic protection, and replace the back cover, if needed.**

17. **Complete any remaining Recovery Guru procedures, if needed.**

18. **Check the status of all the modules in the storage array.**

19. **Do any components have a Needs Attention status?**

    - Yes - Select the Recovery Guru toolbar button in the Array Management Window and complete the recovery procedure. If the problem persists, contact technical support.

- No - Go to step 20.

20. **Create, save, and print a new storage array profile.**

    End Of Procedure

7.7.6 Replacing an SFP Transceiver

Use the following procedure to replace a Small Form-factor Pluggable (SFP) transceiver in a command module. The SFP transceiver shown in this procedure may look different from those you are using, but the difference will not affect transceiver performance. Figure 7-24 illustrates connecting an SFP and a cable.

7.7.6.1 Tools and Equipment

- Antistatic protection
- Replacement SFP transceiver

7.7.6.2 Procedure

1. **If needed, use the storage management software to create, save, and print a new storage array profile.**

2. **Did the Recovery Guru direct you to replace a failed SFP transceiver?**
   - Yes - Go to step 3.
   - No - Run the Recovery Guru to identify the failed component, and then go to step 3.

3. **Mute the alarm, and remove the back cover, if needed.**

> **Caution –** Electrostatic discharge damage to sensitive components. To prevent electrostatic discharge damage to the module, use proper antistatic protection when handling the module components.

4. **Put on antistatic protection.**

5. **Unpack the new SFP transceiver. Verify that it is the same type of transceiver you are replacing. Set the new SFP transceiver on a dry, level surface near the command module. Save all packing materials in case you need to return the SFP transceiver.**

6. **Check the Fault lights to locate the failed SFP transceiver. If a transceiver has a fault, its light will be illuminated and amber.**

Caution – Potential data loss or degraded performance. To prevent data loss or damage to a cable, do not twist, fold, pinch, or step on a fibre optic cable, and do not bend a cable tighter than a 2-inch radius.

Caution – Potential data loss. Removing an SFP transceiver that has not failed can cause data loss. To prevent data loss, remove only the component that has a Fault light on or a failed status in the storage management software.

7. **Disconnect the interface cable from the SFP transceiver.**
8. **Remove the failed SFP transceiver from the controller.**
9. **Install the new SFP transceiver into the controller.**
10. **Reconnect the interface cable.**
11. **Check the Bypass and Fault lights for the new SFP transceiver.**

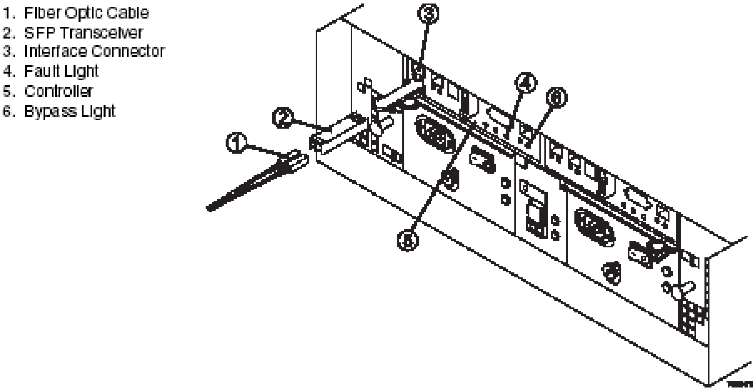

FIGURE 7-24 Replacing an SFP Transceiver

12. **Choose one of the following steps, based on the status of the Bypass and Fault lights.**
    - Bypass light or Fault light is illuminated -The SFP transceiver and cables maybe installed incorrectly, or the cable may not be securely connected. Reinstall the SFP transceiver and cable, check the cable connection, and then go to step 13.
    - Bypass light and Fault light are off - Go to step 14.
13. **Is the problem corrected?**

- Yes - Go to step 14.
- No - Contact technical support.

14. **Remove the antistatic protection, and replace the back cover, if needed.**

15. **Complete any remaining Recovery Guru procedures, if needed.**

16. **Check the status of each module.**

17. **Do any components have a Needs Attention status?**
    - Yes - Select the Recovery Guru toolbar button in the Array Management Window and complete the recovery procedure. If the problem persists, contact technical support.
    - No - Go to step 18.

18. **Create, save, and print a new storage array profile.**

    End Of Procedure