# Intel® Storage System Software User Manual

## Intel® Storage System SSR212MA

Intel Order Number: D26451-004

### Disclaimer

Information in this document is provided in connection with Intel® products. No license, express or implied, by estoppel or otherwise, to any intellectual property rights is granted by this document. Except as provided in Intel's Terms and Conditions of Sale for such products, Intel assumes no liability whatsoever, and Intel disclaims any express or implied warranty, relating to sale and/or use of Intel® products including liability or warranties relating to fitness for a particular purpose, merchantability, or infringement of any patent, copyright or other intellectual property right. Intel products are not designed, intended or authorized for use in any medical, life saving, or life sustaining applications or for any other application in which the failure of the Intel product could create a situation where personal injury or death may occur. Intel may make changes to specifications and product descriptions at any time, without notice.

Intel® server boards contain a number of high-density VLSI and power delivery components that need adequate airflow for cooling. Intel's own chassis are designed and tested to meet the intended thermal requirements of these components when the fully integrated system is used together. It is the responsibility of the system integrator that chooses not to use Intel developed server building blocks to consult vendor datasheets and operating parameters to determine the amount of airflow required for their specific application and environmental conditions. Intel Corporation can not be held responsible if components fail or the server board does not operate correctly when used outside any of their published operating or non-operating limits.

# Contents

# List of Figures

# List of Tables

# 1 Getting Started

Welcome to the Intel® Storage System Console (Console). The Console is used to configure and manage storage volumes spanning clustered Storage System Modules (SSMs).

After you have installed your SSMs and have installed the Console on the system administrator's PC, you must take certain steps to prepare for creating storage clusters and volumes.

The Console is the storage administrator's tool for:

- Configuring and managing the SSM
- Creating and managing clusters and volumes

This user manual provides instructions for installing the Console, configuring individual SSMs, as well as creating volumes that span a cluster of multiple SSMs. Topics in this manual include the following:

- Installing the Console
- Configuring individual SSMs by:
  - Configuring monitoring and reporting
- Creating volumes that span a cluster of SSMs by:
  - Creating management groups and clusters
  - Creating volumes that span multiple SSMs
  - Controlling client access to volumes
  - Creating and using snapshots of volume

## Using the Storage System Console

Use the Console to:

- Configure and manage storage modules
- Create and manage clusters and volumes

# The Console

The Console is divided into three sections, as shown in Figure 1.

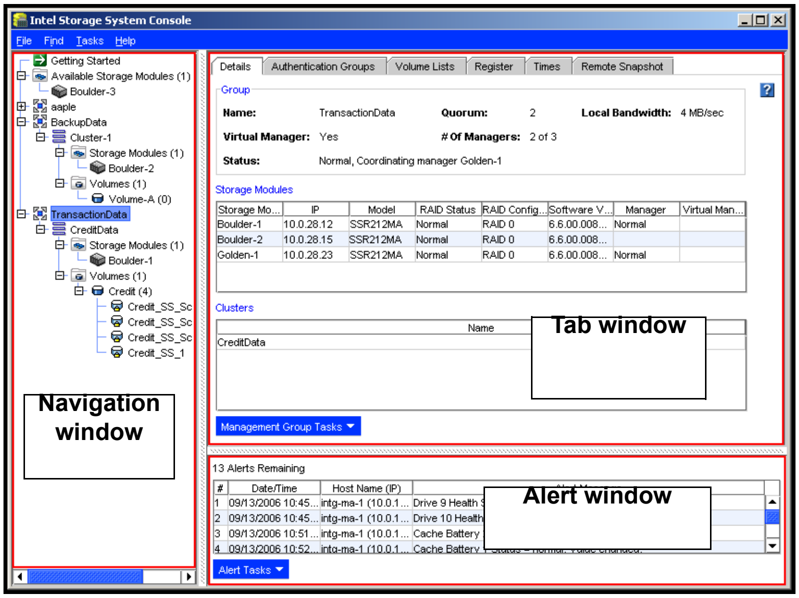

Group

**Name:** TransactionData **Quorum:** 2 **Local Bandwidth:** 4 MB/sec

**Virtual Manager:** Yes **# Of Managers:** 2 of 3

**Status:** Normal, Coordinating manager Golden-1

Storage Modules

| Storage Mo... | IP | Model | RAID Status | RAID Config... | Software V... | Manager | Virtual Man... |
|---|---|---|---|---|---|---|---|
| Boulder-1 | 10.0.28.12 | SSR212MA | Normal | RAID 0 | 6.6.00.008... | Normal | |
| Boulder-2 | 10.0.28.15 | SSR212MA | Normal | RAID 0 | 6.6.00.008... | | |
| Golden-1 | 10.0.28.23 | SSR212MA | Normal | RAID 0 | 6.6.00.008... | Normal | |

Clusters

| Name |
|---|
| CreditData |

13 Alerts Remaining

| # | Date/Time | Host Name (IP) | |
|---|---|---|---|
| 1 | 09/13/2006 10:45... | intg-ma-1 (10.0.1... | Drive 9 Health S |
| 2 | 09/13/2006 10:45... | intg-ma-1 (10.0.1... | Drive 10 Health |
| 3 | 09/13/2006 10:51... | intg-ma-1 (10.0.1... | Cache Battery |
| 4 | 09/13/2006 10:52... | intg-ma-1 (10.0.1... | Cache Battery |

**Figure 1. Viewing the Three Parts of the Console**

- **Navigation window**—the left vertical pane displays the architecture of your network. The items found in the navigation window include physical and logical elements of your network:
  - — Management groups
  - — Virtual managers
  - — Clusters
  - — Volumes
  - — Storage Modules
  - — Snapshots
  - — Remote snapshots
- **Tab window—**For each item selected in the navigation window, the tab window on the right displays information about it, and presents commands related to it.
- **Alert window—**The bottom window lets you view and manage the alerts that display there.

## Menu Bar

Other features of the Console include the following:

- **Menu Bar—**The menu bar provides access to the following menus:
  - — File
  - — Find—find modules on the network.
  - — Tasks — access all storage configuration tasks. The tasks in this menu are grouped by logical or physical items. Tasks are also accessible through right-click menus and from the Tasks button in the tab window.
  - — Help — Access online help and other information about the Console and Storage System Software.

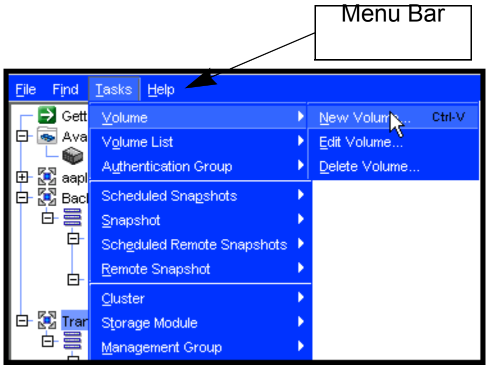

**Figure 2. Viewing the Menu Bar in the Navigation Window**

# Using the Navigation Window

The navigation window displays the components of your network architecture based the criteria you set in the Find item in the menu bar, or by using the Find Storage Modules wizard. That is, you can choose to display just the storage module you are interested in. You do not have to display all storage module at your site.

The navigation window reflects the architecture or topography of your network. Certain rules of hierarchy apply. For example, you may not create a volume until you create a management group.

## Logging In

After opening the Storage System Console, you must log in the first time you try to access a management group or an available storage module. After you have logged in the first time, the Console attempts to log you in automatically to other management groups or storage modules, using the first login.

# Traversing the Navigation Window

As you move through the items in the navigation window, the tab window changes to display the information and tabs about that item.

## Double-Clicking

Double-click on an item in the navigation window to open the hierarchy under that item. Double-click again to close it.

## Right-Clicking

Right-clicking on an item in the navigation window displays a menu of commands useful for that item.

## Getting Started

The first item in the navigation window is always the Getting Started Launch Pad. Click this to display the links to the three wizards you may use to begin your work.

## Available Storage Modules

The second item in the navigation window is Available Storage Modules, if you have storage modules that are not in management groups. This item contains the pool of available storage modules that you can add to management groups.

Other graphical information in the navigation window depicts the storage architecture you create on your system. An example setup is shown in Figure 1.

# Hierarchy

The items in the navigation window obey a specific hierarchy.

- **Management Groups**—Management groups are groups of storage modules within which one or more storage modules are designated as managers. Management groups are a logical item and provide a way to manage the clustered storage modules.
- **Clusters**. Clusters are a logical sub-groupings of storage modules within a management group. Clusters contain the data volumes and snapshots.
- **Volumes**. Volumes are the logical data storage entities that are presented to application servers as disks.
- **Snapshots**. Snapshots are point-in-time copies of volumes. Snapshots can be created manually, as necessary, or scheduled to occur regularly. Snapshots of a volume can be stored on the volume itself, or on volume different from the one where the snapshot was taken, that is, a remote volume.

# Icons

Each item is the navigation window has an icon depicting what type of item it is. If an icon is faded, it means that the item is remote, and not local or primary. A description is available of all the icons used in the Console.

1. Click Help on the menu bar.
2. Select Graphical Legend from the menu.

   The icon display window opens.

3. View the Items tab and the Hardware tab.

   - The Items tab, shown in Figure 3, displays the icons that represent items, activities, and status in the navigation window.

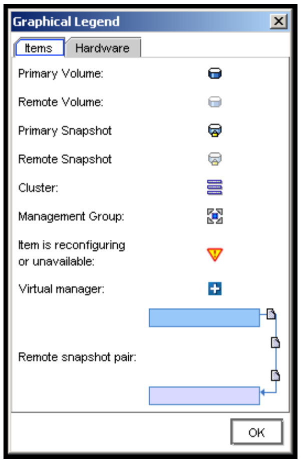

Primary Volume:

Remote Volume:

Primary Snapshot

Remote Snapshot

Cluster:

Management Group:

Item is reconfiguring or unavailable:

Virtual manager:

Remote snapshot pair:

## Figure 3. Viewing the Graphical Legend Items Tab

   - The Hardware tab, shown in Figure 4, displays the icons that represent the different models of physical storage modules that display in the navigation window.

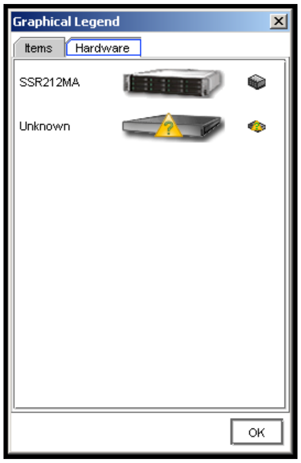

SSR212MA

Unknown

## Figure 4. Viewing the Graphical Legend Hardware Tab

# Using the Tab Window

The tab window displays information about an item selected in the navigation window, as well as tabs for functions related to that item. For example, Figure 5 shows the tabs that display when a management group is selected in the navigation window.

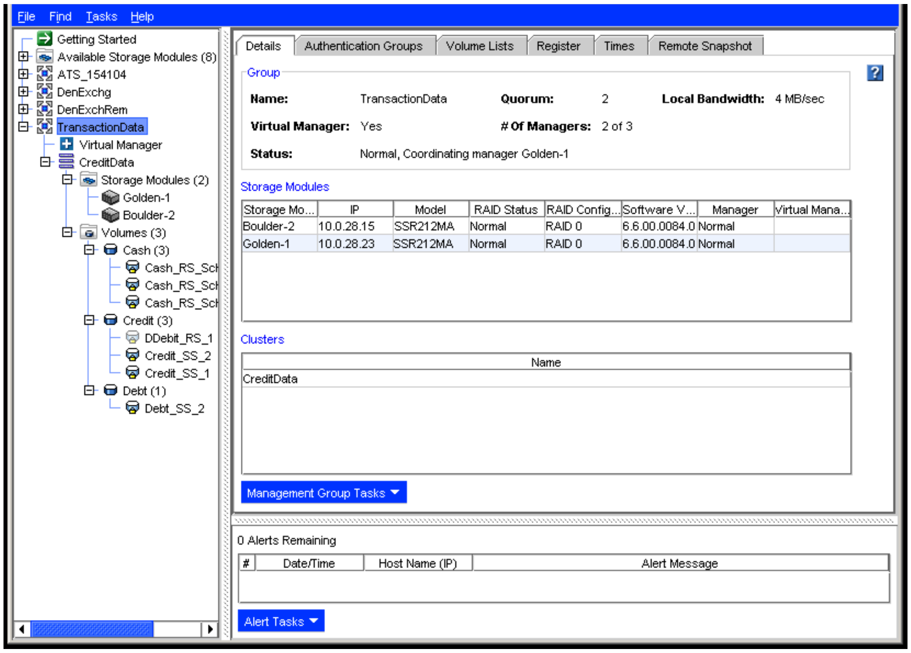

Group

**Name:** TransactionData **Quorum:** 2 **Local Bandwidth:** 4 MB/sec

**Virtual Manager:** Yes **# Of Managers:** 2 of 3

**Status:** Normal, Coordinating manager Golden-1

Storage Modules

| Storage Mo... | IP | Model | RAID Status | RAID Config... | Software V... | Manager | Virtual Mana... |
|---|---|---|---|---|---|---|---|
| Boulder-2 | 10.0.28.15 | SSR212MA | Normal | RAID 0 | 6.6.00.0084.0 | Normal | |
| Golden-1 | 10.0.28.23 | SSR212MA | Normal | RAID 0 | 6.6.00.0084.0 | Normal | |

Clusters

| Name |
|---|
| CreditData |

0 Alerts Remaining

| # | Date/Time | Host Name (IP) | Alert Message |
|---|---|---|---|

Figure 5. Viewing tab Windows in the Console

# Tab Window

The tab windows provide access to the functions of the iSCSI SAN that is configured on your network. Select a tab to perform functions related to the selected item.

## Conventions

Tab windows have certain similarities:

- Tabs—provide access to functions that relate to the item selected in the navigation window. For example, when a management group is selected, the tabs in the tab window reflect items and activities related to management groups.
- Lists—When presented with a list, such as a list of storage modules as seen in the management group Details tab, you may select an item in the list to perform an action on.
- Lists and right-click—right-click on an item in a list and a drop-down list of commands appropriate for that item appears.

- Tasks buttons—at the bottom of a tab window, the tasks button opens a menu of commands available for the item or function of that tab.
- Sortable columns—click on a column head to sort the list on that column
- Sizable columns—drag a column boundary to the right or left to widen the column for better reading.

# Using the Alert Window

Alert messages appear in the alert window as they are generated and are removed when the alert situation resolves on its own, or when you remove them with the Alert Tasks commands.

- Right-click on an alert to expand just that alert in a separate, larger window.
- Clear all alerts, or clear selected alerts using the Alert Tasks menu.

# Creating Storage By Using the Getting Started Launch Pad

Follow the steps in this section to set up a volume quickly. Using the wizards on the Getting Started Launch Pad, you will work through these steps with one storage module, and with one strategy. The rest of this User Manual describes other methods to create storage, as well as detailed information on features of the iSCSI SAN.

**Prerequisites**

- Install the storage modules on your network.
- Know the IP address or host name you configured with the serial Configuration Interface when you installed the storage module.
- Install the Storage System Console software on the system administrator’s PC and familiarize yourself with the interface.
- Install an iSCSI initiator such as the latest version of the Microsoft® iSCSI initiator.

## Finding the Storage Modules

Open the Console, and using the Wizard Launch Pad, start the Find Storage Modules Wizard.

To work through the wizard, you need to know either the

- The subnet and mask of your storage network or
- The IP addresses or host names of the storage modules

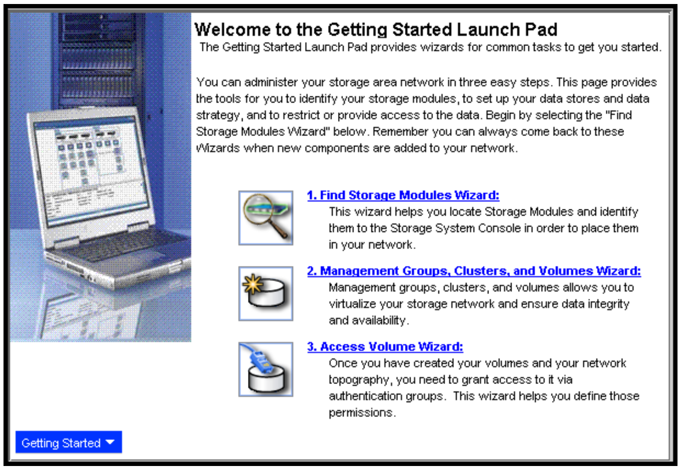

Welcome to the Getting Started Launch Pad

The Getting Started Launch Pad provides wizards for common tasks to get you started.

You can administer your storage area network in three easy steps. This page provides the tools for you to identify your storage modules, to set up your data stores and data strategy, and to restrict or provide access to the data. Begin by selecting the "Find Storage Modules Wizard" below. Remember you can always come back to these Wizards when new components are added to your network.

1. Find Storage Modules Wizard:
This wizard helps you locate Storage Modules and identify them to the Storage System Console in order to place them in your network.

2. Management Groups, Clusters, and Volumes Wizard:
Management groups, clusters, and volumes allows you to virtualize your storage network and ensure data integrity and availability.

3. Access Volume Wizard:
Once you have created your volumes and your network topography, you need to grant access to it via authentication groups. This wizard helps you define those permissions.

Figure 6. The Wizard Launch Pad

# Configuring Storage Modules

Configure the storage module next. If you plan to use multiple storage modules, they must all be configured individually before you use them for clustered storage.

1. From the navigation window, select a storage module in the Available pool.
2. In the tab window, click Storage Module Tasks and select Log In.
3. Click Storage Module Tasks and this time select Edit Configuration. The Edit Configuration window appears, as shown in Figure 7.

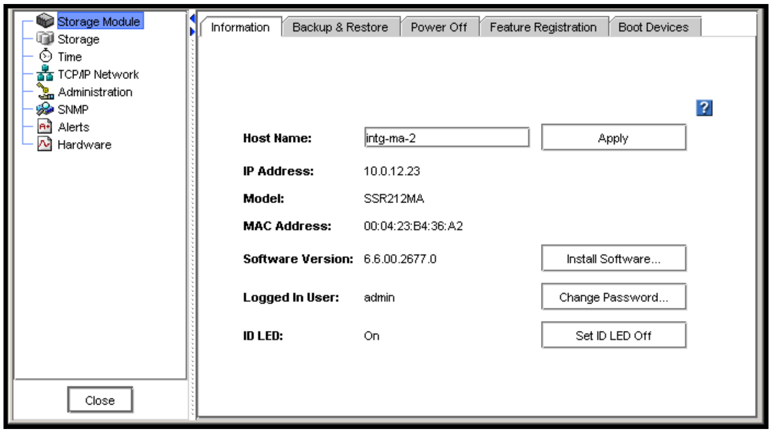

Host Name: intg-ma-2

IP Address: 10.0.12.23

Model: SSR212MA

MAC Address: 00:04:23:B4:36:A2

Software Version: 6.6.00.2677.0

Logged In User: admin

ID LED: On

Figure 7. Viewing the Edit Configuration window

4. Select the Storage category on the left, and verify the RAID settings

   The storage module may be shipped with RAID already configured and operational. Instructions for ensuring that drives in the storage module are properly configured and operating are in "Storage" on page 53.

5. Select the TCP/IP Network category and make sure your network configuration is correct.

   Read detailed network configuration instructions in "Managing the Network" on page 85.

6. Select the SNMP and/or Alerts categories to configure the monitoring for your IP SAN.

   Detailed information for setting up SNMP and alerts can be found in "Using SNMP" on page 141 and "Using Active Monitoring" on page 149.

Later on, you can return to this window and refine your configuration.

You can copy the configuration of the first storage module to others in the management group, as described in "Configuring Multiple Storage Modules" on page 11.

# Creating a Volume

Your storage module may now be assigned to a management group. Use the Management Groups, Clusters, and Volumes Wizard wizard on the Wizard Launch Pad to create the hierarchy of your Storage System Software storage network in the navigation window.

Volumes fit into the hierarchy depicted in Figure 8.

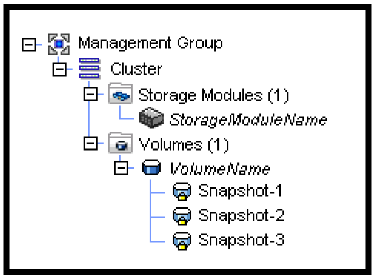

**Figure 8. Storage System Software Network Hierarchy**

To create a volume, the wizard first creates a management group, and next creates a cluster with the storage module, and finally creates a volume.

While working through the wizard, you need to know the following information:

- A name for your management group
- A storage module that you identified with the Find wizard and then configured
- A name for the cluster
- A name for the volume
- The size of the volume

## Enabling Client Access to Volumes

Use the Access wizard to grant read/write permissions and to begin using the volume.

Authentication groups and volume lists control which users, clients, or machines may access a volume. See "Controlling Client Access to Volumes" on page 265 for a complete discussion of these functions.

To work through the Access wizard, you will need to have these things:

- A cluster with a volume in it
- A descriptive name for a volume list
- A meaningful name for the authorization group
- The iSCSI initiator node name (you can copy and paste the name from the iSCSI Initiator Properties, General tab).

The wizard creates a volume list and an authentication group for the volume.

# Continuing with Storage System Software

This section describes useful methods of working with the Console on an ongoing basis. It also describes how you can copy the configuration of one storage module to others on your storage network.

## Finding Storage Modules on an Ongoing Basis

The Find settings from your first search are saved in the Console. Every time you open the Console, the search automatically takes place and the navigation window is populated with all storage modules that are found.

***Note:*** *You can control which storage modules appear in the navigation window by entering only specific IPs or Host Names in the IP and Host Name List window. Then, when you open the Console, only those IPs or Host Names will appear in the navigation window.*

### Modules Not Found

If the network has lots of traffic, or if a module is busy reading or writing data, it may not be found when a search is performed. Try the following steps to find the storage module.

1. If the storage module you are looking for does not appear in the navigation window pane, or in the IP and Host Name List window, search again using the Find menu.
2. If you have searched by Subnet and Mask, try using the Find by IP or Host Name search.
3. If searching again does not work, try the following:

   — Check the physical connection of the module.

— Wait a few minutes and try the search again. If activity to the storage module was high, the storage module might not have responded to the search.

***Note:*** *Other problems can occur that prevent connection, such as a bad cable connection.*

## Changing Which Storage Modules Appear in the Navigation Window

1. Click the Find menu
2. Select Clear All Found Items to remove all storage modules from the navigation window.
3. Perform a Find using either method, Subnet and Mask, or Module IP or Host Name, to find the desired set of storage modules.

# Configuring Multiple Storage Modules

As you add storage modules to the management group you can copy the reporting, monitoring, and time configuration of the first storage module to additional storage modules. Copying these configurations makes it easy to ensure that those storage modules have exactly the same configuration.

1. On the navigation window, select the storage module that has the configuration that you want to copy.
2. Right-click and select Copy Configuration.

   The Copy Configuration window opens, shown in Figure 9.

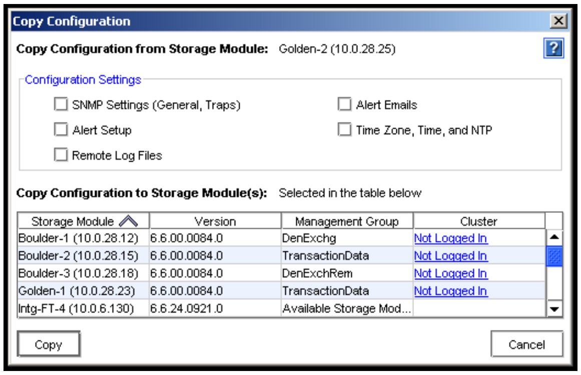

**Copy Configuration from Storage Module:** Golden-2 (10.0.28.25)

Configuration Settings

☐ SNMP Settings (General, Traps)
☐ Alert Emails
☐ Alert Setup
☐ Time Zone, Time, and NTP
☐ Remote Log Files

**Copy Configuration to Storage Module(s):** Selected in the table below

| Storage Module | Version | Management Group | Cluster |
|---|---|---|---|
| Boulder-1 (10.0.28.12) | 6.6.00.0084.0 | DenExchg | Not Logged In |
| Boulder-2 (10.0.28.15) | 6.6.00.0084.0 | TransactionData | Not Logged In |
| Boulder-3 (10.0.28.18) | 6.6.00.0084.0 | DenExchRem | Not Logged In |
| Golden-1 (10.0.28.23) | 6.6.00.0084.0 | TransactionData | Not Logged In |
| Intg-FT-4 (10.0.6.130) | 6.6.24.0921.0 | Available Storage Mod... | |

### Figure 9. Opening the Copy Configuration Window

3. In the Configuration Settings section, select which configurations you want to copy.

   For information about the configuration settings, see the following:

   - "Using SNMP" on page 141
   - "Using Active Monitoring" on page 149
   - "Remote Log Files" on page 165
   - "Setting Email Notification" on page 157

4. In the Copy Configurations to Storage Modules section, select the storage modules to which you want to copy the configurations.
5. Click Copy.

   The configuration settings are copied to the selected storage modules.
6. Click OK on the confirmation window.

   The Copy Configuration window closes.

# Installing the Console

1. Before you install the Console software, ensure that you have set up an IP address and password on your storage system. Instructions for setting up an IP address and password are available on the *Intel® Storage System Quick Start User's Guide* that shipped with your storage system.
2. Insert the Resource CD that shipped with your storage system into the system from which you will install the Console software.
3. Scroll down and click on the “Agree” to accept the license agreement.
4. Select “Microsoft Windows Selection” under “Software Heading”.
5. Select “Run” when prompted.
6. Click on “Next” at the introduction screen.
7. Read the license agreement and click on “I Accept” to accept the terms of the license agreement. Click on “Next”.
8. Select the “Typical” option and click on “Next”.

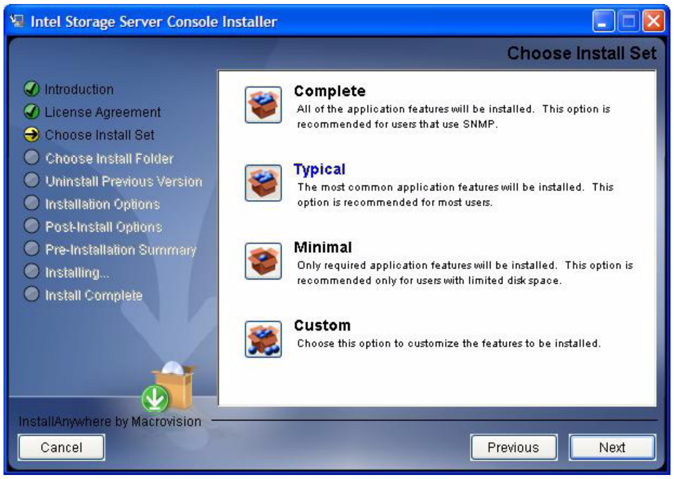

Choose Install Set

Complete
All of the application features will be installed. This option is recommended for users that use SNMP.

Typical
The most common application features will be installed. This option is recommended for most users.

Minimal
Only required application features will be installed. This option is recommended only for users with limited disk space.

Custom
Choose this option to customize the features to be installed.

9. Enter the path of the folder you wish to use to install the Console software.

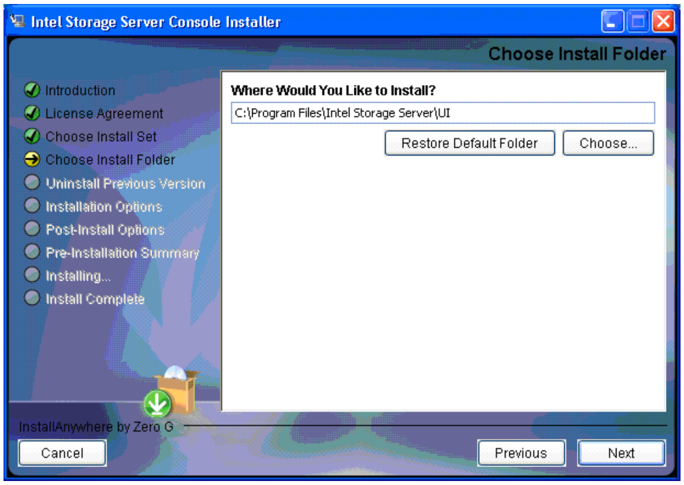

Intel Storage Server Console Installer

Choose Install Folder

Where Would You Like to Install?

C:\Program Files\Intel Storage Server\UI

10. Leave “Desktop” checked if you want the install routine to automatically create a shortcut to the Console on your desktop. Click on “Next” to continue.

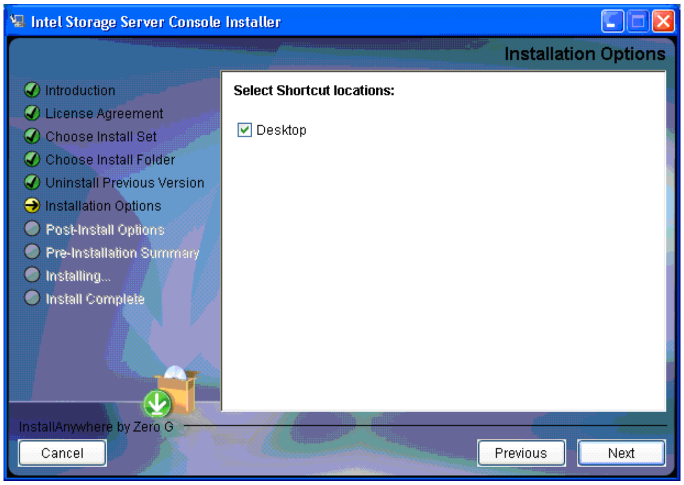

Intel Storage Server Console Installer

Installation Options

Select Shortcut locations:

☑ Desktop

11. Select “Yes” and click on “Next” to start the installation process.

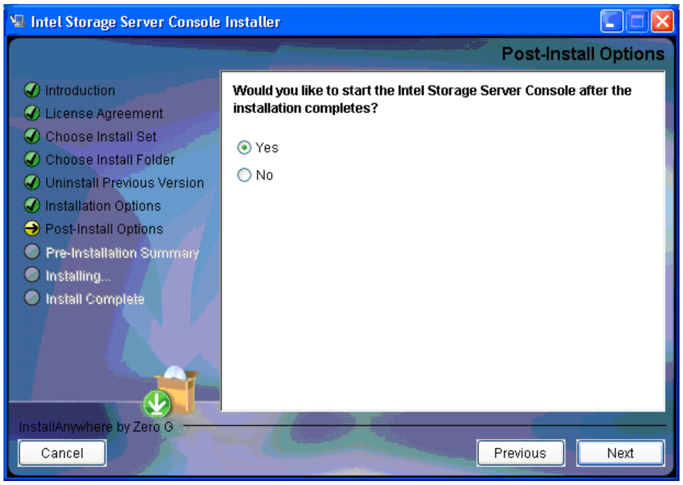

Intel Storage Server Console Installer

Post-Install Options

Would you like to start the Intel Storage Server Console after the installation completes?

● Yes

○ No

12. Review the Pre-Installation Summary and click on “Install”.

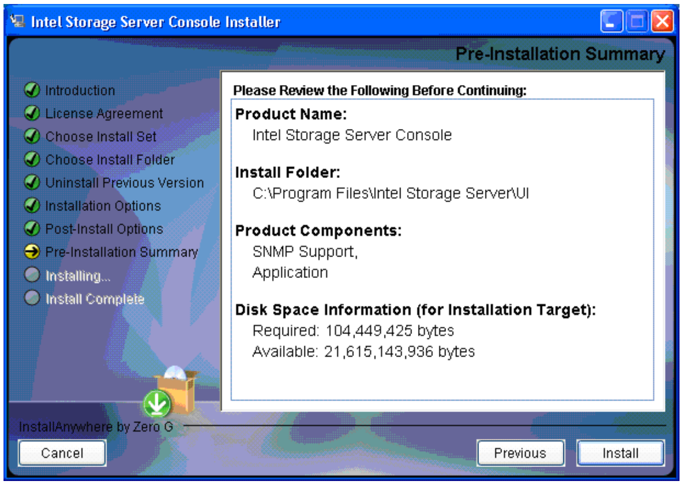

Intel Storage Server Console Installer

Pre-Installation Summary

Please Review the Following Before Continuing:

**Product Name:**
Intel Storage Server Console

**Install Folder:**
C:\Program Files\Intel Storage Server\UI

**Product Components:**
SNMP Support,
Application

**Disk Space Information (for Installation Target):**
Required: 104,449,425 bytes
Available: 21,615,143,936 bytes

13. The software will take several minutes to load.

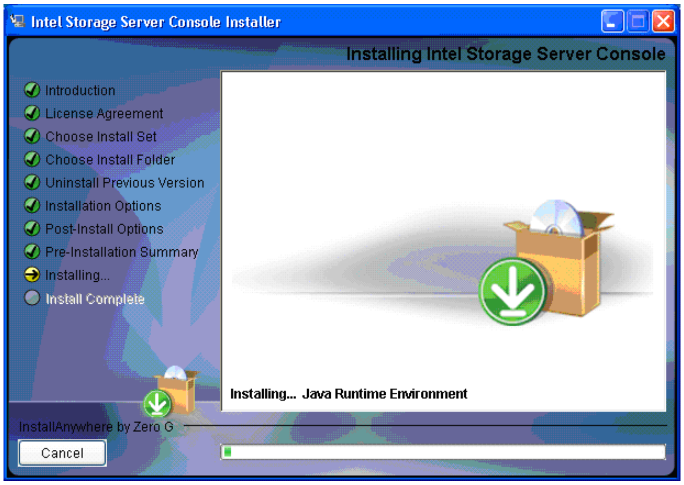

Installing... Java Runtime Environment

14. Click on “Done” once the software completely loads.

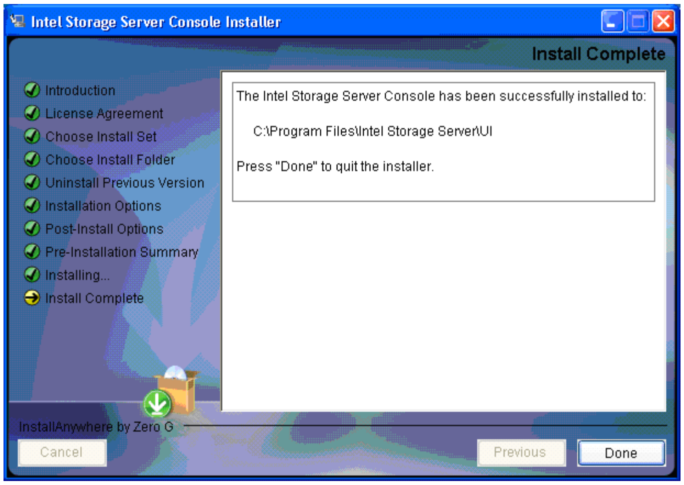

The Intel Storage Server Console has been successfully installed to:

C:\Program Files\Intel Storage Server\UI

Press "Done" to quit the installer.

15. Close out of the browser window of the Resource CD. If you selected Yes to run the Console, the Wizard Launch Pad will display. Select “Find SSMs Wizard” from the Wizard Launch Pad.

**Note:** *Ensure that you have set up an IP address on your storage system before you run the “Find SSMs Wizard”. Instructions for setting up an IP address and password are available from the Intel® Storage System Quick Start User’s Guide that shipped with your storage system.*

16. Select the Find Storage Modules Wizard.

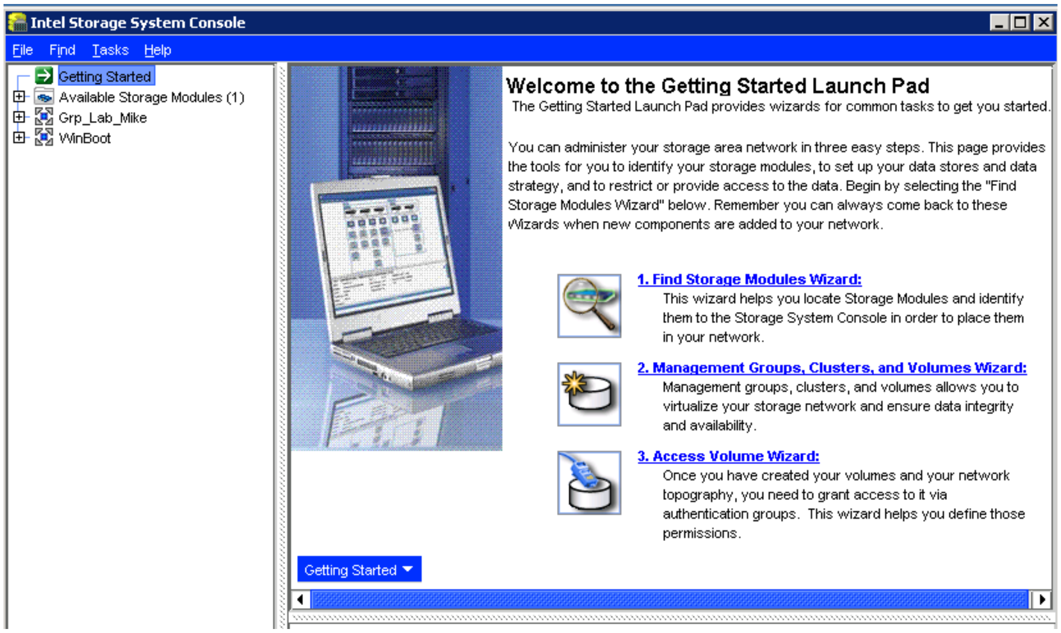

Welcome to the Getting Started Launch Pad

The Getting Started Launch Pad provides wizards for common tasks to get you started.

You can administer your storage area network in three easy steps. This page provides the tools for you to identify your storage modules, to set up your data stores and data strategy, and to restrict or provide access to the data. Begin by selecting the "Find Storage Modules Wizard" below. Remember you can always come back to these Wizards when new components are added to your network.

1. Find Storage Modules Wizard:
This wizard helps you locate Storage Modules and identify them to the Storage System Console in order to place them in your network.

2. Management Groups, Clusters, and Volumes Wizard:
Management groups, clusters, and volumes allows you to virtualize your storage network and ensure data integrity and availability.

3. Access Volume Wizard:
Once you have created your volumes and your network topography, you need to grant access to it via authentication groups. This wizard helps you define those permissions.

17. Select Next at the Find Storage Modules Wizard.

18. Select the “By Subnet and Mask (global broadcast)” or “By Module IP or Host Name (individual search)” option to search for your storage system. If you are using a fixed IP address, select the “By Module IP or Host Name (individual search)” option. Press “Next” to continue.

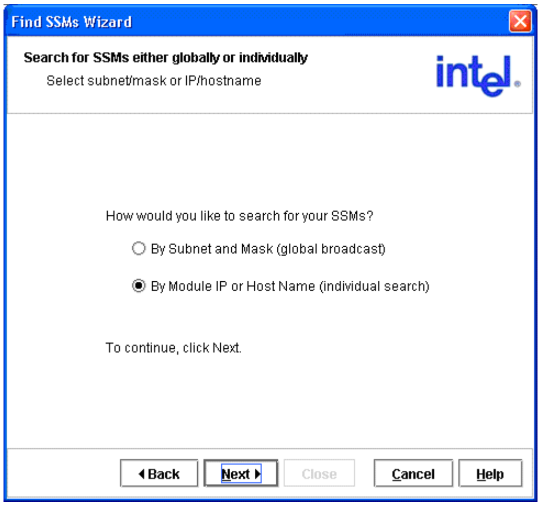

Find SSMs Wizard

**Search for SSMs either globally or individually**

Select subnet/mask or IP/hostname

How would you like to search for your SSMs?

○ By Subnet and Mask (global broadcast)

● By Module IP or Host Name (individual search)

To continue, click Next.

19. If you selected the “By Module IP or Host Name (individual search)” option, a Search by SSM IP or Hostname list will display. Click on the “Add” button.

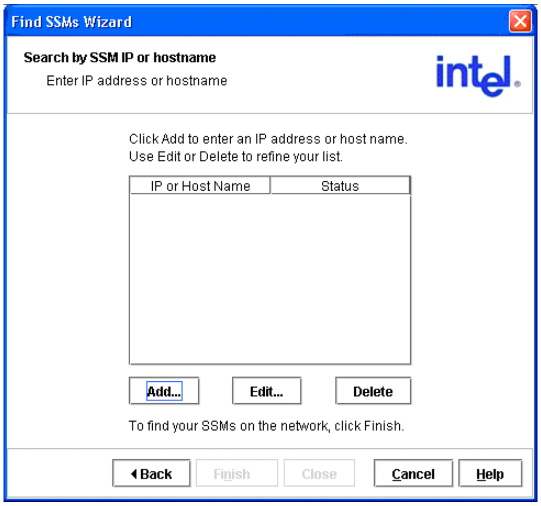

Find SSMs Wizard

**Search by SSM IP or hostname**

Enter IP address or hostname

Click Add to enter an IP address or host name.
Use Edit or Delete to refine your list.

| IP or Host Name | Status |
|---|---|

To find your SSMs on the network, click Finish.

20. At the Add IP or Host Name screen, enter the static IP address of your storage system. Click on “OK” to continue.

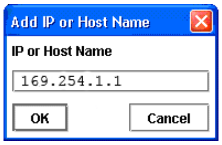

Add IP or Host Name

IP or Host Name: 169.254.1.1

21. Click on “Finish” to search for the specified IP address. The storage system will display a “Found” under “Status” if the unit is correctly detected.

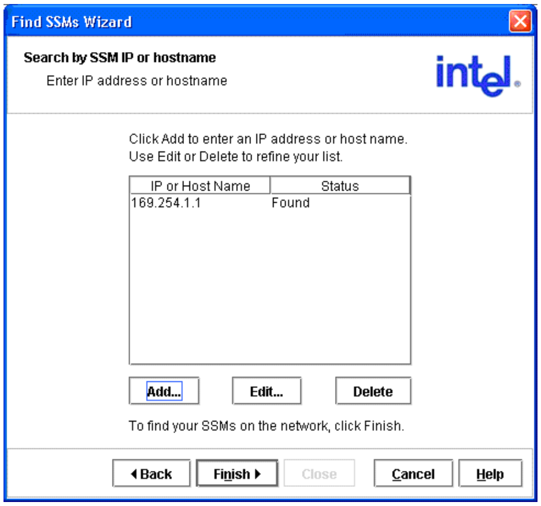

**Search by SSM IP or hostname**

Enter IP address or hostname

Click Add to enter an IP address or host name.
Use Edit or Delete to refine your list.

| IP or Host Name | Status |
|---|---|
| 169.254.1.1 | Found |

To find your SSMs on the network, click Finish.

22. Click on “Close” to return to the Wizard Launch Pad.

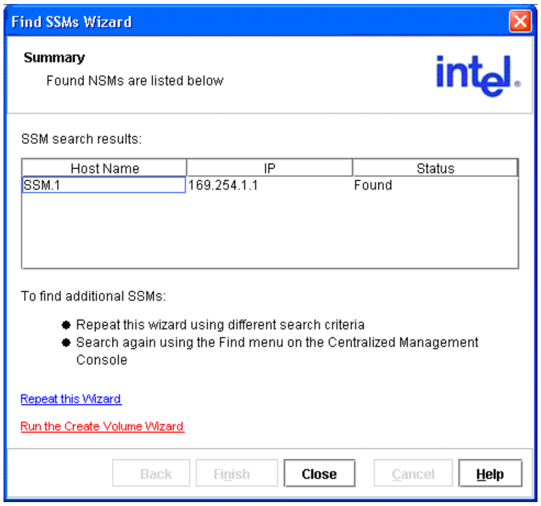

**Summary**

Found NSMs are listed below

SSM search results:

| Host Name | IP | Status |
|---|---|---|
| SSM.1 | 169.254.1.1 | Found |

To find additional SSMs:

- Repeat this wizard using different search criteria
- Search again using the Find menu on the Centralized Management Console

Repeat this Wizard

Run the Create Volume Wizard

23. Select the SSM in the main window by double clicking on Available Storage Modules and then selecting the SSM.

24. Click on Storage Module Tasks and then select Edit Configuration.

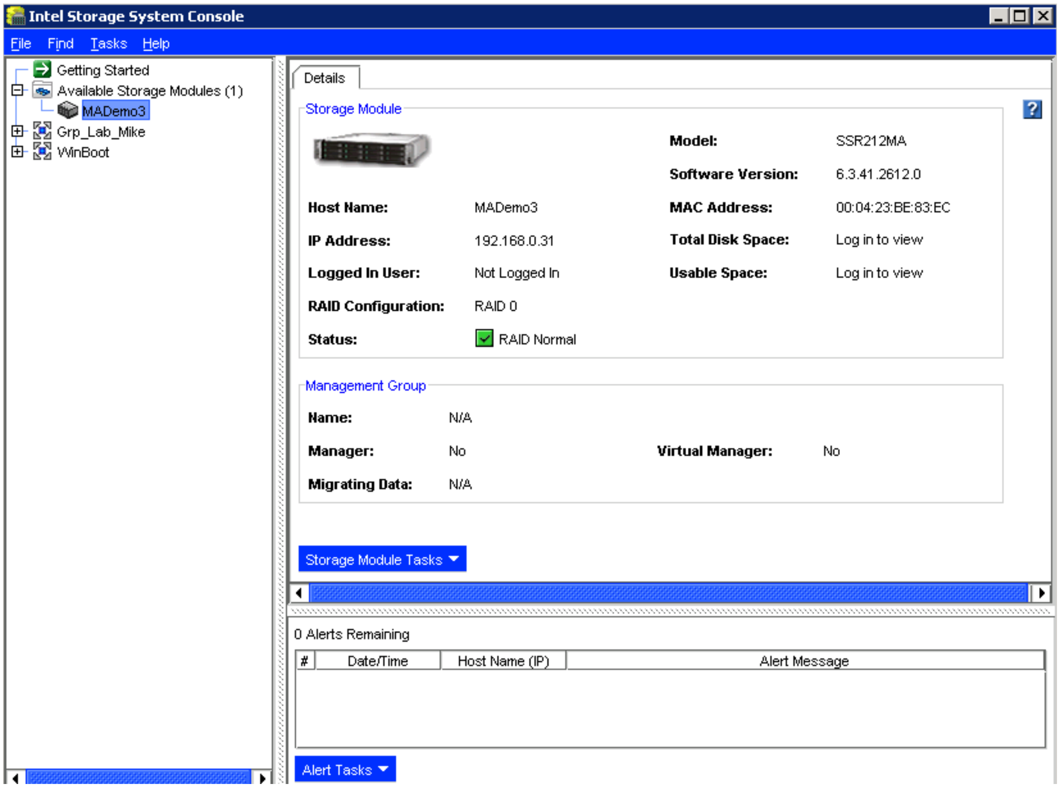

Model: SSR212MA

Software Version: 6.3.41.2612.0

Host Name: MADemo3

MAC Address: 00:04:23:BE:83:EC

IP Address: 192.168.0.31

Total Disk Space: Log in to view

Logged In User: Not Logged In

Usable Space: Log in to view

RAID Configuration: RAID 0

Status: ✓ RAID Normal

Name: N/A

Manager: No

Virtual Manager: No

Migrating Data: N/A

0 Alerts Remaining

| # | Date/Time | Host Name (IP) | Alert Message |
|---|---|---|---|

25. Enter the user name and password for the storage system. The Edit Configuration window should display.

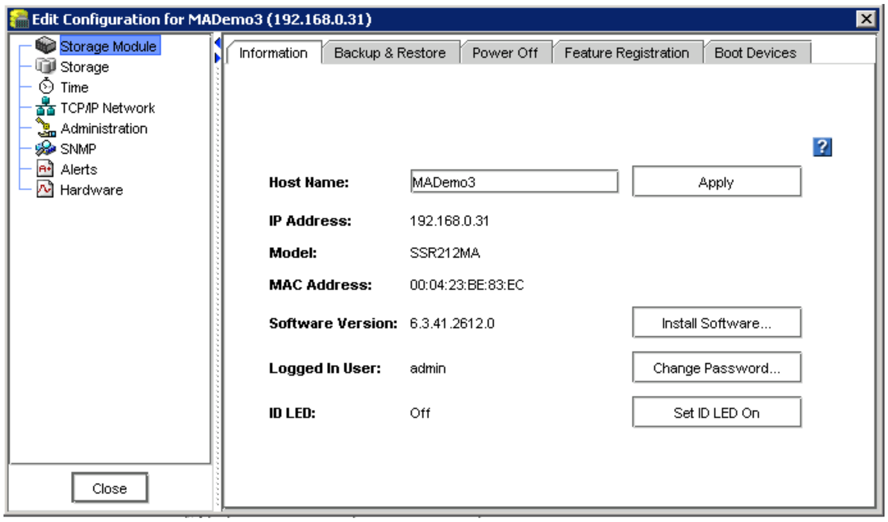

Edit Configuration for MADemo3 (192.168.0.31)

**Host Name:** MADemo3

**IP Address:** 192.168.0.31

**Model:** SSR212MA

**MAC Address:** 00:04:23:BE:83:EC

**Software Version:** 6.3.41.2612.0

**Logged In User:** admin

**ID LED:** Off

26. Select the storage link in the left-hand navigation screen

27. Enter the RAID level and rebuild rate. See "Storage" on page 53 for more information on RAID levels.

28. Click on "Configure RAID" to set up the RAID level.

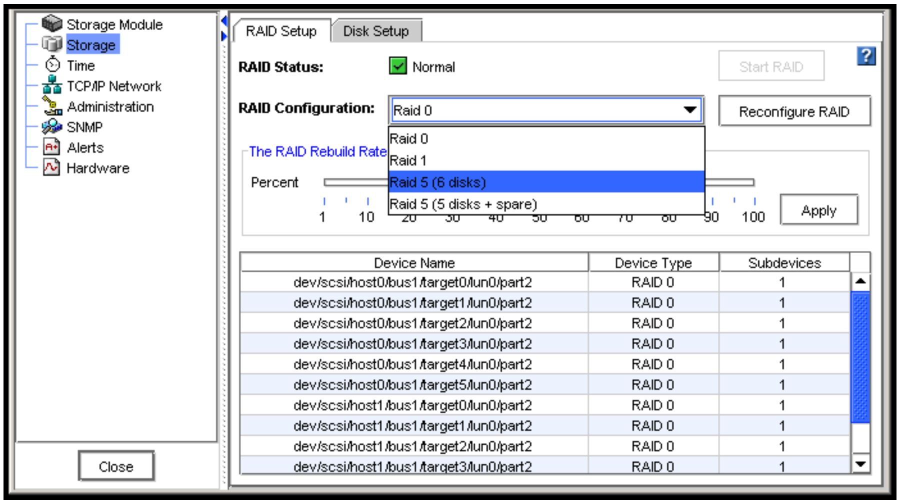

**RAID Status:** Normal

**RAID Configuration:** Raid 0

The RAID Rebuild Rate

Percent

| Device Name | Device Type | Subdevices |
|---|---|---|
| dev/scsi/host0/bus1/target0/lun0/part2 | RAID 0 | 1 |
| dev/scsi/host0/bus1/target1/lun0/part2 | RAID 0 | 1 |
| dev/scsi/host0/bus1/target2/lun0/part2 | RAID 0 | 1 |
| dev/scsi/host0/bus1/target3/lun0/part2 | RAID 0 | 1 |
| dev/scsi/host0/bus1/target4/lun0/part2 | RAID 0 | 1 |
| dev/scsi/host0/bus1/target5/lun0/part2 | RAID 0 | 1 |
| dev/scsi/host1/bus1/target0/lun0/part2 | RAID 0 | 1 |
| dev/scsi/host1/bus1/target1/lun0/part2 | RAID 0 | 1 |
| dev/scsi/host1/bus1/target2/lun0/part2 | RAID 0 | 1 |
| dev/scsi/host1/bus1/target3/lun0/part2 | RAID 0 | 1 |

29. The RAID Status should indicate normal as illustrated in the following figure.

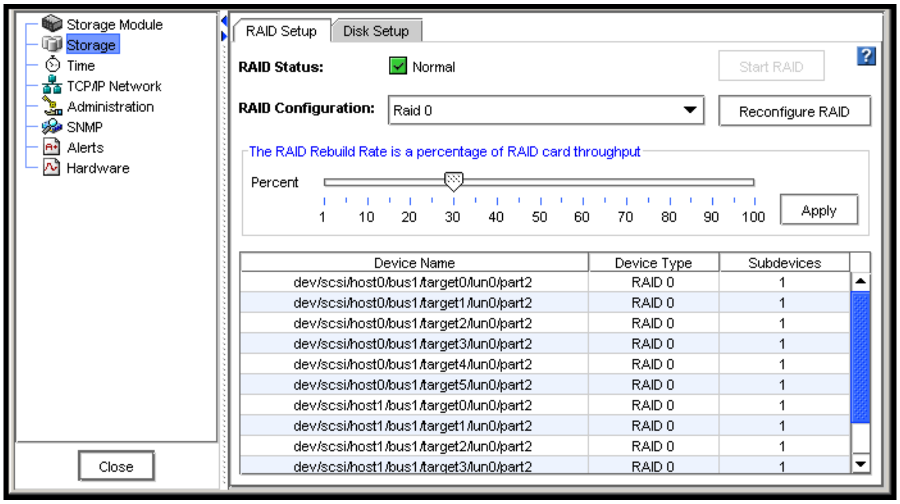

**RAID Status:** ☑ Normal

**RAID Configuration:** Raid 0

The RAID Rebuild Rate is a percentage of RAID card throughput

Percent: 1 10 20 30 40 50 60 70 80 90 100

| Device Name | Device Type | Subdevices |
|---|---|---|
| dev/scsi/host0/bus1/target0/lun0/part2 | RAID 0 | 1 |
| dev/scsi/host0/bus1/target1/lun0/part2 | RAID 0 | 1 |
| dev/scsi/host0/bus1/target2/lun0/part2 | RAID 0 | 1 |
| dev/scsi/host0/bus1/target3/lun0/part2 | RAID 0 | 1 |
| dev/scsi/host0/bus1/target4/lun0/part2 | RAID 0 | 1 |
| dev/scsi/host0/bus1/target5/lun0/part2 | RAID 0 | 1 |
| dev/scsi/host1/bus1/target0/lun0/part2 | RAID 0 | 1 |
| dev/scsi/host1/bus1/target1/lun0/part2 | RAID 0 | 1 |
| dev/scsi/host1/bus1/target2/lun0/part2 | RAID 0 | 1 |
| dev/scsi/host1/bus1/target3/lun0/part2 | RAID 0 | 1 |

30. Click on “Close”.

31. From the Help menu, select Getting Started->Management Groups, Clusters and Volumes Wizard.

32. Click on “Next” to continue.

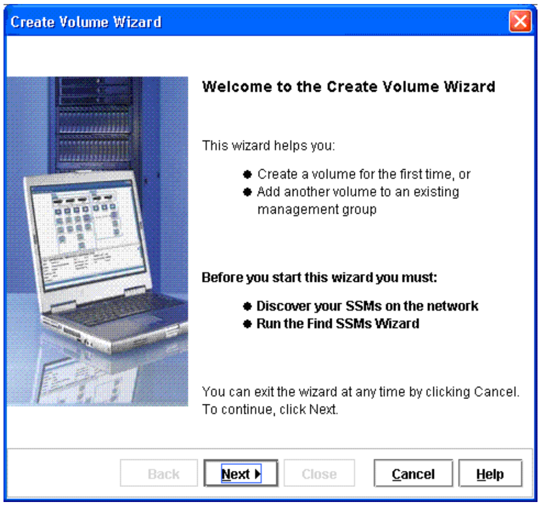

**Welcome to the Create Volume Wizard**

This wizard helps you:

- Create a volume for the first time, or
- Add another volume to an existing management group

**Before you start this wizard you must:**

- **Discover your SSMs on the network**
- **Run the Find SSMs Wizard**

You can exit the wizard at any time by clicking Cancel. To continue, click Next.

33. Select “New Management Group” and click on “Next”. Management groups allow you to manage and configure storage systems as a group. See “Working with Management Groups” on page 169 for additional information on management groups.

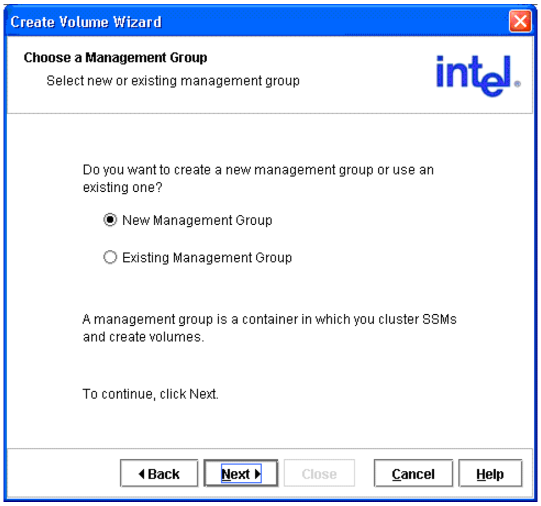

**Choose a Management Group**

Select new or existing management group

Do you want to create a new management group or use an existing one?

● New Management Group

○ Existing Management Group

A management group is a container in which you cluster SSMs and create volumes.

To continue, click Next.

34. Enter the name of your new management group and click on “Next”.

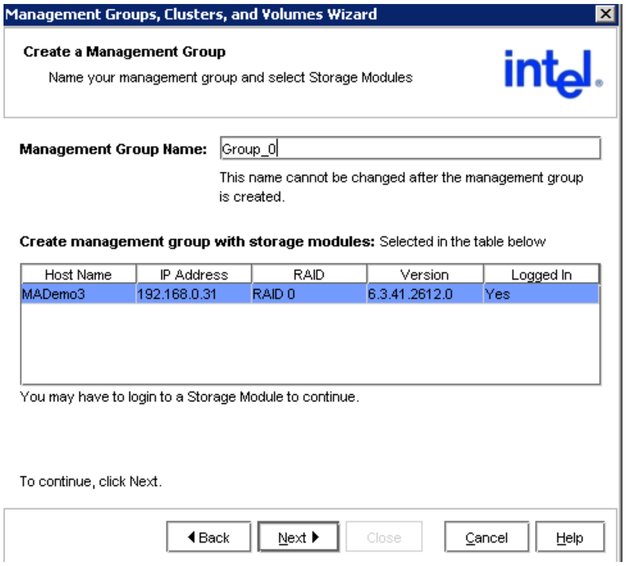

**Create a Management Group**

Name your management group and select Storage Modules

**Management Group Name:** Group_0

This name cannot be changed after the management group is created.

**Create management group with storage modules:** Selected in the table below

| Host Name | IP Address | RAID | Version | Logged In |
|---|---|---|---|---|
| MADemo3 | 192.168.0.31 | RAID 0 | 6.3.41.2612.0 | Yes |

You may have to login to a Storage Module to continue.

To continue, click Next.

35. Enter a cluster name in the “Cluster Name” field and select the SSM. Ensure that the “Use a virtual IP address for this cluster” checkbox is unchecked. Click on “Next” to continue. See “Working with Clusters” on page 201 for additional information on clusters.

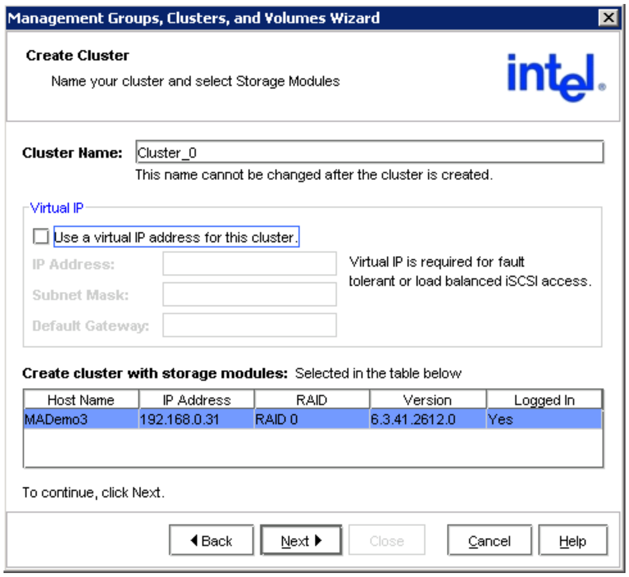

Management Groups, Clusters, and Volumes Wizard

**Create Cluster**

Name your cluster and select Storage Modules

**Cluster Name:** Cluster_0

This name cannot be changed after the cluster is created.

Virtual IP

□ Use a virtual IP address for this cluster.

IP Address:

Subnet Mask:

Default Gateway:

Virtual IP is required for fault tolerant or load balanced iSCSI access.

**Create cluster with storage modules:** Selected in the table below

| Host Name | IP Address | RAID | Version | Logged In |
|---|---|---|---|---|
| MADemo3 | 192.168.0.31 | RAID 0 | 6.3.41.2612.0 | Yes |

To continue, click Next.

36. Enter the “Volume Name”, “Description”, “Size”, “Hard Threshold”, “Soft Threshold”, and “Replication Level” for your storage system. Click on “Finish” to continue. See “Working with Volumes” on page 215 for additional details on these settings.

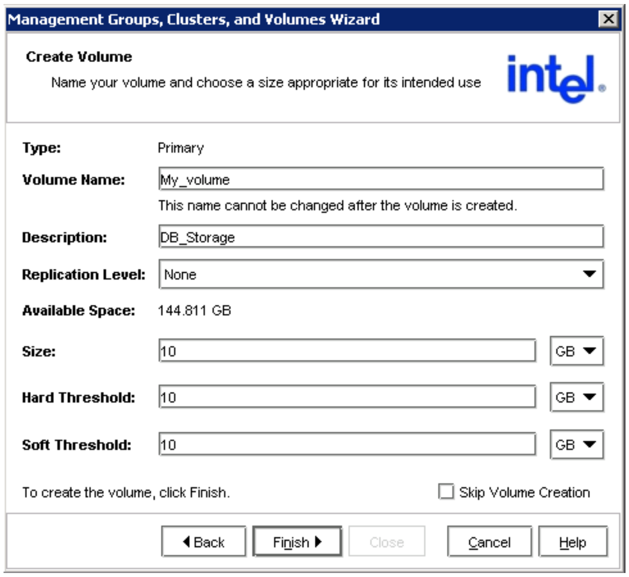

Management Groups, Clusters, and Volumes Wizard

**Create Volume**

Name your volume and choose a size appropriate for its intended use

**Type:** Primary

**Volume Name:** My_volume

This name cannot be changed after the volume is created.

**Description:** DB_Storage

**Replication Level:** None

**Available Space:** 144.811 GB

**Size:** 10 GB

**Hard Threshold:** 10 GB

**Soft Threshold:** 10 GB

To create the volume, click Finish.

□ Skip Volume Creation

37. Select Finish.

38. A summary screen displays showing you that the volume has been successfully created. Click on “Close” to continue.

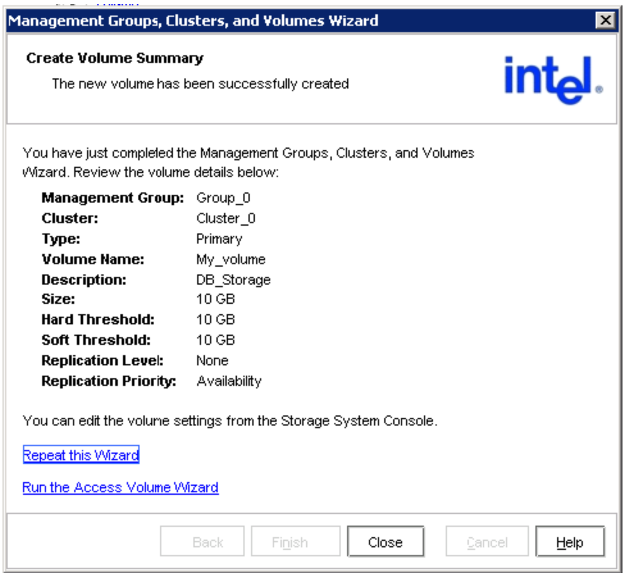

**Create Volume Summary**

The new volume has been successfully created

You have just completed the Management Groups, Clusters, and Volumes Wizard. Review the volume details below:

| | |
|---|---|
| **Management Group:** | Group_0 |
| **Cluster:** | Cluster_0 |
| **Type:** | Primary |
| **Volume Name:** | My_volume |
| **Description:** | DB_Storage |
| **Size:** | 10 GB |
| **Hard Threshold:** | 10 GB |
| **Soft Threshold:** | 10 GB |
| **Replication Level:** | None |
| **Replication Priority:** | Availability |

You can edit the volume settings from the Storage System Console.

Repeat this Wizard

Run the Access Volume Wizard

39. You should see the volume listed in the navigation window. This means that it has been successfully created. Refer to the remainder of this manual for detailed information on managing and configuring your storage system.

# Setting up an Intel® Management Module (IMM) Password for the Intel® Storage System SSR212MA

1. Install `dpcproxy` from the ISM CD that shipped with your Intel® Storage System SSR212MA.

   For Microsoft* Windows*:

   ```
   dpcproxy -install
   ```

   ```
   net start dpcproxy
   ```

2. Bring up an MSDOS window and enter the following commands:

   ```
   c:> telnet localhost 623
   ```

   or

   ```
   dpccli
   ```

   Enter the IP address for your system and press <Enter>.

   For example, `Server: 111.112.113.20`

   Press <Enter> for the user name.

Press <Enter> for the password.

3. Enter the following (all on one line) to set the user name and password for Intel® Management Module Professional Edition:

```
dpccli> set -T BMC/user UserName=YourUserName
Password=YourPassword
```

To permanently set the values, use the `commit` command:

```
dpccli> commit
```

# Wizards

The first time you open the Console, the Wizard Launch Pad opens, shown in Figure 10.

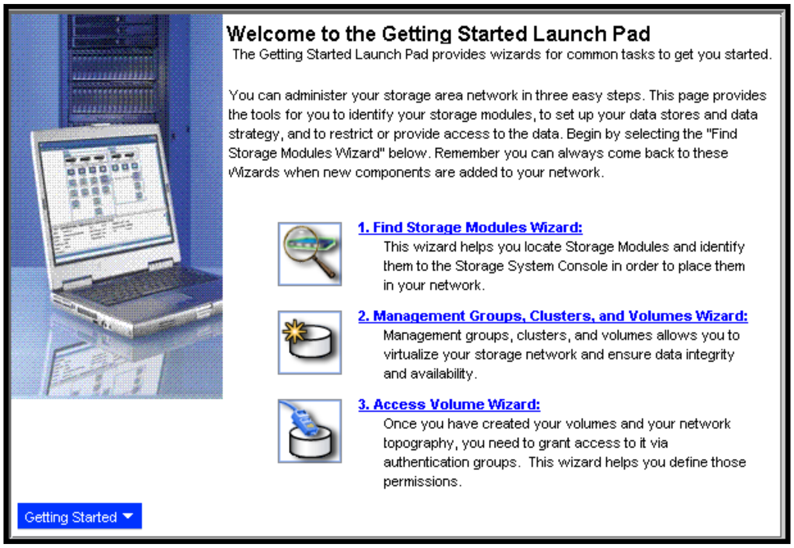

**Welcome to the Getting Started Launch Pad**

The Getting Started Launch Pad provides wizards for common tasks to get you started.

You can administer your storage area network in three easy steps. This page provides the tools for you to identify your storage modules, to set up your data stores and data strategy, and to restrict or provide access to the data. Begin by selecting the "Find Storage Modules Wizard" below. Remember you can always come back to these Wizards when new components are added to your network.

1. Find Storage Modules Wizard:
   This wizard helps you locate Storage Modules and identify them to the Storage System Console in order to place them in your network.

2. Management Groups, Clusters, and Volumes Wizard:
   Management groups, clusters, and volumes allows you to virtualize your storage network and ensure data integrity and availability.

3. Access Volume Wizard:
   Once you have created your volumes and your network topography, you need to grant access to it via authentication groups. This wizard helps you define those permissions.

**Figure 10. Viewing the Wizard Launch Pad**

## Find Storage Module Wizard

The Find Storage Module Wizard guides you through the process for finding SSMs on your network. See also “Finding Storage System Modules on the Network” on page 26.

## Management Groups, Clusters, and Volumes Wizard

The Management Groups, Clusters, and Volumes Wizard guides you through the process for creating management groups, clusters, and volumes.

## Access Volume Wizard

The Access Volume Wizard guides you through the process for configuring client access to your volumes. See also Chapter 15, “Controlling Client Access to Volumes.” .

# Finding Storage System Modules on the Network

After opening the Console, you must find the SSMs you want to manage. Find these SSMs by one of two methods:

- Use a mask to search subnets to find all available SSMs on a network.

  See “Finding by Subnet and Mask” on page 27 for more information about completing the List of Subnets to Search window.

- Enter specific IPs or host names to find individual SSMs.

  See “Finding by Module IP or Host Name” on page 29 for more information about completing the IP and Host Name List window.

Once you have found SSMs the first time, the Find settings are saved. Every time you open the Console, the search takes place and the SSMs are listed in the navigation window.

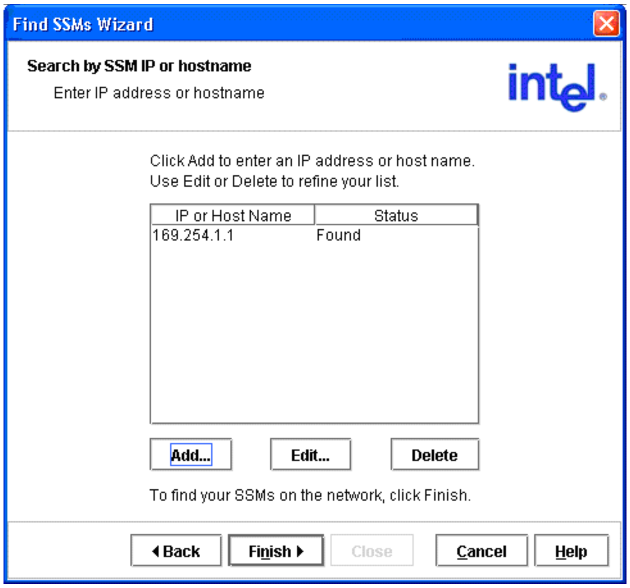

Find SSMs Wizard

Search by SSM IP or hostname

Enter IP address or hostname

Click Add to enter an IP address or host name.
Use Edit or Delete to refine your list.

| IP or Host Name | Status |
|---|---|
| 169.254.1.1 | Found |

To find your SSMs on the network, click Finish.

Figure 11. SSMs Found Message

**Note:** *You can control which SSMs appear in the Network View by entering only specific IPs or Host Names in the IP and Host Name List window. Then, when you open the Console, only those IPs or Host Names will appear in the Network View.*

### Modules Not Found

If the network has a lot of traffic, or if a module is busy reading or writing data, it may not be found when a search is performed. If a module is not found, try the following steps to find it.

1. Search again using the Find menu.
2. If you have searched by Subnet and Mask, try using the Find by IP or Host Name search.
3. If searching again does not work, try the following:
   - Check the physical connection of the module.
   - Wait a few minutes and try the search again. If activity to the module was high, the module might not have responded to the search.

***Note:*** *Other problems can occur that prevent connection, such as a bad cable connection.*

# Finding by Subnet and Mask

Find all the SSMs on the network by searching subnets with masks. To do this: Click the Find menu and click By Subnet and Mask. The List of Subnets to Search window opens, shown in Figure 12.

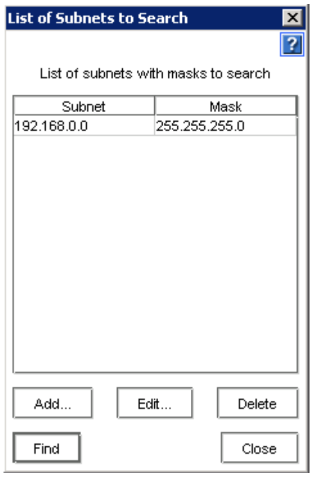

List of subnets with masks to search

| Subnet | Mask |
|---|---|
| 192.168.0.0 | 255.255.255.0 |

**Figure 12. Using Subnet and Mask to Search**

## Adding Subnets and Masks

1. Click Add to enter a subnet and mask. The Add Subnet and Mask window opens.
2. Type in the Subnet.
3. Select the appropriate mask from the list.
4. Click OK to close the Add Subnet and Mask window.

5. Click Find. The Active Search window opens, tracking the search process. When the search is complete, the Active Search window closes. The Console window opens, listing all the SSMs that were found on the network.
6. Click OK to close the Console window.
7. Click Close on the List of Subnets to Search window. The modules appear in the Configuration Categories under Available Storage Modules. Click on Available Storage Modules to view a list of available SSMs.

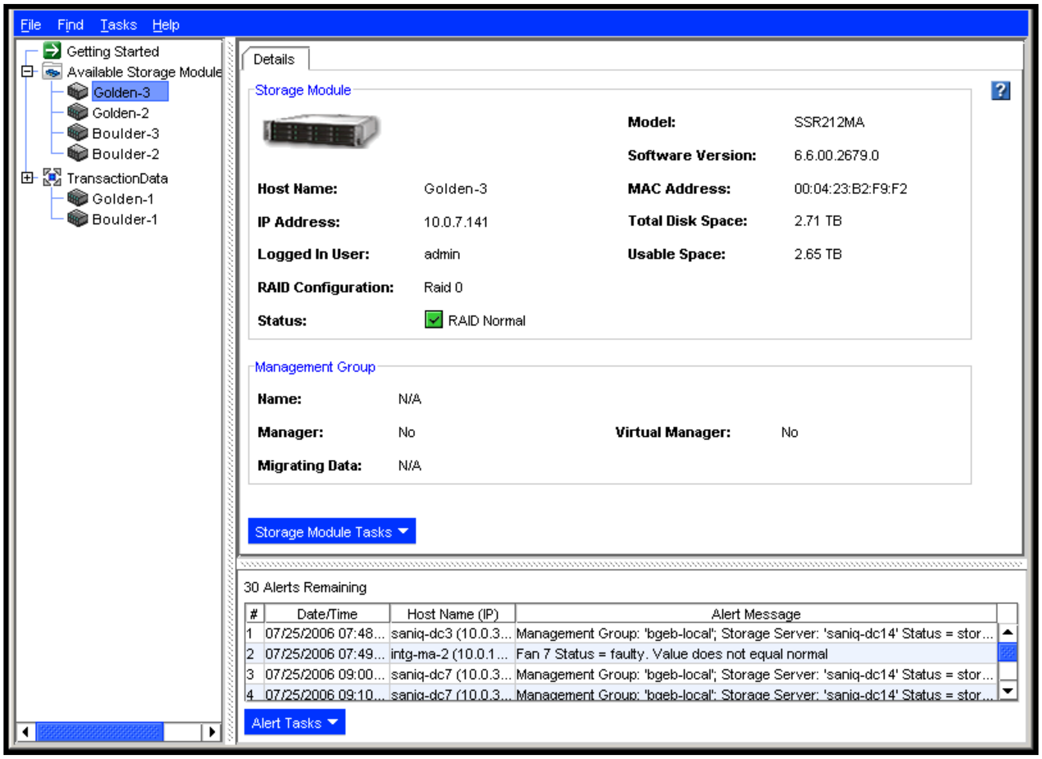

Storage Module

Model: SSR212MA

Software Version: 6.6.00.2679.0

Host Name: Golden-3

MAC Address: 00:04:23:B2:F9:F2

IP Address: 10.0.7.141

Total Disk Space: 2.71 TB

Logged In User: admin

Usable Space: 2.65 TB

RAID Configuration: Raid 0

Status: ✓ RAID Normal

Management Group

Name: N/A

Manager: No

Virtual Manager: No

Migrating Data: N/A

30 Alerts Remaining

| # | Date/Time | Host Name (IP) | Alert Message |
|---|---|---|---|
| 1 | 07/25/2006 07:48... | saniq-dc3 (10.0.3... | Management Group: 'bgeb-local'; Storage Server: 'saniq-dc14' Status = stor... |
| 2 | 07/25/2006 07:49... | intg-ma-2 (10.0.1... | Fan 7 Status = faulty. Value does not equal normal |
| 3 | 07/25/2006 09:00... | saniq-dc7 (10.0.3... | Management Group: 'bgeb-local'; Storage Server: 'saniq-dc14' Status = stor... |
| 4 | 07/25/2006 09:10... | saniq-dc7 (10.0.3... | Management Group: 'bgeb-local'; Storage Server: 'saniq-dc14' Status = stor... |

**Figure 13. SSMs in the Network View Pane**

**Note:** *The subnet and mask are saved in the list. Every time you open the Console, the search takes place automatically and all SSMs on the network are listed in the Network View. See "Deleting Subnets and Masks" on page 29 if you want to disable this search.*

## Editing Subnets and Masks

Change the subnets and masks used to search for modules.

1. Click the Find menu.
2. Click By Subnet and Mask. The List of Subnets to Search window opens.
3. Select the subnet you want to edit.
4. Click Edit. The subnet and mask window opens.
5. Change information as necessary.

6. Click OK.

## Deleting Subnets and Masks

You can delete a subnet and mask from the search list if you remove modules from that network, or if you do not want to view those modules in the Network View.

1. Click the Find menu.
2. Click By Subnet and Mask. The List of Subnets to Search window opens.
3. Select the subnet and mask to delete.
4. Click Delete. A confirmation message opens.
5. Click OK.
6. Click Close.

# Finding by Module IP or Host Name

Identify SSMs by listing module IP or host names and searching for those SSMs. You can connect to one specific IP or host name, or find all the SSMs in the list.

## Network Configuration and Find by IP or Host Name

The way your network is configured may affect the results of finding SSMs by IP address. An example of the effect of network configuration is detailed below.

- You configure both NICs in an SSM (eth0 and eth1).
- The NICs are on separate subnets.
- You open the Console on a system on the same subnet as the eth0 NIC on the SSM.
- The Console Find function is set to Module IP or Host Name using only the IP address of the eth1 NIC.

The SSM is discovered and appears in the Console. However, the IP address returned to the Console is that of the eth0 NIC. The eth1 IP address is not discovered.

This is normal behavior controlled by the way networking is configured. The SSM receives the broadcast and replies through eth0, regardless of which NIC received the broadcast. The Console picks up the address from the packet that was sent through eth0 and displays it as representative of the SSM.

## To Find by IP or Host Name

1. If this is the first time you have opened the Console, select Find Storage Module Wizard, then click OK.

   **or**

   Click the Find menu and click By Storage Module IP or Host Name.

   The IP and Host Name List window opens, as shown in Figure 14.

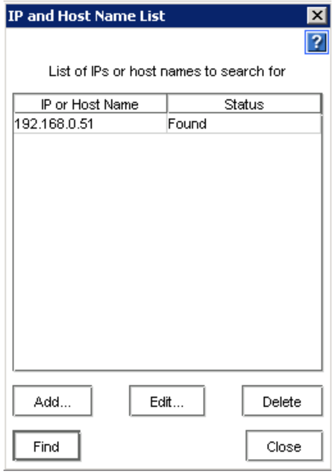

| IP or Host Name | Status |
|---|---|
| 192.168.0.51 | Found |

**Figure 14. Using IP or Host Name to Search**

## Adding IPs or Host Names

Use the following steps to add specific IP addresses or host names to the list.

1. Click Add. The Add IP or Host Name window opens.
2. Type in the IP or Host Name for the module.
3. Click OK.
4. Repeat steps 1 through 3 for each module you want to find.
5. Click Find.

## Editing the IP or Host Name in the Search List

Use the following steps to change the IP or Host Name of an SSM in the list used to search for modules.

1. Click the Find menu.
2. Click By Storage Module IP or Host Name. The IP and Host Name List window opens, shown in Figure 14.
3. Select the IP/Host Name you want to edit.
4. Click Edit. The Edit IP or Host Name window opens.
5. Change the necessary information.
6. Click OK to return to the IP and Host Name List window.

## Deleting the IP or Host Name in the Search List

Once you enter an IP or host name in the IP and Host Name List, that entry is saved. Every time you open the Console, a search for all the IPs and host names occurs.

You can delete an IP from the list if you no longer want to search for that SSM.

1. Click the Find menu.
2. Click By Storage Module IP or Host Name. The IP and Host Name List window opens, shown in Figure 14.
3. Select the IP/Host Name to delete.
4. Click Delete. A confirmation message opens.
5. Click OK. The IP or host name is removed from the list.
6. Click Close.

# 2 Working with Storage System Modules

The SSM configuration  window opens when you log into an individual SSM. From the configuration window you have access to all the configuration tasks for individual SSMs.

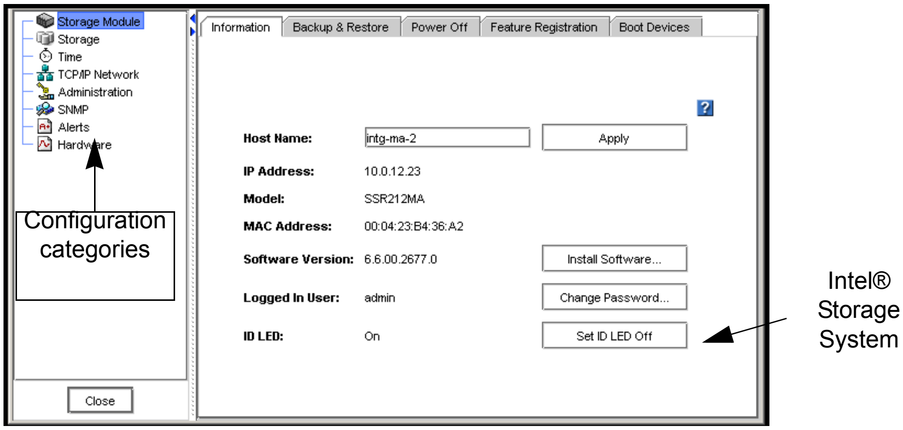

Host Name: intg-ma-2

IP Address: 10.0.12.23

Model: SSR212MA

MAC Address: 00:04:23:B4:36:A2

Software Version: 6.6.00.2677.0

Logged In User: admin

ID LED: On

**Figure 15. SSM Configuration Window for the Intel® Storage System SSR212MA**

## Configuration Categories

The left pane lists the configuration categories. The right pane contains a set of tabs, which you use to configure different functions, for each specific category. The configuration categories are described below.

- **Storage Module** - Use the module category to change the host name and login password for the SSM. You can also install software, backup and restore the Storage System Software  configuration, reboot or power off the SSM, register the SSM for add-on features, and activate the flash cards used for booting the SSM.
- **Storage** - Manage RAID and the individual disks in the SSM.
- **Time** - Configure the time zone and set the date and time on the SSM. The date and time settings are used to create a time stamp on volumes and snapshots. The date and time settings also affect schedules for snapshots and remote copies.
- **TCP/IP Network** - For each SSM you can configure and manage the network settings, including TCP/IP interfaces, DNS servers, and the routing table.
- **SSM Administration** - The SSM comes configured with 2 default groups and 2 default users. All administrative users and groups are added and managed locally.
- **SNMP** - Monitor the SSM using an SNMP Agent. You can also enable SNMP traps.

- **Alerts** - Configure active monitoring settings of selected monitored variables and notification methods for receiving alerts.
- **Hardware** - Use the hardware category to run hardware diagnostic tests, to view current hardware status and configuration information, and to save log files.

# Logging In to the SSM

After finding all the SSMs on the network you must log in to each SSM individually to configure, modify or monitor the functions of that SSM.

1. Select the SSM in the navigation window.
2. Click Storage Module Tasks.
3. Select Log in, shown in Figure 16.

   The Log In window opens, shown in Figure 17.

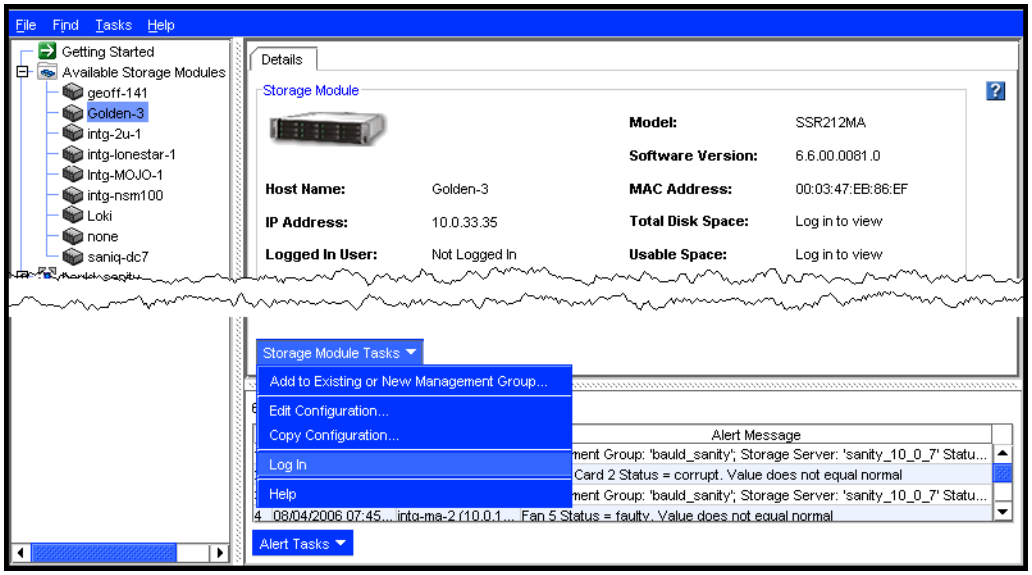

Model: SSR212MA
Software Version: 6.6.00.0081.0
Host Name: Golden-3
MAC Address: 00:03:47:EB:86:EF
IP Address: 10.0.33.35
Total Disk Space: Log in to view
Logged In User: Not Logged In
Usable Space: Log in to view

ment Group: 'bauld_sanity'; Storage Server: 'sanity_10_0_7' Statu...
Card 2 Status = corrupt. Value does not equal normal
ment Group: 'bauld_sanity'; Storage Server: 'sanity_10_0_7' Statu...
4 08/04/2006 07:45... intg-ma-2 (10.0.1... Fan 5 Status = faulty. Value does not equal normal

**Figure 16. Logging in from Storage Module Tasks**

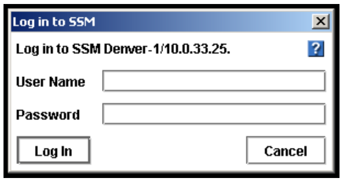

Log in to SSM Denver-1/10.0.33.25.

User Name

Password

**Figure 17. Logging in to an SSM**

4. Type the User Name and Password.
5. Click Log In.

   Alternatively, you may log in to an SSM by one of three other methods:

- By right-clicking an SSM in the navigation window and selecting Log In
- By selecting the Available Storage Modules pool in the navigation window and then right-clicking a storage module name in the tab window and selecting Log In.
- By selecting the SSM in the navigation window and then double-clicking the SSM icon on the Details tab.

## Logging In to Additional SSMs

Once you are logged in to an SSM, you can log in automatically to more SSM.s. Subsequent SSMs must be configured with the same user name and password. Log in by double-clicking those SMMs.

If you try to log in to an SSM that uses a different user name or password, the Log In window opens

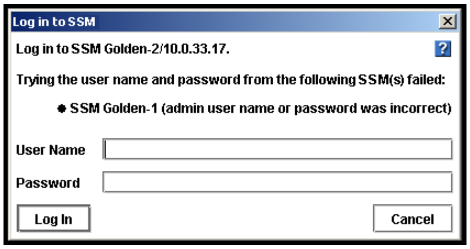

Log in to SSM Golden-2/10.0.33.17.

Trying the user name and password from the following SSM(s) failed:

- SSM Golden-1 (admin user name or password was incorrect)

User Name

Password

**Figure 18. Automatic Log In Fails because SSM User Name and Password are Different**

1. Type the correct User Name and Password.
2. Click Log In.

## Logging Out of the SSM

Log out to prevent access to an SSM without closing the Console. This provides security if you are leaving the management workstation but do not want to close the Console.

1. Select an SSM in the navigation window.
2. Click Storage Module Tasks in the tab window.
3. Select Log Out.

Alternately, you may log out of an SSM by one of three other methods:

- By right-clicking an available SSM in the navigation window and selecting Log Out.
- By selecting the Available Storage Modules pool in the navigation window and then right-clicking an SSM in the tab window and selecting Log Out.

- By selecting Tasks, Storage Module, and Log Out.

***Note:*** *If you are logged in to multiple SSMs, you need to log out of each SSM individually.*

## Opening the Edit Configuration Window

Open the Edit Configuration window to configure the SSM.

1. Select an SSM in the navigation window.
2. Click Storage Module Tasks in the tab window.
3. Select Edit Configuration.

Alternately, you may edit the configuration of an SSM by one of four other methods:

- By double-clicking the SSM icon in the Details tab.
- By right-clicking an available SSM in the navigation window and selecting Edit Configuration.
- By selecting the Available Storage Modules pool in the navigation window and then right-clicking an SSM in the tab window and selecting Edit Configuration.
- By selecting Tasks, Storage Module, and Edit Configuration.

## Closing the Edit Configuration Window

Clicking Close on the Edit Configuration window closes the Edit Configuration window and leaves you logged in to the SSM.

# Module Configuration Overview

The module configuration category provides access to detailed information about the SSM, backing up and restoring SSM configuration files, the software reboot or power off function, boot devices (if present) and feature registration.

The SSM information category window for the Intel® Storage System SSR212MA is shown in the following figure.

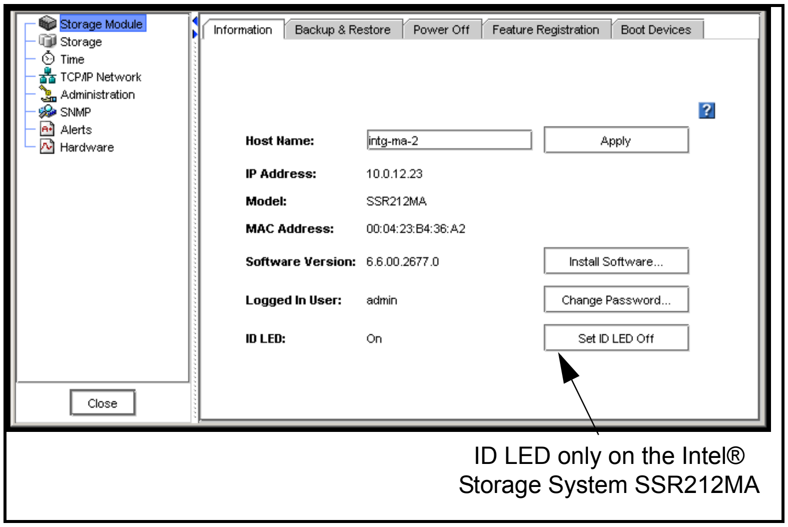

Host Name: intg-ma-2

IP Address: 10.0.12.23

Model: SSR212MA

MAC Address: 00:04:23:B4:36:A2

Software Version: 6.6.00.2677.0

Logged In User: admin

ID LED: On

ID LED only on the Intel®
Storage System SSR212MA

Figure 19. Viewing the Module Configuration Category

# Changing the SSM Host Name

Change the host name of the SSM on the Information tab.

The SSM arrives configured with a default host name.

1. Log in to the SSM.
2. On the Information tab, click the Host Name field and type the new name.
3. Click Apply.

   A confirmation message opens.
4. Click OK.

***Note:*** *Add the host name and IP pair to the host name resolution methodology employed in your environment, e.g., DNS or WINS.*

# Changing Passwords

Change the password for the user who is logged in to an SSM on the Information tab.

1. Log in to the SSM.
2. On the Information tab, click Change Password.

   The Change Password window opens.
3. Type in the User Name and Old Password.

4. Type in the New Password.
5. Retype the New Password for confirmation.
6. Click OK.

Change any other user's password in the SSM Administration configuration category.

# Locating the Module in a Rack (Intel® Storage System SSR212MA only)

The Set ID LED On turns on lights on the physical module so that you can find that SSM in a rack.

1. Log in to the SSM.
2. On the Information tab, click Set ID LED On.

   The ID LED on the left front of the SSM illuminates a bright blue. Another ID LED is located on the back of the SSM on the right side under the empty slot.

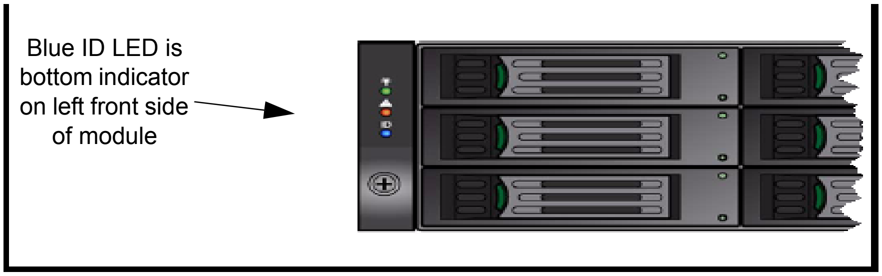
Blue ID LED is bottom indicator on left front side of module

**Figure 20. Viewing ID LED Indicator on Front of SSM**

When you click Set ID LED On, the status changes to On and the button changes to Set ID LED Off.

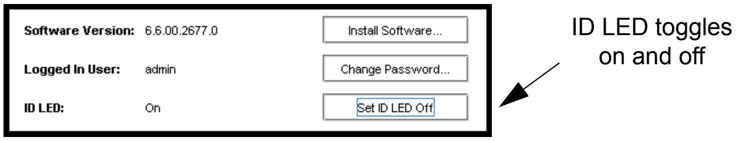
**Software Version:** 6.6.00.2677.0

**Logged In User:** admin

**ID LED:** On

**Figure 21. ID LED Indicator**

3. Click Set ID LED Off when you are finished.

   The LED on the module turns off and the button returns to Set ID LED On.

# Upgrading  the Storage System Software

When you upgrade the Storage System Software, the version number changes. View the current software version in the Edit Configuration window, Storage Module category Information tab. The version number is also displayed when the SSMs are selected in the navigation window.

**Prerequisites**

- Stop any applications that are accessing volumes that reside on the SSM you are upgrading.

To check for available upgrades to download, go to the Intel support site at http://www.intel.com/support/motherboards/server/SSR212MA.

## Copying the Upgrade Files from CD or FTP Site

Upgrade the Storage System Software on the SSM when an upgrade or a patch is released. The Storage System Software upgrade/installation takes about 10 to 15 minutes, including the SSM reboot.

***Note:*** *The 88M must contain both boot flash cards in order to upgrade the storage system software.*

Download the upgrade file from the Intel support web site at http://www.intel.com/support/motherboards/server/SSR212MA or from a CD.

## Upgrading the SSM

You can install upgrades on SSMs individually, which is recommended. If you are upgrading multiple SSMs that are not in a management group, you can upgrade them simultaneously.

1. Log in to the first SSM you want to upgrade.
2. Click Storage Module Tasks and select Edit Configuration

3. On the Information tab, click Install Software.

   The Install Software window opens.

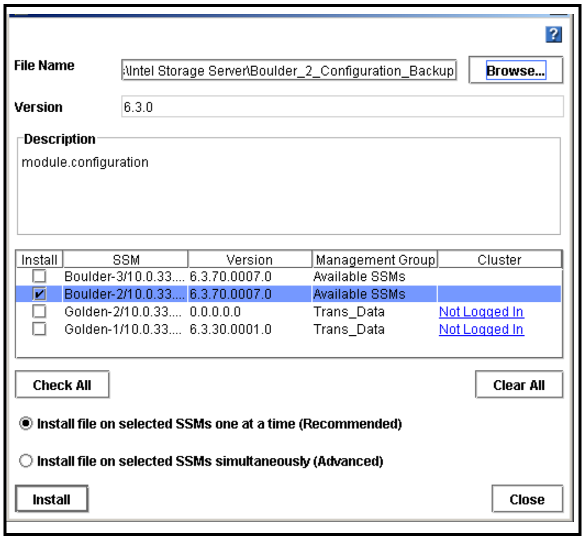

File Name: :\Intel Storage Server\Boulder_2_Configuration_Backup

Version: 6.3.0

Description: module.configuration

| Install | SSM | Version | Management Group | Cluster |
|---|---|---|---|---|
| ☐ | Boulder-3/10.0.33.... | 6.3.70.0007.0 | Available SSMs | |
| ☑ | Boulder-2/10.0.33.... | 6.3.70.0007.0 | Available SSMs | |
| ☐ | Golden-2/10.0.33.... | 0.0.0.0 | Trans_Data | Not Logged In |
| ☐ | Golden-1/10.0.33.... | 6.3.30.0001.0 | Trans_Data | Not Logged In |

● Install file on selected SSMs one at a time (Recommended)

○ Install file on selected SSMs simultaneously (Advanced)

Figure 22. Upgrading the SSM Software

4. From the list, select the SSM that you want to upgrade. Select multiple SSMs to upgrade from the list.
5. Select Install file on selected SSMs one at a time (Recommended).
6. Click Browse to navigate to the folder on the Console computer where you copied the upgrade or patch file.

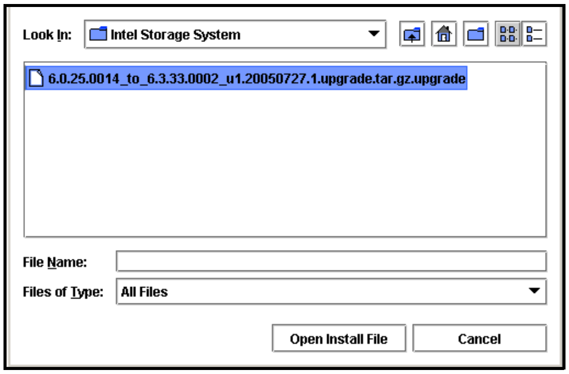

Look In: Intel Storage System

6.0.25.0014_to_6.3.33.0002_u1.20050727.1.upgrade.tar.gz.upgrade

File Name:

Files of Type: All Files

Figure 23. Browsing for Upgrade or Patch File

7. Select the file and click Open Install File.

   Focus returns to the Install Software window. When the file name is present, the Install button becomes enabled.

8. Click Install.

   The install status window opens. Status messages scroll on the window. These messages can be saved to a file.

   — [Optional] After the installation completes, click Save To File and choose a name and location for the file.

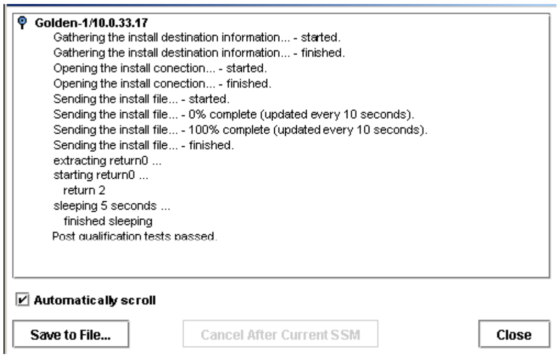

```
Golden-1/10.0.33.17
    Gathering the install destination information... - started.
    Gathering the install destination information... - finished.
    Opening the install conection... - started.
    Opening the install conection... - finished.
    Sending the install file... - started.
    Sending the install file... - 0% complete (updated every 10 seconds).
    Sending the install file... - 100% complete (updated every 10 seconds).
    Sending the install file... - finished.
    extracting return0 ...
    starting return0 ...
      return 2
    sleeping 5 seconds ...
      finished sleeping
    Post qualification tests passed.
```

Figure 24. Upgrade Status Messages

   After the installation completes, the system reboots. After the system comes back online, it conducts a post-install qualification. After the system passes the post-install qualification, the upgrade process is complete.

9. Click Close when the installation has completed.

# Backup and Restore of SSM Configuration

Backup and restore provides the capability to save the SSM configuration file for use in case of an SSM failure. When you back up an SSM configuration, all of the configuration information about the SSM is stored in a file on the computer where the Console is installed. If an SSM failure occurs, you can restore the configuration file to a new SSM. The new SSM will be configured identically to the SSM when it was backed up.

Backing up the configuration file for an SSM does not save information about the configuration of any management groups, clusters, volume lists or authentication groups that the SSM belongs to. It also does not back up license key entries for registered features. To preserve a record of management group configuration information and license keys, see "Backing Up a Management Group Configuration" on page 181.

***Note:*** *Back up the SSM configuration every time you change SSM settings. This ensures that you can restore an SSM to its most recent configuration.*

# Backing Up the SSM Configuration File

Use Backup to save the SSM configuration file to a directory on your local machine.

1. Click Backup from the Backup and Restore tab.

   The Save window opens.

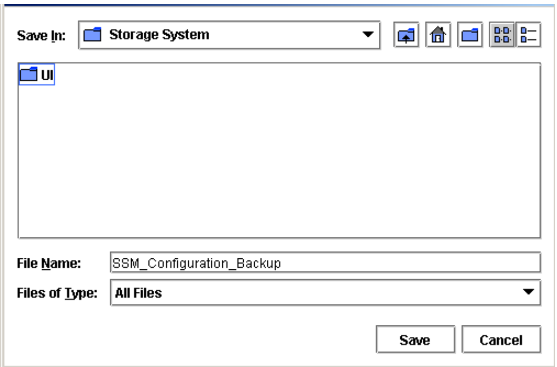

Save In: Storage System

UI

File Name: SSM_Configuration_Backup

Files of Type: All Files

**Figure 25. Backing up the SSM Configuration File**

2. Navigate to a folder on the Console computer to contain the SSM configuration backup file.
3. Enter a meaningful name for the backup file or accept the default name (SSM_Configuration_Backup).

***Note:*** *The configuration files for all SSMs that you back up are stored on the computer running the Console. If you back up multiple SSMs, be sure to give each SSM configuration file a unique and descriptive name. This will make it easier to locate the correct configuration file if you need to restore the configuration of a specific SSM.*

4. Click Save.

# Restoring the SSM Configuration from a File

Use Restore to restore the configuration of an SSM.

1. On the Backup and Restore tab, click Restore.

   The Restore SSM window opens.

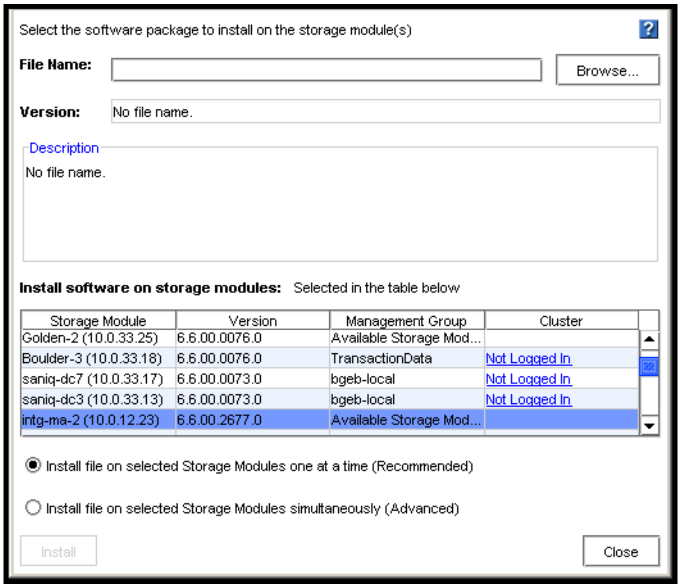

Select the software package to install on the storage module(s)

**File Name:**

**Version:** No file name.

Description

No file name.

**Install software on storage modules:** Selected in the table below

| Storage Module | Version | Management Group | Cluster |
|---|---|---|---|
| Golden-2 (10.0.33.25) | 6.6.00.0076.0 | Available Storage Mod... | |
| Boulder-3 (10.0.33.18) | 6.6.00.0076.0 | TransactionData | Not Logged In |
| saniq-dc7 (10.0.33.17) | 6.6.00.0073.0 | bgeb-local | Not Logged In |
| saniq-dc3 (10.0.33.13) | 6.6.00.0073.0 | bgeb-local | Not Logged In |
| intg-ma-2 (10.0.12.23) | 6.6.00.2677.0 | Available Storage Mod... | |

● Install file on selected Storage Modules one at a time (Recommended)

○ Install file on selected Storage Modules simultaneously (Advanced)

Figure 26. Restoring the SSM Configuration File

2. In the table, select the SSM you want to restore.

   You can select multiple SSMs to restore from the list.

3. Select Install file on selected SSMs one at a time (Recommended).

4. Click Browse to navigate to the folder on the Console computer where the configuration backup file is saved.

5. Select the file to restore and click Open Backup File.

6. Review the version and description to ensure you are restoring the correct file.

7. Click Install.

   The Install Status window opens. When the restoration is complete, the Save to File and Close buttons become enabled.

   — To save a log file of the restore operation before rebooting, click Save to File.

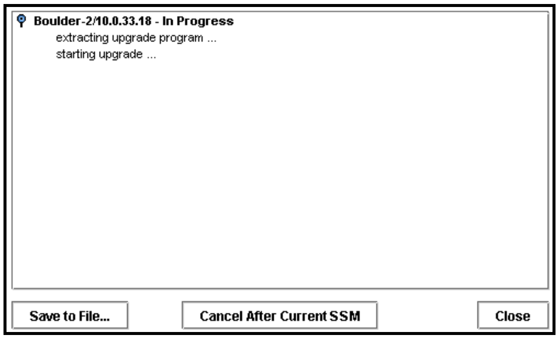

```
Boulder-2/10.0.33.18 - In Progress
    extracting upgrade program ...
    starting upgrade ...
```

Figure 27. Restoring the SSM Configuration File

8. Click Close to finish restoring the configuration.

   The SSM reboots and the configuration is restored to the identical configuration as that in the backup file.

## Completing the Restore

After you restore the SSM configuration from a file, up to three manual configuration steps are required:

- You must manually configure RAID on the SSM.
- You must manually add network routes after the restoration. Restoring a configuration file from one SSM to a second SSM does not restore network routes that were configured on the SSM.
- If you restore multiple SSMs from one configuration file, you must manually change the IP address on the additional SSMs. For example, if you back up the configuration of an SSM with a static IP address, and then restore that configuration to a second SSM, the second SSM will have the same IP address.

# Rebooting the SSM

Reboot the SSM from the Console without powering off. Set the amount of time before the reboot begins to ensure that any activity to the module has stopped.

1. Log in to the SSM.
2. Click Storage Module Tasks and select Edit Configuration.
3. Select the Power Off tab.

   The Power Off window opens.

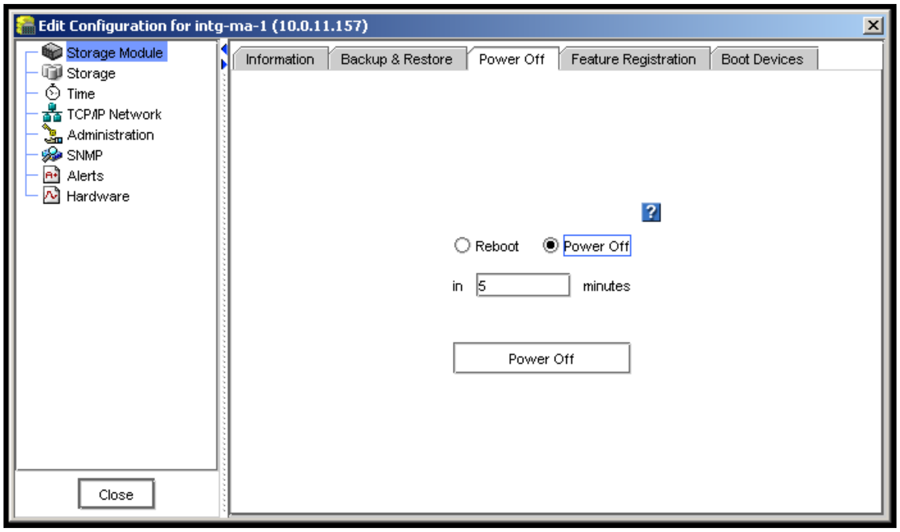
Edit Configuration for intg-ma-1 (10.0.11.157)

○ Reboot ● Power Off

in 5 minutes

**Figure 28. Shutting Down or Rebooting the SSM**

4. Select Reboot.
5. In the minutes field, type the number of minutes before the reboot should begin.

   You can enter any whole number greater than or equal to 0. If you enter 0 the SSM will reboot as soon as you complete step 7.
6. Click Reboot.

   A confirmation message appears.
7. Click OK.

   The SSM reboots in the specified number of minutes. When reboot actually begins, the SSM disappears from the navigation window. The reboot takes 3 to 4 minutes.
8. Search for the SSM to reconnect the Console to the SSM once it has finished rebooting. See "Finding Storage System Modules on the Network" on page 26.

# Powering Off the SSM

Powering off the SSM through the Console physically powers off the SSM. The Console controls the power down process so that data are protected.

Powering off an individual SSM is appropriate for servicing or moving that SSM. However, if you want to shut down more than one SSM in a management group, you should shut down the management group instead of individually powering off the SSMs in that group. See “Shutting Down a Management Group” on page 184.

1. Log in to the SSM.
2. Select the Power Off tab.

   The Power Off window opens.

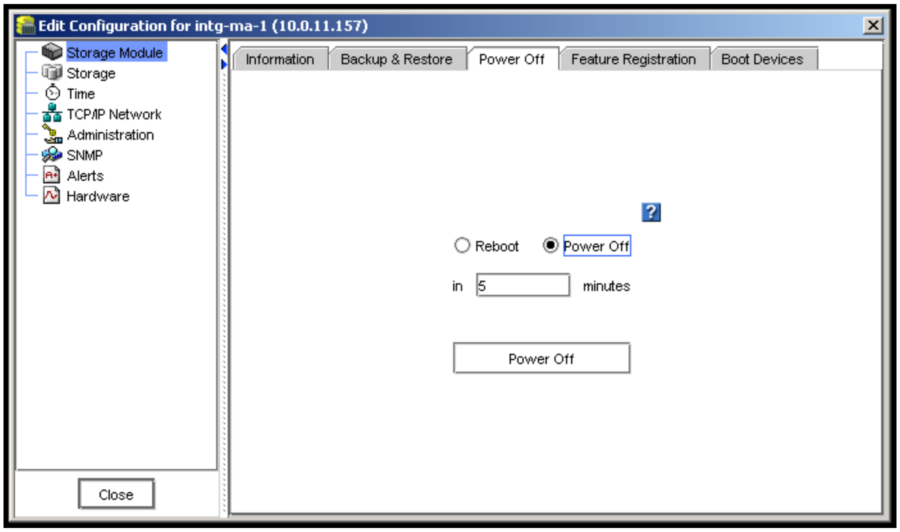

○ Reboot ● Power Off

in 5 minutes

Figure 29. Powering Off the SSM

3. Select Power Off.

   The button changes to Power Off.
4. In the minutes field, type the number of minutes before the powering off should begin. Enter any whole number greater than or equal to 0. If you enter 0, the SSM powers off as soon as you complete step 5.

5. Click Power Off.

   A confirmation message appears.

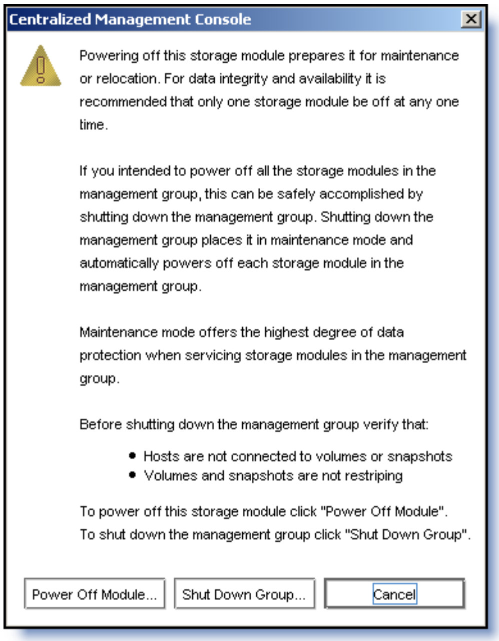

Powering off this storage module prepares it for maintenance or relocation. For data integrity and availability it is recommended that only one storage module be off at any one time.

If you intended to power off all the storage modules in the management group, this can be safely accomplished by shutting down the management group. Shutting down the management group places it in maintenance mode and automatically powers off each storage module in the management group.

Maintenance mode offers the highest degree of data protection when servicing storage modules in the management group.

Before shutting down the management group verify that:

- Hosts are not connected to volumes or snapshots
- Volumes and snapshots are not restriping

To power off this storage module click "Power Off Module".
To shut down the management group click "Shut Down Group".

Figure 30. Confirming Storage Module Power Off

### If you choose Power Off Module

- Click Power Off Module.

  The storage module will power down in the specified number of minutes. Depending on the configuration of the management group and volumes, your volumes and snapshots remain available.

### If you choose Shut Down Group

- Click Shut Down Group.

  The management group shuts down in the allotted time and disappears from the Console.

To restart the management group, see “Starting Up a Management Group” on page 185.

***Note:*** *For information about powering off the module manually, see the Technical Product Specification (TPS) provided with the SSM.*

# Registering Features for an SSM

Using the Feature Registration tab, you can register individual SSMs for add-on features and applications such as Remote Copy. You can also register SSMs when they are in a management group by using the Management Group Register tab.

For detailed information about registering, see "Registering Features and Applications" on page 292.

## Using the Feature Registration Tab

The Feature Registration tab displays the following information:

- The SSM feature key, used to obtain a license key
- The license key for that SSM
- Which, if any, add-on features or applications have been licensed

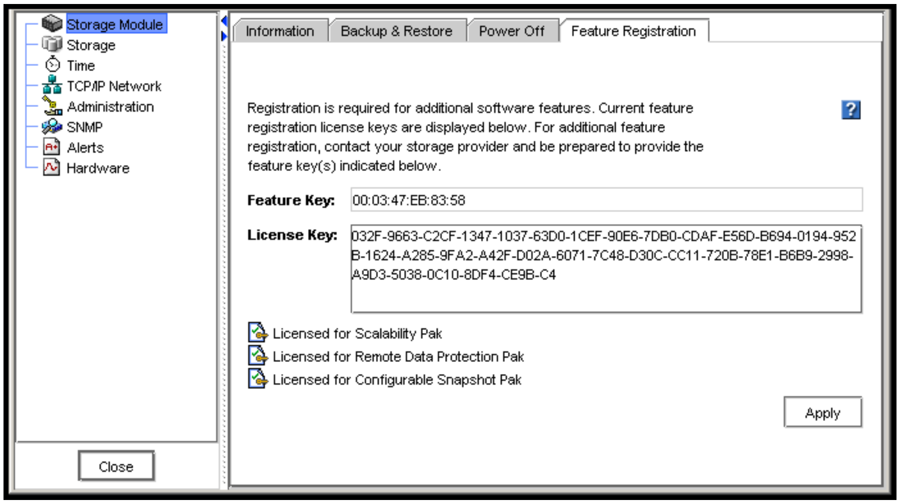

Registration is required for additional software features. Current feature registration license keys are displayed below. For additional feature registration, contact your storage provider and be prepared to provide the feature key(s) indicated below.

**Feature Key:** 00:03:47:EB:83:58

**License Key:** 032F-9663-C2CF-1347-1037-63D0-1CEF-90E6-7DB0-CDAF-E56D-B694-0194-952B-1624-A285-9FA2-A42F-D02A-6071-7C48-D30C-CC11-720B-78E1-B6B9-2998-A9D3-5038-0C10-8DF4-CE9B-C4

Licensed for Scalability Pak
Licensed for Remote Data Protection Pak
Licensed for Configurable Snapshot Pak

### Figure 31. Viewing the Feature Registration Tab

Register the SSM and purchase a license key to apply in this window. First, submit the feature key as instructed. Then, when you receive a license key, copy and paste it into the License Key field.

## Evaluating Features

Add-on features and applications are available when you begin using the Storage System Software. If you begin using an add-on feature or application without registering, a 30-day evaluation period begins. Throughout the evaluation period you receive reminders to register and purchase a license for the add-on feature and applications you want to continue using.

The Feature Registration tab lists the status of add-on features and applications on your SSM. An example of an unlicensed application in the evaluation period is shown in the following figure.

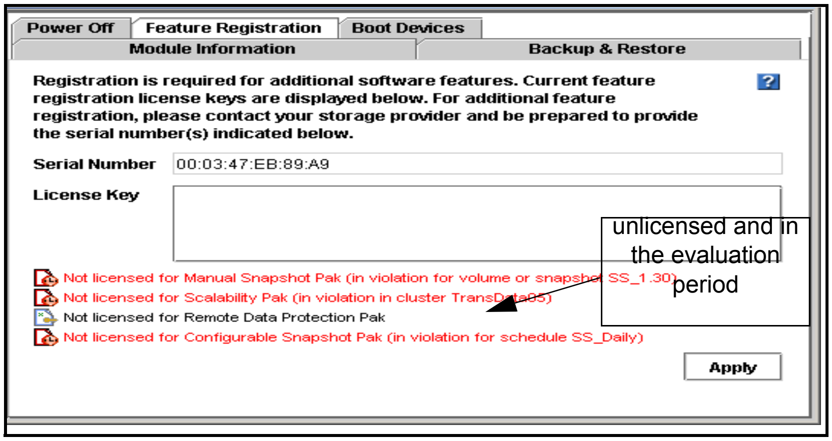

Registration is required for additional software features. Current feature registration license keys are displayed below. For additional feature registration, please contact your storage provider and be prepared to provide the serial number(s) indicated below.

Serial Number 00:03:47:EB:89:A9

License Key

Not licensed for Manual Snapshot Pak (in violation for volume or snapshot SS_1.30)

Not licensed for Scalability Pak (in violation in cluster TransData05)

Not licensed for Remote Data Protection Pak

Not licensed for Configurable Snapshot Pak (in violation for schedule SS_Daily)

**Figure 32. Using Remote Copy without a License**

For more detailed information about the evaluation process, see“Evaluating Features” on page 285.

# Configuring Boot Devices

The Intel® Storage System SSR212MA has a single boot device.

## Checking Boot Device Status in an SSM

You can view the compact flash card status on the Boot Devices window.

Open the Edit Configuration window to access the Storage Module configuration category.

1. Select an SSM in the navigation windows.
2. Click Storage Module Tasks in the tab window.
3. Select the Edit configuration category.

4. Select the Boot Devices tab.

   The Boot Devices window opens

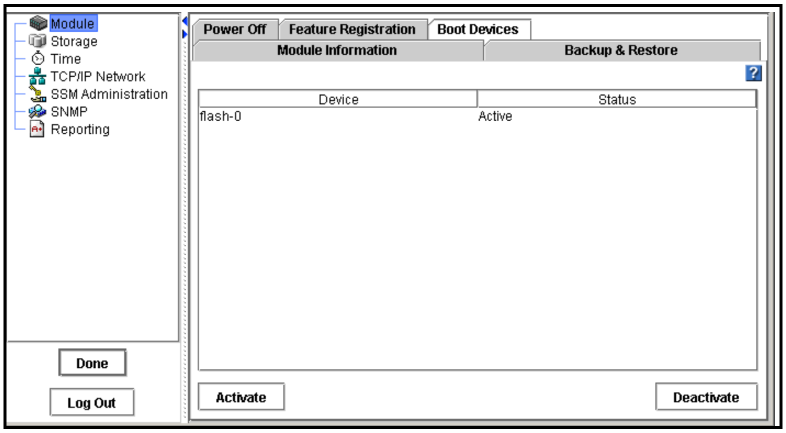

| Device | Status |
|---|---|
| flash-0 | Active |

**Figure 33. Viewing Single Boot Device Status**

The status of each compact flash card is listed in the Status column. The following table describes the possible statuses for compact flash cards.

### Table 1. Boot Flash Card Status[1]

| Flash Card Status | Description |
|---|---|
| Active | The device is synchronized and ready to be used. |
| Inactive | The device is ready to be removed from the SSM. It will not be used to boot the SSM. |
| Failed | The device encountered an I/O error and is not ready to be used. |
| Unformatted | The device has not yet been used in an SSM. It is ready to be activated. |
| Not Recognized | The device in the flash card bay is not recognized as a boot flash device. |
| Unsupported | The flash card in the flash card bay cannot be used. (For example, it is the wrong size or card type.) |

[1] Some statuses listed above only occur in a system with two boot devices.

***Note:*** *When the status of a flash card changes, an alert is generated. See "Using Active Monitoring" on page 149.*

## Replacing a Disk on Module (DOM) (Intel® Storage System SSR212MA only)

1. Power down the SSM.
2. Remove the DOM from the SSM. Refer to the *Intel® Storage System SSR212MA User Guide* for instructions on removing the DOM.
3. Install the new DOM in the SSM. Refer to the *Intel® Storage System SSR212MA User Guide* for instructions on installing the DOM.
4. Attach a serial cable to the storage system and connect to a laptop. Open a terminal emulation program to run a text interface, such as HyperTerminal* or ProComm Plus*.

   Use the following settings to configure your session:

   — Bits per second = 19200

   — Data bits = 8

   — Parity = None

   — Stop bits = 1

   — Flow control = None

   — Backspace key sends = Del

   — Emulation = ANSI

   If using HyperTerminal, set the properties for the backspace key and emulation after the session is established. If you exit the session and return to the session in order to use the Configuration Interface, the screen will not open correctly.
5. Power up the SSM with the replacement DOM. From the laptop, you should be able to observe two boot cycles. A boot cycle is indicated by a “Welcome to SAN IQ” message displayed on the screen. On the second boot, the cycle should end with a “DOM replacement logic: OS was restored to DOM on previous boot cycle” message. The logon screen will display, indicating a proper restoration process.

***Caution:*** *Do not execute any keyboard commands, such as <ESC> to view diagnostic messages, <F2> to enter setup, <F12> for a network boot, <CTRL> <G> for running the RAID BIOS Console or login to the storage system, during reboot.*

***Note:*** *Disregard any failed statuses and failure messages during reboot. These statuses / messages are normal and are not an indication of a failure.*

The entire restoration process, if successful, will take about 30 minutes. Once the two boot cycles have executed, ensure the system has been restored.

6. From the laptop or text interface, log in and verify that the IP address and host name of the storage system have not changed. If the storage system uses DHCP, the IP address may have changed.
7. Login to the Intel® Storage System Console and select Edit Config -> Storage -> RAID Setup and ensure all disks are online and in their original RAID configuration.

All volumes should be available with data restored to all volumes, and your host should be able to perform an iSCSI login.

**Note:** *In most cases, DOM replacement should result in no issues. However, the following two conditions may occur if there is another hardware problem present. In both cases, refer to the Intel® Storage System SSR212MA User Guide for instructions on removing the DOM and replacing it with the original. If the new DOM could not access data on the disks because of another system fault (e.g., RAID is seriously degraded or no longer configured, or the RAID controllers, midplane or server board have a failure) then the replacement DOM will boot a single time and appear to be a newly manufactured system. Check the network settings and if they are set to factory defaults, the restoration process has failed because the DOM could not detect a coherent RAID configuration. In this case, the DOM cannot be used again to attempt a system restoration. Replace with the original DOM because the problem is not a bad DOM. If the original RAID array is intact, but the restoration process is unsuccessful because the new DOM can't be written or verified, then the system will remain in a reboot cycle attempting to recover the configuration. If the DOM has not recovered after several reboot cycles or exceeded an hour without completing the process then the system cannot recover the original configuration. Power down the system if the system is continuously rebooting.*

# 3 Storage

## Storage Overview

The Storage configuration category is where you configure and manage RAID and individual disks for storage modules.

For each storage module, you can select the RAID configuration, the RAID rebuild options, and monitor the RAID status. You can also manage individual disks, including powering them on or off, and reviewing disk information.

## Storage Requirement

RAID must be configured for data storage. Table 2 lists the available RAID levels for the storage module.

### Table 2. RAID Levels

| Model | Available RAID Levels |
|---|---|
| Intel® Storage System SSR212MA | 0, 10, 5/50 |

## Getting There

Open the Edit Configuration window to access the Storage configuration category.

1. Select a SSM in the navigation window.
2. Click Storage Module Tasks in the tab window.
3. Select Edit Configuration.

4. Select the Storage configuration category.

   The RAID Setup tab opens.

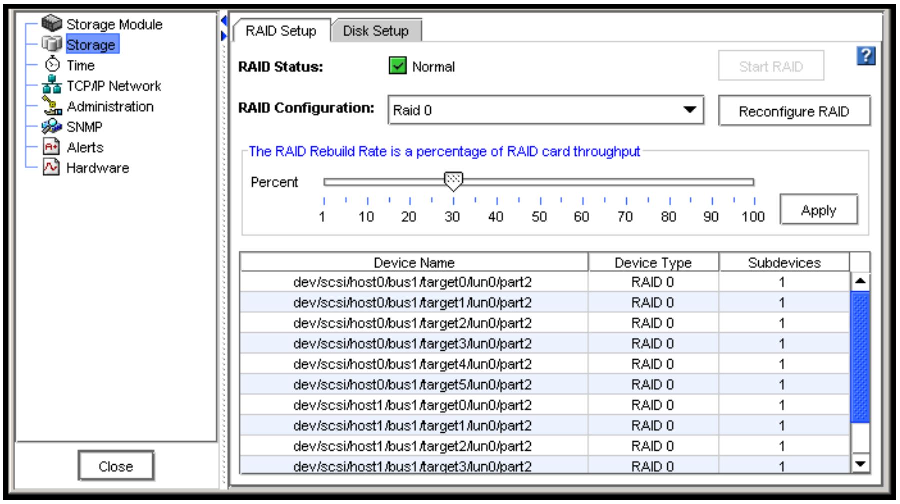

**RAID Status:** Normal

**RAID Configuration:** Raid 0

The RAID Rebuild Rate is a percentage of RAID card throughput

Percent 1 10 20 30 40 50 60 70 80 90 100

| Device Name | Device Type | Subdevices |
|---|---|---|
| dev/scsi/host0/bus1/target0/lun0/part2 | RAID 0 | 1 |
| dev/scsi/host0/bus1/target1/lun0/part2 | RAID 0 | 1 |
| dev/scsi/host0/bus1/target2/lun0/part2 | RAID 0 | 1 |
| dev/scsi/host0/bus1/target3/lun0/part2 | RAID 0 | 1 |
| dev/scsi/host0/bus1/target4/lun0/part2 | RAID 0 | 1 |
| dev/scsi/host0/bus1/target5/lun0/part2 | RAID 0 | 1 |
| dev/scsi/host1/bus1/target0/lun0/part2 | RAID 0 | 1 |
| dev/scsi/host1/bus1/target1/lun0/part2 | RAID 0 | 1 |
| dev/scsi/host1/bus1/target2/lun0/part2 | RAID 0 | 1 |
| dev/scsi/host1/bus1/target3/lun0/part2 | RAID 0 | 1 |

**Figure 34. Viewing the Storage Configuration Category for an Intel® Storage System SSR212MA**

# Configuring and Managing RAID

Managing the RAID settings of an SSM includes:

- Choosing the right RAID configuration for your storage needs
- Setting or changing the RAID configuration
- Setting the rate for rebuilding RAID
- Monitoring the RAID status for the SSM
- Starting or reconfiguring RAID when necessary

## Benefits of RAID

RAID combines several physical disks into a larger virtual disk. This larger virtual disk can be configured to improve both read/write performance and data reliability for the module.

# RAID Configurations Defined

The RAID configuration you choose depends upon how you plan to use the storage module. The module can be reconfigured with RAID 0, RAID 1/10 or RAID 5/50, depending on the model.

## Number of Disks and RAID

The number of disks in the SSM affects the RAID configurations available, as illustrated in Table 2 .

### RAID 0

RAID 0 is available for any number of disks in an SSM.

### RAID 1/10

RAID 1/10 requires pairs of disks. Therefore, an SSM must contain an even number of disks to configure RAID 1/10.

### RAID 5/50

RAID 5/50 requires sets of disks to configure.

- Intel® Storage System SSR212MA requires sets of 6 disks up to a total of 12 disks.

**Table 3. Number of Disks Required for Each RAID Level**

| | **Number of disks required** |
|---|---|
| **RAID Level** | **Intel® Storage System SSR212MA** |
| RAID 0 | from 1 to 12 |
| RAID 1/10 | 2, 4, 6, 8, 10, or 12 |
| RAID 5/50 | 6 or 12 |

## RAID Set Size

The RAID set size is limited to 2 TB. This means that the combined capacity of the disks participating in the RAID set cannot exceed 2 TB.

## RAID 5/50 in the Intel® Storage System SSR212MA

Table 4 illustrates RAID 5/50 capacity calculations for three different disk capacities in the Intel® Storage System SSR212MA. Since 2 TB equals 2048 GB, the RAID 5/50 configuration available for 400 GB disks is 5 plus a spare. RAID 5/50 is supported for 500 GB disks. but your usable disk capacity is limited, as shown in Table 7.

**Table 4. Calculating RAID Set size for RAID 5/50 in the Intel® Storage System SSR212MA**

| For Intel® Storage System SSR212MA | Using Disk Capacity of | | |
|---|---|---|---|
| RAID 5/50 Set Size | 250 GB | 400 GB | 500 GB |
| 5 plus a spare | 1250 GB | 2000 GB | (2500 GB) |
| 6 disks | 1500 GB | (2400 GB) | (3000 GB) [1] |

[1] Parentheses indicate RAID set size greater than 2 TB; therefore, your usable disk capacity will be smaller than expected.

# RAID 0

RAID 0 creates a striped disk set. Data will be stored across all disks in the RAID which increases performance. However, RAID 0 does not provide fault tolerance. If one disk in the set is powered down or fails, all data on the set will be lost.

SSM capacity in RAID 0 is equal to the total capacity of all disks in the module.

# RAID 1 and RAID 10

### RAID 1

RAID 1 provides data redundancy by mirroring the data from one disk onto a second disk.

### RAID 10

RAID 10 combines mirroring data within pairs of disks with striping data across pairs. RAID 10 combines data redundancy with the performance boost of RAID 0.

## Configuring RAID 1 or RAID 10

Whether the SSM is configured in RAID 1 or RAID 10 depends on the number of disks in the module.

- If the SSM contains only 2 disks configured for RAID, then the mirrored disk pair is RAID 1.
- If the SSM contains 4 or more disks configured for RAID, then the 2 or more mirrored disk pairs are RAID 10.

## Storage Capacity in RAID 10

SSM capacity in RAID 10 is the total capacity of all mirrored disk pairs in the module. The capacity of a single disk pair is equal to the capacity of one of the disks, as shown in Figure 35.

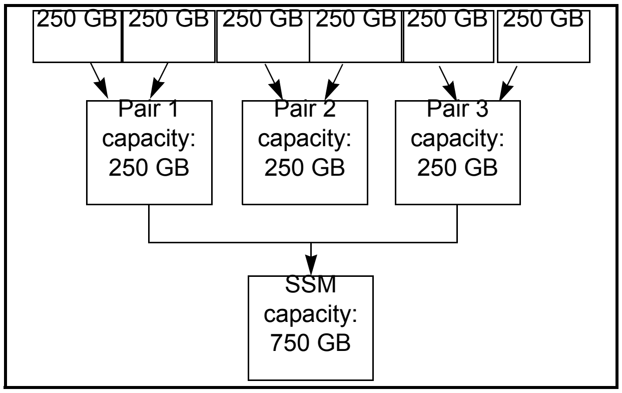

**Figure 35. Capacity of Disk Pairs in RAID 10**

# RAID 5 and RAID 50

RAID 5 provides data redundancy by distributing data blocks across all disks in a RAID set. Redundant information is stored as parity distributed across the disks. Figure 36 shows an example of the distribution of parity across 4 disks in a RAID 5 set.

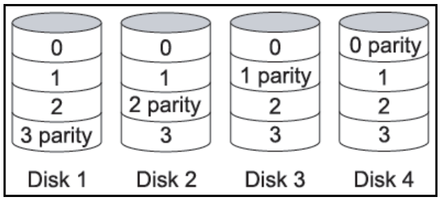

**Figure 36. Parity Distributed Across a RAID 5 Set Using Four Disks**

Parity allows the storage module to use more disk capacity for data storage than RAID 10 allows.

## Parity and Storage Capacity in RAID 5

Parity in a RAID 5 set equals the capacity of one disk in the set. Therefore, the capacity of any RAID 5 set is *n* - 1, as illustrated in Table 36 .

### Table 5. Storage Capacity of RAID 5 Sets in SSMs

| Model | RAID 5 Configuration | Storage Capacity |
|---|---|---|
| Intel® Storage System SSR212MA | RAID 5 - 5 disks plus a spare<br>RAID 5 - 6 disks | 4 x single disk capacity<br>5 x single disk capacity |

## RAID 5 and Hot Spare Disks

RAID 5 configurations that use a spare designate as a hot spare the remaining disk of the RAID set. With a hot spare disk, if any one of the disks in the RAID 5 set fails, the hot spare disk is automatically added to the set.

Table 6 lists the RAID 5 configurations by model, and indicate which configurations include a hot spare.

### Table 6. RAID 5 Configurations in Storage Modules

| SSM Model | RAID 5 Configuration |
|---|---|
| Intel® Storage System SSR212MA | RAID 5 - 5 disks plus a spare<br>RAID 5 - 6 disks |

## Configuring RAID 5 or RAID 50 on the Intel® Storage System SSR212MA

RAID 5 and RAID 50 can only be configured on completely populated sets of disks. This means the Intel® Storage System SSR212MA must contain either 6 or 12 disks.

Whether the Intel® Storage System SSR212MA is configured in RAID 5 or RAID 50 depends on the number of disk sets in the module.

- If the Intel® Storage System SSR212MA contains 1 disk set, then that set is RAID 5.
- If the Intel® Storage System SSR212MA contains 2 disk sets, then both sets are RAID 50.

## Storage Capacity in RAID 50 on the Intel® Storage System SSR212MA

The total capacity of the Intel® Storage System SSR212MA in RAID 50 is the combined capacity of each RAID 5 set in the module.

For example, suppose the Intel® Storage System SSR212MA is configured for RAID 50 and contains 2 sets of 6 250 GB disks. The total capacity for that equals 2500 GB.

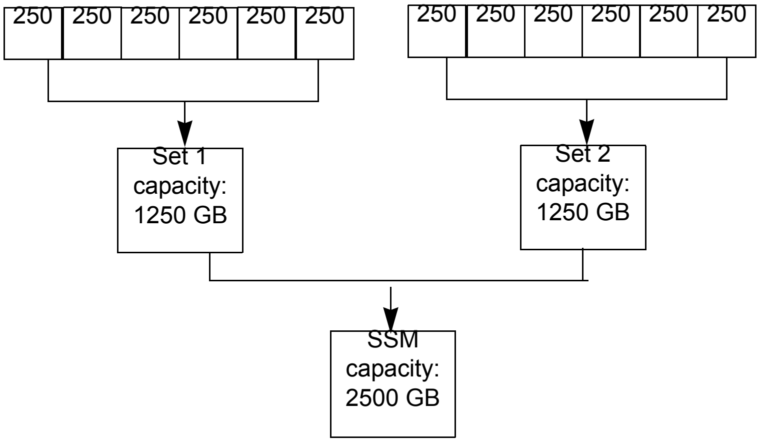

**Figure 37. Capacity of Disk Sets in RAID 50**

# RAID Set Size Limits on Capacity in the Intel® Storage System SSR212MA

The maximum size of a RAID set is 2 TB per RAID device. Certain hardware platforms (currently the Intel® Storage System SSR212MA only) and RAID configurations may not yield expected storage capacities given this limitation. The Intel® Storage System SSR212MA has two RAID devices, so the maximum unformatted capacity available from the Intel® Storage System SSR212MA in a RAID 5 configuration is 4 TB. All other platforms at this time yield capacity in every RAID configuration.

## RAID 5/50 in the Intel® Storage System SSR212MA

The following table illustrates RAID 5/50 capacity calculations for two different supported disk capacities (250 GB and 500 GB) in the Intel® Storage System SSR212MA. Capacity yields are affected by overhead in disk type; however, capacity yields are affected by overhead in disk yield and the partition.

Table 7. Calculating RAID Capacity in the Intel® Storage System SSR212MA

| RAID Type | Usable Disk Capacity 250 GB | Usable Disk Capacity 500 GB |
|---|---|---|
| RAID 0 | 2.61 TB | 5.31 TB |
| RAID 10 | 1.30 TB | 2.66 TB |
| RAID 5 plus hot spare | 1.33 TB | 3.55 TB |
| RAID 5 | 2.20 TB | 3.88 TB |

# Understanding RAID Devices in the RAID Setup Report

In the Storage category, the RAID Setup tab lists the RAID devices in the storage module and provides information about them. Information listed in the report is described in Table 8 .

## RAID Devices by RAID Type

Each RAID type creates different sets of RAID devices. What follows is a description of the variety of RAID devices created by the different RAID types as implemented on various platform models.

Table 8. Information in the RAID Setup Report

| This Item | Describes This |
|---|---|
| Device Name | The disk sets used in RAID. The number and names of devices varies by platform and RAID level. |
| Device Type | The RAID level of the device.<br>If the device is not functioning properly, the RAID Level reads “failed” and the level. For example “failed 5.” |
| Subdevices | The number of disks included in the device. For example, in an Intel® Storage System SSR212MA with 12 drives configured for RAID 5 (5 disks plus a spare) displays a Device Type of “RAID 5” and subdevices as “6.” |

## Devices Configured in RAID 0

If RAID 0 is configured, each physical disk operates as a separate RAID 0 disk, as shown in Figure 38.

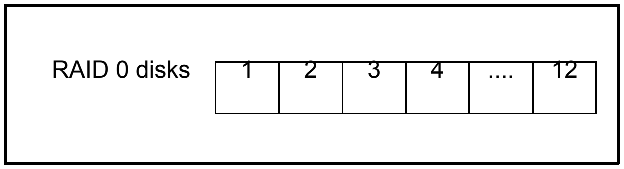

**Figure 38. RAID 0 on an Intel® Storage System SSR212MA**

## Devices Configured in RAID 1/10

If RAID 1 or 10 is configured, the physical disks are combined into mirrored pairs of disks, as shown in Figure 39. RAID 1 uses only one pair of disks. RAID 10 uses up to 8 pairs of disks, depending on the platform.

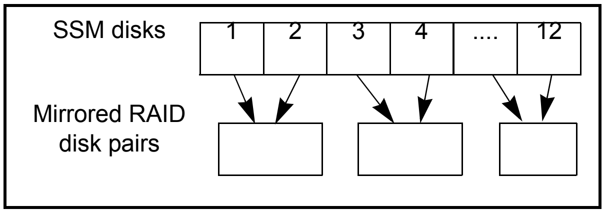

**Figure 39. RAID 10 on an Intel® Storage System SSR212MA**

## Devices Configured in RAID 5/50

If RAID 5 is configured, all the physical disks are grouped into one array of disks. If RAID 5 or 50 is configured, the physical disks are grouped into sets. RAID 5 uses one set of disks. RAID 50 uses multiple sets of disks in each SSM.

## RAID 50 in the Intel® Storage System SSR212MA

RAID 50 in the Intel® Storage System SSR212MA consists of sets using either all 12 disks in 6-disk sets, shown in Figure 40., or *n-1* disks so that the single disk acts as a hot spare for the RAID set, shown in Figure 41. The RAID 50 *n-1* configuration is 5 disks plus a spare.

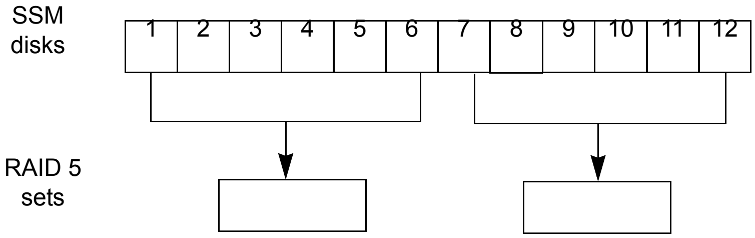

**Figure 40. Intel® Storage System SSR212MA RAID 50 Using 6-disk Sets**

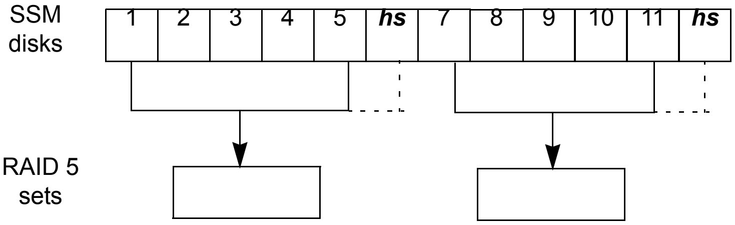

**Figure 41. Intel® Storage System SSR212MA RAID 50 Using 5 Disks Plus a Hot Spare**

# Planning RAID Configuration

The RAID configuration you choose for the storage module depends on your plans for data safety, data availability, and capacity growth. If you plan to expand your network of storage modules and create clusters, choose your RAID configuration carefully.

***Warning:*** *Once RAID is configured, you cannot change the RAID configuration without deleting all data on the storage module.*

## Data Replication

Keeping multiple copies of your data can ensure that data will be safe and will remain available in the case of disk failure. There are two ways to achieve data replication:

- Configure RAID 1, 10, 5, or 50 within each storage module.
- Replicate data volumes across clusters of storage modules.

## Using RAID for Data Replication

Within each storage module, RAID 1 or RAID 10 can ensure that 2 copies of all data exist. If one of the disks in a RAID pair goes down, data reads and writes can continue on the other disk. Similarly, RAID 5 or RAID 50 provides redundancy by spreading parity evenly across the disks in the set. If one disk in the set goes down, data reads and writes continue on the remaining disks in the set.

RAID protects against failure of disks within a module, but not against failure of an entire SSM. For example, if network connectivity to the SSM is lost, then data reads and writes to the SSM cannot continue.

***Note:*** *If you plan to create all data volumes on a single storage module, use RAID 1/10 or 5/50 to replicate data within that module.*

## Using Volume Replication in a Cluster

A cluster is a group of storage modules across which data can be replicated. Volume replication across a cluster of storage modules protects against disk failures within a storage module and failure of an entire storage module. For example, if a single disk or an entire module in a cluster goes down, data reads and writes can continue because an identical copy of the volume exists on other storage modules in the cluster.

Clustering is part of the Scalability Pak feature upgrade. See ”Working with Clusters” for more information.

***Note:*** *If you plan to create data volumes that span 2 or more storage modules, use replication in a cluster to ensure data safety and availability.*

# Using RAID with Replication in a Cluster

If you use replication in a cluster to replicate volumes across storage modules, then the redundancy provided by RAID 10 uses excess capacity and may not be necessary. For example,

- Using replication, up to 3 copies of a volume can be created on a cluster of 3 storage modules. The replicated configuration ensures that 2 of the 3 storage modules can go down and the volume will still be accessible.
- Configuring RAID 10 on these storage modules means that each of these 3 copies of the volume is stored on 2 disks within the storage module, for a total of 6 copies of each volume. For a 50 GB volume, 300 GB of disk capacity is used.
- In this case, data safety and availability are ensured more efficiently by configuring RAID 0 on the storage module and then achieving 2-way volume replication on clustered storage modules. For a 50 GB volume, 100 GB of disk capacity is used.

RAID 5/50 uses less disk capacity than RAID 1/10, so it can be combined with replication and still use capacity efficiently. One benefit of configuring RAID 5/50 in storage modules that use replication in a cluster is that if a single disk goes down, the data on that module can be rebuilt using RAID instead of requiring a complete copy from another module in the cluster. Rebuilding the disks within a single set is faster and creates less of a performance hit to applications accessing data than copying data from another module in the cluster.

**Note:** *If you are replicating volumes across a cluster:*
*- Configuring the SSM for RAID 0 allows you to use all of the disk capacity on the module while protecting against failure of individual disks or failure of an entire SSM.*
*- Configuring the SSM for RAID 5/50 provides redundancy within each SSM while allowing most of the disk capacity to be used for data storage.*

Table 9 summarizes the differences between running RAID 1 or 10 on a stand-alone SSM and running RAID 0 or RAID 5 on SSMs in a cluster.

### Table 9. Data Availability and Safety in RAID 1/10 Configuration and in a Clustered RAID 0 or RAID 5/50 Configuration

| Configuration | Safety and Availability During Disk Failure | Data Availability If Entire SSM Fails | Data Availability If Network Connection to SSM Lost | Hot Spare To Replace Failed Hardware |
|---|---|---|---|---|
| Stand-alone storage modules, RAID 1/10 | Yes. In any configuration, 1 disk per mirrored pair can fail, but there is no redundancy in pairs with a failed disk. | No | No | No hot spare disk within the storage module |
| Replicated volumes on clustered storage modules, RAID 0 | Yes. However, if any disk in the storage module fails, the entire storage module must be copied from another storage module in the cluster. | Yes | Yes | Yes |
| Replicated volumes on clustered storage modules, RAID 5/50 | Yes. 1 disk per RAID set can fail without copying from another storage module in the cluster. | Yes | Yes | Yes (select a hot spare disk RAID configuration) |

## Planning RAID for Capacity Growth

If you plan to add more storage modules to your network as your storage needs grow, remember that all storage modules in a cluster must have the same RAID configuration. For example, if you configure RAID 10 now, and later decide to replicate volumes through clustering, then any new storage modules must also be configured for RAID 10. Alternately, you can remove all data from your existing storage modules, configure RAID 0, and then cluster the storage modules.

**Warning:** *Once RAID is configured, you cannot change the RAID configuration without deleting all data on the storage module.*

# Setting RAID Rebuild Rate for RAID 1/10 or RAID 5/50

Choose the rate at which the RAID configuration rebuilds if a disk is replaced.

[Intel® Storage System SSR212MA] RAID rebuild rate is set as a priority against other operating system tasks.

General guidelines

- Setting the rate high is good for rebuilding RAID quickly and protecting data; however it will slow down user access.
- Setting the rate low allows users quicker access to data during the rebuild, but slows the rebuild rate.

## Setting RAID Rebuild Rate

1. Select the Storage configuration category.
2. Click the RAID Setup tab.
3. Set the slider for the desired rebuild rate.
4. Click Apply.

   The settings are then ready when and if RAID rebuild takes place.

# Starting RAID

If RAID has been configured on the storage module, and RAID is off, it must be started before other RAID tasks can be started.

Normally, once you start RAID, you will not have to restart it. However, in some cases, replacing disks requires that you start RAID after the disk is replaced.

**Example**

In a storage module, two disks were removed and replaced with two new disks. However, the disks that were removed caused the RAID quorum to break. (See "RAID Quorum" on page 69.) To prevent losing quorum, you replace one of the original disks and start RAID. Finally, you add the replacement disk to RAID.

## To Start RAID

1. Select the Storage configuration category.
2. Click the RAID Setup tab.
3. Click Start RAID.

   A confirmation message opens.
4. Click OK.

   RAID starts.

# Reconfiguring RAID

Reconfiguring RAID on a storage module destroys any data stored on that storage module. Requirements for reconfiguring RAID are listed below.

## Requirements for Reconfiguring RAID

**Changing preconfigured RAID on a new storage module**

RAID must be reconfigured on individual storage modules before they are added to a management group. If you want to change the preconfigured RAID level of a storage module, you must make the change before you add the module to a management group.

**Placement of disks in the storage module**

All disks must be in contiguous drive bays, from left to right and, for the Intel® Storage System SSR212MA, from top to bottom as shown in Figure 45, for RAID to be configured. If there are empty drive bays, only the disks to the left of the empty drive bay will be included in RAID. The remaining disks will be inactive.

Because RAID 1 and RAID 10 create mirrored disk pairs, there must be an even number of disks in the SSM. If you configure RAID 1 or RAID 10 on an SSM that contains an odd number of disks, RAID will be configured, but the odd disk will not be included in RAID. For example, if the SSM contains 9 disks, then disks 1-8 will be included in 4 disk pairs. Disk 9 will be inactive. If you add a 10th disk later, you can add disks 9 and 10 to RAID.

RAID 5 and RAID 50 can only be configured on completely populated sets of disks.

- The Intel® Storage System SSR212MA must contain either 6 or 12 disks.

**Management Groups and RAID**

You cannot reconfigure RAID on an storage module that is already in a management group. If you want to change the RAID configuration for an storage module that is in a management group, you must first remove it from the management group.

### Clusters and RAID

All storage modules in a cluster must have the same RAID configuration. However, you can have mixed versions of RAID 5/50 within a cluster.

Before you configure RAID, make sure that the disks in the SSM are inserted in contiguous disk bays, from left to right and, for the Intel® Storage System SSR212MA, from top to bottom, as shown in Figure 47.

- If you are configuring RAID 1 or RAID 10, the SSM must contain an even number of disks.
- If you are configuring RAID 5 or RAID 50:
  - — The Intel® Storage System SSR212MA must contain 6 or 12 disks

***Warning:*** *Changing the RAID configuration will erase all the data on the drives.*

## Reconfigure RAID

1. Click the RAID Setup tab to bring it to the front.

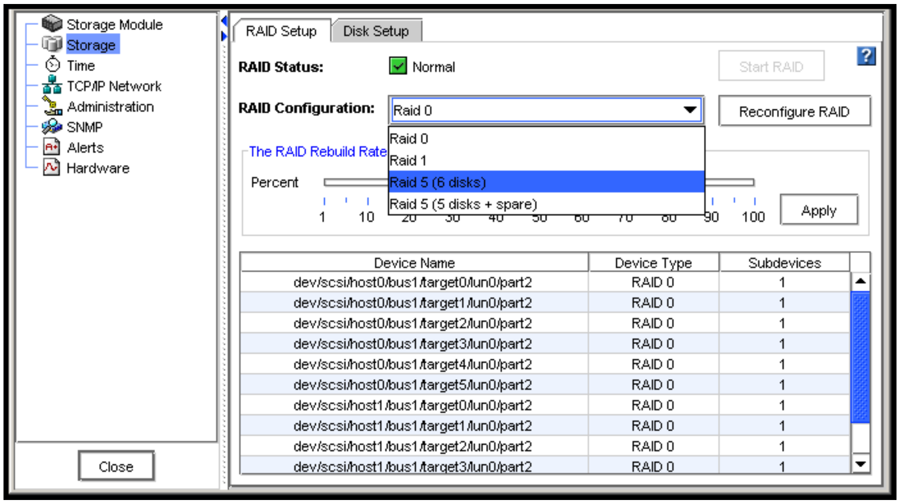

RAID Status: ☑ Normal

RAID Configuration: Raid 0

| Device Name | Device Type | Subdevices |
|---|---|---|
| dev/scsi/host0/bus1/target0/lun0/part2 | RAID 0 | 1 |
| dev/scsi/host0/bus1/target1/lun0/part2 | RAID 0 | 1 |
| dev/scsi/host0/bus1/target2/lun0/part2 | RAID 0 | 1 |
| dev/scsi/host0/bus1/target3/lun0/part2 | RAID 0 | 1 |
| dev/scsi/host0/bus1/target4/lun0/part2 | RAID 0 | 1 |
| dev/scsi/host0/bus1/target5/lun0/part2 | RAID 0 | 1 |
| dev/scsi/host1/bus1/target0/lun0/part2 | RAID 0 | 1 |
| dev/scsi/host1/bus1/target1/lun0/part2 | RAID 0 | 1 |
| dev/scsi/host1/bus1/target2/lun0/part2 | RAID 0 | 1 |
| dev/scsi/host1/bus1/target3/lun0/part2 | RAID 0 | 1 |

**Figure 42. Reconfiguring RAID on an Intel® Storage System SSR212MA**

2. Select the RAID configuration from the list.
3. Click Reconfigure RAID.

   A confirmation message opens.
4. Click OK.

   A warning message opens.
5. Click OK.

   RAID starts configuring.

**Note:** *If the SSM contains a large number of disks, it may take several hours for the disks to synchronize in a RAID 10 configuration.*
*When the RAID status on the RAID Setup tab shows Normal, the disks provide fully operational data redundancy with the mirror in place. The storage module is ready for data transfer at this point.*

# Monitoring RAID Status

RAID is critical to the operation of the storage module. If RAID has not been configured, the storage module cannot be used. Monitor the RAID status of a storage module to ensure that it remains normal. If the RAID status changes, a Console alert is generated. You can configure additional alerts to go to an email address or to an SNMP trap.

## Data Transfer and RAID Status

RAID status of Normal, Rebuild, or Degraded all allow data transfer. The only time data cannot be transferred to the storage module is if the RAID status shows Off.

## Data Redundancy and RAID Status

In a RAID 1/10 or RAID 5/50 configuration, when RAID is degraded there is not full data redundancy. Therefore, data is at risk if there is a disk failure when RAID is degraded.

**Warning:** *In a degraded RAID 1/10 configuration, loss of a second disk within a pair will result in data loss. In a degraded RAID 5/50 configuration, loss of a second disk will result in data loss.*

The RAID Status is located at the top of the RAID Setup tab in Storage. RAID status also displays in the Details Tab on the main Console window when an SSM is selected in the navigation window.

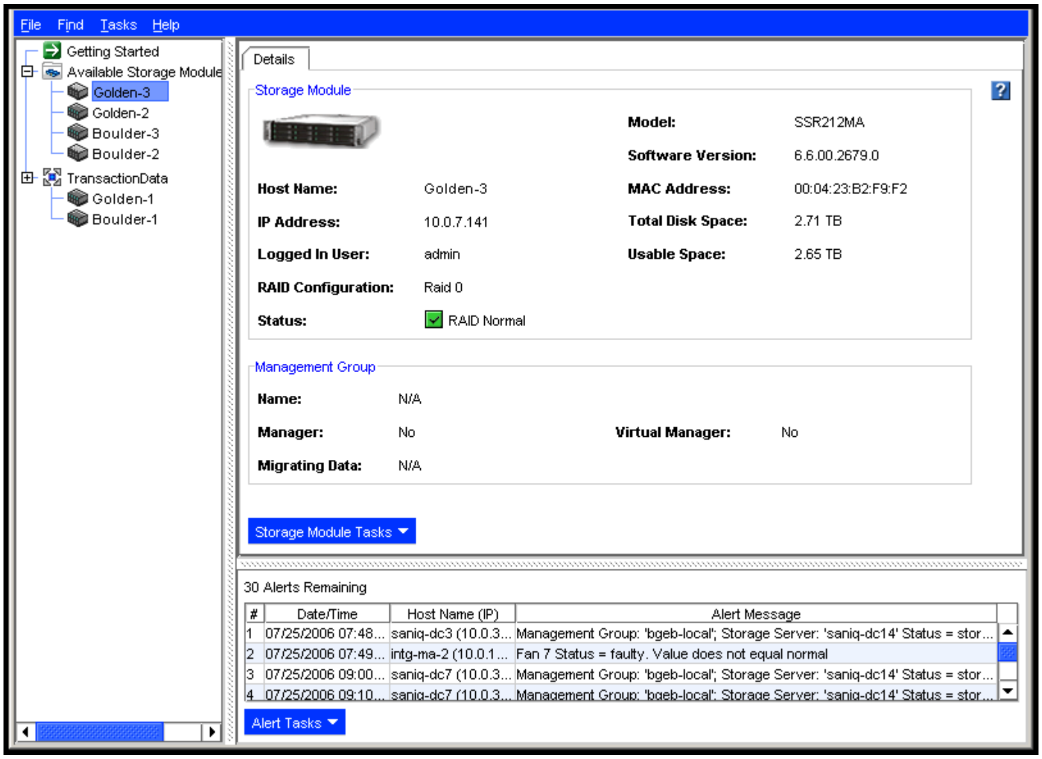

Storage Module

**Model:** SSR212MA

**Software Version:** 6.6.00.2679.0

**Host Name:** Golden-3

**MAC Address:** 00:04:23:B2:F9:F2

**IP Address:** 10.0.7.141

**Total Disk Space:** 2.71 TB

**Logged In User:** admin

**Usable Space:** 2.65 TB

**RAID Configuration:** Raid 0

**Status:** ✓ RAID Normal

Management Group

**Name:** N/A

**Manager:** No

**Virtual Manager:** No

**Migrating Data:** N/A

30 Alerts Remaining

| # | Date/Time | Host Name (IP) | Alert Message |
|---|---|---|---|
| 1 | 07/25/2006 07:48... | saniq-dc3 (10.0.3... | Management Group: 'bgeb-local'; Storage Server: 'saniq-dc14' Status = stor... |
| 2 | 07/25/2006 07:49... | intg-ma-2 (10.0.1... | Fan 7 Status = faulty. Value does not equal normal |
| 3 | 07/25/2006 09:00... | saniq-dc7 (10.0.3... | Management Group: 'bgeb-local'; Storage Server: 'saniq-dc14' Status = stor... |
| 4 | 07/25/2006 09:10... | saniq-dc7 (10.0.3... | Management Group: 'bgeb-local'; Storage Server: 'saniq-dc14' Status = stor... |

Figure 43. Monitoring RAID Status on the Main Console Window

The status displays one of four RAID states.

- **Normal** - RAID is synchronized and running. No action is required.
- **Rebuild** - A new disk has been inserted in a drive bay and RAID is currently rebuilding. No action is required.
- **Degraded** - RAID is degraded. Either a disk needs to be replaced or a replacement disk has been inserted in a drive.

  You must add a disk to RAID on Disk Setup if you are inserting a replacement disk.
- **Off** - Data cannot be stored on the storage module. The storage module is down and flashes red in the Network view.

# RAID Quorum

RAID quorum must be maintained for RAID 1/10 or RAID 5/50 to operate and for data to be preserved.

## Quorum for RAID 1 or RAID 10

For RAID 1/10, quorum requires that at least one disk pair in the storage module and one disk in each remaining pair be intact. This means that a storage module configured for RAID 10 can tolerate the loss of up to 5 disks in the Intel® Storage System SSR212MA. An SSM configured in RAID 1 contains only one pair of disks, so if one of the disks in the pair fails, quorum is broken and RAID cannot be rebuilt.

Data is safe as long as both disks in one of the mirrored pairs are operating normally. In order for RAID to rebuild when disks are replaced, at least one complete pair of disks must be in the storage module to ensure that data is rebuilt correctly. See "Managing Disks" for detailed information about the ordering of disks in the storage module.

**Disk Pairs**

Disks are usually paired from left to right, and for the Intel® Storage System SSR212MA, from top to bottom starting with the first disk in the SSM.

- Disks 1 and 2
- Disks 3 and 4
- Disks 5 and 6

and so on.

## Quorum for RAID 5 or RAID 50

For RAID 5/50, quorum requires that at least *n-1* in each RAID set be intact. If too many disks fail within one set, quorum is broken and RAID cannot be rebuilt. Table 10 shows the number of intact disks required to maintain quorum for each configuration of RAID 5/50 sets in the Intel® Storage System SSR212MA.

**Table 10. Disk Requirements for Maintaining RAID Quorum**

| Model | RAID Set Configuration | Number of Intact Disks Required to Maintain Quorum |
|---|---|---|
| Intel® Storage System SSR212MA | 5 plus a spare<br>6 | 4 of 5<br>5 of 6 |

### *RAID Sets*

Disks are grouped into RAID 5/50 sets from left to right, starting with the first disk in the storage module.

In the Intel® Storage System SSR212MA

- Disks 1-6 and 7-12

## Replacing Disks and RAID

Disk failure in a storage module affects RAID for that module. First, replace the failed disk. Then re-establish RAID. More information about replacing a disk and minimizing data restriping can be found in "Repairing a Storage Module" on page 209.

- When using RAID 0, you must reconfigure RAID 0. If the storage module is in a cluster, you must first remove the storage module from the management group and then reconfigure RAID 0.
- When using RAID 1/10 or RAID 5/50, RAID must be rebuilt. As long as RAID quorum was not lost, you can replace disks in an storage module and rebuild RAID while the storage module remains in the cluster. See "RAID Quorum" on page 69.

You can view the status of the disks in the storage module on the Disk Setup tab, shown in Figure 44. The RAID states are reported on the RAID Setup tab.

For detailed information about replacing disks, see "Managing Disks" on page 71.

# Managing Disks

***Note:*** *SSMs do not support hot-swap disk drives.*

Use the Disk Setup tab to

- monitor information about the disks in the selected storage module,
- power on a disk that you have replaced or added to the storage module, and
- add disks to RAID. You can also power off disks on this tab.

## Getting There

1. Open the Edit Configuration window.
2. Select the Storage category.
3. Select the Disk Setup tab.

## Reading the Disk Report on the Disk Setup Tab

The Disk Setup tab provides a report of the individual disks in the module and provides information about them.

**Table 11. Description of Items on the Disk Report**

| This Item | Describes This |
|---|---|
| Disk | Corresponds to the physical slot in the storage module. |
| Model | The model of the disk. |

### Table 11. Description of Items on the Disk Report

| This Item | Describes This |
|---|---|
| Serial Number | The serial number of the disk. |
| Class | The class (type) of disk. The SSM uses SATA disks. |
| Capacity | The data storage capacity of the disk. |
| Status | Whether the disk is<br>• Active (on and participating in RAID).<br>• Uninitialized or Inactive (On but not participating in RAID).<br>• Off or Missing (Not on).<br>• DMA Off (disk unavailable due to faulty hardware or improperly seated) |

# Verifying Disk Status

Check the Disk Setup window to verify that all the disks in the storage module are active and participating in RAID.

Any drive bays that do not contain disks are labelled “Off or Missing” in the Status column. Disks that have been inserted in the SSM but not yet added to RAID are labelled “Uninitialized” in the Status column.

## Disk Setup Tab for the Intel® Storage System SSR212MA

For the Intel® Storage System SSR212MA, the drives are labelled 1 through 12 in the Disk Setup window and correspond with the disk drives from left to right and top to bottom when you are looking at the front of the Intel® Storage System SSR212MA.

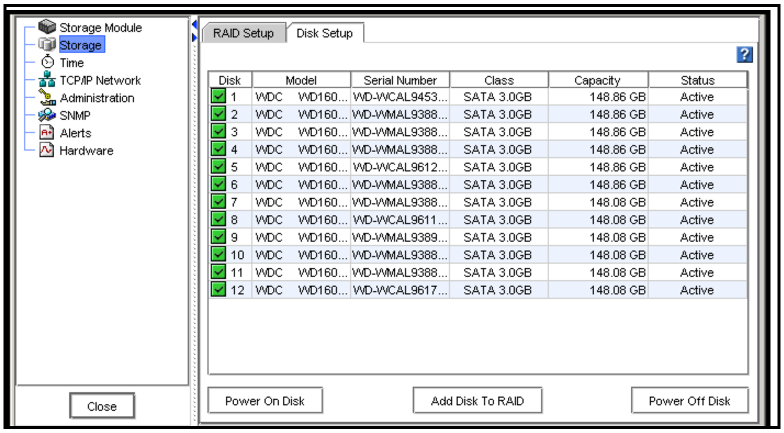

| Disk | Model | Serial Number | Class | Capacity | Status |
|---|---|---|---|---|---|
| ✓ 1 | WDC WD160... | WD-WCAL9453... | SATA 3.0GB | 148.86 GB | Active |
| ✓ 2 | WDC WD160... | WD-WMAL9388... | SATA 3.0GB | 148.86 GB | Active |
| ✓ 3 | WDC WD160... | WD-WMAL9388... | SATA 3.0GB | 148.86 GB | Active |
| ✓ 4 | WDC WD160... | WD-WMAL9388... | SATA 3.0GB | 148.86 GB | Active |
| ✓ 5 | WDC WD160... | WD-WCAL9612... | SATA 3.0GB | 148.86 GB | Active |
| ✓ 6 | WDC WD160... | WD-WMAL9388... | SATA 3.0GB | 148.86 GB | Active |
| ✓ 7 | WDC WD160... | WD-WMAL9388... | SATA 3.0GB | 148.08 GB | Active |
| ✓ 8 | WDC WD160... | WD-WCAL9611... | SATA 3.0GB | 148.08 GB | Active |
| ✓ 9 | WDC WD160... | WD-WMAL9389... | SATA 3.0GB | 148.08 GB | Active |
| ✓ 10 | WDC WD160... | WD-WMAL9388... | SATA 3.0GB | 148.08 GB | Active |
| ✓ 11 | WDC WD160... | WD-WMAL9388... | SATA 3.0GB | 148.08 GB | Active |
| ✓ 12 | WDC WD160... | WD-WCAL9617... | SATA 3.0GB | 148.08 GB | Active |

**Figure 44. Viewing the Disk Setup Tab in an Intel® Storage System SSR212MA**

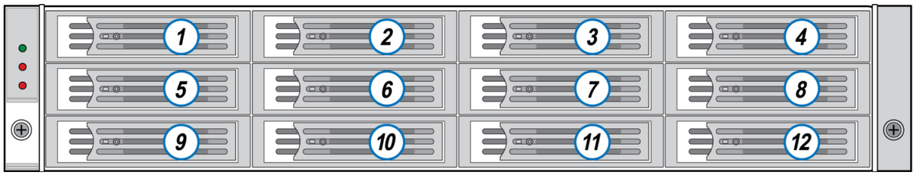

Figure 45. Diagram of the Drive Bays in the Intel® Storage System SSR212MA

# Replacing a Disk

Use this section if you are replacing a single disk under the following conditions:

- You know which disk needs to be replaced either through Intel® Storage Console monitoring or HealthCheck.
- The disk has not yet failed.
- RAID is still on, though it may be degraded.

Use the instructions in "Replacing Disks" on page 76:

- If RAID has gone off.
- If you have more than one disk to replace on the same or different storage modules.
- If you are unsure which disk to replace.

The instructions under"Replacing Disks" on page 76 include contacting Customer Support for assistance in either identifying the disk that needs to be replaced or, for replacing more than one disk, the order in which they should be replaced.

## Overview of Replacing a Disk

The correct procedure for replacing a disk in a storage module depends upon a number of factors, including the RAID configuration, the replication level of volumes and snapshots, and the number of disks you are replacing. Unless you are replacing a disk in a storage module that is not in a cluster, data must be rebuilt either just on the replaced disk or, in the case of RAID 0, on the entire storage module.
Replacing a disk in a storage module includes

- Planning for rebuilding data on either the disk or the entire storage module
- Powering the disk off in the Console
- Replacing the disk in the storage module
- Powering the disk on in the Console
- Rebuilding RAID on the disk or on the storage module

# Preparing for a Disk Replacement

How you prepare for a disk replacement differs according to the RAID level of the storage module.

## Checklist for Single Disk Replacement in RAID 0

RAID 0 provides no fault tolerance by itself. If you remove a disk from a RAID 0 configuration, all the data on the storage module will be lost. Therefore, if you need to replace a disk in a RAID 0 configuration, we recommend the following:

- All volumes and snapshots should have a minimum of 2-way replication.
- If volumes or snapshots are not replicated, change them to 2-way replication before replacing the disk.
- If the cluster does not have enough space for the replication, take a backup of the volumes or snapshots and then delete them from the cluster. After the disk replacement is complete, recreate the volumes and restore the data from the backup.
- All volumes and snapshots should show a status of Normal.
- Any volumes or snapshots that were being deleted should have finished deleting.
- Use the instructions in "Replacing Disks" on page 76if you have more than one disk to replace, or if you are unsure which disk needs replacing.

## Checklist for Single Disk Replacement in RAID 1/10 and 5/50

There are no prerequisites for this case; however we do recommend that:

- All volumes and snapshots should show a status of Normal.
- Any volumes or snapshots that were being deleted have completed deletion.
- Use the instructions in "Replacing Disks" on page 76 if you have more than one disk to replace, or if you are unsure which disk needs replacing.

# Replace the Disk

**Prerequisites**

- Know the name and physical location of the storage module that needs the disk replacement.
- Know the physical position of the disk in the storage module. See "Verifying Disk Status" on page 72 for diagrams of disk locations in the various platforms.
- Have the replacement disk ready and confirm that it is of the right size and has the right carrier.

## Powering Off a Disk in the Console

Powering off a single disk in RAID 1/10 or in RAID 5/50 causes RAID to run in a degraded state. Powering off a single disk in RAID 0 causes RAID to go off.

1. In the navigation window, select the storage module in which you want to replace the disk.
2. Click Storage Module Tasks and select Edit Configuration.

   The Edit Configuration window opens.
3. Select the Storage configuration category.
4. Click the Disk Setup tab.
5. Select the disk in the list to power off.
6. Click Power Off Disk.
7. A confirmation message opens.
8. Click OK.

## Replacing the Disk Drive in the Storage Module

See your hardware documentation for information about physically replacing disk drives in the storage module.

## Powering On a Disk in the Console

When you must insert a new disk into a storage module that is on, the disk must be powered on. If the disk is not powered on, it is listed as Off or Missing in the Status column and the other columns display dotted lines, like this ---------.

1. In the navigation window, select the storage module in which you replaced the disk drive.
2. Click Storage Module Tasks and select Edit Configuration.

   The Edit Configuration window opens.
3. Select the Storage configuration category.
4. Click the Disk Setup tab.
5. Select the disk in the list to power on
6. Click Power On Disk.

   A confirmation message opens.
7. Click OK.

## RAID Rebuilding

After the disks are powered on, RAID starts rebuilding. In RAID 0 the entire storage module starts rebuilding. In RAID 1/10 or 5/50, the replaced disk starts rebuilding.

# Replacing Disks

This section describes the disk replacement procedures for cases in which you do not know which disk to replace and/or you must rebuild RAID on the entire storage module. For example, if RAID has gone off unexpectedly, you need Customer Support to help determine the cause, and if it is a disk failure, to identify which disk must be replaced.

## RAID Levels and Disk Replacements

Single disk replacements in storage modules where RAID is running, but may be degraded, can be accomplished following the procedures described in "Replacing a Disk" on page 73. The following situations may require consulting with Customer Support to identify bad disks and then following the procedures below to rebuild the data (when replicated) on the storage module.

- RAID 0 (Stripe)—RAID is off due to a failed disk.
- RAID 5/50 (Stripe with parity)—if multiple disks need to be replaced, then those disks must be identified and replaced, and the data on the entire storage module rebuilt.
- RAID 10 (Mirror and Stripe) can sustain multiple disk replacements. However Customer Support must identify if any two disks are from the same mirror set, and then the data on the entire storage module needs to be rebuilt.

## Before You Begin

1. Know the name and physical location of the storage module that needs the disk replacement.
2. Know the physical position of the disk in the storage module.
3. Have the replacement disk ready and confirm that it is the right size and has the right carrier.
4. For confirmation on which disks need to be replaced, contact customer support.

## Replacing Disks and Rebuilding Data

Use this procedure for any one of these cases:

- RAID 0 goes off without warning of a bad disk,

  or
- When multiple disks needs to be replaced on a storage module with RAID 5/50,

  or
- When multiple disks on the same mirror set need to be replaced on a storage module with RAID 10.

**Prerequisites**

- All replicated volumes and snapshots should show a status of Normal. Non-replicated volumes may be blinking.
- If volumes or snapshots are not replicated, change them to 2-way replication before replacing the disk.
- If the cluster does not have enough space for the replication, take a backup of the volumes or snapshots and then delete them from the cluster. After the disk replacement is complete, recreate the volumes and restore the data from the backup.
- Any volumes or snapshots that were being deleted should have finished deleting.
- Write down the order in which the storage modules are listed in the Edit Cluster window. You must ensure that they are all returned to that order when the repair is completed.

## Storage Module not Running a Manager

Verify that the storage module that needs the disk replacement is not running a manager

1. Log in to the management group.
2. Select the storage module in the navigation window and review the Details tab information. If the Storage Module Status shows Manager Normal, and the Management Group Manager shows Normal, then a manager is running and needs to be stopped.
3. To stop a manager, right-click the storage module in the navigation window and select Stop Manager.

   When the process completes successfully, the manager is removed from the Status line in the Storage Module box and the Manager changes to "No" in the Management Group box.

If you stop a manager, the cluster will be left with an even number of managers. To ensure the cluster has an odd number of managers, do one of these tasks:

- Start a manager on another storage module or
- Add a virtual manager to the management group by right-clicking on the management group name in the navigation window and select Add Virtual Manager.

## Power off Disk

Power off the disk(s) that needs to be replaced.

1. Double-click the storage module icon on the Details tab to open the Edit Configuration window.
2. Select the Storage category and then select the Disk Setup tab.
3. Highlight the disk(s) that needs to be replaced and click the Power Off Disk button.
4. Select the RAID Setup tab to verify the RAID Status as OFF.

5. Click the Close button to close the Edit Configuration window.

   The storage module flashes red and the status in the navigation window, Storage Module Details tab changes to Storage Server Down.

## Mark the Storage Module

Mark the storage module for repair.

**Prerequisite**

If there are non-replicated volumes that are blinking red, you must either replicate them or delete them before you can proceed with this step.

1. Right-click on the storage module in the navigation window and select Repair Storage Module. A "ghost" image takes the place of the storage module in the cluster with the IP address serving as a place holder. The storage module itself moves from the management group to the Available Storage Modules pool.

***Note:*** *If the storage module does not appear in the Available Storage Modules area, use the “Find” menu option to re-locate it.*

## Replace the Disk

Replace the disk(s) in the storage module

1. Right-click on the storage module in the navigation window and select Edit Configuration.
2. In the Edit Configuration window, select the Storage category and then select the Disk Setup tab. Verify that the status of the disk(s) is Off or Missing.
3. Physically remove the disk(s) from the storage module and insert the replacement disk(s). When inserting the drive, be firm but not too forceful.

## Activate the Replacement Disk

1. Next, select the Storage category and then select the Disk Setup tab.
2. Highlight the disk(s) and click the Power On Disk button.

   The Disk Setup tab should show the status for that disk(s) as Inactive.

**Troubleshooting**

If the disk status remains as Off or Missing even though the disk(s) is physically there, do this:

1. Select the Module category and select the Power Off tab.
2. Shut down the storage module.

3. Physically reseat the disk.
4. Power on the storage module and check the disk status again.

# Re-create the RAID Array

1. Select the Storage category and then select the RAID Setup tab.
2. Verify the RAID configuration in the list and click the Reconfigure RAID button.

   The RAID Status changes from Off to Normal.

## Checking Time for RAID Array to Rebuild

Use the Hardware Information report to check the estimated time remaining for RAID to finish rebuilding.

1. Select the Hardware category and then select the Hardware Information tab.
2. Click on Refresh and scroll down to the RAID section of the Hardware report, shown in Figure 46.

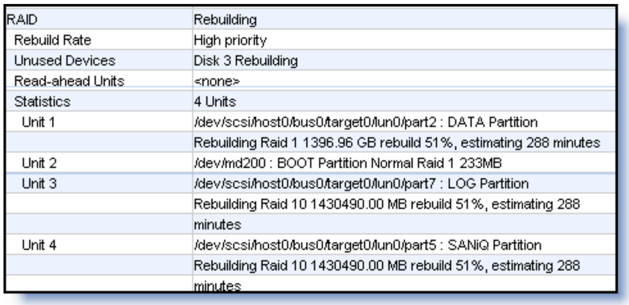

| | |
|---|---|
| RAID | Rebuilding |
| Rebuild Rate | High priority |
| Unused Devices | Disk 3 Rebuilding |
| Read-ahead Units | <none> |
| Statistics | 4 Units |
| Unit 1 | /dev/scsi/host0/bus0/target0/lun0/part2 : DATA Partition |
| | Rebuilding Raid 1 1396.96 GB rebuild 51%, estimating 288 minutes |
| Unit 2 | /dev/md200 : BOOT Partition Normal Raid 1 233MB |
| Unit 3 | /dev/scsi/host0/bus0/target0/lun0/part7 : LOG Partition |
| | Rebuilding Raid 10 1430490.00 MB rebuild 51%, estimating 288 minutes |
| Unit 4 | /dev/scsi/host0/bus0/target0/lun0/part5 : SANiQ Partition |
| | Rebuilding Raid 10 1430490.00 MB rebuild 51%, estimating 288 minutes |

**Figure 46. Checking Time for RAID to Rebuild**

   The Statistics show the number of RAID arrays, and each RAID array will be listed as Unit 1, Unit 2, etc. The Units show the RAID rebuild progress and an estimated time to complete.

3. Use the Refresh button to monitor the progress.

# Return to Cluster

Return the repaired storage module to the cluster.

1. In the navigation window, right-click the storage module and select Add to New or Existing Management Group.
2. Under the Existing Management Group, select the Group Name that the storage module used to belong to and click Add. The storage module appears in the management group and flashes red for a few minutes as it initializes.

### Restarting a manager

If necessary, ensure that after the repair you have the appropriate configuration of managers. If there was a manager running on the storage module before you began the repair process, you may start a manager on the repaired storage module, as necessary to finish with the correct number of managers in the management group.

If you added a virtual manager to the management group you must first delete the virtual manager before you can start a regular manager.

- First, right-click on the virtual manager and select Stop Virtual Manager.
- Next, right-click on the virtual manager and select Delete Virtual Manager.
- Finally, right-click on the storage module and select Start Manager.

1. After the initialization completes, right-click on the cluster and select Edit Cluster. The list of the storage modules in the cluster should include the ghost IP address.

   You now need to add the repaired storage module to the cluster in the spot held by the ghost IP address.
2. In the Edit Cluster window, first note the order of the storage modules in the list.
3. Next, remove the ghost storage module from the cluster.
4. Return the repaired storage module to the cluster in the position of the ghost storage module.

   Use the arrows to return the storage modules in the list to their original order.

   **Example:** If the list of the storage modules in the cluster had storage module A, <IP address>, storage module C, and storage module B is the repaired storage module, then after re-arranging, the list of the storage modules in the cluster should be storage module A, storage module B, storage module C and the storage module <IP address> will be in the management group.

**Table 12. Replacing the ghost storage module with the repaired storage module**

| | Storage Modules in Cluster |
|---|---|
| Before rearranging | storage module A |
| | <IP Address> |
| | storage module C |
| After rearranging | storage module A |
| | storage module B |
| | storage module C |

**Note:** *If you do not arrange the storage modules to match their original order, the data in the cluster is rebuilt across all the storage modules instead of just the repaired storage module. This total data rebuild takes longer to complete and increases the chance of a second failure during this period.*

To ensure that only the repaired storage module goes through the rebuild, before you click the OK button in the Edit Cluster window, double-check that the order of the storage modules in the cluster list matches the original order.

# Rebuild Volume Data

After the storage module is successfully added back to the cluster, the adjacent storage modules start rebuilding data on the repaired storage module.

1. Select the cluster and select the Disk Usage tab.
2. Verify that the disk usage on the repaired storage module starts increasing.
3. Verify that the status of the volumes and snapshots is Restriping.

   Depending on the usage, it may take anywhere from a few hours to a day for the data to be rebuilt on the repaired storage module.

## Control Client Access

Use the Local Bandwidth Priority setting to control client access to data during the rebuild process.

- When the data is being rebuilt, the clients that are accessing the data on the volumes might experience slowness. Reduce the Local Bandwidth Priority to half of its current value for immediate results.
- Alternatively, if client access performance is not a concern, raise the Local Bandwidth Priority to increase the data rebuild speed.

### Change Local Bandwidth Priority

1. Right-click the management group and select Edit Management Group. The current Bandwidth Priority value indicates that each manager in that management group will use that much bandwidth to transfer data to the repaired storage module. Make a note of the current value so it can be restored once the data rebuild completes.
2. Change the bandwidth value as desired and click OK.

# Remove ghost

Remove the ghost storage module after the data is rebuilt.

The data is rebuilt on the storage module when

- the repaired storage module's disk usage matches the usage of the other storage modules in the cluster, and
- the status of the volume and snapshots goes back to Normal.

The ghost IP address showing outside the cluster can now be removed from the management group.

1. Right-click the ghost IP address and select Remove from Management Group.

2. If you have adjusted/reduced the Local Bandwidth Priority of the Management Group while the data was being rebuilt, change it back to the original value.

At this point, the disk(s) in the storage module are successfully replaced, the data will be fully rebuilt on that storage module, and the management group configuration (like number of managers, quorum, local bandwidth etc.) will be restored to the original settings.

# Adding Disks to the SSM

If the SSM is configured for RAID 1 or RAID 10, you must add an even number of disks to include all the disks in the RAID configuration. See "Requirements for Reconfiguring RAID" on page 66.

If the SSM is configured for RAID 5 or RAID 50, you must add disks in complete sets, as follows

- Intel® Storage System SSR212MA - 6 disks at a time

## Diagrams of Disk Bays

Figure 47 illustrate the placement of the drive bays in the Intel® Storage System SSR212MA.

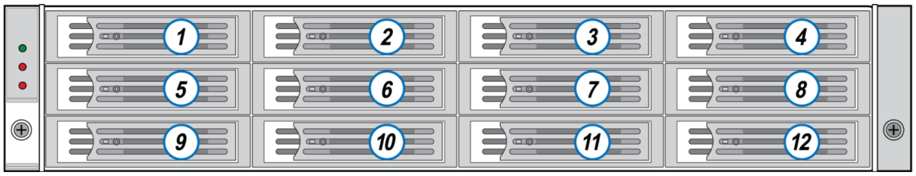

Figure 47. Diagram of the Drive Bays in the Intel® Storage System SSR212MA

## Adding Disks and SSM Capacity

If you are using clustering, all SSMs in a cluster will operate at a capacity equal to that of the smallest capacity SSM. Adding capacity to all SSMs in the cluster will prevent stranded storage.

You cannot reduce the capacity of an SSM that is part of a management group. If you want to reduce the capacity of an SSM, first remove it from the management group. Then remove disks from the SSM and reconfigure RAID.

***Note:*** *You must add disk in contiguous disk bays, from left to right, and for the Intel® Storage System SSR212MA, from top to bottom, as shown in Figure 47.*

# Memory Requirements for Adding Disks

Before you add disks to the SSM, confirm that the SSM has enough memory to use the additional disks. Table 13 summarizes the memory requirements by disk capacity of fully populated SSMs with RAID 0, RAID 1/10, and RAID 5/50 configurations. Contact your SSM supplier for additional memory.

### Table 13. Memory Requirements for Fully Populated SSM

| Intel® Storage System SSR212MA | Memory Requirement for 12 or 16 Disks (Fully Populated) | |
|---|---|---|
| **RAID Level** | **For 250 GB Disks** | **For 400 GB Disks** |
| RAID 0 | 1 GB | 2 GB |
| RAID 1 / 10 | 1 GB | 1 GB |
| RAID 5 / 50 | 1 GB | 2 GB |

# Adding Disks

**Prerequisite**

Before you add disks to the SSM, be sure that the SSM has enough memory to use the additional disks. See “Memory Requirements for Adding Disks” on page 83.

***Warning:*** *Adding a disk to RAID deletes any existing data on that disk.*

1. Add the new disks to the SSM.

   You must add disks in contiguous disk bays, from left to right and, for the Intel® Storage System SSR212MA, from top to bottom, as shown in Figure 47.

2. Using the Console, log in to the SSM.
3. Select the Storage configuration category.
4. Click the Disk Setup tab to bring it to the front.

   The new disks will show a red X and be listed as Off.

5. Select the new disks and click Power On Disk.

   The disk status of the new disks becomes Uninitialized.

6. Select the new disks and click Add Disk to RAID.

   Shift-click to select multiple disks to add to RAID

   - pairs of disks to RAID 1/10 or
   - sets of 6 disks [Intel® Storage System SSR212MA].

   RAID begins to rebuild on the new disks according to the RAID Rebuild rate configured on the RAID Setup tab.

As soon as the RAID Status shows Normal, the disks provide fully operational data redundancy with the mirror in place. The SSM is ready for data transfer at this point. The newly added disks display on the RAID Setup tab.

# 4 Managing the Network

## Managing the Network Overview

The storage module has two integrated TCP/IP network interfaces.In addition, the Intel® Storage System SSR212MA can include one add-on card, with 2 or 4 interfaces. For each storage module you can

- Configure the TCP/IP interfaces
- Set up and manage a DNS server
- Manage the routing table
- View and configure the TCP interface speed, duplex, and frame size
- Update the list of managers running in the management group to which a storage module belongs
- Bond NICs to ensure continuous network access or to improve bandwidth

***Note:*** *Anytime you make a change to the network configuration, you must manually reconnect iSCSI sessions.*

## Getting There

Open the Edit Configuration window to access the TCP/IP Network configuration category.

1. Select an SSM in the navigation window.
2. Click Storage Module Tasks in the tab window.
3. Select Edit Configuration.
4. Select TCP/IP Network configuration category.

   The window opens with the TCP/IP tab on top, shown in Figure 48.

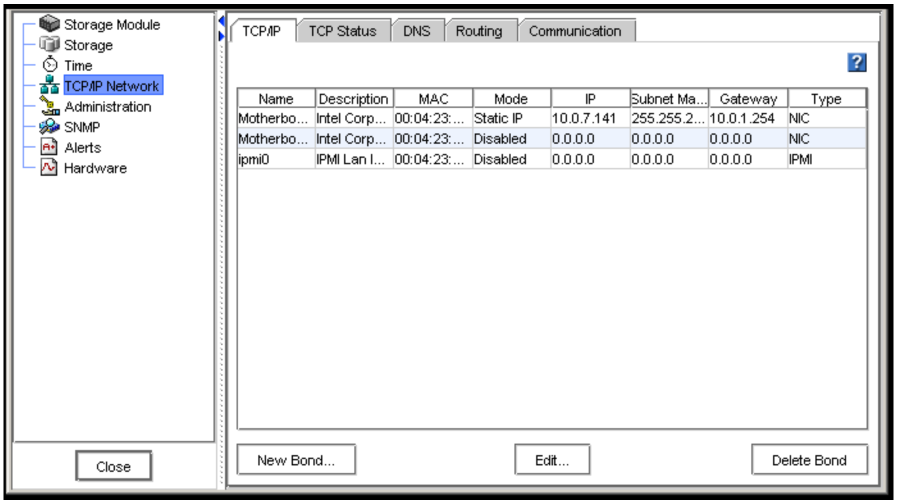

| Name | Description | MAC | Mode | IP | Subnet Ma... | Gateway | Type |
|---|---|---|---|---|---|---|---|
| Motherbo... | Intel Corp... | 00:04:23:... | Static IP | 10.0.7.141 | 255.255.2... | 10.0.1.254 | NIC |
| Motherbo... | Intel Corp... | 00:04:23:... | Disabled | 0.0.0.0 | 0.0.0.0 | 0.0.0.0 | NIC |
| ipmi0 | IPMI Lan I... | 00:04:23:... | Disabled | 0.0.0.0 | 0.0.0.0 | 0.0.0.0 | IPMI |

**Figure 48. Viewing the Network Configuration**

# The TCP/IP Tab

The TCP/IP tab lists the network interfaces on the storage module. You can configure each of these interfaces.

**Table 14. Network Interfaces Displayed on the TCP/IP Tab**

| Name | Description |
|---|---|
| **NICs Embedded in the SSM Motherboard** | |
| Motherboard:Port1 | 1000BASE-T interface |
| Motherboard:Port0 | 1000BASE-T interface |
| IPMI | Intelligent Platform Management Interface |
| **Add-on NICs in PCI Slots** | |
| Slot1:Port0<br>Slot1:Port1<br>and so on | Multiple add-in PCI cards, each containing up to 4 Ethernet interfaces. |
| **Bonded Interfaces** | |
| bondN | [Optional] You can create multiple bonded interfaces, each consisting of 2 or 4 physical interfaces. |

Use the TCP/IP tab to manage the network configurations for each network interface and to bond the network interfaces.

# Identifying the Network Interfaces

## Identifying Ports on the Back of the Storage Module

The SSM comes with two onboard Gigabit Ethernet ports. These ports are named Motherboard:Port0 and Motherboard:Port1, and are labelled on the back of the SSM as listed in Table 15 .

In addition, the SSM can include multiple add-on PCI cards, each with 2 or 4 Gigabit Ethernet ports. These add-on ports are named according to the card's slot and the port number, such as Slot1:Port0.

### Table 15. Identifying the NICs in the Motherboard

| Motherboard Interfaces | |
|---|---|
| **Where labelled** | **What the label says** |
| TCP/IP Network Configuration Category in the Console<br>• TCP/IP tab<br>• TCP Status tab | Name - Motherboard:Port0, Motherboard:Port1<br>Description - Intel Gigabit Ethernet |
| Configuration Interface Name | Motherboard:Port1<br>Motherboard:Port0 |
| Label on the back of the SSM | NICs 1 & 2 |

### Table 16. Identifying Add-on NICs

| Add-on Interfaces | |
|---|---|
| **Where labelled** | **What the label says** |
| TCP/IP Network Configuration Category in the Console<br>• TCP/IP tab<br>• TCP Status tab | Name - Slot1:Port0, Slot1:Port1, and so on<br>Description - Intel Gigabit Ethernet |
| Configuration Interface Name | Slot1:Port0<br>Slot1:Port1<br>and so on |
| Label on the back of the SSM | Port A<br>Port B<br>Port C<br>Port D |

The motherboard interfaces are labelled NICs 1 and 2 on the back of the SSM. Figure 49 illustrates the Intel® Storage System SSR212MA. The PCI slots for add-on interfaces are located to the right of the motherboard ports.

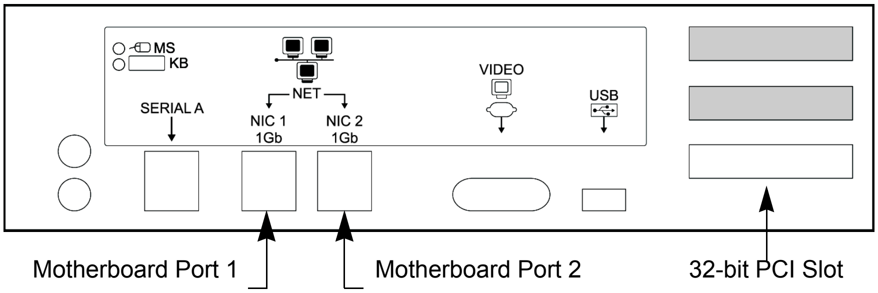

**Figure 49. Network Interface Ports and Open PCI Slot on the Back of the Intel® Storage System SSR212MA**

## Adding Interfaces to PCI Slots

You can add interface cards to the PCI slots located to the right of the motherboard NIC ports on the back of the SSM. These cards can contain Ethernet ports. The Intel® Storage System SSR212MA does not have support for Fibre Channel.

The other three covered slots are occupied by Serial ATA cards.

- The 64-bit PCI slot can hold a quad (4-port) card.
- The 32-bit slots can hold dual (2-port) cards.

To distribute bandwidth and to ensure fault tolerance, connect to ports across more than one PCI slot. For example, connect to the first port in the first (64-bit) PCI slot. Then connect to the next port in the second (32-bit) slot, and connect to the third port in the third (32-bit) slot. Connect to the fourth port in the first slot, and so on. The figure below shows the optimal configuration of add-on ports.

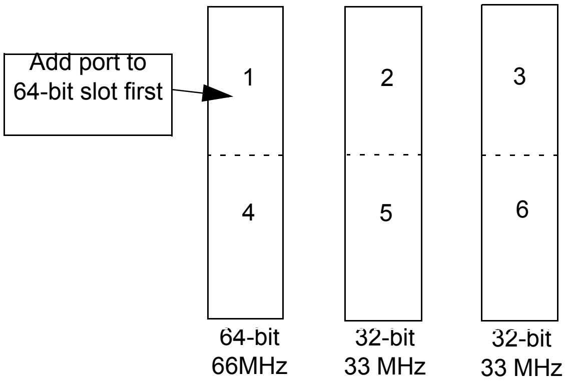

**Figure 50. Distributing Bandwidth and Ensuring Fault Tolerance of Add-on Ports Across PCI Slots**

**Note:** *When adding more than one port to the SSM, you can distribute bandwidth and to ensure fault tolerance by distributing the ports across more than one PCI slot. Start with the first (64-bit) slot.*

The Intel® Storage System SSR212MA contains one open 32-bit / 66 MHz PCI slot. This open 32-bit slot can hold a dual (2-port) or a quad (4-port) NIC card. The other two covered slots are occupied by SATA RAID controller cards.

# Configuring the IP Address Manually

Use the TCP/IP Network category in the edit configuration window to configure the IP address for an interface.

**Note:** *Any time you change an IP address of a storage module that is running a manager, the volumes on the storage module may become inaccessible to hosts configured to access the volume. You must reconfigure all hosts that are using that IP address.*

1. Select TCP/IP Network from the edit configuration categories.

   The window opens with the TCP/IP tab on top.

2. On the TCP/IP tab, select the interface from the list for which you want to configure or change the IP address.
3. Click Edit.

   The Edit TCP/IP Configuration window opens, shown in Figure 51.

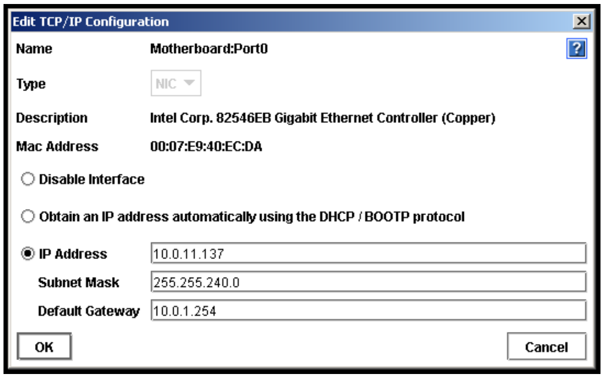

Edit TCP/IP Configuration

Name Motherboard:Port0

Type NIC

Description Intel Corp. 82546EB Gigabit Ethernet Controller (Copper)

Mac Address 00:07:E9:40:EC:DA

○ Disable Interface

○ Obtain an IP address automatically using the DHCP / BOOTP protocol

● IP Address 10.0.11.137

Subnet Mask 255.255.240.0

Default Gateway 10.0.1.254

**Figure 51. Configuring the IP Address Manually**

4. Select IP Address and complete the fields for IP Address, Subnet mask, and Default gateway.
5. Click OK.

   A confirmation message opens.

6. Click OK.

   A message notifying you of an automatic log out opens.

7. Click OK.

   The automatic log out occurs.

***Note:*** *Wait a few moments for the IP address change to take effect.*

8. Log in to the newly addressed storage module.

If you are changing the IP address of a storage module which is a manager in a management group, a window opens which displays all the IP addresses of the managers in the management group and a reminder to reconfigure the application servers that are affected by the change.

# Using DHCP

A DHCP server becomes a single point of failure in your system configuration. If the DHCP server goes down, then IP addresses may be lost.

***Warning:*** *If you use DHCP, be sure to reserve statically assigned IP addresses for all storage modules on the DHCP server. This is required because management groups use unicast communication.*

1. Select from the list the interface you want to configure for use with DHCP.
2. Click Edit.

   The Edit TCP/IP Configuration window opens, shown in Figure 51.
3. Select Obtain an address automatically using the DHCP/BOOTP protocol.
4. Click OK.

# Configuring NIC Bonding

Network interface bonding provides high availability, fault tolerance, and/or bandwidth aggregation for the network interface cards in the storage module. Bonds are created by “bonding” NICs into a single logical interface. This logical interface acts as the “master” interface, controlling and monitoring the physical “slave” interfaces.

Bonding two interfaces for failover provides fault tolerance at the local hardware level for network communication. Failures of NICs, Ethernet cables, individual switch ports, and/ or entire switches can be tolerated while maintaining data availability. Bonding two interfaces for aggregation provides bandwidth aggregation and localized fault tolerance.

Depending on your storage module hardware, network infrastructure design and Ethernet switch capabilities, you can bond NICs in one of three ways:

- **Active Backup.** You specify a preferred NIC for the bonded logical interface to use. If the preferred NIC fails, then the logical interface begins using another NIC in the bond until the preferred NIC resumes operation. When the preferred NIC resumes operation, data transfer resumes on the preferred NIC.

- **NIC Aggregation.** The logical interface uses both NICs simultaneously for data transfer. This configuration increases network bandwidth, and if one NIC fails, the other continues operating normally. NIC aggregation cannot be used on the 100.
- **Adaptive Load Balancing.** Adaptive Load Balancing (ALB), also known as asymmetric port aggregation, is a method of ensuring more throughput and transparent backup connections by using multiple network connections and balancing the data transmissions across them. ALB dynamically manages the NIC bond and evenly distributes the load among the network connections. It also delivers fault tolerance benefits; if one link fails, the other link continues to ensure network connectivity.

***Warning:*** *NIC aggregation requires plugging both NICs into the same switch. This bonding method does not protect against switch failure.*

You can create bonds of 2 or 4 NICs. A NIC can only be in one bond.

## Best Practices

NIC aggregation provides bandwidth gains because data is transferred over both NICs simultaneously. For NIC aggregation, both NICs must be plugged into the same switch, and that switch must be LACP-capable and support 802.3ad aggregation. Because both NICs are plugged into the same switch, NIC aggregation does not protect against switch failure.

For active backup, plug the two NICs on the storage module into separate switches. While NIC aggregation will only survive a port failure, active backup will survive a switch failure.

### NIC Bonding and Speed, Duplex, and Frame Size Settings

These settings are controlled on the TCP Status tab of the TCP/IP Network configuration category. If you change these settings, you must ensure that BOTH sides of the NIC cable are configured in the same manner. For example, if the storage module is set for Auto/Auto, the switch must be set the same. See "The TCP Status Tab" on page 108 for more information.

### Table 17. Comparison of Active-Passive, Link Aggregation Dynamic Mode and Adaptive Load Balancing Bonding

| Feature | Active-Passive | NIC Aggregation Dynamic Mode | Adaptive Load Balancing |
|---|---|---|---|
| Bandwidth | Use of 1 NIC at a time provides normal bandwidth. | Simultaneous use of both NICs increases bandwidth. | Simultaneous use of both NICs increases bandwidth. |
| Protection during port failure | Yes | Yes | Yes |
| Protection during switch failure | Yes (NICs are plugged into different switches) | No (Both NICs are plugged into the same switch) | Yes, NICs can be plugged into different switches |
| Requires support for 802.3ad link aggregation | No | Yes | No |

Allocate a static IP address for the logical bond interface (bond0). You cannot use DHCP for the bond IP.

# How Active Backup Works

Bonding NICs for active backup allows you to specify a preferred interface that will be used for data transfer. This is the active interface. The other interface acts as a backup, and its status is “Passive (Ready).”

## Physical and Logical Interfaces

The NICs in the SSM are labelled Motherboard:PortN and SlotN:PortN (where N is a number), depending on whether the NIC is located in the motherboard or in a PCI slot.

If 2 or 4 physical interfaces are bonded, the logical interface is labelled bondN and acts as the master interface. As the master interface, bondN controls and monitors the two physical slave interfaces.

### Table 18. Bonded Network Interfaces

| Interface Name | Description |
|---|---|
| bond0 | Logical Interface acting as master. |
| Motherboard:Port0 | Physical interface in the motherboard. This interface acts as a slave. |
| Slot1:Port0 | Physical interface in a PCI slot. This interface acts as a slave. |

The logical master bond interface monitors each physical slave interface to determine if its link to the device to which it is connected, such as a router, switch, or repeater, is up. As long as the interface link remains up, the interface status is preserved.

Table 19. Description of NIC Status in an Active Backup Configuration

| If the NIC Status is | The NIC is |
|---|---|
| Active | Currently enabled and in use |
| Passive (Ready) | Slave to a bond and available for failover |
| Passive (Failed) | Slave to a bond and no longer has a link |

If the active NIC fails, or if its link is broken due to a cable failure or a failure in a local device to which the NIC cable is connected, then the status of the NIC becomes Passive (Failed) and the other NIC in the bond, if it has a status of Passive (Ready), becomes active.

This configuration remains until the failed preferred interface is brought back online. When the failed interface is brought back online, it becomes Active. The other NIC returns to the Passive (Ready) state.

## Requirements for Active Backup

To configure active backup:

- Both NICs should be enabled.
- NICs should be connected to separate switches.

## Which Physical Interface is Preferred

A preferred interface is an interface within an active backup bond that is used for data transfer during normal operation. When you create an active backup bond, one of the interfaces becomes the preferred interface in the bond. You can change the preferred setting after creating the bond. See “Creating a NIC Bond” on page 100.

## Which Physical Interface is Active

When the active backup bond is created, if both NICs are plugged in, the preferred interface becomes the active interface. The other interface is Passive (Ready).

For example, suppose you create an active backup bond consisting of 2 NICs: Motherboard:Port0 and Slot1:Port0. If Motherboard:Port0 is the preferred interface, it will be active and Slot1:Port0 will be Passive (Ready). Then, if Motherboard:Port0 fails, Slot1:Port0 changes from Passive (Ready) to active. Motherboard:Port0 changes to Passive (Failed).

Once the link is fixed and Motherboard:Port0 is operational, there is a 30 second delay and then Motherboard:Port0 becomes the active interface. Slot1:Port0 returns to the Passive (Ready) state.

***Note:*** *When the preferred interface comes back up, there is a 30 second delay before it becomes active.*

### Table 20. Example Active Backup Failover Scenario and Corresponding NIC Status

| Example Failover Scenario | NIC Status |
|---|---|
| 1. Active backup bond0 is created. The active (preferred) interface is Motherboard:Port0. | • Bond0 is the master logical interface.<br>• Motherboard:Port0is Active.<br>• Slot1:Port0is connected and is Passive (Ready). |
| 2. Active interface fails. Bond0 detects the failure and Slot1:Port0takes over. | • Motherboard:Port0status becomes Passive (Failed).<br>• Slot1:Port0status changes to Active. |
| 3. The Motherboard:Port0link is restored. | • Motherboard:Port0status changes to Active after a 30 second delay.<br>• Slot1:Port0 status changes to Passive Ready). |

## Summary of NIC Status During Failover

Table 21 shows the states of Motherboard:Port0and Slot1:Port0when configured for Active Backup.

### Table 21. NIC Status During Failover with Active Backup

| Failover Status | Status of Motherboard: Port0 | Status of Slot1: Port0 |
|---|---|---|
| Normal Operation<br>↓ | Preferred: Yes<br>Status: Active<br>Data Transfer: Yes | Preferred: No<br>Status: Passive (Ready)<br>Data Transfer: No |

### Table 21. NIC Status During Failover with Active Backup

| Failover Status | Status of Motherboard: Port0 | Status of Slot1: Port0 |
|---|---|---|
| Motherboard: Port0 Fails,<br>Data Transfer Fails Over to Slot1: Port0<br>↓ | Preferred: Yes<br>Status: Passive (Failed)<br>Data Transfer: No | Preferred: No<br>Status: Active<br>Data Transfer: Yes |
| Motherboard: Port0 Restored | Preferred: Yes<br>Status: Active<br>Data Transfer: Yes | Preferred: No<br>Status: Passive (Ready)<br>Data Transfer: No |

## Example Network Configurations with Active Backup

Two simple network configurations using active backup in high availability environments are illustrated.

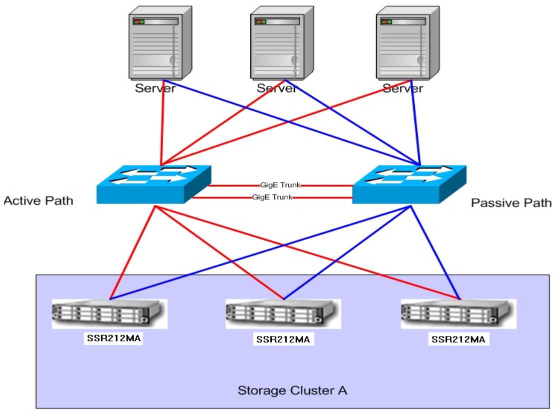

**Figure 52. Active Backup in a Two-switch Topology with Server Failover**

The two-switch scenario in Figure 52 is a basic, yet effective, method for ensuring high availability. If either switch failed, or a cable or NIC on one of the storage modules failed, the active backup bond would cause the secondary connection to become active and take over.

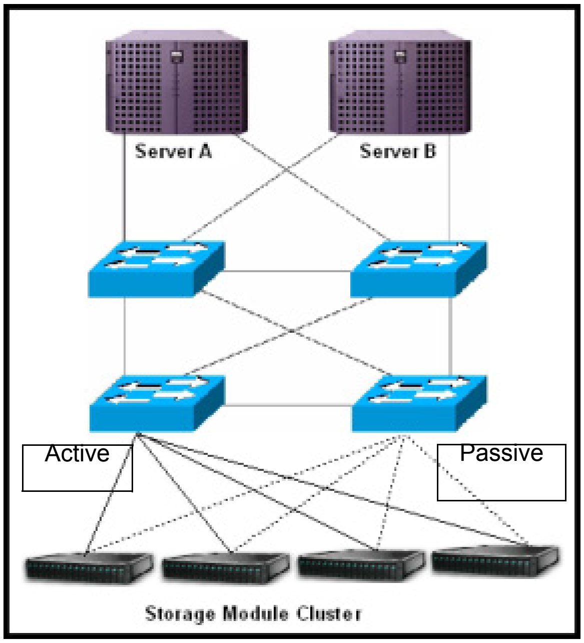

**Figure 53. Active Backup Failover in a Four-switch Topology**

Figure 53 illustrates the active backup configuration in a four-switch topology.

# How NIC Aggregation Works

NIC aggregation allows the storage module to use both interfaces simultaneously for data transfer. Both interfaces have an active status. If the interface link to one NIC goes down, the other interface continues operating. Using both NICs also increases network bandwidth.

## Requirements for NIC Aggregation

To configure NIC aggregation:

- Both NICs should be enabled.
- NICs must be configured to the same subnet.
- NICs must be connected to a single switch that is LACP-capable and supports 802.3ad link aggregation. If the storage module is directly connected to a server, then the server must support 802.3ad link aggregation.

## Which Physical Interface is Preferred

Because the logical interface uses both NICs simultaneously for data transfer, neither of the NICs in an aggregation bond are designated as preferred.

## Which Physical Interface is Active

When the NIC aggregation bond is created, if both NICs are plugged in, both interfaces are active. If one interface fails, the other interface continues operating. For example, suppose Motherboard:Port0 and Slot1:Port0 are bonded in a NIC Aggregation bond. If Motherboard:Port0fails, then Slot1:Port0 remains active.

Once the link is fixed and Motherboard:Port0is operational, it becomes active again. Slot1:Port0 remains active.

### Table 22. NIC Aggregation Failover Scenario and Corresponding NIC Status

| Example Failover Scenario | NIC Status |
|---|---|
| 1. NIC aggregation bond0 is created. Motherboard:Port0 and Slot1:Port0 are both active. | • Bond0 is the master logical interface.<br>• Motherboard:Port0 is Active.<br>• Slot1:Port0 is Active. |
| 2. Motherboard:Port0 interface fails. Because NIC aggregation is configured, Slot1:Port0 continues operating. | • Motherboard:Port0 status becomes Passive (Failed).<br>• Slot1:Port0 status remains Active. |
| 3. Motherboard:Port0 link failure is repaired. | • Motherboard:Port0 resumes Active status.<br>• Slot1:Port0 remains Active. |

## Summary of NIC States During Failover

Table 23 shows the states of Motherboard:Port0 and Slot1:Port0 when configured for NIC aggregation.

### Table 23. NIC Status During Failover with NIC Aggregation

| Failover Status | Status of Motherboard: Port0 | Status of Slot1: Port0 |
|---|---|---|
| Normal Operation<br>↓ | Preferred: No<br>Status: Active<br>Data Transfer: Yes | Preferred: No<br>Status: Active<br>Data Transfer: Yes |
| Motherboard: Port0 Fails,<br>Data Transfer Continues on Slot1: Port0<br>↓ | Preferred: No<br>Status: Passive (Failed)<br>Data Transfer: No | Preferred: No<br>Status: Active<br>Data Transfer: Yes |
| Motherboard: Port0 Restored | Preferred: No<br>Status: Active<br>Data Transfer: Yes | Preferred: No<br>Status: Active<br>Data Transfer: Yes |

## Example Network Configurations with NIC Aggregation

A simple network configuration using NIC aggregation in a high availability environment is illustrated.

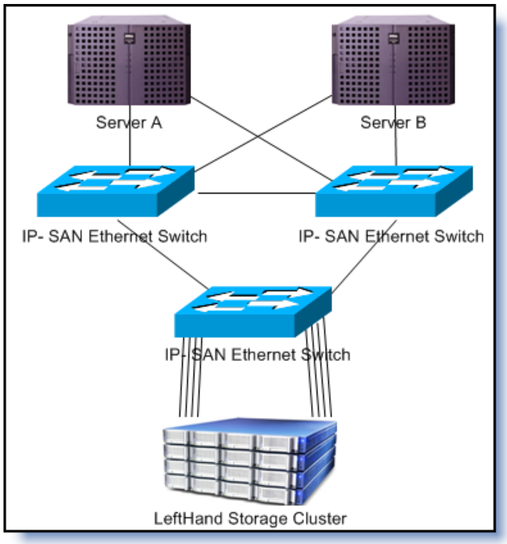

**Figure 54. NIC Aggregation in a Single-switch Topology**

# How Adaptive Load Balancing Works

Adaptive Load Balancing allows the storage module to use both interfaces simultaneously for data transfer. Both interfaces have an active status. If the interface link to one NIC goes down, the other interface continues operating. Using both NICs also increases network bandwidth.

## Requirements for Adaptive Load Balancing

To configure Adaptive Load Balancing:

- Both NICs must be enabled.
- NICs must be configured to the same subnet.
- NICs can be connected to separate switches.

## Which Physical Interface is Preferred

Because the logical interface uses both NICs for data transfer, neither of the NICs in an Adaptive Load Balancing bond are designated as preferred.

## Which Physical Interface is Active

When the Adaptive Load Balancing bond is created, if both NICs are plugged in, both interfaces are active. If one interface fails, the other interface continues operating. For example, suppose Motherboard:Port1 and Motherboard:Port2 are bonded in an Adaptive Load Balancing bond. If Motherboard:Port1 fails, then Motherboard:Port2 remains active.

Once the link is fixed and Motherboard:Port1 is operational, it becomes active again. Motherboard:Port2 remains active.

### Table 24. Example Adaptive Load Balancing Failover Scenario and Corresponding NIC Status

| Example Failover Scenario | NIC Status |
|---|---|
| 1. Adaptive Load Balancing bond0 is created. Motherboard:Port1 and Motherboard:Port2 are both active. | • Bond0 is the master logical interface.<br>• Motherboard:Port1 is Active.<br>• Motherboard:Port2 is Active. |
| 2. Motherboard:Port1 interface fails. Because Link Aggregation Dynamic Mode is configured, Motherboard:Port2 continues operating. | • Motherboard:Port1 status becomes Passive (Failed).<br>• Motherboard:Port2 status remains Active. |
| 3. Motherboard:Port1 link failure is repaired. | • Motherboard:Port1 resumes Active status.<br>• Motherboard:Port2 remains Active. |

## Summary of NIC States During Failover

The following table shows the states of Motherboard:Port1 and Motherboard:Port2 when configured for Adaptive Load Balancing.

### Table 25. NIC Status During Failover with Adaptive Load Balancing

| Failover Status | Status of Motherboard: Port1 | Status of Mother-board: Port2 |
|---|---|---|
| Normal Operation<br>↓ | Preferred: No<br>Status: Active<br>Data Transfer: Yes | Preferred: No<br>Status: Active<br>Data Transfer: Yes |
| Motherboard: Port1 Fails,<br>Data Transfer Fails Over to Motherboard: Port2<br>↓ | Preferred: No<br>Status: Passive (Failed)<br>Data Transfer: No | Preferred: No<br>Status: Active<br>Data Transfer: Yes |
| Motherboard: Port1 Restored | Preferred: No<br>Status: Active<br>Data Transfer: Yes | Preferred: No<br>Status: Active<br>Data Transfer: Yes |

## Example Network Configurations with Adaptive Load Balancing

A simple network configuration using Adaptive Load Balancing in a high availability environment is illustrated.

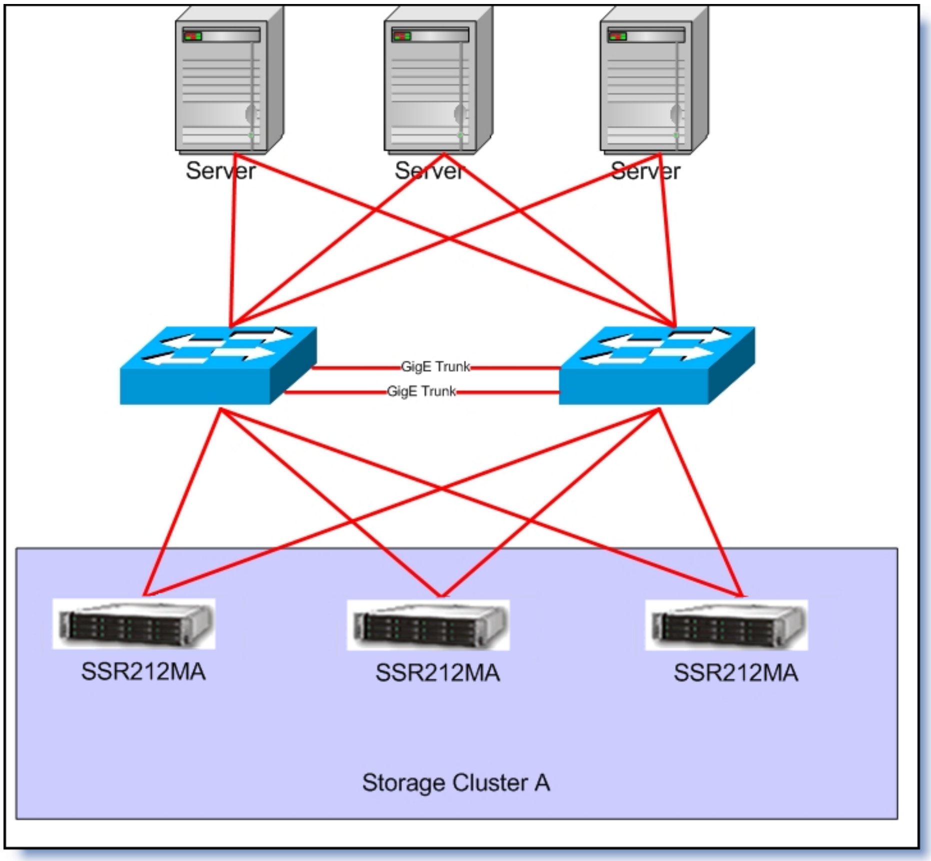

**Figure 55. Adaptive Load Balancing in a Two-switch Topology**

# Creating a NIC Bond

Follow these guidelines when creating NIC bonds:

- Create bonds of 2 or 4 interfaces.
- You can create more than one bond on an SSM.
- An interface can only be in one bond.
- To provide failover capability in the event of a PCI card failure, you should bond interfaces that are located in the motherboard with interfaces that are in PCI slots. This ensures that if an entire PCI card fails, then the bonded interface will use an interface in the motherboard to continue operating.
- Record the configuration information of each interface before you create the bond. Then, if you delete the bond, you can return to the original configuration if desired.

- — When you delete an active backup bond, the preferred interface assumes the IP address and configuration of the deleted logical interface.
- — When you delete a NIC aggregation bond or an Adaptive Load Balancing bond, one of the interfaces retains the IP address of the deleted logical interface. The IP address of the other interface is set to 0.0.0.0.

- Create a bond on an storage module before you add the storage module to a management group.
- Ensure that the bond has a static IP address for the logical bond interface. The default values for the IP address, subnet/mask and default gateway are those of one of the physical interfaces.
- Verify on the Communication tab that the Storage System Software interface is communicating with the bonded interface.

***Warning:*** *To ensure that the bond works correctly, you should configure it as follows:*
*- Create the bond on the storage module before you add it to a management group.*
*- Verify that the bond is created.*
*If you create the bond on the storage module after it is in a management group, and if it does not work correctly, you might*
*- lose the storage module from the network.*
*- lose quorum in the management group for a while.*
*See the chapter "Using the Configuration Interface" for information about deleting NIC bonds through the Configuration Interface.*

## Creating the Bond

1. Remove the storage module from the management group.
2. Log in to the storage module.
3. On the TCP/IP tab, shown in Figure 56, select 2 or 4 NICs to bond.

   The NICs that you select do not have to be consecutive NICs in the list.

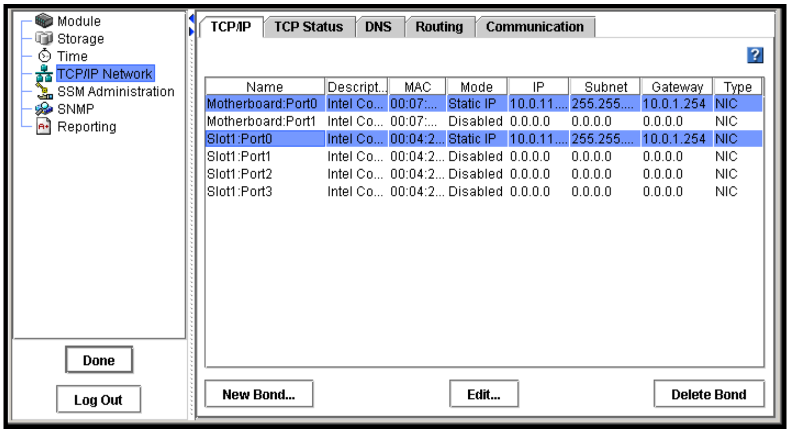

| Name | Descript... | MAC | Mode | IP | Subnet | Gateway | Type |
|---|---|---|---|---|---|---|---|
| Motherboard:Port0 | Intel Co... | 00:07:... | Static IP | 10.0.11.... | 255.255.... | 10.0.1.254 | NIC |
| Motherboard:Port1 | Intel Co... | 00:07:... | Disabled | 0.0.0.0 | 0.0.0.0 | 0.0.0.0 | NIC |
| Slot1:Port0 | Intel Co... | 00:04:2... | Static IP | 10.0.11.... | 255.255.... | 10.0.1.254 | NIC |
| Slot1:Port1 | Intel Co... | 00:04:2... | Disabled | 0.0.0.0 | 0.0.0.0 | 0.0.0.0 | NIC |
| Slot1:Port2 | Intel Co... | 00:04:2... | Disabled | 0.0.0.0 | 0.0.0.0 | 0.0.0.0 | NIC |
| Slot1:Port3 | Intel Co... | 00:04:2... | Disabled | 0.0.0.0 | 0.0.0.0 | 0.0.0.0 | NIC |

**Figure 56. Selecting Motherboard:Port0 and Slot1:Port0 for a New Bond**

4. Click New Bond.

   The Create Bond Configuration window opens, shown in Figure 57.

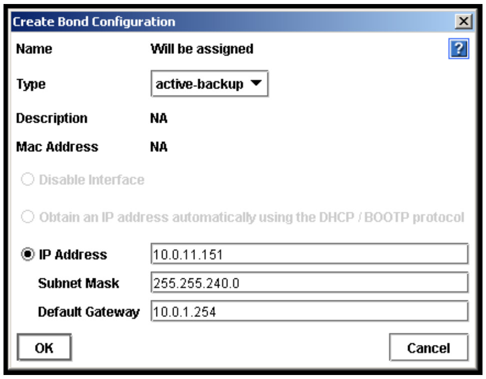

Name: Will be assigned

Type: active-backup

Description: NA

Mac Address: NA

○ Disable Interface

○ Obtain an IP address automatically using the DHCP / BOOTP protocol

● IP Address: 10.0.11.151

Subnet Mask: 255.255.240.0

Default Gateway: 10.0.1.254

**Figure 57. Creating a NIC Bond**

5. Select a bond type from the list.
6. Enter an IP address for the bond.
7. Enter the Subnet mask.
8. [Optional] Enter the default gateway.
9. Click OK.

   A confirmation message opens.

10. Click OK to confirm the TCP/IP changes.

    A message opens, shown in Figure 58, prompting you to search for the bonded storage module on the network.

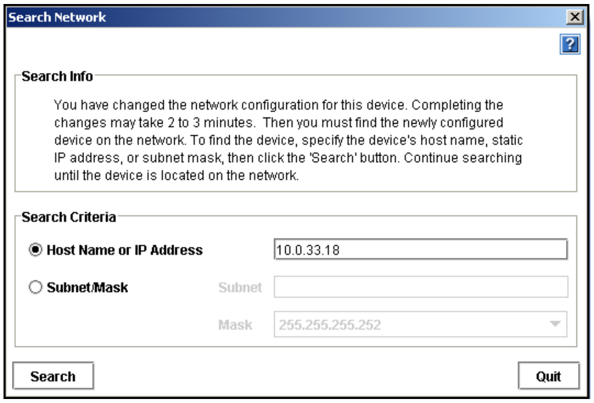

Search Info

You have changed the network configuration for this device. Completing the changes may take 2 to 3 minutes. Then you must find the newly configured device on the network. To find the device, specify the device's host name, static IP address, or subnet mask, then click the 'Search' button. Continue searching until the device is located on the network.

Search Criteria

● Host Name or IP Address: 10.0.33.18

○ Subnet/Mask

Subnet:

Mask: 255.255.255.252

**Figure 58. Searching for the Bonded Storage Module on the Network**

11. Search for the storage module by Host Name or IP address, or by Subnet/mask.

    ***Note:*** *Because it can take a few minutes for the storage module to reinitialize, the search may fail the first time. If the search fails, wait a minute or two and choose Try Again on the Network Search Failed message.*

12. Verify the new bond interface.

    The TCP/IP tab displays the new list of interfaces, as shown in Figure 59.

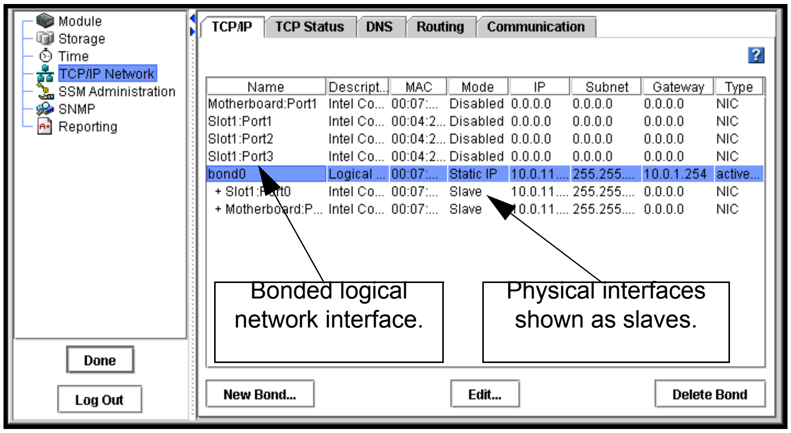

| Name | Descript... | MAC | Mode | IP | Subnet | Gateway | Type |
|---|---|---|---|---|---|---|---|
| Motherboard:Port1 | Intel Co... | 00:07:... | Disabled | 0.0.0.0 | 0.0.0.0 | 0.0.0.0 | NIC |
| Slot1:Port1 | Intel Co... | 00:04:2... | Disabled | 0.0.0.0 | 0.0.0.0 | 0.0.0.0 | NIC |
| Slot1:Port2 | Intel Co... | 00:04:2... | Disabled | 0.0.0.0 | 0.0.0.0 | 0.0.0.0 | NIC |
| Slot1:Port3 | Intel Co... | 00:04:2... | Disabled | 0.0.0.0 | 0.0.0.0 | 0.0.0.0 | NIC |
| bond0 | Logical ... | 00:07:... | Static IP | 10.0.11.... | 255.255.... | 10.0.1.254 | active... |
| + Slot1:Port0 | Intel Co... | 00:07:... | Slave | 10.0.11.... | 255.255.... | 0.0.0.0 | NIC |
| + Motherboard:P... | Intel Co... | 00:07:... | Slave | 10.0.11.... | 255.255.... | 0.0.0.0 | NIC |

Bonded logical network interface.

Physical interfaces shown as slaves.

**Figure 59. Viewing a New Active Backup Bond**

The bond interface shows as “bond0” and has a static IP address. The two physical NICs now show up as slaves in the Mode column.

13. [Optional, for active backup bonds only] To change which interface is the preferred interface in an active backup bond, on the TCP Status tab select one of the NICs in the bond and click Set Preferred.

## Verify Communication Setting for New Bond

1. Select the Communication tab, shown in Figure 60.

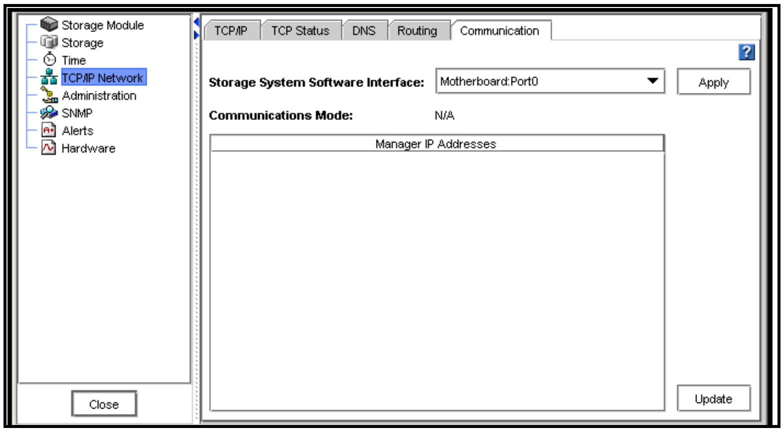

Storage System Software Interface: Motherboard:Port0

Communications Mode: N/A

Manager IP Addresses

**Figure 60. Verifying Interface Used for Storage System Software Communication**

2. Verify that the Storage System Software communication port is correct.

# Viewing the Status of a NIC Bond

You can view the status of the interfaces on the TCP Status tab. Notice that in the active backup bond, one of the NICs is the preferred NIC. In both the NIC aggregation bond, neither physical interface is preferred.

Figure 61 shows the status of interfaces in an active backup bond. Figure 62 shows the status of interfaces in a NIC aggregation bond.

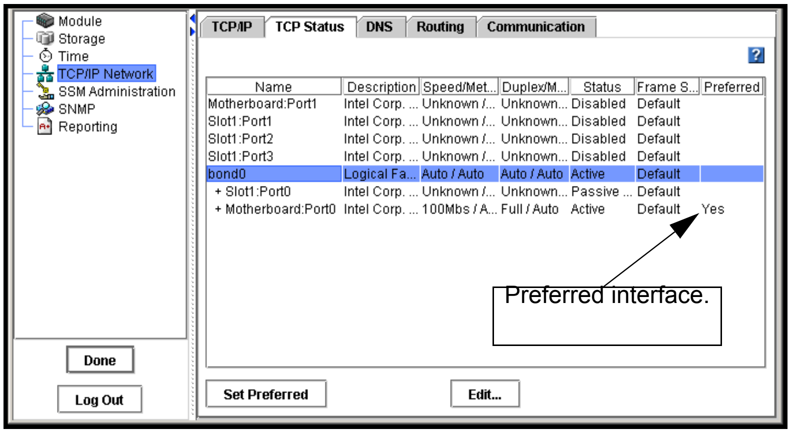

| Name | Description | Speed/Met... | Duplex/M... | Status | Frame S... | Preferred |
|---|---|---|---|---|---|---|
| Motherboard:Port1 | Intel Corp. ... | Unknown /... | Unknown... | Disabled | Default | |
| Slot1:Port1 | Intel Corp. ... | Unknown /... | Unknown... | Disabled | Default | |
| Slot1:Port2 | Intel Corp. ... | Unknown /... | Unknown... | Disabled | Default | |
| Slot1:Port3 | Intel Corp. ... | Unknown /... | Unknown... | Disabled | Default | |
| bond0 | Logical Fa... | Auto / Auto | Auto / Auto | Active | Default | |
| + Slot1:Port0 | Intel Corp. ... | Unknown /... | Unknown... | Passive ... | Default | |
| + Motherboard:Port0 | Intel Corp. ... | 100Mbs / A... | Full / Auto | Active | Default | Yes |

**Figure 61. Viewing the Status of an Active Backup Bond**

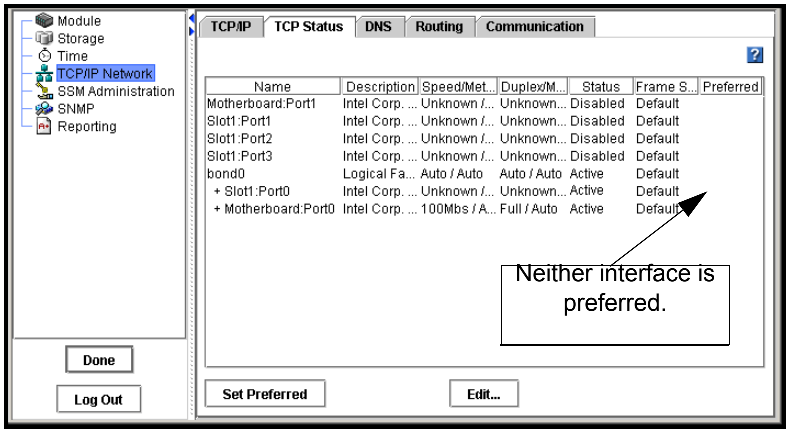

| Name | Description | Speed/Met... | Duplex/M... | Status | Frame S... | Preferred |
|---|---|---|---|---|---|---|
| Motherboard:Port1 | Intel Corp. ... | Unknown /... | Unknown... | Disabled | Default | |
| Slot1:Port1 | Intel Corp. ... | Unknown /... | Unknown... | Disabled | Default | |
| Slot1:Port2 | Intel Corp. ... | Unknown /... | Unknown... | Disabled | Default | |
| Slot1:Port3 | Intel Corp. ... | Unknown /... | Unknown... | Disabled | Default | |
| bond0 | Logical Fa... | Auto / Auto | Auto / Auto | Active | Default | |
| + Slot1:Port0 | Intel Corp. ... | Unknown /... | Unknown... | Active | Default | |
| + Motherboard:Port0 | Intel Corp. ... | 100Mbs / A... | Full / Auto | Active | Default | |

Neither interface is preferred.

**Figure 62. Viewing the Status of a NIC Aggregation Bond**

***Note:*** *If the bonded NIC experiences rapid, sequential Ethernet failures, the Console may display the storage module as failed (flashing red) and access to data on that storage module fails. However, as soon as the Ethernet connection is re-established, the storage module and the Console display the correct information.*

# Deleting a NIC Bond

When you delete an active backup bond, the preferred interface assumes the IP address and configuration of the deleted logical interface. All other interfaces are disabled and the IP addresses are set to 0.0.0.0.

When you delete either a NIC aggregation bond, one of the active interfaces in the bond retains the IP address of the deleted logical interface. All other interfaces are disabled and the IP addresses are set to 0.0.0.0.

1. On the TCP/IP tab, select the bond that you want to delete.

2. Click Delete Bond.

   Because the IP addresses change, the Search Network window opens.

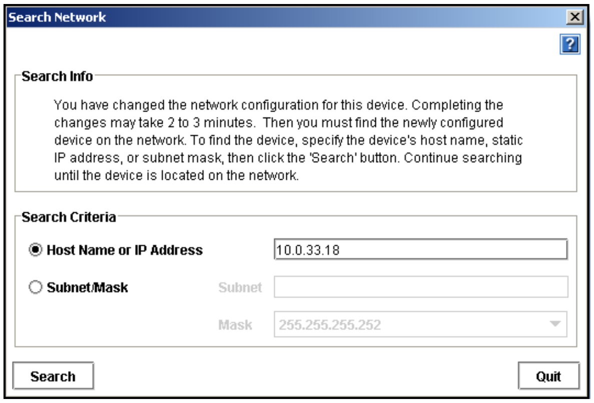

Search Network

Search Info

You have changed the network configuration for this device. Completing the changes may take 2 to 3 minutes. Then you must find the newly configured device on the network. To find the device, specify the device's host name, static IP address, or subnet mask, then click the 'Search' button. Continue searching until the device is located on the network.

Search Criteria

● Host Name or IP Address: 10.0.33.18

○ Subnet/Mask

Subnet:

Mask: 255.255.255.252

**Figure 63. Searching for the Unbonded Storage Module on the Network**

3. Search for the storage module by Host Name or IP Address or Subnet/Mask.

***Note:*** *Because it can take a few minutes for the storage module to reinitialize, the search may fail the first time. If the search fails, wait a minute or two and choose Try Again on the Network Search Failed message.*

## Verify Communication Setting After Deleting a Bond

1. Select the Communication tab, shown in Figure 60.

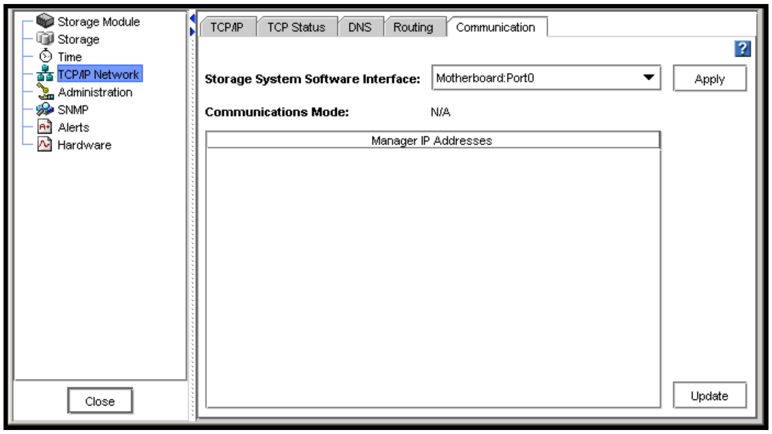

Storage System Software Interface: Motherboard:Port0

Communications Mode: N/A

Manager IP Addresses

**Figure 64. Verifying Interface Used for Storage System Software Communication**

2. Verify that the Storage System Software communication port is correct.

# Disabling a Network Interface

You can disable the network interfaces on the storage module.

- You can only disable top-level interfaces. This includes bonded interfaces and NICs that are not part of bonded interfaces.
- To ensure that you always have access to the storage module, do not disable the last interface. If you want to disable the last interface, first enable another interface.

***Warning:*** *If you disable an interface, be sure you enable another interface first. That way you always have access to the storage module.*
*If you disable all the interfaces, you must reconfigure at least one interface using the Configuration Interface to access the storage module. See "Configuring a Network Connection" on page 300.*

## Disabling a Network Interface

1. Select from the list on the TCP/IP window the interface to disable.
2. Click Edit.

   The Edit TCP/IP Configuration window opens, shown in Figure 51.
3. Click Disable Interface.
4. Click OK.

   A confirmation message opens. If you are disabling the only interface, the message warns that the storage module may be inaccessible if you continue.
5. Click OK.

### If the Storage Module is in a Management Group

If the storage module for which you are disabling the interface is a manager in a management group, a window opens which displays all the IP addresses of the managers in the management group and a reminder to reconfigure the application servers that are affected by the update.

# Configuring a Disabled Interface

If one interface is still connected to the storage module but another interface is disconnected, you can reconnect to the second interface using the Console. See "Configuring the IP Address Manually" on page 89.

If both interfaces to the storage module are disconnected, you must attach a terminal, or PC or laptop to the storage module with a null modem cable and configure at least one interface using the Configuration Interface. See 'Using the Configuration Interface," on page 297.

# TCP Status

Review the status of the TCP interfaces. Change the speed and duplex method of an interface.

## The TCP Status Tab

Review the status of the network interfaces on the TCP Status tab, shown in Figure 65.

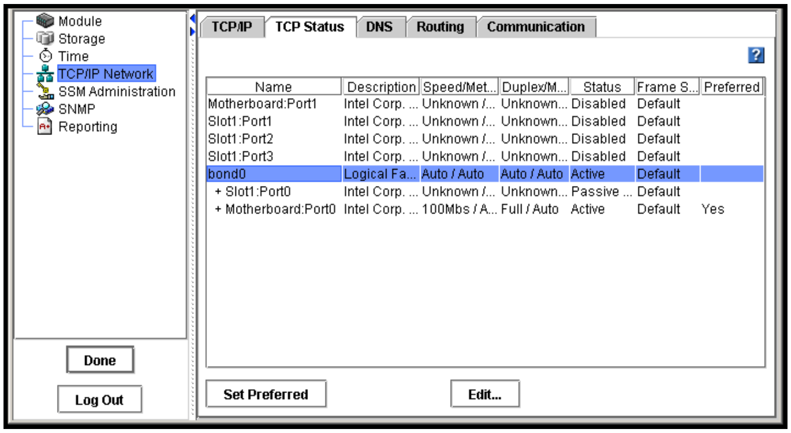

| Name | Description | Speed/Met... | Duplex/M... | Status | Frame S... | Preferred |
|---|---|---|---|---|---|---|
| Motherboard:Port1 | Intel Corp. ... | Unknown /... | Unknown... | Disabled | Default | |
| Slot1:Port1 | Intel Corp. ... | Unknown /... | Unknown... | Disabled | Default | |
| Slot1:Port2 | Intel Corp. ... | Unknown /... | Unknown... | Disabled | Default | |
| Slot1:Port3 | Intel Corp. ... | Unknown /... | Unknown... | Disabled | Default | |
| bond0 | Logical Fa... | Auto / Auto | Auto / Auto | Active | Default | |
| + Slot1:Port0 | Intel Corp. ... | Unknown /... | Unknown... | Passive ... | Default | |
| + Motherboard:Port0 | Intel Corp. ... | 100Mbs / A... | Full / Auto | Active | Default | Yes |

### Figure 65. Viewing the TCP Status

### Table 26. Status Of and Information About Network Interfaces

| Column | Description |
|---|---|
| Name | Name of the interface. Entries are<br>• Motherboard:Port0<br>• Motherboard:Port1<br>• Slot1:Port0<br>◎ bond0 - the bonded interface(s) [displays only if storage module configured for bonding] |
| Description | Describes each interface listed. For example, the bond0 is the Logical Failover Device. |
| Speed/Method | Lists the actual operating speed reported by the device. |
| Duplex/Method | Lists duplex as reported by the device. |
| Status | Describes the state of the interface. See Table 19 for a detailed description of individual NIC status. |
| Frame Size | Lists the frame size setting for the device. |
| Preferred | [For active backup bonds] Indicates whether the device is set as preferred. The preferred interface is the interface within an active backup bond that is used for data transfer during normal operation. |

# Editing the TCP Speed and Duplex

Change  the speed and duplex of the 1000BASE-T TCP interfaces.

## Requirements

- If you change these settings, you must ensure that BOTH sides of the NIC cable are configured in the same manner. For example, if the storage module is set for Auto/ Auto, the switch must be set the same.
- If you edit the speed or duplex on a disabled or failed NIC, the new setting will not be applied until the NIC is enabled or connectivity is restored.

## Best Practice

Change the speed and duplex settings while the storage module is in the Available mode and not in a management group.

**Table 27.   Setting Storage Module Speed and Duplex Settings**

| Storage Module Setting Speed/Duplex | Switch Setting Speed/Duplex |
|---|---|
| Auto/Auto | Auto/Auto |
| 1000/Full | 1000/Full |
| 100/Full | 100/Full |
| 100/Half | 100/Half |
| 10/Full | 10/Full |
| 10/Half | 10/Half |

1. On the TCP Status tab, select the interface you want to edit.
2. Click Edit.

   The Edit Speed and Duplex window opens, shown in Figure 66.

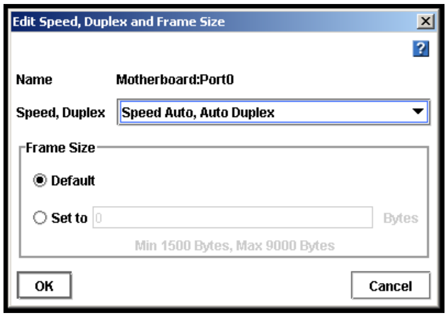

Edit Speed, Duplex and Frame Size

Name Motherboard:Port0

Speed, Duplex Speed Auto, Auto Duplex

Frame Size

● Default

○ Set to 0 Bytes

Min 1500 Bytes, Max 9000 Bytes

**Figure 66. Editing TCP Speed, Duplex, and Frame Size**

3. Select the combination of speed and duplex that you want.

4. Click OK.

   A series of status messages display. Then the changed setting displays in the TCP status report.

**Note:** *You can also use the Configuration Interface to edit the TCP speed and duplex. See "Setting the TCP Speed, Duplex, and Frame Size" on page 302.*

# Editing the NIC Frame Size

The frame size specifies the size of data packets that are transferred over the network. The default Ethernet standard frame size is 1500 bytes. The maximum allowed frame size is 9000 bytes.

Increasing the frame size improves data transfer speed by allowing larger packets to be transferred over the network and by decreasing the CPU processing time required to transfer data. However, increasing the frame size requires that routers, switches, and other devices on your network support that frame size.

**Note:** *Increasing the frame size can cause decreased performance and other network problems if routers, switches, or other devices on your network do not support frame sizes greater than 1500 bytes. If you are unsure about whether your routers and other devices support larger frame sizes, keep the frame size at the default setting.*

**Note:** *If you edit the frame size on a disabled or failed NIC, the new setting will not be applied until the NIC is enabled or connectivity is restored.*

## Best Practice

To avoid potential connectivity and performance problems with other devices on your network, keep the frame size at the default setting. If you decide to change the frame size, set the same frame size on all storage modules on the network, and set compatible frame sizes on all clients.

The frame size on the storage module should correspond to the frame size on Windows and Linux application servers. Table 28  shows recommended storage module frame sizes and the corresponding frame sizes in bytes for Windows and Linux clients.

**Table 28. Setting Corresponding Frame Sizes in Bytes on Storage Modules and Windows* or Linux Clients**

| Storage Module Frame Size | Windows Client Frame Size | Linux Client Frame Size |
|---|---|---|
| 1500 (Default) | 1542 (Default) | 1500 (Default) |
| 4046 | 4088 | 4046 |
| 8972 | 9014 | 8972 |

Frame sizes greater than 1500 bytes, called jumbo frames, can co-exist with 1500 byte frames on the same subnet if the following conditions are met:

- Every device downstream of the storage module on the subnet must support jumbo frames.
- If you are using 802.1q virtual LANs, jumbo frames and non-jumbo frames must be segregated into separate VLANs.

# Best Practice

Change the speed and duplex settings while the storage module is in the Available mode and not in a management group.

**Note:** *The frame size for a bonded logical interface must be equal to the frame size of the NICs in the bond.*

## Editing the Frame Size

To edit the frame size:

1. On the TCP Status tab, select the interface you want to edit.
2. Click Edit.

   The Edit Speed, Duplex, and Frame Size window opens, shown in Figure 67.

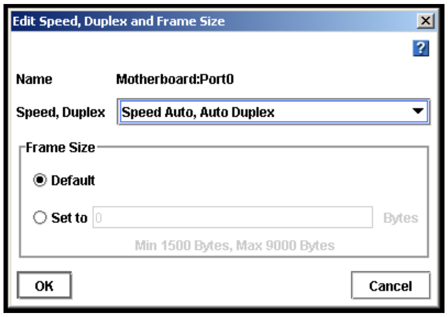

Edit Speed, Duplex and Frame Size

Name Motherboard:Port0

Speed, Duplex Speed Auto, Auto Duplex

Frame Size

● Default

○ Set to 0 Bytes

Min 1500 Bytes, Max 9000 Bytes

Figure 67. Editing TCP Speed, Duplex, and Frame Size

3. Select Set To in the Frame Size section.
4. Enter a value between 1500 and 9000 bytes in the Set To field.
5. Click OK.

   A series of status messages display. Then the changed setting displays in the TCP status report.

**Note:** *You can also use the Configuration Interface to edit the frame size. See "Setting the TCP Speed, Duplex, and Frame Size" on page 302.*

# Using a DNS Server

The storage module can use a DNS server to resolve host names. For example, if you enter a host name to specify an NTP time server, the storage module will use DNS to resolve the host name to its IP address. For example, the time server in Boulder, Colorado has a host name of `time.nist.gov`. DNS resolves this host name to its IP address of 192.43.244.18.

## DNS and DHCP

If you configure the storage module to use DHCP to obtain an IP address, and if the DHCP server is configured to provide the IP addresses of the DNS servers, then a maximum of three DNS servers will automatically be added to the storage module. These DNS servers are listed as IP addresses in the storage module configuration window in the TCP/IP Network category on the DNS tab. You can remove these DNS servers, but the storage module will not be able to resolve host names until you enter a new DNS server.

## DNS and Static IP Addresses

If you assigned a static IP address to the storage module and you want the storage module to recognize host names, you must manually add a DNS server to the Network DNS tab.

***Note:*** *If you initially set up the storage module to use DHCP and then change the configuration to use a static IP address, the DNS server provided by DHCP will remain on the DNS tab. You can remove or change this DNS server.*

1. On the Network View, double-click the storage module and log in, if necessary.
2. The Edit Configuration window opens. Select TCP/IP Network from the configuration categories.

3. Click the DNS tab to bring it to the front, shown in Figure 68.

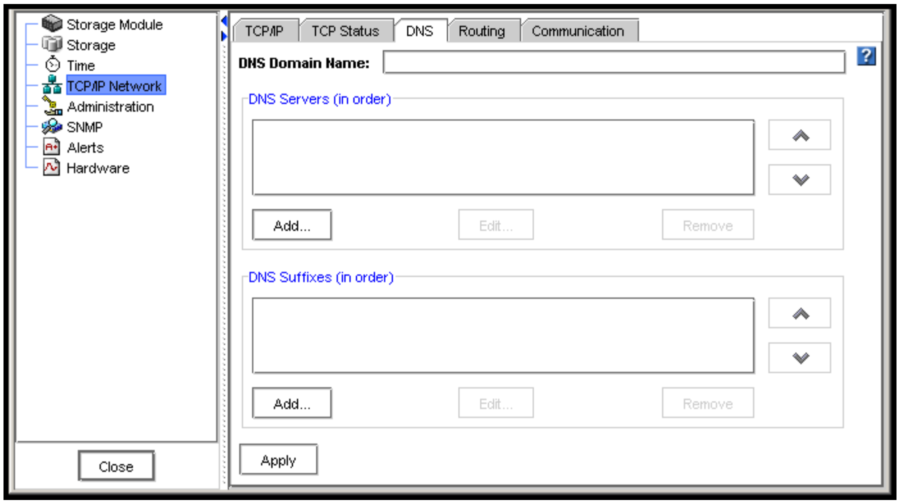

**Figure 68. Adding DNS Servers**

# Adding the DNS Domain Name

Add the name of the DNS domain in which the storage module resides.

1. On the DNS tab, type the DNS domain name.
2. Click Apply when you are finished.

# Adding a DNS Server

Add up to three DNS servers for use with the storage module.

1. Click Add in the DNS Server panel.

   The Add IP Address dialog opens.
2. Type the IP address for the DNS server.
3. Click OK.
4. Repeat steps 1 through 3 to add up to three servers.
5. Use the arrows on the DNS Server panel to order the servers.

   The servers will be accessed in the order they appear in the list.
6. Click Apply when you are finished.

## Adding Domain Names to the DNS Suffixes

Add up to six domain names to the DNS suffix list (also known as the look up zone). The storage module searches the suffixes first and then uses the DNS server to resolve host names.

1. Click Add in the DNS Suffixes panel.

   The Add DNS Suffix window opens.
2. Type the DNS suffix name. Use the domain name format.
3. Click OK.
4. Repeat steps 1 through 3 to add up to six domain names.
5. Click Apply when you are finished.

## Editing a DNS Server

Change the IP address for a DNS Server in the list.

1. Select the server to edit.
2. Click Edit.

   The Edit IP Address window opens.
3. Type the new IP address for the DNS server.
4. Click OK.
5. Click Apply when you are finished.

## Editing a Domain Name in the DNS Suffixes List

Change a domain name in the DNS Suffixes list.

1. Select the domain name to edit.
2. Click Edit.

   The Edit DNS Suffix window opens.
3. Enter the change to the domain name.
4. Click OK.
5. Click Apply when you are finished.

## Removing a DNS Server

Remove a DNS server from the list.

1. Select the server you want to remove from the DNS Servers list.

2. Click Remove.

   A confirmation message opens.
3. Click OK to remove the DNS server from the list.
4. Click Apply when you are finished.

## Removing a Domain Name from the DNS Suffixes List

1. Select the domain name you want to remove from the DNS Suffixes list.
2. Click Remove.

   A confirmation message opens.
3. Click OK to remove the domain name from the list.
4. Click Apply when you are finished.

# Routing Overview

The Routing tab displays the routing table. You can specify static routes and/or a default route. If you specify a default route here, it will not survive a reboot or shut down of the storage module. To create a route that will survive a storage module reboot or shut down, you must enter a default gateway on the TCP/IP tab. See "Configuring the IP Address Manually" on page 89.

Information for each route listed includes the device, the network, gateway, and mask, and flags.

## Adding Routing Information

Open the Edit Configuration window to access the TCP/IP Network configuration category.

1. Select a storage module in the navigation window.
2. Click Storage Module Tasks in the tab window.
3. Select Edit Configuration.

   The Edit Configuration window opens.
4. Select TCP/IP Network from the configuration categories.

5. Click the Routing tab to bring it to the front, shown in Figure 69.

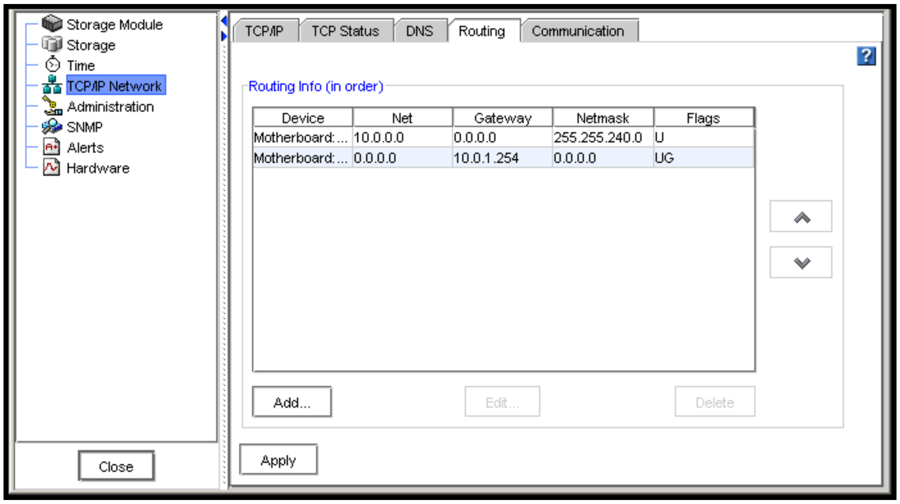

Routing Info (in order)

| Device | Net | Gateway | Netmask | Flags |
|---|---|---|---|---|
| Motherboard:... | 10.0.0.0 | 0.0.0.0 | 255.255.240.0 | U |
| Motherboard:... | 0.0.0.0 | 10.0.1.254 | 0.0.0.0 | UG |

Figure 69. Adding Network Routing Information

6. Click Add.

   The Add Routing Information dialog opens, shown in Figure 70.

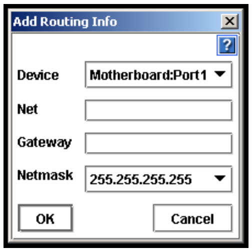

Device: Motherboard:Port1

Net:

Gateway:

Netmask: 255.255.255.255

Figure 70. Adding Routing Information

7. Select the port to use for routing in the Device list.
8. Type the IP address portion of the network address in the Net field.
9. Type the IP address of the router in the Gateway field.
10. Select the netmask.
11. Click OK.
12. Use the arrows on the routing table panel to order devices according to the needs of your network.

    The storage module will attempt to use the routes in the order in which they are listed.
13. Click Apply when you are finished.

# Editing Routing Information

You can only edit optional routes you have added.

1. On the routing tab, select the optional route you want to change.
2. Click Edit.

   The Edit Routing Information dialog opens, shown in Figure 71.

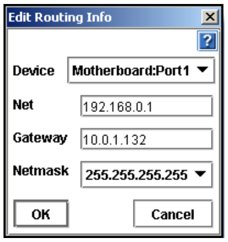

Device: Motherboard:Port1
Net: 192.168.0.1
Gateway: 10.0.1.132
Netmask: 255.255.255.255

Figure 71. Editing Routing Information

3. Change the relevant information.
4. Click OK.
5. Click Apply.

# Deleting Routing Information

You can only delete optional routes you have added.

1. On the Routing tab, select the optional route you want to delete.
2. Click Delete.

   A confirmation message opens.
3. Click OK.
4. Click Apply when you are finished.

# Configuring Storage Module Communication

Use the Communication tab to configure the network interface used by the storage module to communicate with other storage modules on the network and to update the list of managers that the storage module can communicate with.

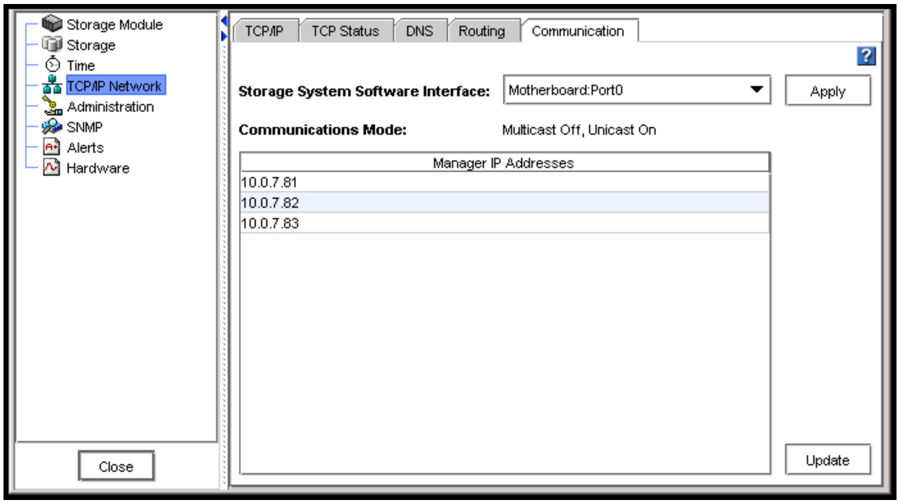

Storage System Software Interface: Motherboard:Port0

Communications Mode: Multicast Off, Unicast On

| Manager IP Addresses |
|---|
| 10.0.7.81 |
| 10.0.7.82 |
| 10.0.7.83 |

### Figure 72. Selecting the Storage System Software Network Interface and Updating the List of Managers

## Selecting the Interface Used by the Storage System Software

The Storage System Software uses one network interface for communication with other storage modules on the network. In order for clustering to work correctly, the Storage System Software communication interface must be designated on each storage module. The interface can be

- a single NIC that is not part of a bond
- a bonded interface consisting of 2 or 4 bonded NICs

***Note:*** *Only NICs that are in the Active or Passive (Ready) state can be designated as the communication interface. You cannot make a disabled NIC the communication interface.*

When you initially set up a storage module using the Configuration Interface, the first interface that you configure becomes the interface used for Storage System Software communication.

***Warning:*** *To change the communication interface, first remove the storage module from the management group.*

To select a different communication interface:

1. Select TCP/IP Network from the configuration categories.

2. Click the Communication Mode tab to bring it to the front, shown in Figure 72.
3. Select an interface from the Storage System Software Interface drop-down list.
4. Click Apply.

# Updating the List of Manager IP Addresses

Update the list of manager IP addresses to ensure that a manager running on this storage module is communicating correctly with all managers in the management group.

## Requirements

- Each time you update the list of managers, you must reconfigure application servers that use the management group to which this storage module belongs. Only update the list mode if you have reason to believe that there is a problem with the communication between the other managers in the group and the manager on this storage module.

1. Select TCP/IP Network from the configuration categories.
2. Click the Communication Mode tab to bring it to the front, shown in Figure 72.
3. Click Update.

   The list is updated with the current storage module in the management group and a list of IPs with every manager’s enabled network interfaces.

   A window opens which displays the IP addresses in the management group and a reminder to reconfigure the application servers that are affected by the update.

***Note:*** *For more information on unicast, see “Guide to Creating Management Groups” on page 171.*

# 5 Setting the Date and Time

The Storage System Module (SSM) uses the date and time settings to create a time stamp when data is stored. You must set the date and time on each SSM.

- **Setting the Time Zone**

  Set the time zone where the SSM is located. This time zone controls the time stamp on volumes and snapshots. You must set the SSM time zone whether you set the time of day manually or use NTP.

- **Using Network Time Protocol (NTP)**

  Configure the SSM to use an external time service.

- **Setting Date and Time**

  Set the date and time on the SSM if not using an external time service.

## Reset Management Group Time

If you change the time on an SSM that is running a manager, you must reset the management group time. If the management group time is different than a manager SSM, you run the risk of inconsistent or unexpected creation time stamps on volumes and snapshots, and also that scheduled snapshots will not start at the intended time. See "Resetting the Management Group Time" on page 179.

## Getting There

1. In the navigation window, select the SSM and log in, if necessary.
2. Click Storage Module Tasks and select Edit Configuration.

   The Time window appears

3. Select Time from the SSM configuration categories. The Time window opens, shown in Figure 73.

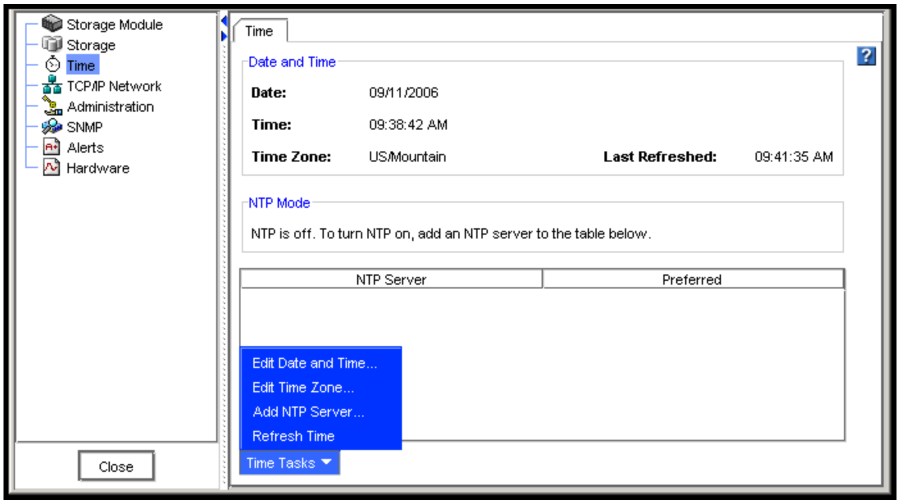

Date and Time

**Date:** 09/11/2006

**Time:** 09:38:42 AM

**Time Zone:** US/Mountain **Last Refreshed:** 09:41:35 AM

NTP Mode

NTP is off. To turn NTP on, add an NTP server to the table below.

| NTP Server | Preferred |
|---|---|

Figure 73. Setting the Time Zone

# Setting the SSM Time Zone

You must set the time zone whether or not you use NTP. Set the time zone for the physical location of the SSM. HTTP files display the time stamp according to this local time zone.

1. In the Edit Configuration window, click Time Tasks and select Edit Time Zone.

   The Time Zone Configuration window appears.

2. From the drop-down list, select the time zone in which this storage module resides.

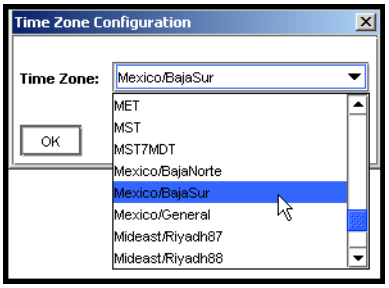

**Time Zone:** Mexico/BajaSur

3. Click OK.

   Note the Time Zone field in the Date and Time group box.

# Setting SSM Date and Time

If using NTP, the NTP server controls the date and time for the SSM. See “Using NTP” on page 124.

**Note:** *Even if you are using an NTP server, you can set the date and time manually. If the difference between the date and time on the SSM and the date and time on the NTP server is too large, the NTP server will not change the date and time on the SSM. To ensure that the NTP server is able to control the SSM date and time, first set the date and time manually.*

## Setting the Date and Time

1. If you are not using an NTP server, and want to set date and time manually, remove all NTP servers from the NTP list.

   See “Deleting an NTP Time Server” on page 126. This sets NTP mode to OFF.

2. In the Edit Configuration window, click Time Tasks and select Edit Date and Time.

   The Date and Time Configuration window appears.

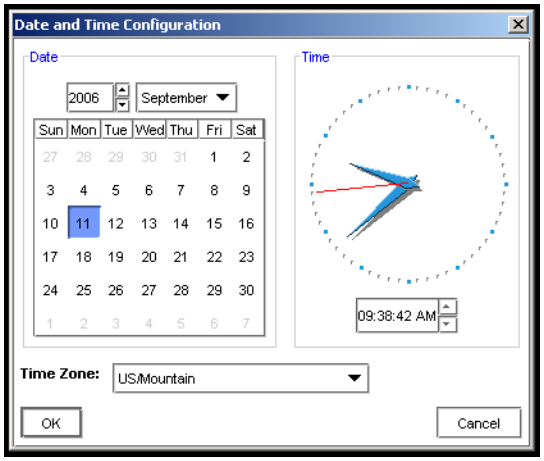

Date and Time Configuration

Date

2006 September

| Sun | Mon | Tue | Wed | Thu | Fri | Sat |
|---|---|---|---|---|---|---|
| 27 | 28 | 29 | 30 | 31 | 1 | 2 |
| 3 | 4 | 5 | 6 | 7 | 8 | 9 |
| 10 | 11 | 12 | 13 | 14 | 15 | 16 |
| 17 | 18 | 19 | 20 | 21 | 22 | 23 |
| 24 | 25 | 26 | 27 | 28 | 29 | 30 |
| 1 | 2 | 3 | 4 | 5 | 6 | 7 |

Time

09:38:42 AM

Time Zone: US/Mountain

### Figure 74. Setting the Time Zone, Date and Time

3. Change the date and time to the correct date and time for that time zone.

   — In the Date group box, set the year, month, and day.

   — In the Time group box, highlight a portion of the time and increase or decrease it with the arrows. You may also type in the time directly.

   — Select a time zone for the Time Zone drop-down list.

4. Click OK.

# Using NTP

Network time protocol servers (NTP) can manage the time for the SSM instead of using the local system time. NTP updates occur at 5 minute intervals. You still must set the time zone for the SSM. See "Setting the SSM Time Zone" on page 122.

**Note:** *When using a Microsoft Windows* server as an external time source for a storage module, you must configure W32Time (the Windows Time service) to also use an external time source. The storage module does not recognize the windows server as an NTP server if W32Time is configured to use an internal hardware clock.*

1. In the Edit Configuration window, click Time Tasks and select Add NTP Server.

   The Add NTP Server window appears.

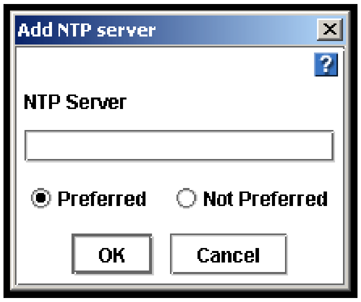

**Figure 75. Adding an NTP Server**

2. Type the IP address of the NTP server you want to use.
3. Click whether you want the NTP server to be designated preferred or not preferred.

**Note:** *A* ***preferred*** *NTP server is one that is more reliable, such as a server that is on a local network. An NTP server on a local network would have a reliable and fast connection to the SSM.*
***Not preferred*** *designates an NTP server to be used as a back up if a preferred NTP server is not available. An NTP server that is not preferred might be located further away and have a less reliable connection.*

4. Click OK. The NTP server is added to the list on the NTP tab, shown in Figure 76.

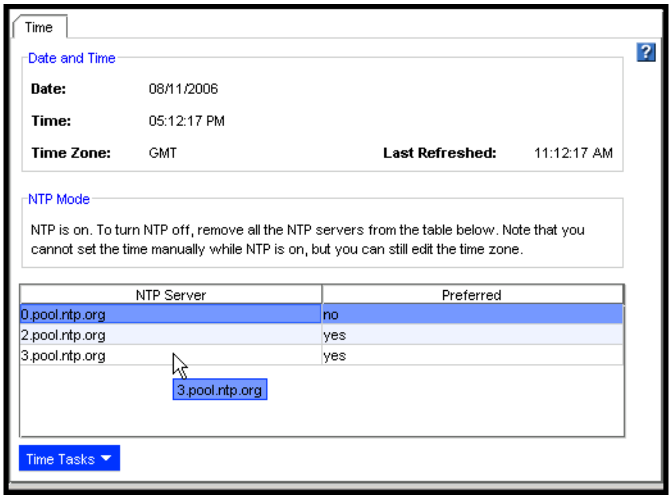

Date: 08/11/2006

Time: 05:12:17 PM

Time Zone: GMT

Last Refreshed: 11:12:17 AM

NTP is on. To turn NTP off, remove all the NTP servers from the table below. Note that you cannot set the time manually while NTP is on, but you can still edit the time zone.

| NTP Server | Preferred |
|---|---|
| 0.pool.ntp.org | no |
| 2.pool.ntp.org | yes |
| 3.pool.ntp.org | yes |

### Figure 76. Viewing the List of NTP Servers

The NTP servers are accessed in the order you added them. The first server you add, if it is marked preferred, has the highest order of precedence. The second server you add takes over as a time server if the first added server fails.

## Editing NTP Servers

You can change the preference properties of NTP servers. To change the IP address of an NTP server, you must remove the one no longer in use and add a new NTP server.

1. Observe the NTP servers in the list of the Time tab of the Edit Configuration window.

   Having servers in the list sets the NTP Mode to On.

2. In the Edit Configuration window, select an NTP server in the list and right-click.

3. Click Edit NTP Server, shown in Figure 77.

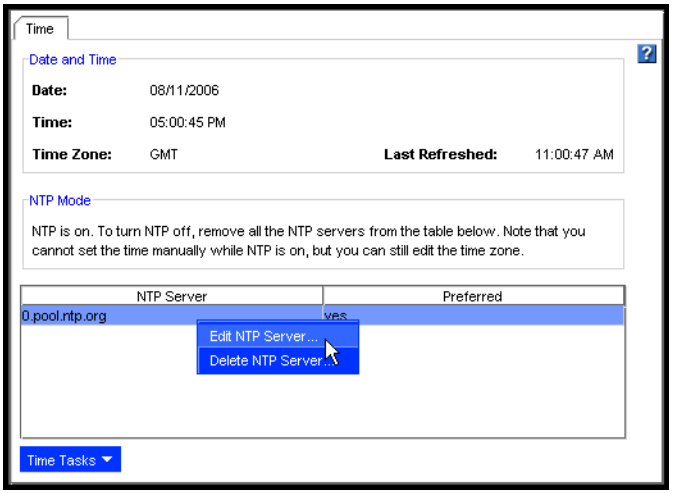

Time

Date and Time

Date: 08/11/2006

Time: 05:00:45 PM

Time Zone: GMT

Last Refreshed: 11:00:47 AM

NTP Mode

NTP is on. To turn NTP off, remove all the NTP servers from the table below. Note that you cannot set the time manually while NTP is on, but you can still edit the time zone.

| NTP Server | Preferred |
|---|---|
| 0.pool.ntp.org | yes |

Figure 77. Selecting an NTP Server to Edit

4. The Edit NTP Server window appears.

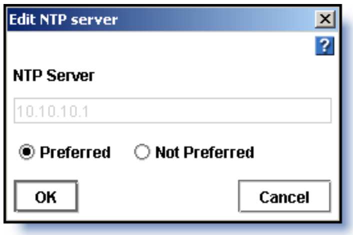

Edit NTP server

NTP Server

10.10.10.1

● Preferred ○ Not Preferred

Figure 78. Editing an NTP Server

5. Change the preference of the NTP server.
6. Click OK. The list of NTP servers displays the changed NTP server in the list.

***Note:*** *To change the IP address of an NTP server, you must remove the server no longer in use and add a new NTP server.*

# Deleting an NTP Time Server

You may need to delete an NTP time server for these reasons:

- If the IP address of that server becomes invalid
- If you no longer want to use that server
- If you want to change the order of servers in the list.

1. Observe the NTP servers in the list of the Time tab of the Edit Configuration window.

   Having servers in the list sets the NTP Mode to On.

2. In the Edit Configuration window, select an NTP server in the list and right-click.
3. Select Delete NTP Server.

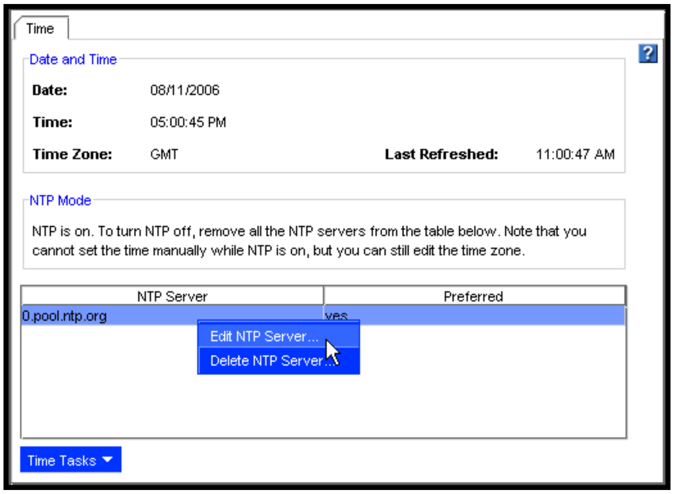

**Date:** 08/11/2006

**Time:** 05:00:45 PM

**Time Zone:** GMT **Last Refreshed:** 11:00:47 AM

NTP is on. To turn NTP off, remove all the NTP servers from the table below. Note that you cannot set the time manually while NTP is on, but you can still edit the time zone.

| NTP Server | Preferred |
|---|---|
| 0.pool.ntp.org | yes |

**Figure 79. Selecting an NTP Server to Delete**

4. A confirmation window appears.
5. Click OK.

   The list of NTP servers displays a freshened list of available servers.

## Changing the Order of NTP Servers

To change the order of access for time servers, delete the one whose order you want to change. Add that same server back into the list. It is placed at the bottom of the list, and is the last to be accessed.

The screen displays the NTP servers in the order you added them, from first added to last.

The server you added first is the one accessed first when time needs to be established. If this NTP server is not available for some reason, the next NTP server that was added, and is preferred, is used for time serving.

Click the headings in the list to change the visual order of the NTP servers on the screen. This does not change the order of precedence when serving time.

# 6 Administrative Users and Groups

The Storage System Software comes configured with two default administrative groups and one default administrative user. You can add, edit, and delete administrative users and groups. All administrative users and groups must be added and managed locally.

***Note:*** *The user who is created during SSM configuration using the Configuration Interface becomes a member of the Full Administrator group by default.*

## Getting There

Open the Edit Configuration window to access the Administration configuration category.

1. Select an SSM in the navigation window.
2. Click Storage Module Tasks in the tab window.
3. Select Edit Configuration.
4. Select Administration from the configuration categories.

   The Groups tab opens.

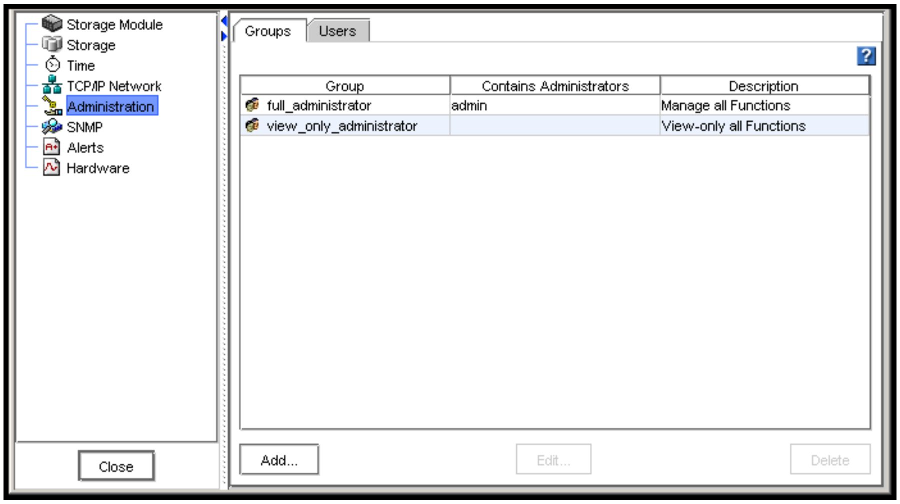

| Group | Contains Administrators | Description |
|---|---|---|
| full_administrator | admin | Manage all Functions |
| view_only_administrator | | View-only all Functions |

Figure 80. Viewing the Administration Groups Tab

# Managing Administrative Groups

The SSM comes configured with a set of default administrative groups. You can use these groups or create new ones.

## Default Administrative Groups

If you assign an administrative user to one of the following groups, that user will have the privileges associated with the group.

### Table 29. Using Default Administrative Groups

| Name of Group | Management Capabilities Assigned to Group |
|---|---|
| Full_Administrator | Manage all functions (read, write access to all functions) |
| View_Only_Administrator | View-only capability to all functions (read only) |

## Adding Administrative Groups

Administrative groups are listed on the SSM Administration window on the Groups tab, shown in Figure 81.

### Adding a Group

1. Select SSM Administration from the configuration categories.

2. Click Add on the Groups tab. The Create Administrative Group window opens, shown in Figure 81.

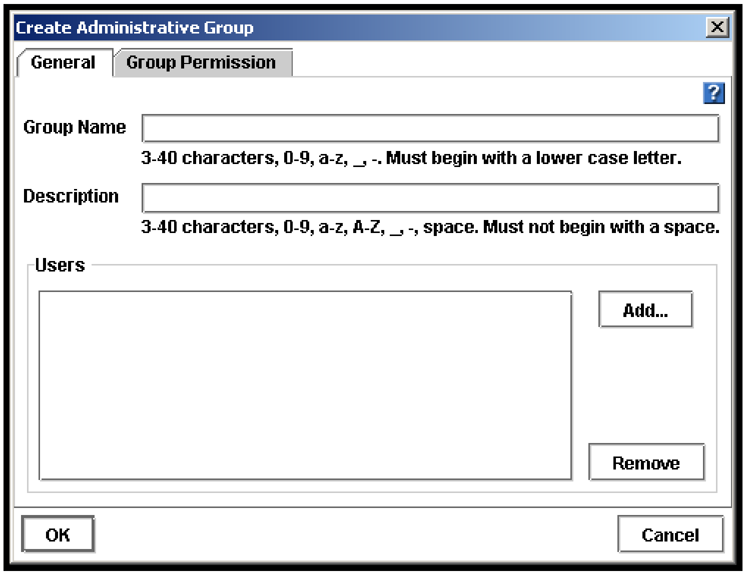

3-40 characters, 0-9, a-z, _, -. Must begin with a lower case letter.

3-40 characters, 0-9, a-z, A-Z, _, -, space. Must not begin with a space.

Figure 81. Adding an Administrative Group

3. Type a Group Name and Description. Both are required.

## Table 30. Administrative Group Name Requirements

| Group Name Requirements | Example |
|---|---|
| • 3 to 40 characters<br>• start with a letter<br>• Use letters a-z, A-Z, numbers 0-9, or characters _, - | • Software_Admins<br>• Region11_Managers |

## Adding a User to the Group

1. Click Add in the Users section. The Add Users window opens with a list of administrative users, shown in Figure 82.

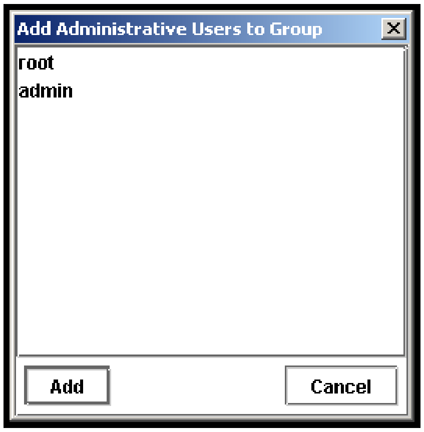

Figure 82. Adding an Administrative User to a Group

2. Select one or more users you want to add to the group.
3. Click Add.

# Adding Administrative Group Permissions

Administrative groups can have

- Different levels of access to the SSM, such as read/write
- Access to different management capabilities for the SSM, such as creating volumes

When you are creating a group, you also set the management capabilities available to members of a group. The default setting for a new group is Read Only for each category.

1. From the Create Administrative Group window, click the Group Permission tab to bring it to the front, shown in Figure 83.

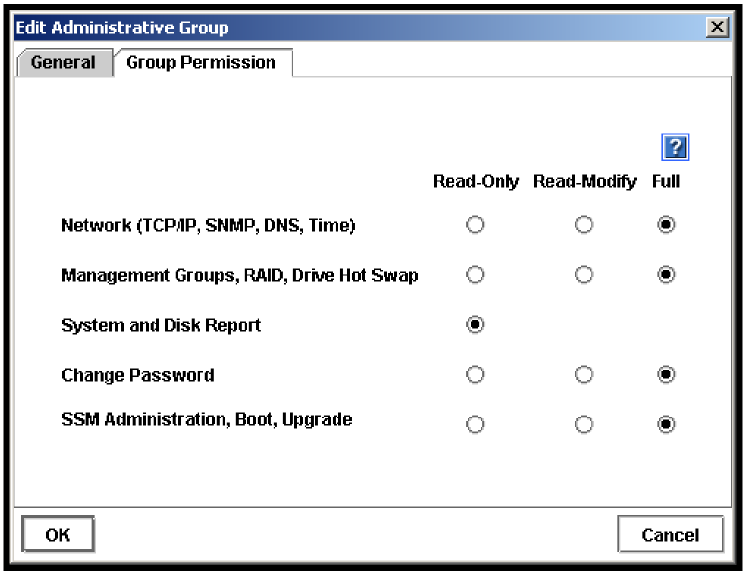

| | Read-Only | Read-Modify | Full |
|---|---|---|---|
| Network (TCP/IP, SNMP, DNS, Time) | ○ | ○ | ● |
| Management Groups, RAID, Drive Hot Swap | ○ | ○ | ● |
| System and Disk Report | ● | | |
| Change Password | ○ | ○ | ● |
| SSM Administration, Boot, Upgrade | ○ | ○ | ● |

**Figure 83. Adding Permissions to Administrative Groups**

2. Click the permission level for each function for the group you are creating.
3. Click the General tab and complete the rest of the information if you have not already done so.
4. Click OK to finish adding the group. The SSM Administration window opens with the Groups tab in front. The new group is added to the list.

# Description of Administrative Group Permissions

### Table 31. Descriptions of Group Permissions

| Management Area | Activities Controlled by This Area |
|---|---|
| Network | Choose type of network connection, set the time and time zone for the SSMs, identify the Domain Name Server, and use SNMP. |
| Management Groups, RAID, Drive Hot Swap | Set the RAID configuration for the SSM. Shut down disks, restart RAID, and hot swap disks. Create management groups. |
| System and Disk Report | View reports about the status of the SSM. |
| Change Password | Change administrative users' passwords. |
| SSM Administration and Upgrade | Add administrators and upgrade the Storage System Software. |

**What the Permission Levels Mean:**

- **Read Only:** User can only view the information about these functions.
- **Read-Modify:** User can view and modify existing settings for these functions.
- **Full**: Users can perform all actions (view, modify, add new, delete) in all functions.

## Sorting Columns in the Administrative Group Window

The columns in the Administrative Group window can be sorted in ascending or descending order.

- Click on the column header to sort.
- Click again to reverse the sort.

  The arrow next to the column title indicates which column is the sorted column, and whether the sorting order is ascending (up arrow) or descending (down arrow).

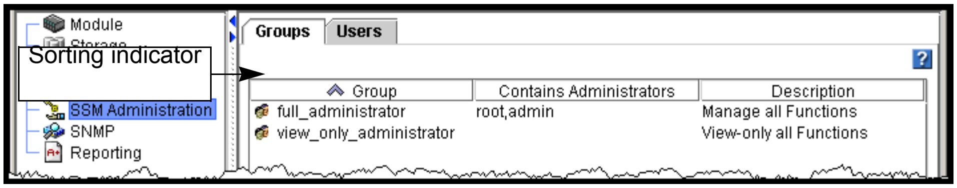

| ⌃ Group | Contains Administrators | Description |
|---|---|---|
| full_administrator | root,admin | Manage all Functions |
| view_only_administrator | | View-only all Functions |

**Figure 84. Sorting Administrative Groups**

# Editing Administrative Groups

Change information about administrative groups. Administrative groups are listed on the SSM Administration window on the Groups tab.

1. Select SSM Administration from the configuration categories.
2. Select the group you want to edit.

3. Click Edit. The Edit Administrative Group window opens, shown in Figure 85.

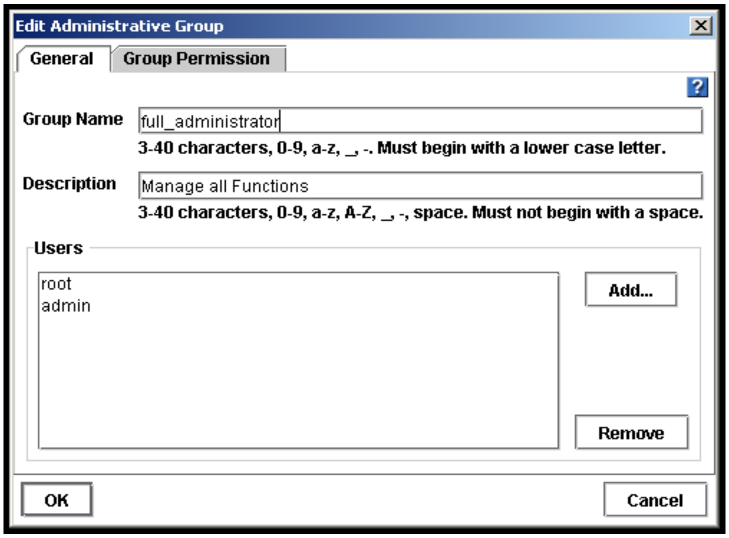

Group Name: full_administrator
3-40 characters, 0-9, a-z, _, -. Must begin with a lower case letter.

Description: Manage all Functions
3-40 characters, 0-9, a-z, A-Z, _, -, space. Must not begin with a space.

Users:
root
admin

Figure 85. Editing an Administrative Group

4. Change the name and description as necessary.

## Adding or Removing Administrative Users in an Existing Group

### Adding New Users to the Group

1. Click Add in the Users section. The Add Users window opens with a list of administrative users.
2. Select one or more users to add to the group.
3. Click Add. The users are added to the list.
4. Click OK when you are finished adding users.

### Removing Users from a Group

1. Select the user to remove in the Users section.
2. Click Remove. The user is removed from the list.

## Changing Administrative Group Permissions

Change the management capabilities available to members of a group. The default setting is Read Only for each category.

1. Click the Groups tab to bring it to the front.
2. Select a group and click Edit. The Edit Administrative Group window opens.
3. Click the Group Permission tab to bring it to the front.
4. Click the management capabilities you want for the group you are editing.

5. Click OK when you are finished.

## Deleting Administrative Groups

Delete all users from a group before you delete the group.

1. Select SSM Administration from the configuration categories.
2. Click the Groups tab to bring it to the front.
3. Select the group to delete.
4. Click Delete. A confirmation message opens.
5. Click OK.

***Note:*** *When you delete a group, the users who are members of that group are NOT deleted.*

# Managing Administrative Users

Add administrative users as necessary to provide access to the management functions of Storage System Software.

***Note:*** *The user who is created during SSM configuration using the Configuration Interface becomes a member of the Full Administrator group by default.*

# Adding Administrative Users

Administrative users are listed on the SSM Administration window on the Users tab along with their group membership and a description.

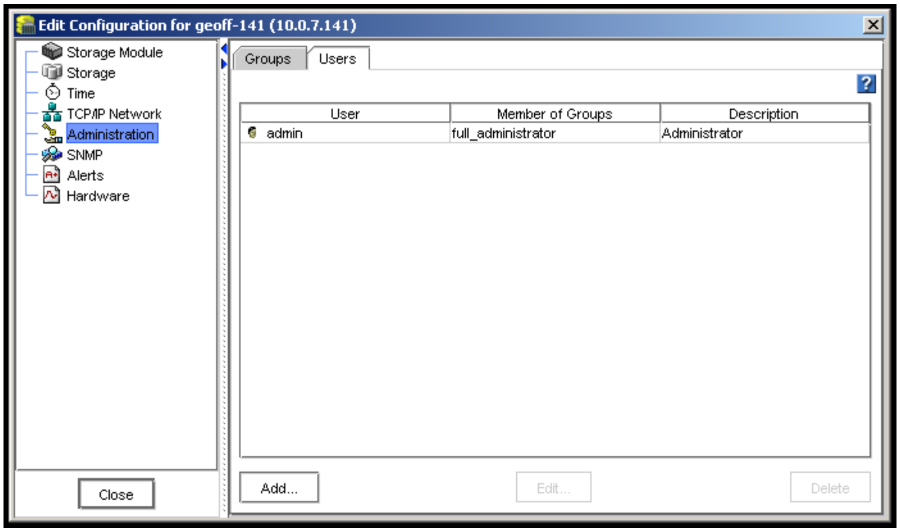

| User | Member of Groups | Description |
|---|---|---|
| admin | full_administrator | Administrator |

**Figure 86. Adding Administrative Users**

## Adding an Administrative User

1. Select SSM Administration from the configuration categories.
2. Click the Users tab to bring it to the front, shown in Figure 86.

3. Click Add. The Create Administrative User window opens, shown in Figure 87.

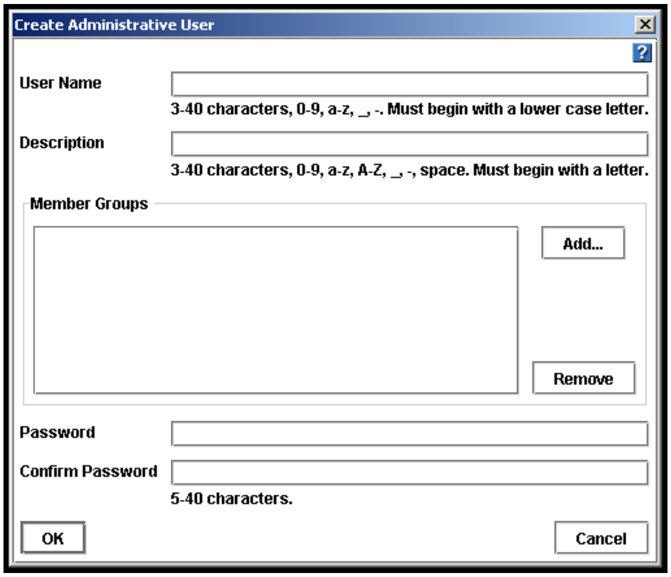

User Name

3-40 characters, 0-9, a-z, _, -. Must begin with a lower case letter.

Description

3-40 characters, 0-9, a-z, A-Z, _, -, space. Must begin with a letter.

Member Groups

Password

Confirm Password

5-40 characters.

**Figure 87. Adding an Administrative User**

4. Type a User Name and Description.
5. Type a password and confirm that password.

## Adding a Member Group

1. Click Add in the Member Groups section. The Add Administration Groups window opens, shown in Figure 88.

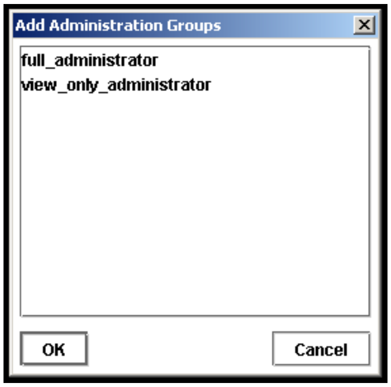

full_administrator

view_only_administrator

**Figure 88. Adding a Group to an Administrative User**

2. Select one or more groups you want to add.
3. Click OK. The Create Administrative User window opens.
4. Click OK to finish adding the administrative user.

## Sorting Columns in the Administrative Users Window

The columns in the Administrative Users window can be sorted in ascending or descending order.

- Click on the column header to sort.
- Click again to reverse the sort.
- The arrow next to the column title indicates which column is the sorted column, and whether the sorting order is ascending (up arrow) or descending (down arrow).

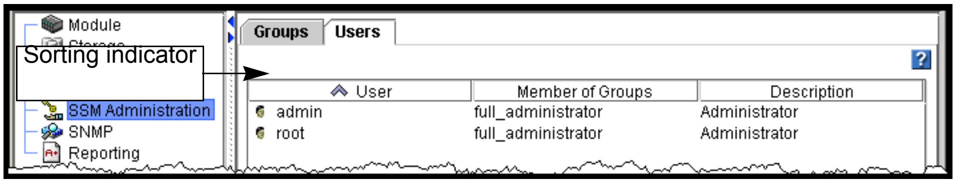

| User | Member of Groups | Description |
|---|---|---|
| admin | full_administrator | Administrator |
| root | full_administrator | Administrator |

Figure 89. Sorting Administrative Users

# Editing Administrative Users

1. Select SSM Administration from the configuration categories.
2. Click the Users tab to bring it to the front.
3. Select the user to edit from the list of users.

4. Click Edit. The Edit Administrative User window opens, shown in Figure 90.

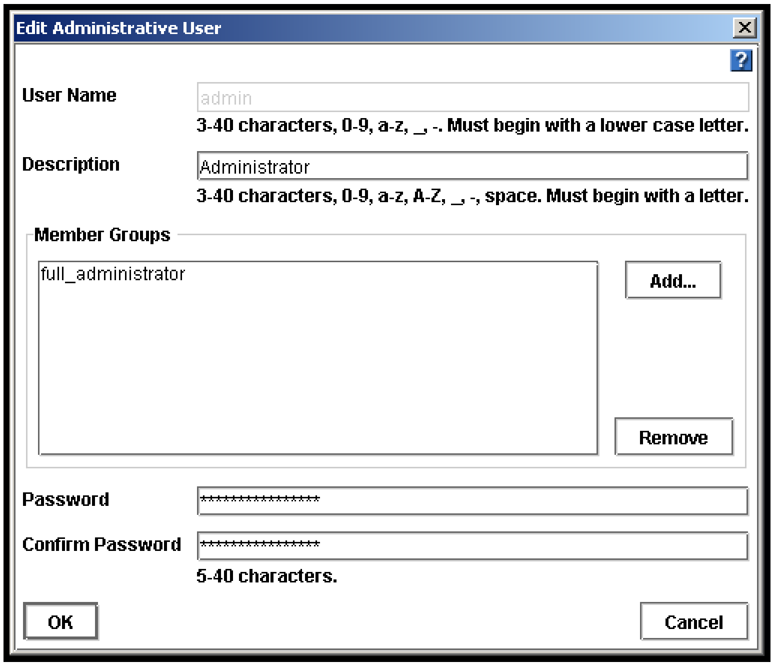

3-40 characters, 0-9, a-z, _, -. Must begin with a lower case letter.

3-40 characters, 0-9, a-z, A-Z, _, -, space. Must begin with a letter.

5-40 characters.

**Figure 90. Editing an Administrative User**

5. Change the necessary information.
6. Click OK.

## Deleting Administrative Users

1. Select SSM Administration from the configuration categories.
2. Click the Users tab to bring it to the front.
3. Select the user to delete from the list of users.
4. Click Delete. A confirmation message opens.
5. Click OK

***Note:*** *If you delete an administrative user, that user is automatically removed from any administrative groups.*

# 7 Using SNMP

The  SSM  can be monitored using an SNMP Agent. You can also enable SNMP traps.

The SSM Management Information Base (MIB) is read-only and supports SNMP versions 1 and 2c. See “Installing the Storage System MIB” on page 145 for a list of Storage System MIBs.

## Getting There

Open the Edit Configuration window to access the SNMP configuration category.

1. Select a storage module in the navigation window.
2. Click Storage Module Tasks in the tab window.
3. Select Edit Configuration.
4. Select the SNMP configuration category.
5. The SNMP General tab opens.

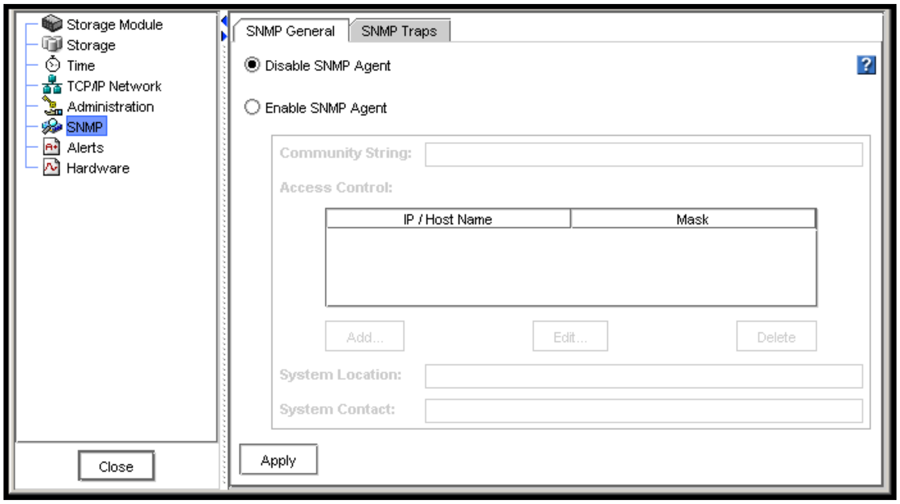

● Disable SNMP Agent

○ Enable SNMP Agent

Figure 91. Opening the SNMP Configuration Category

# Using SNMP

The storage modules allow enabling and disabling of SNMP agents.

Adding an SNMP agent includes these tasks:

- Enabling the SNMP Agent
- Adding a community string.
- The community string acts as an authentication password. It identifies hosts that are allowed read-only access to the SNMP data. The community "public" typically denotes a read-only community. This string is entered into an SNMP management tool (not included with SSMs) when attempting to access the system.
- Adding access control for SNMP clients.

  You can either enter a specific IP address and the IP Netmask as None to allow a specific host to access SNMP, or you can specify the Network Address with its netmask value so that all hosts matching that IP and netmask combination can access SNMP.

Additional agents and traps can be added and modified.

# Enabling the SNMP Agent

1. Click Enable SNMP Agent. The Enable.

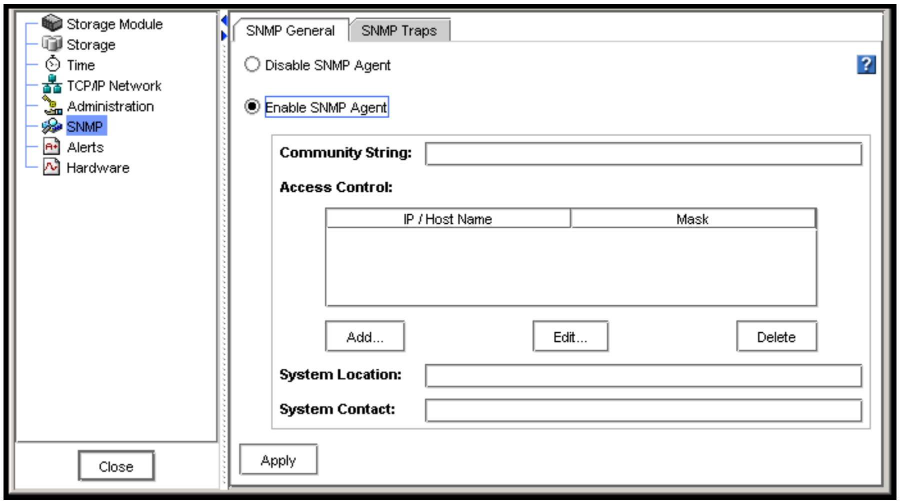

**Figure 92. Enabling the SNMP Agent**

2. Type the Community String.

3. Under Access Control, click Add to add an SNMP client.

   The Add SNMP Client window opens. You can add SNMP clients by specifying either IP addresses or host names.

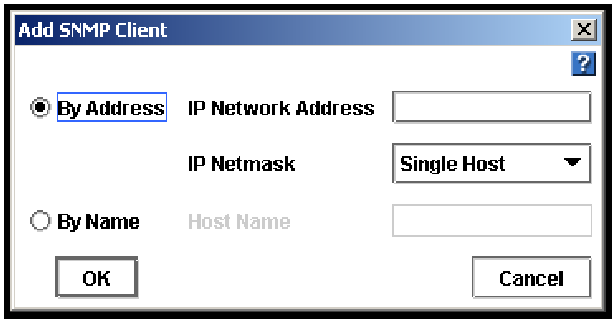

Figure 93. Adding an SNMP Client

## By IP Address

1. Click By Address and type the IP address.
2. Type an IP Netmask from the list. Select Single Host if adding only one SNMP client.
3. Click OK.

   The IP address and netmask appear in the Access Control list.

## By Host Name

1. Click By Name and type a host name.

   That host name must exist in DNS and the SSM must be configured with DNS for the client to be recognized by the host name.
2. Click OK. The host name appears in the Access Control list.

# Entering System Information (Optional)

1. Enter any desired System Location information for the storage module.

   For example, this information may include the address, building name, room number, and so on.
2. Enter System Contact information.

   Normally, this will be network administrator information, such as email address or phone number for the person you would contact if you could not connect to SNMP clients.

# Editing Access Control Entries

You can change the information for the hosts granted access.

1. Select a host listed in the Access Control list, shown in Figure 94.

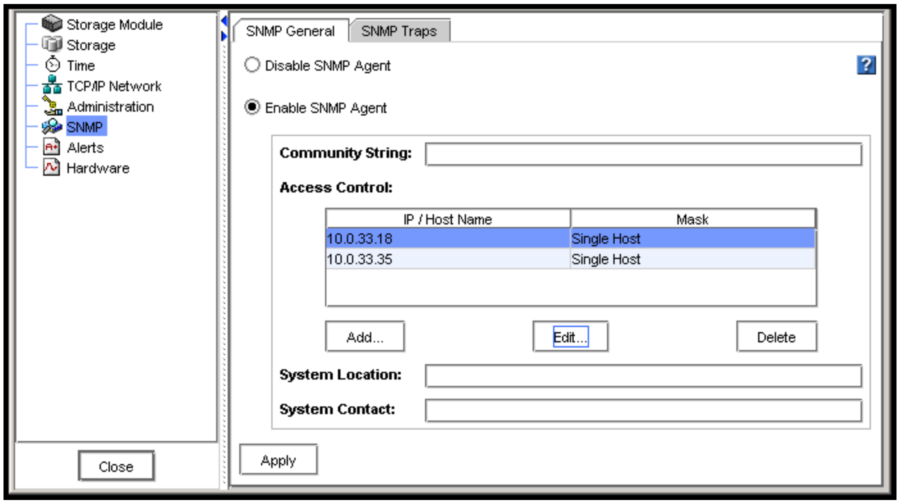

○ Disable SNMP Agent

● Enable SNMP Agent

Community String:

Access Control:

| IP / Host Name | Mask |
|---|---|
| 10.0.33.18 | Single Host |
| 10.0.33.35 | Single Host |

System Location:

System Contact:

Figure 94. Editing a Host in the Access Control List

2. Click Edit. The Edit SNMP Client window opens, shown in Figure 95.

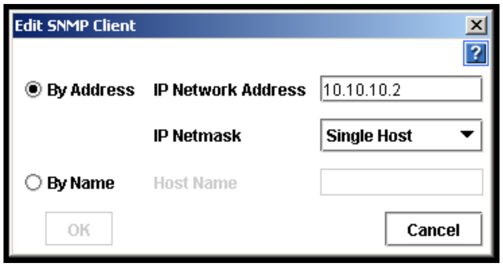

● By Address IP Network Address 10.10.10.2

IP Netmask Single Host

○ By Name Host Name

Figure 95. Editing SNMP Client from the Access Control List

3. Change the appropriate information.
4. Click OK.

# Deleting Access Control Entries

Delete an SNMP client from the list.

1. Select a host listed in the Access Control list, shown in Figure 94.
2. Click Delete. A confirmation message opens.
3. Click OK.

# Using the SNMP MIB

The Storage System MIB provides read-only access to the SSM. The SNMP implementation in the SSM supports MIB-II compliant objects.

In addition, MIB files have been developed for SSM-specific information. These files, when loaded in the Network Management System, allow you to see SSM specific information such as model number, serial number, hard disk capacity, network parameters, RAID configuration, DNS server configuration details, and more. See “Installing the Storage System MIB” on page 145.

# Installing the Storage System MIB

The Storage System MIB files are installed when you install the Storage System Console. Load the Storage System MIB in the Network Management System as outlined below.

1. Load STORAGE – SYSTEMS – GLOBAL – REG
2. Load STORAGE–SYSTEMS–SSM–COMMON – MIB
3. The following MIB files can be loaded in any sequence:
   - STORAGE–SYSTEMS–SSM–COMMON–DNS–MIB
   - STORAGE–SYSTEMS–SSM–COMMON-CLUSTERING-MIB
   - STORAGE–SYSTEMS–SSM–COMMON–INFO–MIB
   - STORAGE–SYSTEMS–SSM–COMMON– NETWORK–MIB
   - STORAGE–SYSTEMS–SSM–COMMON–NIS–MIB
   - STORAGE-SYSTEMS-SSM-COMMON-NOTIFICATION-MIB
   - STORAGE–SYSTEMS–SSM–COMMON–NTP–MIB
   - STORAGE–SYSTEMS–SSM–COMMON–STATUS–MIB
   - STORAGE–SYSTEMS–SSM–COMMON–STORAGE–MIB

***Note:*** *Any variable that is labeled “Counter64” in the MIB requires version 2c or later of the protocol.*

***Note:*** *Other standard version 2c compliant MIB files are also provided on the resource CD. Load these MIB files in the Network Management System if required.*

# Disabling the SNMP Agent

Disable the SNMP Agent if you no longer want to use SNMP applications to monitor your network of SSMs.

Open the Edit Configuration window to access the SNMP configuration category.

1. Select a storage module in the navigation window.
2. Click Storage Module Tasks in the tab window.
3. Select Edit Configuration
4. Select the SNMP configuration category. The SNMP General window opens.

## Disabling SNMP

1. On the SNMP General window, select Disable SNMP Agent.
2. Click Apply.

# Enabling SNMP Traps

You must have first enabled the SNMP agent in order to use SNMP traps.

Enable SNMP Traps to have an SNMP tool send alerts when a monitoring threshold is reached.

Open the Edit Configuration window to access the SNMP configuration category.

1. Select a storage module in the navigation window.
2. Click Storage Module Tasks in the tab window.
3. Select Edit Configuration.
4. Select the SNMP configuration category. The SNMP General window opens.

5. Select the SNMP Traps tab. The SNMP Traps window opens, shown in Figure 96.

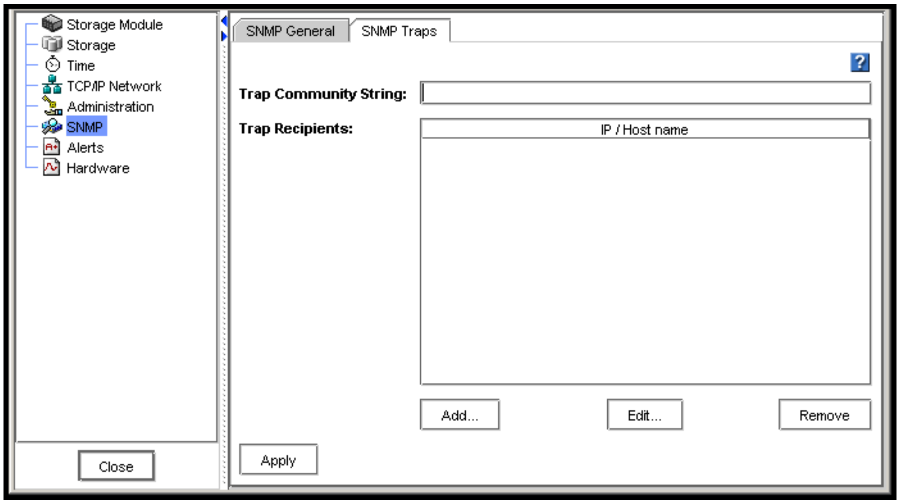

Figure 96. Enabling SNMP Traps

## Enable SNMP Traps

1. Enter the Trap Community String. This is required if you want to use SNMP traps.

***Note:*** *The Trap Community String is used for client-side authentication.*

2. Click Add in the Trap Recipients area to add specific trap recipients. The Trap Recipient window opens.

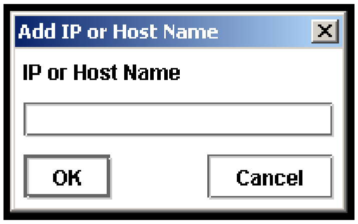

Figure 97. Adding an SNMP Trap Recipient

3. Enter the host name or IP address for the SNMP client that is receiving the traps.
4. Click OK.
5. Repeat steps 2 through 4 for each host in the trap community.
6. Click Apply on the SNMP Traps tab when you are finished adding hosts.

## Editing Trap Recipients

1. Select the host you want to change from the list of Trap Recipients and click Edit. The Trap Recipient window opens.
2. Change the host name or IP address.
3. Click OK.
4. Click Apply when you are finished editing trap recipients.

## Removing Trap Recipients

1. Select the host you want to remove from the list of Trap Recipient and click Remove.

   A confirmation window opens.
2. Click OK to remove the trap recipient. The host is removed from the list.
3. Click Apply on the SNMP Traps tab when you are finished removing trap recipients.

# Disabling SNMP Traps

To disable SNMP traps, you must delete all of the settings in the SNMP Traps window.

1. Remove the Trap Recipient hosts.
2. Delete the Trap Community String.
3. Click Apply.

# 8 Reporting

## Reporting Overview

Reporting capabilities are divided into two categories, the Alerts category and the Hardware category.

The Alerts category allows you to set up email alerting and to review alerts generated automatically by the operating system.

The Hardware category provides a report of system statistics, hardware, and configuration information; allows you to run diagnostic tests on the hardware; and save log files.

## Alerts Overview

The Alerts category includes multiple types of information and reporting capabilities. Review configuration information, save log files, set up email alerting, and review alerts generated automatically by the operating system.

Use alerts to:

- View real-time statistical information about the SSM
- View and save log files
- Set up active monitoring of selected variables
- Set up email notification
- View alerts.

## Using Active Monitoring

Use active monitoring to track the health of the SSM. Active monitoring allows you to set up notification through emails, alerts in the Console, and SNMP traps. You can choose which variables to monitor and choose the notification methods for alerts related to the monitored variables.

Critical variables, such as the CPU temperature and motherboard temperature, have thresholds that trigger a shutdown of the SSM.

Open the Edit Configuration window to access the Alerts configuration category.

1. Select a storage module in the navigation window.

2. Click Storage Module Tasks in the tab window.
3. Select Edit Configuration.
4. Select the Alerts configuration category.
5. The Alerts Setup tab opens.

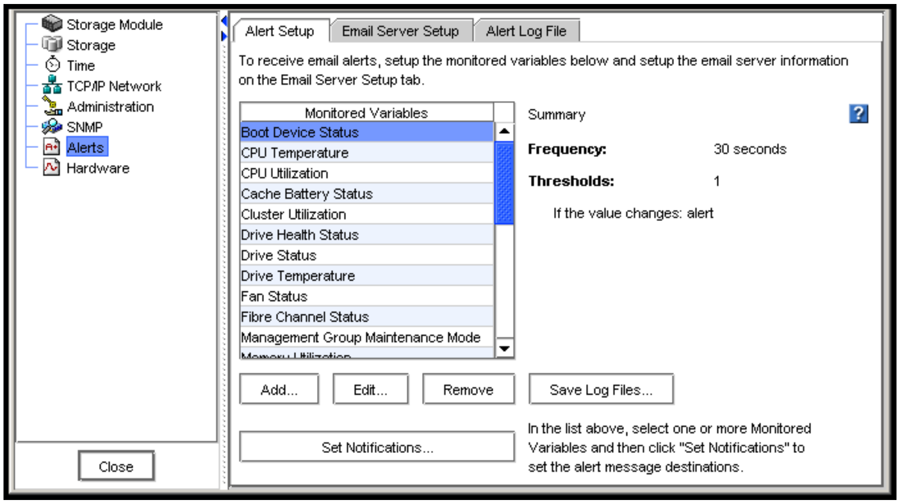

To receive email alerts, setup the monitored variables below and setup the email server information on the Email Server Setup tab.

Monitored Variables

- Boot Device Status
- CPU Temperature
- CPU Utilization
- Cache Battery Status
- Cluster Utilization
- Drive Health Status
- Drive Status
- Drive Temperature
- Fan Status
- Fibre Channel Status
- Management Group Maintenance Mode
- Memory Utilization

Summary

**Frequency:** 30 seconds

**Thresholds:** 1

If the value changes: alert

In the list above, select one or more Monitored Variables and then click "Set Notifications" to set the alert message destinations.

Figure 98. Setting Active Monitoring Variables

# Setting Notification Methods for Monitored Variables

Use Edit to configure notification methods and change the frequency that the variable is monitored, if allowed.

## Configuring Notification Methods for All Variables

You can configure alert actions for one variable and then apply those settings to all the variables in the list. Use the Apply Threshold Actions button to globally apply the settings. Then you can customize alerting actions for a particular variable by editing that variable.

1. Select from the list the variable you want to edit.

2. Click Edit.

   The Configure Variable wizard opens to Step 1.

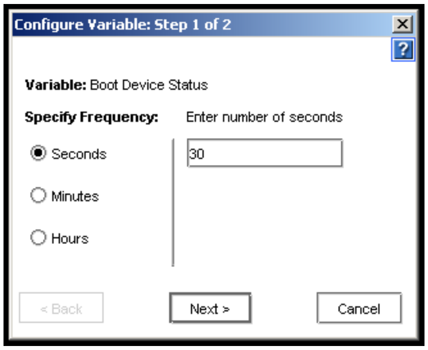

Configure Variable: Step 1 of 2

**Variable:** Boot Device Status

**Specify Frequency:** Enter number of seconds

● Seconds

○ Minutes

○ Hours

30

Figure 99. Editing a Variable, Step 1

***Note:*** *For some variables, only the notification method can be changed. For example, the frequency for the motherboard temperature variable is set to 15 seconds and cannot be changed.*

3. [Optional] If allowed, change the frequency for the variable and click Next.

   The Configure Variable wizard opens to Step 2.

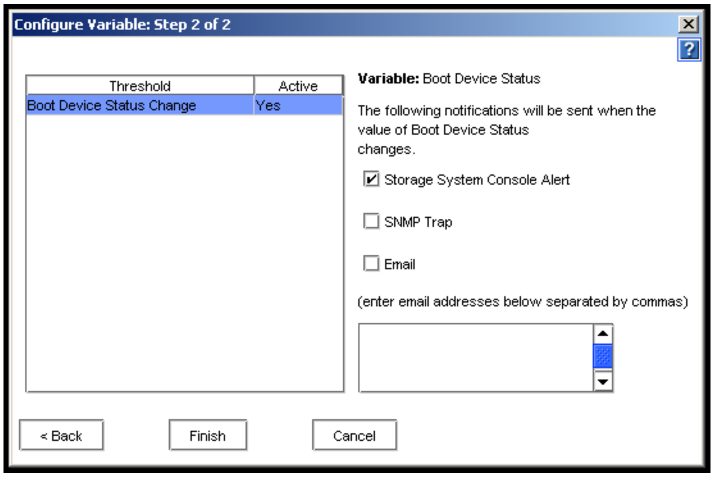

Configure Variable: Step 2 of 2

| Threshold | Active |
|---|---|
| Boot Device Status Change | Yes |

**Variable:** Boot Device Status

The following notifications will be sent when the value of Boot Device Status changes.

☑ Storage System Console Alert

☐ SNMP Trap

☐ Email

(enter email addresses below separated by commas)

Figure 100. Editing a Variable, Step 2

4. [Optional] Change the alert notification method.

5. [Optional] To apply the alert actions (including the email addresses) that you selected in step 4 to all variables that are monitored on the SSM, select the Apply Threshold Actions to All Monitored Variables checkbox.

6. Click Finish.

**Note:** *If you are requesting email notification, be sure to set up the SMTP settings on the Email tab.*

## Removing a Variable from Active Monitoring

Use Remove to remove variables to stop active monitoring. You can return a variable to active monitoring at any time. Permanent variables, such as motherboard temperature, cannot be removed.

1. Select the variable you want to remove.
2. Click Remove.

   A confirmation message opens.
3. Click OK.

   The variable is removed.

**Note:** *Variables are not deleted when they are removed from active monitoring. You can add them back to active monitoring at any time.*

## Adding Variables to Monitor

You can only add variables that have been previously removed. The variables that the SSM is currently monitoring are listed in the box. All variables in the list are configured and set for Console alerts.

1. Click Add.

   The Configure Variable wizard opens to Step 1.

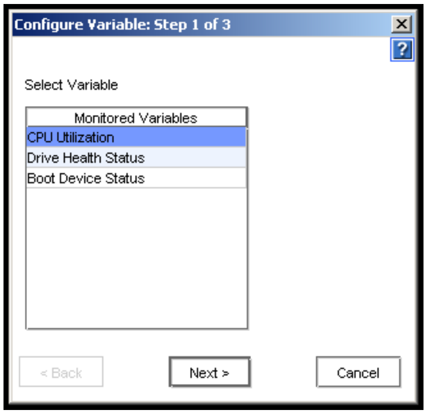

Configure Variable: Step 1 of 3

Select Variable

| Monitored Variables |
|---|
| CPU Utilization |
| Drive Health Status |
| Boot Device Status |

Figure 101. Adding a Variable, Step 1

2. Select the variable that you want to monitor and click Next.

   The Configure Variable wizard, Step 2, opens.

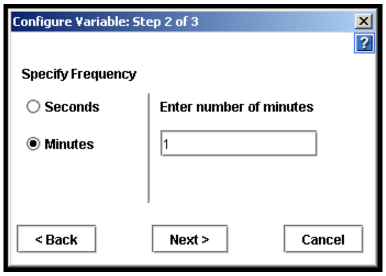

Configure Variable: Step 2 of 3

Specify Frequency

○ Seconds

● Minutes

Enter number of minutes: 1

Figure 102. Adding a Variable, Step 2

3. Specify the frequency for monitoring the variable and click Next.

   The Configure Variable wizard, Step 3, opens.

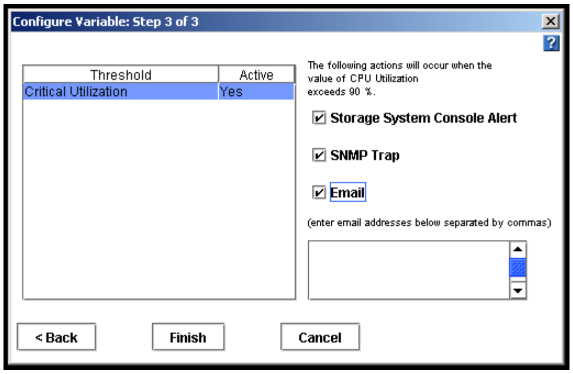

Configure Variable: Step 3 of 3

| Threshold | Active |
|---|---|
| Critical Utilization | Yes |

The following actions will occur when the value of CPU Utilization exceeds 90 %.

☑ Storage System Console Alert

☑ SNMP Trap

☑ Email

(enter email addresses below separated by commas)

Figure 103. Setting Alerts for Monitored Variables

4. For each threshold listed, select the type of alert you want to receive.

### Table 32. Types of Alerts Available for Active Monitoring

| Type of Alert | Where Alerts Are Sent |
|---|---|
| Console alerts | To the Alert Message area of the Console and the Alerts tab in Reporting. |
| SNMP traps | To the SNMP trap community managers. You must have configured the SSM to use SNMP. |
| Email | To specified email addresses. Type the email addresses to receive the notification, separated by commas. Then configure Email Notification on the Email tab. |

5. [Optional] To apply the alert actions (including the email addresses) that you selected in step 4 to all variables that are monitored on the SSM, select all of the variables, then click Set Notifications.

**Note:** *To save time while setting up active monitoring, specify alert actions for one variable and then check the box to apply those actions to all variables on the SSM. This setting applies* *the same email address and other alert settings to all SSMs. Then, if you need to customize alert actions for a particular variable, you can edit that variable.*

6. Click Finish when you have configured all the threshold items in the list.

## Downloading a Variable Log File

To save the history of a variable, download a copy of the log file.

1. In the list of monitored variables, click the variable for which you want to save the log file.
2. Click Save Log Files on the Active Reporting window.

   The Save Variable Log File window opens.
3. Choose a location for the file.
4. [Optional] Change the name of the log file.
5. Click Save.

   The file is saved to the location you specified.

## Viewing the Variable Summary

You can review the frequency settings and the triggers for a variable in the Monitored Variables list without editing the variable.

1. In the list of monitored variables, select a variable.

   The frequency, thresholds, and notification settings display to the right of the list.

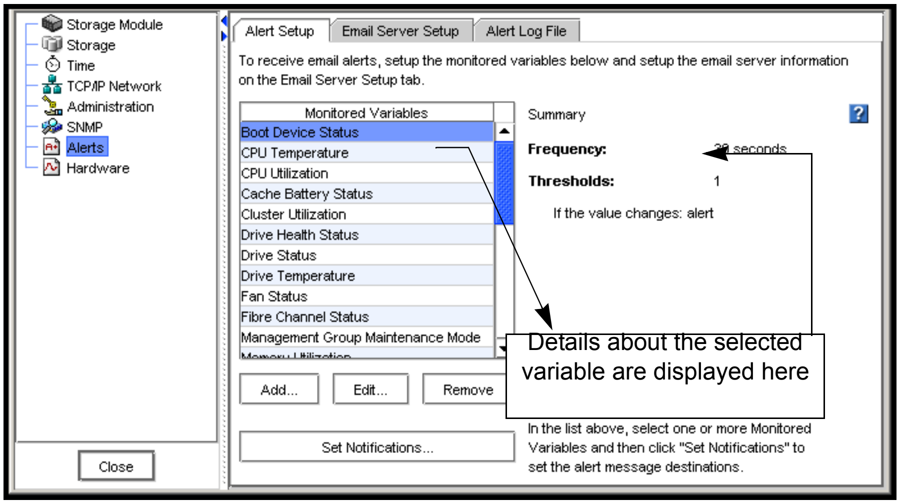

To receive email alerts, setup the monitored variables below and setup the email server information on the Email Server Setup tab.

Boot Device Status
CPU Temperature
CPU Utilization
Cache Battery Status
Cluster Utilization
Drive Health Status
Drive Status
Drive Temperature
Fan Status
Fibre Channel Status
Management Group Maintenance Mode
Memory Utilization

**Frequency:** 30 seconds

**Thresholds:** 1

If the value changes: alert

Details about the selected variable are displayed here

In the list above, select one or more Monitored Variables and then click "Set Notifications" to set the alert message destinations.

**Figure 104. Viewing the Monitoring Variable Summary on the Active Window**

# List of Monitored Variables

The following table shows the variables that are monitored for the SSM. For each variable, the table lists the following information:

- The units of measurement.
- Whether the variable is permanent. (Permanent variables cannot be removed from active reporting.)
- Whether you can change the frequency with which the measurements are taken.
- The default frequency of measurements.
- The default action that occurs if the measured value of the variable reaches a threshold.

### Table 33. List of Variables Available for Active Monitoring

| Variable Name | Units | Perm. Variable | Specify Freq. | Default Freq. | Default Action/ Threshold |
|---|---|---|---|---|---|
| Boot Devices Status | Active,<br>Inactive,<br>Failed,<br>Empty,<br>Unformatted,<br>Not recognized,<br>Unsupported | No | Yes | 30 seconds | Console alert if not normal |
| Cache Battery Status | Normal, Charging, Missing, Faulty | No | Yes | 1 minute | Console alert if changes |
| CPU Utilization | Percent | No | Yes | 1 minute | Console alert at > 90% |
| CPU Temperature | Celsius | Yes | No | 15 seconds | Warning at 65º, Console alert [SSR316MJ2] Critical at 70º, Console alert, Shutdown [SSR212MA] Critical at 90º, Console alert, Shutdown |
| Drive Health Status | Normal,<br>Marginal,<br>Faulty | Yes | Yes | 1 minute | Console alert if not normal |
| Drive Status | On and secured, Off and secured, Off or removed | No | Yes | 1 minute | Console alert if changes |

### Table 33. List of Variables Available for Active Monitoring

| Variable Name | Units | Perm. Variable | Specify Freq. | Default Freq. | Default Action/ Threshold |
|---|---|---|---|---|---|
| Drive Temperature | Celsius | Yes | No | 1 minute | Warning at 60º, Console alert Critical at 65º, Console alert, Drive Power Off |
| Fan Status | Normal,<br>Faulty | No | Yes | 1 minute | Console alert if changes |
| Memory Utilization | Percent | No | Yes | 1 minute | Console alert at > 90% |
| Motherboard Temperature | Celsius | Yes | No | 15 seconds | Warning at 65º, Console alert Critical at 70º, Console alert, Shutdown |
| Network Interface Status | - | No | Yes | 1 minute | Console alert if NIC status changes |
| NVRAM Status | - | Yes | Yes | 1 minute | Console alert if not normal |
| Power Supply Status | -- | No | Yes | 1 minute | Console alert if status changes |
| RAID Status | -- | Yes | Yes | 15 seconds | Console alert if changes |
| Remote Copy Complete | - | No | Yes | 15 minutes | Console alert if true |
| Remote Copy Failovers | - | No | Yes | 15 minutes | Console alert if true |
| Remote Copy Status | - | No | Yes | 15 minutes | Console alert if fails |
| Remote Management Group Status | - | No | Yes | 1 minute | Console alert if changes |
| Snapshot Status | - | No | Yes | 1 minute | Console alert if snapshot status changes |
| Storage Server Status | - | No | Yes | 1 minute | Console alert if not up |
| Volume Restripe Complete | - | No | Yes | 1 minute | Console alert if completed |
| Volume Status | - | No | Yes | 15 minutes | Console alert if volume status changes |

**Table 33. List of Variables Available for Active Monitoring**

| Variable Name | Units | Perm. Variable | Specify Freq. | Default Freq. | Default Action/ Threshold |
|---|---|---|---|---|---|
| Volume Thresholds | -- | No | Yes | 15 minutes | Consolealert if threshold exceeded for any volume or snapshot in the mgt. group |

# Setting Email Notification

If you request email notification on the Active tab, you set the email addresses to receive the notifications there. You then use the Email tab to configure the SMTP settings for email communication. For more information on configuring active monitoring, see “Using Active Monitoring” on page 149.

To complete the request for email notification that you configured for monitored variables:

1. In the Alerts category, select the Email Server Setup tab.

   The Email Server Setup window opens.

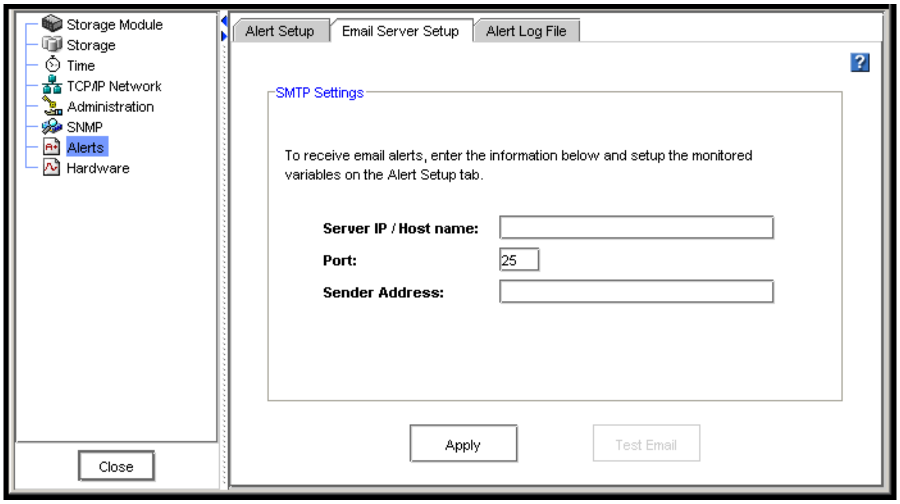

To receive email alerts, enter the information below and setup the monitored variables on the Alert Setup tab.

**Server IP / Host name:**

**Port:** 25

**Sender Address:**

**Figure 105. Configuring Email Settings for Email Alert Notifications**

2. Enter the IP address or host name of the email server.
3. Enter the email port.

   The standard port is 25.

4. (Optional) If your email server is selective about valid sender addresses on incoming emails, enter a sender address, for example, “username@company.com.”

   If you do not enter a sender address, the From field of email notifications will display “root@hostname,” where hostname is the name of the SSM.

5. Click Apply.

**Note:** *Notification of undeliverable email messages are sent to the sender address.*

**Note:** *If you are requesting email notification, be sure to set up the email notification in Active monitoring.*

# Viewing Alerts

Any time that an actively monitored variable causes an alert, the alert is logged by the storage module. If the Console is open, alerts display in the Alert window on the Console main window.

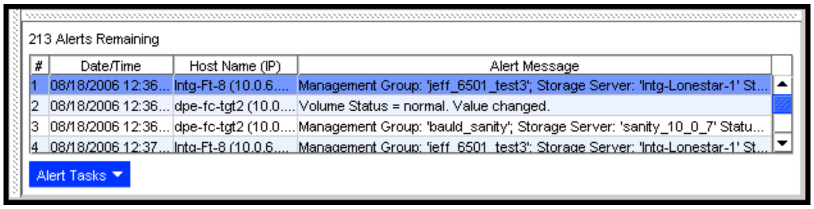

213 Alerts Remaining

| # | Date/Time | Host Name (IP) | Alert Message |
|---|---|---|---|
| 1 | 08/18/2006 12:36... | Intg-Ft-8 (10.0.6.... | Management Group: 'jeff_6501_test3'; Storage Server: 'Intg-Lonestar-1' St... |
| 2 | 08/18/2006 12:36... | dpe-fc-tgt2 (10.0.... | Volume Status = normal. Value changed. |
| 3 | 08/18/2006 12:36... | dpe-fc-tgt2 (10.0.... | Management Group: 'bauld_sanity'; Storage Server: 'sanity_10_0_7' Statu... |
| 4 | 08/18/2006 12:37... | Intg-Ft-8 (10.0.6.... | Management Group: 'jeff_6501_test3'; Storage Server: 'Intg-Lonestar-1' St... |

**Figure 106. Alert Messages on Console Main Window**

If the Console is not open, these alerts are still logged, and you can view them in the Alert Log File tab in the Alerts category of the Edit Configuration window the next time you open the Console.

The Alerts tab in the Reporting category displays the most recent alerts, up until the alert list reaches 1 MB in size. To view alerts older than those displayed on the Alerts tab, save the Alerts log on the Log Files tab.Open the Edit configuration window to access the Alerts configuration category.

1. Select a storage module in the navigation window.
2. Click Storage Module Tasks in the task window.
3. Select Edit Configuration.
4. Select the Alerts configuration category.

5. Select the Alert Log File tab.

   The Alert Log File window opens.

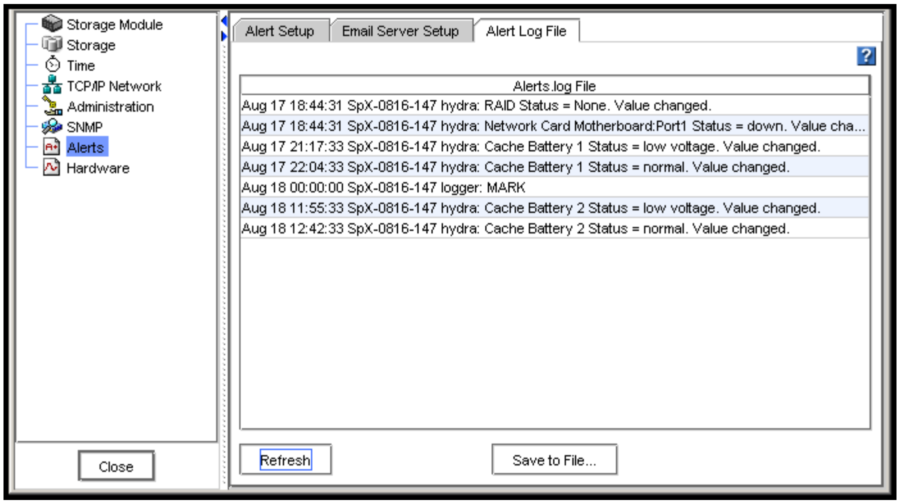

```
Alerts.log File
Aug 17 18:44:31 SpX-0816-147 hydra: RAID Status = None. Value changed.
Aug 17 18:44:31 SpX-0816-147 hydra: Network Card Motherboard:Port1 Status = down. Value cha...
Aug 17 21:17:33 SpX-0816-147 hydra: Cache Battery 1 Status = low voltage. Value changed.
Aug 17 22:04:33 SpX-0816-147 hydra: Cache Battery 1 Status = normal. Value changed.
Aug 18 00:00:00 SpX-0816-147 logger: MARK
Aug 18 11:55:33 SpX-0816-147 hydra: Cache Battery 2 Status = low voltage. Value changed.
Aug 18 12:42:33 SpX-0816-147 hydra: Cache Battery 2 Status = normal. Value changed.
```

**Figure 107. Viewing Alerts**

6. To refresh the list of alerts, click Refresh.
7. (Optional) To save the list of alerts, click Save to File. Then select a location for the file.

# Hardware Overview

The Hardware category includes multiple types of information and reporting capabilities. Review a passive report of system statistics, hardware, and configuration information.

Use the Hardware category to:

- View real-time statistical information about the storage module.
- View and save log files.
- Run hardware diagnostics.

# Running Diagnostics

Use diagnostics to check the health of the SSM hardware.

**Note:** *Running diagnostics can help you to monitor the health of the SSM or to troubleshoot hardware problems.*

To run diagnostic tests:

Open the Edit Configuration window to access the Hardware configuration category.

1. Select a storage module in the navigation window.
2. Click Storage Module Tasks in the tab window.
3. Select Edit Configuration.
4. Select the Hardware configuration category.

   The Diagnostics window opens.

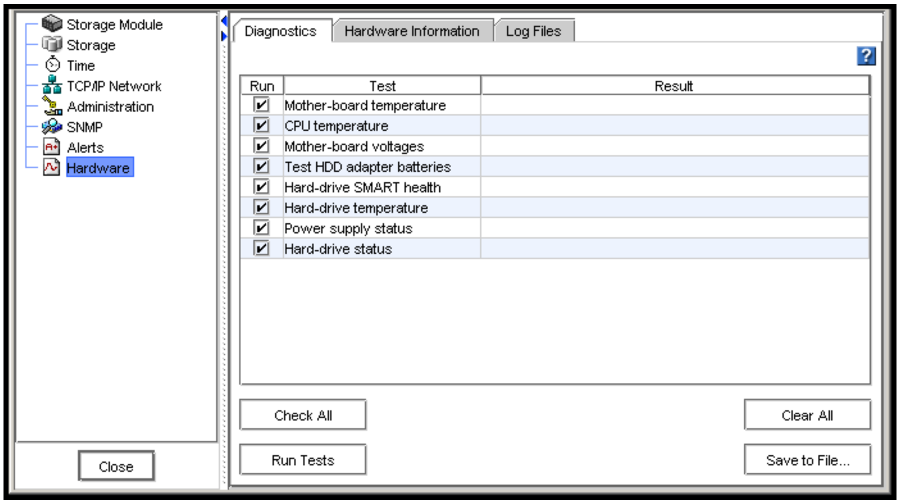

| Run | Test | Result |
|---|---|---|
| ☑ | Mother-board temperature | |
| ☑ | CPU temperature | |
| ☑ | Mother-board voltages | |
| ☑ | Test HDD adapter batteries | |
| ☑ | Hard-drive SMART health | |
| ☑ | Hard-drive temperature | |
| ☑ | Power supply status | |
| ☑ | Hard-drive status | |

Figure 108. Viewing the List of Diagnostics

5. Select the diagnostic tests that you want to run.

   The default setting is to run all tests. Clear any tests that you do not want to run. To clear all selections, click Clear All.

**Note:** *Running all of the diagnostic tests will take several minutes. To shorten the time required to run tests, clear the checkboxes for any tests that you do not need.*

6. Click Run Tests.

   A progress message displays. When the tests complete, the results of each test display in the Result column.

7. [Optional] When the tests complete, if you want to view a report of test results, click Save to File. Then select a location for the diagnostic report file and click Save.

   The diagnostic report is saved as a ".txt" file in the designated location.

## Viewing the Diagnostic Report

The results of diagnostic tests are written to a report file. For each diagnostic test, the report lists whether the test was run and whether the test passed, failed, or issued a warning.

***Note:*** *If any of the diagnostics show a result of "Failed," call your Technical Support representative.*

To view the report file:

1. After the diagnostic tests complete, save the report to a file.
2. Browse to the location where you saved the diagnostics report (.txt) file.
3. Open the report file.

## List of Diagnostic Tests

The following table shows the diagnostic tests that are available for the SSM. For each test, the table lists the following information:

- A description of the test
- Pass / fail criteria.

| Diagnostic Test | Description | Pass Criteria | Fail Criteria |
|---|---|---|---|
| Motherboard temperature | Compares the mother board temperature against the accepted temperature range for normal operation. | Within range | Outside range |
| CPU temperature | Compares the processor temperature against the accepted temperature range for normal operation. | Within range | Outside range |
| Mother board voltages | Compares the power supply voltages against the accepted voltage range for normal operation. | All voltages are within the range | One or more voltages outside range |
| Enclosure<br>communication | Sends a passive command to the backplane and verifies that the response from the backplane matches criteria. | Backplane returns expected string | Backplane times out or does not return expected string |

| Diagnostic Test | Description | Pass Criteria | Fail Criteria |
|---|---|---|---|
| Hard drive status | Checks the status of all installed drives. | All drives are “On and Secured” | One or more drives not “On and Secured” |
| Hard drive temperature | For each of the drives, compares the temperature against an accepted range for normal operation. | Temp. of all drives are within range | One or more drives out of range |
| Fan status | Checks the fan status. | Fan is normal. | Fan is faulty. |
| Power supply status | Checks the power supply status. | Power supply is normal. | Power supply is faulty. |
| Hard drive SMART health | S.M.A.R.T. (Self-Monitoring, Analysis, and Reporting Technology) is implemented in all modern disks. A program inside the disk constantly tracks a range of the vital parameters, including driver, disk heads, surface state, and electronics. This information may be used to predict hard drive failures. | All drives pass health test | Warning or Failed if one or more drives fails health test |
| RAID controller BBU Status | Checks the status of the RAID controller Battery Backup Unit (BBU). | BBU is Normal | Failed if Charging, Faulty or Missing |

# Using Passive Reports

Passive reports display statistics about the performance of the SSM, its drives and configuration. Statistics in the passive reports are point-in-time data, gathered when you click the Refresh button on the Hardware Information tab.

Open the Edit Configuration window to access the Hardware configuration category.

1. Select a storage module in the navigation window.
2. Click Storage Module Tasks in the tab window.
3. Select Edit Configuration.
4. Select the Hardware configuration category.
5. Select the Hardware Information tab.
6. Click Refresh to display statistics on the Hardware Information tab.

# Saving the Report to a File

1. On the Hardware Information tab, click Save To Save to File to download a text file of the reported statistics.

   The Save dialog opens.

2. Choose the location and name for the report.
3. Click Save.

   The report is saved with a .txt extension.

# Passive Reporting Detail

This list details selected information available on the Hardware Information window. Not all items are listed here.

### Table 34. Selected Details of the Passive Report

| This Term | Means This |
|---|---|
| Hardware Information | Date and time report created. |
| Host Name | Host name of the SSM. |
| IP Number | IP address of the SSM. |
| SSM Software | Full version number for SSM software. |
| GUI Software | Full version number for the Console. |
| Support Key | Support Key is used by a Technical Support representative to log in to the SSM. |
| Boot Devices<br>flash-0, flash-1 | Status information about the compact flash card(s) used to boot the SSM. |
| NIC Data<br>Card<br>Description<br>MAC Address<br>Address<br>Mask<br>Gateway<br>Mode | Information about NICs in the SSM.<br>Indicates which NIC in the list is being described.<br>Card name/manufacturer and capable speed of the NIC.<br>Physical address of the NIC. Each card has a unique MAC (media access control) address.<br>IP address of the NIC.<br>Network mask for NIC.<br>Gateway that the SSM is using.<br>Shows manual/auto/disabled. Manual equals a static IP, auto equals DHCP, disabled means the interface is disabled. |
| DNS Data<br>Server 1, Server 2 | Information about DNS, if a DNS server is being used.<br>IP address of the DNS servers. |
| Memory<br>Total<br>Free<br>Shared | Information about RAM memory in the SSM.<br>Total amount of memory in KB.<br>Total amount of free memory in KB.<br>Total amount of free memory in KB. |

### Table 34. Selected Details of the Passive Report

| This Term | Means This |
|---|---|
| CPU<br>Model Name<br>Speed | Information about the CPU.<br>Model name/manufacturer and capable speed of the CPU.<br>Clock speed of the microprocessor. |
| Load Average | Information about the average load on the system. |
| Machine Uptime | The total time the SSM has been running from initial boot up. |
| Enclosure Firmware Version | Firmware version number for the midplane. |
| Drive Temperature | The temperature of the drive in centigrade. |
| Drive Status<br>Drive Number | Information about the drives in the SSM.<br>[Intel® Storage System SSR212MA] Drive 1 through 12. Indicates a specific drive in the SSM. |
| RAID<br>Rebuild Rate<br>Statistics<br>Unit Number | Information about RAID.<br>RAID Rebuild Rate is a percentage of RAID card throughput.<br>Information about the RAID for the SSM.<br>Identifies disks that make up the RAID configuration, their RAID level, chunk size, and device name. |
| Power Supplies<br>Number 1, Number 2 | Status information about the power supplies. |
| Cache Battery Items | Status information about the batteries. |
| IDE Statistics | Lists the drive number and capacity. |

# Saving Log Files

If Technical Support requests that you send a copy of a log file, the Log Files tab is where you can save that log file as a text file.

The Log Files tab lists two types of logs:

- Log files that are stored locally on the SSM (displayed on the left side of the tab).
- Log files that are written to a remote log server (displayed on the right side of the tab).

**Note:** *Save the log files that are stored locally on the SSM. For more information about remote log files, see "Remote Log Files" on page 165.*

Open the Edit Configuration window to access the Hardware configuration category.

1. Select a storage module in the navigation window.
2. Click Storage Module Tasks in the tab window.
3. Select Edit Configuration.
4. Select the Hardware configuration category.

5. Select the Log Files tab.

   The Log Files window opens.

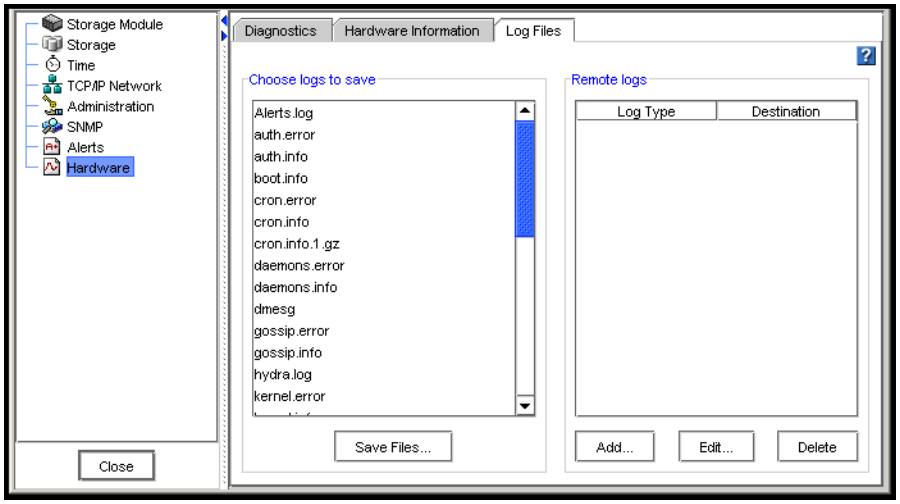

**Figure 109. Saving Log Files to a Local Machine**

6. Scroll down the Choose Logs to Save list.
7. Select the file or files you want to save.

   To select multiple files, use the Ctrl key.
8. Click Save Files.

   The Save dialog opens.
9. Select a location for the file or files.
10. Click Save In Directory.

    The file or files are saved to the designated location.

# Remote Log Files

Use remote log files to automatically write log files to a computer other than the SSM. For example, you can direct the log files for one or more SSMs to a single log server in a remote location. The computer that receives the log files is called the Remote Log Target.

You must also configure the target computer to receive the log files.

## Adding a Remote Log

Open the Edit Configuration window to access the Hardware configuration category.

1. Select a storage module in the navigation window.
2. Click Storage Module Tasks in the tab window.
3. Select Edit Configuration.

4. Select the Hardware configuration category.
5. Select the Log Files tab.

   The Log Files window opens.
6. Click Add below the list of remote logs.

   The Add Remote Log Target window opens.

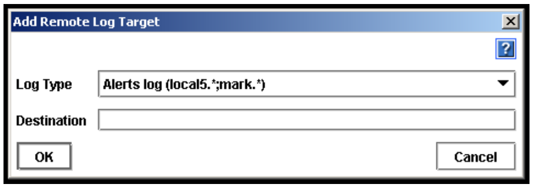

Add Remote Log Target

Log Type: Alerts log (local5.*;mark.*)

Destination:

Figure 110. Adding a Remote Log

7. In the Log Type list, select the log that you want to direct to a remote computer.

   The Log Type list only contains logs that support syslog.
8. In the Destination field, type the IP address or host name of the computer that will receive the logs.
9. Click OK.

   The remote log displays in the Remote logs list on the Log Files tab.

## Configuring the Remote Log Target Computer

Configure syslog on the remote log target computer. Refer to the syslog product documentation for information about configuring syslog.

***Note:*** *The string in parentheses next to the remote log name on the Log Files tab includes the facility and level information that you will configure in syslog. For example, in the log file name:*

*auth error (auth.warning)*

*the facility is "auth" and the level is "warning."*

## Editing Remote Log Targets

You can select a different log file or change the target computer for a remote log:

1. On the Log Files tab, select the log in the Remote logs list.
2. Click Edit.

   The Edit Remote Log Target window opens.
3. Change the log type or destination.
4. Click OK.

## Deleting Remote Logs

To delete a remote log:

1. On the Log Files tab, select the log in the Remote logs list.
2. Click Delete.

   A confirmation message opens.
3. Click OK.

***Note:*** *After deleting a remote log file from the SSM, remove references to this log file from the syslog configuration on the target computer.*

# 9 Working with Management Groups

A management group is a collection of one or more SSMs. It is the container within which you cluster SSMs and create volumes for storage. Creating a management group is the first step towards maximizing the clustering capacity of the SSM.

Management groups serve several purposes:

- **To organize your SSMs into different groups for different functional areas of your organization.** For example, you might create a management group for your Oracle applications and a separate management group for user file share storage.
- **To ensure added administrative security.** For example, you could give one storage administrator access to the SSMs in one management group but not in another management group.
- **To prevent some storage resources from being used unintentionally.** If an SSM is not in a management group, the management group cannot use that SSM as a storage resource. For example, all of the SSMs in management group 1 can be pooled together for use by volumes in that group, if you purchase the Scalability Pak upgrade. To prevent a new SSM from being included in this pool of storage, you would put it in a separate management group.
- **To contain clustering managers.** Within a management group, one or more of the SSMs will act as the managers that control data transfer and replication.

This chapter discusses:

- Managers
- Quorum
- Setting the management group time
- Setting the local bandwidth
- Backing up the management group configuration

## Managers Overview

Managers are SSMs within a management group that you designate to govern the activity of all of the SSMs in the group. All SSMs contain the management software, but you must designate which SSMs in the management group you want to act as managers. These SSMs “run” managers, much like a PC runs various services.

# Functions of Managers

Managers have the following functions:

- Control data replication.
- Manage communication between storage modules in the cluster.
- Coordinate reconfigurations as SSMs are brought up and taken offline.
- Re-synchronize data when SSMs change states.

# Managers and Quorum

Managers use a voting technology to coordinate SSM behavior. In this voting technology, a strict majority of managers (a "quorum") must be running and communicating with each other in order for the Storage System Software to function. An even number of managers can get into a situation where no majority exists - one-half of the managers do not agree with the other one-half. This may cause the management group to become unavailable.

A SAN state and configuration database quorum must agree on the state of the system.

For optimal fault tolerance, you should have three or five managers in your management group. Three or five managers provide the best balance between fault tolerance and performance.

| Number of Managers | Number for a Quorum | Management Fault Tolerance | Explanation |
|---|---|---|---|
| 1 | 1 | None | If the manager fails, no data control takes place. This arrangement is not recommended. |
| 2 | 2 | None | Even number of managers; not recommended, except in specific configurations. |
| 3 | 2 | High | If one manager fails, 2 remain, so there is still a quorum. (Note: 2 managers are not fault tolerant. See above.) |
| 4 | 3 | High | Even number of managers; not recommended, except in specific configurations. |
| 5 | 3 | High | If one or two managers fail, 3 remain so there is still a quorum. |

# Guide to Creating Management Groups

When creating a management group, you must configure the parameters in the following table.

### Table 35. Management Group Requirements

| Management Group Requirement | Throughput (MB/sec) |
|---|---|
| Configure storage modules | Before you create a management group, make sure all the storage modules for the cluster have the same RAID type. |
| Log in to storage modules | Log in to at least one storage module to create a management group. |
| Starting a manager | A management group must have at least one manager running. So, when you create a new management group, the first storage module added to the group has the manager started automatically. Start managers on other storage modules later. |
| Assigning manager IP addresses | The storage modules that are running managers must have static IP addresses (or reserved IP addresses if using DHCP). That is, the IP address must not change while it is a manager. |

# Creating a Management Group

Creating a management group is the first step in the process of creating clusters and volumes for storage. Tasks included in creating a management group are:

- Adding SSMs to the management group
- Starting managers on selected SSMs

# Getting There

You may access the Create Management Group window in these ways:

- By right-clicking an available storage module in the navigation window.

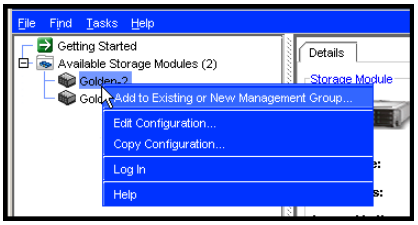

**Figure 111. Create a Management Group from the Navigation Window**

- By selecting the Available Storage Modules level in the navigation window and then right-clicking a storage module name from the list in the tab window

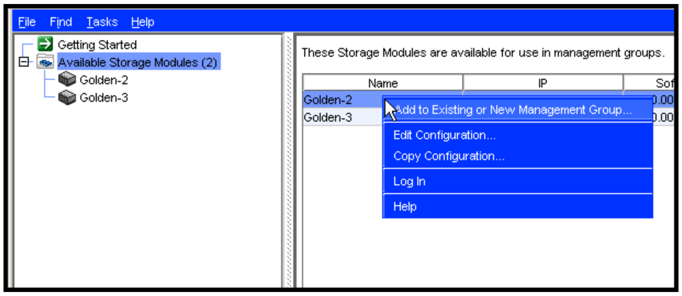
These Storage Modules are available for use in management groups.

**Figure 112. Create a Group from the Tab Window**

- From the Management Groups wizard on the Console window.
- From the menu bar with Tasks > Management Group > Log In.

# Adding the First Storage Module to a New Management Group

1. In the navigation window, select a storage module from the Available Storage Modules group.
2. In the tab window, click Storage Module Tasks and select Log In.

3. Click Storage Module Tasks again and select Add to New or Existing Management Group.

   The Add To or Create New Management Group window opens, shown in Figure 113.

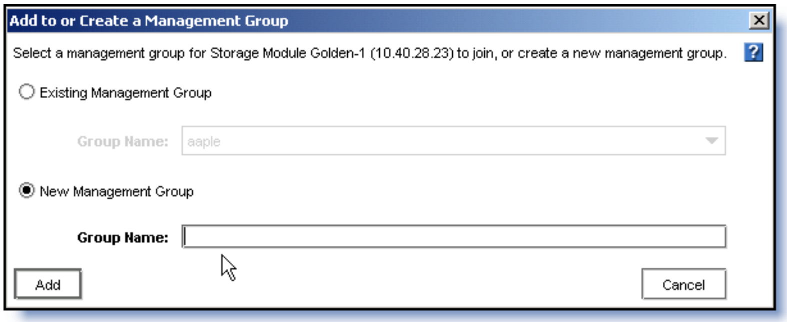

Add to or Create a Management Group

Select a management group for Storage Module Golden-1 (10.40.28.23) to join, or create a new management group.

○ Existing Management Group

Group Name: aaple

● New Management Group

Group Name:

**Figure 113. Creating a New Management Group**

4. Select the radio button for New Management Group.
5. Type a name for the new management group.

   This name cannot be changed later.
6. Click Add.

   The navigation window shows the placement and relationships of the new management group, shown in Figure 114.

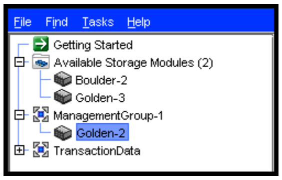

**Figure 114. A New Management Group in the Navigation Window**

You may wish to add more storage modules to the management group or to create a cluster.

## Adding Additional Storage Modules to a Management Group

1. In the navigation window, select an available storage module and log in.
2. Click Storage Module Tasks and select Add to New or Current Management Group.

   The Add to or Create a Management Group window opens.
3. Select the desired management group and click Add.

   The storage module is added to the specified management group.
4. Repeat steps 1 through 3 to add additional storage modules.

# Logging In to a Management Group

You must log in to a management group to administer the functions of that group.

1. In the navigation window, select a management group.

   The management group tab window opens.

2. Click Management Group Tasks and select Log In.

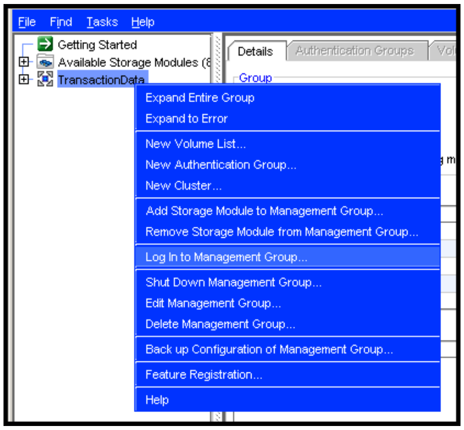
File
Find
Tasks
Help
Getting Started
Available Storage Modules (8
TransactionData
Details
Authentication Groups
Vol
Group
Expand Entire Group
Expand to Error
New Volume List...
New Authentication Group...
New Cluster...
Add Storage Module to Management Group...
Remove Storage Module from Management Group...
Log In to Management Group...
Shut Down Management Group...
Edit Management Group...
Delete Management Group...
Back up Configuration of Management Group...
Feature Registration...
Help

Figure 115. Logging in to a Management Group

# Starting and Stopping Managers

After adding the storage modules to the management group, start managers on the additional storage modules in the management group. The number of managers you start depends upon the overall design of your storage system. See “Managers Overview” on page 169 for more information about how many managers to add.

## Starting Additional Managers

1. In the navigation window, in the management group, select a storage module on which to start a manager.

2. Click Storage Module Tasks and select Start Manager.

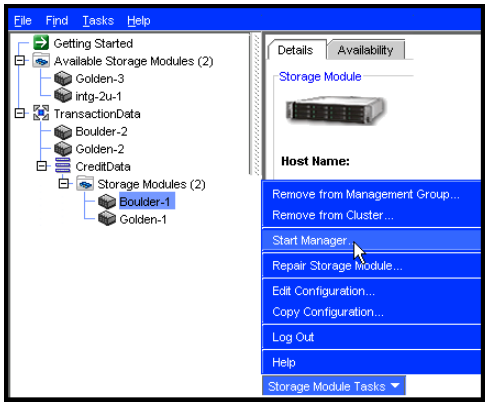

Figure 116. Starting a Manager

Repeat steps 1 and 2 to start managers on additional storage modules.

To start or stop managers on storage modules already in a management group.

1. Log in to the management group.
2. Select the storage module on which you want to start or stop a manager.
3. In the tab window, click Management Group Tasks and select Start Manager/Stop Manager.

   The manager starts on the selected storage module.

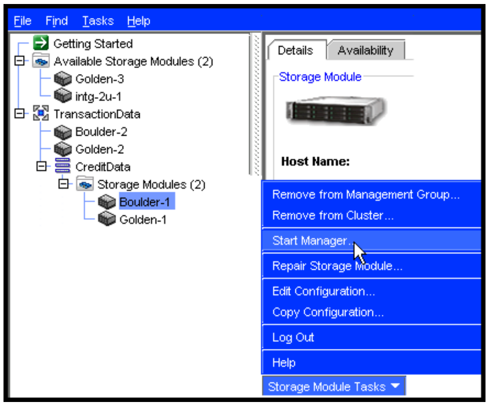

Figure 117. Starting a Manager

## Stopping Managers

Under normal circumstances, you stop a manager when you are removing a storage module from a management group. You cannot stop the last manager in a management group. Deleting the management group is the only way to stop the last manager.

**Implications of stopping managers**

- Quorum of the storage modules may be decreased
- Fewer copies of configuration data are maintained
- Fault tolerance of the configuration data may be lost
- Data integrity and availability may be compromised

***Warning:*** *Stopping a manager can result in the loss of fault tolerance.*

1. Log in to the management group.
2. Select a storage module.
3. Click Storage Module Tasks and select Stop Manager.

   A confirmation message opens
4. Click OK to confirm stopping the manager.

# Editing a Management Group

Editing a management group includes

- Changing Local Bandwidth Priority
- Editing Remote Bandwidth
- Disassociating Management Groups
- Setting Group Mode To Normal.

## Setting or Changing the Local Bandwidth Priority

After a management group has been created, edit the management group to change the local bandwidth priority. This is the maximum rate per second that a manager devotes to non-application processing, such as moving data. The default rate is 4 MB per second. You cannot set the range below .25MB/sec.

## Local Bandwidth Priority Settings

The bandwidth setting is in MB per second. Use Table 36 as a guide for setting the local bandwidth.

### Table 36. Guide to Local Bandwidth Priority Settings

| Network Type | Throughpu t (MB/sec) | Throughput Rating |
|---|---|---|
| Minimum | 0.25 | 2 Mbps |
| Ethernet | 1.25 | 10 Mbps |
| Factory Default | 4.00 | 32 Mbps |
| Fast-Ethernet | 12.50 | 100 Mbps |
| Half Gigabit-Ethernet | 62.50 | 500 Mbps |
| Gigabit-Ethernet | 128.00 | 1 Gbps |
| Bonded Gigabit-Ethernet (2) | 256.00 | 2 Gbps |
| Bonded Gigabit-Ethernet (4) | 512.00 | 4 Gbps |

## Set or Change Local Bandwidth Priority

1. In the navigation window, select a management group and log in.
2. Click Management Group Tasks and select the Edit Management Group command.

   The Edit Management Group window opens.

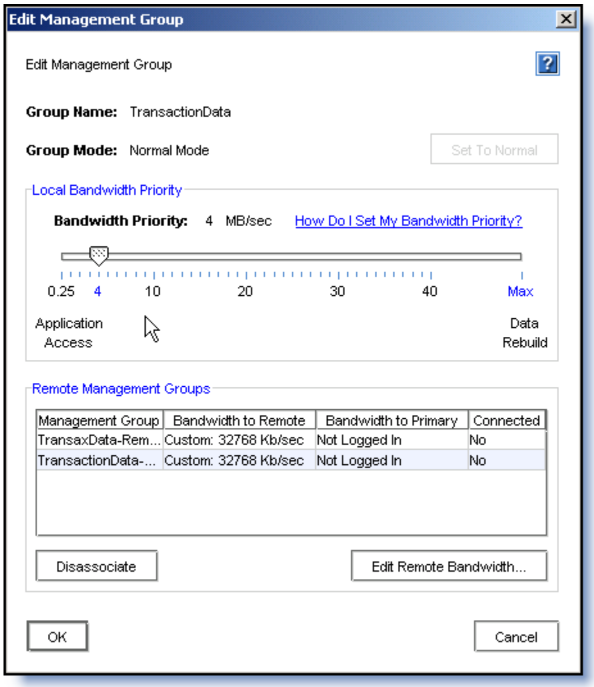

Edit Management Group

**Group Name:** TransactionData

**Group Mode:** Normal Mode

**Bandwidth Priority:** 4 MB/sec

| Management Group | Bandwidth to Remote | Bandwidth to Primary | Connected |
|---|---|---|---|
| TransaxData-Rem... | Custom: 32768 Kb/sec | Not Logged In | No |
| TransactionData-... | Custom: 32768 Kb/sec | Not Logged In | No |

Figure 118. Editing a Management Group

3. Change the local bandwidth priority using the slider.

   A default setting of 4, at the Application Access of the slider, is more appropriate for everyday situations where many clients are busy with the volume. A setting of 40, at the Data Rebuild end of the slider, is most commonly used for quick data migration or copies when rebuilding or moving damaged volumes.

4. Click OK.

   The new rate displays on the Details tab in the management group tab window.

# Logging Out of a Management Group

Logging out of a management group prevents unauthorized access to that management group and the storage modules in that group.

1. In the navigation window, select a management group to log out of.
2. On the Details tab, click Management Group Tasks and select Log Out of Management Group.

# Adding a Storage Module to an Existing Management Group

Storage modules can be added to management groups at any time. Adding a storage module to a management group increases the storage space available to the management group, even though the storage is not used. This addition may or may not have an effect at the cluster level.

**Note:** *All storage modules in a cluster must have the same RAID type.*

1. In the navigation window, select an available storage module that you want to add to a management group.
2. Click Storage Module Tasks and select Add to New or Existing Management Group.

   The Add to or Create a Management Group window opens.

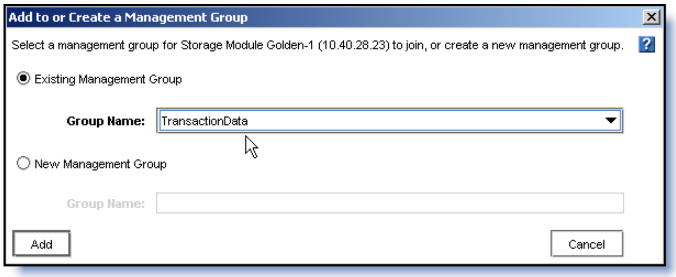

Add to or Create a Management Group

Select a management group for Storage Module Golden-1 (10.40.28.23) to join, or create a new management group.

● Existing Management Group

**Group Name:** TransactionData

○ New Management Group

Group Name:

**Figure 119. Adding a Storage Module to Existing Management Group**

3. Select the desired management group from the drop-down list of existing management groups.
4. Click Add.
5. (Optional) If you want the storage module to run a manager, select the storage module in the management group, right-click, and select Start Manager.

# Resetting the Management Group Time

Whenever you change the time setting of a storage module that is running a manager, you must reset the time of the management group as well. If the manager storage module time is different from the management group time, then

- File creation times on volumes and snapshots might be affected
- Scheduled snapshots might not start at the intended time

**Note:** *Use network time protocol (NTP) to ensure closely synchronized times on the storage modules in the management group.*

## Reset Management Group Time

First, verify the time settings of the storage modules running managers. If necessary, change time settings to ensure all the manager storage modules have the same time. For information about setting the time on the storage module, see "Setting the Date and Time" on page 121.Then refresh the management group time.

1. Log in to the management group.
2. Click the Times tab.
3. Click Time Tasks and select Refresh All.

   Verify the time settings on the storage modules running managers.
4. Click Time Tasks and select Reset Management Group Time.

   A confirmation message opens.
5. Click OK.

   The management group time is now updated.

# Backing Up a Management Group Configuration

Use Backup Configuration of Management Group to save one or both of the following on your local machine:

- Back up configuration—creates a binary file of the management group configuration from which you can restore the management group
- Save the configuration description—creates a text file listing the configuration parameters of the management group

The binary file enables you to automatically recreate a management group with the same configuration. Use the text file for support information or to manually reconstruct the configuration of a management group.

***Note:*** *Backing up the management group configuration does not save the configuration information for the individual storage modules in that management group nor the data. To back up storage module configurations, see ""Backing Up the SSM Configuration File" on page 42.*

1. In the navigation window, select the management group and log in.
2. Click Management Group Tasks and select Back up Configuration of Management Group.

   The Back up Configuration of Management Group window opens.

3. Click
   - — Save Configuration to create a text file that describes your system
   - — Back Up Configuration to create a binary file of your data
   - — OK to exit when you are finished

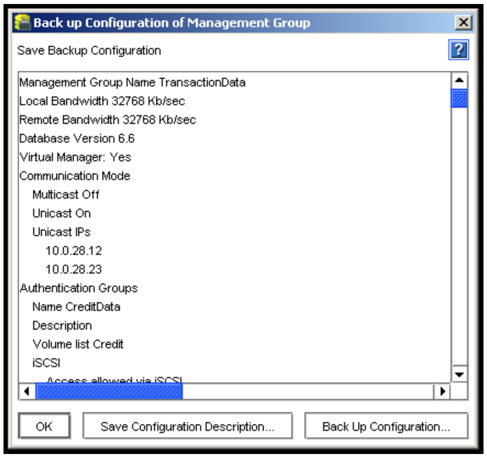
Back up Configuration of Management Group
Save Backup Configuration
?
OK
Save Configuration Description...
Back Up Configuration...

```
Management Group Name TransactionData
Local Bandwidth 32768 Kb/sec
Remote Bandwidth 32768 Kb/sec
Database Version 6.6
Virtual Manager: Yes
Communication Mode
    Multicast Off
    Unicast On
    Unicast IPs
        10.0.28.12
        10.0.28.23
Authentication Groups
    Name CreditData
    Description
    Volume list Credit
    iSCSI
        Access allowed via iSCSI
```

**Figure 120. Backing up the Management Group Configuration**

# Backing Up a Management Group Configuration

1. Click Back Up Configuration.

   The Save window opens.

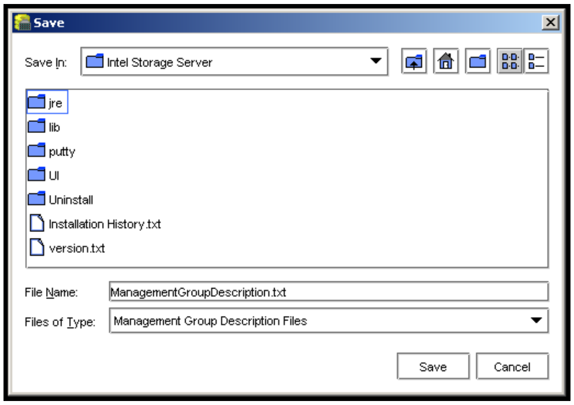

Save In: Intel Storage Server

File Name: ManagementGroupDescription.txt

Files of Type: Management Group Description Files

**Figure 121. Save Window for Backing Up the Management Group Configuration**

2. Navigate to the location in which to store the configuration binary file.
3. Use the default name or type a new name for the file.

4. Click Save.

   The configuration file is saved as a binary file in the folder you selected.

5. Click OK to close the Back Up Configuration window.

### Backing Up a Management Group With Remote Copy Relationships

If you back up a management group that is participating in Remote Copy, it is important to back up the other Remote Copy management groups at the same time. If you back them up at different times, and then try to restore one of the groups, the back up files will not match. This mismatch could cause problems with the restore.

## Saving the Management Group Configuration Description

1. Click Save Configuration Description.

   The Save window opens.

2. Navigate to the location in which to store the configuration description text file.
3. Use the default name or type a new name for the file.
4. Click Save.

   The configuration description is saved as a text file in the folder you selected.

5. Click OK to close the Backup Configuration window.

# Restoring a Management Group

For disaster recovery, use the management group binary file to recreate a management group. From the console, completely disassemble the corrupted management groups and return the storage modules to the available pool.

The restore procedure restores the complete configuration except snapshots, because the data stored in volumes and snapshots is gone.

## Requirements for Restoring a Management Group

- **Hardware**—You must have the same number of storage modules available that are the same capacity or greater.
- **IP Addresses**—You must use the same IP addresses for the replacement storage modules that were assigned to the original storage modules. If you do not have a record of those IP addresses, you can retrieve them when performing the restore. As part of the restore process, the configuration description is displayed and it lists the IP addresses.

### To Restore a Management Group

1. Make sure that the storage modules you are using to restore your management group are in the Available pool in the navigation window.

   If possible, be sure they already have the same IP addresses as the originals.

2. From the menu bar at the top, select Tasks > Management Group > Restore Management Group.

   A standard Open window opens.

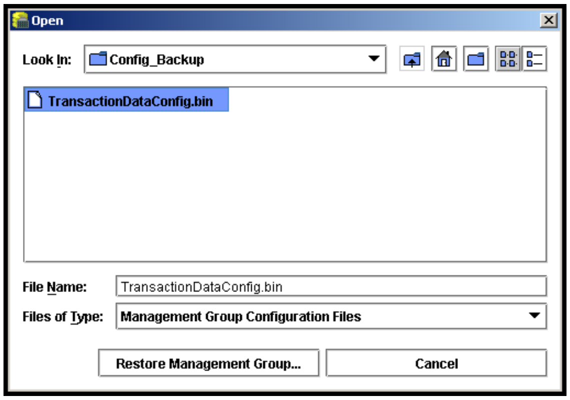

Look In: Config_Backup

TransactionDataConfig.bin

File Name: TransactionDataConfig.bin

Files of Type: Management Group Configuration Files

**Figure 122. Opening the Configuration Binary File**

3. Navigate to the location of the configuration binary file.

4. Select the file and click Restore Management Group.

   The Verify Management Group Configuration window opens.

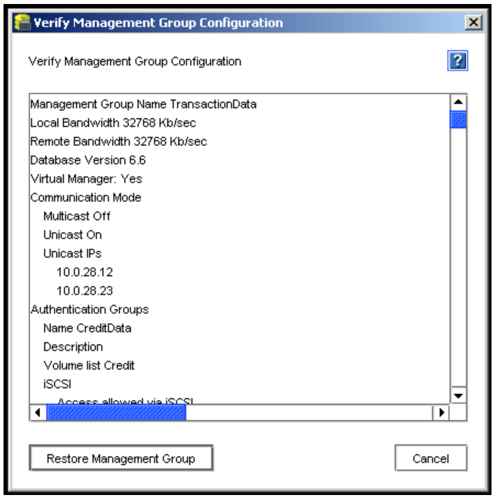

Verify Management Group Configuration

```
Management Group Name TransactionData
Local Bandwidth 32768 Kb/sec
Remote Bandwidth 32768 Kb/sec
Database Version 6.6
Virtual Manager: Yes
Communication Mode
   Multicast Off
   Unicast On
   Unicast IPs
      10.0.28.12
      10.0.28.23
Authentication Groups
   Name CreditData
   Description
   Volume list Credit
   iSCSI
      Access allowed via iSCSI
```

**Figure 123. Verifying the Management Group Configuration**

5. After you have reviewed the configuration parameters, click Restore Management Group.

# Shutting Down a Management Group

Safely shut down a management group to ensure the safety of your data. Shutting down lets you:

- Perform maintenance on storage modules in that group.
- Move storage modules around in your data center.
- Perform maintenance on other equipment such as switches or UPS units.
- Prepare for pending natural disaster.

Also, use a script to configure a safe shut down in the case of a controlled power down by a UPS. See "Working with Scripting" on page 257.

Shutting down a management group also relates to powering off individual storage modules and maintaining access to volumes. See "Powering Off the SSM" on page 46.

**Prerequisites**

- Disconnect any hosts or clients that are accessing volumes in the management group.
- Wait for any restriping of volumes or snapshots to complete.

## Shut Down a Management Group

1. Log in to the management group that you want to shut down.
2. Click Management Group Tasks and select Shut Down Management Group.

   A confirmation window opens, shown below.

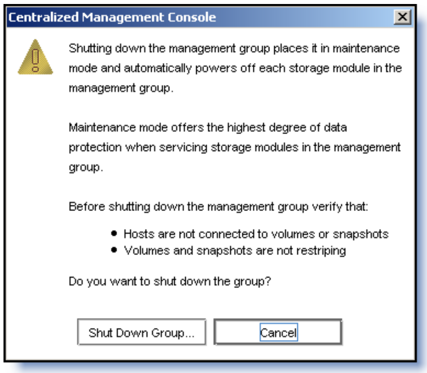

Shutting down the management group places it in maintenance mode and automatically powers off each storage module in the management group.

Maintenance mode offers the highest degree of data protection when servicing storage modules in the management group.

Before shutting down the management group verify that:

- Hosts are not connected to volumes or snapshots
- Volumes and snapshots are not restriping

Do you want to shut down the group?

**Figure 124. Confirm Shutting Down a Management Group**

3. Click Shut Down Group.

   The management group shuts down and disappears from the Console.

**If volumes are still connected to clients or hosts**

After you click Shut Down Group, a confirmation window opens, listing volumes that are still connected and that will go offline if you continue shutting down the management group.

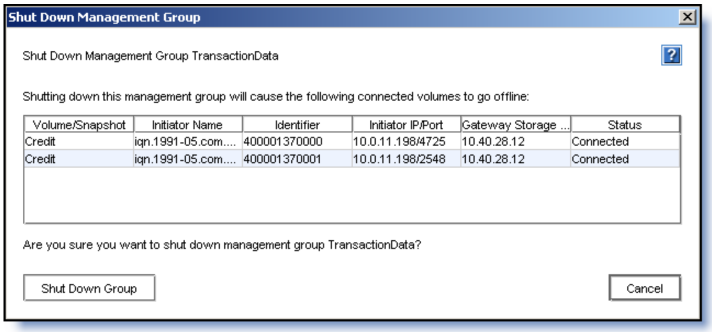

Shut Down Management Group TransactionData

Shutting down this management group will cause the following connected volumes to go offline:

| Volume/Snapshot | Initiator Name | Identifier | Initiator IP/Port | Gateway Storage ... | Status |
|---|---|---|---|---|---|
| Credit | iqn.1991-05.com.... | 400001370000 | 10.0.11.198/4725 | 10.40.28.12 | Connected |
| Credit | iqn.1991-05.com.... | 400001370001 | 10.0.11.198/2548 | 10.40.28.12 | Connected |

Are you sure you want to shut down management group TransactionData?

Figure 125. Notification of Taking Volumes Offline

1. Stop client or host access to the volumes in the list.
2. Click Shut Down Group.

   The management group shuts down and disappears from the Console.

# Starting Up a Management Group

When you are ready to start up the management group, simply power on the storage modules for that group.

1. Power on the storage modules that were shut down.
2. Use Find in the Console to discover the storage modules.

   When the storage modules are all operating properly the volumes come back online and can be reconnected with the hosts or clients.

## Restarted Management Group in Maintenance Mode

In certain cases the management group may start up in maintenance mode. Maintenance mode usually indicates that either the management group is not completely restarted, or the volumes are resynchronizing. When the management group is completely operational and the resyncronizing is complete, the management group changes to normal mode. Figure 126 shows the management group in maintenance mode status.

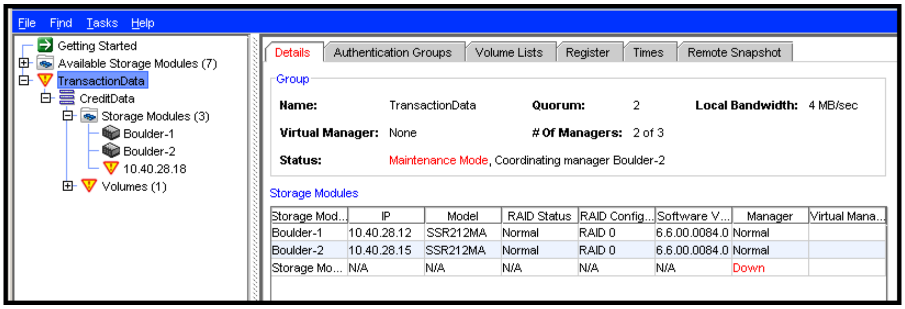

**Name:** TransactionData **Quorum:** 2 **Local Bandwidth:** 4 MB/sec

**Virtual Manager:** None **# Of Managers:** 2 of 3

**Status:** Maintenance Mode, Coordinating manager Boulder-2

| Storage Mod... | IP | Model | RAID Status | RAID Config... | Software V... | Manager | Virtual Mana... |
|---|---|---|---|---|---|---|---|
| Boulder-1 | 10.40.28.12 | SSR212MA | Normal | RAID 0 | 6.6.00.0084.0 | Normal | |
| Boulder-2 | 10.40.28.15 | SSR212MA | Normal | RAID 0 | 6.6.00.0084.0 | Normal | |
| Storage Mo... | N/A | N/A | N/A | N/A | N/A | Down | |

**Figure 126. Identifying Management Group Status in Maintenance Mode**

Some situations which might cause a management group to start up in maintenance mode include these:

- A storage module becomes unavailable, and the management group is shut down while that storage module is repaired or replaced. After the storage module is repaired or replaced and the management group is started up, the management group is in maintenance mode while the repaired or replaced storage module is resynchronizing with the rest of the management group.
- After a management group is shut down, a subset of storage modules is powered on. The management group remains in maintenance mode until the remaining storage modules are powered on and rediscovered in the Console.

## Change Management Group to Normal Mode

While the management group is in maintenance mode, volumes and snapshots are unavailable. You may get volumes and snapshots back online if you manually change the status from maintenance mode to normal mode, depending upon how your cluster and volumes are configured. However, manually changing from maintenance mode to normal mode causes the management group to run in degraded mode while it continues resynchronizing, potentially placing your data at risk.

***Warning:*** *If you are not certain that manually setting the management group to normal mode will bring your data online, or if it is not imperative to gain access to your data, do not change this setting.*

1. In the navigation window, select the management group and log in.

2. Click Management Group Tasks and select Edit Management Group.

   The Edit Management Group window opens, as shown in Figure 127.

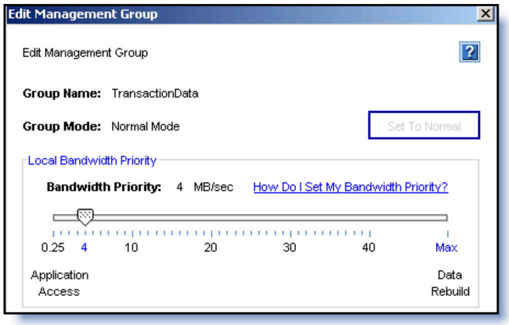

Edit Management Group

Edit Management Group

Group Name: TransactionData

Group Mode: Normal Mode

Bandwidth Priority: 4 MB/sec

**Figure 127. Manually Setting Management Group to Normal Mode**

3. Click Set To Normal.

   The management group is reset to normal mode.

# Removing a Storage Module from a Management Group

**Prerequisites**

- Stop or remove the storage module from data storage activities.
- Let any restripe operations finish completely.
- Remove the storage module from the cluster. See "Removing a Storage Module from a Cluster" on page 207.
- Stop the manager on the storage module if it is running a manager. You may want to start a manager on a different storage module to maintain quorum and the best fault tolerance. See "Stopping Managers" on page 176.

## Removing the Storage Module

1. Log in to the management group from which you want to remove a storage module.
2. In the navigation window, select the storage module to remove.
3. Click Storage Module Tasks and select Remove from Management Group.

   A confirmation message opens.
4. Click OK.

   In the navigation window, the storage module is removed from the management group and moved to Available pool.

# Deleting a Management Group

Delete a management group when you are completely reconfiguring your SAN and you intend to delete all data.

***Warning:*** *When a management group is deleted, all data stored on storage modules in that management group are lost.*

**Prerequisites**

- Log in to each storage module in the management group.
- Remove all volumes and snapshots.
- Delete all clusters.

## Deleting a Management Group

1. In the navigation window, log in to the management group.
2. Click Management Group tasks and select Delete Management Group.

   A confirmation message opens.
3. Click OK.

   After the management group is deleted, the storage modules in it return the Available pool.

# 10 Disaster Recovery Using A Virtual Manager

A virtual manager provides disaster recovery for two specific system configurations. A virtual manager is a manager that is added to a management group, but is not started on an SSM until it is needed to regain quorum. A virtual manager provides disaster recovery for one of two configurations:

- Configurations with only 2 storage modules, or
- Configurations in which a management group spans 2 geographic sites.

See "Managers and Quorum" on page 170 for detailed information about quorum, fault tolerance, and the number of managers.

Virtual manager is part of the add-on module, Scalability Pak. See Chapter 16, "Feature Registration" for information about add-on modules and registering features.

The following are definitions of the terms used in this chapter.

- **Virtual Manager:** A manager which is added to a management group but is not started on an SSM until a failure in the system causes a loss of quorum. The virtual manager is designed to be used in specific system configurations which are at risk for a loss of quorum.
- **Regular Manager**: A manager which is started on an SSM and operates according to the description of managers found in "Managers Overview" on page 169.
- **Manager**: Either a virtual manager or a regular manager.

## When to Use a Virtual Manager

Use a virtual manager in the following configurations:

- A management group across two sites with shared data
- A management group in a single location with two storage modules.

### A management group across two sites

Using a virtual manager allows continuing operation by one site if the other site fails. The virtual manager provides the ability to regain quorum in the operating site if one site goes down, or in one selected site if communication between the sites is lost. Such capability is necessary if volumes in the management group reside on SSMs in both locations.

## A management group in a single location with two SSMs

If you create a management group with two managers in the same location, that management group is in a non-fault tolerant configuration. One manager provides no fault tolerance. Two managers also provide no fault tolerance, due to loss of quorum if one manager goes down. See "Managers and Quorum" on page 170 for more information.

Running two managers and adding a virtual manager to this management group provides the capability of regaining quorum if one manager becomes unavailable.

# Benefits of a Virtual Manager

Running a virtual manager supports implementation of disaster recovery configurations to support full site failover. The virtual manager ensures that, in the event of either a failure of an SSM running a manager, or of communication breakdown between managers (as described in the two-site scenario), quorum can be recovered and, hence, data remains accessible.

# Requirements for Using a Virtual Manager

It is critical to use a virtual manager correctly. A virtual manager is added to the management group, but not started on an SSM until the management group experiences a failure and a loss of quorum. To regain quorum, you start the virtual manager on an SSM that is operating and in the site that is operational or primary, depending upon your situation.

| Requirement | What it Means | | |
|---|---|---|---|
| | Disaster Recovery Scenario | # of SSMs Running Regular Managers | Total # of Managers Including the Virtual Manager |
| Use a Virtual Manager with an Even Number of Regular Managers Running on SSMs | Two sites with shared data | 4 | 5 |
| | Two SSMs in Management Group | 2 | 3 |

| Requirement | What it Means |
|---|---|
| Add a Virtual Manager When Creating Management Group | You cannot add a virtual manager after quorum has been lost. The virtual manager must be added to the management group before any failure occurs. |
| A Virtual Manager Must Only Be Started Once, and Run Only Until the Site is Restored or Communication is Restored | Only one instance of a virtual manager must run at a time. Once you start a virtual manager, you must not start that virtual manager a second time. The virtual manager should run only until the site is restored and data is resynchronized, or until communication is restored and data is resynchronized. |

Illustrations of correct uses of a virtual manager are shown in the following example of two-site failure scenarios.It is important to only start a virtual manager once.

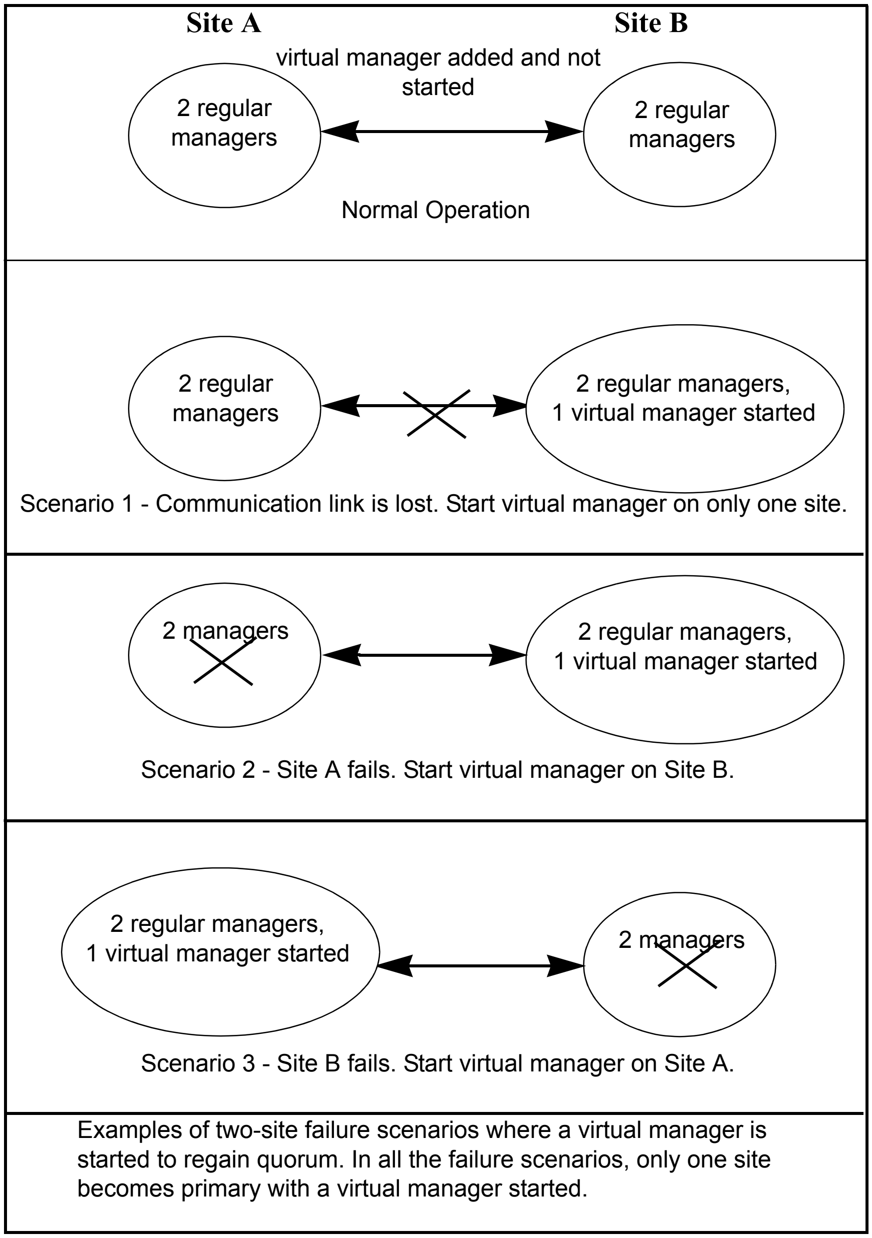
Site A
Site B
virtual manager added and not started
2 regular managers
2 regular managers
Normal Operation
2 regular managers
2 regular managers, 1 virtual manager started
2 managers
2 regular managers, 1 virtual manager started
2 regular managers, 1 virtual manager started
2 managers

Scenario 1 - Communication link is lost. Start virtual manager on only one site.

Scenario 2 - Site A fails. Start virtual manager on Site B.

Scenario 3 - Site B fails. Start virtual manager on Site A.

Examples of two-site failure scenarios where a virtual manager is started to regain quorum. In all the failure scenarios, only one site becomes primary with a virtual manager started.

**Figure 128. Correct Two-site Failure Scenarios Using Virtual Managers**

# Configuring a Cluster for Disaster Recovery

In addition to using a virtual manager, you must configure your cluster and volumes correctly for disaster recovery. This section describes how to configure your system, including the virtual manager.

## Best Practice

The following configuration steps ensure that you have all the data replicated at each site and the managers configured correctly to handle disaster recovery.

For the following example, we are configuring two sites with two SSMs at each site, for an even number of SSMs. The management group contains one cluster. The cluster contains four SSMs and one volume with 2-way replication that spans both sites. That volume must contain all the data in each site.

## Configuration Steps

### Step 1: Name SSMs with Site-Identifying Host Names

To ensure that you can easily identify in the Console which SSMs reside at each site, use host names that identify where each SSM is located. See “Changing the SSM Host Name” on page 37.

- **Management Group Name - Transaction_Data**
- **SSM names**
  - — Denver-1
  - — Boulder-1
  - — Denver-2
  - — Boulder-2

### Step 2: Create Management Group - Plan the Managers and Virtual Manager

When you create the management group in the 2-site scenario, plan to start two managers per site and add a virtual manager to the management group. You then have five managers for fault tolerance. See “Managers Overview” on page 169.

## Step 3: Add SSMs to the Cluster in Alternating Order

Create the cluster. When adding SSMs to the cluster, add them in alternating order, as shown in Figure 129. The order in which the SSMs are added to the cluster determines the order in which copies of data are written to the volume. Alternating the addition of SSMs by site location ensures that data is written to each site as part of the 2-way replication you configure when you create the volume. See "Creating a Cluster" on page 202.

Add SSMs to the cluster in the following order:

1. **1st SSM:** Denver-1
2. **2nd SSM:** Boulder-1
3. **3rd SSM:** Denver-2
4. **4th SSM:** Boulder-2

***Warning:*** *If SSMs are added to the cluster in any order different than alternating order by site, you will not have a complete copy of data on each site.*

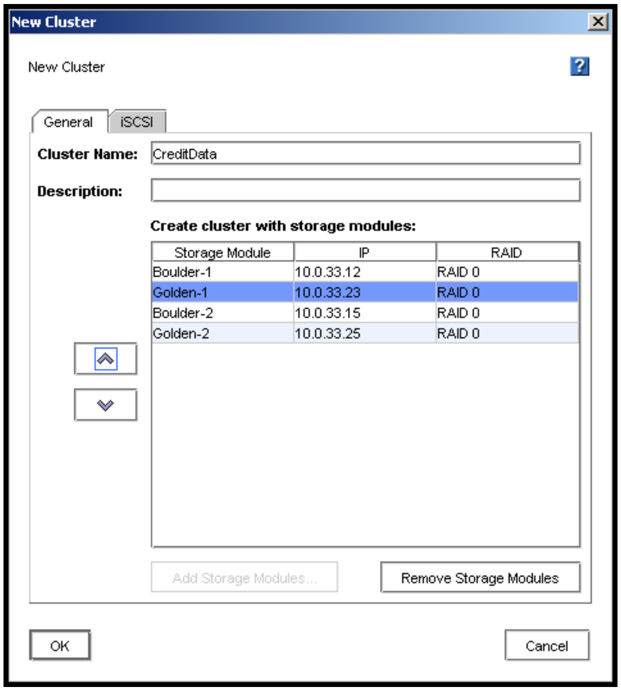

**Cluster Name:** CreditData

**Description:**

**Create cluster with storage modules:**

| Storage Module | IP | RAID |
|---|---|---|
| Boulder-1 | 10.0.33.12 | RAID 0 |
| Golden-1 | 10.0.33.23 | RAID 0 |
| Boulder-2 | 10.0.33.15 | RAID 0 |
| Golden-2 | 10.0.33.25 | RAID 0 |

**Figure 129. Adding SSMs to Cluster in Alternating Site Order**

## Step 4: Create the Volume with 2-way Replication

Two way replication causes two copies of the data to be written to the volume. The fact that you added the SSMs to the cluster in alternating order ensures that a complete copy of the data exists on each site. See “Planning Data Replication” on page 219.

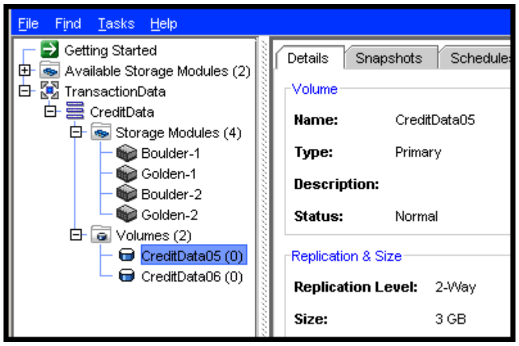

Volume

Name: CreditData05

Type: Primary

Description:

Status: Normal

Replication & Size

Replication Level: 2-Way

Size: 3 GB

Figure 130. 2-Way Replicated Volume on 2-Site Cluster

# Configuring a Virtual Manager

In order to use a virtual manager in a management group beyond the 30-day evaluation period, you must purchase the Scalability Pak. See Chapter 16, “Feature Registration.”

## Adding a Virtual Manager

1. Select the management group in the navigation window.
2. Click Management Group Tasks and select Add Virtual Manager.

   A confirmation message opens.

3. Click OK to continue.

   The virtual manager is added to the management group. The Details tab lists the virtual manager as added and the virtual manager icon appears in the management group as shown in Figure 131.

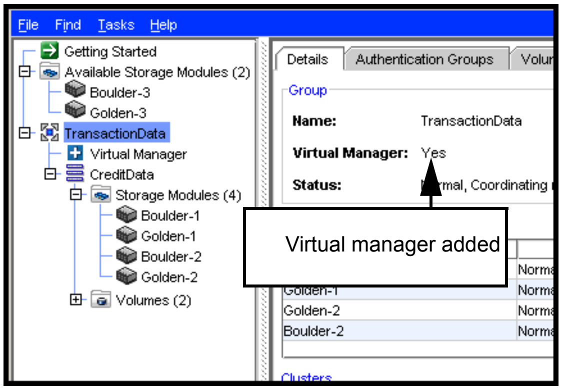

**Name:** TransactionData

**Virtual Manager:** Yes

**Status:** ...mal, Coordinating

Virtual manager added

**Figure 131. Management Group with Virtual Manager Added**

The virtual manager remains added to the management group until needed.

# Starting a Virtual Manager to Regain Quorum

Only start a virtual manager when it is needed to regain quorum in a management group. Figure 128 illustrates the correct way to start a virtual manager when necessary to regain quorum.

- Two-site Scenario, One Site Goes Down

  For example, in the two-site disaster recovery model, one of the sites goes down. On the site that is still up, all managers must be running. Select one of the SSMs at that site and start the virtual manager on it. That site then regains quorum and can continue to operate until the other site is recovered. Once the other site is recovered, the managers in both sites reestablish communication and they ensure that the data in both sites is resynchronized. When the data is resynchronized, stop the virtual manager to return to the disaster recovery configuration.

  ***Note:*** *If the downed site is not recoverable, you can create a new site with new SSMs and reconstruct the cluster. Call your technical support representative.*

- Two-site Scenario, Communication Between the Sites is Lost

  In this scenario, the sites are both operating independently. On the appropriate site, depending upon your configuration, select one of the SSMs and start the virtual manager on it. That site then recovers quorum and operates as the primary site. Once communication between the sites is restored, the managers in both sites reestablish communication and they ensure that the data in both sites is resynchronized. When the data is resynchronized, stop the virtual manager to return to the disaster recovery configuration.

## Starting a Virtual Manager

A virtual manager must be started on an SSM, ideally one that is not already running a manager. However, if necessary, you can start a virtual manager on an SSM that is already running a manager. The following figure shows a management group with an unavailable manager.

1. Select the storage module on which you want to start the virtual manager.
2. Click Storage Module Tasks and select Start Virtual Manager.

   The virtual manager starts on the selected storage module.

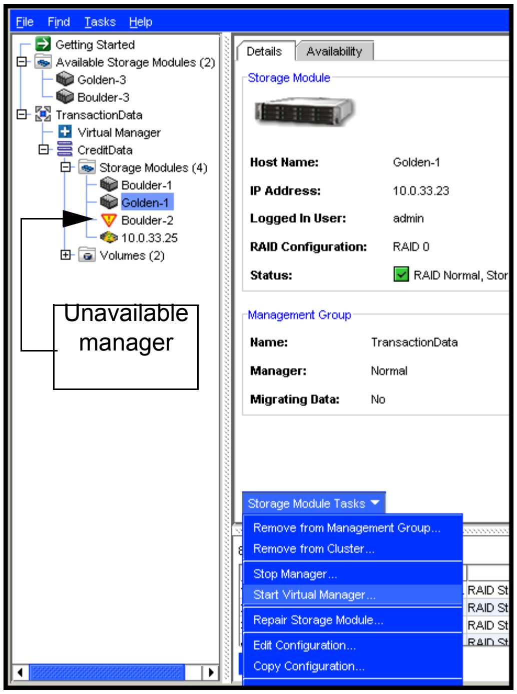

Host Name: Golden-1
IP Address: 10.0.33.23
Logged In User: admin
RAID Configuration: RAID 0
Status: ✓ RAID Normal, Stor

Name: TransactionData
Manager: Normal
Migrating Data: No

**Figure 132. Starting a Virtual Manager When Storage Module Running a Manager Becomes Unavailable**

The virtual manager starts on that storage module.

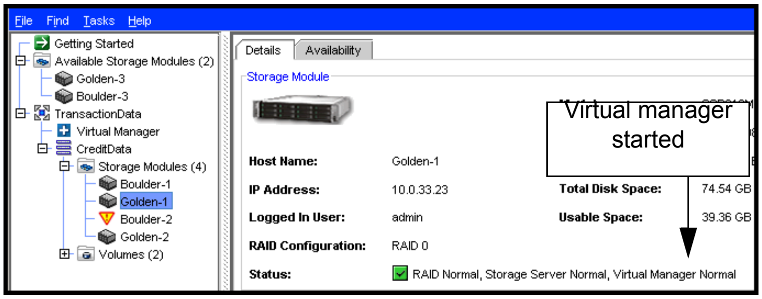

Host Name: Golden-1
IP Address: 10.0.33.23
Logged In User: admin
RAID Configuration: RAID 0
Status: ✓ RAID Normal, Storage Server Normal, Virtual Manager Normal
Total Disk Space: 74.54 GB
Usable Space: 39.36 GB

**Figure 133. Verifying the Virtual Manager**

***Note:*** *If you attempt to start a virtual manager on an SSM that appears to be up in the Console, and you receive a message that the SSM is down, start the virtual manager on a different SSM. This situation can occur when quorum is lost because the Console may still display the SSM in a normal state, even though the SSM is unavailable.*

## Verifying Virtual Manager Status

Verify whether a virtual manager has been started, and if so, which storage module it is started on.

1. Select the virtual manager icon in the navigation window.

   The virtual manager Details tab opens.

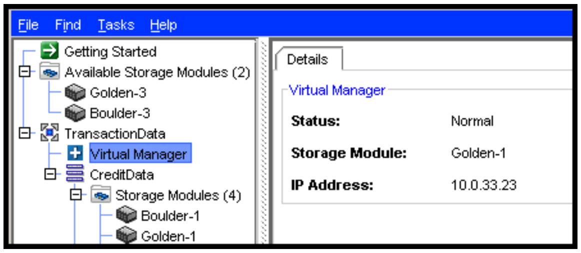

**Status:** Normal

**Storage Module:** Golden-1

**IP Address:** 10.0.33.23

**Figure 134. Identifying the Status of a Virtual Manager**

The Details tab displays the location and status of the virtual manager.

# Stopping a Virtual Manager

When the situation requiring the virtual manager is resolved—either the down site recovers or the communication link is restored—you must stop the virtual manager. Stopping the virtual manager returns the management group to a fault tolerant configuration.

1. Select the storage module with the virtual manager.
2. Click Storage Module Tasks and select Stop Virtual Manager.

   A confirmation message opens.
3. Click OK. The virtual manager is stopped. However, it remains part of the management group and part of the quorum.

## Removing a Virtual Manager

You can remove the virtual manager from the management group altogether.

1. Select the management group from which you want to remove the virtual manager.
2. Click Management Group Tasks and select Delete Virtual Manager.

   A confirmation window opens.
3. Click OK. The virtual manager is removed.

***Note:*** *The Console will not allow you to delete a manager or virtual manager if that deletion causes a loss of quorum.*

# 11 Working with Clusters

## Clusters Overview

Within a management  group you create sub-groups of storage modules called clusters. A cluster is a grouping of storage modules from which you create volumes.

Think of a cluster as a pool of storage. You add storage to the pool by adding storage modules. You then carve volumes out of the pool. Volumes seamlessly span the storage modules in the cluster.

### Topics Covered in This Chapter

- Mixing storage modules in a cluster
- iSCSI and clusters
- Creating and managing clusters
- Repairing a storage module

## Clusters and Storage Module Capacity

Clusters can contain storage modules with different capacities. However, all storage modules in a cluster will operate at a capacity equal to that of the smallest capacity storage module.

**Prerequisites**

All the storage modules in a cluster must be configured the same as the RAID type.

You may find it helpful to have the same settings for reporting, monitoring and time settings for all storage modules in a cluster.

Before you create a cluster, you must have created a management group. See "Creating a Management Group" on page 171.

## Clusters and iSCSI

If you plan to use iSCSI with the Storage System Software, there are iSCSI features you configure at the cluster level, either when you create the cluster or by editing the cluster to configure these items.

- Virtual IP Address - A virtual IP address is a highly available address that ensures that if a storage module in a cluster becomes unavailable, clients can still access the volume through the other storage modules in the cluster.

  A virtual IP address is required for a fault tolerant SAN.
- iSNS Server - An iSNS server simplifies the discovery of iSCSI targets on multiple clusters on a network. If you use an iSNS server, configure your cluster to register the iSCSI target with the iSNS server. You can use up to 3 iSNS servers.

Requirements for Using a Virtual IP Address

The following table lists the requirements for using a virtual IP address.

| Requirements for Using a Virtual IP |
|---|
| Storage modules must be in same subnet address range as the virtual IP. |
| The virtual IP must be routable regardless of which storage module it is assigned to. |
| iSCSI clients must be able to ping the virtual IP when it is enabled in a cluster. |
| Must be different than storage module IPs on the network. |
| Must be a specific IP reserved for this purpose. If you use DHCP, you must use a static IP. |
| All iSCSI initiators must be configured to connect to this IP for failover. |

## iSCSI Load Balancing

Use the iSCSI load balancing feature to improve iSCSI performance and scalability by distributing iSCSI sessions for different volumes evenly across storage modules in a cluster. iSCSI load balancing uses iSCSI Login-Redirect. Only initiators that support Login-Redirect should be used.

Requirements

- Virtual IP address
- Authentication group configured for iSCSI load balancing. See “Client Access Using iSCSI” on page 266.

# Creating a Cluster

Creating a cluster is the first step in designating space for storage in a management group.

**Note:** *If you plan to have two clusters, each with one storage module, the most reliable configuration is to create two management groups with one storage module in each group.*

1. Log in to the management group for which you want to create a cluster.
2. Click Management Group Tasks and select New Cluster.

   The New Cluster window opens with the General tab on top.

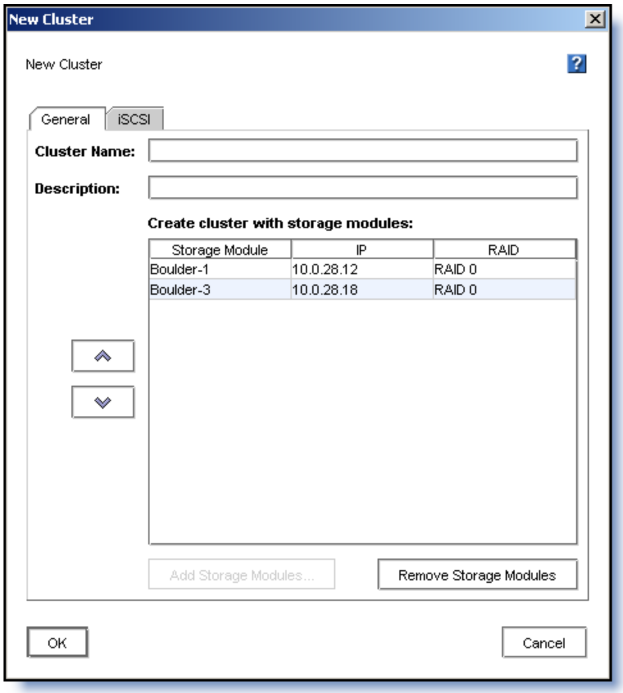

| Storage Module | IP | RAID |
|---|---|---|
| Boulder-1 | 10.0.28.12 | RAID 0 |
| Boulder-3 | 10.0.28.18 | RAID 0 |

**Figure 135. Creating a New Cluster**

3. Type a meaningful name for the cluster.

   A cluster name is case sensitive and must be from 1 to 127 characters. It is also uneditable after this.
4. (Optional) Type a description of the cluster.
5. Select one or more storage modules from the list in the middle of the window.

   Use the up and down arrows on the left to promote and demote storage modules in the list to set the logical order in which they appear. For information about the specific disaster recovery configuration when the order matters, see "Configuring a Cluster for Disaster Recovery" on page 193.

   ***Note:*** *The storage modules in the list are all those included in the management group that are not already in a cluster.*
6. Click Add.

## Configure Virtual IP and iSNS for iSCSI

A virtual IP address is required for iSCSI load balancing and fault tolerance.

1. Click the iSCSI tab to bring it to the front, shown in Figure 136.

   The choice to use a virtual IP is enabled by default.

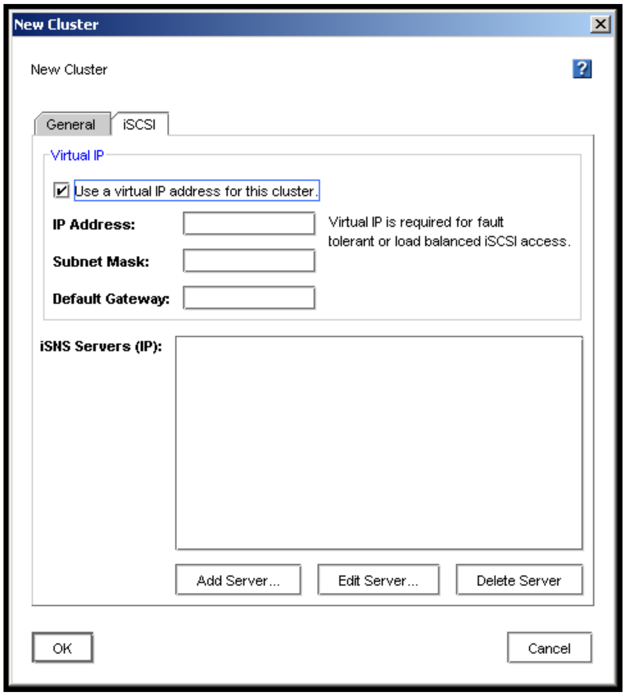

☑ Use a virtual IP address for this cluster.

Virtual IP is required for fault tolerant or load balanced iSCSI access.

**Figure 136. Configuring a Virtual IP Address for iSCSI**

2. Add the IP address, subnet mask and default gateway if required.

## Adding an iSNS Server

[Optional] Add an iSNS server.

***Note:*** *If you use an iSNS server, you may not need to add Target Portals in the Microsoft iSCSI Initiator.*

3. Click Add.

   The Add iSNS Server window opens, shown in Figure 137.

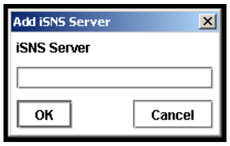

**Figure 137. Adding an iSNS Server**

4. Type the IP address of the iSNS server.
5. Click OK.

The server is added to the list, shown in Figure 138.

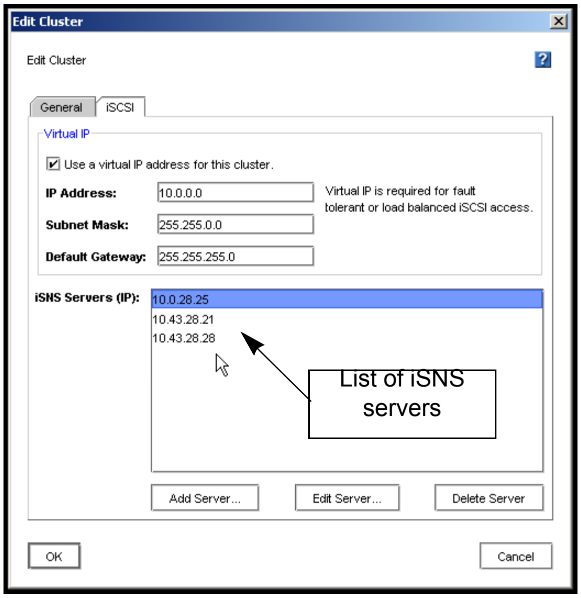

☑ Use a virtual IP address for this cluster.

**IP Address:** 10.0.0.0

**Subnet Mask:** 255.255.0.0

**Default Gateway:** 255.255.255.0

Virtual IP is required for fault tolerant or load balanced iSCSI access.

**iSNS Servers (IP):**
10.0.28.25
10.43.28.21
10.43.28.28

**Figure 138. Viewing the List of iSNS Servers**

6. Click OK when you have finished.

   The cluster is created and displayed inside the management group, shown in Figure 139.

7. Select the cluster to open the Clusters tab window, also shown in Figure 139.

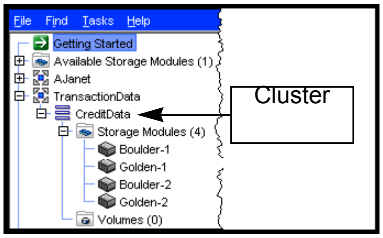

**Figure 139. Viewing a Cluster and the Navigation Window**

# Editing a Cluster

When editing a cluster, you can change the description, and add or remove storage modules. You can also edit or remove the virtual IP and iSNS servers associated with the cluster.

**Prerequisite**

You must log in to the management group before you can edit any clusters within that group.

## Getting There

1. In the navigation window, select the cluster you want to edit.
2. Click Cluster Tasks and select Edit Cluster.

   The Edit Cluster window opens.

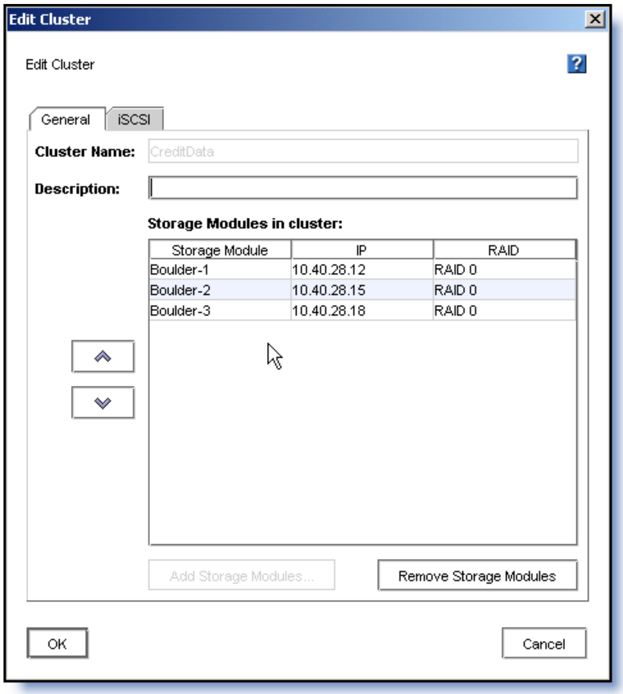

**Cluster Name:** CreditData

**Description:**

**Storage Modules in cluster:**

| Storage Module | IP | RAID |
|---|---|---|
| Boulder-1 | 10.40.28.12 | RAID 0 |
| Boulder-2 | 10.40.28.15 | RAID 0 |
| Boulder-3 | 10.40.28.18 | RAID 0 |

Figure 140. Editing a Cluster

## Adding a New Storage Module to an Existing Cluster

Add a new storage module to an existing cluster to expand the storage for that cluster.

Adding a new storage module is not the same as replacing a repaired storage module with a new one. If you have repaired a storage module and want to place it in a cluster, see “Repairing a Storage Module” on page 209.

### Prerequisites

- Configure the new storage module to match the RAID type already in the cluster. See “Configuring Multiple Storage Modules” on page 11 for information about what features must be configured.
- Add the storage module to the management group that contains the cluster.

# Cluster Capacity

All storage modules in a cluster operate at a capacity equal to that of the smallest capacity storage module.

Be certain that the capacity of the storage module you add to the cluster matches the capacity of the storage modules already in the cluster. If you add a storage module with a smaller capacity, the capacity of the entire cluster might be reduced. If you have three storage modules, two of which have a capacity of 1 TB and one with a capacity of 2 TB, all three SSNs will operate at a 1 TB capacity.

# Adding Storage to a Cluster

1. Select the cluster in the navigation window, and display the Edit Cluster window.
2. Click Add Storage Modules.

   The system displays the Add Storage Modules to Cluster window.

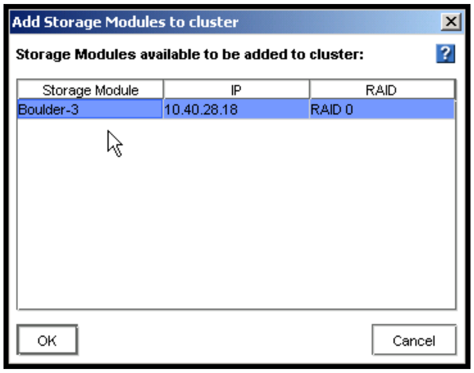

Storage Modules available to be added to cluster:

| Storage Module | IP | RAID |
|---|---|---|
| Boulder-3 | 10.40.28.18 | RAID 0 |

Figure 141. Adding a Storage Module to a Cluster Window

3. Select a storage module from the list.

   The storage modules in this list are in the management group, but not assigned to any cluster.
4. Click OK.
5. Click OK again in the Edit Clusters window.

# Removing a Storage Module from a Cluster

You can remove a storage module from a cluster only if the cluster contains sufficient storage modules to maintain the existing volumes and replication level. See Chapter 12, “Working with Volumes,” on page 215 for details about editing volumes.

1. In the Edit Cluster window, select a storage module from the list.
2. Click Remove Storage Module.

In the navigation window, that storage module moves out of the cluster, but is still in the management group.

3. Click OK when you are finished.

***Note:*** *Changing the order of the storage module list causes a full cluster restripe.*

# Changing or Removing the Virtual IP

Anytime you change or remove the virtual IP address for iSCSI volumes, you are changing the configuration that clients are using. Before making this change, disconnect any clients.

## Preparing Clients

- Quiesce any applications that are accessing volumes in the cluster.
- Log off the active sessions in the server’s iSCSI initiator for those volumes.

## Changing the Virtual IP Address

1. In the Edit Cluster window, click the iSCSI tab to bring it to the front.
2. Change the entries in the IP Address, Subnet Mask and Default Gateway fields.

## Removing the Virtual IP Address

1. In the Edit Cluster window, click the iSCSI tab.
2. Clear the Use a Virtual IP check box.

## Finishing Up

1. Click OK when you are finished changing or removing the virtual IP.
2. Reconfigure the server’s iSCSI initiator with the changes.
3. Reconnect to the volumes.
4. Restart the applications that use the volumes.

# Changing or Removing an iSNS Server

If you change the IP address of an iSNS server, or remove the server, you may need to change the configuration that clients are using. Therefore, you may need to disconnect any clients before making this change.

### Preparing Clients

- Quiesce any applications that are accessing volumes in the cluster.
- Log off the active sessions in the initiator for those volumes.

### Changing an iSNS Server

1. Select the iSNS server to change.
2. Click Edit Server.

   The Edit iSNS Server window opens.
3. Change the IP address.
4. Click OK.

### Deleting an iSNS Server

1. Select the iSNS server to delete.
2. Click Delete Server.

   A confirmation message opens.
3. Click OK.

### Finishing Up

1. Click OK when you are finished changing or removing an iSNS server.
2. Reconfigure the server's iSCSI initiators with the changes.
3. Reconnect to the volumes.
4. Restart the applications that use the volumes.

# Repairing a Storage Module

Repairing a storage module allows you to replace a failed disk in a storage module that is in a cluster configured for 2-way or 3-way replication and only trigger one resync of the data stored on storage modules in that cluster, rather than restriping. Resyncing the data is a shorter operation than a restripe.

## Prerequisites for Using Repair Storage Module

Volume must have 2-way or 3-way replication.

- Storage module must have the blinking red and yellow triangle the navigation window.

- If the storage module is running a manager, stopping that manager must not break quorum.
- Write down the order in which the storage modules are listed in the cluster. You must ensure that they are all returned to that order when the repair is completed.

# How Repair Storage Module Works

Replacing a failed disk requires this:

- Removing the storage module from the cluster and then the management group
- Replacing the disk
- Returning the storage module to the cluster

Because of the replication level, removing and returning the storage module to the cluster would normally cause the remaining storage modules in the cluster to restripe the data twice—once when the storage module is removed from the cluster and once when it is returned. The Repair Storage Module command creates a placeholder in the cluster, in the form of a "ghost" storage module. This ghost storage module keeps the cluster intact while you remove the storage module, replace the disk, configure RAID, and return the storage module to the cluster. The returned storage module only has to resynchronize with the other 2 storage modules in the cluster.

## Using the Repair Storage Module Command

When a storage module in a cluster has a disk failure, the navigation window displays the storage module and the cluster with a blinking triangle next to them in the tree, shown in Figure 142.

When a storage module fails, an alert appears in the alert window, and the Status label in the tab window shows the failure.

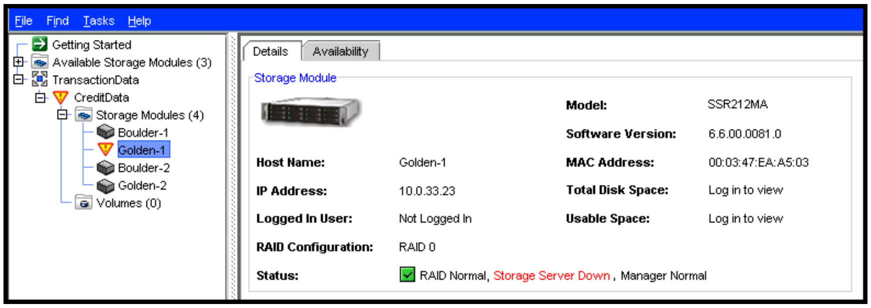

| | | | |
|---|---|---|---|
| | | **Model:** | SSR212MA |
| | | **Software Version:** | 6.6.00.0081.0 |
| **Host Name:** | Golden-1 | **MAC Address:** | 00:03:47:EA:A5:03 |
| **IP Address:** | 10.0.33.23 | **Total Disk Space:** | Log in to view |
| **Logged In User:** | Not Logged In | **Usable Space:** | Log in to view |
| **RAID Configuration:** | RAID 0 | | |
| **Status:** | ✓ RAID Normal, Storage Server Down , Manager Normal | | |

**Figure 142. Finding a Cluster with Failed Storage Module**

1. If the storage module is running a manager, stop the manager. See "Stopping Managers" on page 176.
2. Select the storage module in the navigation window.
3. Right-click and select Repair Storage Module.

A message opens, shown in Figure 143.

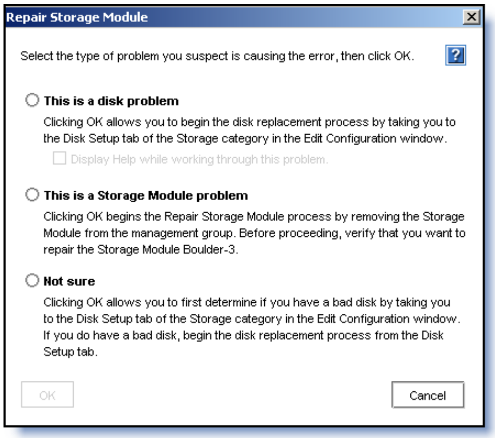

Repair Storage Module

Select the type of problem you suspect is causing the error, then click OK.

○ **This is a disk problem**

Clicking OK allows you to begin the disk replacement process by taking you to the Disk Setup tab of the Storage category in the Edit Configuration window.

□ Display Help while working through this problem.

○ **This is a Storage Module problem**

Clicking OK begins the Repair Storage Module process by removing the Storage Module from the management group. Before proceeding, verify that you want to repair the Storage Module Boulder-3.

○ **Not sure**

Clicking OK allows you to first determine if you have a bad disk by taking you to the Disk Setup tab of the Storage category in the Edit Configuration window. If you do have a bad disk, begin the disk replacement process from the Disk Setup tab.

**Figure 143. Selecting the Problem to Repair**

From this window, select the item that describes the problem you want to solve.

- Repair a disk problem

  If the storage module has a bad disk, be sure to read “Replacing a Disk” on page 73 before you begin the process.

- Storage Module problem

  Select this choice if you have verified that the storage module must be removed from the management group to fix the problem.

- Not sure

  This choice allows you to confirm whether the storage module has a disk problem by taking you directly to the Disk Setup window so that you can verify disk status. As in the first choice, be sure plan carefully for a disk replacement.

4. Click OK.

   The storage module leaves the management group and moves to the Available group. A placeholder, or “ghost” storage module remains in the cluster, shown in Figure 144. It is labelled with an IP address instead of a host name.

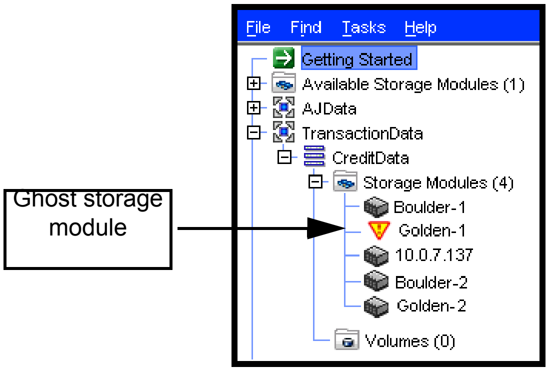

**Figure 144. Viewing the Ghost Storage Module**

Note the ghost in the cluster. Note that the original storage module is now present in the Available pool.

5. Replace the disk in the storage module and perform any other physical repairs.
6. Depending on the model of the storage module, you may need to power on the disk and reconfigure RAID. See “Replacing Disks and RAID” on page 71.
7. Using the navigation window, return the repaired storage module to the management group.

   The ghost storage module remains in the cluster, shown in Figure 144.

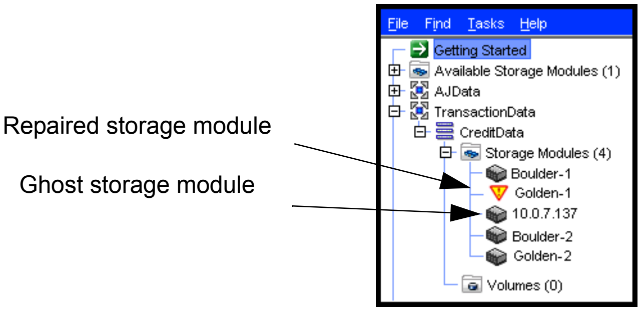

**Figure 145. Returning the Storage Module to the Management Group**

8. Start the manager on the repaired storage module.
9. Edit the cluster and add the repaired storage module to the cluster. It must be returned to the cluster in its former position.

***Warning:*** *The repaired storage module must be returned to the cluster in the same place it originally occupied to have the cluster resync, rather than restripe.*

### To return the repaired storage module to the cluster in the original order

1. In the Edit Cluster window, shown in Figure 146, remove any storage modules in the list that are **below** the ghost storage module.

   The removed storage modules become available in the management group.

2. Remove the ghost storage module.

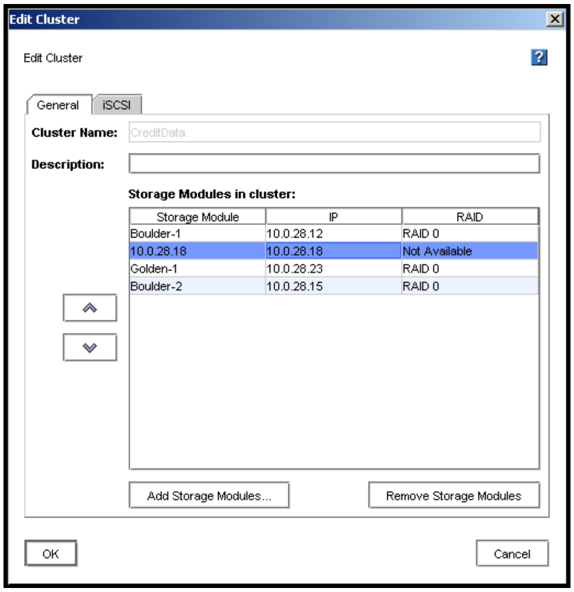

Edit Cluster

Cluster Name: CreditData

Description:

Storage Modules in cluster:

| Storage Module | IP | RAID |
|---|---|---|
| Boulder-1 | 10.0.28.12 | RAID 0 |
| 10.0.28.18 | 10.0.28.18 | Not Available |
| Golden-1 | 10.0.28.23 | RAID 0 |
| Boulder-2 | 10.0.28.15 | RAID 0 |

Figure 146. Removing the Ghost Storage Module from the Cluster

3. Return the removed storage modules to the cluster in the same order they once had.

   Use the arrows to change the order of the storage modules in the list, if necessary.

4. Click OK.

   The storage modules are in the cluster in their original order. The ghost storage module is moved from the cluster to the management group.

5. Select the ghost storage module and remove it from the management group.

   A confirmation message opens, warning that the storage module cannot be found on the network.

6. Click OK to confirm removing storage module from the management group.

   Another confirmation message opens.

7. Click OK.

   The ghost storage module disappears from the navigation window.

# Deleting a Cluster

Volumes and snapshots must be deleted or moved to a different cluster before you can delete the cluster. For more information, see "Deleting a Volume" on page 235**,** and "Deleting a Snapshot" on page 254.

**Prerequisite**

You must log in to the management group before you can delete any clusters within that group.

1. Log in to the management group that contains the cluster you want to delete.
2. In the navigation window, select the cluster you want to delete.

   The Cluster tab window opens.
3. From Cluster Tasks, select Delete Cluster.

   A confirmation message opens. If the message says that the cluster is in use, you must delete the snapshots and volumes on the cluster first.
4. Click OK.

   The cluster is deleted and the storage modules return to the management group as available.

# 12 Working with Volumes

## Volume Overview

A volume is a logical entity that is made up of storage on one or more storage modules. It can be used as raw data storage or it can be formatted with a file system and used by a host or file server. You create volumes on clusters of one or more storage modules.

Before creating volumes, plan your strategies for using the volume: how you plan to use it, its size, how clients will access it, and how you will manage backups of the data, whether through Remote Copy or third-party applications, or both.

## Volumes and Client Access

### Topics Covered in This Chapter

- Planning volume size and thresholds
- Planning data replication and data priority
- Creating and managing volumes

After you create a volume, you must add it to a volume list which is associated with an authentication group. Volume lists and authentication groups control access to volumes by application servers. For detailed information, see Chapter 15, “Controlling Client Access to Volumes.”

**Prerequisite**

Before you create a volume, you must have created a management group and at least one cluster. See Chapter 9, “Working with Management Groups.”  and Chapter 11, “Working with Clusters.”

## Planning Volumes

Planning volumes takes into account multiple factors.

- How many volumes do you need?
- What type of volume are you creating - primary or remote?
- What size do you want the volume to be?
- Do you plan to use snapshots?

- Do you plan to use data replication?
- Do you plan to grow the volume or to keep it the same size?

***Note:*** *If you plan to mount file systems, create a volume for each file system you plan to mount. You can then grow each file system independently.*

## Planning Volume Type

- Primary volumes are volumes used for data storage.
- Remote volumes are used with Remote Copy for business continuance, backup and recovery, and data mining/migration configurations.

## Planning Volume Size

Volume size is the size (or length) of the virtual storage disk device presented to the operating system and to the applications. Configure volume size based on data needs and how you plan to provision your volumes, and whether you plan to use snapshots.

- **Full provisioning**—sets the volume size equal to the hard threshold.
- **Thin provisioning**—sets the hard threshold to less than the volume size.

**Table 37. Volume Provisioning Methods**

| Method | Settings |
|---|---|
| Full provisioning | Volume size = hard threshold = soft threshold |
| Thin provisioning | Volume size > hard threshold > soft threshold |

Provisioning is explained in more detail in “Planning Hard Thresholds” on page 218.

### The Effect of Snapshots on Volume Size

Snapshots take up space on the cluster. Planning how much space, and planning the use and scheduling of snapshots, impacts the hard threshold you should set for the volume. If you are planning to use snapshots, Table 147  lists approximate data change rates for common applications.

**Figure 147. Common Applications’ Daily Change Rates**

| Application | Daily Change Rates |
|---|---|
| Fileshare | 1 - 3% |
| Email/Exchange | 10 - 20% |
| Database | 10% |

***Note:*** *Volume size, volume thresholds, and using snapshots should be planned in conjunction. If you intend to use snapshots, review “Working with Snapshots” on page 237*

## Best Practice for Setting Volume Size

Create the volume with the size that you currently need. Later, when you need to make the volume bigger, increase the volume size in the Console and then grow the file system on the operating system. In Microsoft Windows*, you expand a basic disk using Windows Logical Disk Manager and Diskpart. For detailed instructions, see the Support Note, “Expanding a Windows Basic Disk,” on the Customer Resource Center.

# Measuring Disk Capacity and Volume Size

All operating systems that are capable of connecting to the SAN via iSCSI deal with two disk space accounting systems-the block system and the native file system (on Windows, this is usually NTFS).

**Table 38. Common File Systems for Different Operating Systems**

| OS | File System Names |
|---|---|
| Windows | NTFS, FAT, |
| Linux | EXT2, EXT3 |
| Netware | NWFS |
| Solaris | UFS |
| VMWare | VMFS |

## Block Systems and File Systems

Operating systems see hard drives (both physical and virtual) as abstractions known as "block devices": arbitrary arrays of storage space, which can be read from and written to at will.

Files on disks are handled by a different abstraction: the "file system." File systems are placed on block devices. File systems are given authority over reads and writes to block devices.

iSCSI do not operate at the file system level of abstraction. Instead, they present the Storage System Software volume to an OS such as Windows as a block device. Typically, then, a file system is created on top of this block device so that it can be used for storage. In contrast, an Oracle database uses a Storage System Software volume as a raw block device.

## Storing File System Data on a Block System

The Windows file system treats this block device as simply another hard drive; that is, it is an array of blocks which the file system can use for storing data. As the iSCSI initiator passes writes from the file system, the Storage System Software simply writes those blocks into the volume. So when you look at the Console, the allocation percentage displayed is based on how many physical blocks have been written for this volume.

Now, when you delete a file, typically the file system updates the directory information which removes that file. Then the file system notes that the blocks which that file previously occupied are now freed up. Subsequently, when you query the file system about how much free space is available, the space occupied by the deleted files appears as part of the free space, since the file system knows it can overwrite that space.

However, the file system does not inform the block device underneath (the Storage System Software volume) that there is freed up space. In fact, no mechanism exists to transmit that information. There is no SCSI command which says "block 198646 can be safely forgotten"; at the block device level there are only reads and writes.

So, to ensure that our iSCSI network block devices work correctly with file systems, any time a block is written to, that block is forever marked as allocated. Then, when all blocks are allocated up to the full size of the storage volume, the file system takes over. The file system reviews its "available blocks" list and reuses blocks that have been freed up.

# Planning Hard Thresholds

The hard threshold is the amount of application data that can actually be written to the volume. This size is the actual physical space reserved for data on the disks in the cluster. Therefore, it is the limit beyond which data can no longer be written to the volume. The hard threshold can be increased up to the volume size, if they are not set as equal.

## Best Practice if Not Using Snapshots - Full Provisioning

For volumes that will not be used with snapshots, hard thresholds should be set equal to the volume size. This setting ensures that the hard threshold cannot be exceeded, which prevents clients from accessing the volume. If you intend to use snapshots, see “Managing Capacity Using Volume and Snapshot Thresholds” on page 239.

## Best Practice if Using Snapshots

For volumes that will be used with snapshots, the minimum recommended hard threshold is 512 MB (0.5 GB) with Auto Grow enabled.

# Planning Soft Thresholds

Soft thresholds trigger alerts to system administrators to help ensure that hard thresholds are not exceeded. Upon receiving an alert, the system administrator can take steps to increase capacity according to planned capacity management. See “Managing Volume Growth Capacity” on page 223 for strategies to manage volume growth.

## Best Practice If Not Using Snapshots

If the hard threshold is equal to the volume size, set the soft threshold equal to the volume size as well. Use application-level monitoring to manage capacity growth.

## Best Practice If Using Snapshots

If the hard threshold is less than the volume size, set the soft threshold to a minimum of 256 MB. There should be at least a 256 MB gap between the hard threshold and the soft threshold.

| | |
|---|---|
| Minimum recommended hard threshold | 512 MB |
| Recommended gap between hard and soft thresholds | 256 MB |
| Minimum soft threshold | 256 MB |

Reaching the soft threshold triggers auto grow. When a soft threshold alert is received, take these actions:

- Provision more storage for the cluster (if required).
- Monitor the hard threshold, and increase it if necessary.
- Re-adjust the soft threshold to be a larger percentage of the new hard threshold.

# Planning Data Replication

Data replication creates redundant copies of a volume. You can create up to three copies using 3-way replication. Because these copies reside on different storage modules, replication levels are tied to the number of available storage modules in a cluster.

The Storage System Software and the Storage System Console provide flexibility when planning data replication through two features.

- Replication level allows you to choose how many copies of data you want to keep in the cluster.
- Replication priority allows you to choose whether availability or redundancy is more important in your configuration.

## Replication Level

Three replication levels are available depending upon the number of available storage modules in the cluster. The level of replication you choose also affects the Replication Priority you can set.

## Table 39. Setting a Replication Level for a Volume

| With This Number of Available Storage Modules in Cluster | Select This Replication Level | For This Number of Copies |
|---|---|---|
| 1 | • None | • 1 copy of data in the cluster. No replica is created. |
| 2<br>(not a recommended configuration for high availability) | • None<br>• 2-Way | • 1 copy of data in the cluster, no replication.<br>• 2 copies of data in the cluster. One replica is created. |
| 3 or more | • None<br>• 2-Way<br>• 3-Way | • 1 copy of data in the cluster (no replication).<br>• 2 copies of data in the cluster (one replica).<br>• 3 copies of data in the cluster. Two replicas are created. |

***Note:*** *The system calculates the actual amount of storage resources needed if the replication level is greater than none.*

## How Replication Levels Work

When you choose 2-way or 3-way replication, data is written to either 2 or 3 consecutive storage modules in the cluster. For example:

**2-Way Replication**

A cluster with 3 storage modules, configured for 2-way replication. There have been 5 writes to the cluster. Figure 148 illustrates the write patterns on the 3 storage modules.

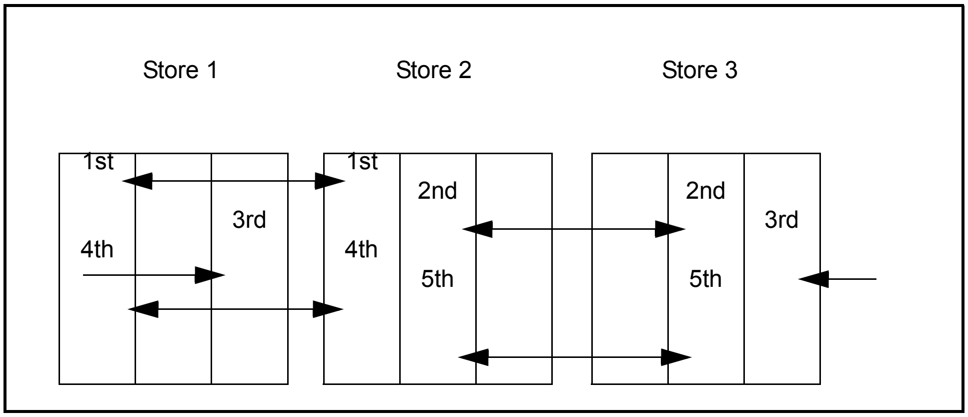

Figure 148. Write Patterns in 2-way Replication

## Replication Priority

Set the replication priority according to your requirements for data availability and/or data redundancy for the volume.

Choose the redundancy mode if you require that the volume be replicated in order to be available. The redundancy mode ensures fault-tolerance.

Choose the availability mode (which is the default) if you want your data to be available even if it cannot be replicated.The availability mode ensures that data may remain available to clients even if a storage module becomes unavailable.

**Table 40. Storage Module Availability and Volume Access by Replication Level and Priority Setting**

| Volume is available to a client with | | | |
|---|---|---|---|
| **a priority setting of:** | **a replication level of:** | | |
| | **None** | **2-way** | **3-way** |
| **Availability** | All storage modules must be up. | 1 of every 2 adjacent storage modules must be up. | 1 of every 3 adjacent storage modules must be up. |
| **Redundancy** | N/A | All storage modules must be up. | 2 of every 3 adjacent storage modules must be up. |

***Warning:*** *A management group with 3 storage modules is the minimum configuration for fault tolerant operation. Although the system allows you to configure 2-way replication on 2 storage modules, this does not guarantee data availability in the event that 1 storage module becomes unavailable. See "Managers Overview" on page 169.*

## Best Practice

For mission-critical data, choose 3-way replication and redundancy priority. This configuration sustains the first fault and ensures that the volume is redundant and available.

If your volumes contain critical data, configure them for 2-way replication and a priority of redundancy.

# Requirements for Volumes

When creating a volume, you define the following parameters.

### Table 41. Parameters for Volumes

| Volume Parameter | Configurable for Primary or Remote Volume | What it means |
|---|---|---|
| Type | Any | Whether the volume is primary or remote.<br>• Primary volumes are used for data storage.<br>• Remote volumes are used for configuring Remote Copy for business continuance, backup and recovery, or data mining/migration.<br>◎ Note: Remote Copy is a feature upgrade. You must purchase a Remote Data Protection Pak license to use remote volumes past the 30-day trial period. |
| Volume Name | Any | The name of the volume that is displayed in the Console. A volume name must be from 1 to 127 characters and is case sensitive. |
| Description | Any | [Optional] A description of the volume. |
| Cluster | Any | If the management group contains more than one cluster, you must specify the cluster on which the volume resides. |
| Replication Level | Any | The number of copies of the data to create on storage modules in the cluster. The replication level must be at most the number of storage modules in the cluster or 3, whichever is smaller. See “Planning Data Replication” on page 219. |
| Replication Priority | Any | • Availability - Default setting. These volumes will remain available as long as at least one storage module out of every n (n = replication level) remains active. When the unavailable storage module returns to active status in the cluster, than the volume re synchronizes across the replicas. A replication of None must have a replication priority setting of Availability.<br>• Redundancy - Choose this setting to ensure that the volume will go offline if it cannot maintain 2 replicas. For example, if 2-way replication is selected, and a storage module in the cluster becomes unavailable, thereby preventing 2-way replication, the volume goes offline until the storage module is again available. |

**Table 41. Parameters for Volumes**

| Volume Parameter | Configurable for Primary or Remote Volume | What it means |
|---|---|---|
| Size | Primary | The logical block storage size of the volume. Hosts and file systems will operate as if storage space equal to the volume size is available in the cluster. This volume size may exceed the true allocated disk space on the cluster for data storage, which facilitates adding more storage modules to the cluster later for seamless storage growth. However, if the volume size does exceed true allocated disk space, the ability to make snapshots may be impacted. See Chapter 13, "Working with Snapshots." Remote volumes contain no data and therefore do not have a size. |
| Hard Threshold | Primary | The amount of physical space allocated for actual data storage. Reaching the hard threshold triggers an alert and data can no longer be written to the volume. The hard threshold is less than or equal to the volume size, but the hard threshold cannot exceed remaining storage space in the cluster.<br>Remote volumes contain no data and do not have a size. You cannot set a hard threshold for a remote volume. |
| Soft Threshold | Primary | The amount of space used on the volume that triggers a warning alert. This alert notifies the storage administrator that the volume is approaching the hard threshold. The soft threshold must be less than or equal to the hard threshold.<br>Because remote volumes have no size, and cannot have a hard threshold, they also cannot have a soft threshold. |
| Auto Grow | Both | Auto grow is on by default. Auto grow automatically increases the hard and soft volume thresholds in the event that data written to the volume exceeds the soft threshold. |

# Managing Volume Growth Capacity

When creating a volume for which you plan to use snapshots, you can set the soft threshold value to help manage capacity growth. This threshold value triggers an alert, providing you the opportunity to increase the capacity of the volume before it is full. The soft threshold value triggers an alert and also triggers auto grow if you have it configured.

***Note:*** *Volume size, replication level, and snapshots should be planned in conjunction. if you intend to use Snapshots, review Chapter 13, "Working with Snapshots."*

# Creating the Volume and Setting Thresholds

- First, create the volume and designate the size. This size is the logical size on the cluster. For example, you have a 750 GB cluster and you create a 500 GB volume.
- Second, set the hard threshold to some size smaller than the actual volume size. For our example 500 GB volume, you set the hard threshold at 300 GB.
- Third, set the soft threshold lower than the hard threshold. The soft threshold triggers an alert to the system administrator, notifying that the soft threshold has been reached. This alert gives you time to increase the volume size and hard threshold. For our example, set the soft threshold at 485 GB.

***Warning:*** *If the hard threshold is set lower than the volume size and the hard threshold is reached, then other applications that are accessing the volume will hang until you increase the hard threshold. In this scenario, system resources will be exhausted. Therefore, if there are other volumes in the cluster, accessed by other applications, those volumes will hang as well, even though those volumes' hard thresholds have NOT been reached.*

*To avoid halting writes to the volume, enable auto grow on the volume.*

## Managing the Volume Growth Capacity

You have two choices for managing volume threshold growth:

- Use auto grow (recommended)
- Manually monitor and increase the hard and soft thresholds

See "Editing a Volume" on page 227 for information about changing the volume size, and the soft and hard thresholds.

Over time, as you near the capacity of the cluster, you can increase the storage capacity of the cluster by adding more storage modules.

***Note:*** *If you have file systems mounted on the host volume, and you reach the soft or hard threshold, deleting files from the volume does not create space on the storage module volume.*

# Using Auto Grow

Auto grow automatically increase the hard and soft thresholds of a volume.

***Note:*** *Auto grow is also available in the application-based scripting described in Chapter 14, "Working with Scripting."*

## How Auto Grow Works

Auto grow is triggered when a soft threshold is reached. Auto grow then raises both the soft and hard thresholds by a calculated increment. The thresholds will only increase

- when there is sufficient room in the cluster to accommodate the increases or
- to the point where the hard threshold equals the volume length (at which point the soft threshold also is increased to equal the volume length so that alerts are not triggered)

whichever of these conditions occur first.

### Auto Grow Options

Auto grow enabled in the Console uses a dynamic algorithm that allows volumes to grow smoothly and "just-in-time," conserving space and limiting the total number of auto grow events to a reasonable number. Also, auto grow causes the volume to grow before the client(s) hit the hard threshold, which would cause the client(s) to stall for a few seconds.

Using auto grow through application-based scripting uses a static algorithm, increasing the hard and soft thresholds by the same amount every time. For more information, see "Working with Scripting" on page 257.

# Creating a Volume

A volume resides on the storage module(s) contained in a cluster.

1. Log in to the management group for which you want to create a volume.
2. In the navigation window, select the cluster on which you want to create a volume.

   The Cluster tab window opens, shown in Figure 149.

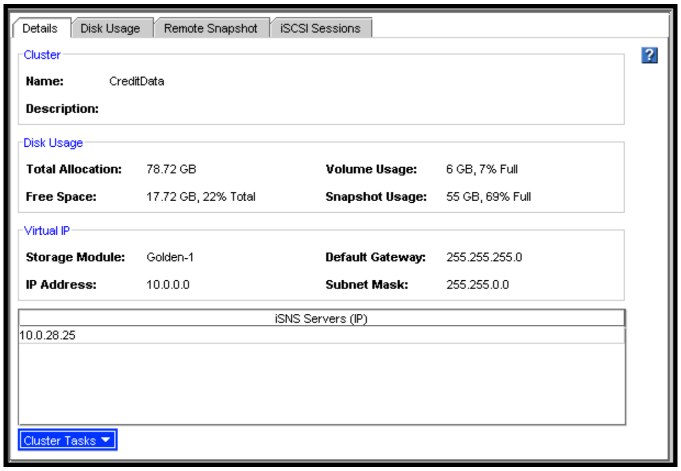

Cluster

**Name:** CreditData

**Description:**

Disk Usage

| | | | |
|---|---|---|---|
| **Total Allocation:** | 78.72 GB | **Volume Usage:** | 6 GB, 7% Full |
| **Free Space:** | 17.72 GB, 22% Total | **Snapshot Usage:** | 55 GB, 69% Full |

Virtual IP

| | | | |
|---|---|---|---|
| **Storage Module:** | Golden-1 | **Default Gateway:** | 255.255.255.0 |
| **IP Address:** | 10.0.0.0 | **Subnet Mask:** | 255.255.0.0 |

| iSNS Servers (IP) |
|---|
| 10.0.28.25 |

**Figure 149. Viewing the Cluster tab window**

3. Click Cluster Tasks and select New Volume.

4. The New Volume window opens, shown in Figure 150. See Table 41 on page 222 for detailed information about setting volume parameters See ”Requirements for Volumes” for a detailed description of each item.

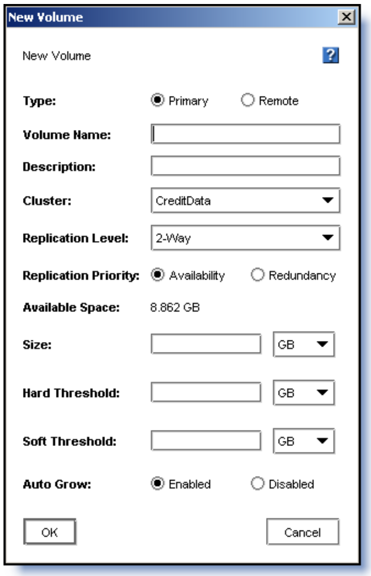

New Volume

New Volume

Type: ● Primary ○ Remote

Volume Name:

Description:

Cluster: CreditData

Replication Level: 2-Way

Replication Priority: ● Availability ○ Redundancy

Available Space: 8.862 GB

Size: GB

Hard Threshold: GB

Soft Threshold: GB

Auto Grow: ● Enabled ○ Disabled

Figure 150. Creating a New Primary Volume

5. Make sure Primary is set as the volume type.

   The window for a new primary volume is shown in Figure 150.

6. Type a name for the volume.

7. [Optional] Type a description of the volume.

8. Assign the volume to a cluster named in the drop-down list.

9. Select a replication level.
   You must purchase the Scalability Pak to use the N-way replication feature beyond the 30-day evaluation period.

10. Select a replication priority.

11. Based on the Available Space shown, type a size and select the units.

    When you enter a size, the hard and soft thresholds automatically fill in with the same values. This is the recommended configuration; size = hard threshold = soft threshold

    ***Note:*** *The system automatically factors replication levels into the settings. For example, if you create a 500 GB volume and the replication level is 2, the system automatically allocates 1000 GB for the volume.*

12. [Optional] Select the Auto Grow setting you want.

    Auto grow is enabled by default. This is the recommended configuration.

13. Click OK.

The Storage System Software creates the volume and attaches it to the cluster, shown in Figure 151.

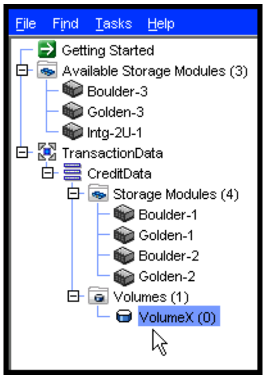

**Figure 151. New Volume in Cluster**

14. Select the new volume in the navigation window to display the information about it in the tab view, shown in Figure 152.

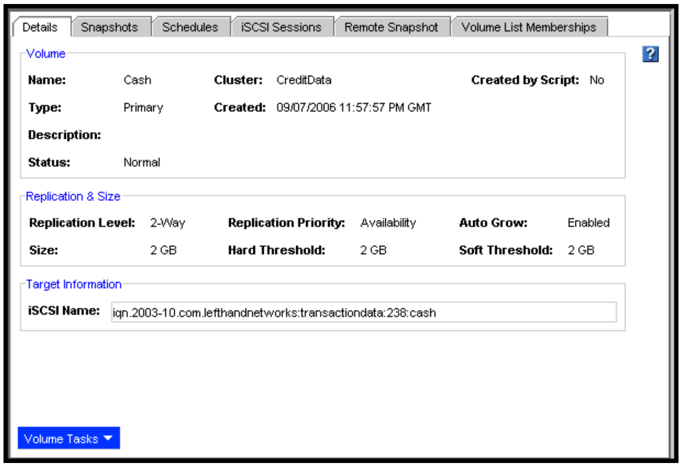

Volume

Name: Cash Cluster: CreditData Created by Script: No

Type: Primary Created: 09/07/2006 11:57:57 PM GMT

Description:

Status: Normal

Replication & Size

Replication Level: 2-Way Replication Priority: Availability Auto Grow: Enabled

Size: 2 GB Hard Threshold: 2 GB Soft Threshold: 2 GB

Target Information

iSCSI Name: iqn.2003-10.com.lefthandnetworks:transactiondata:238:cash

**Figure 152. New Volume displayed in the Tab View**

# Editing a Volume

When editing a primary volume, that is, not a remote volume, you can change the description, replication level, replication priority, size, hard and soft thresholds, the cluster that contains the volume, and whether the volume is configured for auto grow.

### Table 42. Requirements for Changing Volume Parameters

| Item | Requirements for Changing |
|---|---|
| Description | Must be from 0 to 127 characters. |
| Cluster | The target cluster must<br>• Reside in the same management group.<br>• Have sufficient unallocated space for the hard threshold and replication level of the volume being moved.<br>• Use a Virtual IP if the originating cluster has a Virtual IP<br>When moving a volume to a different cluster, that volume temporarily exists on both clusters. |
| Replication Level | The cluster must have sufficient storage modules and unallocated space to support the new replication level. |
| Replication Priority | To change the replication priority, the replication level must support the change. You can always go from Redundancy to Availability. However, you cannot go from Availability to Redundancy unless a sufficient number of storage modules are in the cluster to make the volume available. For a detailed explanation, see Table 40 on page 221.<br>For example, if you have 2-way replication with 3 storage modules in the cluster, you can change from Availability to Redundancy if all the storage modules in the cluster are available and have enough space for replicating the data. |
| Size | • To increase the size of the volume:<br>— If you have enough free space in the cluster, simply type in the new size amount<br>— If you do not have enough free space in the cluster, add a storage module to the cluster<br>• To decrease the size of the volume:<br>— If the volume has been or is mounted by any operating system, you must shrink the file system on the volume before shrinking the volume in the Console.<br>— The size entered must be greater than or equal to the hard threshold. Do not decrease the volume size to a value less than the hard threshold.<br>— You also should not decrease the size of the volume below the size needed for data currently stored on the volume.<br>**WARNING:**Decreasing the volume size or hard threshold is not recommended. If you shrink the volume in the Storage System Console before shrinking it from the client file system, your data will be corrupted or lost. |

**Table 42. Requirements for Changing Volume Parameters (Cont'd)**

| Item | Requirements for Changing |
|---|---|
| Hard Threshold | • There must be sufficient unallocated space in the cluster.<br>• Increase the hard threshold to turn off an alert generated when the threshold is exceeded. The hard threshold must be equal to or less than the size of the volume and there must be sufficient space on the cluster.<br>• To decrease the hard threshold, first decrease the size of the volume and decrease the hard threshold to the same value as the size.<br>**WARNING:**Decreasing the volume size or hard threshold is not recommended. If you shrink the volume in the Storage System Console before shrinking it from the client file system, your data will be corrupted or lost. |
| Soft Threshold | • The soft threshold must be equal to or less than the hard threshold.<br>• Decrease the soft threshold to a value less than or equal to the hard threshold. |
| Auto Grow | • Change the auto grow setting as desired. Enable or disable auto grow. |

***Warning:*** *Decreasing the volume size or hard threshold is not recommended. If you shrink the volume in the Console before shrinking it from the client file system, your data will be corrupted or lost.*

## Getting There

1. In the navigation window, select the volume you want to edit
2. Click Volume Tasks and select Edit Volume.

   The Edit Volume window opens, shown in Figure 153. See Table 42 for detailed information about making changes to the volume parameters.

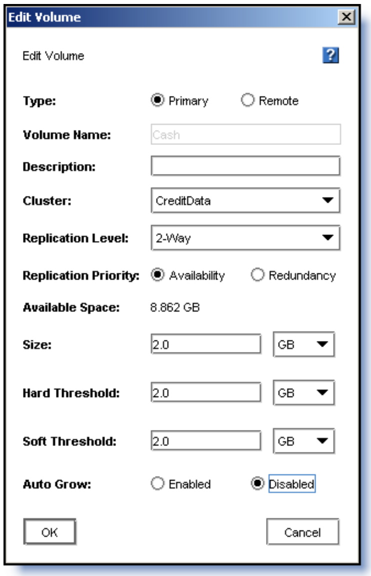

Edit Volume

Edit Volume

Type: ● Primary ○ Remote

Volume Name: Cash

Description:

Cluster: CreditData

Replication Level: 2-Way

Replication Priority: ● Availability ○ Redundancy

Available Space: 8.862 GB

Size: 2.0 GB

Hard Threshold: 2.0 GB

Soft Threshold: 2.0 GB

Auto Grow: ○ Enabled ● Disabled

**Figure 153. Editing a Volume**

## Changing the Volume Description

1. In the Description field, change the description.
2. Click OK when you are finished.

## Changing the Cluster

**Prerequisite**

- Log off the volume from the iSCSI client

1. In the navigation window, select the volume you want to move.
2. Click the Volume Tasks drop-down list and select Edit Volume.

3. In the Cluster drop-down list on the Edit Volume window, select a different cluster.

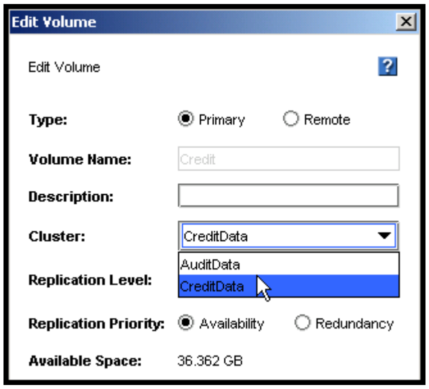

Edit Volume

Edit Volume

Type: ● Primary ○ Remote

Volume Name: Credit

Description:

Cluster: CreditData

AuditData<br>CreditData

Replication Level:

Replication Priority: ● Availability ○ Redundancy

Available Space: 36.362 GB

Figure 154. Moving a Volume to a Different Cluster

4. Click OK.

   The volume resides on both clusters until all of the data are moved to the new cluster.

5. Add the new Virtual IP.
6. Discover and mount the volumes while the migration is under way.

   Even if using MPIO for DSM, log off the volumes from the server, start the migration, add the IPs of the storage modules in the other cluster, discover and mount volumes.

## Changing the Replication Level

1. In the Replication Level drop-down list, select the level of replication you want.
2. Click OK when you are finished.

## Changing the Replication Priority

1. Select the replication priority you want.
2. Click OK when you are finished.

## Changing the Size

1. In the Size field, change the number and change the units if necessary.
2. Click OK when you are finished.

***Warning:*** *Decreasing the volume size or hard threshold is not recommended. If you shrink the volume in the Console before shrinking it from the client file system, your data will be corrupted or lost.*

## Changing the Volume Size on the Client

When you increase the size of the volume on the SAN, you must next increase the corresponding volume or LUN on the client side.

### Increasing the Volume Size in Microsoft Windows

Once you have increased the volume size on the SAN, you must next expand the Windows partition to use the full space available on the disk.

Windows Logical Disk Manager, the default disk management program that is included in any Windows installation, uses a tool called Diskpart.exe to grow volumes from within Windows. Diskpart.exe is an interactive command line executable which allows administrators to select and manipulate disks and partitions. This executable and its corresponding documentation can be downloaded from Microsoft if necessary. Follow the steps below to extend the volume you just increased in the SAN.

1. Launch Windows Logical Disk Manager to rescan the disk and present the new volume size.
2. Open a Windows command line and run diskpart.exe.

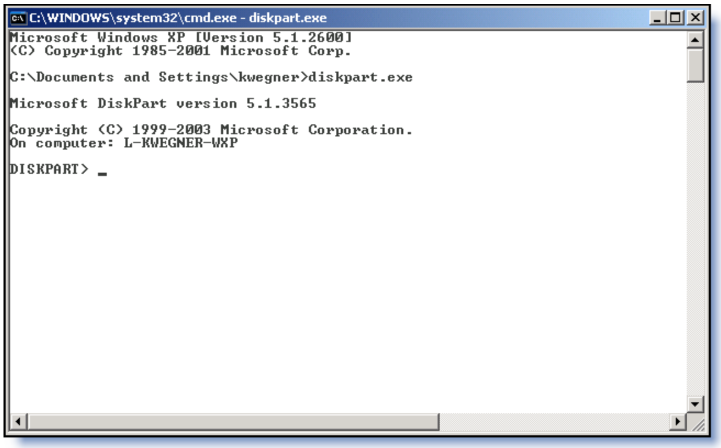

```
Microsoft Windows XP [Version 5.1.2600]
(C) Copyright 1985-2001 Microsoft Corp.

C:\Documents and Settings\kwegner>diskpart.exe

Microsoft DiskPart version 5.1.3565

Copyright (C) 1999-2003 Microsoft Corporation.
On computer: L-KWEGNER-WXP

DISKPART> _
```

**Figure 155. Running diskpart.exe**

3. List the volumes that appear to this host by typing the command "list volume".

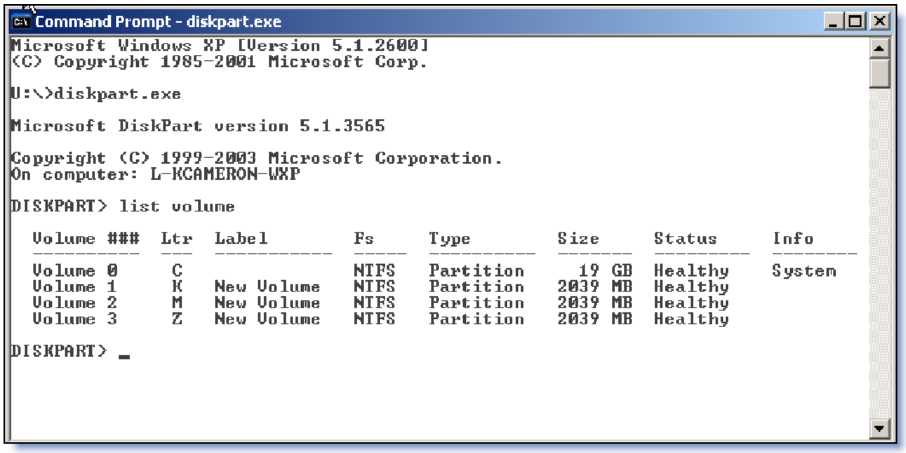

```
Microsoft Windows XP [Version 5.1.2600]
(C) Copyright 1985-2001 Microsoft Corp.

U:\>diskpart.exe

Microsoft DiskPart version 5.1.3565

Copyright (C) 1999-2003 Microsoft Corporation.
On computer: L-KCAMERON-WXP

DISKPART> list volume

  Volume ###  Ltr  Label        Fs     Type        Size     Status     Info
  ----------  ---  -----------  -----  ----------  -------  ---------  --------
  Volume 0     C                NTFS   Partition     19 GB  Healthy    System
  Volume 1     K   New Volume   NTFS   Partition   2039 MB  Healthy
  Volume 2     M   New Volume   NTFS   Partition   2039 MB  Healthy
  Volume 3     Z   New Volume   NTFS   Partition   2039 MB  Healthy

DISKPART> _
```

Figure 156. Diskpart “list volumes”

4. Select the volume to extend by typing "select volume #" (where # is the volume's corresponding number).

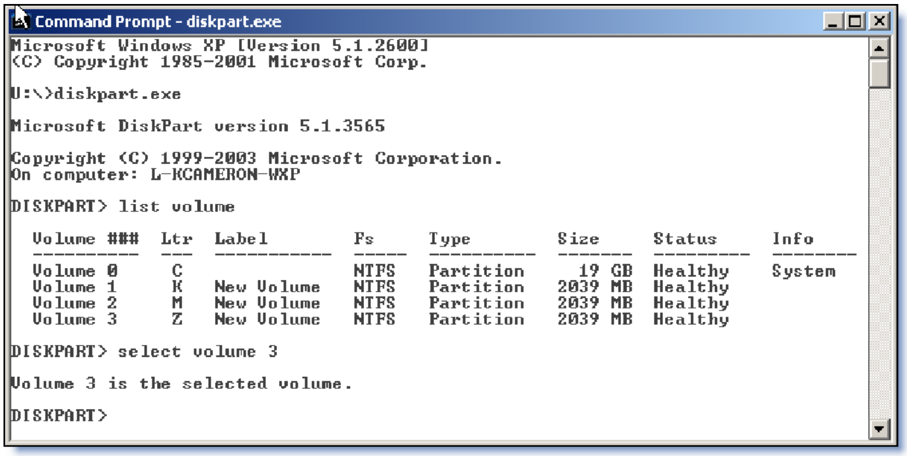

```
Microsoft Windows XP [Version 5.1.2600]
(C) Copyright 1985-2001 Microsoft Corp.

U:\>diskpart.exe

Microsoft DiskPart version 5.1.3565

Copyright (C) 1999-2003 Microsoft Corporation.
On computer: L-KCAMERON-WXP

DISKPART> list volume

  Volume ###  Ltr  Label        Fs     Type        Size     Status     Info
  ----------  ---  -----------  -----  ----------  -------  ---------  --------
  Volume 0     C                NTFS   Partition     19 GB  Healthy    System
  Volume 1     K   New Volume   NTFS   Partition   2039 MB  Healthy
  Volume 2     M   New Volume   NTFS   Partition   2039 MB  Healthy
  Volume 3     Z   New Volume   NTFS   Partition   2039 MB  Healthy

DISKPART> select volume 3

Volume 3 is the selected volume.

DISKPART>
```

Figure 157. Diskpart “select volume”

5. Type "extend" to grow the volume to the size of the full disk that has been expanded.

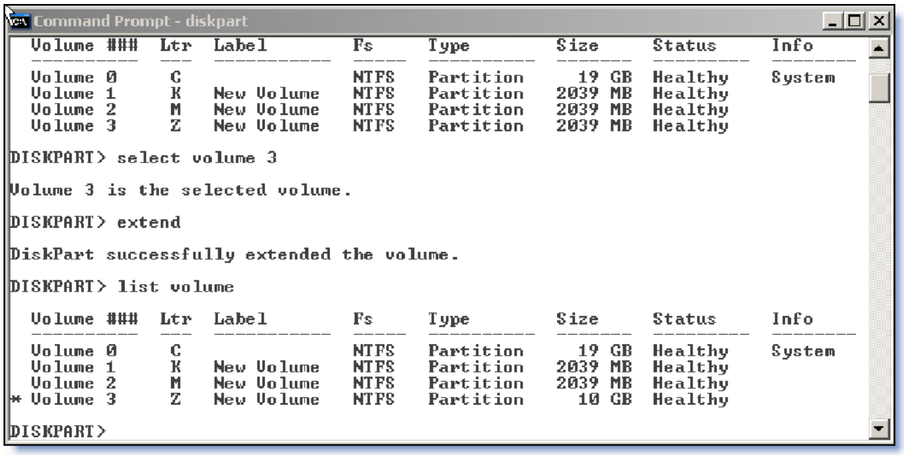

```
  Volume ###  Ltr  Label        Fs     Type        Size     Status     Info
  ----------  ---  -----------  -----  ----------  -------  ---------  --------
  Volume 0     C                NTFS   Partition     19 GB  Healthy    System
  Volume 1     K   New Volume   NTFS   Partition   2039 MB  Healthy
  Volume 2     M   New Volume   NTFS   Partition   2039 MB  Healthy
  Volume 3     Z   New Volume   NTFS   Partition   2039 MB  Healthy

DISKPART> select volume 3

Volume 3 is the selected volume.

DISKPART> extend

DiskPart successfully extended the volume.

DISKPART> list volume

  Volume ###  Ltr  Label        Fs     Type        Size     Status     Info
  ----------  ---  -----------  -----  ----------  -------  ---------  --------
  Volume 0     C                NTFS   Partition     19 GB  Healthy    System
  Volume 1     K   New Volume   NTFS   Partition   2039 MB  Healthy
  Volume 2     M   New Volume   NTFS   Partition   2039 MB  Healthy
* Volume 3     Z   New Volume   NTFS   Partition     10 GB  Healthy

DISKPART>
```

Notice the asterisk by the volume and the new size of the volume. The disk has been extended and is now ready for use. All of the above operations are performed while the volumes are on-line and available to users.

**Other Disk Management Tools**

Some environments use alternative tools, such as Dell Array Manager and VERITAS Volume Manager. Both of these disk management tools use a utility called Extpart.exe instead of Diskpart.exe. Extpart.exe commands are similar o those of Diskpart.exe. The only major difference is that instead of selecting the volume number, as in Diskpart.exe, you select the drive letter instead. Extpart.exe and corresponding documentation can be downloaded from www.dell.com.

## Changing the Hard Threshold

1. In the Hard Threshold field, change the number and change the units if necessary.
2. Click OK when you are finished.

***Warning:*** *Decreasing the volume size or hard threshold is not recommended. If you shrink the volume in the Console before shrinking it from the client file system, your data will be corrupted or lost.*

## Changing the Soft Threshold

1. In the Soft Threshold field, change the number and change the units if necessary.
2. Click OK when you are finished.

# Fixing an Unavailable Priority Volume

If a storage module becomes unavailable and needs to be repaired or replaced, and a replicated volume that is configured for redundancy becomes unavailable to clients, the following procedure allows you to safely return the volume to fully operational status.

1. Stop any clients from accessing the volume.
2. Select the volume in the navigation window.
3. Right-click and select Edit Volume.
4. Change the data priority from data redundancy to data availability.
5. Remove the storage module from the cluster.

   Repair or replace the storage module.
6. [Optional] Add the new or repaired storage module to the cluster.
7. Wait for the restripe of the volume to finish.
8. Edit the volume.
9. Change the data priority from data availability to data redundancy.
10. Restore the clients' access to the volume.

# Deleting a Volume

Delete a volume to remove that volume's data from the storage module and make that space available. When deleting volumes, you must delete all snapshots of that volume before you delete the volume itself.

***Warning:*** *Deleting a volume permanently removes that volume's data from the storage module.*

**Prerequisites**

- Delete all snapshots of the volume that you want to delete.
- Stop applications from accessing the volume.
- Disconnect the iSCSI clients from the volume..

## To Delete the Volume

1. In the navigation window, select the volume you want to delete.

   The Volume tab window opens.

2. Click Volume Tasks and select Delete Volume.

   A confirmation window opens.

3. Click OK.

   The volume is removed from the cluster.

# 13 Working with Snapshots

## Snapshots Overview

Snapshots provide  a point-in-time copy of a volume for use with backup and other applications.

**Snapshots and Backups**

Backups are typically stored on different physical devices or tapes. Snapshots are stored on the same cluster as the volume. Therefore, snapshots protect against data corruption or deletion, but not device or storage media failure. Use snapshots along with backups to improve your overall data backup strategy.

**Prerequisites**

Before you create a snapshot, you must have created a management group, a cluster, and a volume to receive it.

For information, see

- "Creating a Management Group" on page 171
- "Creating a Cluster" on page 202
- "Creating a Volume" on page 225

### Topics Covered in This Chapter

- Single snapshots and scheduled snapshots
- Managing capacity using volume and snapshot thresholds
- Creating scheduled snapshots

## Using Snapshots

You create snapshots from a volume on the cluster. At any time you can roll back to a specific snapshot. When you do roll back, you must delete all the snapshots created after that snapshot. Also, using an iSCSI initiator, you can mount a snapshot to a different server and recover data from the snapshot to that server.

Snapshots can be used for these cases:

- Source volumes for data mining and other data use

- Source volumes for creating backups
- Data or file system preservation before upgrading software
- Protection against data or file system corruption
- File-level restore without tape or backup software

## Single Snapshots vs. Scheduled Snapshots

Some snapshot scenarios call for creating a single snapshot and then deleting it when it is no longer needed. Other scenarios call for creating a series of snapshots up to a specified number or for a specified time period, after which the earliest snapshot is deleted when the new one is created (scheduled snapshots).

For example, you plan to keep a series of daily snapshots, up to four. After creating the fifth snapshot, the earliest snapshot is deleted, thereby keeping the number of snapshots on the cluster at four.

Scheduled snapshots are an add-on feature. You must purchase the Configurable Snapshot Pak to use snapshot schedules beyond the 30-day evaluation period.

# Guide for Snapshots

Review in “Planning Volumes” on page 215 to ensure that you configure snapshots correctly. When creating a snapshot, you define the following parameters.

### Table 43. Snapshot Parameters

| Snapshot Parameter | What it means |
|---|---|
| Snapshot Name | The name of the snapshot that is displayed in the Console. A snapshot name must be from 1 to 127 characters and is case sensitive. |
| Description | (Optional) A description of the snapshot. |

### Table 43. Snapshot Parameters (Cont'd)

| Snapshot Parameter | What it means |
|---|---|
| Hard Threshold | This becomes the hard threshold of the writable volume and defines the amount of space allocated for changes to the original volume. When reached, the hard threshold triggers an alert and data can no longer be written to the volume. The hard threshold must be less than, or equal to, the volume size, and cannot exceed available space in the cluster.<br>**Warning**: If the hard threshold is less than the volume size, and the hard threshold gets exceeded, then data writes to the volume may stop until the hard threshold is increased by auto grow (recommended) or manually increased in the Console. |
| Soft Threshold | The amount of space actually used on the writable volume that triggers a warning alert. This alert notifies the storage administrator that the writable volume is approaching the hard threshold. The soft threshold must be less than, or equal to, the hard threshold. |

## Best Practice

Create snapshots with the following parameters

- Minimum Hard Threshold = 512 MB
- Minimum Soft Threshold = 256 MB
- Minimum Gap between Hard and Soft Threshold = 256 MB
- Auto grow enabled on volume

# Managing Capacity Using Volume and Snapshot Thresholds

**How Snapshots are Created**

When you create a snapshot of a volume, the original volume is actually saved as the snapshot, and a new volume (the "writable" volume) with the original name is created to record any changes made to the volume's data after the snapshot was created. Subsequent snapshots record only changes made to the volume since the previous snapshot.

### Hard Thresholds and Snapshots

One implication of the relationship between volumes and snapshots is that the space used by the writable volume can become very small when it records only the changes that have occurred since the last snapshot was taken. This means that less space—or a smaller hard threshold—may be required for the writable volume.

You can save space on your cluster of storage modules by estimating the size required for the changes in data between snapshots and decreasing the hard threshold of each snapshot accordingly. This planning is particularly important if you plan to use a series of snapshots to protect against data corruption. For more information about hard thresholds and volumes, see "Planning Hard Thresholds" on page 218.

### Deleting Snapshots

One important factor in planning capacity is the fact that when a snapshot is deleted, the snapshot's hard and soft thresholds are added to the snapshot or volume directly above it (the next newer snapshot). Hard and soft thresholds of the volume or snapshot directly above the deleted snapshot will increase by the hard and soft thresholds of the deleted snapshot, up to the size of the volume. Adding hard and soft thresholds into the next volume or snapshot insures that all changes to data are accounted for and saved. Therefore, if you plan a protocol where you routinely delete snapshots, there are two considerations you should take into account.

- You must calculate the effect on the volume size of adding the hard thresholds back into the volume.
- You may want to plan for deleting snapshots to take place at off-peak hours due to the impact on performance of migrating data during the deletion.

For a detailed explanation of disk capacity allocation in a cluster and its relationship to disk or volume size in a file system, see "Measuring Disk Capacity and Volume Size" on page 217.

## Full Provisioning Method for Planning Capacity

Make the snapshot hard threshold equal to the volume size, and the soft threshold equal to the hard threshold.

## Thin Provisioning Method for Planning Capacity

Make the hard threshold less than the volume size, and the soft threshold less than the hard threshold. Then, increase the volume size, hard threshold, and soft threshold as necessary to manage capacity growth.

Table 44 , Table 45 , and Table 46  illustrate how storage can be effectively managed by setting the hard threshold below the original volume size in a series of snapshots.

### Table 44. Space Used by Snapshots when Hard Threshold is set to the Original Volume Size

| Day | Volume/Snapshot | Data Stored or Changed | Snapshot Size w/ No Threshold Change | Total Space Used on Cluster |
|---|---|---|---|---|
| Mon. | Original Volume = 50 GB | N/A | 50 GB | 50 GB |
| Tue. | Snapshot 1 | < 15 GB | 50 GB | 100 GB |
| Wed. | Snapshot 2 | < 10 GB | 50 GB | 150 GB |
| Thur | Snapshot 3 | < 8 GB | 50 GB | 200 GB |

### Table 45. Space used by Snapshots when Hard Threshold is Reduced. Note the Dramatic Savings in Storage Space

| Day | Volume/Snapshot | Data Stored or Changed | Snapshot Size w/ Hard Threshold Reduced | Total Space Used on Cluster |
|---|---|---|---|---|
| Mon. | Original Volume = 50 GB | N/A | 50 GB | 50 GB |
| Tue. | Snapshot 1 | < 15 GB | 15 GB | 65 GB |
| Wed. | Snapshot 2 | < 10 GB | 15 GB | 80 GB |
| Thur | Snapshot 3 | < 8 GB | 15 GB | 95 GB |

### Table 46. Common Data Change Rates

| Data Type | Average Change Rate per Day |
|---|---|
| File share | 1 – 3% |
| Email/Exchange | 10 – 20% |
| Database | 10%+ [1] |

1. Depends on transaction rate.

**Note:** *Deleting files on a file system does not free up space on the volume. For file-level capacity management, use application or file system-level tools.*

# Planning Snapshots

When planning to use snapshots, consider their purpose and size.

**Note:** *When considering the size of snapshots in the cluster, remember that the replication level of the volume is duplicated in the snapshot.*

## Source Volumes for Data Mining or Tape Backups

**Best Practice**

Plan to use a single snapshot and delete it when you are finished. Consider the following questions in your planning.

- Is space available on the cluster to create the snapshot?
- Is space available in the cluster to accommodate the increase in the volume's hard threshold when the snapshot is deleted? Remember that the hard threshold can never exceed the volume size.

## Data Preservation Before Upgrading Software

**Best Practice**

Plan to use a single snapshot and delete it when you are finished. Consider the following questions in your planning.

- Is space available on the cluster to create the snapshot?
- Is space available in the cluster to accommodate the increase in the volume's hard threshold when the snapshot is deleted? Remember that the hard threshold can never exceed the volume size.

## Protection Against Data Corruption

**Best Practice**

Plan to use a series of snapshots, deleting the oldest on a scheduled basis. Consider the following questions in your planning.

- What is the minimum size you can set for the hard threshold that will accommodate the changes likely to occur between snapshots?
- Is space available on the cluster to create the snapshots?
- Is space available in the cluster to accommodate the increase in the volume's hard threshold when the snapshot is delete.

# Creating a Snapshot

Create a snapshot to preserve a version of a volume at a specific point in time.

1. Log into the management group that contains the volume for which you want to create a new snapshot.
2. Right-click on the volume for which you want to create a snapshot.

3. Select New Snapshot.

   The New Snapshot window opens, shown in Figure 158.

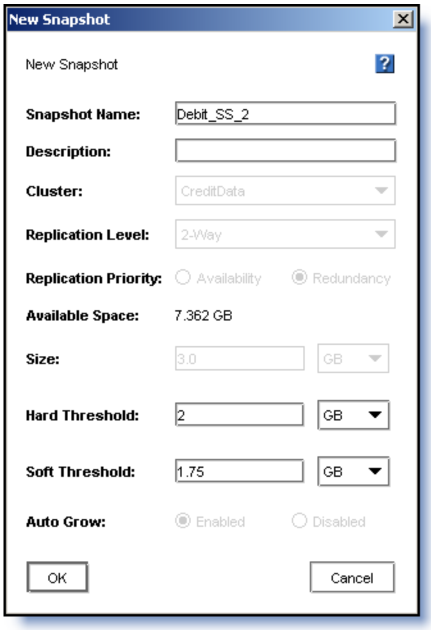

Snapshot Name: Debit_SS_2

Description:

Cluster: CreditData

Replication Level: 2-Way

Replication Priority: ○ Availability ● Redundancy

Available Space: 7.362 GB

Size: 3.0 GB

Hard Threshold: 2 GB

Soft Threshold: 1.75 GB

Auto Grow: ● Enabled ○ Disabled

**Figure 158. Creating a New Snapshot**

4. Type a name for the snapshot.

   Names are case sensitive. They cannot be changed after the snapshot is created.

5. (Optional) Type in a description of the snapshot.
6. (Optional) Enter the hard and soft thresholds for the snapshot.

# Best Practice

Create snapshots with the following parameters

- Minimum Hard Threshold = 512 MB
- Minimum Soft Threshold = 256 MB
- Minimum Gap between Hard and Soft Threshold = 256 MB
- Auto grow enabled on volume

**Note:** *The hard threshold of the snapshot becomes the hard threshold of the writable volume and defines the amount of space allocated for changes to the original volume.*

**Note:** *Setting the hard threshold smaller than the size of the original volume allows you to create snapshots that require less space on the cluster.*

7. Click OK when you are finished.

   The Snapshots tab refreshes with the new snapshot listed, as shown in Figure 159. The new snapshot also displays in the navigation window, as shown in Figure 160.

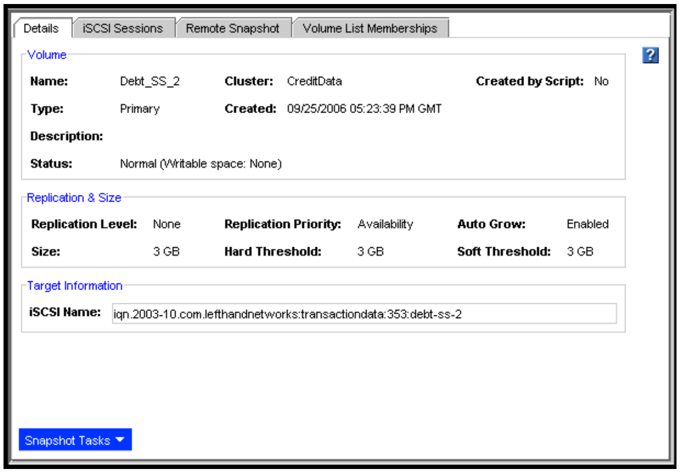

Volume

**Name:** Debt_SS_2 **Cluster:** CreditData **Created by Script:** No

**Type:** Primary **Created:** 09/25/2006 05:23:39 PM GMT

**Description:**

**Status:** Normal (Writable space: None)

Replication & Size

**Replication Level:** None **Replication Priority:** Availability **Auto Grow:** Enabled

**Size:** 3 GB **Hard Threshold:** 3 GB **Soft Threshold:** 3 GB

Target Information

**iSCSI Name:** iqn.2003-10.com.lefthandnetworks:transactiondata:353:debt-ss-2

Figure 159. Viewing the New Snapshot in the Tab Window

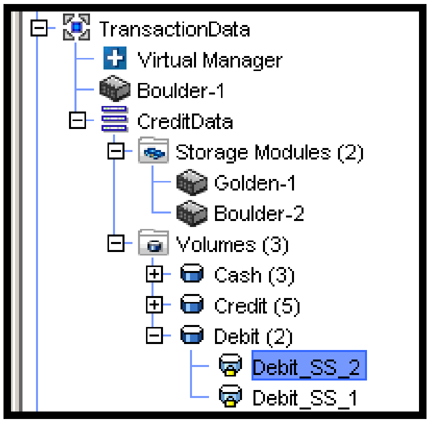

Figure 160. Viewing the New Snapshot in the Navigation Window

***Note:*** *In the navigation window, snapshots are listed below the volume in descending date order, from newest to oldest.*

# Mounting or Accessing a Snapshot

A snapshot is a point-in-time picture of a volume. To mount the snapshot for backing up or making the data available for other uses such as data recovery, data mining or testing, you can configure the snapshot as a read/write volume and connect to it with an iSCSI initiator.

# Snapshot Writable Space

When you configure a snapshot as read/write, additional space is created in the cluster for use by applications and operating systems that need to write to the snapshot when they access it. For example, MS Windows performs a write when the snapshot is mounted via iSCSI. Microsoft Volume Shadow Service (VSS) and other backup programs write to the snapshot when backing it up. You can see how much writable space is being used for a snapshot on the Disk Usage tab in the Cluster tab window, as shown in Figure 161.

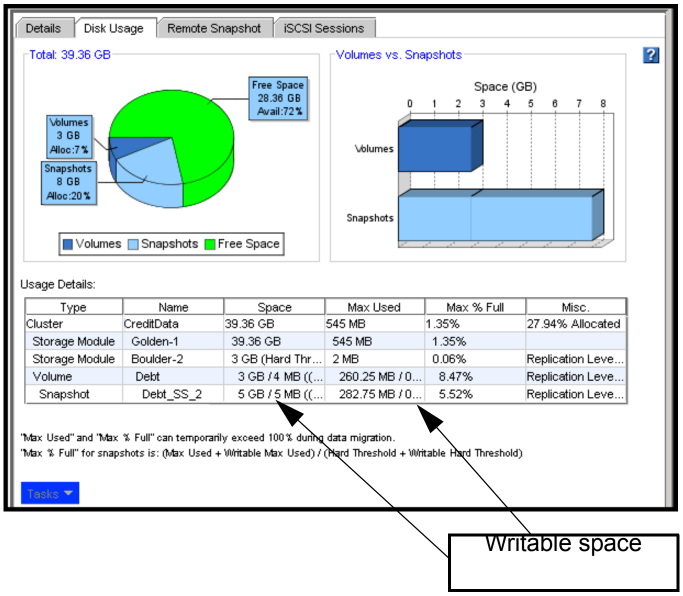

Usage Details:

| Type | Name | Space | Max Used | Max % Full | Misc. |
|---|---|---|---|---|---|
| Cluster | CreditData | 39.36 GB | 545 MB | 1.35% | 27.94% Allocated |
| Storage Module | Golden-1 | 39.36 GB | 545 MB | 1.35% | |
| Storage Module | Boulder-2 | 3 GB (Hard Thr... | 2 MB | 0.06% | Replication Leve... |
| Volume | Debt | 3 GB / 4 MB ((... | 260.25 MB / 0... | 8.47% | Replication Leve... |
| Snapshot | Debt_SS_2 | 5 GB / 5 MB ((... | 282.75 MB / 0... | 5.52% | Replication Leve... |

"Max Used" and "Max % Full" can temporarily exceed 100% during data migration.
"Max % Full" for snapshots is: (Max Used + Writable Max Used) / (Hard Threshold + Writable Hard Threshold)

**Figure 161. Viewing the Writable Space used for a Snapshot**

The additional writable space is deleted when the snapshot is deleted. If you need to free up the extra space before the snapshot is deleted, you can do so manually in the Console or through your snapshot scripts. The next time an application or operating system accesses the snapshot, the writable space will be recreated.

## Deleting Snapshot Writable Space

**Prerequisite**

- Stop any applications from accessing the volume.

## Deleting the Writable Space

1. In the navigation window, select the snapshot for which you want to delete the writable space.
2. Right-click and select Delete Writable Space.

   A warning message opens.
3. Click OK to confirm the delete.

# Editing a Snapshot

You can edit the description of a snapshot. You can also change the hard and soft thresholds. See "Creating a Snapshot" on page 242.

1. Log into the management group that contains the snapshot that you want to edit.
2. In the navigation window, select the snapshot you want to edit.
3. Click Snapshot Tasks and select Edit Snapshot.

   The Edit Snapshot window opens, shown in Figure 162.

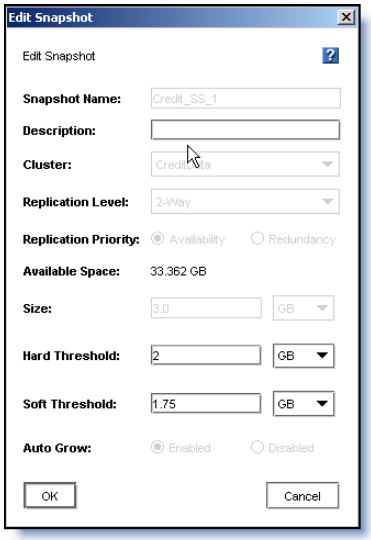

Edit Snapshot

Edit Snapshot

Snapshot Name: Credit_SS_1

Description:

Cluster: CreditData

Replication Level: 2-Way

Replication Priority: ● Availability ○ Redundancy

Available Space: 33.362 GB

Size: 3.0 GB

Hard Threshold: 2 GB

Soft Threshold: 1.75 GB

Auto Grow: ● Enabled ○ Disabled

## Figure 162. Editing a Snapshot

4. Navigate to the field you want to change and change the information.

## Table 47. Data Requirements for Editing a Snapshot

| Item | Requirements for Changing |
|---|---|
| Description | Must be from 0 to 127 characters. |
| Hard Threshold | Hard threshold size must be equal to or less than the size of the volume and available storage in the cluster. You cannot decrease the hard threshold. |
| Soft Threshold | Soft threshold size must be equal to or less than the hard threshold size. |

5. Click OK when you are finished.

   The snapshot tab window opens.

# Manually Copying a Volume from a Snapshot

When you have mounted the snapshot on a host, you can do the following:

- Recover individual files or folders and restore to an alternate location
- Copy the snapshot data to a read/write volume
- Back up the data

To mount the snapshot on a host

1. Create an authentication group for the client that you want to mount the snapshot on. See "Creating an Authentication Group" on page 272.
2. Create a volume list for the snapshot, and configure the snapshot for read/write access. See "Creating a Volume List" on page 278.
3. Configure client access to the snapshot volume.

Now you can access the snapshot in these ways:

- As a source volume for data mining and other data use
- As a source volume for creating backups
- For data and file system preservation before upgrading software
- For protection against data and file system corruption
- For file-level restore without tape or backup software

# Creating Scheduled Snapshots

You can schedule recurring snapshots. Recurring snapshots can be scheduled in a variety of frequencies and with a variety of retention policies. You can schedule a snapshot every 30 minutes or more, and retain up to 50 snapshots.

**Note:** *Scripting snapshots can also take place on the client side. Scripted snapshots offer greater flexibility for quiescing hosts while taking snapshots, and for automating tasks associated with volumes and their snapshots.*

Scripting of snapshots is an add-on feature. You must purchase the Configurable Snapshot Pak to use snapshot scripting beyond the 30-day evaluation period.

## Requirements for Scheduling Snapshots

Scheduled snapshots require particular attention to capacity management. Additionally, you must ensure that the time settings on thestorage modules running managers and the time setting of the management group are synchronized.

## Best Practice

Schedule snapshots during off-peak hours. If setting scheduled snapshots for multiple volumes, stagger the schedules with at least an hour between start times for best results.

***Note:*** *Use NTP to ensure that all the storage modules in the management group have synchronized time settings.*

### Table 48. Requirements for Scheduling Snapshots

| Requirement | What it means |
|---|---|
| Plan for capacity management | Scheduling snapshots should be planned with careful consideration for capacity management as described in "Managing Capacity Using Volume and Snapshot Thresholds" on page 239.<br>Pay attention to how you want to retain snapshots and the capacity in the cluster. If you want to retain <n> snapshots, the cluster should have space for <n+1>.<br>It is possible for the new snapshot and the one to be deleted to coexist in the cluster for some period of time.<br>If there is not sufficient room in the cluster for both snapshots, the scheduled snapshot will not be created, and the schedule will not continue until an existing snapshot is deleted. |
| Synchronize storage module times with management group time | The time setting on the storage modules running managers and the time setting of the management group must be synchronized. If they are not synchronized, then the scheduled snapshot might run incorrectly.<br>Be sure to configure the correct time on the storage modules and then reset the management group time. See "Setting the Date and Time" on page 121. Also, see "Resetting the Management Group Time" on page 179. |
| Plan scheduling and retention policies | The minimum recurrence you can set for snapshots is 30 minutes.<br>The maximum number of snapshots (scheduled and manual combined) you can retain is 50 snapshots per volume.<br>The **recommended** number of snapshots of all types per management group is no more than, and preferably less than, 200 total. |

# Creating Scheduled Snapshots

You can create one or more scheduled snapshots for a volume. For example, one schedule could be for daily snapshots intended for backup and recovery. A second schedule could be for weekly snapshots used for data mining.

1. In the navigation window, select the volume for which you want to schedule snapshots.

   The Volume tab window opens.

2. Click the Schedules tab to bring it to the front.
3. Click Schedule Tasks and select New Scheduled Snapshot.

   The New Scheduled Snapshot window opens, shown in Figure 163.

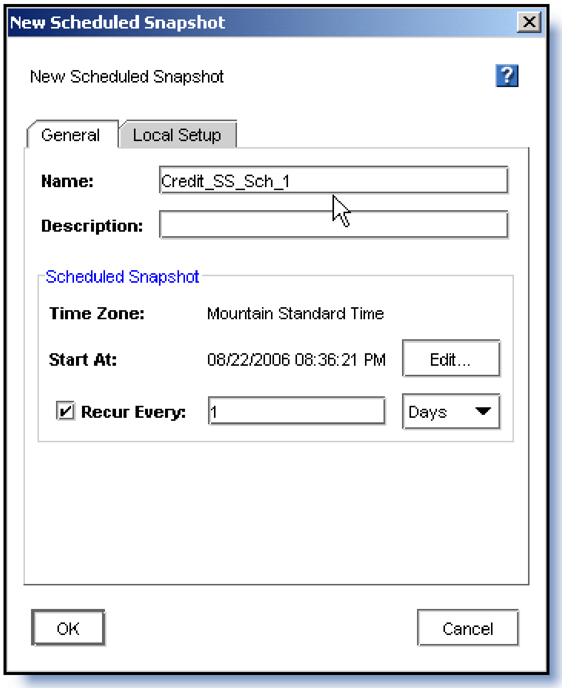

Name: Credit_SS_Sch_1

Description:

Time Zone: Mountain Standard Time

Start At: 08/22/2006 08:36:21 PM

☑ Recur Every: 1 Days

### Figure 163. Creating a Scheduled Snapshot, General Setup

4. Type a name for the snapshots.

   The name will be used with sequential numbering. For example, if the snapshot name is Backup, the list of scheduled snapshots will be named Backup.1, Backup.2, Backup.3.

5. (Optional) Enter a snapshot description.
6. Enter a start date and time.

   The date and time must be valid, but they can occur in the past.

7. Select a recurrence schedule.

8. Click the Local Setup tab to bring it to the front, as shown in Figure 164.

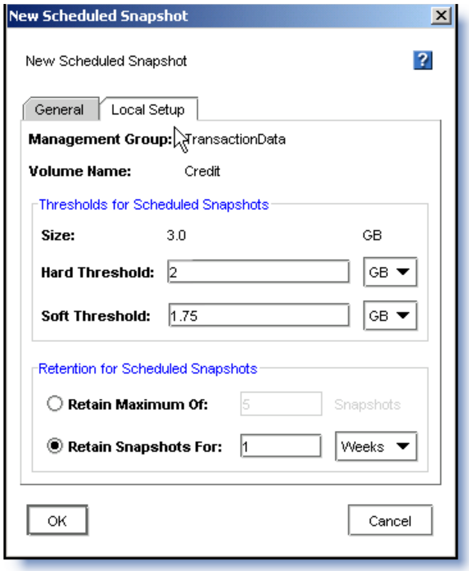

Management Group: TransactionData

Volume Name: Credit

Size: 3.0 GB

Hard Threshold: 2 GB

Soft Threshold: 1.75 GB

○ Retain Maximum Of: 5 Snapshots

● Retain Snapshots For: 1 Weeks

**Figure 164. Creating a Scheduled Snapshot, Local Setup**

9. (Optional) Change the hard and soft thresholds for the snapshots.

***Note:*** *Setting the hard threshold smaller than the size of the original volume allows you to create snapshots that require less space on the cluster. See "Managing Capacity Using Volume and Snapshot Thresholds" on page 239.*

10. Set a retention schedule.

    The retention schedule can be for specified number of snapshots or for a designated period of time. You can retain up to 50 snapshots.

11. Click OK.

    The New Scheduled Snapshot window closes and the new scheduled snapshot itself appears in the tab view, shown in Figure 165.

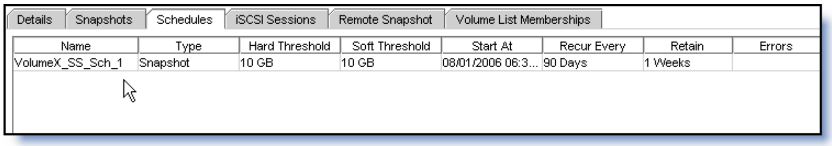

| Name | Type | Hard Threshold | Soft Threshold | Start At | Recur Every | Retain | Errors |
|---|---|---|---|---|---|---|---|
| VolumeX_SS_Sch_1 | Snapshot | 10 GB | 10 GB | 08/01/2006 06:3... | 90 Days | 1 Weeks | |

**Figure 165. List of Scheduled Snapshots**

# Editing Scheduled Snapshots

You can edit everything in the scheduled snapshot window except the name.

1. In the navigation window, select the volume for which you want to edit the scheduled snapshot.

2. In the tab window, click the Schedules tab to bring it to the front.
3. Select the schedule you want to edit.
4. Click Schedule Tasks and select Edit Schedule.

   The Edit Scheduled Snapshot window opens, shown in Figure 166.

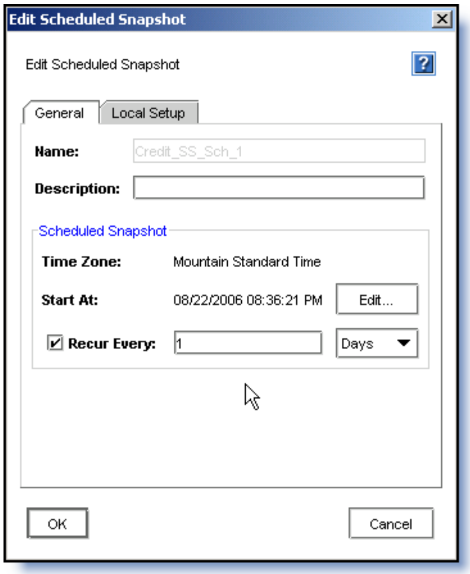

Name: Credit_SS_Sch_1

Description:

Time Zone: Mountain Standard Time

Start At: 08/22/2006 08:36:21 PM

☑ Recur Every: 1 Days

Figure 166. Editing a Scheduled Snapshot

5. Change the desired information.
6. Click OK.

**Note:** *If you change the hard threshold, be sure to review the information about snapshot thresholds and their effect on volume thresholds in "Managing Capacity Using Volume and Snapshot Thresholds" on page 239.*

## Deleting Scheduled Snapshots

1. In the navigation window, select the volume for which you want to delete the scheduled snapshot.
2. Click the Schedules tab to bring it to the front.
3. Select the schedule you want to delete.
4. Click Schedule Tasks and select Delete Schedule.
5. To confirm the deletion, click OK.

   The tab view refreshes without the deleted scheduled snapshot.

# Scripting Snapshots

Application-based scripting is available for taking snapshots. Using application-based scripts allows automatic snapshots of a volume. For detailed information, see "Working with Scripting" on page 257.

The Customer Resource Center has samples of snapshot scripts available for downloading.

# Rolling Back a Volume to a Snapshot

Rolling back a volume to a snapshot replaces the original volume with a read/write copy of the selected snapshot. The new volume has a different name than the original and the original volume is deleted.

**Prerequisites**

- Stop applications from accessing the volume.
- Disconnect iSCSI initiator(s) from the volume.

## Requirements for Rolling Back a Volume

Many of the parameters for the new volume must be configured as if you had created this volume for the first time.

**Warning:** *After rolling back a volume to a snapshot, you lose all data stored after the rolled back snapshot. Consider doing a Remote Copy before the roll back. You may have to delete snapshots that were made later than the one you intend to roll back.*

| Parameter | Requirements for Changing |
|---|---|
| New Volume Name | You must choose a new name for the volume. The name must be from 1 to 127 characters. Names are case sensitive. |
| New Hard Threshold | Hard threshold size must be equal to or less than the size of the volume. See "Managing Capacity Using Volume and Snapshot Thresholds" on page 239. |
| New Soft Threshold | Soft threshold size must be equal to or less than the hard threshold size. |
| Authentication Groups | You must include the new volume in a volume list. See "Volume Lists Overview" on page 277. |
| Hosts | You must reconfigure hosts to connect to the new volume. |

**Prerequisites**

- Stop applications from accessing the volume.
- Delete all snapshots that are newer than the snapshot you are rolling back. If you need to preserve the interim snapshots, use Remote Copy to create a copy of the snapshots before deleting them.
- If you need to preserve the original volume, use Remote Copy to first copy the snapshot to a new volume.

## Rolling Back the Volume

1. Log in to the management group that contains the volume that you want to roll back.
2. In the navigation window, select the snapshot to which you want to roll back.
3. Review the snapshot Details tab to ensure you have selected the correct snapshot.
4. Click Snapshot Tasks on the Details tab and select Roll Back Volume.

   The Roll Back Volume window opens, shown in Figure 167.

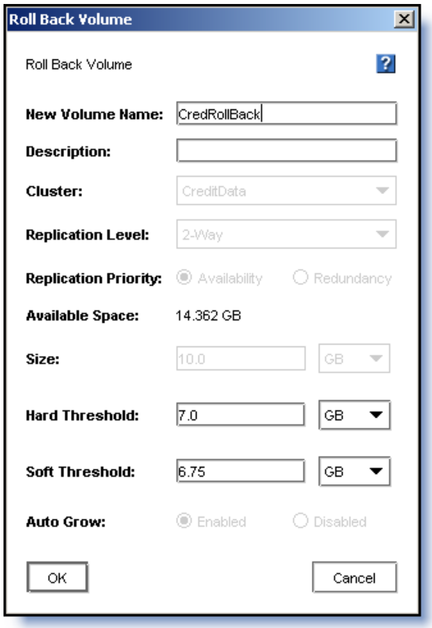

Roll Back Volume

Roll Back Volume

New Volume Name: CredRollBack

Description:

Cluster: CreditData

Replication Level: 2-Way

Replication Priority: ● Availability ○ Redundancy

Available Space: 14.362 GB

Size: 10.0 GB

Hard Threshold: 7.0 GB

Soft Threshold: 6.75 GB

Auto Grow: ● Enabled ○ Disabled

**Figure 167. Rolling Back a Volume**

5. Type a new name for the rolled back volume.

   You can also change the hard and soft thresholds if necessary.

6. Click OK.

   The Roll Back Volume confirmation message, shown in Figure 168, explains that the original volume will be deleted.

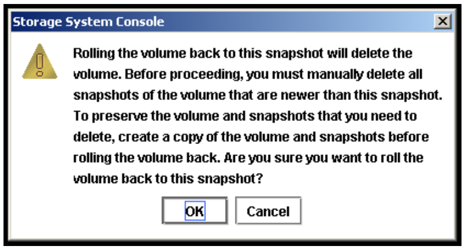

Rolling the volume back to this snapshot will delete the volume. Before proceeding, you must manually delete all snapshots of the volume that are newer than this snapshot. To preserve the volume and snapshots that you need to delete, create a copy of the volume and snapshots before rolling the volume back. Are you sure you want to roll the volume back to this snapshot?

Figure 168. Verifying the Volume Roll Back

7. Click OK.

   The new volume that has the snapshot information is read/write now.

8. Add the restored volume to the original volume list.
9. Reconfigure hosts to access the new volume.

***Warning:*** *The original volume is deleted as part of the rollback.*

# Deleting a Snapshot

Deleting a snapshot moves that snapshot’s incremental data from the SAN. The data necessary to maintain volume consistency are moved up to the next snapshot and the snapshot is removed from the navigation window. The writable space associated with the snapshot is deleted.

***Warning:*** *Deleting a snapshot causes that snapshot’s data to be unavailable from the storage module.*

**Prerequisites**

- Stop applications from accessing the snapshot.
- Disconnect iSCSI initiator(s) from the snapshot.

# Delete the Snapshot

1. Log into the management group that contains the snapshot that you want to delete.
2. In the navigation window, select the snapshot that you want to delete.
3. Review the Details tab to ensure you have selected the correct snapshot.
4. Click Snapshots Tasks and select Delete Snapshot.

   A confirmation message opens.
5. Click OK.

# 14 Working with Scripting

## Scripting Overview

The Storage System Software provides application-based scripting for taking snapshots. Using application-based scripts allows automatic snapshots of a volume and automatic increases in the volume thresholds. Scripting also provides access to Remote Copy, the ability to maintain multiple copies of data across multiple facilities. See "Working with Snapshots" on page 237 for detailed information about snapshot requirements.

The tasks supported by scripting includes

- Taking a snapshot of the volume
- Mounting the snapshot
- (Optional) Unmounting or deleting the snapshot
- Increasing volume thresholds

A scripting tool, named `commandline.CommandLine` provided to access the Console functionality.

## Tools for Scripting

The scripting tool, `java commandline.CommandLine`, creates and deletes snapshots, and automatically increases volume thresholds.

### Java commandline. CommandLine

`Java commandline.CommandLine` is the program that actually invokes the snapshot function in the Console for creating and deleting snapshots. In addition, the program can respond when a soft threshold is reached on a volume and can automatically increase the hard and soft thresholds on that volume.

- First, set the environment.

Table 49. Setting the Environment for Using Scripting Tools

| Operating System | Syntax | Example |
|---|---|---|
| Windows | set CLASSPATH <full path to UI.jar> | set CLASSPATH C:\Program Files\ |

### Table 49. Setting the Environment for Using Scripting Tools

| Operating System | Syntax | Example |
|---|---|---|
| Unix<br>(C Shell type) | setenv CLASSPATH <full path to UI.jar> | setenv CLASSPATH /opt/ |
| Unix<br>(Bourne or Kshell or Bash) | export CLASSPATH=<full path to UI.jar> | export CLASSPATH=/opt/ |

- Then, run the tool.
  `commandline.CommandLine`

*Note:* *Run this program twice to take a snapshot of both the journaling data and the application data if you have them stored in separate volumes.*

### Table 50. Parameters for `commandline.CommandLine`

| Parameter | What It Is |
|---|---|
| userName | Value = text<br>Name of the administrator with full administrative privileges. Can be either the primary or remote administrator, if they are different. |
| password | Value = text<br>The administrator's Storage System Console password. Can be either primary or remote password, if they are different. |
| manager ip addr | Value = IP address<br>IP address of an storage module running a manager in the management group containing either the source volume or the remote volume. |
| volume name | Value = text<br>Name of the volume.<br>For volume_autogrow and FC_LUN commands, this also could be a snapshot. |
| snapshot name | Value = text<br>Name of the snapshot to create. |
| primary volume name | Value = text<br>Name of the primary volume to make remote. |
| remote volume name | Value = text<br>Name of the remote volume created in the Console. |
| remote snapshot name | Value = text<br>Name of the remote snapshot. |
| remote snapshot description | Value = text<br>Description of the remote snapshot. |

### Table 50. Parameters for `commandline.CommandLine (Cont'd)`

| Parameter | What It Is |
|---|---|
| soft quota* | Value = number<br>Size of the volume's new soft threshold in megabytes (MB). May be the soft threshold of the new primary volume if using Remote Copy. |
| hard quota* | Value = number<br>Size of the source volume's new hard threshold in megabytes (MB). May be the hard threshold of the new primary volume if using Remote Copy. |
| description* | Value = text<br>(Optional) Description associated with the snapshot. |
| failure timeout seconds | Value = number<br>The number of seconds to wait until exiting with a failure. |
| grow size | Value = number<br>The size in MegaBytes by which to increase the volume thresholds. |
| volume_snapshot | Use this value as written. (This is verbatim.)<br>Creates a snapshot from a volume. |
| volume_snapshot_no_acl | Use this value as written. (This is verbatim.)<br>Creates a snapshot from a volume, however, the snapshot will not inherit the volume list information from the parent volume. |
| volume_delete | Use this value as written. (This is verbatim.)<br>Deletes a volume. |
| volume_autogrow_set | Use this value as written. (This is verbatim.)<br>Sets the value by which to increase a volume threshold. |
| volume_autogrow_get | Use this value as written. (This is verbatim.)<br>Returns the value currently in the volume_autogrow_set command. |
| volume_rpc_checksums_set | Enables or disables checksum on the volume. |
| volume_rpc_checksums_get | Gets the value of the checksum setting on the volume. |
| get_snapshot_name | Gets the name of the most recent snapshot for the volume. |
| volume_remote_snapshot | Use this value as written. (This is verbatim.)<br>Makes a remote snapshot of a volume. |
| volume_make_primary | Use this value as written. (This is verbatim.)<br>Makes a remote volume into a primary volume. |
| volume_make_remote | Use this value as written. (This is verbatim.)<br>Makes a primary volume into a remote volume. |
| print_management_group_information | Gets the management group description. |

Table 50. Parameters for `commandline.CommandLine (Cont'd)`

| Parameter | What It Is |
|---|---|
| delete_writable_space | Deletes additional space created in the cluster when any activity takes place on the snapshot. |
| print_volumes_snapshot_list | Gets list of all snapshots for the volume along with the snapshot creation timestamps. |
| volume_rollback | Rolls back a volume from a snapshot. |
| acl_add_vlist | Add the volume list to the snapshot/volume in order to script the snapshot connection to a server. |
| acl_delete_vlist | Removes a volume list from a snapshot/volume. |
| get_available_mode | Gets the shutdown mode of the management group. |
| set_available_mode | Sets the shutdown mode of the management group. |
| * You must provide either all three items, or none of them. For example, you cannot provide only a soft threshold value. | |

# Scripted Commands for Volumes and Snapshots

Below are examples of the Storage System Software functions that can be accomplished using application-based scripts.

## Creating a Snapshot

Create a snapshot using **`commandline.CommandLine`**

```
commandline.CommandLine <admin name> <admin password> <manager ip>
volume_snapshot <source volume name> <snapshot name> [<soft threshold
(Megabytes)> <hard threshold (Megabytes)> <description>] [<failure
timeout seconds>]
```

**Example**

Joe Jones is creating a snapshot for his management group Images, volume named X-Rays, and he wants the snapshot name to be XRayReview. The size of the thresholds for the snapshot is a 100 MB hard threshold and a 98 MB soft threshold. So Joe's use of java commandline.CommandLine will look as follows

```
java commandline.CommandLine jjones trumpet 10.0.111.212 volume_snapshot
X-Rays XRayReview 98 100 "review volume for xray storage" 10
```

## Deleting a Snapshot

Delete a snapshot using **`java commandline.CommandLine`**

```
java commandline.CommandLine <admin name> <admin password> <manager ip>
volume_delete <snapshot name> [<failure timeout seconds>]
```

**Example**

Joe Jones plans to retain the snapshot for a review period, so he writes a script to delete the snapshot after 5 weeks.

```
java commandline.CommandLine jjones trumpet 10.0.111.212 volume_delete
XRayReview 45
```

# Increasing Volume Hard and Soft Thresholds

You can create a script that automatically increases the hard and soft volume thresholds by a specific amount.

The operation is triggered when a soft threshold is reached. It then raises both the soft and hard thresholds by the amount you specify in the script. The thresholds will only increase

- when there is sufficient room in the cluster to accommodate the increases or
- to the point where the hard threshold equals the volume length

whichever of these conditions occur first. To increase space in the cluster by adding more storage modules or to increase the volume length, follow instructions as described in Chapter 11, “Working with Clusters” or Chapter 12, “Working with Volumes.”

## Scripting Automatic Threshold Increases

Below is an example of scripting automatic threshold increases using **java commandline.CommandLine**

```
 java commandline.CommandLine <admin name> <admin password> <manager ip>
volume_autogrow_set <volume name> <grow size (Megabytes)> [<failure
timeout seconds>]
```

**Example**

Joe Jones creates a script to automatically increase the hard and soft thresholds for his X-Rays volume The volume length is 10 GB with a hard threshold of 2 GB and a soft threshold of 1 GB. Joe scripts the increases for increments of 512 MB.

```
java commandline.CommandLine jjones trumpet 10.0.111.212
volume_autogrow_set X-Rays 512 600
```

## Reviewing the Increment Size for Increasing the Thresholds

You can run an operation to review the setting for automatic threshold increases using **java commandline.CommandLine**

```
 java commandline.CommandLine <admin name> <admin password> <manager ip>
volume_autogrow_get <volume name> [<failure timeout seconds>]
```

**Example**

```
java commandline.CommandLine jjones trumpet 10.0.111.212
volume_autogrow_get X-Rays 60
```

# Scripted Commands for Remote Copy

Scripting operations for Remote Copy use the same tools that are available for scripting snapshots, with the addition of parameters specific to Remote Copy. Using the command line parameters you can create scripts for these cases:

- Creating a primary snapshot
- Creating a remote snapshot
- Making a primary volume into a remote volume
- Failing over to a remote snapshot

***Note:*** *You must first use the Storage System Console to create a remote snapshot between the desired management groups before creating a scripted remote snapshot.*

## Creating a Remote Snapshot in a Different Management Group

First, create the primary snapshot.

```
java commandline.CommandLine <primary admin name> <primary admin
password> <primary manager ip> volume_snapshot <primary volume name>
<primary snapshot name> [<soft threshold (Megabytes)> <hard threshold
(Megabytes)> <description>] [<failure timeout seconds>]
```

Next, create the remote snapshot.

```
java commandline.CommandLine <remote admin name> <remote admin password>
<remote manager ip> volume_remote_snapshot <remote volume name> <remote
snapshot name> <remote snapshot description> <primary admin name> <primary
admin password> <primary manager ip> <primary snapshot name> [<failure
timeout seconds>]
```

**Example**

Joe Jones plans to create a remote snapshot of his X-Rays volume in the backup management group in the corporate backup site. He is naming this new remote snapshot RSS2_xrays and the new primary snapshot PSS2_xrays. He created his remote volume RemVolX_Rays using the Console and named his first primary snapshot PSS1_xrays and his first remote snapshot RSS1_xrays. The size of the thresholds for the new primary and remote snapshots are the same — 500 MB hard thresholds and 500 MB soft thresholds. The script is as follows:

```
java commandline.CommandLine jjones trumpet 10.0.111.212 volume_snapshot
X-Rays PSS2_xrays 500 500 “first primary snapshot” 15
java commandline.CommandLine jjones saxophone 10.10.45.72
volume_remote_snapshot RemVolX_Rays RSS2_xrays “second remote snapshot”
jjones trumpet 10.0.111.212 PSS2_xrays 15
```

# Creating a Remote Snapshot in the Same Management Group

First, create the primary snapshot.

```
java commandline.CommandLine <primary admin name> <primary admin
password> <primary manager ip> volume_snapshot <primary volume name>
<primary snapshot name> [<failure timeout seconds>]
```

Next, create the remote snapshot.

```
java commandline.CommandLine <primary admin name> <primary admin
password> <primary manager ip> volume_remote_snapshot <remote volume
name> <remote snapshot name> <remote snapshot description> <primary
snapshot name> [<failure timeout seconds>]
```

**Example**

If Joe Jones was creating his remote snapshot in the same management group, the script would look like this:

```
java commandline.CommandLine jjones trumpet 10.0.111.212 volume_snapshot
X-Rays PSS2_xrays 500 500 "first primary snapshot" 30
java commandline.CommandLine jjones trumpet 10.0.111.212
volume_remote_snapshot RemVolX_Rays RSS2_xrays "second remote snapshot"
PSS2_xrays 30
```

# Converting a Remote Volume to a Primary Volume and Back to a Remote Volume

Convert a remote volume into a primary volume to gain read/write access to the most recently completed Remote Copy snapshot. However, if that remote volume is the target for scheduled remote snapshots, those snapshots cannot take place if the remote volume is not present. Therefore, you use the operation for returning the primary volume back to its remote status to allow the scheduled remote snapshots to continue.

## Make Remote Volume into Primary Volume

```
java commandline.CommandLine <remote admin name> <remote admin password>
<remote manager ip> volume_make_primary <remote volume name> [<soft quota
(Megabytes)> <hard quota (Megabytes)>] [<failure timeout seconds>]
```

## Make Primary Volume into Remote Volume

```
 java commandline.CommandLine <primary admin name> <primary admin
password> <primary manager ip> volume_make_remote <primary volume name>
<snapshot name> <snapshot description> [<failure timeout seconds>]
```

**Example**

Joe has scripted an operation to make his remote volume into a primary volume once a week so that he can access the data from the most recently completed scheduled snapshot. Since he is running scheduled remote snapshots to that volume, he then needs to convert that primary volume back into a remote volume so that the remote snapshot schedule is maintained.

```
java commandline.CommandLine jjones saxophone 10.10.45.72
volume_make_primary RemVolX_Rays 512000 512000 30
```

```
java commandline.CommandLine jjones trumpet 10.3.11.19 volume_make_remote
RemVolX_Rays snapshot_convert “snapshot from making vol remote” 30
```

# Scripting Failover

Scripting failover uses a **`java commandline.CommandLine`** script along with an iSCSI initiator.

## Make Remote Volume into Primary Volume

```
java commandline.CommandLine <remote admin name> <remote admin password>
<remote manager ip> volume_make_primary <remote volume name> [<soft quota
(Megabytes)> <hard quota (Megabytes)>] [<failure timeout seconds>]
```

**Example**

Joe’s script for failing over to his remote volume would include the following commands to make the remote volume into a primary volume and mount it in the local network to make it available to the backup application servers.

```
java commandline.CommandLine jjones saxophone 10.10.45.72
volume_make_primary RemVolX_Rays 512000 512000 30
```

# 15 Controlling Client Access to Volumes

## Client Access Overview

Access to storage volumes by application servers, also called “clients” or “hosts,” is controlled using volume lists and authentication groups.

- Authentication groups identify the client or entity accessing the volume.
- Volume lists provide the association between authentication groups and volumes. They are created at the management group level and they link designated volumes with the authentication groups that can access those volumes.

The relationship between authentication groups, volume lists, and volumes is shown in Figure 169.

## Topics in this Chapter

- ”Client Access Overview”
- ”Creating Access with Authentication Groups”
- ”Volume Lists Overview”

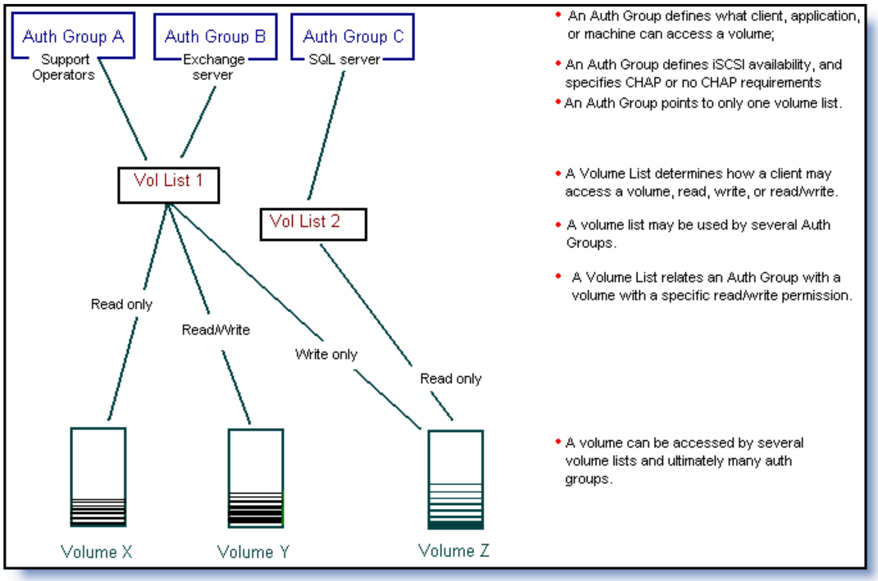

- An Auth Group defines what client, application, or machine can access a volume;
- An Auth Group defines iSCSI availability, and specifies CHAP or no CHAP requirements
- An Auth Group points to only one volume list.
- A Volume List determines how a client may access a volume, read, write, or read/write.
- A volume list may be used by several Auth Groups.
- A Volume List relates an Auth Group with a volume with a specific read/write permission.
- A volume can be accessed by several volume lists and ultimately many auth groups.

**Figure 169. Understanding the Authentication Group and Volume List Relationship**

- Authentication groups may be user groups, software applications, or other servers.
- An authentication group may only access one volume list.
- Volume lists specify a volume name and a read/write permission to that volume.
- A volume list may be accessed by several authentication groups.

# Creating Access with Authentication Groups

After you have configured storage and created volumes, you then create access to the volumes using authentication groups and volume lists.

1. Create an authentication group.
2. Create a volume list and associate the authentication group to the specific volume(s) it can access. See "Volume Lists Overview" on page 277.

## Types of Client Access

The Storage System Software software supports iSCSI initiator access to volumes.

Authentication using an iSCSI initiator can be based on the initiator node name (single hosts) or CHAP-based authentication (single or multiple hosts).

## Client Access Using iSCSI

Client access using iSCSI can be authenticated via the initiator node name (single host) or via CHAP (Challenge-Handshake Authentication Protocol), which can support single or multiple hosts.

You may also enable load balancing using compliant iSCSI initiators. iSCSI load balancing distributes iSCSI sessions across storage modules in a cluster.

***Note:*** *The iSCSI terminology in this discussion is based on the Microsoft iSCSI Initiator terminology.*

# Configuring Authentication Groups for iSCSI

When configuring client access using iSCSI, you create an authentication group that allows iSCSI access. The New Authentication Group window with the iSCSI tab is shown in Figure 170.

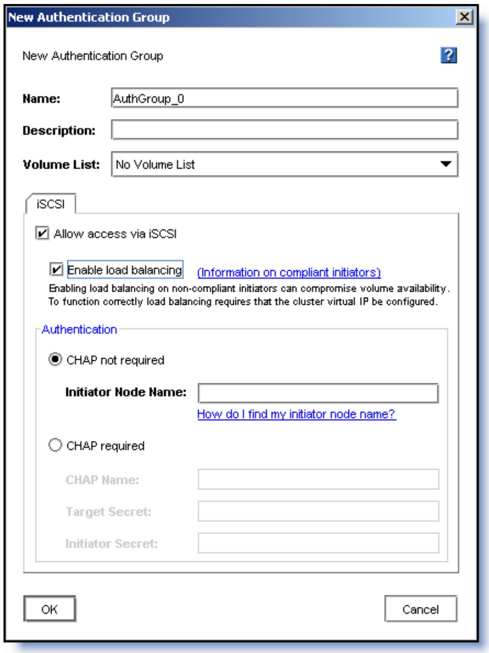

New Authentication Group

Name: AuthGroup_0

Description:

Volume List: No Volume List

iSCSI

☑ Allow access via iSCSI

☑ Enable load balancing (Information on compliant initiators)

Enabling load balancing on non-compliant initiators can compromise volume availability. To function correctly load balancing requires that the cluster virtual IP be configured.

Authentication

● CHAP not required

Initiator Node Name:

How do I find my initiator node name?

○ CHAP required

CHAP Name:

Target Secret:

Initiator Secret:

### Figure 170. Creating a New Authentication Group for iSCSI Access

Planning iSCSI access requires planning for load balancing and configuring authentication:

- Single host with or without CHAP

  or

- Multiple hosts, with 1-way or 2-way CHAP

# Planning iSCSI and Load Balancing

Use the iSCSI load balancing feature to improve iSCSI performance and scalability. iSCSI load balancing distributes iSCSI sessions for different volumes evenly across storage modules in a cluster.

### Requirements

- Cluster configured with a Virtual IP address. See "Clusters and iSCSI" on page 201.
- A compliant iSCSI initiator.

## Compliant iSCSI Initiators

A compliant initiator is one that supports iSCSI Login-Redirect AND has passed test criteria for iSCSI failover in a load balanced configuration.

Find information about which iSCSI initiators are compliant by clicking the link in the New or Edit Authentication Group window, shown in Figure 171.

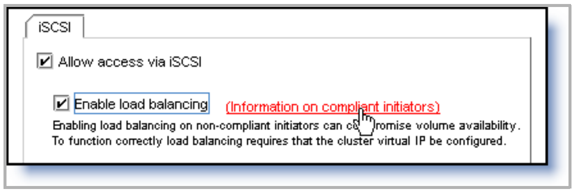

☑ Allow access via iSCSI

☑ Enable load balancing (Information on compliant initiators)

Enabling load balancing on non-compliant initiators can compromise volume availability. To function correctly load balancing requires that the cluster virtual IP be configured.

**Figure 171. Finding Compliant Initiator Information**

The link opens the iSCSI initiator information window, shown in Figure 172. Scroll down for a list of compliant initiators.

If your initiator is not on the list, continue to use a traditional VIP.

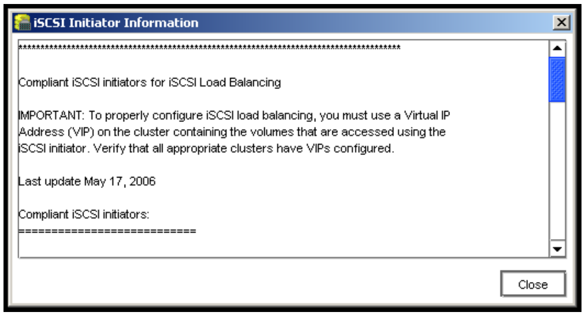

```
*************************************************************************

Compliant iSCSI initiators for iSCSI Load Balancing

IMPORTANT: To properly configure iSCSI load balancing, you must use a Virtual IP
Address (VIP) on the cluster containing the volumes that are accessed using the
iSCSI initiator. Verify that all appropriate clusters have VIPs configured.

Last update May 17, 2006

Compliant iSCSI initiators:
==========================
```

**Figure 172. Viewing Compliant iSCSI Initiators**

# Planning iSCSI and CHAP

CHAP is a standard authentication protocol. The Storage System Software supports no CHAP, 1-way CHAP, and 2-way CHAP, as shown in Figure 173 on page 269.

## CHAP Glossary

**Target secret**

The target secret is required. It is used in both 1-way and 2-way CHAP configurations when the target volume challenges the iSCSI initiator.

### Initiator secret

The initiator secret is optional. It is used in 2-way CHAP when the iSCSI initiator challenges the target volume.

## How CHAP Works

- No CHAP—authorized initiators can log in to the volume without proving their identity. The target does not challenge the client.
- 1-way CHAP—initiators must log in with a target secret to access the volume. This secret proves the identity of the initiator to the target.
- 2-way CHAP—initiators must log in with a target secret to access the volume as in 1-way CHAP. In addition, the target must prove its identity to the initiator using the initiator secret. This second step prevents target spoofing.

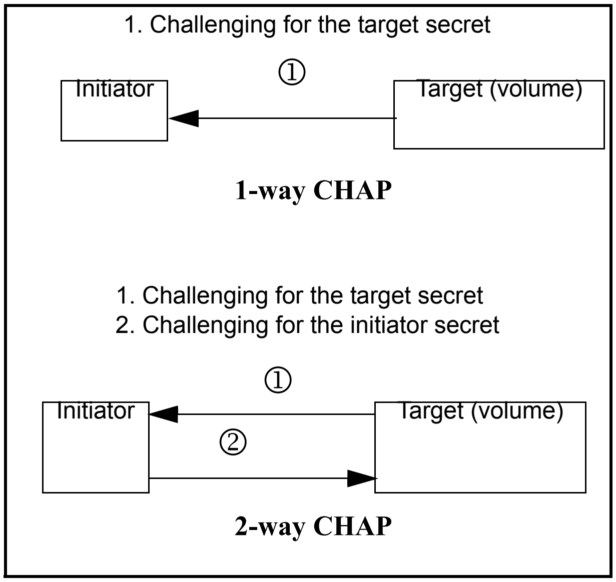
1. Challenging for the target secret
①
Initiator
Target (volume)
1-way CHAP
1. Challenging for the target secret
2. Challenging for the initiator secret
①
Initiator
②
Target (volume)
2-way CHAP

**Figure 173. Differentiating Types of CHAP**

CHAP is optional. However, if you configure 1-way or 2-way CHAP, you must remember to configure both the authentication group and the iSCSI initiator with the appropriate parameters. Table 51 lists the requirements for configuring CHAP.

## Requirements for Configuring CHAP

### Table 51. Configuring iSCSI CHAP

| CHAP Level | What to Configure in the Authentication Group | What to Configure in the iSCSI Initiator |
|---|---|---|
| CHAP not required | • initiator node name only | • No configuration requirements |
| 1-way CHAP | • CHAP Name*<br>• Target Secret | • Enter the target secret when logging on to available target. |
| 2-way CHAP | • CHAP Name*<br>• Target Secret<br>• Initiator Secret | • Enter the initiator secret.<br>• Enter the target secret. |
| * If using CHAP with a single node only, use the initiator node name as the CHAP name. | | |

## Sample iSCSI Configurations

Figure 174 illustrates the configuration for a single host authentication with CHAP not required with Microsoft iSCSI.

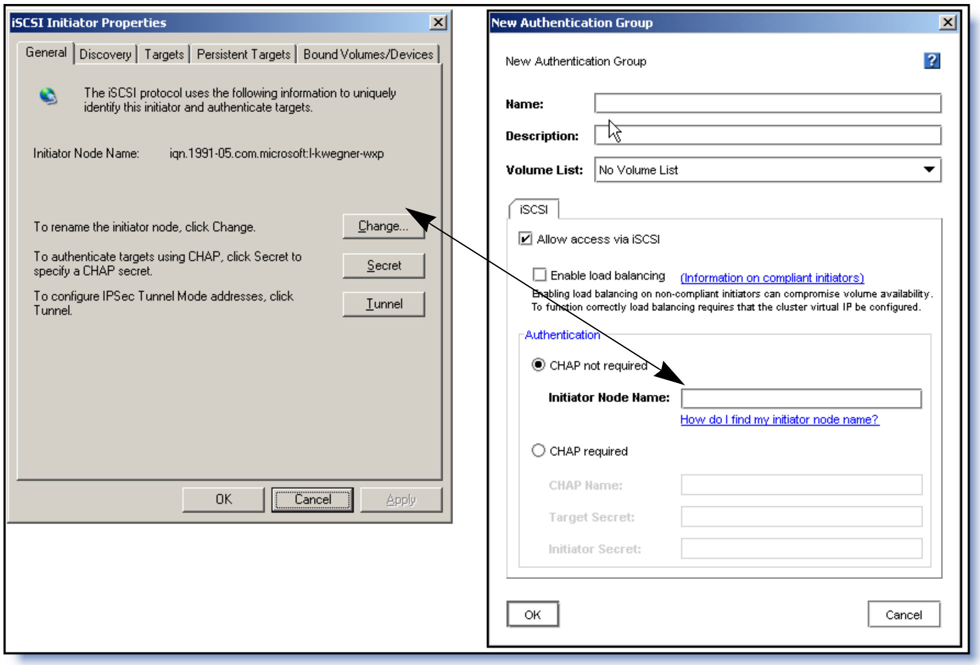

iSCSI Initiator Properties

The iSCSI protocol uses the following information to uniquely identify this initiator and authenticate targets.

Initiator Node Name: iqn.1991-05.com.microsoft:l-kwegner-wxp

To rename the initiator node, click Change.

To authenticate targets using CHAP, click Secret to specify a CHAP secret.

To configure IPSec Tunnel Mode addresses, click Tunnel.

New Authentication Group

New Authentication Group

Name:

Description:

Volume List: No Volume List

☑ Allow access via iSCSI

□ Enable load balancing (Information on compliant initiators)

Enabling load balancing on non-compliant initiators can compromise volume availability. To function correctly load balancing requires that the cluster virtual IP be configured.

Authentication

● CHAP not required

Initiator Node Name:

How do I find my initiator node name?

○ CHAP required

CHAP Name:

Target Secret:

Initiator Secret:

Figure 174. Viewing the MS iSCSI Initiator to Copy the Initiator Node Name

Figure 175 illustrates the configuration for a single host authentication with 1-way CHAP required.

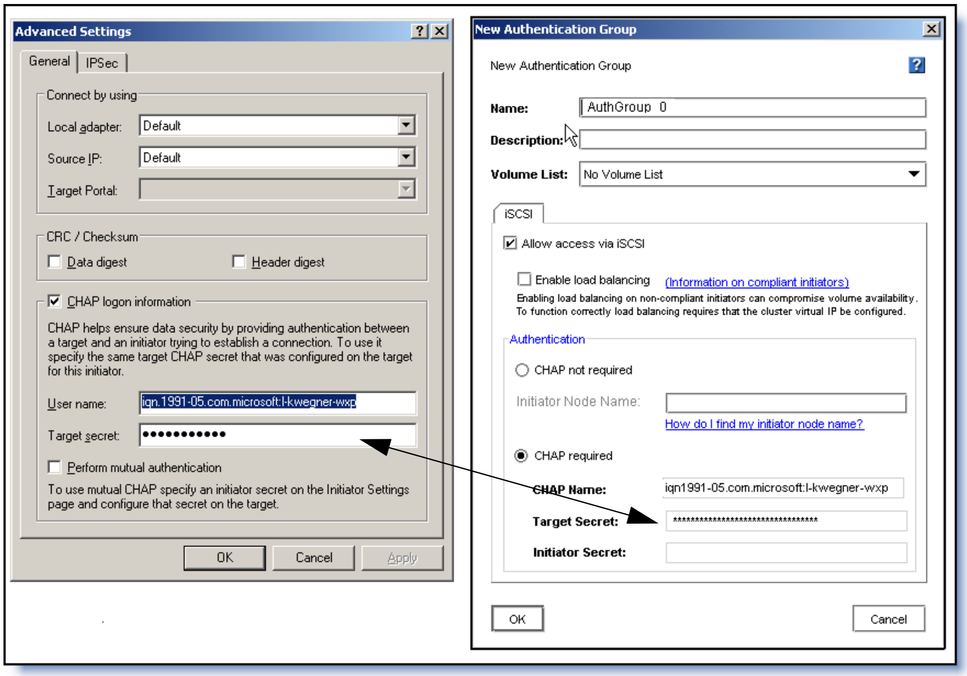

Advanced Settings

General | IPSec

Connect by using

Local adapter: Default

Source IP: Default

Target Portal:

CRC / Checksum

□ Data digest □ Header digest

☑ CHAP logon information

CHAP helps ensure data security by providing authentication between a target and an initiator trying to establish a connection. To use it specify the same target CHAP secret that was configured on the target for this initiator.

User name: iqn.1991-05.com.microsoft:l-kwegner-wxp

Target secret: ●●●●●●●●●●●

□ Perform mutual authentication

To use mutual CHAP specify an initiator secret on the Initiator Settings page and configure that secret on the target.

OK | Cancel | Apply

New Authentication Group

New Authentication Group

**Name:** AuthGroup_0

**Description:**

**Volume List:** No Volume List

iSCSI

☑ Allow access via iSCSI

□ Enable load balancing (Information on compliant initiators)

Enabling load balancing on non-compliant initiators can compromise volume availability. To function correctly load balancing requires that the cluster virtual IP be configured.

Authentication

○ CHAP not required

Initiator Node Name:

How do I find my initiator node name?

● CHAP required

**CHAP Name:** iqn1991-05.com.microsoft:l-kwegner-wxp

**Target Secret:** ********************************

**Initiator Secret:**

OK | Cancel

## Figure 175. Configuring iSCSI (shown in the MS iSCSI initiator) for a single host with CHAP

Figure 176 illustrates the configuration for a single host authentication with 2-way CHAP required.

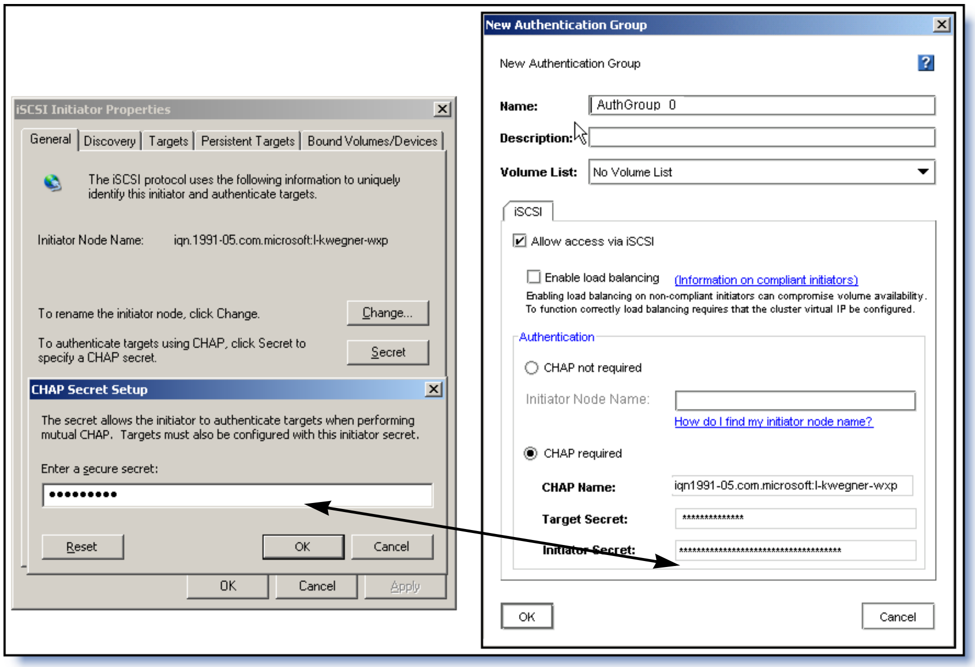

The iSCSI protocol uses the following information to uniquely identify this initiator and authenticate targets.

Initiator Node Name: iqn.1991-05.com.microsoft:l-kwegner-wxp

To rename the initiator node, click Change.

To authenticate targets using CHAP, click Secret to specify a CHAP secret.

The secret allows the initiator to authenticate targets when performing mutual CHAP. Targets must also be configured with this initiator secret.

Enter a secure secret: •••••••••

New Authentication Group

Name: AuthGroup_0

Description:

Volume List: No Volume List

☑ Allow access via iSCSI

☐ Enable load balancing (Information on compliant initiators)

Enabling load balancing on non-compliant initiators can compromise volume availability. To function correctly load balancing requires that the cluster virtual IP be configured.

○ CHAP not required

Initiator Node Name:

How do I find my initiator node name?

● CHAP required

CHAP Name: iqn1991-05.com.microsoft:l-kwegner-wxp

Target Secret: *************

Initiator Secret: ************************************

Figure 176. Adding an Initiator Secret for 2-way CHAP (shown in the MS iSCSI initiator)

***Warning:*** *Without the use of shared storage (clustering) technology, allowing more than one iSCSI application server to connect to a volume in read/write mode could result in data corruption.*

# Creating an Authentication Group

Have a volume list already set up.

1. In the navigation window, log in to the management group.
2. Click the Authentication Groups tab to bring it to the front.

3. Click Authentication Groups Tasks and select New Authentication Group.

   The New Authentication Group window opens, shown in Figure 177.

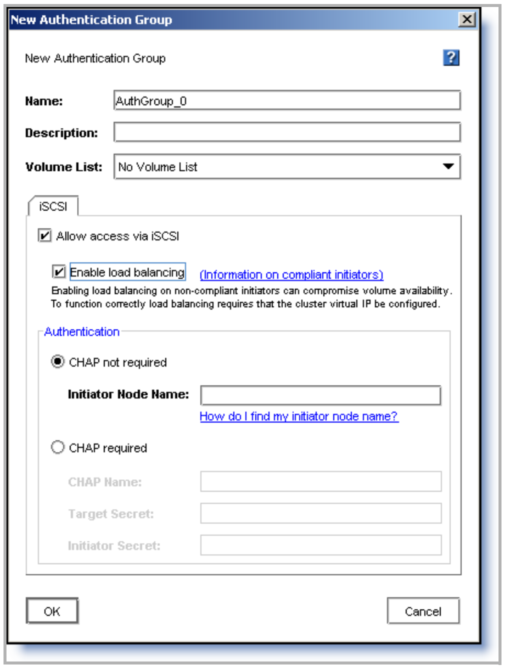

New Authentication Group

Name: AuthGroup_0

Description:

Volume List: No Volume List

☑ Allow access via iSCSI

☑ Enable load balancing (Information on compliant initiators)

Enabling load balancing on non-compliant initiators can compromise volume availability. To function correctly load balancing requires that the cluster virtual IP be configured.

● CHAP not required

Initiator Node Name:

How do I find my initiator node name?

○ CHAP required

CHAP Name:

Target Secret:

Initiator Secret:

Figure 177. Creating an Authentication Group

4. Type a name and description for the authentication group.

   The authentication group name is case sensitive. It cannot be changed later.

5. Select the volume list, if appropriate.

   The volume list can be added later.

# Configuring iSCSI

1. In the New Authentication Group window, shown in Figure 178, select the check box to allow access via iSCSI.
2. Click the link to review the list of compliant iSCSI initiators.

   Scroll down to see the entire list.

3. Select the check box to enable load balancing.
4. Select the authentication method, either CHAP not required or CHAP required.

**Warning:** *Using a non-compliant iSCSI initiator for load balancing can compromise volume availability during iSCSI failover events.*

**Warning:** *Without the use of shared storage (clustering) technology, allowing more than one iSCSI application server to connect to a volume without cluster-*

*aware applications and /or file systems in read/write mode could result in data corruption.*

5. In the Initiator Node name field, enter the correct string.

   To do this, open your iSCSI initiator software an look for that string there.

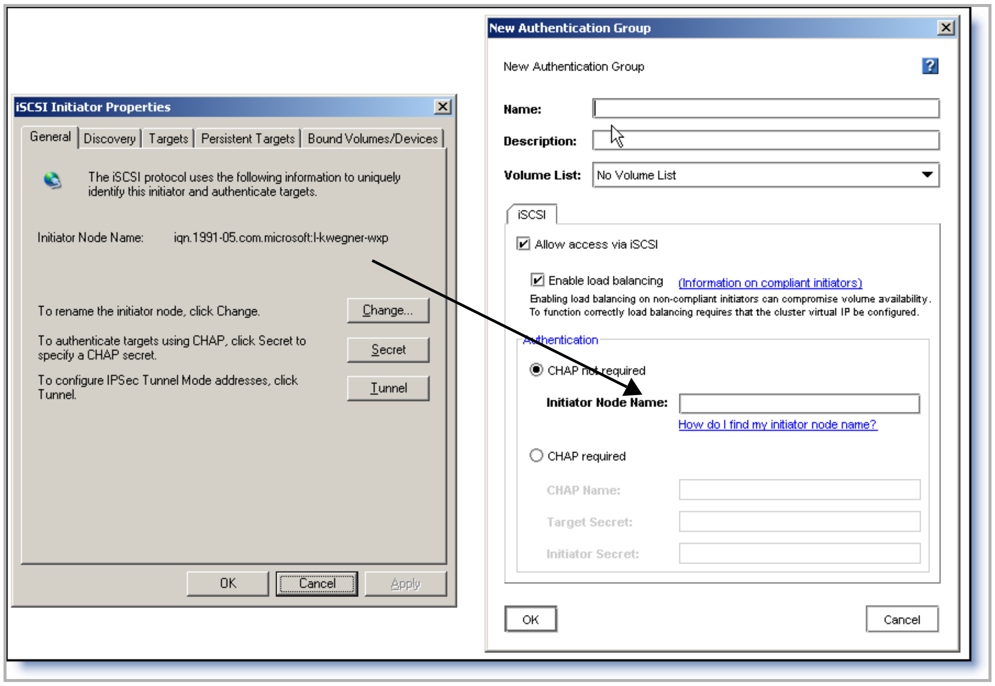

The iSCSI protocol uses the following information to uniquely identify this initiator and authenticate targets.

Initiator Node Name: iqn.1991-05.com.microsoft:l-kwegner-wxp

To rename the initiator node, click Change.

To authenticate targets using CHAP, click Secret to specify a CHAP secret.

To configure IPSec Tunnel Mode addresses, click Tunnel.

Volume List: No Volume List

☑ Allow access via iSCSI

☑ Enable load balancing (Information on compliant initiators)

Enabling load balancing on non-compliant initiators can compromise volume availability. To function correctly load balancing requires that the cluster virtual IP be configured.

● CHAP not required

Initiator Node Name:

How do I find my initiator node name?

○ CHAP required

CHAP Name:

Target Secret:

Initiator Secret:

Figure 178. Creating iSCSI Access in a New Authentication Group

## Authenticate with CHAP Not Required

For detailed illustrations of the relationship between the authentication group fields and the MS iSCSI Initiator, see "Planning iSCSI and CHAP" on page 268.

1. In the Authentication Group window, select the radio button CHAP not required.
2. Copy the initiator node name, provided by MS iSCSI Initiator, into the Initiator Node Name field.

The MS iSCSI Initiator is a third-party software, available on your desk top as an icon.

## Authenticate with CHAP Required

1. In the Authentication Group window, select the radio button CHAP required.
2. Complete the fields necessary for the type of CHAP you intend to configure, as shown in Table 52 .

3. Copy the initiator node name, provided by MS iSCSI Initiator, into the Initiator Node Name field.

**Table 52. Entering CHAP Information in a New Authentication Group**

| For this CHAP Mode | Complete these fields |
|---|---|
| 1-way CHAP | • CHAP name<br>• Target Secret—minimum of 12 characters |
| 2-way CHAP | • CHAP name<br>• Target Secret—minimum of 12 characters<br>• Initiator Secret—minimum of 12 characters; must be alphanumeric |

## Best Practice

In the Microsoft iSCSI initiator, target and initiator secrets are not displayed. Keep a separate record of the iSCSI initiator CHAP information and the corresponding authentication group information.

## Finishing iSCSI Configuration

Click OK in the Authentication group window if you are finished configuring the authentication group.

## Finishing the New Authentication Group

1. Click OK on the New Authentication Group window when you are finished.
2. In the management group's tab window, go to the Authentication Groups tab to see the new group displayed in the list, shown in Figure 179.

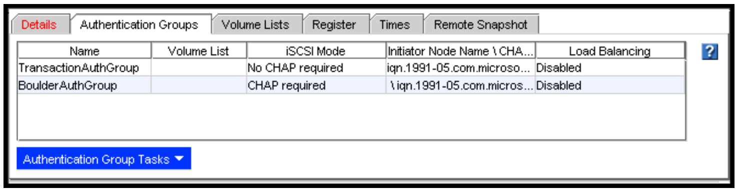

| Name | Volume List | iSCSI Mode | Initiator Node Name \ CHA... | Load Balancing |
|---|---|---|---|---|
| TransactionAuthGroup | | No CHAP required | iqn.1991-05.com.microso... | Disabled |
| BoulderAuthGroup | | CHAP required | \ iqn.1991-05.com.micros... | Disabled |

**Figure 179. Viewing the Authentication Groups**

# Editing an Authentication Group

You may edit the following fields for an authentication group:

- Change the description.
- Add a volume list.
- Change a volume list.
- Change the choice of load balancing.
- Change the types of authentication in either of those modes.

**Warning:** *Editing an authentication group may interrupt client access to volumes. If necessary, or if the client is sensitive to disconnections, stop client access before editing an authentication group.*

**Warning:** *If you change an iSCSI authentication group, you must log off and log back on to the target in the iSCSI initiator for the changes to take effect.*

See “Client Access Using iSCSI” on page 266 before changing iSCSI authentication group parameters.

1. Log in to the management group.
2. In the tab window, click the Authentication Groups tab to bring it to the front.
3. Select from the table the group you want to edit.
4. Click Authentication Groups Tasks and select Edit Authentication Group.

   The Edit Authentication Group window opens, shown in Figure 180.

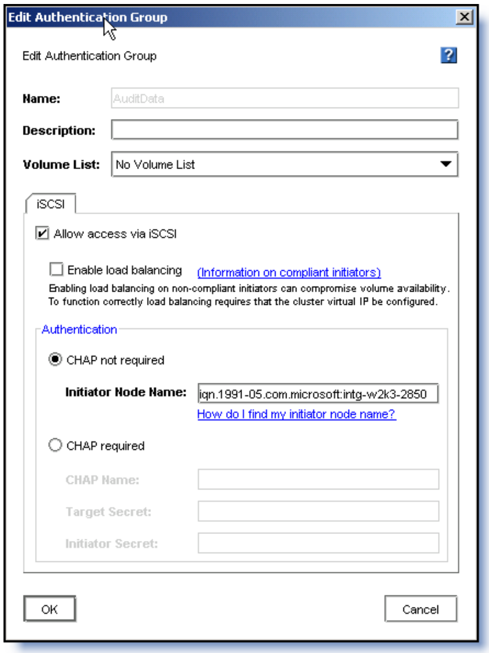

Edit Authentication Group

Edit Authentication Group

Name: AuditData

Description:

Volume List: No Volume List

iSCSI

☑ Allow access via iSCSI

□ Enable load balancing (Information on compliant initiators)

Enabling load balancing on non-compliant initiators can compromise volume availability.
To function correctly load balancing requires that the cluster virtual IP be configured.

Authentication

● CHAP not required

Initiator Node Name: iqn.1991-05.com.microsoft:intg-w2k3-2850

How do I find my initiator node name?

○ CHAP required

CHAP Name:

Target Secret:

Initiator Secret:

**Figure 180. Editing an Authentication Group**

5. Change the appropriate information.
6. Click OK when you are finished.

# Deleting an Authentication Group

Deleting an authentication group stops access to volumes by clients using that group. Access to the same volume by other authentication groups continues.

1. In the navigation window, log in to the management group.
2. Click the Authentication Groups tab to bring it to the front.
3. From the table, select the group you want to delete.
4. Click Authentication Groups Tasks and select Delete Authentication Group.

   A confirmation window opens.
5. Click OK to delete the group.

# Volume Lists Overview

The volume list for an authentication group contains the list of volumes accessible to that group and how it may be accessed with reads and writes.

**Prerequisites**

- At least one management group has been created.
- At least one cluster has been created in that management group.
- At least one volume has been created in that cluster.

## Requirements for Volume Lists

- An authentication group can use only one volume list.
- If you do not use host clustering or shared storage technologies, only one authentication group should have read/write access to a specific volume.

# Planning Volume Lists

Planning volume lists takes into account multiple factors.

- Applications or clients that access the volume
- Permissions that you assign for those clients
- Configuration of iSCSI initiator sessions

# Creating a Volume List

1. Log into the management group.
2. In the management group's tab window, click the Volume Lists tab to bring it to the front, shown in Figure 181.

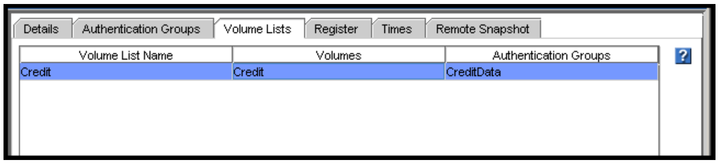

| Volume List Name | Volumes | Authentication Groups |
|---|---|---|
| Credit | Credit | CreditData |

**Figure 181. Viewing Volume Lists Tab Window**

3. Click Volume Lists Tasks and select New Volume List.

   The New Volume List window opens, shown in Figure 182.

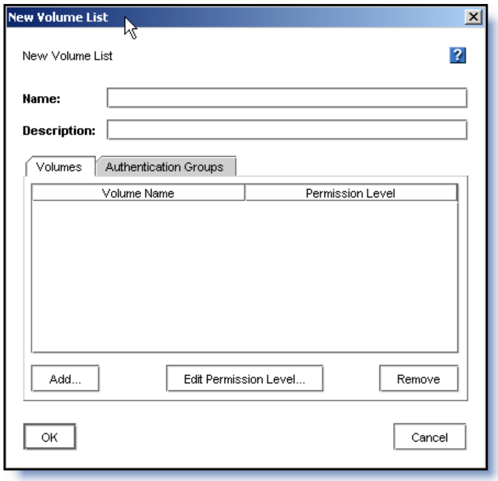

**Figure 182. Creating a New Volume List**

4. Type a name and description for the volume list.

   The volume list name is case sensitive. It cannot be changed later.

# Adding Volumes to the Volume List

1. Select the Volumes tab, shown in Figure 182.
2. Click Add to add one or more volumes or snapshots to the volume list.

3. The Add Volume to Volume List window opens, shown in Figure 183.

   This list shows the volumes and snapshots available for client access via authentication groups.

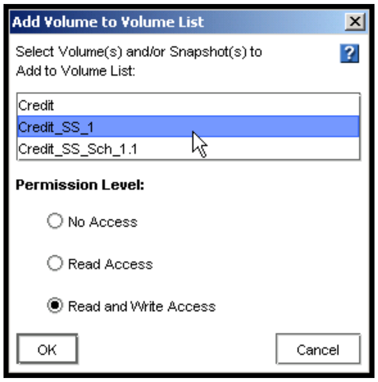

Add Volume to Volume List

Select Volume(s) and/or Snapshot(s) to Add to Volume List:

- Credit
- Credit_SS_1
- Credit_SS_Sch_1.1

**Permission Level:**

○ No Access

○ Read Access

● Read and Write Access

### Figure 183. Adding Volumes and Snapshots to a Volume List

***Warning:*** *Without the use of shared storage (clustering) technology, allowing more than one iSCSI application server to connect to a volume in read/write mode could result in data corruption.*

4. Select the volumes or snapshots to add to the list.
5. Select the access permission level for the volumes and snapshots.

   Characteristics of the permission levels are described in Table 53 .

### Table 53. Characteristics of Permission Levels

| Type of Access | Allows This |
|---|---|
| No Access | Prevents the authentication group from accessing the volume or snapshot. |
| Read Access | Restricts the authentication group to read-only access to the data on the volume or snapshot. |
| Read/Write Access | Allows the authentication group read and write permissions to the volume. |

***Note:*** *Microsoft Windows requires read/write access to volumes.*

6. Click OK when you are finished.

   In the tab view, note that the volumes and snapshots are listed under the Volumes tab.

## Best Practice

Enable read/write permissions on iSCSI volumes and snapshots.

# Associating Authentication Groups with a Volume List

1. Next, select the Authentication Groups tab in the New Volume List window, shown in Figure 184.

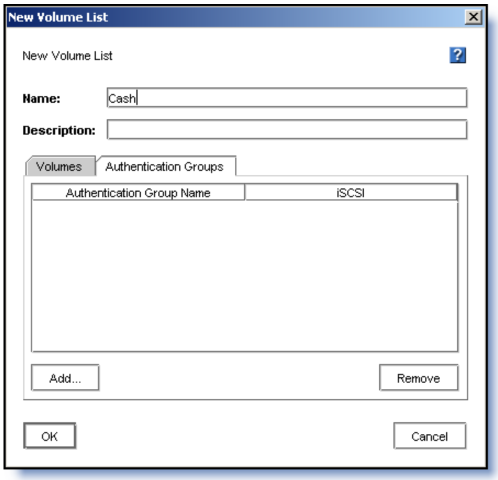

Figure 184. Connecting Authentication Groups to a Volume List

2. Click Add.

   A list of existing authentication groups is displayed, shown in Figure 185.

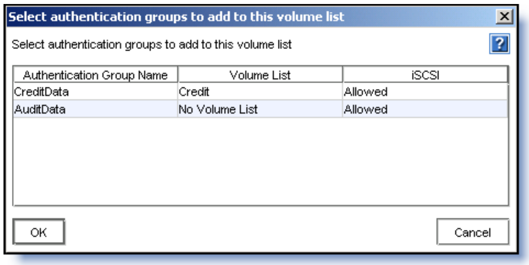
Select authentication groups to add to this volume list
Select authentication groups to add to this volume list
?
OK
Cancel

| Authentication Group Name | Volume List | iSCSI |
|---|---|---|
| CreditData | Credit | Allowed |
| AuditData | No Volume List | Allowed |

Figure 185. List of Authentication Groups to Add

3. Select the authentication group for which you want to give access to the volume or snapshot.

4. Click OK.

   The authentication group is listed on the Authentication Groups tab.

## Completing the Volume List

1. Click OK on the New Volume List window when you are finished.

   Go to the Volume Lists tab of the management group to see the new list displayed.

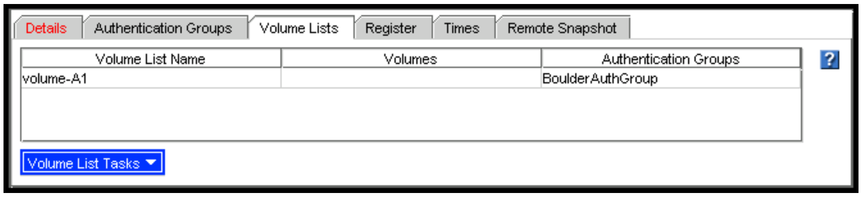

| Volume List Name | Volumes | Authentication Groups |
|---|---|---|
| volume-A1 | | BoulderAuthGroup |

**Figure 186. Viewing the New Volume List**

# Editing a Volume List

Edit the volumes and snapshots in a volume list to:

- Add a volume
- Edit the permissions for a volume
- Remove a volume

Edit the authentication groups in a volume list to:

- Add a group to the list
- Remove a group from the list

***Warning:*** *Before editing the volume list, stop any applications from accessing volumes for which you are restricting permissions or removing authentication groups.*

## Opening the Volume List to Edit

1. Log in to the management group that contains the volume list you want to edit.
2. Select the Volume Lists tab.
3. Select the volume list to edit.

4. Click Volume Lists Tasks and select Edit Volume List.

   The Edit Volume List window opens, shown in Figure 187.

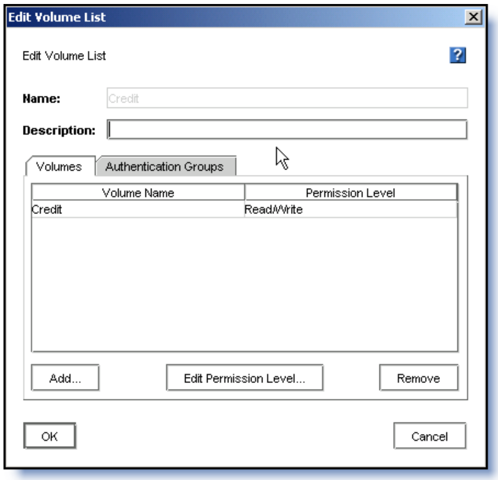

Name: Credit

Description:

| Volume Name | Permission Level |
|---|---|
| Credit | Read/Write |

Figure 187. Editing a Volume List

5. Select either the Volumes tab or the Authentication Groups tab and make the necessary changes.
6. Click OK when you are finished.

# Editing Volume Permission Levels

Changing the permission level for a volume changes the access rights for the authentication groups that connect to that volume.

***Warning:*** *Before editing the volume list, stop any applications from accessing volumes for which you are restricting permissions.*

1. In the Edit Volume List window, select the volume or snapshot for which you want to edit the permissions.
2. Click Edit Permission Level.

   The Edit Volume in Volume List window opens, shown in Figure 188.

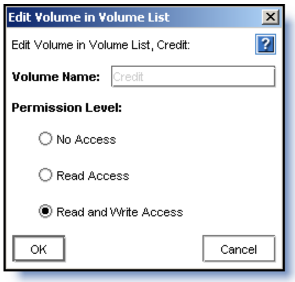

Edit Volume in Volume List, Credit:

Volume Name: Credit

Permission Level:

○ No Access

○ Read Access

● Read and Write Access

Figure 188. Editing Permissions on a Volume in a Volume List

3. Change the permission level.
4. Click OK when you are finished.
5. Click OK in the Edit Volume List window when you are finished editing the volume list.
6. In the volume tab window, select the Volume Lists Membership tab, shown in Figure 189, to review your changes.

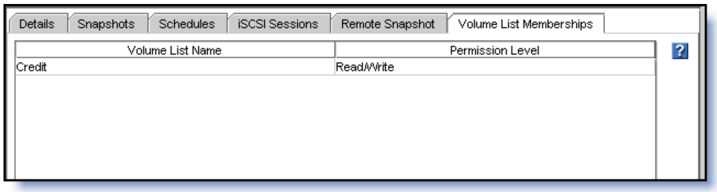

| Volume List Name | Permission Level |
|---|---|
| Credit | Read/Write |

**Figure 189. Confirming Permission Level Changes for a Volume List**

## Changing Authentication Groups in a Volume List

Removing authentication groups from a volume list has these effects:

- Prevents the associated application servers from accessing the volumes in that list
- Is a prerequisite for deleting the volume list

***Warning:*** *Before editing the volume list, stop any applications from accessing volumes for which you are restricting permissions or removing authentication groups.*

1. In the Edit Volume List window, shown in Figure 187, select the Authentication Group tab.
2. Add or remove authentication groups as required.
3. Click OK when you are finished.

   In the management group tab window, select the Volume Lists tab, shown in Figure 190, to review your changes.

## Removing a Volume from a Volume List

Remove volumes and snapshots from a volume list in preparation for deleting the volume list.

***Warning:*** *Before deleting volumes or snapshots from the volume list, stop any applications from accessing those volumes/snapshots you are removing.*

1. Open the Edit Volume List window, shown in Figure 187.
2. Select the volumes and snapshots you want to remove from the volume list.
3. Click Remove.
4. Click OK when you are finished.

5. In the management group tab window, select the Volume Lists tab, shown in Figure 190, to review your changes.

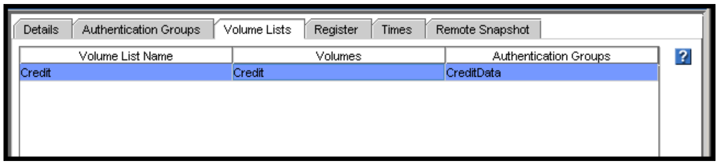

| Volume List Name | Volumes | Authentication Groups |
|---|---|---|
| Credit | Credit | CreditData |

Figure 190. Viewing the Edited Volume List

# Deleting a Volume List

Deleting a volume list removes that list from the management group.

**Prerequisites**

- Remove all volumes and snapshots from the volume list.
- Remove all authentication groups from the volume list.

**Warning:** *Before you delete access to a volume or snapshot, close the host server's Disk Management (or equivalent) window.*

**Note:** *To prevent a group from accessing a volume without deleting the volume list, change the volume permissions to "No Access." See "Editing Volume Permission Levels".*

1. Log into the appropriate management group.
2. Select the Volume Lists tab, shown in Figure 191.

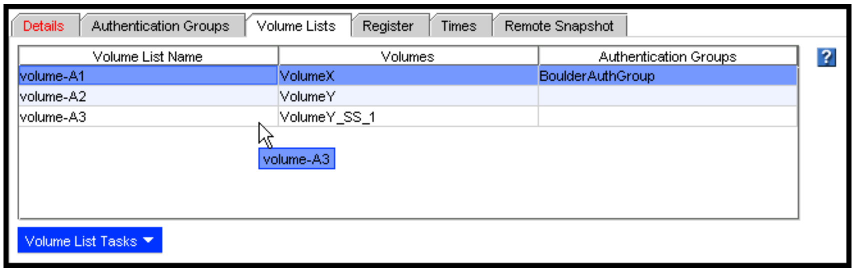

| Volume List Name | Volumes | Authentication Groups |
|---|---|---|
| volume-A1 | VolumeX | BoulderAuthGroup |
| volume-A2 | VolumeY | |
| volume-A3 | VolumeY_SS_1 | |

Figure 191. Viewing the Volume Lists Tab

3. Select the volume lists you want to remove.
4. Click Volume Lists Tasks and select Delete Volume List.

   A confirmation window opens.
5. Click OK to confirm the deletion. In the management group tab window, the Volume Lists tab no longer displays these volume lists.

# 16 Feature Registration

## Add-On Features and Applications Registration Overview

Add-on features and applications expand the capabilities of the Storage System Software. Add-on features and applications include the following:

- Scalability Pak
- Configurable Snapshot Pak
- Remote Data Protection Pak
- Client Server Clustering Pak

All add-on features and applications are available when you begin using the Storage System Software. If you begin using an add-on feature or application without first registering, a 30-day evaluation period begins. Throughout the evaluation period you receive reminders to register and purchase a license for the add-on features and applications you want to continue using.

## Evaluating Features

Add-on features and applications are active and available when you install and configure your system.

### 30-Day Evaluation Period

When you use any feature that requires registration, a message opens, shown in Figure 192, asking you to verify that you want to enter a 30-day evaluation period.

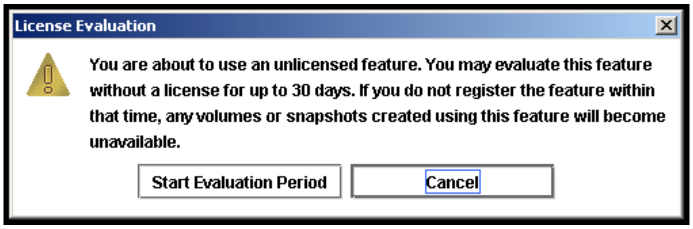

License Evaluation

You are about to use an unlicensed feature. You may evaluate this feature without a license for up to 30 days. If you do not register the feature within that time, any volumes or snapshots created using this feature will become unavailable.

Figure 192. Verifying the Start of the 30-day Evaluation Period

During this evaluation period you may configure, test, and modify any feature. At the end of the 30-day evaluation period, if you do not register and obtain a license key, then all volumes and snapshots associated with the feature or application become unavailable to any clients. The data is safe and you can manage the volumes and snapshots in the Console. Also, the entire configuration can be restored to availability when a license key is obtained and applied to the SSMs in the management group containing the configured features.

***Note:*** *If you know you are not going to purchase the feature, plan to remove any volumes and snapshots created by using the feature before the end of the 30-day evaluation period.*

## Tracking the Time Remaining in the Evaluation Period

Track the time left on your 30-day evaluation period by using either the management group Register tab, shown in Figure 193 or the reminder notices that open periodically, as shown in Figure 194.

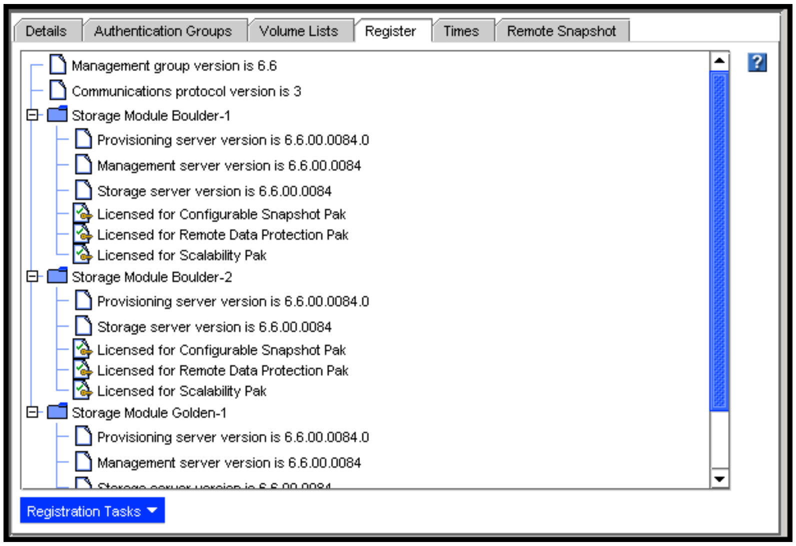

Management group version is 6.6
Communications protocol version is 3
Storage Module Boulder-1
Provisioning server version is 6.6.00.0084.0
Management server version is 6.6.00.0084
Storage server version is 6.6.00.0084
Licensed for Configurable Snapshot Pak
Licensed for Remote Data Protection Pak
Licensed for Scalability Pak
Storage Module Boulder-2
Provisioning server version is 6.6.00.0084.0
Storage server version is 6.6.00.0084
Licensed for Configurable Snapshot Pak
Licensed for Remote Data Protection Pak
Licensed for Scalability Pak
Storage Module Golden-1
Provisioning server version is 6.6.00.0084.0
Management server version is 6.6.00.0084
Storage server version is 6.6.00.0084

### Figure 193. Evaluation Period Countdown on Register Tab

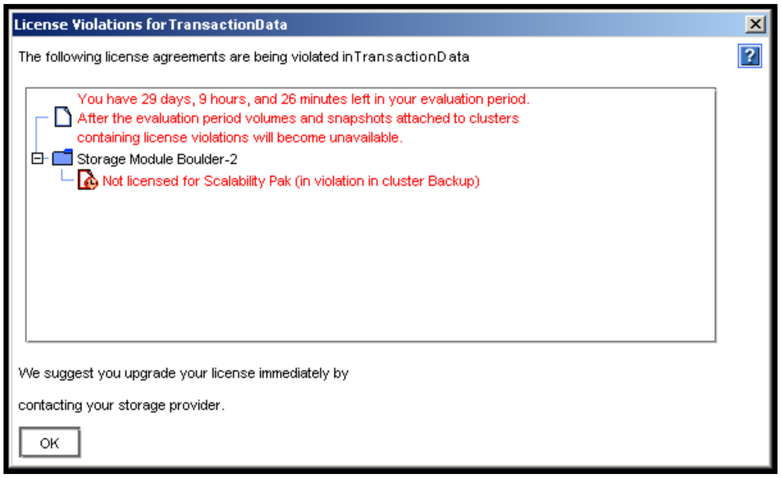

The following license agreements are being violated in TransactionData

You have 29 days, 9 hours, and 26 minutes left in your evaluation period.
After the evaluation period volumes and snapshots attached to clusters
containing license violations will become unavailable.

Storage Module Boulder-2

Not licensed for Scalability Pak (in violation in cluster Backup)

We suggest you upgrade your license immediately by

contacting your storage provider.

**Figure 194. Evaluation Period Countdown Message**

## Viewing Licensing Icons

Icons indicate the status of licensing on individual modules. Figure 195 illustrates the licensing icons. Note that the violation icon displays throughout the 30-day evaluation period.

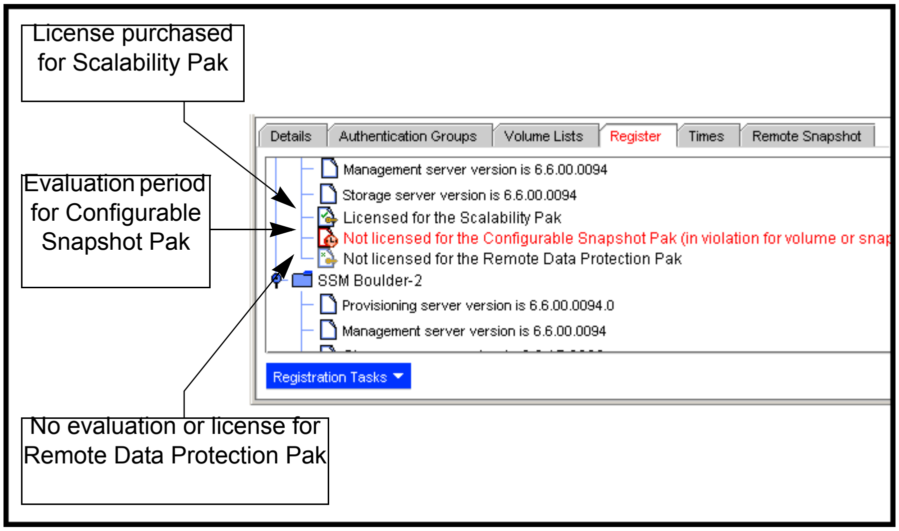
License purchased for Scalability Pak
Evaluation period for Configurable Snapshot Pak
No evaluation or license for Remote Data Protection Pak
Details
Authentication Groups
Volume Lists
Register
Times
Remote Snapshot
Registration Tasks

Management server version is 6.6.00.0094

Storage server version is 6.6.00.0094

Licensed for the Scalability Pak

Not licensed for the Configurable Snapshot Pak (in violation for volume or snap

Not licensed for the Remote Data Protection Pak

SSM Boulder-2

Provisioning server version is 6.6.00.0094.0

Management server version is 6.6.00.0094

**Figure 195. Icons Indicating License Status for Features**

# Evaluating the Scalability Pak

The Scalability Pak includes the following features:

- Multiple nodes in a cluster
- N-way replication
- Virtual manager

## Starting the License Evaluation Period

If you put more than one SSM into a cluster, the 30-day license evaluation begins. During the 30-day evaluation period you can create volumes with 2- or 3-way replication, or configure a virtual manager. Please read "Planning Data Replication" on page 219 and Chapter 10, "Disaster Recovery Using A Virtual Manager" before working with these features.

## Backing Out of the License Evaluation Period

If you decide not to purchase the Scalability Pak, you must reduce the cluster to one SSM. The features you are evaluating dictate the steps required to safely back out of the evaluation configuration, particularly if you want to save any volumes or snapshots in the test configuration.

1. Back up any volumes you plan to retain. The table below describes additional steps to safely back out of the Scalability Pak evaluation.

### Table 54. Safely Backing out of Scalability Pak Evaluation

| Feature Being Evaluated | Steps to Back Out |
|---|---|
| Multiple SSMs with a large volume | If volume is too large to fit on a cluster with one SSM, do one of the following:<br>• Delete the volume<br>• Move the volume to another single note cluster with adequate capacity<br>• Add storage to the SSM |
| 2- or 3-way replication | Set volume replication level to none |
| Virtual manager | Stop virtual manager |

2. Remove the extra SSMs from the cluster.

# Evaluating the Configurable Snapshot Pak

The Configurable Snapshot Pak includes programmable snapshots. Features included are

- Scheduled snapshots
- Scripting for snapshots

## Starting the License Evaluation Period

The Configurable Snapshot Pak 30-day evaluation period begins if you create a snapshot schedule.

## Backing Out of the License Evaluation Period

If you decide not to purchase the Configurable Snapshot Pak, you must delete any snapshot schedules that you have configured.

1. Back up any volumes you plan to retain. Table 55 describes how to safely back out of the Configurable Snapshot Pak evaluation.

**Table 55. Safely Backing out of Configurable Snapshot Pak Evaluation**

| Feature Being Evaluated | Steps to Back Out |
|---|---|
| Scheduled snapshots | Delete the snapshot schedule |

# Evaluating the Remote Data Protection Pak

The Remote Data Protection Pak includes Remote Copy. Features included are

- Remote volumes
- Remote snapshots
- Remote snapshot schedules
- Scripting for remote copy

## Starting the License Evaluation Period

The Remote Data Protection Pak 30-day evaluation period begins if you create a remote volume by

- Making an existing primary volume into a remote volume
- Creating a remote volume in the process of creating a remote snapshot

- Creating a new volume and selecting the "Remote" radio button on the New Volume dialog

When a remote volume is created, the license evaluation period begins on both the primary and remote SSMs. For example, suppose the primary volume is on Cluster 1. You create a remote snapshot of that primary volume to Cluster 2. SSMs in both clusters show the clock ticking for the license evaluation period.

Read the *Remote IP Copy User Manual* before working with these features.

## Backing Out of the License Evaluation Period

If you decide not to purchase the Remote Data Protection Pak, you must delete any remote volumes you have configured. The features you are evaluating dictate the steps required to safely back out of the evaluation configuration, particularly if you want to save any volumes or snapshots in the test configuration.

1. Back up any volumes you plan to retain. Table 56  describes additional steps to safely back out of the Remote Data Protection Pak evaluation.

### Table 56. Safely Backing Out of Remote Data Protection Pak Evaluation

| Feature Being Evaluated | Steps to Back Out |
|---|---|
| Remote snapshots - removing data from the remote target | • Delete any remote snapshots<br>• Delete the remote volume |
| Remote snapshots - retaining the data on the remote target | • Make the remote volume into a primary volume<br>• Disassociate the primary and remote management groups, if the remote copy was between management groups. |

# Scripting Evaluation

Application-based scripting is available for volume and snapshot features as part of the Configurable Snapshot Pak and the Remote Data Protection Pak. Features that can be scripted include

- Creating snapshots and setting hard and soft snapshot thresholds
- Increasing volume hard and soft thresholds
- Scripting automatic threshold increases
- Creating remote volumes and snapshots

Because using scripts with add-on features and applications starts the 30-day evaluation period without requiring you to use the Console, you must first verify that you are aware of starting the 30-day evaluation clock when using scripting. If you do not enable the scripting evaluation period, any scripts you have running (licensed or not) will fail.

## Turn On Scripting Evaluation

To use scripting while evaluating add-on features or applications, enable the scripting evaluation period.

1. In the navigation window, select the management group.
2. Select the Register tab.
3. Click Registration Tasks and select Feature Registrations, shown in Figure 196.

   The Feature Registration window appears.
4. Select the Scripting Evaluation tab, shown in Figure 196.

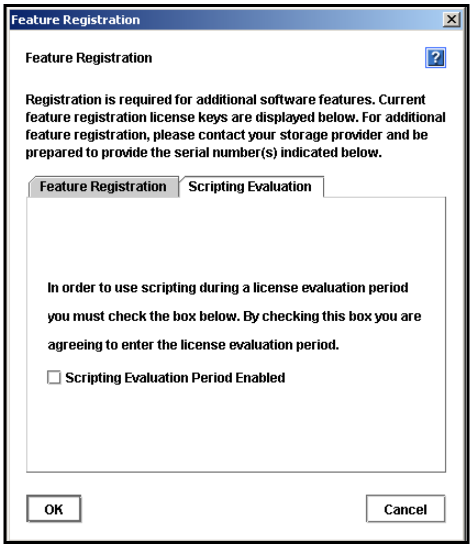
Feature Registration

Feature Registration

Registration is required for additional software features. Current feature registration license keys are displayed below. For additional feature registration, please contact your storage provider and be prepared to provide the serial number(s) indicated below.

In order to use scripting during a license evaluation period you must check the box below. By checking this box you are agreeing to enter the license evaluation period.

□ Scripting Evaluation Period Enabled

Figure 196. Enabling Scripting Evaluation

5. Check the box to enable the use of scripts during a license evaluation period.
6. Click OK.

For more information about scripting, see Chapter 14, “Working with Scripting.”

## Turn Off Scripting Evaluation

The scripting evaluation period is turned off when

- You purchase the add-on feature or application you were evaluating, or
- You complete the evaluation and decide not to purchase any add-on features or applications.

1. Select the management group.
2. Select the Register tab.
3. Click Registration Tasks and select Feature Registration.

4. Select the Scripting Evaluation tab, shown in Figure 196.
5. Clear the check box.
6. Click OK. Table 57 describes additional steps to safely back out of the scripting evaluation.

### Table 57. Safely Backing Out of Scripting Evaluation

| Feature Being Evaluated | Steps to Back Out |
|---|---|
| Any of the items below that are created by an application-based script<br>• Remote copy volumes and snapshots | 1. Back out of any remote copying<br>2. Delete any scripts<br>3. Delete any primary or remote snapshots created by the scripts. You can identify these snapshots by viewing the item  “Created By Script” on the snapshot Details tab. |

**Note:** *Turning off the scripting evaluation ensures that no scripts will continue to push the 30-day evaluation clock.*

# Registering Features and Applications

When registering SSMs for add-on features and applications, you first submit the appropriate SSM serial number(s) to purchase the license key(s). You will then receive the license key(s) to apply to the SSM(s).

## Using License Keys

License keys are assigned to individual SSMs. License keys can be added to SSMs before or after they are in a management group. One license key is issued per SSM and that key licenses all the features requested for that SSM. Therefore, you register each SSM for which you want to use add-on features and applications.

For example, if you wanted to configure multiple node clusters in two locations to use with, you would license the SSMs in both the primary location and the remote location for both the Scalability Pak and the Remote Data Protection Pak.

**Note:** *If you remove the SSM from the management group, the license key remains with that SSM. See the chapter on “Working with Management Groups” on page 169 for more information about removing SSMs from a management group.*

# Registering Available SSMs for License Keys

SSMs that are not in a management group are licensed individually in the module configuration category, see “Registering Features for an SSM” on page 48.

# Registering SSMs in a Management Group

SSMs that are in a management group are licensed through the management group.

## Submitting SSM Serial Numbers

First you must submit the feature keys of all the SSMs that you want to register.

1. In the navigation window, select the management group for which you want to register features.
2. Select the Register tab, shown in Figure 197. The Register tab lists what, if any, licenses have been purchased. If you are evaluating features, the time remaining in the evaluation period is listed on the tab as well.

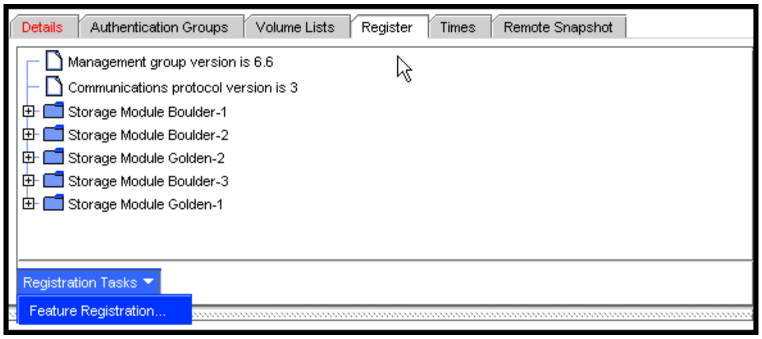

Management group version is 6.6
Communications protocol version is 3
Storage Module Boulder-1
Storage Module Boulder-2
Storage Module Golden-2
Storage Module Boulder-3
Storage Module Golden-1

**Figure 197. Registering Features**

3. Click Registration Tasks and select Feature Registration. The Feature Registration window opens, shown in Figure 198. Listed are all the SSMs in that management group.

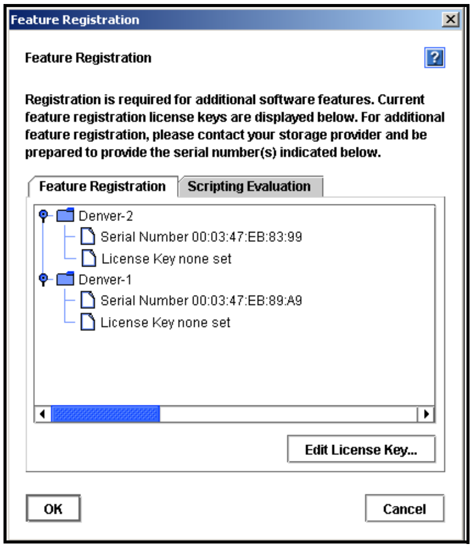

Feature Registration

Registration is required for additional software features. Current feature registration license keys are displayed below. For additional feature registration, please contact your storage provider and be prepared to provide the serial number(s) indicated below.

Denver-2
Serial Number 00:03:47:EB:83:99
License Key none set
Denver-1
Serial Number 00:03:47:EB:89:A9
License Key none set

Figure 198. Opening the Feature Registration Window

4. For each SSM listed in the window that you want to register, submit the feature key as instructed in the Feature Registration window.

   Control + C copies the feature key so that you can paste it into an application such as Notepad or Word.

***Note:*** *Record the host name or IP address of the SSM along with the feature key. This record will make it easier to add the license key to the correct SSM when you receive it.*

## Entering License Keys

When you receive the license keys add them to the SSMs in the Feature Registration window.

1. In the navigation window, select the management group.
2. Select the Register tab.
3. Click Registration Tasks and select Feature Registration.

4. Select an SSM and click Edit License Key. The Edit Feature Registration window opens, shown in Figure 199.

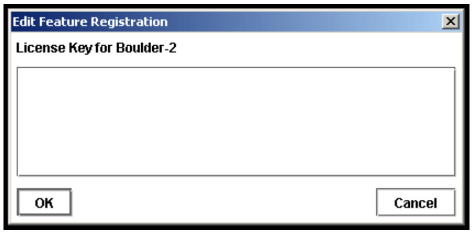

**Figure 199. Entering License Key**

5. Copy and paste the appropriate license key for that SSM into the window.

6. Click OK. The license key information is updated in the Feature Registration window, as shown in Figure 200.

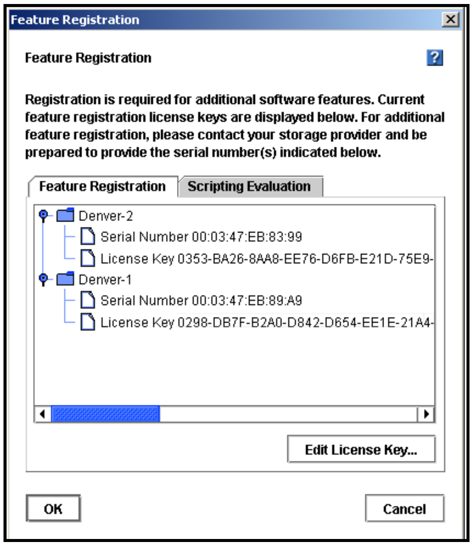

Registration is required for additional software features. Current feature registration license keys are displayed below. For additional feature registration, please contact your storage provider and be prepared to provide the serial number(s) indicated below.

Denver-2
Serial Number 00:03:47:EB:83:99
License Key 0353-BA26-8AA8-EE76-D6FB-E21D-75E9-
Denver-1
Serial Number 00:03:47:EB:89:A9
License Key 0298-DB7F-B2A0-D842-D654-EE1E-21A4-

**Figure 200. Viewing License Keys**

# A Using the Configuration Interface

The Configuration Interface is the command line interface that uses a direct connection with the SSM.

You may need to access the Configuration Interface if all network connections to the SSM are disabled. Use the Configuration Interface to

- Add SSM administrators and change passwords
- Access and configure network interfaces
- Delete a NIC bond
- Set the TCP speed and duplex
- Edit the frame size
- Reset the SSM configuration to factory defaults

## Connecting to the Configuration Interface

Accessing the Configuration Interface is accomplished by attaching a PC or a laptop to the SSM using a null modem cable and connecting to the Configuration Interface with a terminal emulation program.

### Connecting to the Configuration Interface with Windows

On the PC or laptop attached directly to the SSM with a null modem cable, open a session with a terminal emulation program such as HyperTerminal or ProComm Plus.

Use the following settings.

- Bits per second = 19200
- Data bits = 8
- Parity = None
- Stop bits = 1
- Flow control = None
- Backspace key sends = Del
- Emulation = ANSI

When the session is established, the Configuration Interface window opens, shown in Figure 201.

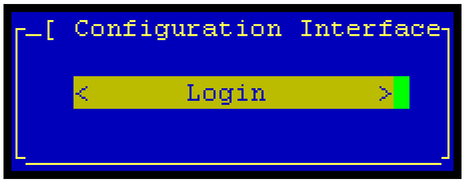

Figure 201. Opening the Configuration Interface

## Connecting to the Configuration Interface with Linux/UNIX

If using Linux, create the following configuration file. You must create the file as root, or root must change permissions for */dev/cua0* in order to create the config file in */etc/*.

1. Create the */etc/minirc.SSM* with the following parameters:

— # Begin SSM configuration

— # Machine-generated file – use “minicom –s” to

— # change parameters

^ pr port = /dev/cua0

^ pu baudrate = 19200

^ pu bits = 8

^ pu parity = N

^ pu stopbits = 1

^ pu mautobaud = Yes

^ pu backspace = DEL

^ pu hasdcd = No

^ pu rtscts = No

^ pu xonxoff = Yes

^ pu askdndir = Yes

— # End SSM configuration

2. Start xterm as follows: $ xterm
3. In the xterm window, start minicom as follows: $ minicom -c on -l Storage System Module
4. Press Enter when the terminal emulation session is established. A prompt appears asking you to type “start” and hit enter at the login prompt.
5. Type start and press Enter. When the session is connected to the SSM, the Configuration Interface window opens, shown in Figure 201.

# Logging in to the SSM

Once you have established a connection to the SSM using a terminal emulation program, log in to the Configuration Interface.

1. From the Configuration Interface entry window, press Enter to start the log in process. The Configuration Interface Login window opens, shown in Figure 202.

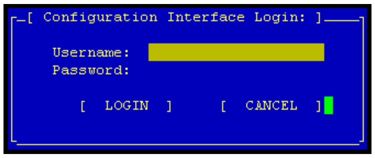

**Figure 202. Enter User Name and Password**

2. Type the user name and password of the administrative user established when the SSM was first configured.

***Note:*** *This user is viewable in the Storage System Console under SSM Administration. Click Users and find the admin user on the list.*

1. Tab to Login and press Enter. The Configuration Interface main menu opens, shown in Figure 203.

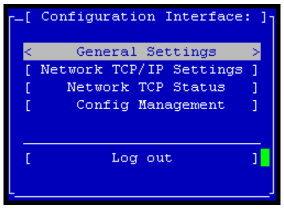

**Figure 203. Configuration Interface Main Menu**

# Configuring Administrative Users

Use the Configuration Interface to add new administrative users or to change administrative passwords. You can only change the password for the administrative user that you used to log in to the Configuration Interface.

1. On the Configuration Interface main menu, tab to General Settings and press Enter. The General window opens, shown in Figure 204.

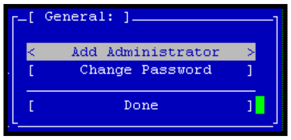

**Figure 204. General Settings Window**

2. To add an administrative user, tab to Add Administrator and press Enter. Then enter the new user’s name and password. Confirm password, tab to Ok and press Enter.
3. To change the password for the user that you are currently logged in as, tab to Change Password and press Enter. Then enter the new password. Confirm password, tab to Ok and press Enter.
4. On the General window, tab to Done and press Enter.

# Configuring a Network Connection

The SSM has two 1000BASE-T (Gigabit Ethernet) NICs in its motherboard. These interfaces are named Motherboard:Port0 and Motherboard:Port1. In addition, the SSM can include multiple add-on PCI cards, each with up to 4 interfaces. These add-on interfaces are named according to the card’s slot and the port number, such as Slot1:Port0. For information about how the interfaces are labelled on the back of the SSM, see “Configuring the IP Address Manually” on page 89.

Once you have established a connection to the SSM using a terminal emulation program, you can configure an interface connection using the Configuration Interface.

1. On the Configuration Interface main menu, tab to Network TCP/IP Settings and press Enter. The Available Network Devices window opens, shown in Figure 205.

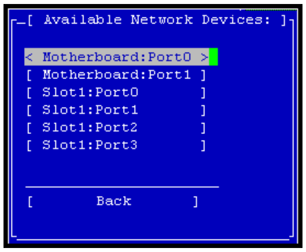

Figure 205. Selecting an Interface to Configure

2. Tab to select the network interface that you want to configure and press Enter. The Network Settings window opens, shown in Figure 206.

   If the interface you selected is a bond, then the Logical Interface Device window displays first. Click Change Settings to open the Network Settings window for the bond.

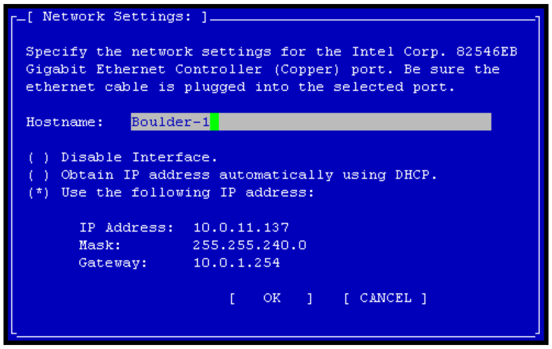

```
[ Network Settings: ]

Specify the network settings for the Intel Corp. 82546EB
Gigabit Ethernet Controller (Copper) port. Be sure the
ethernet cable is plugged into the selected port.

Hostname:    Boulder-1

( ) Disable Interface.
( ) Obtain IP address automatically using DHCP.
(*) Use the following IP address:

      IP Address:  10.0.11.137
      Mask:        255.255.240.0
      Gateway:     10.0.1.254

                        [   OK   ]  [ CANCEL ]
```

**Figure 206. Entering the Host Name and Settings for an Interface**

3. Enter the host name and tab to the next section to configure the network settings.

**Note:** *If you specify an IP address, the Gateway is a required field. If you do not have a Gateway, enter 0.0.0.0 for the Gateway address.*

4. Tab to OK and press Enter to complete the network configuration.

   A second window opens, asking you to confirm the changes.

5. Press Enter. Return to the Storage System Console and locate the SSM using the Find menu to search by subnet and mask, or search by entering the SSM IP address.

# Deleting a NIC Bond

You can delete two types of NIC bonds using the Configuration Interface:

- Active backup bond
- NIC aggregation bond

For more information about creating and configuring NIC aggregation and active backup bonds, see “Configuring NIC Bonding” on page 90.

When you delete an active backup bond, the primary interface assumes the IP address and configuration of the deleted logical interface. The other NIC is disabled and its IP address is set to 0.0.0.0.

When you delete a NIC aggregation bond, one of the active interfaces in the bond retains the IP address of the deleted logical interface. The other NIC is disabled and its IP address is set to 0.0.0.0.

1. On the Configuration Interface main menu, tab to Network TCP/IP Settings and press Enter. The Available Network Devices window opens, shown in Figure 207. The logical bond is listed in the window.

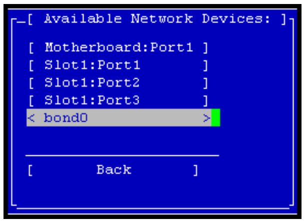

**Figure 207. Selecting a Bonded Interface in the Available Network Devices Window**

2. Tab to select the bond and press Enter. The Logical Failover Device window opens, shown in Figure 208.

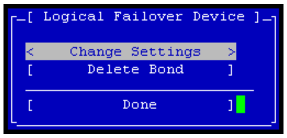

**Figure 208. Deleting a NIC Bond**

3. Tab to Delete Bond and press Enter. A window opens, asking you to confirm the changes.
4. Press Enter.
5. On the Available Network Devices window, tab to Back and press Enter.

# Setting the TCP Speed, Duplex, and Frame Size

You can use the Configuration Interface to set the TCP speed, duplex, and frame size of a network interface.

- **TCP speed and duplex.** You can change the speed and duplex of a 10/100/1000 interface. If you change these settings, you must ensure that BOTH sides of the NIC cable are configured in the same manner. For example, if the SSM is set for Auto/ Auto, the switch must be set the same. For more information about TCP speed and duplex settings, see "Editing the TCP Speed and Duplex" on page 109.
- **Frame size.** The frame size specifies the size of data packets that are transferred over the network. The default Ethernet standard frame size is 1500 bytes. The maximum allowed frame size is 9000 bytes.

Increasing the frame size improves data transfer speed by allowing larger packets to be transferred over the network and by decreasing the CPU processing time required to transfer data. However, increasing the frame size requires that routers, switches, and other devices on your network support that frame size.

For more information about setting a frame size that corresponds to the frame size used by routers, switches, and other devices on your network, see "Editing the NIC Frame Size" on page 110.

1. On the Configuration Interface main menu, tab to Network TCP Status and press Enter. The Available Network Devices window opens, shown in Figure 209.

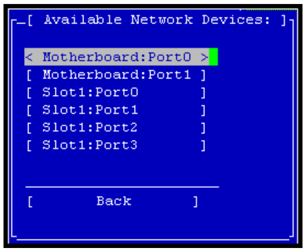

```
[ Available Network Devices: ]

< Motherboard:Port0 >
[ Motherboard:Port1 ]
[ Slot1:Port0       ]
[ Slot1:Port1       ]
[ Slot1:Port2       ]
[ Slot1:Port3       ]

[       Back       ]
```

**Figure 209. Available Network Devices Window**

2. Tab to select the network interface for which you want to set the TCP speed and duplex and press Enter. The Network TCP Status window opens, shown in Figure 210.

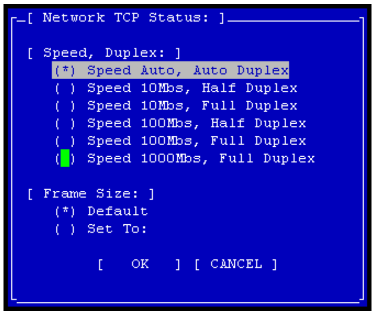

```
[ Network TCP Status: ]

[ Speed, Duplex: ]
   (*) Speed Auto, Auto Duplex
   ( ) Speed 10Mbs, Half Duplex
   ( ) Speed 10Mbs, Full Duplex
   ( ) Speed 100Mbs, Half Duplex
   ( ) Speed 100Mbs, Full Duplex
   ( ) Speed 1000Mbs, Full Duplex

[ Frame Size: ]
   (*) Default
   ( ) Set To:

        [   OK   ] [ CANCEL ]
```

**Figure 210. Setting the Speed, Duplex, and Frame Size**

3. To change the speed and duplex of an interface, tab to a setting in the Speed / Duplex list.
4. To change the frame size, select Set To in the Frame Size list. Then tab to the field to the right of Set To and type a frame size. The frame size value must be between 1500 bytes and 9000 bytes.
5. On the Network TCP Status window, tab to OK and press Enter.
6. On the Available Network Devices window, tab to Back and press Enter.

# Removing an SSM from a Management Group

Removing an SSM from a management group, deletes all data from the SSM, clears all information about the management group from the SSM, and will reboot the SSM.

***Warning:*** *Removing an SSM from a management group deletes all data on the SSM.*

1. On the Configuration Interface main menu, tab to Config Management and press Enter. The Configuration Management window opens, shown in Figure 211.

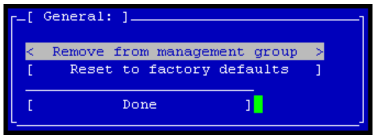
[ General: ]
< Remove from management group >
[ Reset to factory defaults ]
[ Done ]

**Figure 211. Removing the SSM from a Management Group**

2. Tab to Remove from management group and press Enter. A window opens, warning you that removing the SSMfrom the management group will delete all data on the SSM and reboot the SSM.
3. Tab to Ok and press Enter
4. On the Configuration Management window, tab to Done and press Enter.

# Resetting the SSM to Factory Defaults

Resetting the SSM to factory defaults deletes all data and erases the configuration of the SSM, including administrative users and network settings.

***Warning:*** *Resetting the SSM to factory defaults deletes all data on the SSM.*

1. On the Configuration Interface main menu, tab to Config Management and press Enter. The Configuration Management window opens, shown in Figure 212.

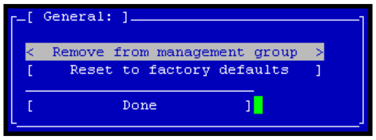
[ General: ]
< Remove from management group >
[ Reset to factory defaults ]
[ Done ]

**Figure 212. Resetting to Factory Defaults**

2. Tab to Reset to factory defaults and press Enter. A window opens, warning you that resetting the SSM configuration will delete all data on the SSM and reboot the SSM.
3. Tab to Ok and press Enter.
4. On the Configuration Management window, tab to Done and press Enter.
5. Use the default User Name “admin” and Password “storage” to log in to the SSM.

# B SNMP MIB Information

## SNMP Agent

The SNMP Agent resides in the Storage System Module. The agent takes SNMP network requests for reading or writing configuration information and translates them into internal system requests. Management Information Base (MIB) files are provided which can enable the system administrator to use their favorite SNMP tool to view or modify configuration information. The SNMP Agent supports versions 1, 2c, and 3 of the protocol. Security can be configured based on the host making the request and a password.

**Note:** *To ensure that all items display properly in your SNMP tool, use version 2c or later of the protocol.*

## Supported MIBs

- MIB II
- Host Resources MIB
- UCD Extensions MIB
- SNMPv3 MIB

## Exceptions

### MIB II

```
system.sysServices
interfaces.ifTable.ifEntry.ifLastChange
interfaces.ifTable.ifEntry.ifInNUcastPkts
interfaces.ifTable.ifEntry.ifInDiscards
interfaces.ifTable.ifEntry.ifInUnknownProtos
interfaces.ifTable.ifEntry.ifOutNUcastPkts
ip.ipAddrTable.ipAddrEntry.ipAdEntReasmMaxSize
ip.ipRouteTable.ipRouteEntry.ipRouteMetric2
ip.ipRouteTable.ipRouteEntry.ipRouteMetric3
ip.ipRouteTable.ipRouteEntry.ipRouteMetric4
ip.ipRouteTable.ipRouteEntry.ipRouteAge
ip.ipRouteTable.ipRouteEntry.ipRouteMetric5
ip.ipForward (MIB Tree)
tcp.tcpInErrs
tcp.tcpOutRsts
tcp.ipv6TcpConnTable (MIB Tree)
```

```
udp.ipv6UdpTable (MIB Tree)
egp (MIB Tree)
transmission (MIB Tree)
snmp.snmpSilentDrops
snmp.snmpProxyDrops
rmon (MIB Tree)
application (MIB Tree)
mta (MIB Tree)
ipv6MIB (MIB Tree)
schedMIB (MIB Tree)
scriptMIB (MIB Tree)
agentxMIB (MIB Tree)
ifInvertedStackMIB (MIB Tree)
```

## Host Resources MIB

```
host.hrDevice.hrDeviceTable.hrDeviceEntry.hr DeviceStatus
host.hrDevice.hrDeviceTable.hrDeviceEntry.hr DeviceErrors
host.hrDevice.hrProcessorTable.hr ProcessorEntry.hrProcessorLoad
host.hrDevice.hrPrinterTable (MIB Tree)
host.hrSWRun.hrSWOSIndex
host.hrSWInstalled (MIB Tree)
host.hrMIBAdminInfo (MIB Tree)
```

## UCD Extensions MIB

```
ucdavis.processes (MIB Tree)
ucdavis.prTable (MIB Tree)
ucdavis.extensible (MIB Tree)
ucdavis.memory.memTotalSwapTXT
ucdavis.memory.memAvailSwapTXT
ucdavis.memory.memTotalRealTXT
ucdavis.memory.memAvailRealTXT
ucdavis.disk (MIB Tree)
ucdavis.loadaves (MIB Tree)
ucdavis.extTable (MIB Tree)
ucdavis.dskTable (MIB Tree)
ucdavis.systemStats.ssCpuRawWait
ucdavis.systemStats.ssCpuRawKernel
ucdavis.systemStats.ssCpuRawInterrupt
ucdavis.systemStats.ssIORawSent
ucdavis.systemStats.ssIORawReceived
ucdavis.systemStats.ssRawInterrupts
ucdavis.systemStats.ssRawContexts
ucdavis.ucdExperimental (MIB Tree)
ucdavis.fileTable (MIB Tree)
```

## SNMPv3 MIB

```
snmpModules.snmpTargetMIB (MIB Tree)
snmpModules.snmpNotificationMIB (MIB Tree)
snmpModules.snmpProxyMIB (MIB Tree)
snmpModules.snmpUsmMIB.usm MIBObjects.usmUser.usm UserTable (MIB Tree)
snmpModules.snmpVacmMIB.vacm MIBObjects.vacmContextTable (MIB Tree)
snmpModules.snmpCommunityMIB (MIB Tree)
```

# C Understanding and Planning Remote Copy

## Remote Copy Overview

Remote Copy provides a powerful and flexible method for replicating data and keeping that replicated data available for disaster recovery, business continuance, backup and recovery, data migration, and data mining.

## How Remote Copy Works

Remote Copy is a feature upgrade. You must purchase a Remote Data Protection Pak license to use Remote Copy beyond the 30-day evaluation period. You must purchase a license for each Storage System Module in a cluster that contains a primary volume or a remote volume. For information about registering Remote Copy licenses, see "Feature Registration" on page 285.

Remote Copy uses the existing volume and snapshot features along with replication across geographic distances to create remote snapshots. The geographic distance can be local (in the same data center or on the same campus), metro (in the same city), or long distance. (cross-country, global).

For example, the accounting department in the corporate headquarters in Chicago runs the corporate accounting application and stores the resulting data. The designated backup site is in Des Moines. Nightly at 11:00 p.m., accounting updates are replicated to the Des Moines backup facility using Remote Copy.

## Glossary for Remote Copy

The following terminology is used in describing the components and processes involved in Remote Copy.

### Table 58. Remote Copy Glossary

| Term | Definition |
|---|---|
| Primary Volume | The volume which is being accessed by the application server. The primary volume is the volume that is backed up with Remote Copy. |
| Primary Snapshot | A snapshot of the primary volume which is created in the process of creating a remote snapshot. The primary snapshot is located on the same cluster as the primary volume. |

Table 58. Remote Copy Glossary

| Term | Definition |
|---|---|
| Remote Volume | The volume that resides in the Remote Copy location where the remote snapshots are created. The remote volume contains no data. It acts as a pointer to tell the system where to make the copy of the primary snapshot. The remote volume can be stored on the same cluster or a different cluster than the primary volume. |
| Remote Snapshot | An identical copy of a primary snapshot. The remote snapshot is located on the same cluster as the remote volume. |
| Remote Copy Pair | The primary volume and its associated remote volume. |
| Failover | The process by which the user transfers operation of the application server over to the remote volume. This can be a manual operation or it can be scripted. |
| Acting Primary Volume | The remote volume, when it assumes the role of the primary volume in a failover scenario. |
| Failback | After failover, the process by which the user restores the primary volume and turns the acting primary back into a remote volume. |
| Failover Recovery | After failover, the process by which the user chooses to fail back to the primary volume or to make the acting primary into a permanent primary volume. |
| Synchronize | The process of copying the most recent snapshot from the primary volume to a new remote snapshot. On failback, synchronization is the process of copying the most recent remote snapshot back to the primary volume. The Console displays the progress of this synchronization. |

# How Remote Copy Works

Replicating data using Remote Copy follows a three-step process.

1. At the production location, you create a snapshot of the primary volume — this is called the primary snapshot.
2. You create a remote volume at the remote location and then create a remote snapshot. The remote snapshot is a snapshot of the empty remote volume, and it is linked to the primary snapshot.

3. The system copies data from the primary snapshot to the remote snapshot.

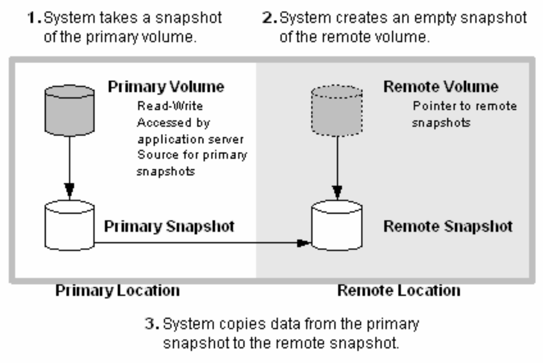
1. System takes a snapshot of the primary volume.
2. System creates an empty snapshot of the remote volume.
Primary Volume
Read-Write
Accessed by application server
Source for primary snapshots
Remote Volume
Pointer to remote snapshots
Primary Snapshot
Remote Snapshot
Primary Location
Remote Location
3. System copies data from the primary snapshot to the remote snapshot.

Figure 213. Basic flow of Remote Copy

***Note:*** *Both primary and completed remote snapshots are the same as regular snapshots.*

***Note:*** *Remote Copy can be used on the same site, even in the same management group and cluster.*

# Graphical Representations of Remote Copy

The Storage System Console displays special graphical representations of Remote Copy.

**Copying the Primary Snapshot to the Remote Snapshot**

When the primary snapshot is copying to the remote snapshot, the Console depicts the process with a moving graphic of pages from the primary to the remote snapshot, as illustrated in Figure 214. The pages move in the direction of the data flow from primary to remote snapshot.

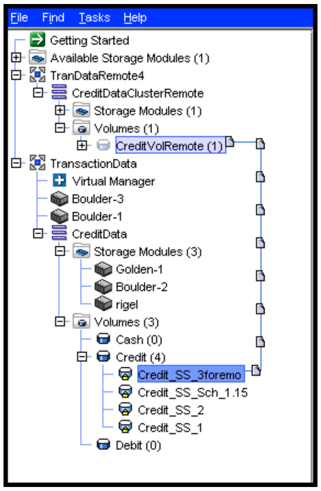

## Figure 214. Icons Depicting the Primary Snapshot Copying to the Remote Snapshot

### Graphical Legend for Remote Copy Icons

The graphical legend available from the Help menu depicts the icons associated with Remote Copy. Figure 215 displays the Remote Copy states icons from the graphical legend.

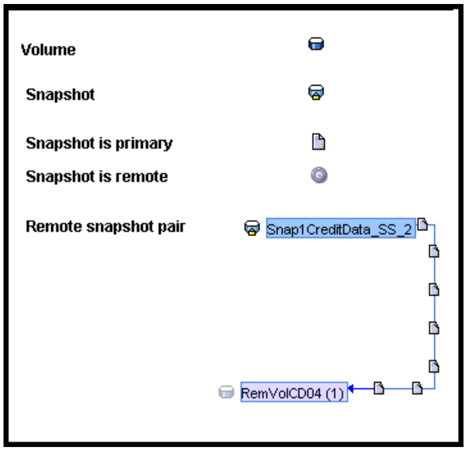

## Figure 215. Icons for Remote Copy in the Graphical Legends Window

# Remote Copy and Volume Replication

Remote Copy is asynchronous replication of data. Volume replication is synchronous replication. Using synchronous volume replication on multiple storage modules within a cluster in combination with asynchronous Remote Copy on a different cluster of storage modules creates a robust, high-availability configuration.

## Uses for Remote Copy

Examine Table 59 to see common applications for the Remote Copy application.

Table 59. Uses for Remote Copy

| Use Remote Copy for | How It Works |
|---|---|
| • Business continuance/ disaster recovery | Using Remote Copy, store remote snapshots on a machine geographically separate. The remote snapshots remain available in the event of a site or system failure at the primary site. |
| • Off-site backup and recovery | Remote Copy eliminates the backup window on an application server by creating remote snapshots on a backup server, either local or remote, and backing up from that server. |
| • Split mirror, data migration, content distribution | Using Remote Copy, make a complete copy of one or more volumes without interrupting access to the original volumes. Move the copy of the volume to the location where it is needed. |
| • Volume clone | Using Remote Copy, create clones or copies of the original volume for use by other application servers. |

## Benefits of Remote Copy

- Remote Copy maintains the primary volume’s availability to application servers. Snapshots on the primary volume are taken instantaneously, and are then copied to remote snapshots in the off-site location.
- Remote Copy operates at the block level, moving large amounts of data much more quickly than file system copying.
- Snapshots are incremental, that is, snapshots save only those changes in the volume since the last snapshot was created. Hence, failover recovery may need to resynchronize only the latest changes rather than the entire volume.
- Remote Copy is robust. If the network link goes down during the process, copying resumes where it left off when the link is restored.

# Planning for Remote Copy

Remote Copy works at the management group, cluster, volume, snapshot, and storage module level. Examine Table 60 for common configurations at various levels.

Table 60. Remote Copy, Storage System Software, and Storage Modules

| Storage System Level | Remote Copy Configuration |
|---|---|
| Management Groups | • Create remote snapshots in the same management group or in a different management group than the primary volume.<br>• If using different management groups, the remote bandwidth setting of the management group containing the remote volume determines the maximum rate of data transfer to the remote snapshot. |
| Clusters | • Create remote snapshots in the same cluster or in a cluster different from the primary volume. |
| Volumes | • Primary volumes contain the data to be copied to the remote snapshot.<br>• Data is copied to the remote snapshot via the remote volume.<br>• The remote volume is a pointer to the remote snapshot. The remote volume has a size of 0. |
| Snapshots | • After data are copied from the primary snapshot to the remote snapshot, the remote snapshot behaves as a regular snapshot. |
| Storage Modules | • Active monitoring of each storage module notifies you when copies complete or fail. Active monitoring also notifies you if a remote volume or snapshot is made primary or if the status of the connection between management groups containing primary and remote volumes changes. |

# Planning the Remote Snapshot

To create a remote snapshot, meet these prerequisites:

- Log in to both the management group that contains the primary volume and the management group that contains the target cluster where the remote snapshot will be created.
- Designate or create a remote volume in that remote management group.
- Have enough space on the target cluster for the remote snapshot.

## Logging in to Primary and Remote

Log in to both the primary and the remote management groups before you begin.

Alternatively, log in to the remote management group as part of the remote copy procedure.

If you are creating the remote volume and remote snapshot in the same management group as the primary volume, then you only need to log in to that management group.

## Designating or Creating the Remote Volume

Create a remote volume by any of the following methods:

- Make an existing volume into a remote volume.
- Create a new remote volume during creation of a remote snapshot.
- Create a new volume from the cluster Details tab window and then select the Remote radio button on the New Volume window.
- From the menu bar, select Tasks >Volume > New Volume.

For more information about these methods of creating remote volumes, see “Creating a Remote Volume” on page 323.

# Using Schedules for Remote Copy

Scheduled remote snapshots provide fault tolerance for business continuance/disaster recovery and provide a consistent, predictable update of data for remote backup and recovery.

## Planning the Remote Copy Schedule

Planning is critical. All of these issues impact the amount of storage available in the system:

- Recurrence
- Thresholds
- Retention

### Recurrence

How often do you want the snapshots created? The recurrence frequency must account for the amount of time it takes to complete a remote snapshot. For example, if your recurrence schedule is set for a new snapshot every 4 hours you should ensure that the time to copy that snapshot to the remote location is less than 4 hours.

Test the time required for copying a snapshot. One way to check the time required to copy a snapshot is to run a test of the actual process. In the test, you take 2 remote snapshots of the primary volume. Because the first remote snapshot copies the entire volume, it takes longer to copy. The second remote snapshot copies only *changes* made to the volume since the first remote snapshot. Because you create the second remote snapshot after the time interval you intend to schedule, the copy time for the second remote snapshot is more representative of the actual time required for copying subsequent remote snapshots.

1. Create a remote snapshot of the primary volume.
2. Wait for the copy to finish.

3. Create another remote snapshot of the primary volume.
4. Track the time required to complete the second remote snapshot. This is the minimum amount of time that you should allow between scheduled copies.

   Be sure to check the remote bandwidth setting for the other for the management group with the Edit Management Group command. This setting affects the time required to copy a remote snapshot.

## Thresholds

Does the cluster that contains the remote snapshots have sufficient space to accommodate scheduled snapshots?

If the cluster does not have sufficient space available, the remote snapshot appears in the Console and it flashes red. On the Details tab of the remote snapshot, the status says "Read only, not enough space in cluster to start copy."

## Retention Policies

How long do you want to retain the primary snapshots? The remote snapshots? Set different retention policies for the primary and remote snapshots. For example, you can choose to retain 2 primary snapshots and 5 remote snapshots. The number of snapshots retained refers to completed snapshots.

## Parameters for Remote Snapshot Schedule Retention Policies

**The system never deletes the last fully synchronized remote snapshot.**

Under some circumstances, such as unpredictable network speeds or varying snapshot size, a scheduled remote snapshot may create primary snapshots more frequently than the remote copy process can keep up with. The retention policies for scheduled remote copies ensure that such factors do not cause primary and remote snapshots to become unsynchronized. Regardless of the retention policy defined for scheduled remote copies, up to 2 additional snapshots may be retained by the system at any given time. These 2 additional snapshots include the snapshot that is in the process of being copied, and the last fully synchronized snapshot. A fully synchronized snapshot is one that has completed copying so that the remote snapshot is a complete mirror of its corresponding primary snapshot.

**Up to 2 additional snapshots may be retained by the system at any given time**

Because the system never deletes the last fully synchronized primary snapshot, a remote copy schedule may retain N+2 copies for a retention policy of N (the currently copying remote snapshot plus the last fully synchronized snapshot). Using the example above, if you have a retention policy for your remote copy schedule of 2 primary and 5 remote snapshots, the system may retain up to 4 primary and 7 remote snapshots for a period of time.

### Table 61. Snapshot retention policy and maximum number of retained snapshots

| Schedule Remote Snapshot Retention Policy | Maximum Number of Snapshots Retained |
|---|---|
| *n* of primary snapshots | *n* + 2 primary snapshots |
| *x* of remote snapshots | *x* + 2 remote snapshots |
| *n* of hours for primary snapshots | *n* + 2 primary snapshots older than *n* |
| *x* of hours for remote snapshots | *x* + 2 remote snapshots older than *xx* |
| *n* of days for primary snapshots | *n* + 2 primary snapshots older than *n* |
| *x* of days for remote snapshots | *x* + 2 remote snapshots older than *xx* |
| *n* of weeks for primary snapshots | *n* + 2 primary snapshots older than *n* |
| *x* of weeks for remote snapshots | *x* + 2 remote snapshots older than *xx* |

**Remote snapshots are deleted only after their corresponding primary snapshot is deleted.**

Additionally, a remote snapshot is deleted only after its counterpart primary snapshot. You cannot retain fewer scheduled remote snapshots than primary snapshots when setting your retention policies.

**Note:** *If you retain more remote snapshots than primary snapshots, the remote snapshots become regular snapshots when their corresponding primary snapshots are deleted. Identify them as remote snapshots by their names, since the naming convention is established as part of creating the remote snapshot schedule.*

# Best Practices

- Retain at least 2 primary snapshots to ensure that only incremental copying is required for primary snapshots.
- Review your remote copy schedule to ensure that the frequency of the remote copies correlates to the amount of time required to complete a copy.

Use the checklist in Table 62  to help plan scheduled remote snapshots.

## Scheduled Remote Copy Planning Checklist

Use Table 62  to help plan your Remote Copy schedule.

Table 62. Scheduled Remote Copy Planning Checklist

| Configuration Category | Parameters |
|---|---|
| **Scheduled Snapshot** | |
| Start Time | • Start date (mm/dd/yyyy)<br>• Start time (mm:hh:ss) for the schedule to begin |
| Recurrence | • Recurrence (✓). Recurrence is a yes/no choice. Schedule a remote snapshot to occur one time in the future and not have it recur.<br>• Frequency (minutes, hours, days or weeks) |
| **Primary Setup** | |
| Hard Threshold<br>Soft Threshold | Set the hard threshold and soft threshold for the primary snapshot. |
| Retention | Retain either<br>• Maximum number of snapshots (#)<br>• Set period of time (minutes, hours, days or weeks) |
| **Remote Setup** | |
| Management Group | The management group to contain the remote snapshot |
| Volume | The remote volume for the remote snapshots |
| Retention | Retain either<br>• Maximum number of snapshots (#). This number equals completed snapshots only. In-progress snapshots take additional space on the cluster while they are being copied. Also, the system will not delete the last fully synchronized snapshot. For space calculations, figure N+2 with N=maximum number of snapshots.<br>• Set period of time (minutes, hours, days or weeks) |

# D Using Remote Copy

## Remote Copy Overview

This chapter provides instructions for registering, configuring, and using Remote Copy for business continuance, backup and recovery, and failover.

Remote Copy provides a powerful and flexible method for replicating data available for disaster recovery, business continuance, backup and recovery, data migration, and data mining.

For information about how Remote Copy works and how to plan capacity for Remote Copy, see "Understanding and Planning Remote Copy" on page 309.

## Registering Remote Copy

Remote Copy is a feature upgrade. You must purchase a Remote Data Protection Pak license to use Remote Copy beyond the 30-day evaluation period. For information about registering Remote Copy licenses, see "Feature Registration" on page 285

### Number of Remote Copy Licenses Required

Register Remote Copy on each management group that contains storage modules that will participate in Remote Copy. If there are storage modules in a management group that will not contain Remote Copy primary or remote volumes, you do not need to purchase licenses for those modules. For example, if your management group contains a cluster of 2 storage modules that will contain a remote volume, and another cluster of 3 storage modules that will not use Remote Copy, you only need 2 Remote Copy licenses.

### Registering Remote Copy

For information about starting the 30-day evaluation period and about registering Remote Data Protection Pak, "Feature Registration" on page 285.

## Working with Remote Snapshots

Remote snapshots are a core component of Remote Copy. Remote Copy uses the existing volume and snapshot capabilities to replicate, or copy, the data across geographic distances.

# Creating a Remote Snapshot

Creating a remote snapshot is the main task when working with Remote Copy. You can either create a one-time remote snapshot or set up a schedule for recurring remote snapshots. Many of the parameters for either case are the same.

Creating a remote snapshot involves these main steps:

- First, log in to the primary management group.
- Log into the remote management group, either at the navigation window, or within the procedure's dialog window.
- Create a primary snapshot of the primary volume manually.
- When doing a scheduled remote snapshot, the software creates a primary snapshot, which is then copied to the remote volume.
- Either create a remote volume on a remote management group, or select an existing remote volume.
- Last, create the remote snapshot.

## Best Practice

The best way to prepare for remote snapshots is to create the management group and volumes that will be remote *before* taking the snapshot. Although the interface allows you to create management groups, volumes, and snapshots as you go, that may be a distraction at the time a crucial snapshot is needed.

## Getting There

This procedure takes you to the New Remote Snapshot window where remote copy procedures start.

1. In the navigation window, log in to the management group that contains the primary volume for which you are creating the remote snapshot.

   You can create remote volumes and snapshots within the same management group. In that case, you only log in to the one management group as in step 1.

2. Log in to the *remote* management group.
3. In the navigation window, go back to the primary volume, the one you want to copy, and select it.

   If you would like to copy an existing snapshot to a remote management group, select a snapshot at this step.

4. Click the tasks button and select New Remote Snapshot.

   The New Remote Snapshot window displays, as shown in Figure 216.

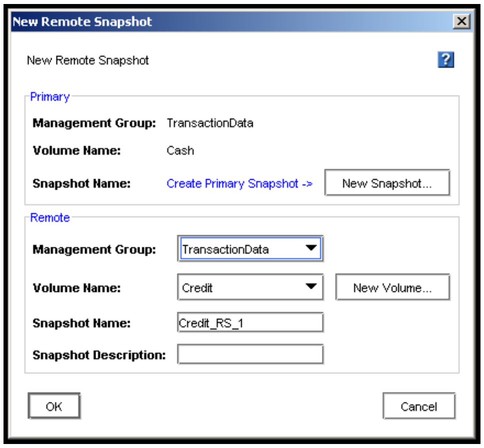

New Remote Snapshot

Primary

Management Group: TransactionData

Volume Name: Cash

Snapshot Name: Create Primary Snapshot ->

Remote

Management Group: TransactionData

Volume Name: Credit

Snapshot Name: Credit_RS_1

Snapshot Description:

Figure 216. Creating a New Remote Snapshot

## Creating the Primary Snapshot

1. In the Primary group box of the New Remote Snapshot window, click New Snapshot.

   The New Snapshot window opens, shown in Figure 217.

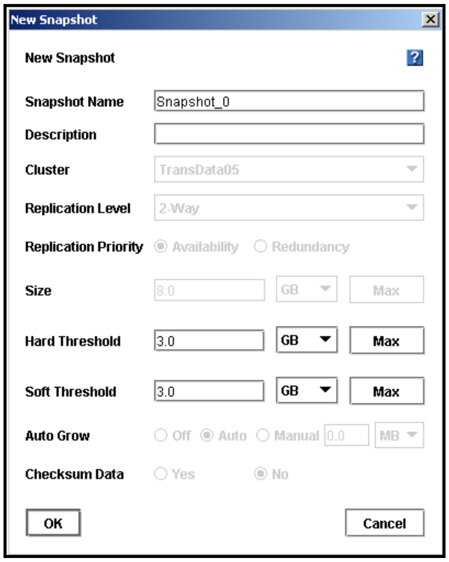

New Snapshot

Snapshot Name: Snapshot_0

Description:

Cluster: TransData05

Replication Level: 2-Way

Replication Priority: ● Availability ○ Redundancy

Size: 8.0 GB

Hard Threshold: 3.0 GB

Soft Threshold: 3.0 GB

Auto Grow: ○ Off ● Auto ○ Manual 0.0 MB

Checksum Data: ○ Yes ● No

Figure 217. Creating a New Primary Snapshot

2. Either accept the suggested name or type a different name for the primary snapshot.

   Names are case sensitive. They cannot be changed after the snapshot is created.

   ***Note:*** *Make the beginning of volume and snapshot names meaningful, for example, "Snap1Exchg_03."*

3. (Optional) Type in a description of the snapshot.
4. (Optional) Change the hard and soft thresholds for the snapshot.
5. Click OK to return to the New Remote Snapshot window.

   The information for the primary snapshot is filled in, as shown in Figure 218. That is, the text for the field Snapshot Name has changed:

   - From 'Create Primary Snapshot'
   - To 'Credit_SS_3forremo'

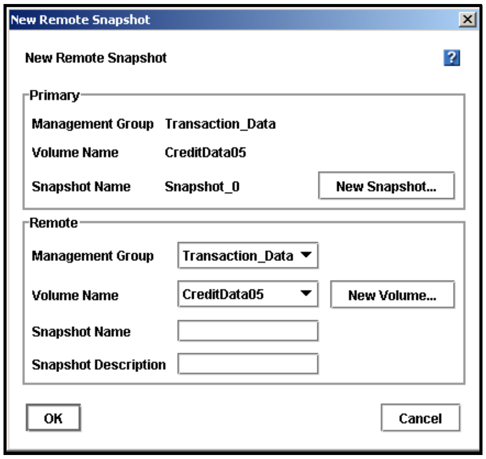

New Remote Snapshot

Primary

Management Group Transaction_Data

Volume Name CreditData05

Snapshot Name Snapshot_0

Remote

Management Group Transaction_Data

Volume Name CreditData05

Snapshot Name

Snapshot Description

Figure 218. New Primary Snapshot Created

In the Remote group box, use the drop-down lists and select the remote site's management group and volume.

The Management Group, Cluster, Volume wizard is available at this point.

6. In the Snapshot Name field, type in the name for this remote snapshot.
7. (Optional) Type in a description for the remote snapshot.

8. Click OK in the New Remote Snapshot window.

   The remote copy of the primary snapshot to the remote volume happens now, as shown in Figure 219.

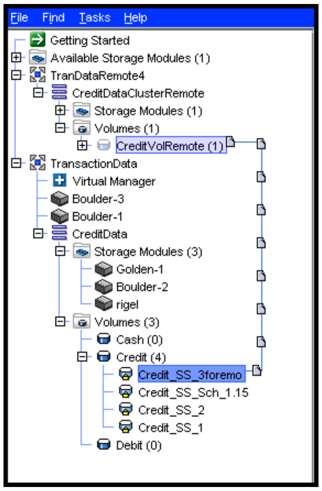

**Figure 219. Remote Copy in Progress**

## Creating a Remote Volume

Create a remote volume by any of the following methods:

- Designate an existing primary volume as a remote volume. See "Designating an existing volume as a remote volume" on page 323.
- Create a new remote volume during creation of a remote snapshot. See "Creating a remote volume on the fly" on page 324.
- Create a new volume manually. See "Creating a new remote volume manually" on page 324.
- Use the Management Group, Clusters, and Volume wizard in Getting Started. See "Getting Started" on page 1 for details on working through the wizards.

### Designating an existing volume as a remote volume

Selecting an existing volume to become a remote volume causes these things to happen:

- A snapshot of all existing data to be created for that volume and then
- All the data in that volume is deleted so that the remote volume has a zero length and zero hard and soft thresholds.

See "Making a Primary Volume Into a Remote Volume" on page 338.

### Creating a new remote volume manually

Create a remote volume as you would any other volume. Be sure to choose the storage modules at the remote site. Because management groups and clusters are logical entities, name them to reflect their remote functionality.

In this method, the primary volume is ready. You create a remote volume at the remote site to receive the snapshot. Then, either take the snapshot and remote copy it, or create the schedule to take remote snapshots.

See "Creating a Volume" on page 225.

### Creating a remote volume on the fly

If you are using the New Remote Snapshot window, you can create a needed cluster and volume as you work through the window.

1. In the Remote group box, select a Cluster Group to contain the remote snapshot.

   You must be logged into the management group you select.

2. Click New Volume.

   The Management Groups, Clusters, and Volumes wizard opens.

   For specific help, see "Getting Started" on page 1 for details on working through the wizards.

3. The volume and cluster names you specified in the wizard fill in the New Remote Snapshot window.

4. (Optional) Type in a description of the remote snapshot and click OK.

   Remote snapshot may take some time. When finished, the Details tab for the new snapshot appears.

## What the System Does

The system creates the remote snapshot in the cluster that contains the remote volume.

The system then copies the primary snapshot onto the remote snapshot. The process of copying the data may take some time.

The remote snapshot appears below the remote volume in the navigation window, as shown in Figure 220.

***Note:*** *If you create a remote snapshot of a volume with a remote snapshot still in progress, the second remote snapshot does not begin copying until the first remote snapshot is complete.*

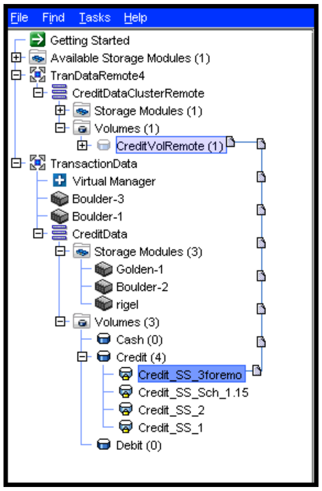

**Figure 220. Viewing the Remote Snapshot**

## Creating the First Copy

Creating the first copy of data is the first step when setting up a Remote Copy solution. Three methods for creating the first copy are described below.

**Copy data directly to the remote site over the WAN.**

Use this method if you are implementing the Remote Copy solution before you have accumulated much data in the primary site and your hardware is already installed in the remote site.

In this method, you create the primary management group and the remote management group in their respective locations. You then create the initial copy of the data directly over the WAN using Remote Copy.

**Use the storage modules intended for the remote site to configure the remote management group on-site and copy data locally. Then ship the remote storage modules to the remote site.**

Use this method if you initially have all the storage modules for the Remote Copy solution at the primary site.

1. Configure both the primary and remote management groups.
2. Create the first copy of the data locally over the Gigabit Ethernet.
3. Ship the storage modules for the remote site and install the remote management group just as you configured it in the primary site.

   The subsequent snapshots from the primary volume to the remote volume are incremental.

**Use the PrimeSync feature of Remote Copy to configure a temporary management group, create the first copy locally, then ship the temporary storage module and again copy locally to the remote target.**

Use this method if you have the primary and remote sites configured and operational.

1. Use a single storage module to create a temporary management group.
2. Copy the primary volume to the temporary management group over Gigabit Ethernet.
3. Ship the temporary system to the remote site and copy the data to the existing remote management group.
4. Remove the temporary system.

   PrimeSync ensures that the proper relationship is established between the original primary volume and the remote site. Subsequent remote snapshots from the primary site to the remote site are incremental.

# Viewing a List of Remote Snapshots

View a list of remote snapshots associated with management groups, clusters, volumes or snapshots.

1. In the navigation window, select the cluster for which you want to view the list of remote snapshots.
2. Click the Remote Snapshot tab to bring it to the front, shown in Figure 221.

   The report in the tab window lists management groups and all the snapshots. The other columns show status information about the remote snapshots as described in detail in “Monitoring Remote Snapshots” on page 330.

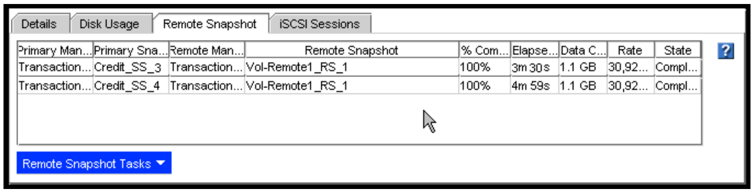

| Primary Man... | Primary Sna... | Remote Man... | Remote Snapshot | % Com... | Elapse... | Data C... | Rate | State |
|---|---|---|---|---|---|---|---|---|
| Transaction... | Credit_SS_3 | Transaction... | Vol-Remote1_RS_1 | 100% | 3m 30s | 1.1 GB | 30,92... | Compl... |
| Transaction... | Credit_SS_4 | Transaction... | Vol-Remote1_RS_1 | 100% | 4m 59s | 1.1 GB | 30,92... | Compl... |

Figure 221. Viewing the List of Remote Snapshots

# Setting the Remote Bandwidth

The remote bandwidth sets the maximum rate for data transfer between management groups. The remote bandwidth setting is the upper limit of the transfer speed. That is, the copy rate is equal to, or less than, the rate set.

To control the maximum rate of data transfer to a remote snapshot, set the remote bandwidth on the management group that contains the remote snapshot—the remote management group. When setting the remote bandwidth, you can choose from a list of common network types, or you can calculate a custom rate, based on your particular requirements.

## Set the Bandwidth

1. In the navigation window, select the remote management group.

1. Click Management Group Tasks and select Edit Management Group.

   The Edit Management Group window opens, as shown in Figure 222.

2. Select the management group.

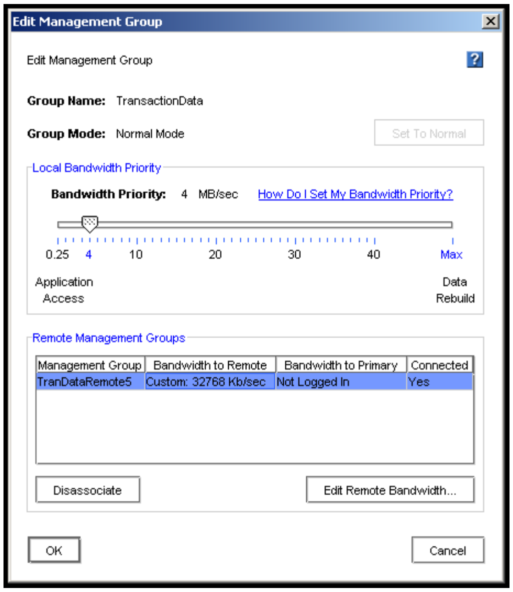

Edit Management Group

Edit Management Group

**Group Name:** TransactionData

**Group Mode:** Normal Mode

Local Bandwidth Priority

**Bandwidth Priority:** 4 MB/sec

| Management Group | Bandwidth to Remote | Bandwidth to Primary | Connected |
|---|---|---|---|
| TranDataRemote5 | Custom: 32768 Kb/sec | Not Logged In | Yes |

**Figure 222. Editing a Remote Management Group**

3. Click Edit Remote Bandwidth.

   The Edit Remote Bandwidth window opens, as shown in Figure 223.

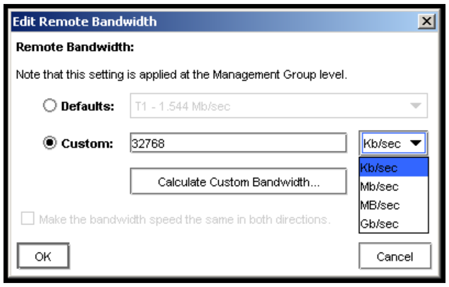

Edit Remote Bandwidth

**Remote Bandwidth:**

Note that this setting is applied at the Management Group level.

○ **Defaults:** T1 - 1.544 Mb/sec

● **Custom:** 32768 Kb/sec

□ Make the bandwidth speed the same in both directions.

**Figure 223. Editing the Remote Bandwidth**

4. Change the bandwidth setting as desired.

## Selecting Remote Bandwidth Rate

Select a preset speed from a list of standard network types. Remember, the speed is the maximum rate at which data will be copied. Or, calculate a custom speed based on your specific requirements.

### Defaults Settings

When setting remote bandwidth, selecting Defaults allows you to choose from a list of common network types, as shown in Figure 224.

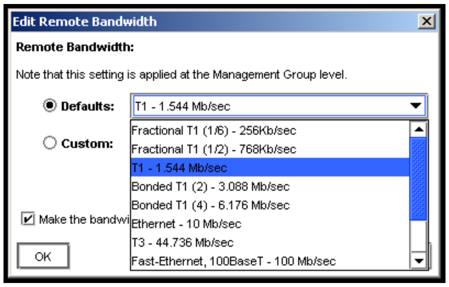

Edit Remote Bandwidth

Remote Bandwidth:

Note that this setting is applied at the Management Group level.

● Defaults: T1 - 1.544 Mb/sec

○ Custom:

☑ Make the bandwi

Fractional T1 (1/6) - 256Kb/sec
Fractional T1 (1/2) - 768Kb/sec
T1 - 1.544 Mb/sec
Bonded T1 (2) - 3.088 Mb/sec
Bonded T1 (4) - 6.176 Mb/sec
Ethernet - 10 Mb/sec
T3 - 44.736 Mb/sec
Fast-Ethernet, 100BaseT - 100 Mb/sec
OK

**Figure 224. Default settings for Remote Bandwidth**

### Custom Settings

The custom setting for remote bandwidth defaults to 32768 Kb, or about 4 Mb. Use the calculation tool to identify a desired bandwidth setting. For example, if you have a T1 line and you want to set the remote bandwidth to 12% of that capacity, you can use the calculation tool to find the correct value, 189 Kb, as shown in Figure 225.

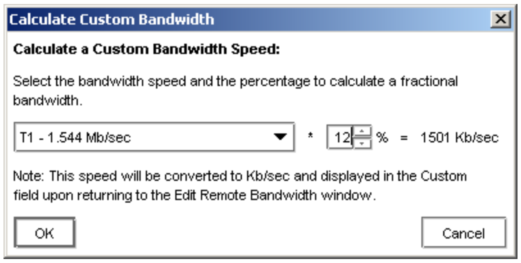

Calculate Custom Bandwidth

Calculate a Custom Bandwidth Speed:

Select the bandwidth speed and the percentage to calculate a fractional bandwidth.

T1 - 1.544 Mb/sec * 12 % = 1501 Kb/sec

Note: This speed will be converted to Kb/sec and displayed in the Custom field upon returning to the Edit Remote Bandwidth window.

**Figure 225. Calculating a Custom Value for Setting Remote Bandwidth**

## Best Practice

Set the bandwidth speed the same in both directions unless you have an asymmetrical WAN link.

# Cancelling a Remote Snapshot

When you cancel a remote snapshot that is in progress, the remote snapshot is deleted, but the primary snapshot remains.

To cancel a remote snapshot that is in progress

1. In the navigation window, select the remote snapshot.
2. Click the Remote Snapshot tab.

3. Select the remote snapshot you want to cancel.
4. Click Snapshot Tasks and select Cancel Remote Snapshot.

   A confirmation message opens.
5. Click OK.

## Editing a Remote Snapshot

You can edit the description of a remote snapshot. You can also change the hard and soft thresholds, but it is not recommended.

1. Log in to the management group that contains the remote snapshot.

1. Select the remote snapshot.
2. Click Snapshots Tasks and select Edit Snapshot.

   The Edit Snapshot window opens, shown in Figure 226.

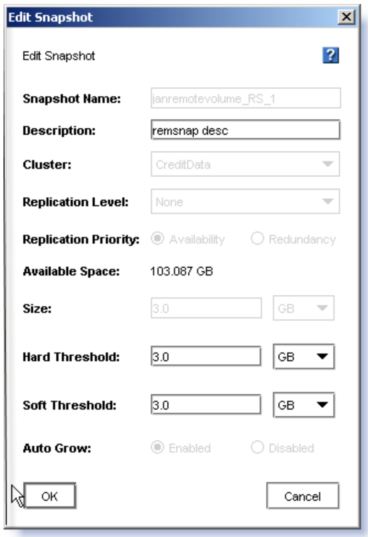

Edit Snapshot

Edit Snapshot

| | | |
|---|---|---|
| **Snapshot Name:** | janremotevolume_RS_1 | |
| **Description:** | remsnap desc | |
| **Cluster:** | CreditData | |
| **Replication Level:** | None | |
| **Replication Priority:** | ● Availability | ○ Redundancy |
| **Available Space:** | 103.087 GB | |
| **Size:** | 3.0 | GB |
| **Hard Threshold:** | 3.0 | GB |
| **Soft Threshold:** | 3.0 | GB |
| **Auto Grow:** | ● Enabled | ○ Disabled |

**Figure 226. Editing a Remote Snapshot**

3. Change the desired information and click OK.

## Deleting a Remote Snapshot

1. Log in to the management group that contains the remote snapshot.
2. Select the remote snapshot in the navigation window.
3. Click Snapshot Tasks and select Delete Snapshot from the menu.

   A confirmation message opens.
4. Click OK, and OK again in the next confirmation window.

# Monitoring Remote Snapshots

Information for monitoring remote snapshots is available from multiple sources. Active monitoring features provide you the capability to configure alerts that you view in the alert window as well as receiving alerts as emails and through SNMP traps. The tab window also provides monitoring information for remote snapshots.

## Monitoring Remote Snapshot Details from the Tab Window

View information about each remote snapshot in both the Remote Snapshot tab and in the Remote Snapshot Details window.

### Viewing Information in the Remote Snapshot Tab

The Remote Snapshot tab displays a list of remote snapshots connected with a selected item in the navigation window. For example, if you select a management group, the Remote Snapshot tab displays the list of remote snapshots associated with that management group. You can view lists of remote snapshots by management group, cluster, volume and snapshot levels.

1. Select the appropriate Remote Snapshot in the navigation window.
2. Click the Remote Snapshot tab to bring it to the front, shown in Figure 227.

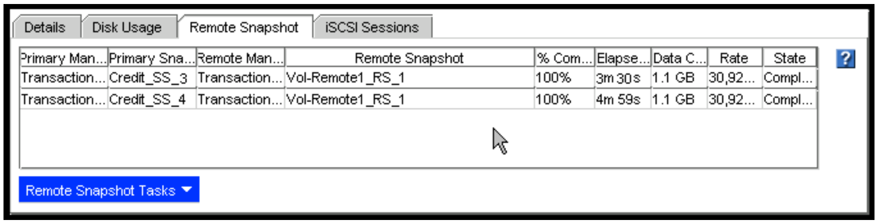

| Primary Man... | Primary Sna... | Remote Man... | Remote Snapshot | % Com... | Elapse... | Data C... | Rate | State |
|---|---|---|---|---|---|---|---|---|
| Transaction... | Credit_SS_3 | Transaction... | Vol-Remote1_RS_1 | 100% | 3m 30s | 1.1 GB | 30,92... | Compl... |
| Transaction... | Credit_SS_4 | Transaction... | Vol-Remote1_RS_1 | 100% | 4m 59s | 1.1 GB | 30,92... | Compl... |

Figure 227. Viewing Remote Snapshot Details in the Remote Snapshot Tab

You may want to check the he remote snapshot details for this information:

- % Complete—the incremental progress of the remote copy operation.
- Elapsed Time—incremental time of the copy operation.
- Data Copied—incremental quantity of data copied.
- Rate—rate at which data is being copied, or, when the remote snapshot is complete, the average rate for the total operation.
- State—status of the operation.

### Viewing Status in the Remote Copy Details Window

The Remote Snapshot Details window displays additional details about a remote snapshot.

1. In the tab window, select the Remote Snapshot tab to bring it to the front.

2. Select a remote snapshot from the list.
3. Click Remote Snapshot Tasks and select View Remote Snapshot Details.

   The Remote Snapshot Details window opens, as shown in Figure 228.

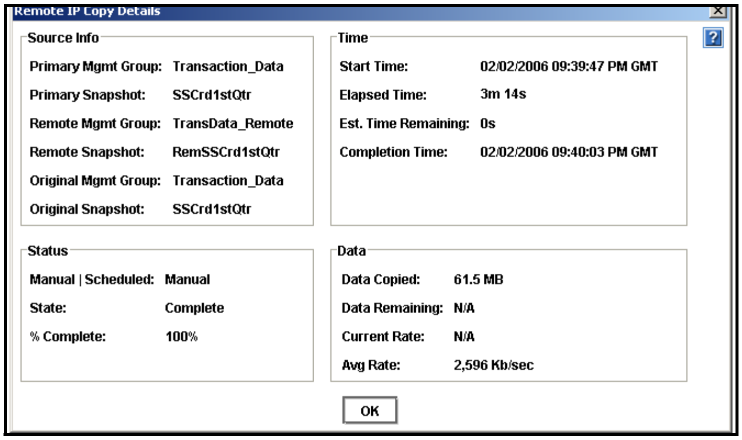

Source Info

Primary Mgmt Group: Transaction_Data
Primary Snapshot: SSCrd1stQtr
Remote Mgmt Group: TransData_Remote
Remote Snapshot: RemSSCrd1stQtr
Original Mgmt Group: Transaction_Data
Original Snapshot: SSCrd1stQtr

Time

Start Time: 02/02/2006 09:39:47 PM GMT
Elapsed Time: 3m 14s
Est. Time Remaining: 0s
Completion Time: 02/02/2006 09:40:03 PM GMT

Status

Manual | Scheduled: Manual
State: Complete
% Complete: 100%

Data

Data Copied: 61.5 MB
Data Remaining: N/A
Current Rate: N/A
Avg Rate: 2,596 Kb/sec

## Figure 228. Viewing Remote Snapshot Details for Completed Remote Copy

During the remote copy process, the Details window reports current data for the statistics. When the copy is completed, the statistics show summary data. Figure 228 shows a completed remote copy. Table 63 lists the values for the statistics reported in the Details window.

## Table 63. Values for Remote Snapshot Details Window

| Statistic | Values |
|---|---|
| **Source Info Section** | |
| Primary Mgmt Group | The management group containing the primary volume and snapshot. |
| Primary Snapshot | The primary snapshot. |
| Remote Mgmt Group | The management group containing the remote snapshot |
| Remote Snapshot | The remote snapshot. |
| Original Mgmt Group | The original management group that contained the original volume and snapshot. Used with PrimeSync feature. |
| Original Snapshot | The first version of the snapshot from which the first copy was created. Used with PrimeSync feature. |
| **Status** | |
| Manual I Scheduled | Whether the snapshot was created using a scheduled snapshot or manually |
| State | Started, Copying, Stalled, Complete<br>Current state of the copy process. |
| Snapshot Scanned (%) | 0-100%<br>Percent of the copy process that is completed. |
| **Time** | |

## Table 63. Values for Remote Snapshot Details Window (Cont'd)

| Statistic | Values |
|---|---|
| Start Time | MM/DD/YY HH:MM [AM/PM]<br>Date and time copy started |
| Elapsed Time | Xd Xh Xm Xs<br>X = a number and the days, hours, minutes, and seconds the copy has been processing.<br>N/A if not yet available. |
| Est. Time Remaining | Xd Xh Xm Xs<br>X = a number and the days, hours, minutes, and seconds estimated to remain in the copy process.<br>N/A for completed copies or in-progress copies not yet calculated. |
| Completion Time | MM/DD/YY HH:MM [AM/PM]<br>Date and time copy completed.<br>N/A for in-progress copies. |
| **Data** | |
| Data Copied | MB, GB, or TB<br>Amount of data copied so far in smallest unit size. |
| Data Remaining | MB, GB, or TB<br>Amount of data remaining to be copied in smallest unit size |
| Current Rate | Kb/sec.<br>Current rate of data being copied in Kb/second. This rate is recalculated regularly throughout the remote copy process.<br>N/A If not yet available or completed. |
| Avg. Rate | Kb/sec.<br>Average rate of copy progress. |

You can leave the Details window open and monitor the progress of the remote copy.

An example of a Details window with a remote copy in progress is shown in Figure 229.

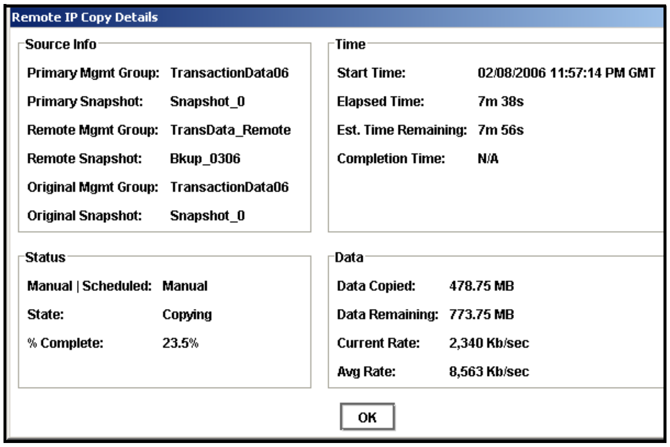

Remote IP Copy Details

Source Info

Primary Mgmt Group: TransactionData06
Primary Snapshot: Snapshot_0
Remote Mgmt Group: TransData_Remote
Remote Snapshot: Bkup_0306
Original Mgmt Group: TransactionData06
Original Snapshot: Snapshot_0

Time

Start Time: 02/08/2006 11:57:14 PM GMT
Elapsed Time: 7m 38s
Est. Time Remaining: 7m 56s
Completion Time: N/A

Status

Manual | Scheduled: Manual
State: Copying
% Complete: 23.5%

Data

Data Copied: 478.75 MB
Data Remaining: 773.75 MB
Current Rate: 2,340 Kb/sec
Avg Rate: 8,563 Kb/sec

**Figure 229. Viewing remote Snapshot Details for Remote Copy in Progress**

## Configuring Active Monitoring Alerts for Remote Copy

There are four variables for remote snapshots for which you can configure alerts. Notification for these variables automatically displays as alert messages in the alert window. You can also configure Active Monitoring to receive email notification or for SNMP traps. The Remote Copy variables that are monitored include these:

- Remote Copy status—an alert is generated if the copy fails.
- Remote Copy complete—an alert is generated when the remote copy is complete.
- Remote Copy failovers—an alert is generated. when a remote volume is made primary.
- Remote management group status—an alert is generated if the connection to a remote management group changes (disconnects and/or reconnects).

To read about configuring Active Monitoring, see "Using Active Monitoring" on page 149.

# Scheduling Remote Snapshots

In addition to taking remote snapshots of a volume manually, one at a time, you can set up a schedule to take snapshots and save them remotely. Scheduled remote snapshots provide for business continuance and disaster recovery, as well as provide a consistent, predictable update of data for remote backup and recovery.

Planning for remote snapshot schedules is a crucial initial step in implementing Remote Copy. The following items require planning in advance for successful deployment of remote snapshot schedules.

- Recurrence (frequency)
- Snapshot thresholds
- Retention policies
- Timing

For detailed information about these issues, see "Planning for Remote Copy" on page 313.

### Best Practices for Scheduling Remote Snapshots

- Create a new remote volume to use with the scheduled remote snapshots.
- If performing daily remote copies, schedule remote snapshots during off-peak hours. If setting scheduled remote snapshots for multiple volumes, stagger the schedules with at least an hour between start times for best results.
- Using NTP, set all storage modules in the management group to the same time zone.

- Reset the management group time before creating a timetable for a scheduled remote snapshot. For detailed information, see "Resetting the Management Group Time" on page 179.

# Creating the Schedule

Create the timetable for future scheduled remote snapshots with these steps.

1. In the navigation window, select the primary volume.

   The primary volume is the one you want to snapshot.

2. Click the Schedules tab to bring it to the front.

3. Click Schedule Tasks and select New Scheduled Remote Snapshot.

   The New Scheduled Remote Snapshot window opens, shown in Figure 230.

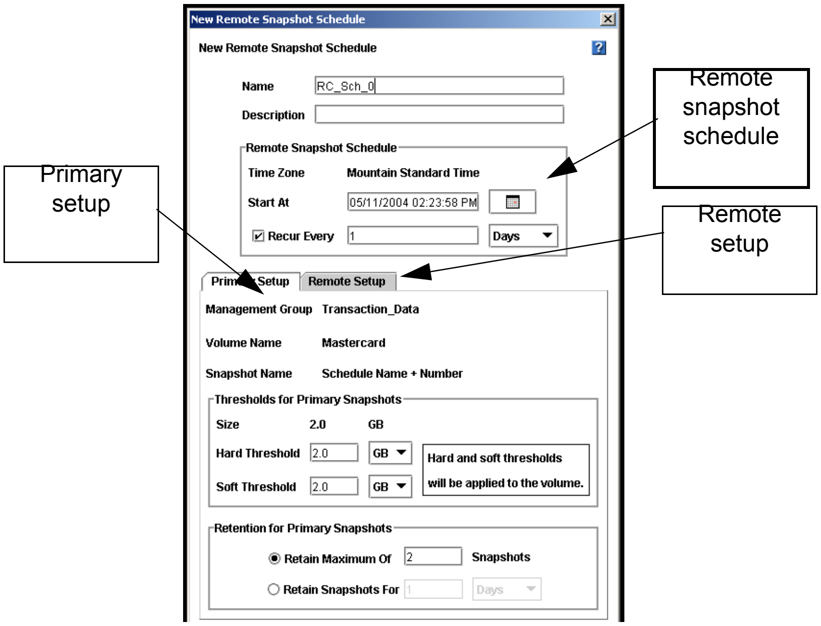

New Remote Snapshot Schedule

Name: RC_Sch_0
Description:

Remote Snapshot Schedule
Time Zone: Mountain Standard Time
Start At: 05/11/2004 02:23:58 PM
☑ Recur Every: 1 Days

Management Group: Transaction_Data
Volume Name: Mastercard
Snapshot Name: Schedule Name + Number

Thresholds for Primary Snapshots
Size: 2.0 GB
Hard Threshold: 2.0 GB
Soft Threshold: 2.0 GB
Hard and soft thresholds will be applied to the volume.

Retention for Primary Snapshots
● Retain Maximum Of 2 Snapshots
○ Retain Snapshots For 1 Days

**Figure 230. Creating a New Timetable for a Remote Snapshot**

## Set General Parameters

4. Click the General tab, and make sure the name and time zone are correct.

   Reset the time zone with the Edit Configuration window.

5. Click Edit to set the time to start the first remote snapshot of the schedule, and select a recurrence interval.

   You can schedule a snapshot every 30 minutes or more.

## Setup Primary Snapshot Parameters

6. Click the Primary Setup tab and make sure the thresholds are satisfactory according to the best practice as recommended in Table 64 .

**Table 64. Best Practice for Setting Thresholds**

| | |
|---|---|
| Minimum recommended hard threshold | 512 MB |
| Recommended gap between hard and soft thresholds | 256 MB |
| Minimum soft threshold | 256 MB |

2. Select a retention interval for the primary snapshot, either number of days or number of snapshots.

   You can retain up to 50 snapshots for a volume, and up to 200 for all volumes.

## Setup Remote Snapshot Parameters

1. Click the Remote Setup tab and select the management group and volume that will hold the *remote* snapshots.

   Log in if you need to.

   Click New Volume to use the wizard to create a volume if you need to make a new one.

2. Set the retention interval for the *remote* snapshot.

   You can retain up to 50 snapshots for a volume, and up to 200 for all volumes.

3. Click OK to close the scheduling windows and return to the navigation and tab windows.

The timetable you just created is listed in the tab view now.

## Timing for a Scheduled Remote Snapshot

When you set up a timetable for scheduled remote snapshots with the previous procedure, you rely on the time. The time zone displayed in the Schedule Remote Snapshot windows is the time zone of the storage module through which you first logged in to the management group. See “Best Practices for Scheduling Remote Snapshots” on page 333.

## What the System Does

**Best Practice: If you created a new volume for the *remote* volume**, the system creates a new primary snapshot of the primary volume and a remote snapshot of the remote volume. See “Best Practices for Scheduling Remote Snapshots” on page 333

**If you selected an existing volume to become the** remote volume, the system alerts you that all the data on the existing volume will be deleted, but that a snapshot of all the existing data will be created first. The snapshot that is then created retains all the volume's data.

1. Type a name for that snapshot in the alert.
2. Click Yes to continue.

The new snapshot is created and the volume becomes a remote volume.

The system creates a new primary snapshot of the primary volume and a remote snapshot of the remote volume. It then copies the data from the primary snapshot to the remote snapshot. This process occurs according to the schedule.

## Editing the Remote Snapshot Schedule

When editing the timetable for a scheduled remote snapshot, you can change the following items.

- **Schedule**—description, start date and time, recurrence policy
- **Primary Setup**—primary snapshot thresholds, retention policy
- **Remote Setup**—retention policy

1. In the navigation window, select the primary volume that has the schedule you want to edit.
2. Click the Schedules tab and select the schedule to edit.

3. Click Schedule Tasks and select Edit Schedule.

   The Edit Schedule Remote Snapshot window opens, shown in Figure 231.

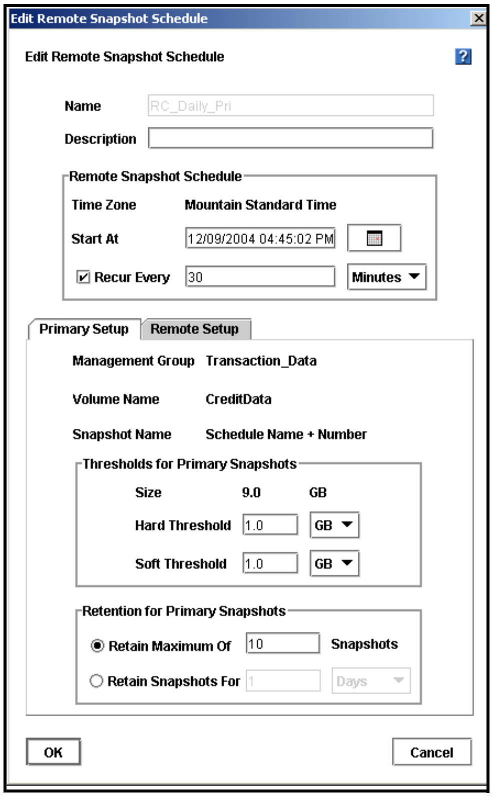

Edit Remote Snapshot Schedule

Name RC_Daily_Pri

Description

Remote Snapshot Schedule

Time Zone Mountain Standard Time

Start At 12/09/2004 04:45:02 PM

☑ Recur Every 30 Minutes

Management Group Transaction_Data

Volume Name CreditData

Snapshot Name Schedule Name + Number

Thresholds for Primary Snapshots

Size 9.0 GB

Hard Threshold 1.0 GB

Soft Threshold 1.0 GB

Retention for Primary Snapshots

● Retain Maximum Of 10 Snapshots

○ Retain Snapshots For 1 Days

**Figure 231. Editing the Timetable for a Scheduled Remote Snapshot**

4. Change the desired information.
5. Click OK.

# Deleting the Remote Snapshot Schedule

1. In the navigation window, select the primary volume that has the schedule you want to delete.
2. Click the Schedule tab to bring it to the front.
3. Select the schedule you want to delete.
4. Click Schedule Tasks and select Delete Schedule.

   A confirmation message opens.
5. Click OK.

# Changing the Roles of Primary and Remote Volumes

Changing the roles of primary and remote volumes may be necessary during failover recovery. Use these procedures when you are resynchronizing data between the acting primary volume and the recovered, or newly configured, production site primary volume.

## Making a Primary Volume Into a Remote Volume

Make any primary volume into a remote volume. First, the system takes a snapshot of the volume to preserve the existing data that are on the volume. The data can then be accessed on that snapshot.

Next, the volume is converted to a remote volume. The remote volume is a placeholder for the remote snapshots and does not contain data itself. So the size, hard threshold and soft threshold change to 0.

1. In the navigation view, select the volume that you want to convert.
2. Click Volume Tasks and select Edit Volume.

   The Edit Volume window opens.

3. Change the type from Primary to Remote.

   Notice that the window changes to the Edit Remote Volume window and all the fields are greyed out, as shown in Figure 232.

   Additionally, the values in the size, hard threshold and soft threshold fields are set to 0.

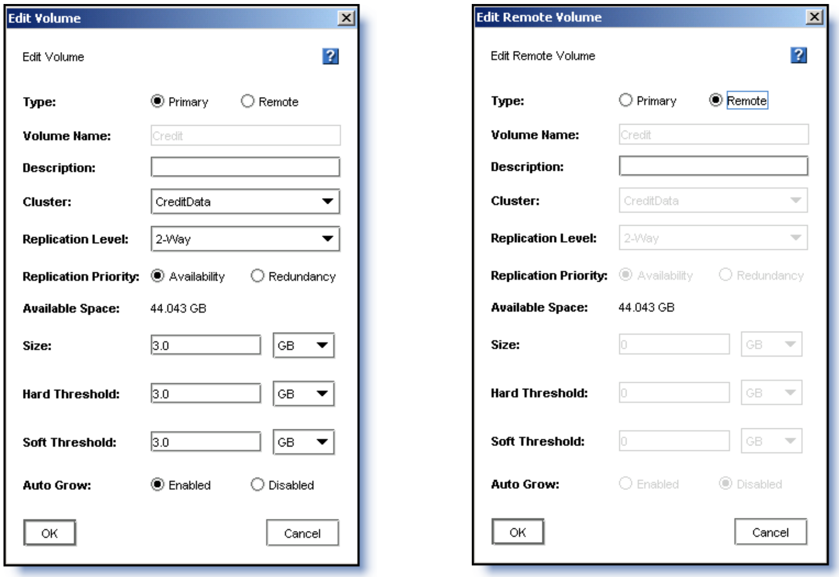

**Edit Volume**

Edit Volume

Type: ● Primary ○ Remote

Volume Name: Credit

Description:

Cluster: CreditData

Replication Level: 2-Way

Replication Priority: ● Availability ○ Redundancy

Available Space: 44.043 GB

Size: 3.0 GB

Hard Threshold: 3.0 GB

Soft Threshold: 3.0 GB

Auto Grow: ● Enabled ○ Disabled

**Edit Remote Volume**

Edit Remote Volume

Type: ○ Primary ● Remote

Volume Name: Credit

Description:

Cluster: CreditData

Replication Level: 2-Way

Replication Priority: ● Availability ○ Redundancy

Available Space: 44.043 GB

Size: 0 GB

Hard Threshold: 0 GB

Soft Threshold: 0 GB

Auto Grow: ○ Enabled ● Disabled

Primary volume with size, hard and soft thresholds defined.

Selecting Remote changes the volume to a remote volume with the size, hard and soft thresholds set to 0.

**Figure 232. Changes to a Volume When Changed from Primary to Remote**

4. Click OK.

   The Make Volume Remote window opens, shown in Figure 233.

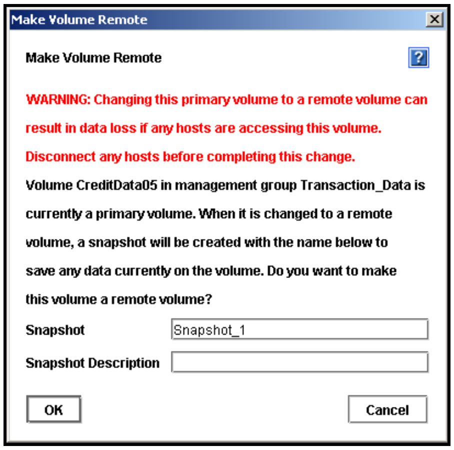

Make Volume Remote

Make Volume Remote

WARNING: Changing this primary volume to a remote volume can result in data loss if any hosts are accessing this volume. Disconnect any hosts before completing this change.

Volume CreditData05 in management group Transaction_Data is currently a primary volume. When it is changed to a remote volume, a snapshot will be created with the name below to save any data currently on the volume. Do you want to make this volume a remote volume?

Snapshot: Snapshot_1

Snapshot Description:

Figure 233. Creating a Snapshot Before Making a Primary Volume into a Remote Volume

5. Type a name for the snapshot that will be created.

   This snapshot preserves any existing data on the volume.

6. (Optional) Type a description for the snapshot.

7. Click OK.

   The snapshot is created and the volume becomes a remote volume.

   The Edit Remote Volume window opens again with the editable fields enabled, as shown in Figure 234.

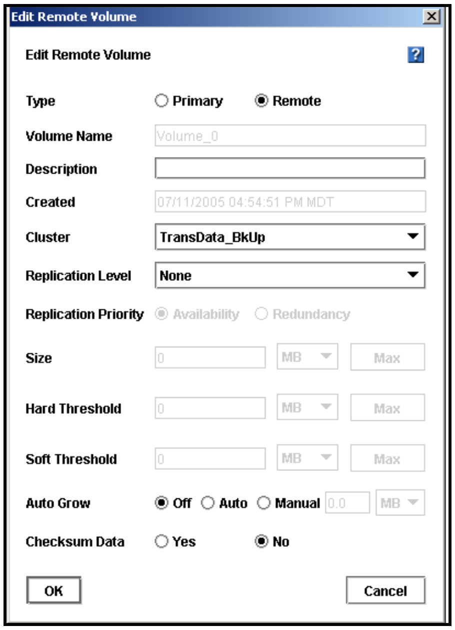

Edit Remote Volume

Edit Remote Volume

Type ○ Primary ● Remote

Volume Name: Volume_0

Description:

Created: 07/11/2005 04:54:51 PM MDT

Cluster: TransData_BkUp

Replication Level: None

Replication Priority ● Availability ○ Redundancy

Size: 0 MB Max

Hard Threshold: 0 MB Max

Soft Threshold: 0 MB Max

Auto Grow ● Off ○ Auto ○ Manual 0.0 MB

Checksum Data ○ Yes ● No

Figure 234. Finalizing the New Remote Volume

8. (Optional) Make any necessary changes to the new remote volume.
9. (Optional) If you do not need to preserve the data in the snapshot, delete it to reclaim space in the cluster.

## Making a Remote Volume Into a Primary Volume

Make a remote volume into a primary volume. Changing the remote volume into a primary volume allows the backup application server to read and write to the volume. This is useful in failover recovery if you want to use the failover site as the acting primary site.

***Note:*** *You cannot make a remote volume into a primary volume while a remote snapshot is in progress. Wait until the remote snapshot copy is complete before making the remote volume into a primary volume, or cancel the in-progress remote copy.*

## Designating Size and Threshold Values for the Converted Volume

1. In the navigation window, select the remote volume you want to convert.

2. Click Volume Tasks and Select Edit Volume.

   The Edit Remote Volume window opens.

3. Change the type from Remote to Primary.

   Notice that the window title changes to the Edit Volume window and some fields are greyed out, as shown in Figure 235.

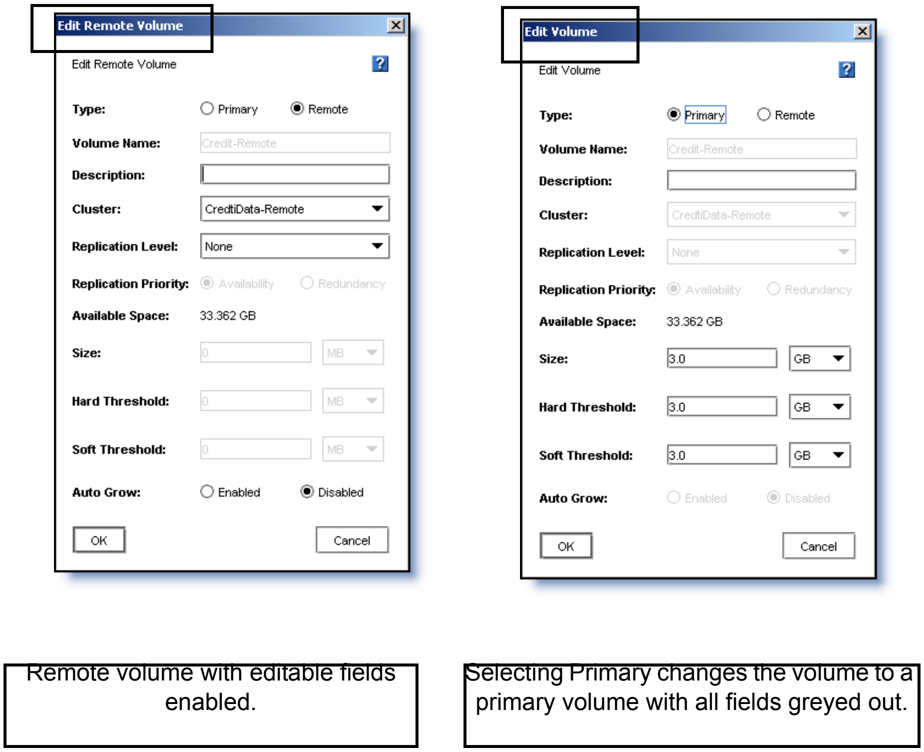

**Edit Remote Volume**

Edit Remote Volume

| | |
|---|---|
| Type: | ○ Primary ● Remote |
| Volume Name: | Credit-Remote |
| Description: | |
| Cluster: | CredtiData-Remote |
| Replication Level: | None |
| Replication Priority: | ● Availability ○ Redundancy |
| Available Space: | 33.362 GB |
| Size: | 0 MB |
| Hard Threshold: | 0 MB |
| Soft Threshold: | 0 MB |
| Auto Grow: | ○ Enabled ● Disabled |

**Edit Volume**

Edit Volume

| | |
|---|---|
| Type: | ● Primary ○ Remote |
| Volume Name: | Credit-Remote |
| Description: | |
| Cluster: | CredtiData-Remote |
| Replication Level: | None |
| Replication Priority: | ● Availability ○ Redundancy |
| Available Space: | 33.362 GB |
| Size: | 3.0 GB |
| Hard Threshold: | 3.0 GB |
| Soft Threshold: | 3.0 GB |
| Auto Grow: | ○ Enabled ● Disabled |

Remote volume with editable fields enabled.

Selecting Primary changes the volume to a primary volume with all fields greyed out.

**Figure 235. Making a Remote Volume into a Primary Volume**

4. Click OK.

   The Edit Volume window displays the editable fields enabled. You can edit everything but the name.

5. Make any required changes, for example, to the size and hard and soft thresholds.

6. Click OK.

   The volume becomes a primary volume.

# Configuring Failover

Configuring Remote Copy for failover provides for business continuance and disaster recovery. When configuring failover, consider both the failover path and the recovery from failover.

## Planning Failover

To achieve failover, consider the following points:

- The location and structure of management groups and clusters
- Configuration of primary and remote volumes and snapshots and scheduling snapshots
- Configuration of application servers and backup application servers
- Task flow for failover recovery (resuming production after failover)

## Using Scripting for Failover

Application-based scripting provides the capability for creating, mounting, and deleting snapshots using scripts. Remote Copy can be scripted as well. Remote snapshots and scheduled remote snapshots can be created and managed using scripts. Detailed information about snapshot scripting can be found in "Working with Scripting" on page 257.

# Resuming Production After Failover

After failover occurs, three scenarios exist for resuming production.

- Failback recovery returns operations to the original primary site once it is restored.
- Make the backup site into the new primary site.
- Set up a new primary site and resume operations at that site.

The task flow for restoring or recovering data and resuming the original Remote Copy configuration is different for each scenario.

## Synchronizing Data After Failover

After a failover, there will usually be 2 snapshots or volumes that have conflicting data. Recovering and synchronizing such data depends on multiple factors, including the application involved. For more detail about synchronizing, see Table 58, "Remote Copy Glossary".

**Example Scenario**

The following example illustrates only one process for synchronizing data. Remember that such synchronization is optional.

## Time Line of Failover

**Table 65. Time Line of Failover**

| Time | Event | What Happens |
|---|---|---|
| 1:00 p.m. | Regular hourly scheduled remote snapshot | Remotess_0 created in remote Management Group |
| 1:10 p.m. | Remote copy finishes | Copying is complete |
| 1:30 p.m. | Primary volume goes offline | OrigPrimaryVol_0 offline |
| 1:33 p.m. | Scripted failover causes remote volume to become the acting primary volume. | ActPrimaryVol_0 active in remote Management Group |
| 2:00 p.m. | Original primary volume comes back online | OrigPrimaryVol_0 online |

## Data that Now Needs to be Synchronized

- Original volume, which contains data from 1:00 to 1:30 p.m.
- Acting primary volume which contains data from 1:33 to 2:00 p.m.

# Returning Operations to Original Primary Site

Once the original primary site is operational again, restore operations to that site. The steps to restore operations depend upon the state of the original primary volume.

- If the primary volume is working

  Synchronize the data between the acting primary volume and the restored primary volume before returning the acting primary volume to its remote volume state.

- If the primary volume is not available

  Create a new primary volume, synchronize the data with the acting primary volume, and then return the acting primary volume to a remote volume.

## Synchronizing the Data Between the Acting Primary Volume and the Original Primary Volume

**1. Create Snapshots of Data**

Create snapshots that contain the data that you need to synchronize. The steps to create those snapshots are described in Table 66 .

### Table 66. Creating Snapshots of Data to Synchronize

| Action | Volumes and Snapshots on Primary Management Group | Volumes and Snapshots on Remote Management Group | What This Step Accomplishes |
|---|---|---|---|
| 1. Stop applications that are accessing the volumes. | | | |
| 2. Make a snapshot of the original volume. | OrigPrimaryVol_0<br>OrigPrimarySS_0 | | Creates a snapshot of the original primary volume that includes the data from 1:00 - 1:30 p.m. |
| 3. Make the acting primary volume into the remote volume. This automatically creates a snapshot of the acting primary volume. | | Remotevol_0<br>ActPrimarySS_0 | Returns the remote management group to its original configuration. |

**2. Synchronize the Data**

Synchronize the snapshots OrigPrimarySS_0 and ActPrimarySS_0 that were created in Steps 2 and 3 of Table 66 .

- To synchronize the snapshots, remote copy the remote snapshot back to the original primary volume. For more detail about synchronizing, see Table 58, “Remote Copy Glossary”.

## Creating a New Primary Volume at the Original Production Site

If the original primary volume is not available, designate a new primary volume, synchronize the data from the acting primary volume, and configure the timetable for the scheduled remote snapshot schedule on the new primary volume.

1. Stop the application that is accessing the acting primary volume.

2. Create a remote snapshot of the acting primary volume and make a new primary volume on the original production site as part of creating that remote snapshot.
3. Convert the remote volume into a primary volume.
4. Make the acting primary volume into the remote volume.

   This creates a snapshot of that volume.
5. Configure a new timetable for the scheduled remote snapshots on the new primary volume.
6. Reconfigure scripts for failover on the application servers.

## Setting Up a New Production Site

Setting up a new production site involves creating a new primary volume and synchronizing the acting primary volume before returning it to its original state as a remote volume. The steps are the same as those for creating a new primary volume at the original production site.

## Making the Backup Site into the New Production Site

Turn the backup site into the new production site and designate a different backup site. The steps are similar to those for initially configuring Remote Copy.

1. Create a remote snapshot or a timetable for a scheduled remote snapshot on the acting primary volume.
2. Make a new remote volume on the new backup site as part of creating that remote snapshot or timetable for a scheduled remote snapshot.
3. Reconfigure scripts for failover on the application servers.

# Rolling Back Primary and Remote Volumes

Rolling back a volume from a snapshot is the method for reverting to an earlier copy of the data on a volume. Rolling back procedures require that you delete any snapshots that were created after the snapshot that is rolled back to.

## Rolling Back a Primary Volume

Rolling back a primary volume to a primary snapshot replaces the original primary volume with a read/write copy of the selected primary snapshot. The new volume has a different name than the original and the original volume is deleted.

**Prerequisites**

- Stop applications from accessing the volume.

- Delete all snapshots that are newer than the snapshot you are rolling back from.

***Warning:*** *Any uncompleted remote copy snapshot that is newer than the snapshot that you are rolling back to is cancelled.*

1. Log in to the management group that contains the primary volume that you want to roll back.

1. Select the snapshot that you want to roll back to.
2. Review the snapshot Details tab to ensure you have selected the correct snapshot.
3. Click Snapshot tasks and select Roll Back Volume.

   The Roll Back Volume window opens, shown in Figure 236.

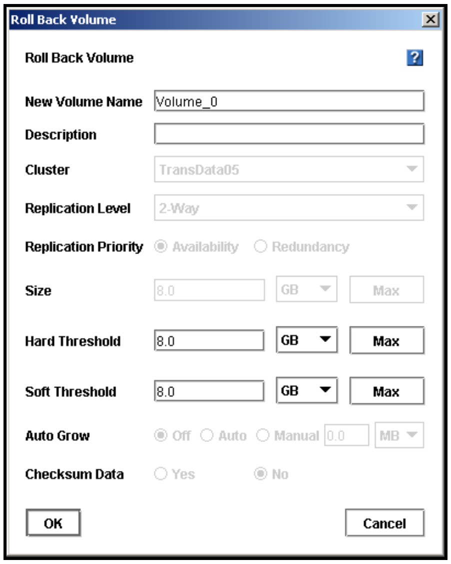

Roll Back Volume

Roll Back Volume

New Volume Name: Volume_0

Description:

Cluster: TransData05

Replication Level: 2-Way

Replication Priority: ● Availability ○ Redundancy

Size: 8.0 GB Max

Hard Threshold: 8.0 GB Max

Soft Threshold: 8.0 GB Max

Auto Grow: ● Off ○ Auto ○ Manual 0.0 MB

Checksum Data: ○ Yes ● No

**Figure 236. Rolling Back a Primary Volume**

4. Type a new name for the rolled back primary volume.
5. You can also change the hard threshold and soft threshold if necessary, using the recommendations listed in Table 67  for setting thresholds. For more information, read "Managing Capacity Using Volume and Snapshot Thresholds" on page 239.

**Table 67. Best Practice for Setting Thresholds**

| | |
|---|---|
| Minimum recommended hard threshold | 512 MB |
| Recommended gap between hard and soft thresholds | 256 MB |
| Minimum soft threshold | 256 MB |

6. Click OK.

   The Roll Back Volume confirmation message opens, as shown in Figure 237.

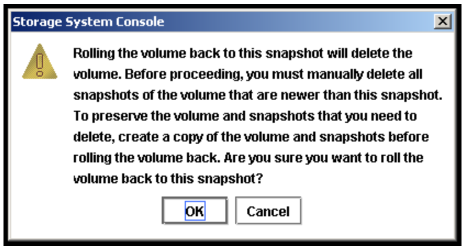

Rolling the volume back to this snapshot will delete the volume. Before proceeding, you must manually delete all snapshots of the volume that are newer than this snapshot. To preserve the volume and snapshots that you need to delete, create a copy of the volume and snapshots before rolling the volume back. Are you sure you want to roll the volume back to this snapshot?

**Figure 237. Verifying the Primary Volume Roll Back**

7. Click OK.

   The primary snapshot version of the primary volume is restored as a read/write volume.

8. Reconfigure application servers to access the new volume.

## Rolling Back a Remote Volume

A remote volume cannot be rolled back. To roll back a remote volume, make the remote volume into a primary volume.

# Using Remote Snapshots for Data Migration and Data Mining

Use remote snapshots to create split mirrors for data mining and data migration. A split mirror is a one-time remote snapshot created from the volume containing the data you want to use or move. Split mirrors are usually created for one-time use and then discarded.

## Creating a Split Mirror

To create a split mirror, perform these general steps:

- Create a remote snapshot
- Create a volume list for that snapshot
- Create an authentication group for client access
- Configure client to access the remote snapshot

# Disassociating Remote Management Groups

Management groups become associated when linked by either remote snapshots or scheduled remote snapshots. Disassociating management groups destroys all the shared knowledge between those groups.

### Best Practice for Disassociating Management Groups

Do this only if a group no longer exists, or if instructed by Customer Support.

1. Log in to both management groups that you want to disassociate.
2. In the navigation window, select the remote management group.
3. Click Management Group Tasks and select Edit Management Group.

   The Edit Management Groups window opens, shown in Figure 238.

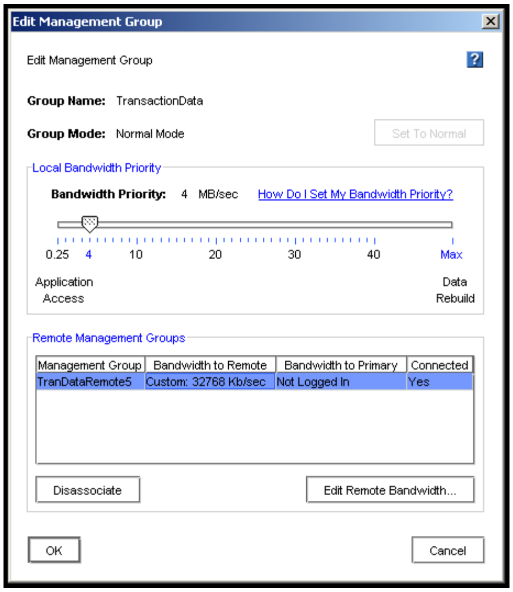

**Group Name:** TransactionData

**Group Mode:** Normal Mode

**Bandwidth Priority:** 4 MB/sec

| Management Group | Bandwidth to Remote | Bandwidth to Primary | Connected |
|---|---|---|---|
| TranDataRemote5 | Custom: 32768 Kb/sec | Not Logged In | Yes |

Figure 238. Editing a Management Group

4. Select the management group or groups you want to disassociate.
5. Click Disassociate.

   A confirmation message opens, describing the results of disassociating the management groups.

***Warning:*** *Disassociating the management group cancels any in-progress remote snapshots and deletes all timetables between the primary and remote management groups.*

6. Click OK.

   The Edit Management Group window opens and the remote management group you disassociated from is gone from the list.

7. Click OK to return to the navigation window.

# E Sample Remote Copy Configurations

## Overview

Because of the flexibility provided by Remote Copy, you can use the functionality in a variety of configurations that are most suitable for your requirements. The sample configurations described in this chapter are only a few possible ways to use Remote Copy for business continuance, backup and recovery, data migration and data mining.

## Using Remote Copy for Business Continuance

Business continuance comprises both disaster recovery and high availability of data. Using Remote Copy for business continuance, data is stored off-site and is continuously available in the event of a site or system failure.

### Achieving High Availability

Creating remote snapshots in remote locations with application-based scripting can ensure that database applications such as SQL Server, Oracle, and Exchange have continual access to data volumes if production application servers or data volumes fail.

Using off-site remote snapshots of your production volumes, you can configure a backup application server to access those remote snapshots. Off-site remote snapshots, particularly when supplemented with synchronous volume replication within a cluster, ensures high availability of critical data volumes.

### Configuration for High Availability

To use remote snapshots for high availability, configure a backup application server to access remote snapshots in the event of a primary system failure. Figure 239 illustrates this simple high availability configuration.

- Configure clustered application servers in both the primary and backup locations.
- During normal operation, the production application server read/writes to the primary volume.
- Set up a schedule for copying remote snapshots to the backup location. If your application server uses multiple volumes that must be in sync, use a script to quiesce the application before creating remote snapshots.

## Configuration Diagram

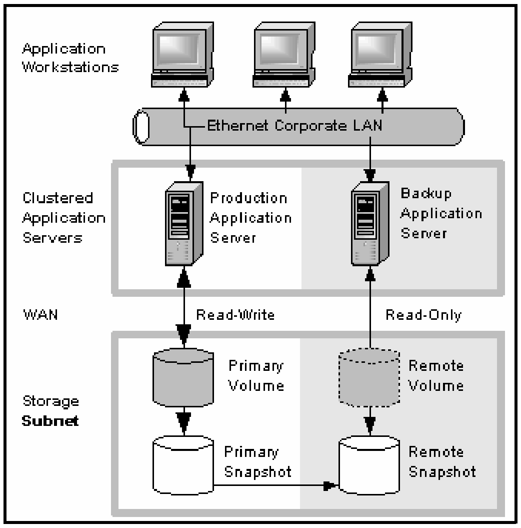

**Figure 239. High Availability Example Configuration**

# How This Configuration Works for High Availability

If the production application server or volumes become unavailable, application processing fails over to the backup application server. As shown in Figure 240, the remote volume and remote snapshots become primary and the backup application server becomes the production application server, accessing data from the acting primary volume.

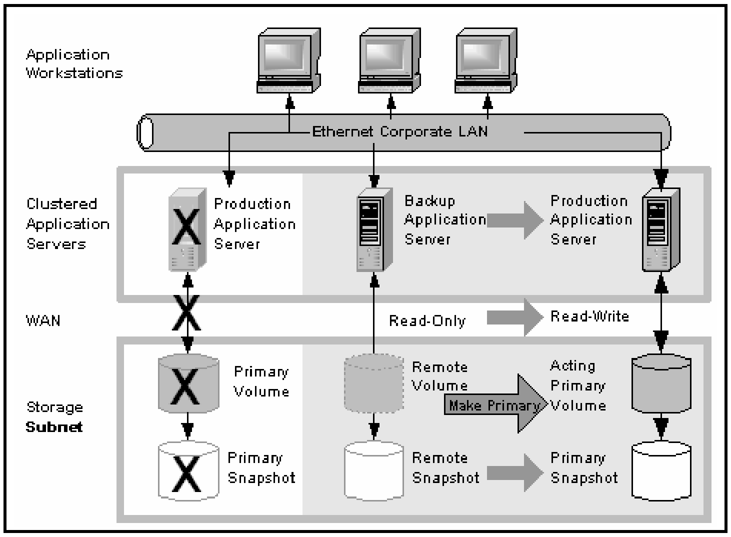

**Figure 240. High Availability Configuration During Failover**

## Data availability if the primary volume or production application server fails

If either the primary volume or production application server in your production site fails, only that data written to the volume since the last remote snapshot was created will be unavailable until the volume or production application server is restored.

## Failover to the backup application server

To maintain availability of the application and the remaining data, the following process occurs:

1. A script or other application monitoring the production application server discovers that primary volume is not available. A script executes to fail over to the backup application server.
2. The backup application server executes a script to convert the remote volume into a primary volume so that the volume can be accessed by the backup application server.
3. Because the backup application server was configured to access the remote (now primary) volume, operation of backup application server begins.

The application continues to operate after the failover to the backup application servers.

## Failback to the production configuration

When the production server and volumes become available again, you have two failback options:

- Resume operations using the original production server, and return the backup volumes to their original remote status, as illustrated in Figure 241. This will require migration back onto the production volumes of data that was written to the backup volumes since the failure.
- Continue operating on the backup application server. When the production server and volumes become available, configure the production server to be the backup server (role reversal).

## Merging data for failback

In the failover scenarios described above there are probably two snapshots with different data. As part of failback, users must make a decision whether to merge the data from the two snapshots and the most effective method for doing so.

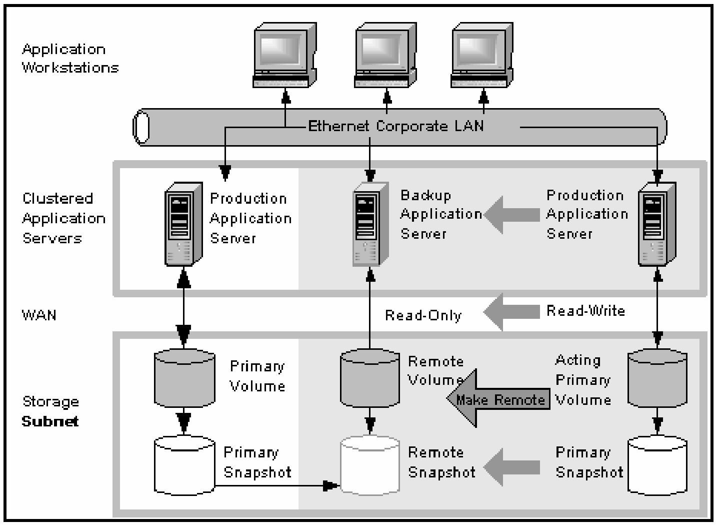

Figure 241. High Availability Configuration During Failback

# Best Practices

## Use remote snapshots in conjunction with local synchronous volume replication

Using remote snapshots alone, any data written to the primary volume since the most recent remote snapshot was created will be unavailable if the primary volume is unavailable.

However, you can lessen the impact of primary volume failure by using synchronous volume replication. Volume replication allows you to create up to 3 copies of a volume on the same cluster of SSMs as the primary volume. The only limitation is that the cluster must contain at least as many SSMs as replicas of the volume. Replicating the volume within the cluster ensures that if an SSM in the cluster goes down, replicas of the volume elsewhere in the cluster will still be available. (For 3-way replication up to 2 SSMs can fail.) For detailed information about volume replication,

## Example configuration

This example, illustrated in Figure 242, uses 3 SSMs per cluster. However, this scenario can use any number of SSMs. Information about creating clusters and volumes can be found at “Working with Clusters” on page 201 and “Working with Volumes” on page 215.

- In the production location, create a management group and a cluster of 3 SSMs.
- Create volumes on the cluster, and set the replication level to 2.
- Configure the production application server to access the primary volume via iSCSI.
- Create a second management group and cluster of 3 SSMs in the backup location.
- Create a schedule for making remote snapshots of the primary volume.

***Note:*** *Volume replication levels are set independently for primary and remote volumes.*

**How It Works:** If one of the SSMs in the primary location fails, the primary volume will still be available. If all of the SSMs fail, or if the application server fails, then failover to the backup application server occurs, and the remote snapshot becomes available.

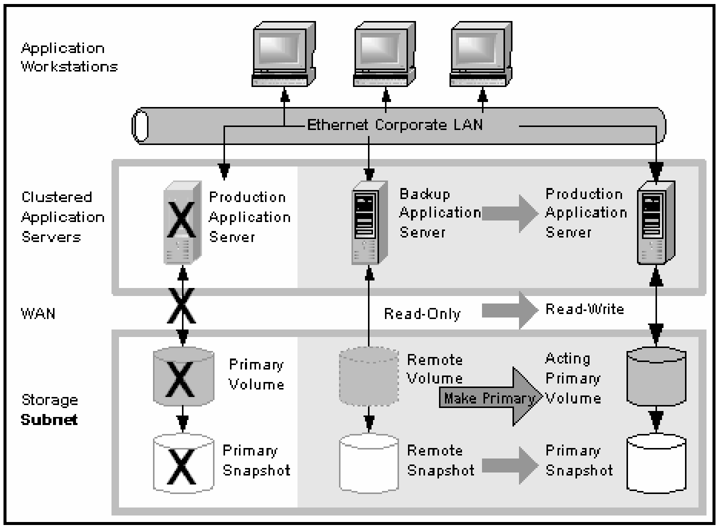

**Figure 242. High Availability During Failover - Example Configuration**

# Achieving Affordable Disaster Recovery

Even if you do not have clustered application servers or network bandwidth required for configuring hot backup sites, you can still use Remote Copy to protect your data during an emergency.

Using remote snapshots, you can maintain copies of your volumes in remote sites. Set up a schedule for creating remote copies, and if your primary storage site becomes unavailable, you can easily access the most recent remote copy of your data volumes. You can also use remote snapshots to transfer data to a backup location where tape backups are then created. This eliminates the backup window on your primary volumes, and ensures that you have copies of your data in the remote site on SSMs as well as on tape.

# Configuration for Affordable Disaster Recovery

To configure affordable disaster recovery, create remote snapshots of your volumes in an off-site location. In addition, you can create tape backups from the remote snapshots in the off-site location:

- Designate one or more off-site locations to be the destination for remote snapshots.
- Set up a schedule for creating remote snapshots in the designated off-site locations. If your application server uses multiple volumes that must be in sync, use a script to quiesce the application before creating remote snapshots.
- Create routine tape backups of the remote snapshots in the off-site locations.

## Configuration Diagram

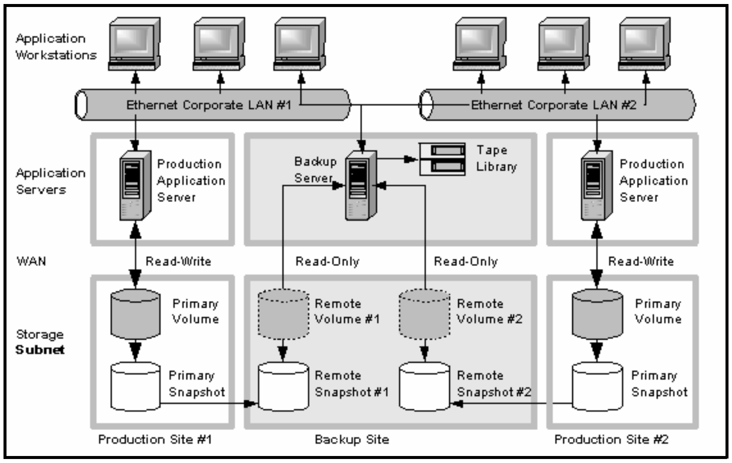

**Figure 243. Affordable Disaster Recovery Example Configuration**

# How this Works for Affordable Disaster Recovery

If the SSMs in your primary location fail or volumes become unavailable, the off-site location contains the most recent remote snapshots.

- Use the remote snapshots to resume operations as shown in Figure 244. If you created tape backups, you can recover data from tape backups, as shown in Figure 245.
- Only data written to the primary volumes since the last remote snapshot was created will be unavailable.
- Application servers that were accessing the down volumes will not be available until you reconfigure them to access recovered data.

To resume operations using the most recent set of remote snapshots:

1. In the backup location, make the remote volume into a primary volume.
2. Configure application servers to access this volume, or if network connections are not fast enough to facilitate reading and writing to the off-site location, copy this volume to a location where application servers can access it more efficiently.

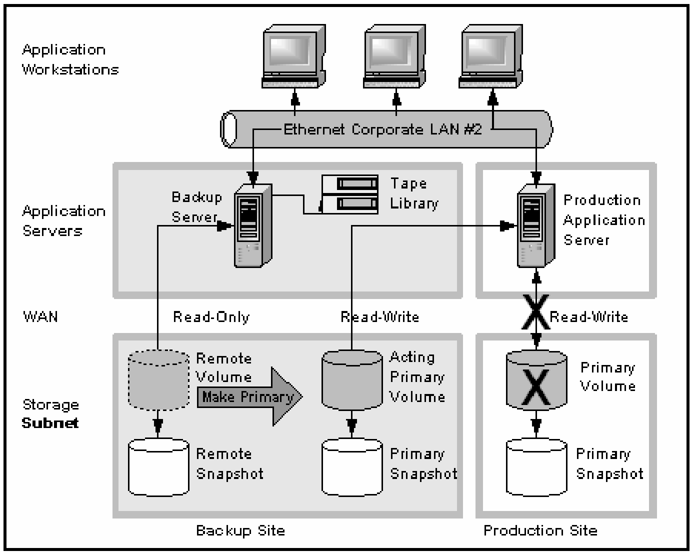

Figure 244. Restoring from a Remote Volume

In Figure 244, note the server labelled Primary Snapshot in the gray area on the left. It originated as a read only back up, but is brought into use as a primary.

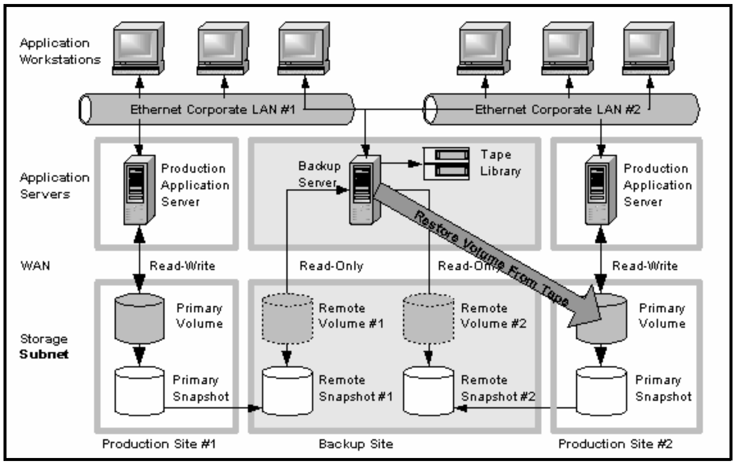

Figure 245. Restoring from Tape Backup

# Best Practices

## Select a Optimum Recurrence Schedule

Select a recurrence schedule for remote snapshots that minimizes the potential for data loss. Any data written to the primary volume since the most recent remote snapshot was created will be unavailable if the primary volume is unavailable. Consider how much data you are willing to lose in the event of an emergency and set the recurrence for creating remote snapshots accordingly.

If you do not want a large number of remote snapshots to accumulate on your remote volume, you can use more than one remote snapshot schedule, each with different retention policies. For example, suppose you want to create remote snapshots every 4 hours to ensure that no more than 4 hours worth of data is lost in an emergency. In addition, you want to retain 1 week's worth of remote snapshots. Retaining 4-hour snapshots for 1 week can result in the accumulation of over 40 remote snapshots. Another approach would be to create 2 remote snapshot schedules for the volume:

- One schedule to create remote snapshots every 4 hours, but only retain the most recent 3 remote snapshots. This will ensure that you do not lose more than 4 hours worth of data in an emergency.
- A second schedule to create remote snapshots every 24 hours and retain 7 remote snapshots.

## Use Remote Snapshots in Conjunction with Local Synchronous Volume Replication

To prevent data loss, reinforce Remote Copy with synchronous replication of the volume within the cluster of SSMs at the primary geographic site. With synchronous replication, a single SSM can be off-line, and your primary volume will remain intact.

At the backup location, you can also use synchronous replication to protect your remote volume against SSM failure.

## Example Configuration

- In the production location, create a cluster of 3 SSMs, all with managers.
- Create volumes on the cluster, and set the replication level to 2.
- Create a schedule for making remote snapshots of the primary volume. Set the recurrence to every 4 hours, and retention of remote snapshots to 2 days.

**Note:** *You can use the same volume replication configuration on the remote volume as well. However, this replication is configured independently of the volume replication configured on the primary volume.*

If one of the SSMs in the primary location fails, the primary volume will still be available. If all of the SSMs fail, or if the application server fails, then you can recover data from the remote snapshots or tape backups in the off-site location.

# Using Remote Copy for Off-site Backup and Recovery

For backup and recovery systems, Remote Copy can eliminate the backup window on an application server. Using iSCSI command line interface commands and scripts, configure the iSCSI initiator to mount remote snapshots on a backup server (either local or remote), and then back up the remote snapshot from the backup server. The remote snapshot is available if the primary volume fails.

## Achieving Off-site Tape Backup

Rather than creating tape backups and then transporting them to a secure off-site location, you can use Remote Copy to create remote snapshots in an off-site location and then create tape backups at the off-site location.

## Configuration for Off-site Backup and Recovery

To use remote snapshots for off-site tape backup, create remote snapshots for access by your tape backup application:

- Create remote volumes in your backup location.
- Configure your backup application to access the remote snapshots.
- Configure schedules to create remote snapshots in the designated off-site locations. If your application server uses multiple volumes that must be in sync, use a script to quiesce the application before creating remote snapshots.
- [Optional] Create routine tape backups of the remote snapshots.

See the example configuration illustrated in Figure 246.

## Configuration Diagram

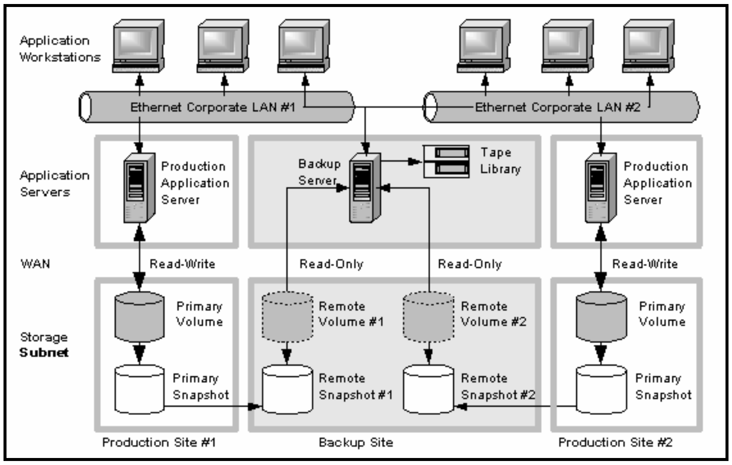

**Figure 246. Off-site Backup and Recovery Example Configuration**

# How This Configuration Works for Off-site Tape Backup

Depending on how long you retain the copies of the remote snapshots, you can retrieve data directly from recent remote snapshots rather than going to tape backups. Otherwise, retrieve data as you normally would from the tape backup.

# Best Practices

## Retain the Most Recent Primary Snapshots in the Primary Cluster

By keeping snapshots on your primary volume, you can quickly roll back a volume to a previous snapshot without accessing off-site backups.

- When you create a schedule for Remote Copy, you specify a number of primary and remote snapshots that you want to retain. You can retain primary snapshots to facilitate easy rollback of the primary volume. (Retention of snapshots will affect the amount of space that is used in the cluster of SSMs, so balance the number of snapshots to retain with the amount of space you are willing to use. To roll back to a snapshot that you did not retain, you can still access remote snapshots or tape backups.)
- Retain remote snapshots in the backup location to facilitate fast recovery of backed up data. If you retain a number of remote snapshots after a tape backup is created, you can access this data without going to the backup tape.

## Example Configuration

- Retain 3 primary snapshots. This enables you to roll the primary volume back, yet it requires a relatively small amount of space on the primary cluster.
- Retain up to a week's worth of remote snapshots on the backup cluster.
- For snapshots older than 1 week, go to the backup tape.

# Achieving Non-Destructive Rollback

As discussed in "Rolling Back a Primary Volume" on page 346, rolling a snapshot back to a volume deletes any snapshots that were created since the snapshot that you roll back to. For example, suppose you created snapshots of a volume on Monday, Tuesday, and Wednesday. On Thursday, if you roll the volume back to Monday's snapshot, then the snapshots from Tuesday and Wednesday will be deleted.

You can use Remote Copy to roll a volume back to an old snapshot without losing the interim snapshots. Because Remote Copy creates two sets of snapshots—primary snapshots and remote copies—you can roll a volume back to a snapshot and still retain the other set of snapshots.

# Configuration for Non-Destructive Rollback

To use remote snapshots for non-destructive rollback:

- Create a remote snapshot schedule.
- In the schedule, specify the same retention policy for the primary and remote snapshots. This ensures that you have copies of the same number of snapshots in your primary and remote locations. Any snapshots destroyed during rollback of one volume will remain intact on the other volume.

See Figure 247 for an illustration of this configuration.

## Configuration Diagram

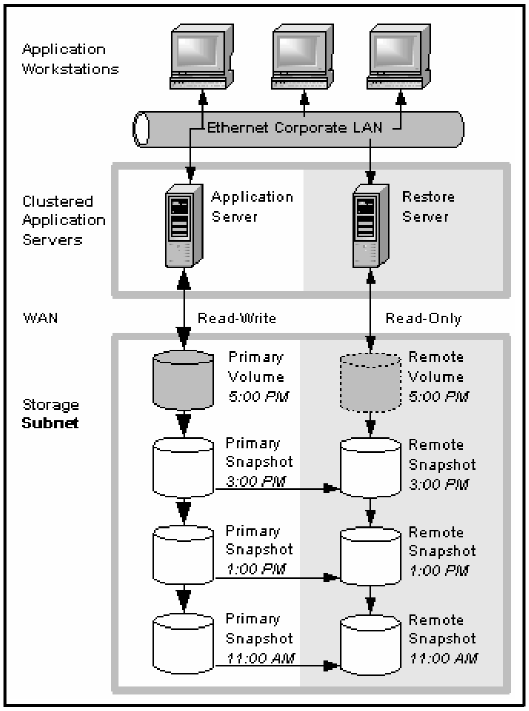

**Figure 247. Non-destructive Rollback Example**

# How This Configuration Works for Non-Destructive Rollback

You can choose to roll back either the primary snapshot or the remote snapshot. Rolling back one of the snapshots will cause all the more recent snapshots of that volume to be deleted. The other volume retains the full set of snapshots. You can continue to make snapshots even though one side was rolled back and the other side was not.

When deciding whether to roll back the primary or remote volume, consider the following:

- When you roll back the primary snapshot to a primary volume, any applications accessing the primary volume will no longer have access to the most current data (as the primary volume has been rolled back to a previous state). If the primary volume must be synchronized with other volumes accessed by the same application, consider rolling back the remote volume instead. Figure 248 shows rollback of the primary snapshot while leaving the remote snapshots intact.

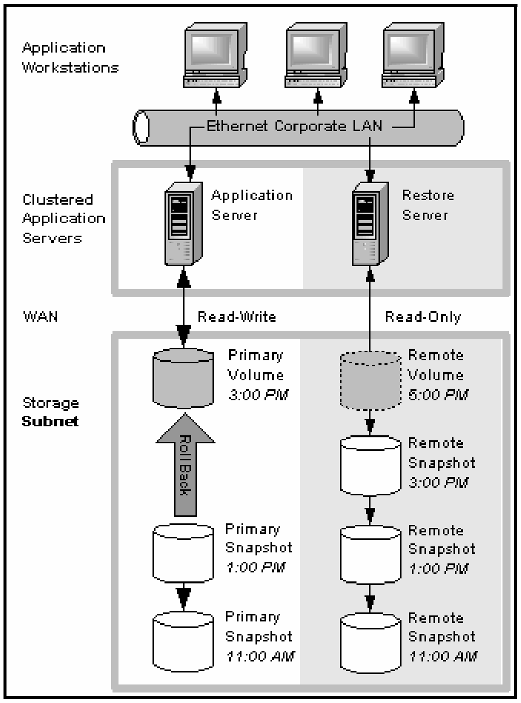

Figure 248. Non-destructive Rollback from the Primary Snapshot

- To roll back the remote snapshot, you must first make the remote volume into a primary volume. This will stop scheduled creation of remote snapshots, which may jeopardize your high availability, disaster recovery, or routine backup strategies. Figure 249 shows rollback of the remote snapshot.

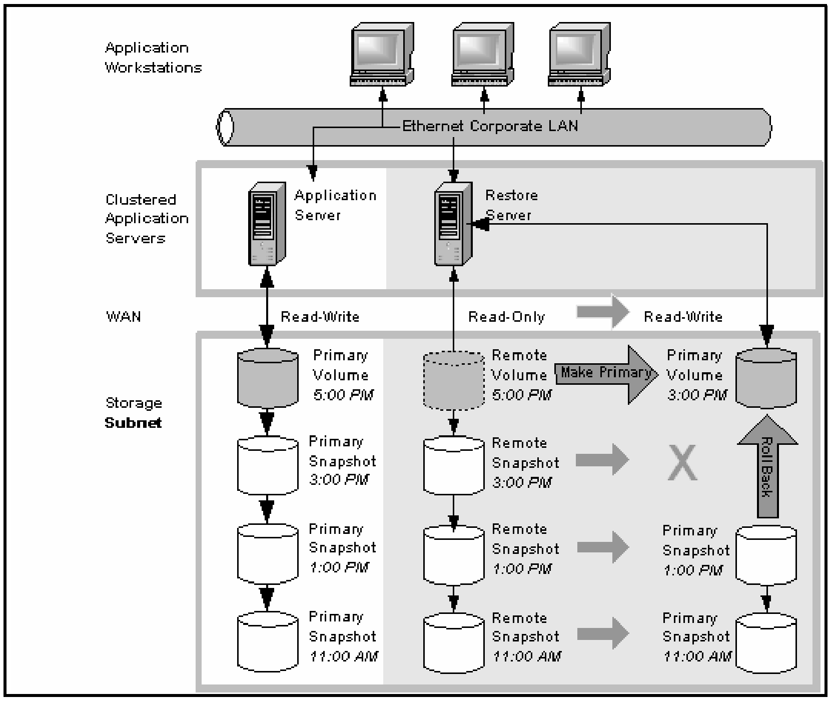

Figure 249. Non-destructive Rollback from the Remote Snapshot

# Best Practices

## Roll Back the Primary Snapshot and Keep the Remote Snapshots as a Backup

To ensure that Remote Copy continues to operate, roll back the primary volume as follows:

1. Preserve the current state of the primary volume that you want to roll back by creating a one-time (manual) remote snapshot of it.
2. Roll back the volume.

   Before roll back, scheduled remote snapshots fail. After the primary volume is rolled back, scheduled creation of remote copies will resume correctly.

   Completed remote snapshots remain intact.

# Using Remote Copy for Data Migration

Remote Copy allows migration of data from one application server to another without interrupting the production application server. This capability supports a number of uses such as data mining or content distribution.

# Achieving Data Migration

You can use Remote Copy to make a complete copy of one or more volumes without interrupting access to the original volumes. This type of data migration allows you to copy an entire data set for use by a new application or workgroup.

To copy data from one location to another, simply create a one-time remote snapshot of the volume. To make the remote snapshot a read/write volume, make it into a primary volume.

# Configuration for Data Migration

To make a copy of a volume in a remote location, configure a cluster of SSMs in the remote location with enough space to accommodate the volume. See the example illustrated in Figure 250.

## Configuration Diagram

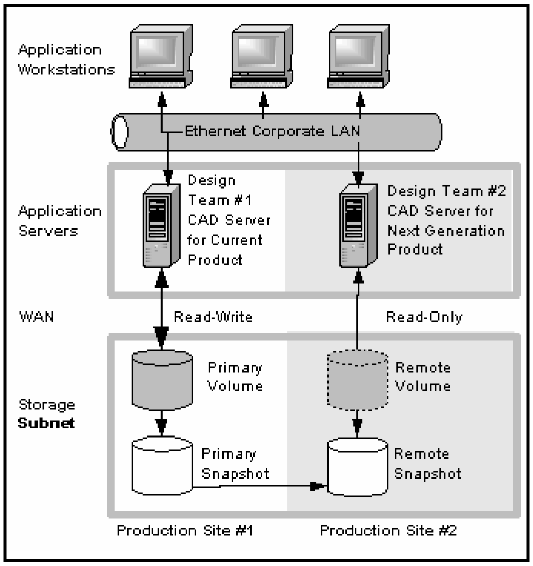

**Figure 250. Data Migration Example Configuration**

# How This Configuration Works for Data Migration

Suppose you want to create a complete copy of a volume for an application to use in different location.

1. Configure a cluster of SSMs in the new location to contain the copied volume.
2. Create either a one-time remote snapshot of the volume onto the cluster in the new location.

If your application server uses multiple volumes that must be in sync, use a script to quiesce the application before creating remote snapshots.

[Optional] You can create regular one-time snapshots and use remote copy to move the snapshots to the remote cluster at your convenience.

3. On the cluster in the new location, make the remote volume into a primary volume.
4. Configure the application server in the new location to access the new primary volume.

   Figure 251 shows migration of data by making a remote volume into a primary volume.

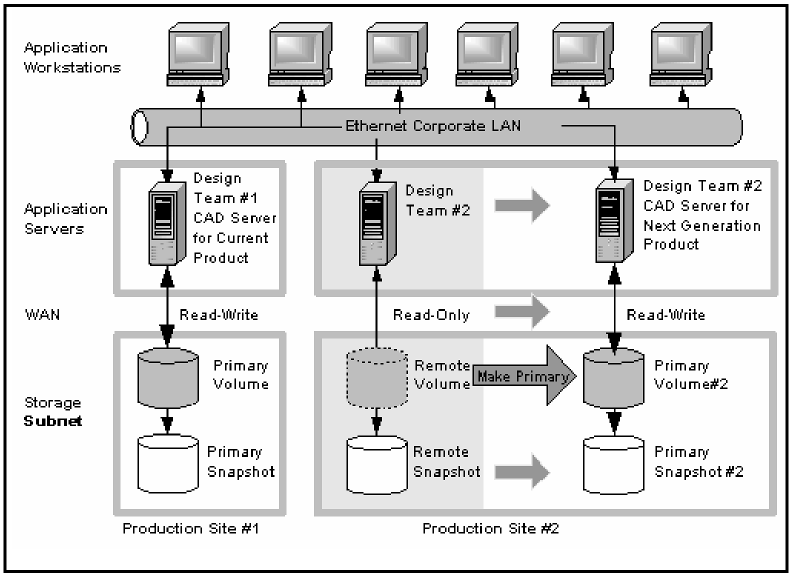

Figure 251. Configuration after Data Migration