Intel® Storage System Module SSR316MJ2

User Manual

September 2004

Order Number: C75954-001

Contact your local Intel sales office or your distributor to obtain the latest specifications and before placing your product order.

Copies of documents which have an ordering number and are referenced in this document, or other Intel literature may be obtained by calling 1-800-548-4725 or by visiting Intel's website at http://www.intel.com.

Contents

Figures

Tables

Revision History

Date	Revision	Description
September 2004	001	Initial release

Getting Started 1

1.1 Storage Server Console Overview

Welcome to the Storage Server Console (Console). Use the Console to configure and manage storage volumes spanning clustered Storage Server Modules (SSMs).

This User Manual provides instructions for configuring individual Storage Server Modules, as well as creating volumes that span a cluster of multiple SSMs. Topics in this manual include:

Configuring individual Storage Server Modules

- Configuring RAID
- Configuring SSMs on the network
- Configuring monitoring and reporting

Creating volumes that span a cluster of SSMs

- Creating management groups and clusters
- Creating volumes that span multiple SSMs
- Designating application server access to volumes
- Creating snapshots of volumes
- Creating and using snapshots of volumes

1.2 Configuration Overview

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

1.2.1 Configuration Tasks

Complete the following tasks to configure SSMs and create clusters and volumes.

Table 1. List of Configuration Tasks

Complete This Task	Find Detailed Information In
Search for one or more SSMs on the network.	"Finding Storage Server Modules on the Network" on page 11.
Log in to the SSMs you want to work with.	"Logging In to the Storage System Module" on page 24.
Configure individual SSMs.	"Configuring Storage Server Modules" on page 20.
Create one or more management groups.	Chapter 9, "Working with Management Groups."

Table 1. List of Configuration Tasks (Continued)

Complete This Task	Find Detailed Information In
Create one or more clusters.	Chapter 11, “Working with Clusters.”
Create one or more volumes.	Chapter 12, “Working with Volumes.”
Configure access to volumes	Chapter 15, “Working with Authentication Groups.”

1.3 Using the Storage Server Console

The Storage Server Console is the storage administrator’s tool for

- configuring and managing the SSM, and
- for creating and managing clusters and volumes.

1.3.1 The Storage Server Console Main Window

The Console main window presents a 3-pane view, as shown in Figure 1.

- **Network View** - the top pane displays all the SSMs on the network. The graphic display indicates the configuration of management groups, clusters and volumes.
- **Tab View** - The center pane presents the functions associated with the selected item in the Network View.
- **Alert Message View** - The bottom pane is a message area where alerts arrive from the active monitoring of the SSM. The most recent alert is at the top, and the messages are continuous while the Console is open. When you close the Console the messages are cleared.

Other features of the Console include the following:

- **Menu Bar** - The menu bar provides access to the following menus:
  - Find - use to find modules on the network.
  - Tasks - access all available storage configuration tasks (tasks are also accessible through right click menus and from the Tab View pane).
  - View - change the icon view in the Console. Large and small versions of the icons are available.
- **Online Help** is provided through the Help menu. Click Help and select Help Topics. Context sensitive Help is available by clicking the question mark icon on individual windows.

Figure 1. Features of the Console Main Window

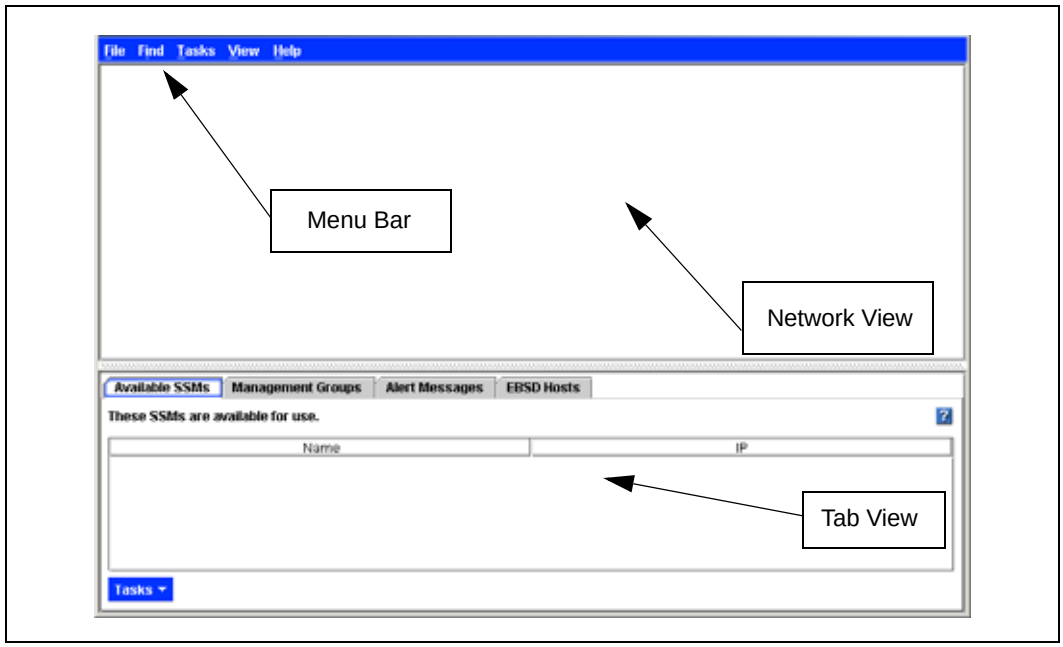

1.3.2 Using the Network View

The Network View displays SSMs according to the criteria you set in the Find function. The graphics displayed in the Network View provide information about the following:

- SSMs
- Management groups
- Clusters
- Volumes
- Snapshots

Figure 2. Viewing all the Features in Network View

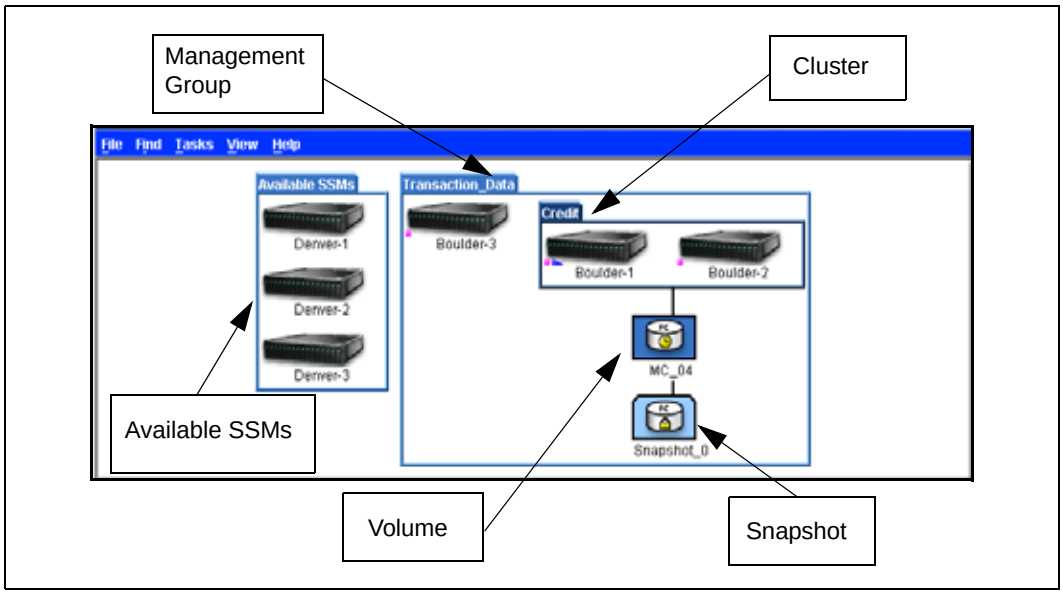

SSM Status

The Network View graphically depicts the status of each SSM. SSMs on the network are either available or part of a management group. If SSMs are running an older version of the Storage System Engine, they will display in the Console as "Downlevel."

Other graphical information in the Network View depicts the storage architecture you create on your system. An example setup is shown in Figure 2.

- **Management Groups.** Management groups are groups of SSMs within which one or more SSMs are designated as managers.
- **Clusters**. Clusters are sub-groupings of SSMs within a management group.
- **Volumes**. Volumes are data storage areas created on clusters.
- **Snapshots**. Snapshots are read-only copies of volumes created at specific points in time.

Log In Status

If you are logged into the module, a pink square displays underneath the SSM in the Network View.

1.3.2.1 Icons Used in the Storage Server Console

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.

Graphical Legend Items

The Graphical Legend has three tabs.

- The Items tab, shown in Figure 3, displays the icons used to represent physical or virtual items displayed in the Console. For example, SSMs are physical items and management groups are virtual items.

Figure 3. Viewing the Graphical Legend Items tab Available from the Help Menu

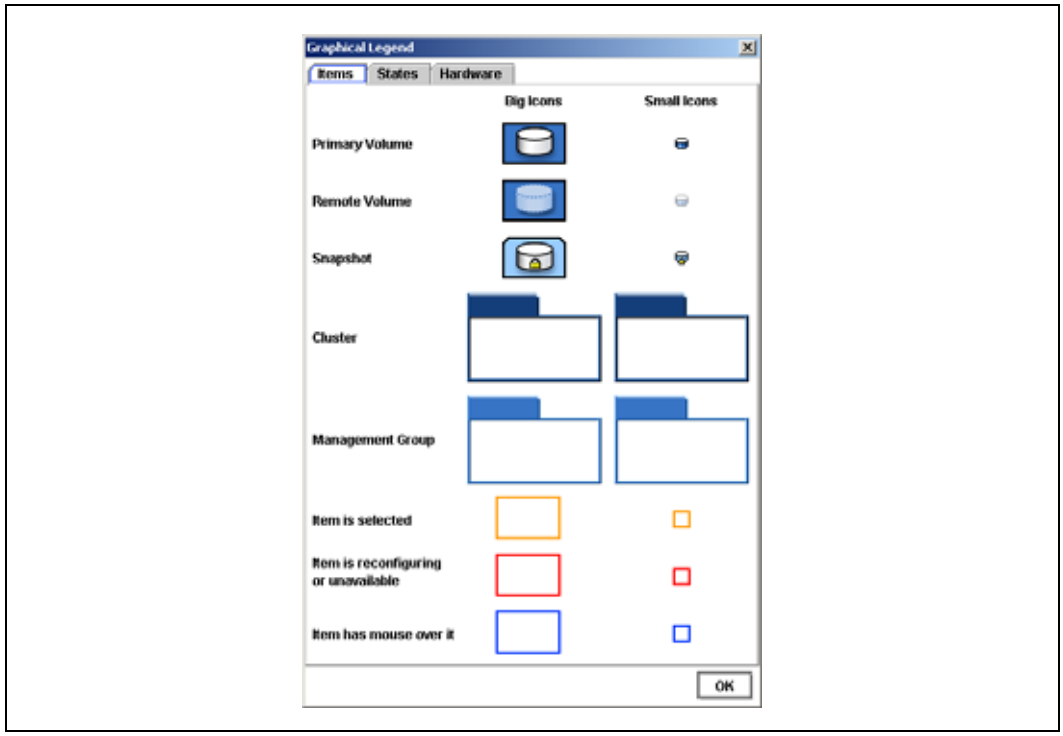

- The States tab, shown in Figure 4, displays the icons used to depict states that the items are in. For example, when you are logged into an SSM, a pink square displays underneath the SSM. When an item such as an SSM or a cluster is selected, it displays a yellow outline.

Figure 4. Viewing the Graphical Legend States Tab Available from the Help Menu

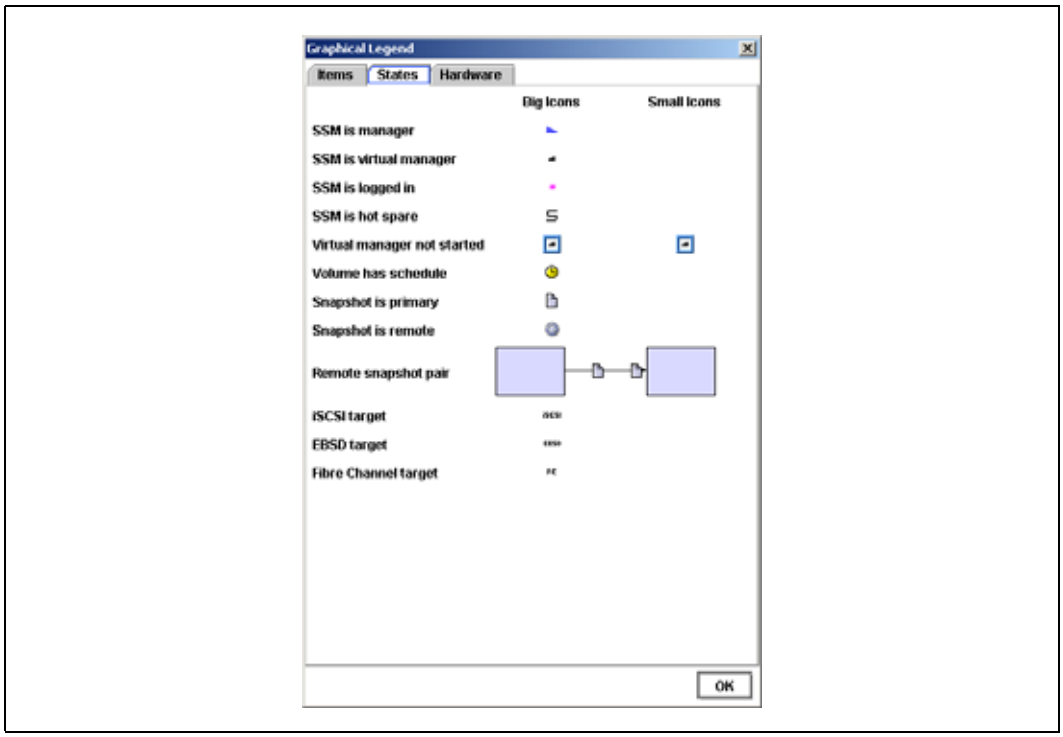

- The Hardware tab, shown in Figure 5, displays the storage unit.

Figure 5. Viewing the Graphical Legend Hardware Tab Available from the Help Menu

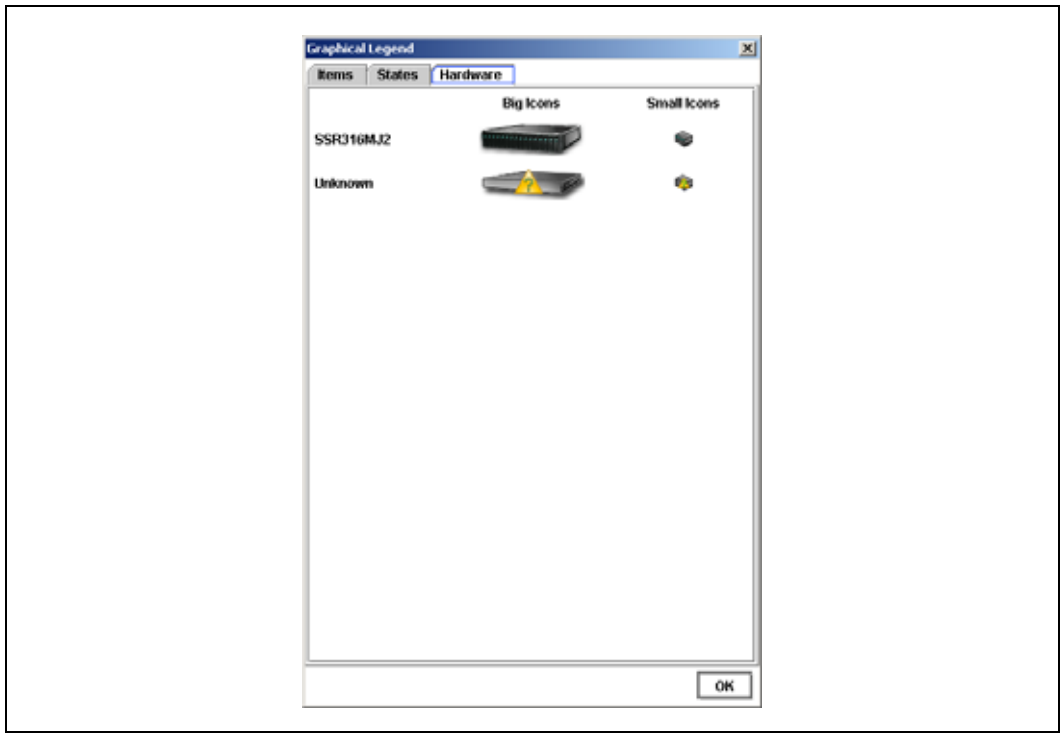

1.3.3 Using the Tab View

The Tab View displays properties of the item selected in the Network View. For example, Figure 6 shows the tabs that display when the SSMs on the network are found.

Select a tab to perform functions related to the selected item.

Figure 6. Tab View in the Main Window

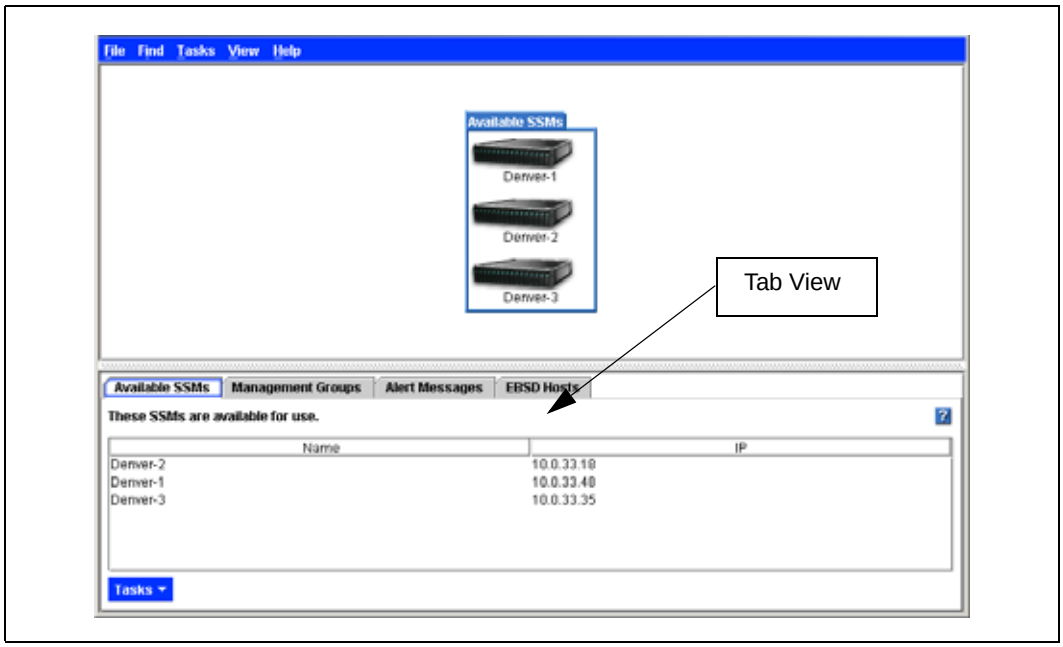

1.3.4 Network Tab View

The network tabs provide access to the groups of SSMs and the management groups that are configured on your network.

1.3.4.1 Available SSMs Tab

The Available SSMs tab, shown in Figure 7, lists the SSMs in the Network View pane that are available — that is, are not part of a management group. The SSMs are listed in ascending order of their respective IP addresses.

Figure 7. Available SSMs Tab with SSMs Listed

1.3.4.2 Management Groups Tab

The Management Groups tab, shown in Figure 8, lists all the management groups currently created with the SSMs that are displayed in the Network View pane.

For information on management groups, see Chapter 9, “Working with Management Groups.”

Figure 8. Management Groups Tab from Main Window

1.3.4.3 Alert Messages View

Review any alert messages that appear here. Figure 9 shows the area in which alert messages display. These messages include alerts from the monitoring parameters you set in Reporting for individual SSMs (see “Using Active Monitoring” on page 114 for detailed information about setting reporting parameters).

Figure 9. Viewing Messages in the Alert Messages Tab

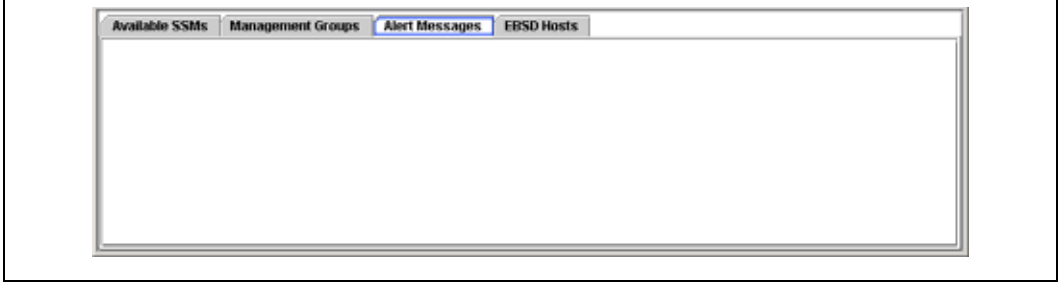

1.3.4.4 EBSD Hosts Tab

The EBSD Hosts tab, shown in Figure 10, lists all versions of the EBSD drivers that are currently installed on the network. It also lists information about the hosts that are using that driver.

See the EBSD Windows Manual for more information about the EBSD drivers.

Figure 10. EBSD Hosts Tab

1.3.5 The SSM Configuration Window

To configure specific settings of an individual module, you use the SSM configuration window, shown in Figure 11.

Figure 11. The SSM Configuration Window

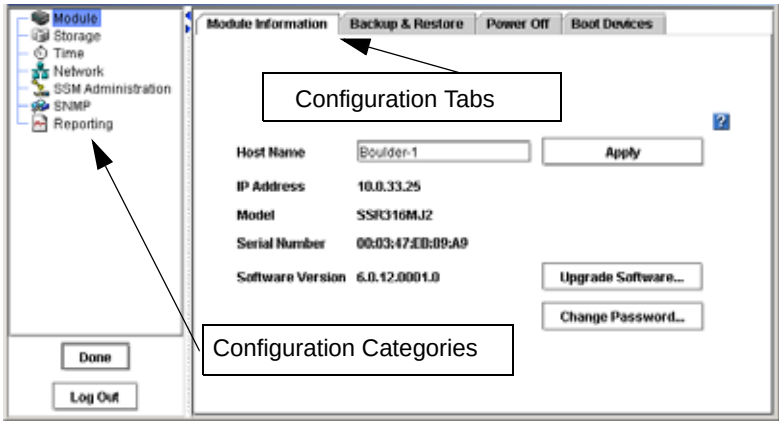

The SSM Configuration window opens when you log into an individual SSM. From the configuration window you have access to all the configuration tasks for individual SSMs.

1.3.5.1 Configuration Categories

The pane on the left of the configuration window lists the configuration categories. Within each category is a set of tabs which you use to configure different functions.

- Module – upgrade the Storage System Engine software, change the password or host name, perform backup and restore of the SSM configuration, reboot or shut down the SSM, and manage boot devices. See Chapter 2, “Working with Storage System Modules.”

- Storage – manage RAID and manage individual drives, including powering them on or off, and reviewing drive information. See Chapter 3, “Storage.”
- Time – use NTP or manually set the time zone, date, and time for the SSM. See Chapter 5, “Setting the Date and Time.”
- Network – specify the TCP/IP settings of the SSM, manage DNS information, manage the routing table, and update the communication mode information if the SSM is running a manager. See Chapter 4, “Managing the Network.”
- SSM Administration – Add, edit, and delete administrative users and groups. See Chapter 6, “Administrative Users and Groups.”
- SNMP – enable SNMP and enable SNMP traps. See Chapter 7, “Using SNMP.”
- Reporting – view real-time statistical information about the SSM and configure selected variables for active monitoring. See Chapter 8, “Reporting.”

1.4 Finding Storage Server Modules on the Network

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

- Search subnets using a mask to find all available SSMs on a network
- Enter specific IPs or host names and connect to modules

1.4.1 Finding Modules the First Time or If No IP Has Been Saved

The first time you open the Console, a search window opens, shown in Figure 12.

Figure 12. Selecting a Search Method for SSMs

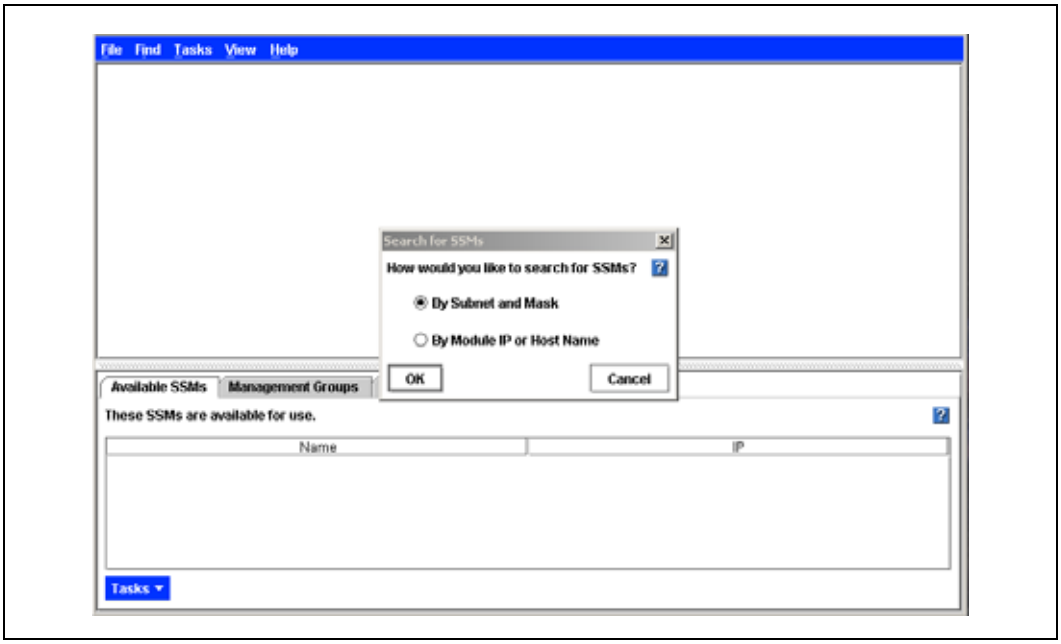

- Select from the list whether you want to find modules by using individual IPs or host names or by using a subnet and mask.

  See "Finding by Subnet and Mask" on page 13 for more information about completing the List of Subnets to Search window.

  See "Finding by Module IP or Host Name" on page 16 for more information about completing the IP and Host Name List window.

1.4.2 Finding Modules On An Ongoing Basis

Once you have entered either a subnet and mask into the List of Subnets to Search window or an IP or host name in the IP and Host Name List window, these settings are saved. Every time you open the Console, the search takes place and a message opens, as shown in Figure 13, listing which SSMs have been found. The window also lists any EBSD hosts that are found on the network.

Figure 13. SSMs Found Message

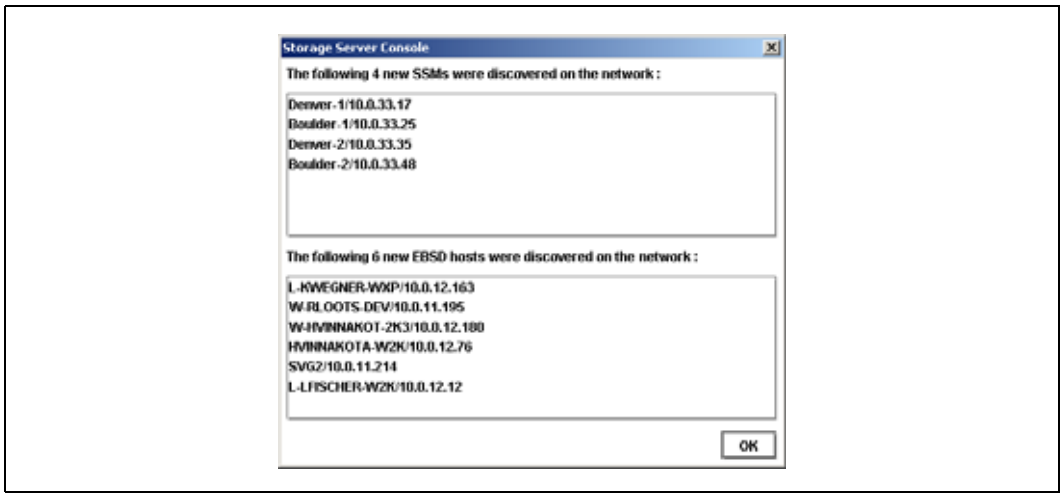

1. Click OK.

   Those SSMs appear in the Network View.

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 lots of traffic, or if a module is busy reading or writing data, it may not be found when a search is performed. This is a function of the search protocol used. See "Search Protocols" on page 13.

- If the module you are looking for does not appear in the Network View pane, perform the subnet and mask search again by selecting that method from the Find menu.
- If the module still does not appear, try using the Find All Modules via TCP option, which is the most predictable method for connecting to an SSM, due to the search protocol used. See "Finding by Module IP or Host Name" on page 16.

Note: If neither of the above methods works, 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.

1.4.2.1 Search Protocols

UDP

When you select the UDP option on either of the Find windows (Subnet and Mask or IP or Host Name), the User Datagram Protocol (UDP) is used. The UDP protocol does not guarantee delivery of the packet to the host. If there is a high volume of activity on a host or network, the UDP packets are liable to get dropped. Therefore, if the host is busy or the network is busy, the UDP search may not find the designated SSM(s). Also, Find by Subnet and Mask uses the IP broadcast address for the selected subnet and may therefore be filtered by routers and possibly switches.

TCP

In contrast to UDP, the Transmission Control Protocol (TCP) guarantees delivery of the packet to the host. Therefore, the Find IP or Host Name via TCP function is able to connect to an SSM that cannot be found by the UDP search when there is a high volume of network traffic or when SSMs are busy reading and writing data.

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

1.4.3 Finding by Subnet and Mask

Find all the SSMs on the network by searching subnets with masks.

1. If this is the first time you have opened the Console, select By Subnet and Mask on the dialog box, then click OK.

   or

   Click the Find menu and click By Subnet and Mask.

   The List of Subnets to Search window opens, shown in Figure 14.

Figure 14. Using Subnet and Mask to Search

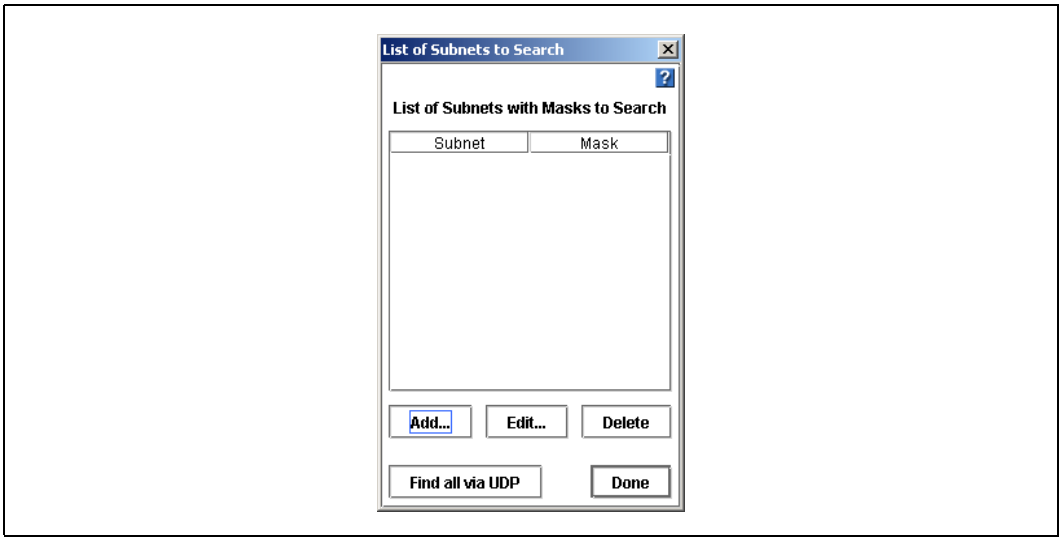

1.4.3.1 Adding Subnets and Masks

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 all via UDP.
   The Active Search window opens, tracking the search process. When the search is complete, the Active Search window closes. The SSM Storage Server Console window opens, listing all the SSMs that were found on the network.
6. Click OK to close the SSM Storage Server Console window.
7. Click Done on the List of Subnets to Search window.
   The modules appear in the Network View, identified by host name.

Figure 15. Viewing SSMs in the Network View Pane

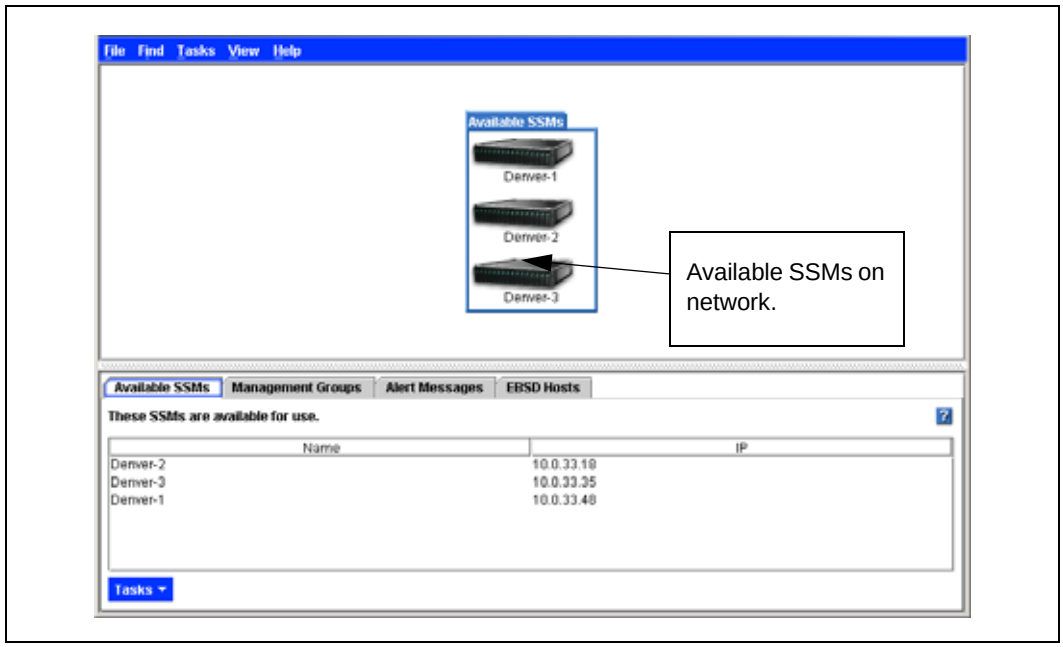

Note: The subnet and mask are saved in the list. Every time you open the Console, the search takes place automatically and all SSMs on the network are listed in the Network View. See "Deleting Subnets and Masks" on page 15 if you want to disable this search.

1.4.3.2 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 the information as necessary.
6. Click OK.

1.4.3.3 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. Exit the Console and open it again to remove the modules from the Network View pane.

1.4.4 Finding by Module IP or Host Name

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

1.4.4.1 Find by IP or Host Name and Network Configuration

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 SSM's eth0 NIC.
- The Console Find is set to Module IP or Host Name using only the IP address of the eth1 NIC.

The SSM is indeed 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 in which networking is configured. The SSM receives the UDP 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.

1.4.4.2 To Find by IP or Host Name

1. If this is the first time you have opened the Storage Server Console, select By Module IP or Host Name at the dialog box, then click OK.
   or
   Click the Find menu and click By Module IP or Host Name.
   The IP and Host Name List window opens, shown in Figure 16.

Figure 16. Using IP or Host Name to Search

1.4.4.3 Adding IPs or Host Names

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.

1.4.4.4 Searching for the Listed Modules

Search using one of two methods to find the listed IP addresses or host names.

1. To find the modules in the IP and Host Name list, click Find all SSMs via TCP.

   or

   Click Find all via UDP.

   Use TCP if your modules reside on a network behind a firewall or if you have other configuration parameters that prevent the UDP protocol from working correctly (see "Search Protocols" on page 13).

1.4.4.5 Editing the IP or Host Name in the Search List

Change the IP or Host Name of a module in the list used to search for modules.

1. Click the Find menu.
2. Click By Module IP or Host Name.

   The IP and Host Name List window opens, shown in Figure 16 on page 17.
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.

1.4.4.6 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 Module IP or Host Name.
   The IP and Host Name List window opens, shown in Figure 16.
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 OK.
7. Exit the Console and open it again to remove the modules from the Network View pane.

1.5 Viewing Storage Server Module Details

Select an SSM from the Network View and the SSM Details tab opens in the Tab View, shown in Figure 17.

Figure 17. Viewing Individual SSM Information

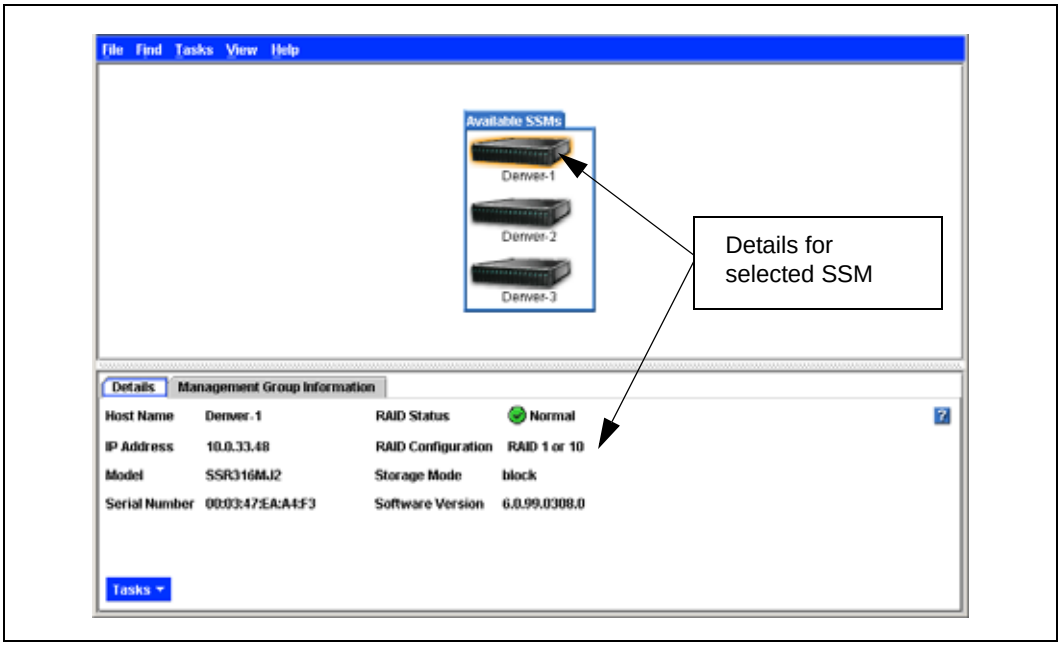

1.5.1 Details Tab

Includes host name, IP address, the model, serial number, RAID status, RAID configuration, storage mode, and software version.

1.5.1.1 RAID States

Three RAID states are reflected on the SSM Details tab.

- If RAID is **normal**, a green circle displays in the SSM configuration details tab when the SSM is selected in the Network View, as shown in Figure 18.

Figure 18. Icon Showing that RAID is Normal

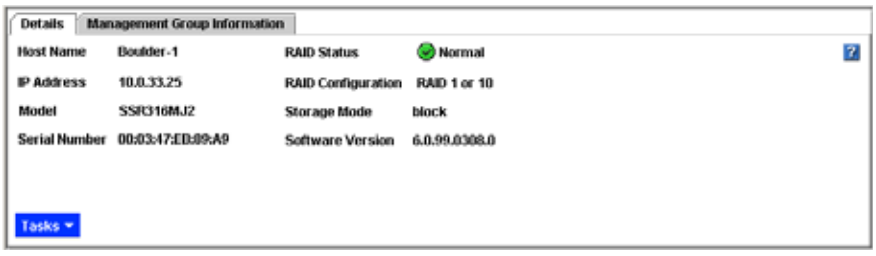

- If RAID is **off**, a red circle displays in the SSM configuration details tab, as shown in Figure 19. For information about turning RAID on, see "Starting RAID" on page 47.

Figure 19. Icon Showing that RAID is Off

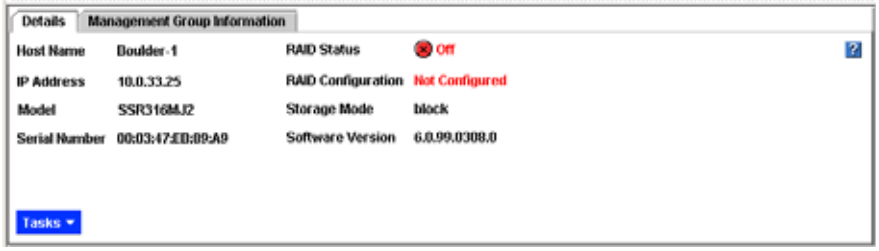

- If RAID is **degraded**, a yellow circle displays, as shown in Figure 20. See "Monitoring RAID Status" on page 48 for information about fixing degraded RAID.

Figure 20. Icon Showing that RAID is Degraded

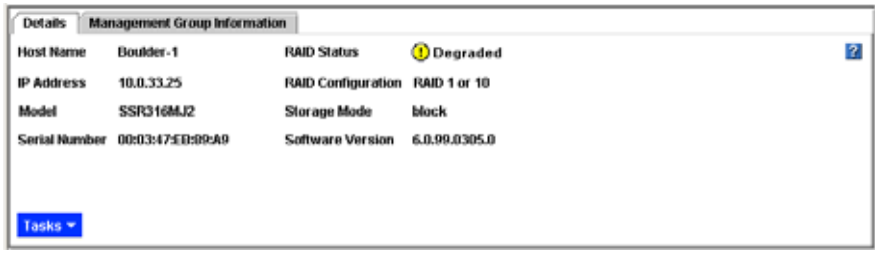

1.5.2 Management Group Information

Includes the name of the management group to which the SSM belongs, the status of the management group, whether the SSM is designated as a hot spare, and whether the SSM is running a manager or a virtual manager.

1.6 Configuring Storage Server Modules

An SSM must be configured before you use it for storage. If you plan to use multiple SSMs, they must all be configured individually before you use them for clustered storage.

The list below details the main items for individual SSMs that must be configured before you create management groups, clusters, and volumes.

Note: When planning the configuration of your SSMs, note that all of the SSMs in a **cluster** must be configured the same way.

- **Changing RAID and verifying SSM disks.** The SSM may be shipped with RAID already configured and operational. Instructions for ensuring that drives in the SSM are properly configured and operating are in Chapter 3, "Storage."

- **Setting date and time for the SSM**. Instructions for setting the date and time are in Chapter 5, “Setting the Date and Time.”
- **Configuring the network.** Configure additional NICs, change IP addresses, and manage DNS. Detailed network configuration instructions are in Chapter 4, “Managing the Network.” .
- **Creating administrative users and groups.** Add administrative users and groups. The SSM comes configured with 2 default groups and 2 default users. Add additional groups and individual users as desired. See Chapter 6, “Administrative Users and Groups.”
- **Using SNMP.** The SSM can be monitored using an SNMP Agent. The SSM Management Information Base (MIB) is read-only and supports SNMP versions 1, 2c, and 3. SNMP instructions are in Chapter 7, “Using SNMP.”
- **Setting monitoring parameters and viewing reports.** The Reporting section provides real-time information about the status of the SSM. You can also set active monitoring with Console alerts as well as e-mail alerts. See Chapter 8, “Reporting.”

Working with Storage System Modules 2

2.1 Storage System Module Configuration Window Overview

The Storage System Module configuration window, shown in Figure 21, opens when you log into an individual Storage System Module. From the configuration window you have access to all the configuration tasks for individual Storage System Modules.

Figure 21. Storage System Module Configuration Window

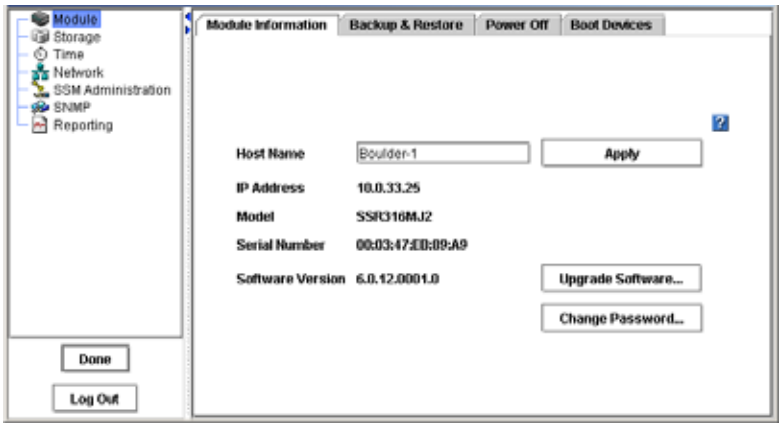

2.1.1 Configuration Categories

The left pane lists the configuration categories. Within each category is a set of tabs which you use to configure different functions.

- **Module** - Use the module category to change the host name and login password for the SSM. You can also backup and restore the Storage System Engine, reboot or shut down the SSM, 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.
- **Network** - For each SSM you can configure and manage the network settings, including TCP/IP interfaces, DNS servers, and the routing table.
- **Storage System Module Administration** - The SSM comes configured with 2 default groups and 2 default users. All administrative users and groups are added and managed locally.
- **SNMP** - The SSM can be monitored using an SNMP Agent. You can also enable SNMP traps.

- **Reporting** - The SSM offers multiple reporting capabilities, including real-time statistical information, active monitoring of selected variables, and diagnostics.

2.2 Module Information Overview

The module category contains tabs that provide access to detailed information about the SSM, backing up and restoring SSM configuration files, and the software reboot or shutdown function.

2.3 Logging In to the Storage System Module

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 module.

1. On the Network view, double-click the SSM that you want to log in to.

   The Log In window opens, shown in Figure 22.

Figure 22. Logging in to an SSM

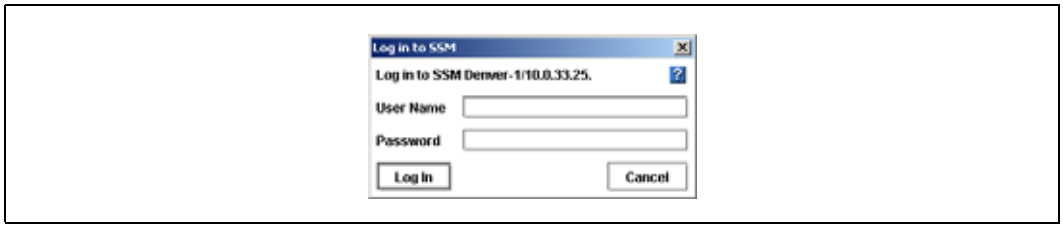

2. Type the User Name and Password.
3. Click Log In.

   When you are successfully logged in to the module, the SSM configuration window opens to the Module Information tab, shown in Figure 23.

Figure 23. The Module Information Tab

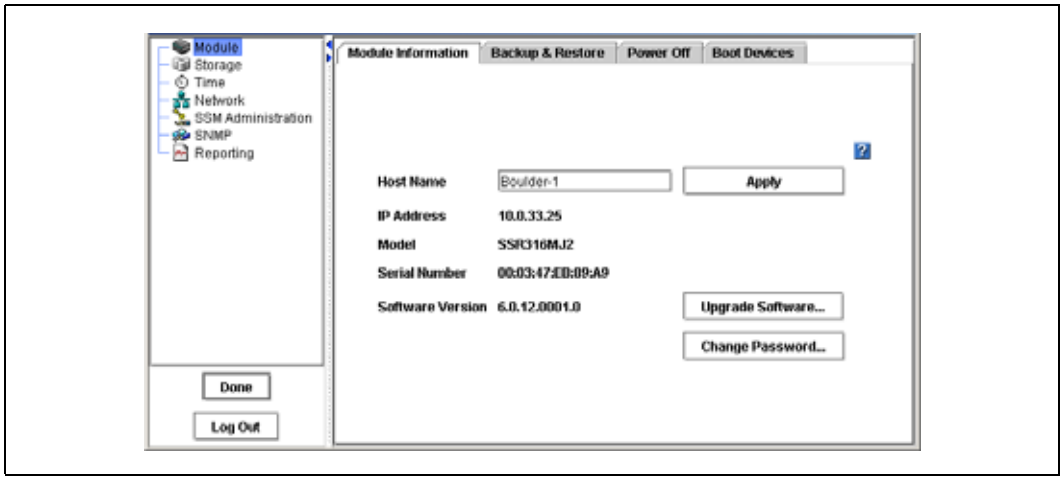

2.4 Logging Out

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

1. When the SSM Configuration window is open, click Log Out.

   The Network view opens and the SSM you logged out of no longer displays the logged in icon — the pink square.

2.4.1 Closing the Storage System Module Configuration Window without Logging Out

Clicking Done on the Module window returns the Console to the Network View and leaves you logged in to the SSM. The Network View window displays the SSM with a pink square underneath, indicating the logged in status.

2.5 Changing the Storage System Module Host Name

The host name on an SSM is the name that displays below the SSM icon in the Network View. It also is the name that end users of the module see when they browse the network. Change the host name of the SSM on the Module Information tab.

The SSM may come configured with a default host name of “none.”

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

   If you are operating in a Windows environment, the host name should be 15 characters or fewer.
3. Click Apply.

   A confirmation message opens.
4. Click OK.

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

2.6 Changing Passwords

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

1. Log in to the SSM.
2. On the Module 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.

2.7 Upgrading the Storage System Engine

When you upgrade the Storage System Engine, the version number will change. You can view the current software version on the Module Information tab in the SSM configuration window or on the Details tab in the Network View.

Upgrade the software on a SSM only when an upgrade or patch is released. The Storage System Engine upgrade/installation takes about 5 to 8 minutes, including the SSM reboot.

Note: The SSM must contain both boot flash cards in order to upgrade the Storage System Engine. See "Configuring Boot Devices" on page 33.

1. Download and install the upgrade file from the web site of your approved supplier or from a CD.

Table 2. Patch and Upgrade Installation Options

Downloading and Installing the Upgrade File	Upgrading from an Installation CD
1. From the website of your approved supplier, copy **install.htm** and the **InstData** folder to your hard drive. 2. Navigate to **install.htm** on your hard drive. 3. Double-click **install.htm**. 4. Complete the installation wizard. 5. Continue with step 2 below.	1. Insert the CD in your CD drive. The installation wizard should automatically begin. • If the wizard does not start, navigate to the CD drive and double-click the installation executable. 2. Complete the installation wizard. 3. Continue with step 2 below.

Installing the upgrade or patch places a file in an upgrade or patch folder. For example,

```
C:\Program Files\Storage System\6.0\Upgrade
```

or

```
C:\Program Files\Storage System\6.0\Patch
```

2. Log in to the first Storage System Module you want to upgrade.
3. On the Module Information tab, click Upgrade Software.
   The Upgrade Software window opens, shown in Figure 24.

Figure 24. Upgrading the Storage System Module Software

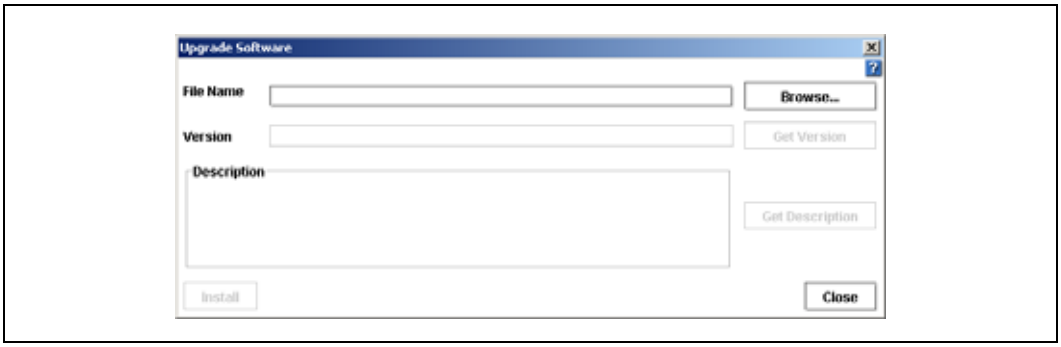

4. Browse for the upgrade or patch file that was installed in Step 1., as shown in Figure 25.

Figure 25. Browsing for the Upgrade or Patch File

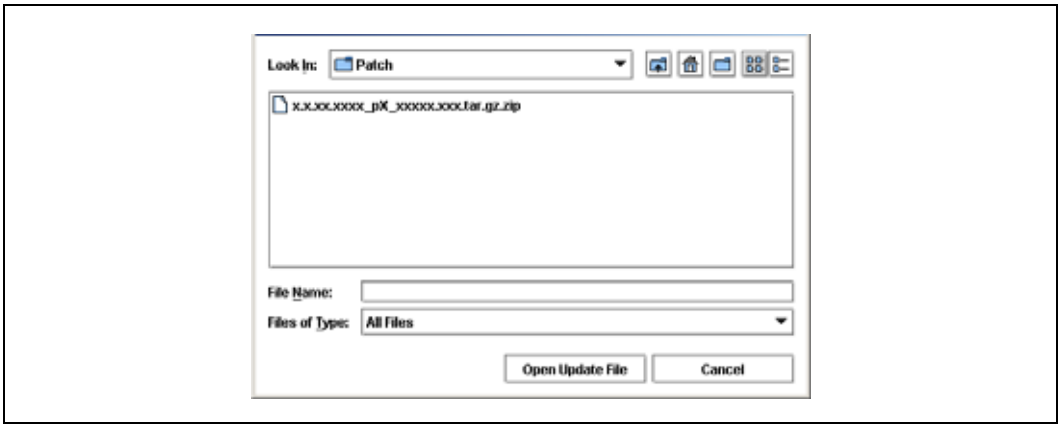

5. Select the file and click Open Update File.

   Focus returns to the Upgrade Software window. When the file name is present, the Install button becomes enabled. The description and version information display in the window.

6. Click Install.

   The upgrade status window opens, shown in Figure 26. Status messages scroll on the window. These messages can be saved to a file.

   — [Optional] Click Save To File and choose a name and location for the file.

Figure 26. Upgrade Status Messages

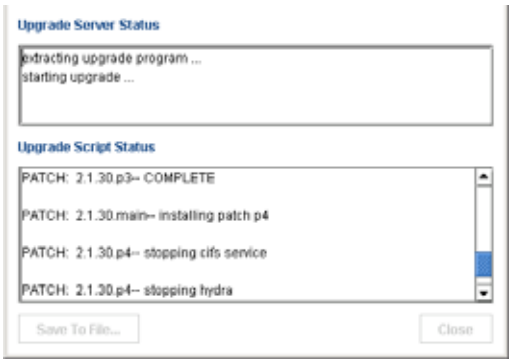

7. Click Close when the installation is completed.
8. Repeat steps 2 through 7 for each Storage System Module you want to upgrade.

Note: The Storage System Module will reboot as part of the upgrade process. Search for the Storage System Module and ensure that it appears in the Console before upgrading the next Storage System Module.

2.8 Backup and Restore of Storage System Module Configuration

Backup and restore provides the capability to save the Storage System Module configuration file for use in case of a Storage System Module 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 Storage Server 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 and 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 Management Group Configuration" on page 140.

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

Note: If you have multiple SSMs with the same configuration, you can create a single configuration backup file and use it to restore the configuration on any of the SSMs. Any SSM that you restore from the backup file will have exactly the same configuration. For example, if you back up the configuration of an SSM that has a static IP address, and then restore that configuration to a second SSM, the second SSM will have the same IP address. You must manually change the IP address on the second SSM.

1. Log in to the Storage System Module.
2. Click the Backup & Restore tab.

The Backup & Restore window opens, shown in Figure 27.

Figure 27. Viewing the Backup and Restore Window

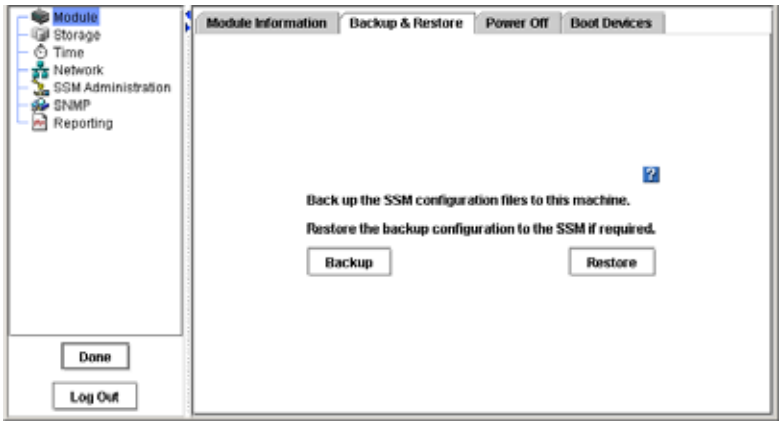

2.8.1 Backing Up the Storage System Module Configuration File

Use Backup to save the Storage System Module configuration file to a directory on your local machine.

1. Click Backup.

   The Save Backup File window opens, shown in Figure 28.

Figure 28. Backing up the Storage System Module Configuration File

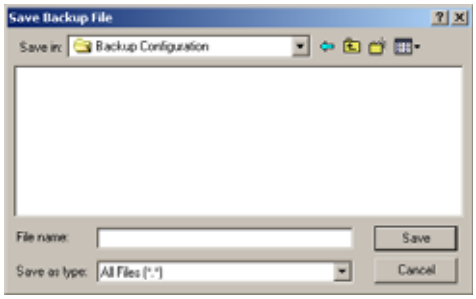

2. Navigate to a folder on the Storage Server Console computer to contain the SSM configuration backup file.
3. Accept the default name (SSM_Configuration_Backup) or enter a new name for the backup file.

Note: The configuration files for all SSMs that you back up are stored on the computer running the Storage Server 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 an SSM.

4. Click Save.

2.8.2 Restoring the Storage System Module Configuration File

Use Restore to restore the configuration of an SSM after it fails or configuration files are lost.

1. On the Backup and Restore tab, click Restore.
   The Restore Storage System Module window opens, shown in Figure 29.

Figure 29. Restoring the Storage System Module Configuration File

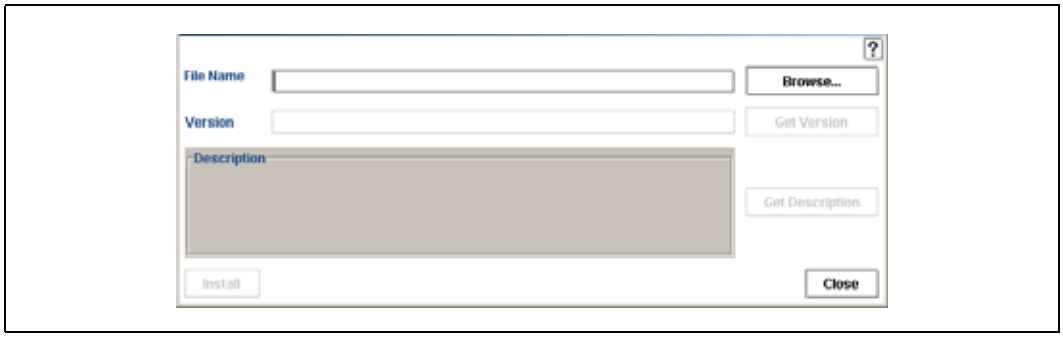

2. Click Browse to navigate to the folder on the Storage Server Console computer where the configuration backup file is saved.
3. Select the file to restore.
4. Review the version and description to ensure you are restoring the correct file.
5. Click Install to install the configuration file.
   The Upgrade 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.

Figure 30. Restoring the Storage System Module Configuration File

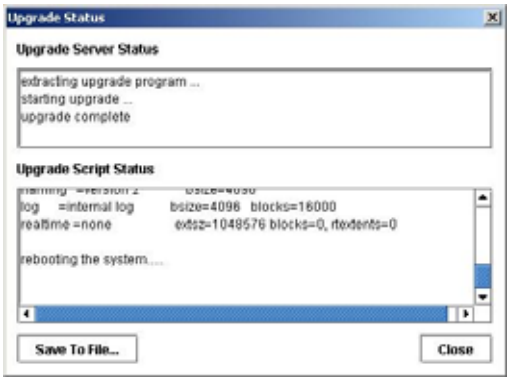

6. Click Close to finish restoring the configuration.
   The SSM reboots and the configuration is restored to the identical configuration as that on the backup file.

2.8.2.1 Completing the Restore

After you restore the SSM configuration from a file, three manual configuration steps may be required:

- You must manually configure RAID on the SSM.
- Restoring an SSM configuration file from one SSM to a second SSM does not restore network routes that were configured on the SSM. You must manually add network routes after the restoration.
- Any Storage System Modules that you restore from the backup file will have exactly the same configuration. For example, if you back up the configuration of an SSM that has a static IP address, and then restore that configuration to a second SSM, the second SSM will have the same IP address. You must manually change the IP address on the second SSM.

2.9 Rebooting the Storage System Module

Reboot the Storage System Module from the Storage Server Console without turning the module off. Set the amount of time before the reboot begins to ensure that the activity to the module has stopped.

1. Log in to the Storage System Module.
2. Click the Power Off tab.

   The Power Off window opens, shown in Figure 31.

Figure 31. Shutting Down or Rebooting the Storage System Module

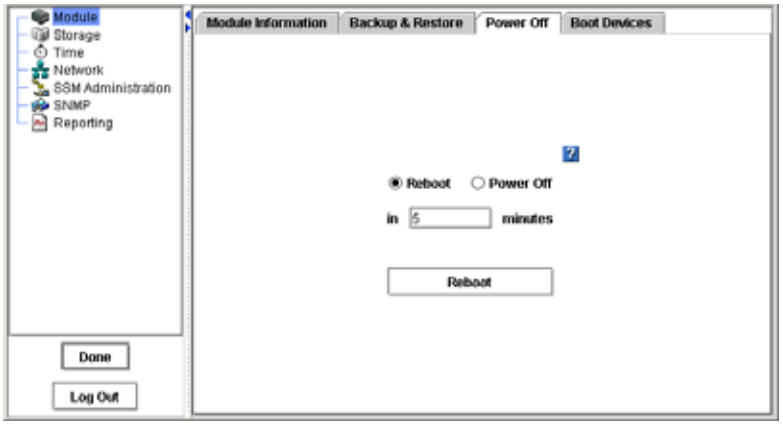

3. Select Reboot.
4. In the minutes field, type the number of minutes before the reboot should begin.

   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 6.
5. Click Reboot.

   A confirmation message appears.
6. Click OK.

The module will restart in the specified amount of time. When reboot actually begins, the module disappears from the Network View. The reboot takes 3 to 4 minutes.

7. Search for the module to reconnect the Storage Server Console to the module once it has finished rebooting.

   See "Finding by Subnet and Mask" on page 13 or "Finding by Module IP or Host Name" on page 16.

2.9.0.1 Canceling a Reboot

1. When you click OK, the Reboot button changes to Cancel, as shown in Figure 32.

Figure 32. Canceling the Storage System Module Reboot

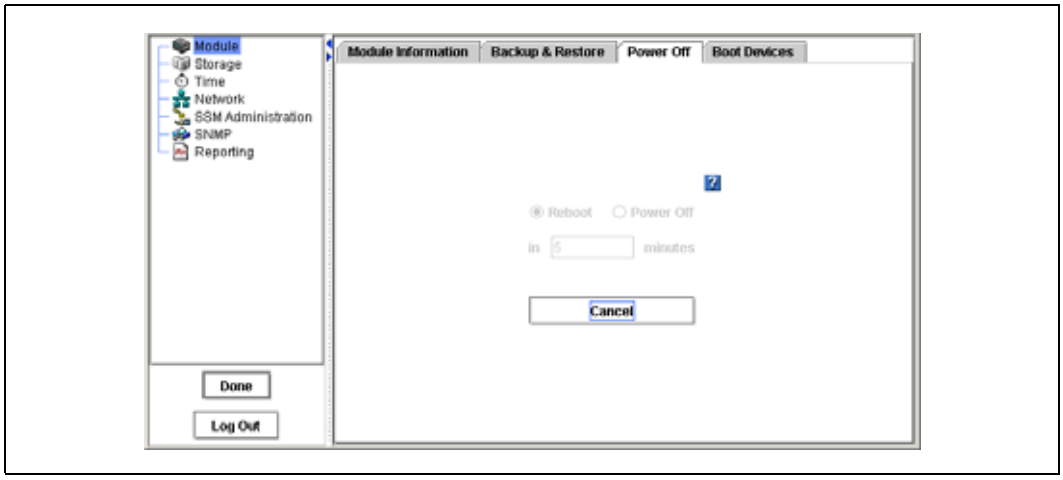

2. Click Cancel to stop the reboot.

   A confirmation message opens.

3. Click OK.

   The Shut Down window returns with the original settings, such as those in Figure 31.

2.10 Powering Off the Storage System Module

Powering off the Storage System Module through the Storage Server Console physically turns the module off. The Console controls the power down process so that data is protected.

1. Log in to the Storage System Module.
2. Click the Shut Down tab.

   The Shut Down window opens, shown in Figure 31.

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.

   You can enter any whole number greater than or equal to 0. If you enter 0 the SSM will power off as soon as you complete step 6.

5. Click Power Off.

A confirmation message appears.

6. Click OK.
   The module will power down in the specified amount of time.

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

2.11 Configuring Boot Devices

When the Storage System Module powers on or reboots, it references boot configuration information from one of two boot flash cards, located on the front of the module.

The SSM boot configuration information is mirrored between the two boot flash cards. If one card fails or is removed, the Storage System Module can still boot. If you remove and replace one of the cards, you must activate the card to synchronize it with the other card.

Note: There must always be at least one active flash card in the Storage System Module. If you are upgrading the Storage System Engine software, the module must contain both flash cards.

2.11.1 Checking Boot Device Status

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

1. Log in to the Storage System Module.
2. Click the Boot Devices tab.
   The Boot Devices window opens, shown in Figure 33.

Figure 33. Activating Boot Devices

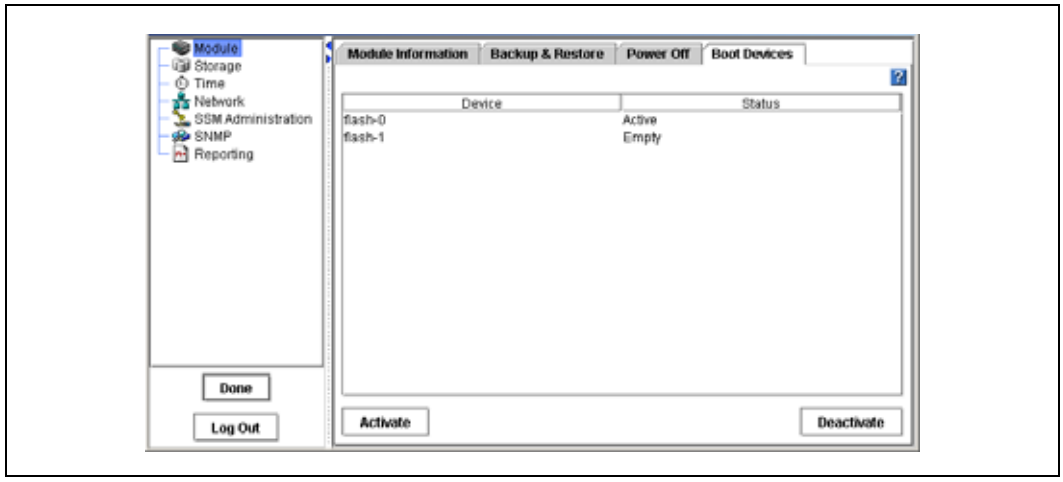

3. The status of each boot flash card is listed in the Status column.

Table 3. Boot Flash Card Status

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. Note: If a flash card has a status of Failed, select the card and click Activate. If the card fails repeatedly, it needs to be replaced.
Empty	The flash card bay on the front of the SSM does not contain a boot flash card, or the card in the slot is unreadable.
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.)

Note: When the status of a flash card changes, an alert is generated. See "Using Active Monitoring" on page 114.

2.11.2 Replacing a Boot Device

If a boot flash card fails, first try to activate it on the Boot Devices window. If the card fails repeatedly, replace it with a new one.

You can also replace a boot flash card if you have removed the original card to store it as a backup in a remote location.

Note: The replacement boot flash card must be a standard 256 M (or larger) compact flash memory card. As of publication, validated cards are:
- SanDisk 256 Mb Compact Flash
- Kingston 256 Mb Compact Flash (P725228X1)

Check with your supplier for an updated list of approved flash cards.

Warning: A flash card from one Storage System Module cannot be used in a different Storage System Module. If a card fails, replace it with a new flash card.

2.11.2.1 Removing a Boot Flash Card

Before you remove one of the boot flash cards from the SSM, deactivate the device in the Storage Server Console.

1. On the Boot Devices window, select the flash card that you want to remove.
2. Click Deactivate.

   The flash card status changes to Inactive. It is now safe to remove the card from the Storage System Module.
3. Power off the Storage System Module.

4. Remove the flash card from the front of the Storage System Module.

2.11.2.2 Replacing and Activating a New Boot Flash Card

If you replace a boot flash card in the SSM, you must activate the card before it can be used. Activating the card erases any existing data on the card and then synchronizes it with the other card in the Storage System Module.

1. Insert the new flash card in the front of the Storage System Module.
2. Power on the SSM.
3. Log in to the SSM.
4. On the Boot Devices window, select the new flash card.
5. Click Activate.

   The flash card begins synchronizing with the other card. When synchronization is complete, 'Active' displays in the Status column.

Storage 3

3.1 Storage Overview

For each SSM, you can set the RAID configuration, set the RAID rebuild options, and check the RAID status. You can also manage individual drives, including powering them on or off, and reviewing drive information.

3.1.1 Getting There

1. On the Network View, double-click the SSM and log in, if necessary.
   The Edit SSM Configuration window opens.
2. Select Storage from the configuration categories.
   The Storage category opens, shown in Figure 34.

Figure 34. Managing Storage, RAID, and SSM Disks

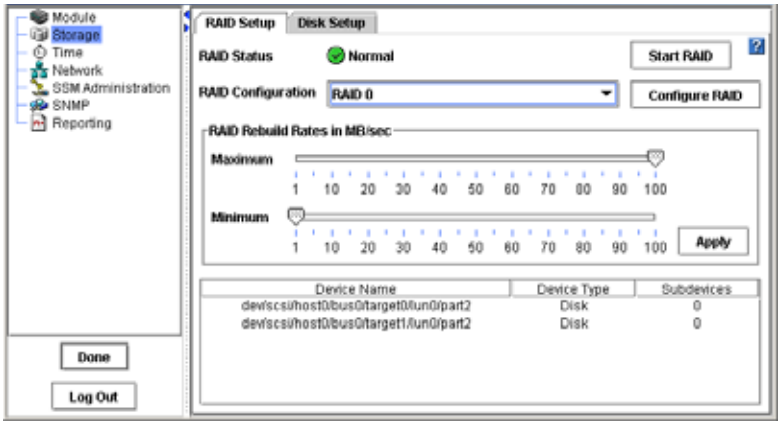

3.2 Configuring and Managing RAID

Managing the RAID settings of a SSM includes:

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

Note: You must configure RAID before you can use a SSM for data storage.

3.2.1 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.

3.3 RAID Configurations Defined

The RAID configuration you choose depends upon how you plan to use the SSM. The SSM can be configured with RAID 0, RAID 1 / 10 or RAID 5 / 50.

3.3.1 RAID 0

RAID 0 creates a striped disk array. 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 array fails, all data on the array may be lost.

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

3.3.2 RAID 1 and RAID 10

RAID 1

RAID 1 provides data redundancy by mirroring the data from one disk onto a second disk. Because data is mirrored between two disks, SSM capacity in RAID 1 is equal to the size of the smallest disk.

RAID 10

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

SSM capacity in RAID 10 is the total capacity of all mirrored disk pairs in the module. The capacity of each mirrored pair is equal to the size of the smallest disk in the pair.

For example, suppose the SSM is configured for RAID 10 and contains 6 disks: 3 100G disks and 3 150G disks. Mirrored pairs are formed by contiguous disks, from left to right. Within each pair, the capacity used is equal to the capacity of the smallest disk. In this example, 3 mirrored pairs are created:

- Pair 1 capacity = 100 G
- Pair 2 capacity = 100 G (because the smallest disk in the pair is 100 G)
- Pair 3 capacity = 150 G
- Total SSM capacity = 350 G

Figure 35. Capacity of Disk Pairs in RAID 10

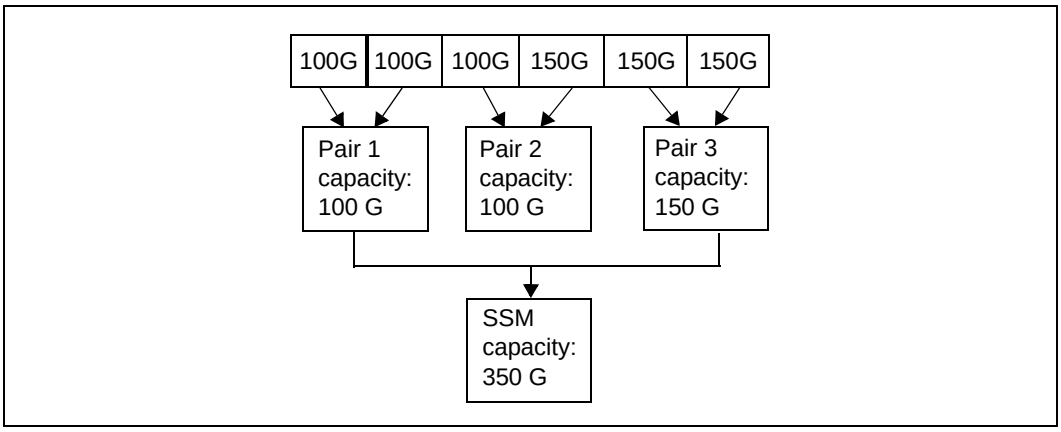

Selecting RAID 1 or RAID 10

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

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

Best Practice for Disk Capacity in RAID 1 / 10

Because disk pairs configured for RAID 1 and RAID 10 use capacity up to the size of the smallest disk in the pair, using different sized disks can result in unused disk capacity. For example, as shown in Figure 35, the second disk pair includes a 100 G and a 150 G disk. Capacity of the pair is limited to 100 G, so 50 G of the larger disk will not be used.

To utilize all of your disk capacity in a RAID 1 or 10 configuration, use the same sized disks throughout the SSM. If this is not possible, use same sized disks within each disk pair. Disk pairs are grouped from left to right in contiguous drive bays.

Note: If you are configuring the SSM for RAID 1 or RAID 10, use disks with the same capacity in all drive bays.

3.3.3 RAID 5 and RAID 50

RAID 5

RAID 5 provides data redundancy by distributing data blocks across disks in an array. Redundant information is stored as parity distributed across the disks. Each disk contains the information that is used to regenerate data if one of the disks in the array fails.

Each RAID 5 array in the SSM contains 4 disks. To maintain redundancy, RAID 5 uses parity equal to the size of one disk in the array. This means that SSM capacity in each RAID 5 array is 3 times the capacity of the smallest disk in the array. Figure 36 shows the distribution of parity across 4 disks in a RAID 5 array.

Figure 36. Parity Distributed Across a RAID 5 Array

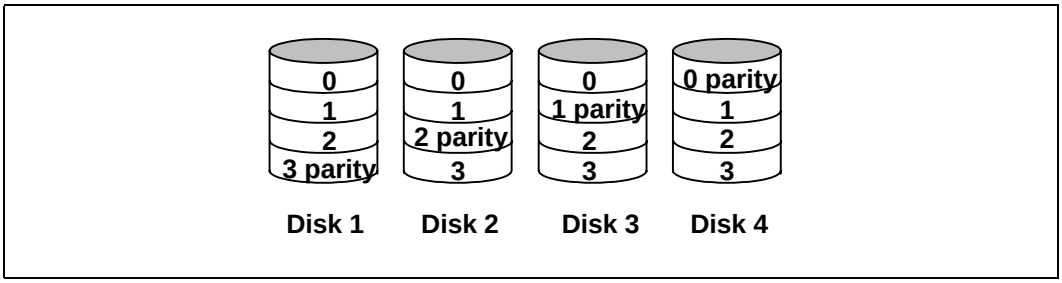

Parity allows the SSM to use more disk capacity for data storage than RAID 10 allows.

RAID 50

RAID 50 combines the redundancy of parity within an array with striping across arrays.

The total capacity of the SSM in RAID 50 is the combined capacity of each RAID 5 array in the module.

For example, suppose the SSM is configured for RAID 50 and contains 8 disks (2 arrays of 4 disks). If the first array contains equal sized disks of 150G, and the second array contains 2 150 G disks and 2 100 G disks, then the total capacity is calculated as follows. Remember that capacity equal to one disk in each array is used for parity.

- Array 1 capacity = 3 disks X 150 G = 450 G
- Array 2 capacity = 3 disks X 100 G (because the smallest disk in the array is 100 G) = 300 G
- Total SSM capacity = 750 G

Figure 37. Capacity of Disk Pairs in RAID 50

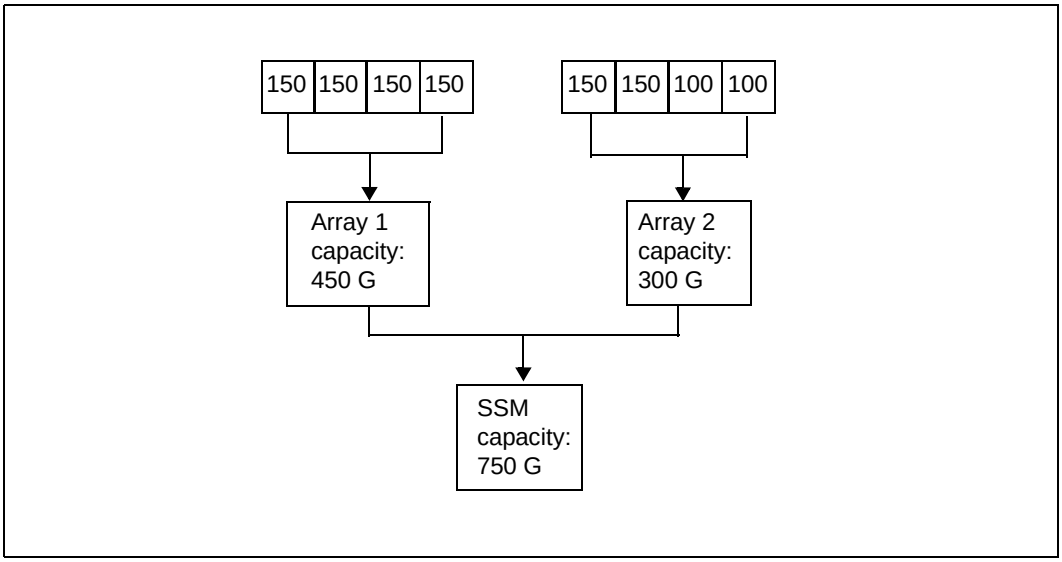

Selecting RAID 5 or RAID 50

RAID 5 and RAID 50 can only be configured on completely populated arrays of disks. This means the SSM must contain either 4, 8, 12, or 16 disks.

Whether the SSM is configured in RAID 5 or RAID 50 depends on the number of disks in the module.

- If the SSM contains 4 disks, then RAID 5 is configured using the first RAID array.
- If the SSM contains 8, 12, or 16 disks, then RAID 50 is configured using 2, 3, or 4 RAID arrays.

Best Practice for Disk Capacity in RAID 5 / 50

Because disk arrays configured for RAID 5 and RAID 50 use capacity up to 3 times the size of the smallest disk in the array, using different sized disks can result in unused disk capacity. For example, as shown in Figure 37, the second disk array includes 2 100 G disks and 2 150 G disks. Capacity of the array is limited to 3 times the size of the smallest disk, or 300 G. This means that 50 G of the two larger disks will not be used.

To utilize all of your disk capacity in a RAID 5 or 50 configuration, use the same sized disks throughout the SSM. If this is not possible, use same sized disks within each disk array.

Note: If you are configuring the SSM for RAID 5 or RAID 50, use disks with the same capacity in all drive bays.

3.3.4 Viewing the RAID Setup Report

In the Storage category, the RAID Setup tab lists the RAID disks in the SSM and provides information about them.The RAID Setup Report is shown in Figure 38. Table 4 describes the information listed in the report.

Figure 38. Viewing the RAID Setup Report

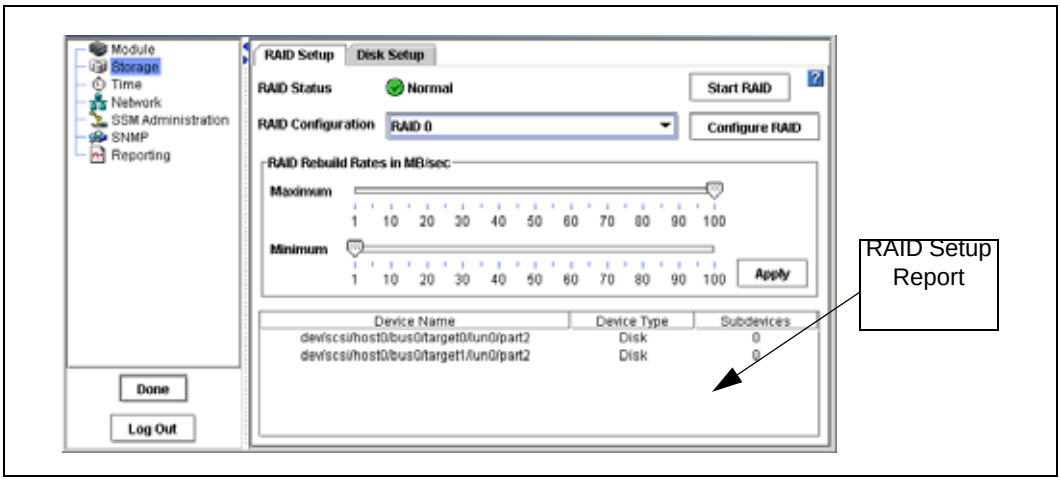

Table 4. The RAID Device Information

This Item	Describes This
Device	The disks, pairs of disks, or arrays used in RAID. • For RAID 0, an entry for each disk in the SSM. • For RAID 1 and RAID 10, one entry for each disk pair. • For RAID 5 and RAID 50, one entry for each array.
Device Type	The RAID level of the device. • For RAID 0, the device type of each disk is Disk. • For RAID 1 and RAID 10, the device type for each disk pair is RAID 1. • For RAID 5 and RAID 50, the device type for each disk array is RAID 5. 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 RAID 0, there are no subdevices because each disk is listed separately in the Device column. • For RAID 1 and RAID 10, there are 2 subdevices per disk pair. • For RAID 5 and RAID 50, there are 4 subdevices per array.

3.3.4.1 Devices Configured in RAID 0

If RAID 0 is configured, each physical disk operates as a separate RAID 0 disk, as shown below.

Figure 39. RAID 0 on an SSM

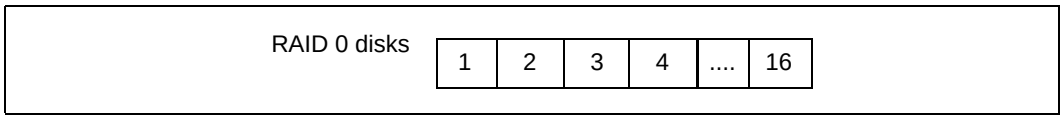

3.3.4.2 Devices Configured in RAID 1 / 10

If RAID 1 or 10 is configured, the physical disks are combined into mirrored pairs of disks. RAID 1 uses only one pair of disks. RAID 10 uses up to 8 pairs of disks.

Figure 40. RAID 10 on an SSM

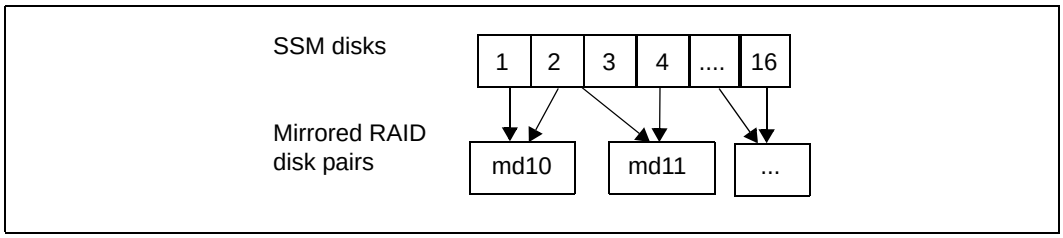

3.3.4.3 Devices Configured in RAID 5/50

If RAID 5 or 50 is configured, the physical disks are grouped into arrays of 4 disks. RAID 5 uses one array of disks. RAID 50 uses up to 4 arrays of disks in each SSM.

Figure 41. Raid 50 on an SSM

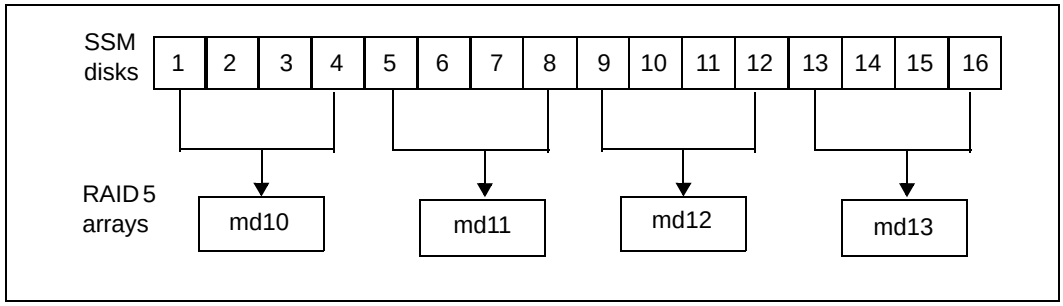

3.4 Planning RAID Configuration

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

Warning: Once RAID is configured, you cannot select a different RAID configuration without deleting all data on the SSM.

3.4.1 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 SSM.
- Replicate volumes across clusters of SSMs.

3.4.1.1 Using RAID for Data Replication

Within each SSM, 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 array. If one disk in the array goes down, data reads and writes continue on the remaining disks in the array.

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 SSM, use RAID to replicate data within the SSM.

3.4.1.2 Using Clustering for Data Replication

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

Clustering is part of the Scalability Pak feature upgrade. See Chapter 11, “Working with Clusters” for more information.

Note: If you plan to create data volumes that span 2 or more SSMs, use clustering to replicate data and to ensure data safety and availability.

3.4.2 Using RAID with Clustering

If you use clustering to replicate volumes across SSMs, then the redundancy provided by RAID 1/10 uses excess capacity and may not be necessary. For example,

- Using clustering, up to 3 copies of a volume can be created on a cluster of 3 SSMs. The clustered configuration ensures that 2 of the 3 SSMs can go down and the volume will still be accessible.
- Configuring RAID 10 on these clustered SSMs means that each of these 3 copies of the volume is stored on 2 disks within the SSM, for a total of 6 copies of each volume. For a 50 G volume, 300 G of disk capacity is used.
- In this case, data safety and availability are ensured more efficiently by configuring RAID 0 on the SSMs and then achieving 3-way replication through clustering. For a 50 G volume, 150 G of disk capacity is used.

RAID 5 / 50 uses less disk capacity than RAID 1 / 10, so it can be combined with clustering and still use capacity efficiently. One benefit of configuring RAID 5 / 50 in clustered SSMs 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 SSM in the cluster. Rebuilding the disks within a single array is faster and creates less of a performance hit to applications accessing data than copying data from another SSM in the cluster.

Note: If you are achieving data replication through clustering:

- — Configuring the SSM for RAID 0 will allow you to utilize 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 or 50 will provide redundancy within each module while allowing most of the disk capacity to be used for data storage.

The table below 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 5. 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 SSMs, RAID 1/10	Yes. 1 less than half the drives can fail, but there is no redundancy in pairs with a failed disk.	No	No	No hot spare disk within the SSM
Clustered SSMs, RAID 0	Yes. However, if any disk in the SSM fails, the entire SSM must be copied from another SSM in the cluster.	Yes	Yes	Yes (configure a hot spare SSM within a cluster)
Clustered SSMs, RAID 5/50	Yes. 1 drive per RAID array can fail without copying from another SSM in the cluster.	Yes	Yes	Yes (configure a hot spare SSM within a cluster)

3.4.3 Planning RAID for Capacity Growth

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

Warning: Once RAID is configured, you cannot select a different RAID configuration without deleting all data on the SSM.

3.5 Requirements for Configuring RAID

Placement of Disks in the SSM

All disks must be in contiguous disk bays, from left to right, for RAID to be configured. If there are empty drive bays, only the disks to the left of the missing 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 arrays of disks. This means the SSM must contain 4, 8, 12, or 16 disks.

Management Groups and RAID

You cannot configure RAID on an SSM that is already in a management group. If you want to change the RAID configuration for an SSM that is in a management group, first delete any volumes and clusters on the SSM and remove it from the management group.

All SSMs in a management group must have the same RAID configuration. See Chapter 9, "Working with Management Groups."

3.6 Configuring RAID

Before you configure RAID, make sure that the disks in the SSM are inserted in contiguous disk bays, from left to right.

- 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 SSM must contain 4, 8, 12, or 16 disks.

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

1. On the Storage configuration category, click the RAID Setup tab to bring it to the front, shown in Figure 38.
2. Select the RAID configuration from the list.
3. Click Configure 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 drives provide fully operational data redundancy with the mirror in place. The SSM is ready for data transfer at this point. See "Monitoring RAID Status" on page 48.

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

Choose the minimum and maximum rates at which the RAID configuration rebuilds if a drive is replaced. The rates are calculated in terms of the amount of data moving onto the new disk per second during the rebuild.

- Setting the rate high is good for rebuilding RAID quickly and protecting data.
- Setting the rate low is good for keeping the SSM available to users during the rebuild.

Example

You have set the RAID rebuild rate to a maximum of 25 and a minimum of 10. A disk fails and you replace it with a new disk. You insert and power on a new disk. When this disk is powered on, RAID rebuild begins on the new disk. Depending upon the network traffic, the maximum and minimum rates for the rebuild range from 10 to 25 MB per second.

Note: When a new or clean disk is installed in an SSM, the SSM automatically starts rebuilding RAID. If a used SSM disk is inserted into a drive bay, you must manually add the disk to RAID on the Disk Setup tab. See "Replacing a Disk" on page 52 and "Adding Disks to the SSM" on page 52 for more information.

Note: RAID rebuild is not available when the SSM is configured for RAID 0. This is because if a disk fails in RAID 0, all of the data on the SSM may be lost. Rebuilding the data on the disk is not possible.

3.6.1.1 Setting RAID Rebuild Rate

1. Select Storage in the configuration categories.
2. Click the RAID Setup tab.
3. Set the first slider for the Maximum Rebuild Rate.
4. Set the second slider for the Minimum Rebuild Rate.
5. Click Apply.
   The settings are then ready when and if RAID rebuild takes place.

3.6.2 Starting RAID

If RAID has been configured on the SSM, 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 starting RAID.

Example

In an SSM, 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 48.) To prevent losing quorum, one of the original disks is replaced. Then RAID is started. Finally, the replacement disk is added to RAID.

3.6.2.1 To Start RAID

1. Select Storage in the configuration categories.
2. Click the RAID Setup tab.
3. Click Start RAID.
   A confirmation message opens.
4. Click OK.

RAID starts.

3.6.3 RAID Quorum

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

3.6.3.1 Quorum for RAID 1 or RAID 10

For RAID 1/10, quorum requires that at least one disk pair in the SSM and one drive in each remaining pair be intact. This means that a SSM configured for RAID 10 can tolerate a loss of up to 1 less than half of the drives. 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 drives in one of the mirrored pairs are not lost. In order for RAID to rebuild when disks are replaced, at least one complete pair of disks must be in the SSM to ensure that data is rebuilt correctly.

Disks are paired from left to right, starting the first disk in the SSM.

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

3.6.3.2 Quorum for RAID 5 or RAID 50

For RAID 5/50, quorum requires that at least 3 of the 4 disks in each RAID array be intact. This means that a SSM that contains 16 drives and is configured for RAID 50 can tolerate a loss of up to 4 disks, provided that each of the failed drives is part of a different array. If 2 or more disks fail within an array, quorum is broken and RAID cannot be rebuilt.

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

- Disks 1, 2, 3, and 4
- Disks 5, 6, 7, and 8
- Disks 9, 10, 11 and 12
- Disks 13, 14, 15, and 16

3.6.4 Monitoring RAID Status

RAID is critical to the operation of the SSM. If RAID has not been configured, the SSM cannot be used. Monitor the RAID status of an SSM to ensure that operation is normal.

3.6.4.1 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 SSM is if the RAID status shows Off.

3.6.4.2 Data Redundancy and RAID Status

In a RAID 10 configuration, when RAID is degraded, there is not full data redundancy. Therefore, data is at risk if there is a drive 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 within an array 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 Network view.

Figure 42. Monitoring RAID Status on the Main Console Window

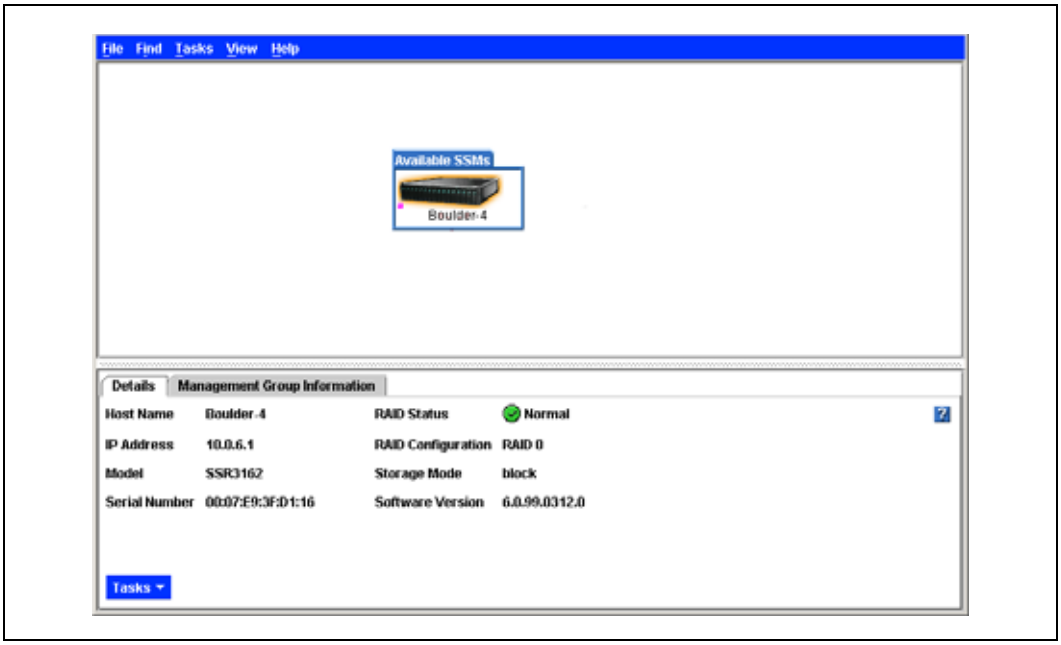

The status displays one of four RAID states.

- **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 used SSM replacement disk has been inserted in a drive. You must add a drive to RAID on Disk Setup if you are inserting a previously used SSM replacement disk.
- **Off** - Data cannot be stored on the SSM. The SSM is down and flashes red in the Network view.

3.6.5 Replacing Disks and RAID

Disk failure in an SSM affects RAID for that module. First, replace the failed disk. Then reestablish RAID on the SSM.

- When using RAID 0, you must reconfigure RAID 0. If the SSM is in a cluster, you must first remove the SSM from the cluster and 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 drives in an SSM and rebuild RAID while the SSM remains in the cluster.

You can view the status of the disks in the SSM on the Disk Setup tab, shown in Figure 43 on page 51. The RAID states are reported on the RAID Setup tab, as shown in Figure 38 on page 41.

Removing and Reinserting the Same Disk

If you pull a disk from its drive bay, and then push it back in that same drive bay,

- RAID rebuild (for RAID 1/10 or RAID 5/50) will work automatically, once you power the disk on.
- RAID 0 will come up in the state it was in when the drive was removed.

See "Setting RAID Rebuild Rate" on page 47.

3.7 Managing Disks

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

3.7.1 Getting There

1. On the Network View, double-click the SSM and log in, if necessary.
   The SSM Configuration window opens to the Module Information tab, shown in Figure 34.
2. Select Storage from the configuration categories.
3. Click the Disk Setup tab to bring it to the front, as shown in Figure 43.
   Any drive bays that do not contain disks are labeled "Off or Missing" in the Status column.

Figure 43. Viewing the Disk Setup Tab in a SSM

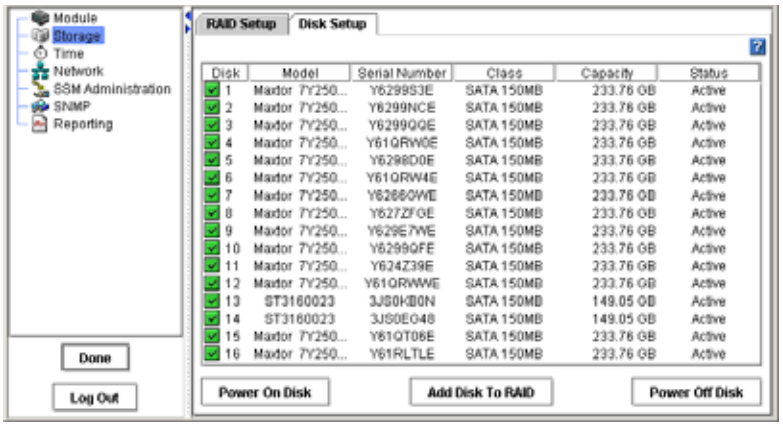

Note: The drives are labeled 1 through 16 in the Disk Setup window and correspond with the disk drives from left to right (#1 on the left and #16 on the right) when you are looking at the SSM.

3.7.2 Using the Disk Report

The Disk Setup tab lists the individual disks in the module and provides information about them.

Table 6. Description of Items on the Disk Report

This Item	Describes This
Numbers	Lists the numbers of the disk drives. The disks are numbered sequentially from left to right as you view the front of the SSM.
Model	The model of the disk in the SSM.
Serial Number	The serial number of the disk.
Class	The class (type) of disk. The SSM uses SATA drives.
Capacity	The data storage capacity of the disk.
Status	Whether the disk is • Active and participating in RAID (Status column Active, other columns with information). • On but not participating in RAID (Status column Inactive, other columns with information). • Not on (Status column Off or Missing, other columns with dashed lines -------).

3.7.3 Verifying Disk Status

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

3.8 Replacing a Disk

If you replace a failed disk with a new, or clean, disk, the SSM will automatically start rebuilding RAID after you power the disk on. See “Powering Drives On” on page 53.

If you replace the failed disk with a used disk, and you start either the rebuild (for RAID 1/10 or RAID 5/50) or the configuration (for RAID 0), a message prompts you to confirm that either operation will remove any existing data from the disk. This safety feature ensures that data on a disk is not accidentally lost due to a RAID rebuild or configuration.

Note: Always power the drive off before removing and replacing a disk. See “Powering Drives Off” on page 53.

Replacing Disks in RAID 0

If you lose a disk in RAID 0, you may lose all of your data. (If RAID 0 is configured, but data redundancy is achieved by replication of data within a cluster of SSMs, then data is not lost. See “Repairing a SSM” on page 166.) In order to make the SSM functional again, you must replace the disk and reconfigure RAID 0.

1. Remove the SSM from the cluster. See “Removing a SSM from a Cluster” on page 164.
2. Replace the disk in the SSM. See “Replacing a Disk” on page 52.
3. Power on the drive. See “Powering Drives On or Off” on page 53.
4. Reconfigure RAID 0.

Replacing Disks in RAID 1/10 or RAID 5/50

To replace a disk in an SSM running RAID 1/10 or RAID 5/50:

1. On the Disk Setup tab, select the old disk and click Power Off Disk.
2. Replace the disk in the SSM. See “Replacing a Disk” on page 52.
3. On the Disk Setup tab, select the new disk and click Power On Disk.

3.9 Adding Disks to the SSM

The SSM can hold up to 16 disks. If the SSM is configured for RAID 1 or RAID 10, you must add an even number of disks. If the SSM is configured for RAID 5 or RAID 50, you must add disks in complete arrays (4 disks at a time).

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.

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

1. Add the new disks to the SSM.
2. Using the Storage System Console, log in to the SSM.
3. Select Storage from the configuration categories.
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 Inactive.

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

   If you are adding pairs of disks to a RAID 1/10 SSM or arrays of 4 disks to a RAID 5/50 SSM, Shift-click to select multiple disks to add to RAID.

   RAID begins to rebuild on the new disk. RAID will rebuild according to the RAID Rebuild rate set on the RAID Setup tab.

   As soon as the RAID Status shows as Normal on the RAID Setup tab, the drives 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.

Note: 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 delete any volumes and clusters on the SSM and remove it from the management group. Then remove disks from the SSM and reconfigure RAID.

3.10 Powering Drives On or Off

Powering drives on and off is part of removing and replacing disks in the SSM. A bad drive should be powered off from the Storage System Console before you remove it from the module. Then, after the replacement disk is inserted in the drive bay, it must be powered on.

Warning: Any time you must remove a drive, you should power it off from the Console before you physically remove it from the module, unless the module itself is powered off.

3.10.1 Powering Drives Off

1. Select Storage from the configuration categories.
2. Click the Disk Setup tab.
3. Select the drive in the list to power off.

   If the drive is on, all the columns are filled in.

4. Click Power Off Disk.

   A confirmation message opens.

5. Click OK.

3.10.2 Powering Drives On

1. When a new disk is inserted into an SSM that is on, the disk must be powered on.
2. Select the drive in the list to power on.

   If the drive is not on, it is listed as Off or Missing in the Status column and the other columns display dotted lines, like this ---------.

3. Click Power On Disk.

   A confirmation message opens.

4. Click OK.

Managing the Network 4

The SSM has two integrated TCP/IP network interfaces. In addition, the SSM can include three add-on cards, each with 2 or 4 interfaces.

For each SSM you can

- Configure the SSM's 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 an SSM's management group
- Bond NICs to ensure continuous network access or to improve bandwidth

4.1 Getting There

1. On the Network View, double-click the SSM and log in.
   The SSM Configuration window opens.
2. Select Network from the SSM configuration categories.
   The window opens with the TCP/IP tab on top, shown in Figure 44.

Figure 44. Viewing the Network Configuration

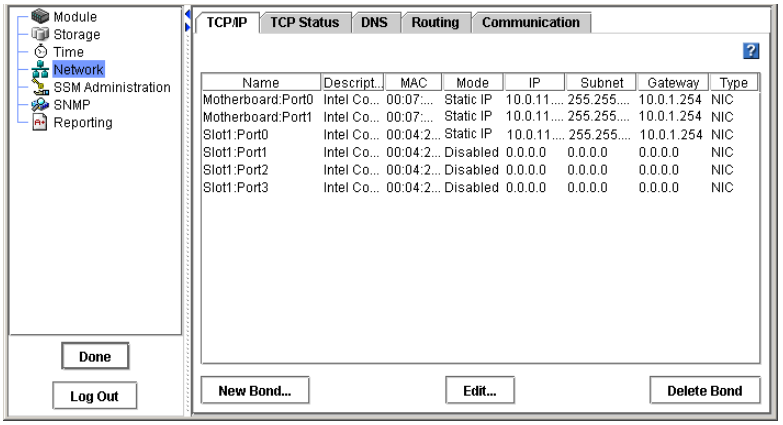

4.1.1 The TCP/IP Tab

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

Table 7. 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
Add-on NICs in PCI Slots	
Slot1:Port0 Slot1:Port1 and so on	Multiple add-in PCI cards, each containing up to 4 ethernet or Fibre Channel 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.

4.2 Identifying the Network Interfaces

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

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

Table 8. Identifying the NICs in the Motherboard

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

Table 9. Identifying Add-on NICs

Add-on Interfaces	
Where labeled	**What the label says**
Network Configuration Category in the Storage System Console • TCP/IP tab • TCP Status tab	Name - Slot1:Port1, Slot1:Port2, and so on Description - Intel Gigabit Ethernet
Configuration Interface Name	Slot1:Port0 Slot1:Port1 and so on
Label on the back of the SSM	Port A Port B Port C Port D

The motherboard interfaces are labeled NICs 1 and 2 on the back of the SSM, shown in Figure 45. The PCI slots for add-on interfaces are located to the right of the motherboard ports.

Figure 45. Network Interface Ports and Open PCI Slots on the Back of the SSM

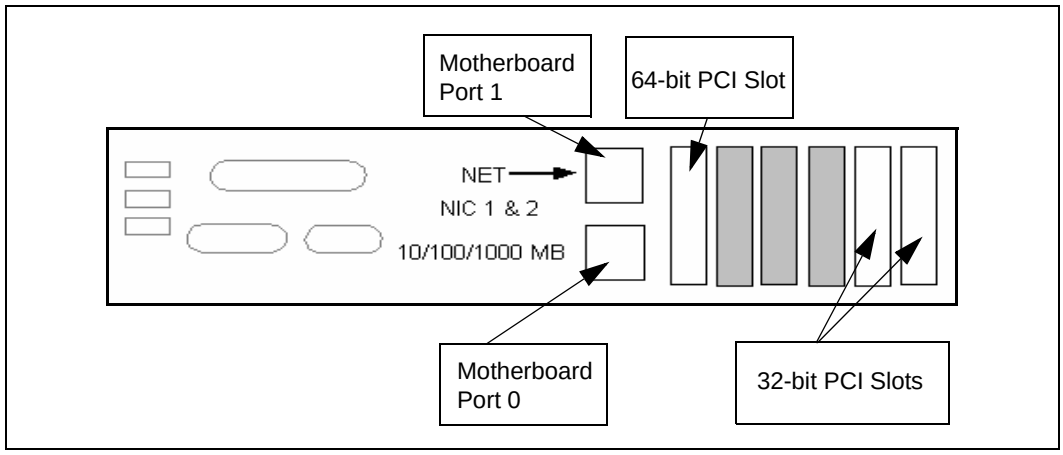

4.2.1 Adding Interfaces to PCI Slots

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 interface cards can contain ethernet or Fibre Channel ports.

The SSM contains one open 64-bit / 66 MHz PCI slot and two open 32-bit / 33 MHz PCI slots. 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 46. Distributing Bandwidth and Ensuring Fault Tolerance of Add-on Ports Across PCI Slots

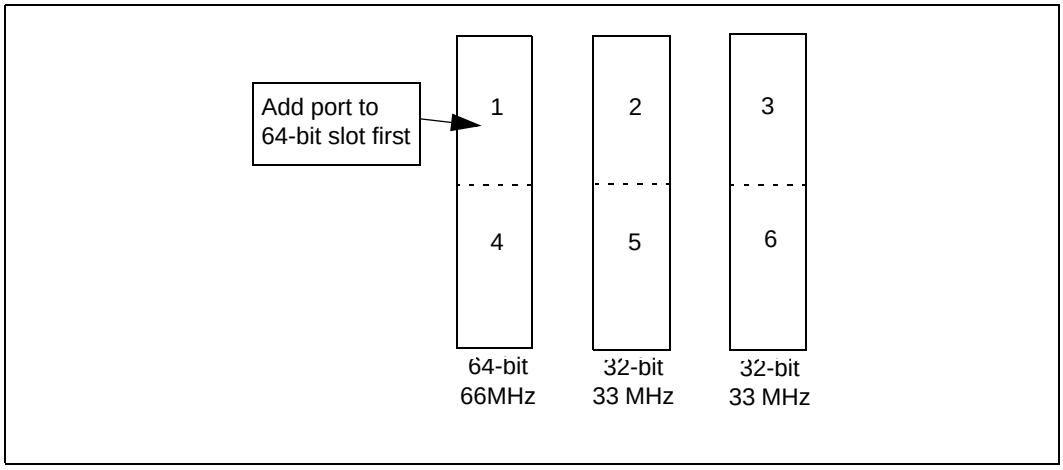

Note: When adding more than one port to the SSM, you can distribute bandwidth and to ensure fault tolerance by distributing the ports across more than one PCI slot. Start with the first (64-bit) slot.

4.2.1.1 Adding Fibre Channel Ports

When you add a card containing Fibre Channel ports to a PCI slot, the Fibre Channel ports do not display on the TCP/IP tab of the Network configuration window.

You can view the status of the Fibre Channel ports and the unique World Wide Name (WWN) of each port in the Passive Report.

1. Select Reporting from the configuration categories.
   The Reporting window opens.
2. Click Refresh to display statistics on the Passive tab.
3. Scroll down to the Fibre Channel statistics.

- The Node WWN is the same for all ports in a management group.
- The Port WWN is unique for each port.

Figure 47. Viewing the WWN of a Fibre Channel Port

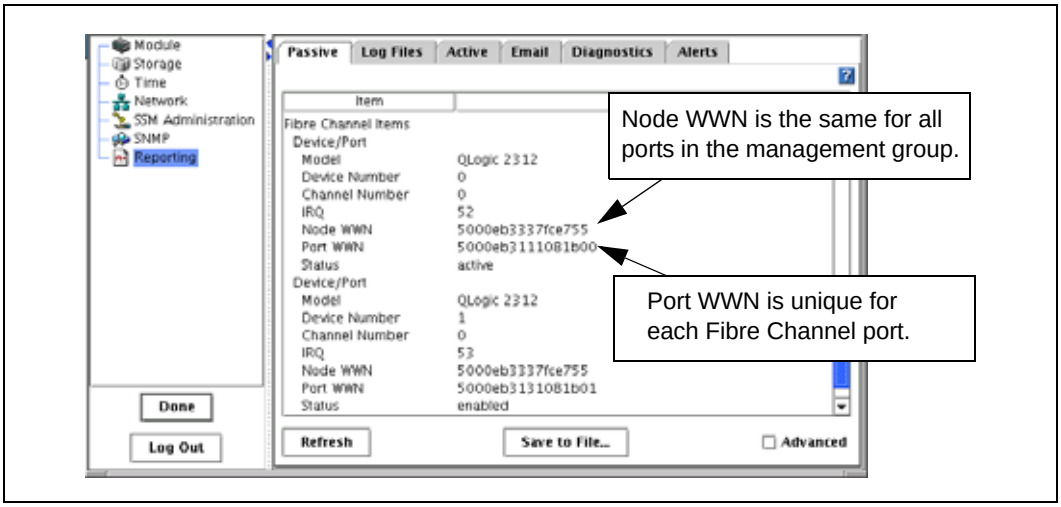

4.3 Configuring the IP Address Manually

Use the Network category in the SSM configuration window to configure the IP address for an interface.

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

1. Select Network from the SSM configuration categories.
2. The window opens with the TCP/IP tab on top.
3. On the TCP/IP tab, select the interface from the list for which you want to configure or change the IP address.
4. Click Edit.

   The Edit TCP/IP Configuration window opens, shown in Figure 48.

Figure 48. Configuring the IP Address Manually

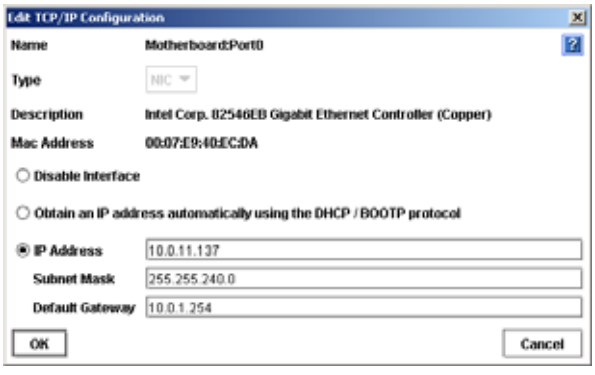

5. Select IP Address and complete the fields for IP Address, Subnet mask, and Default gateway.
6. Click OK.
   A confirmation message opens.
7. Click OK.
   A message notifying you of an automatic log out opens.
8. Click OK.
   The automatic log out occurs.

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

9. Log in to the newly addressed SSM.

If you are changing the IP address of an SSM 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.

4.4 Using DHCP

Warning: If you use DHCP, be sure to reserve statically assigned IP addresses for all SSMs 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 48.
3. Select Obtain an address automatically using the DHCP/BOOTP protocol.
4. Click OK.

4.5 Configuring NIC Bonding

Network interface bonding provides high availability, fault tolerance, and / or bandwidth aggregation for the network interface cards in the SSM. 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.

Depending on your network infrastructure design and Ethernet switch capabilities, you can bond NICs in two 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.

Warning: NIC aggregation requires plugging both NICs into the same switch. This means that NIC aggregation does not protect against switch failure.

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

4.5.1 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 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 SSM into separate switches. While NIC aggregation will only survive a port failure, active backup will survive a switch failure.

Table 10. Comparison of Active Backup and NIC Aggregation Bonding

Feature	Active Backup	NIC Aggregation
Bandwidth	Use of 1 NIC at a time provides normal bandwidth.	Simultaneous use of both NICs increases bandwidth.
Protection during port failure	Yes	Yes
Protection during switch failure	Yes (NICs are plugged into separate switches)	No (Both NICs are plugged into the same switch)
Requires support for 802.3ad link aggregation	No	Yes

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

4.5.2 Physical and Logical Interfaces

The NICs in the SSM are labeled 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 labeled bondN and acts as the master interface. As the master interface, bondN controls and monitors the two physical slave interfaces.

Table 11. Physical and Logical Interfaces in a Bond

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.

4.5.3 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)."

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 12. 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.

4.5.3.1 Requirements for Active Backup

To configure active backup:

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

4.5.3.2 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 68.

4.5.3.3 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 13. SSM 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. • Motherboard:Port0 is Active. • Slot1:Port0 is connected and is Passive (Ready).
2. Active interface fails. Bond0 detects the failure and Slot1:Port0 takes over.	• Motherboard:Port0 status becomes Passive (Failed). • Slot1:Port0 status changes to Active.
3. The Motherboard:Port0 link is restored.	• Motherboard:Port0 status changes to Active after a 30 second delay. • Slot1:Port0 status changes to Passive Ready).

4.5.3.4 Summary of NIC Status During Failover

Figure 49 shows the states of Motherboard:Port0 and Slot1:Port0 when configured for Active Backup.

Figure 49. NIC Status During Failover with Active Backup

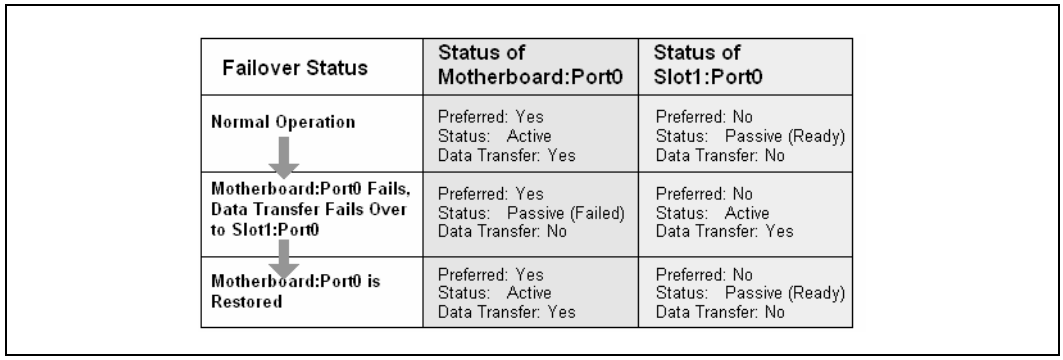

4.5.3.5 Example Network Configurations with Active Backup

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

Figure 50. Active Backup in a Two-switch Topology with Server Failover

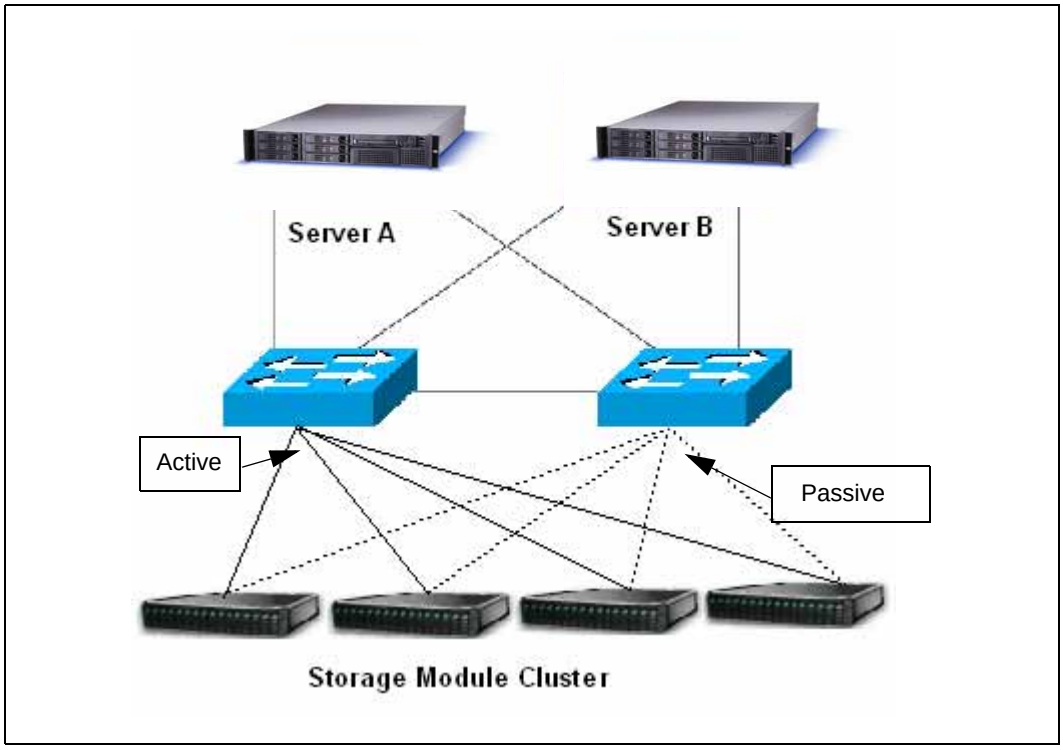

The two-switch scenario in Figure 50 is a basic, yet effective, method for ensuring high availability. If either switch failed, or a cable or NIC on one of the SSMs failed, the active backup bond would cause the secondary connection to become active and take over.

Figure 51. Active Backup Failover in a Four-switch Topology

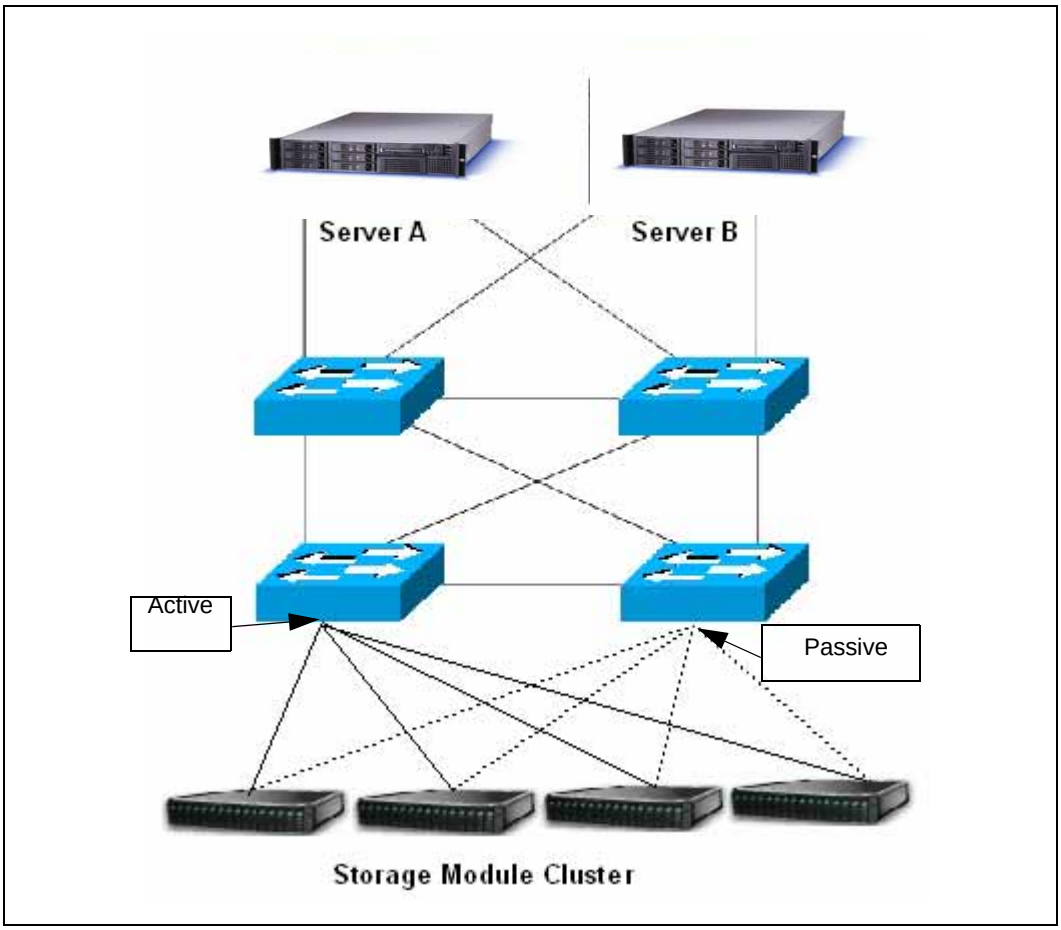

Figure 51 illustrates the active backup configuration in a four-switch topology.

Note: For information about how active backup works in more complex network environments, contact your technical support representative.

4.5.4 How NIC Aggregation Works

NIC aggregation allows the SSM 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.

4.5.4.1 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 supports 802.3ad link aggregation. If SSM is directly connected to a server, then the server must support 802.3ad link aggregation.

4.5.4.2 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.

4.5.4.3 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:Port0 fails, then Slot1:Port0 remains active.

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

Table 14. SSM 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. • Motherboard:Port0 is Active. • Slot1:Port0 is Active.
2. Motherboard:Port0 interface fails. Because NIC aggregation is configured, Slot1:Port0 continues operating.	• Motherboard:Port0 status becomes Passive (Failed). • Slot1:Port0 status remains Active.
3. Motherboard:Port0 link failure is repaired.	• Motherboard:Port0 resumes Active status. • Slot1:Port0 remains Active.

4.5.4.4 Summary of NIC States During Failover

Figure 52 shows the states of Motherboard:Port0 and Slot1:Port0 when configured for NIC aggregation.

Figure 52. NIC Status During Failover with NIC Aggregation

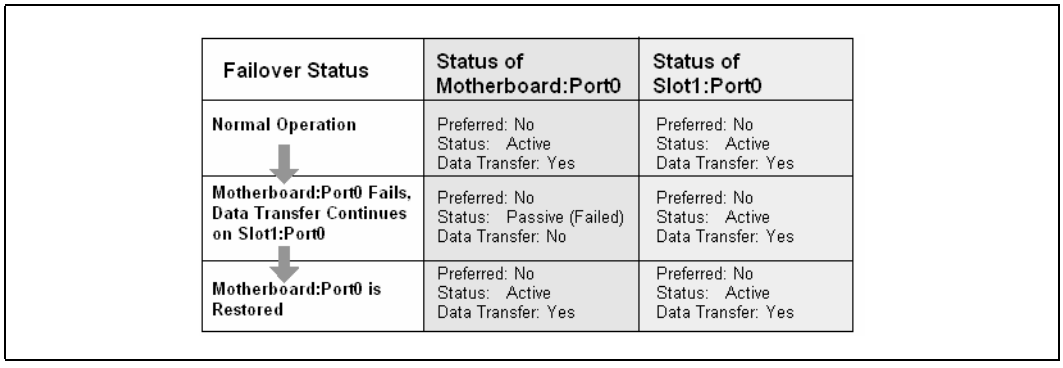

4.5.4.5 Example Network Configurations with NIC Aggregation

Two simple network configurations using NIC aggregation in high availability environments are illustrated.

Figure 53. NIC Aggregation in a Partial-mesh Topology with Server Failover

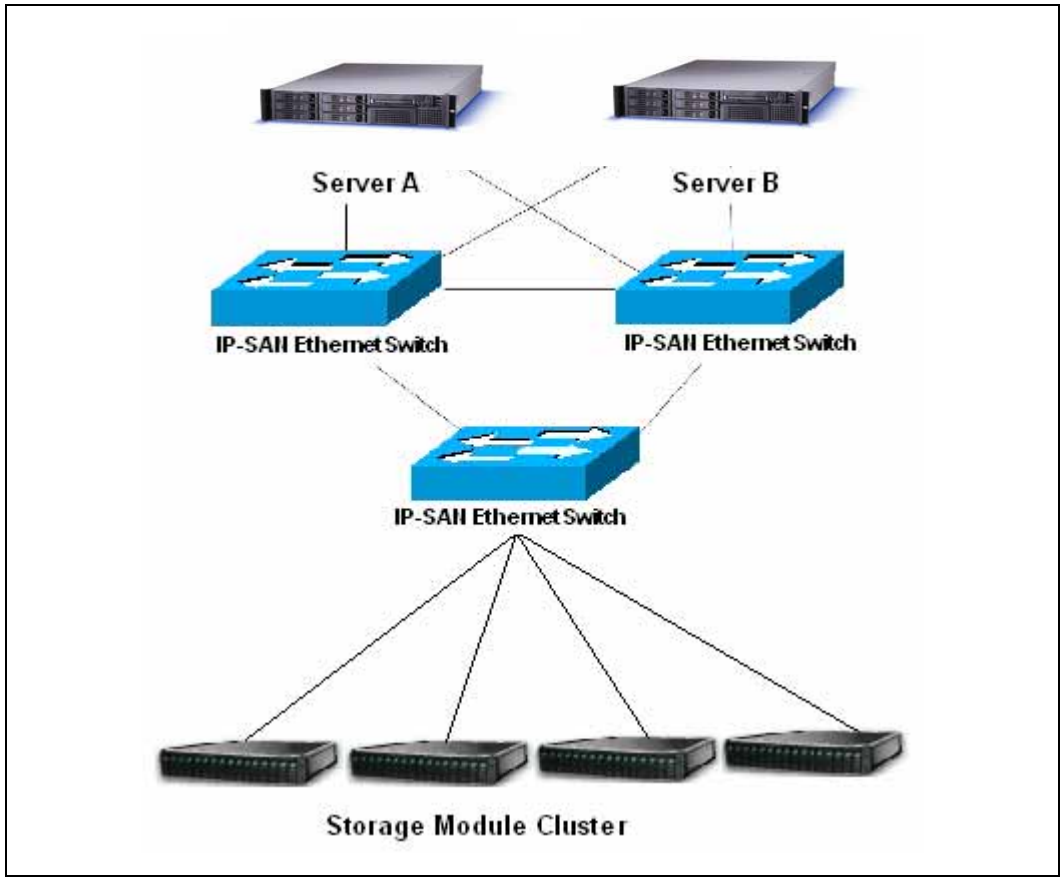

Figure 54. NIC Aggregation in a Single-switch Topology

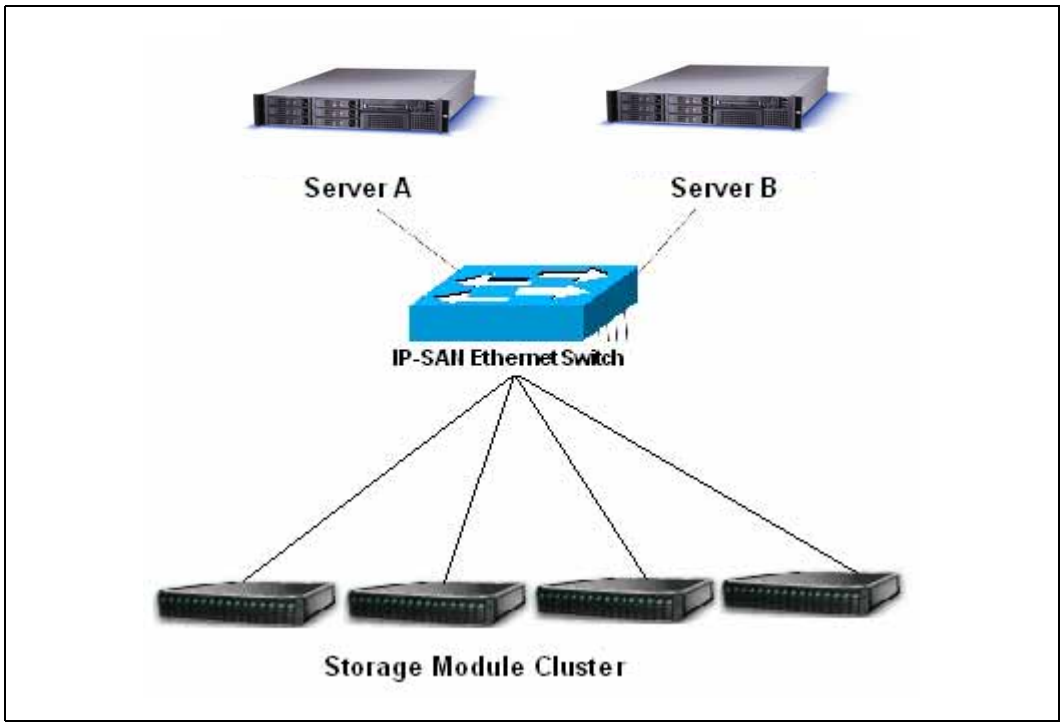

Note: For information about how NIC aggregation works in more complex network environments, contact your technical support representative.

4.5.5 Creating a NIC Bond

Follow these guidelines when creating NIC bonds:

- You can 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, bond interfaces located in the motherboard with interfaces 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.
- Be sure to record the configuration information of each interface before you create the bond. 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, 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 SSM before you add the SSM to a management group.
- Allocate a static IP address for the logical bond interface. You cannot use DHCP for the bond IP.

Warning: To ensure that the bond works correctly, you must configure it as follows:

— Create the bond on the SSM before you add it to a management group.

— Verify that the bond is created.

1. Ensure that the SSM is not in a management group.
2. Log in to the SSM.
3. On the TCP/IP tab, shown in Figure 55, select 2 or 4 NICs to bond.
   The NICs that you select do not have to be consecutive NICs in the list.

Figure 55. Selecting Motherboard: Port0 and Slot1: Port0 for a New Bond

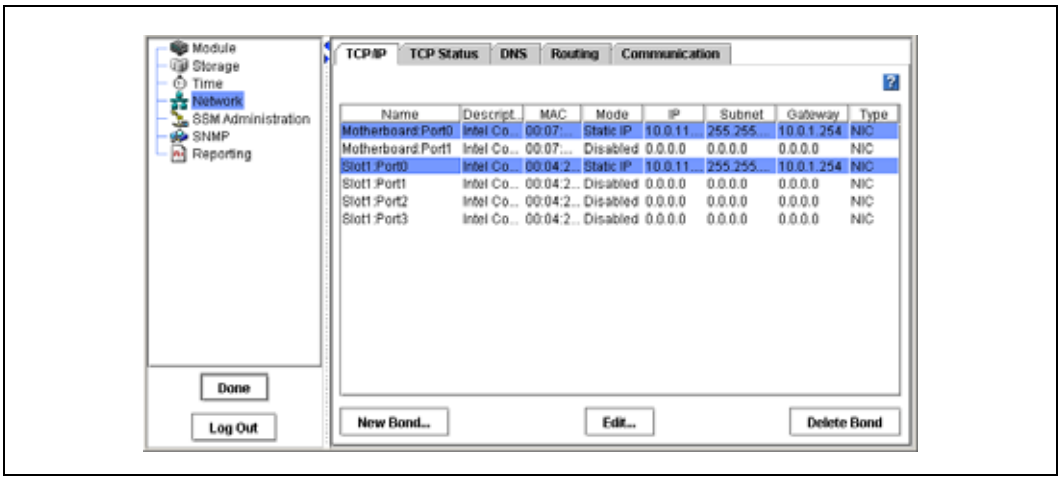

4. Click New Bond.
   The Create Bond Configuration window opens, shown in Figure 56.

Figure 56. Creating a NIC Bond

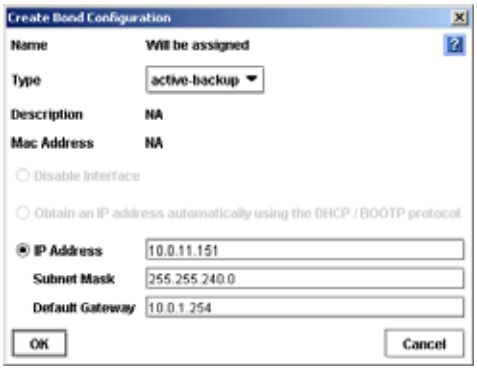

5. **To create an active backup bond,** select active-backup from the Type list.
   or
   To create a NIC aggregation bond, select 802.3ad from the Type list.
6. Enter a static IP address for the bond.

The default value for the bond IP address is the IP address of one of the physical interfaces in the bond.

7. Enter the Subnet mask.

   The default value for the bond subnet mask is the subnet mask of one of the physical interfaces in the bond.

8. [Optional] Enter the default gateway.

   The default value for the bond default gateway is the gateway of the one of the physical interfaces in the bond.

9. Click OK.

   A confirmation message opens.

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

    A message opens prompting you to search for the bonded SSM on the network.

11. Search for the SSM by subnet and mask or by IP address / host name.

    A message opens listing the unicast IP addresses that must be set on the application servers.

12. Click OK.

13. Verify the new bond interface.

    The TCP/IP tab displays the new list of interfaces, as shown in Figure 57.

Figure 57. Viewing a New Active Backup Bond

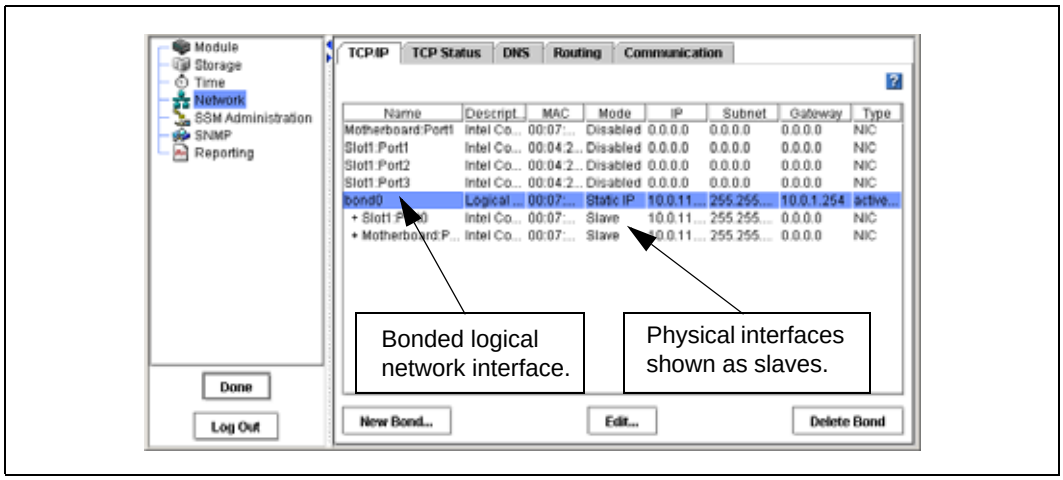

The bond interface shows as “bond0” and has a static IP address. The two physical NICs now show up as slaves in the Mode column.

14. [Optional, for active backup bonds] To change which interface is the preferred interface in an active backup bond, select one of the NICs in the bond and click Set Preferred.

4.5.6 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 the NIC aggregation bond, neither physical interface is preferred.

Figure 58 shows the status of interfaces in an active backup bond. Figure 59 shows the status of interfaces in a NIC aggregation bond.

Figure 58. Viewing the Status of an Active Backup Bond

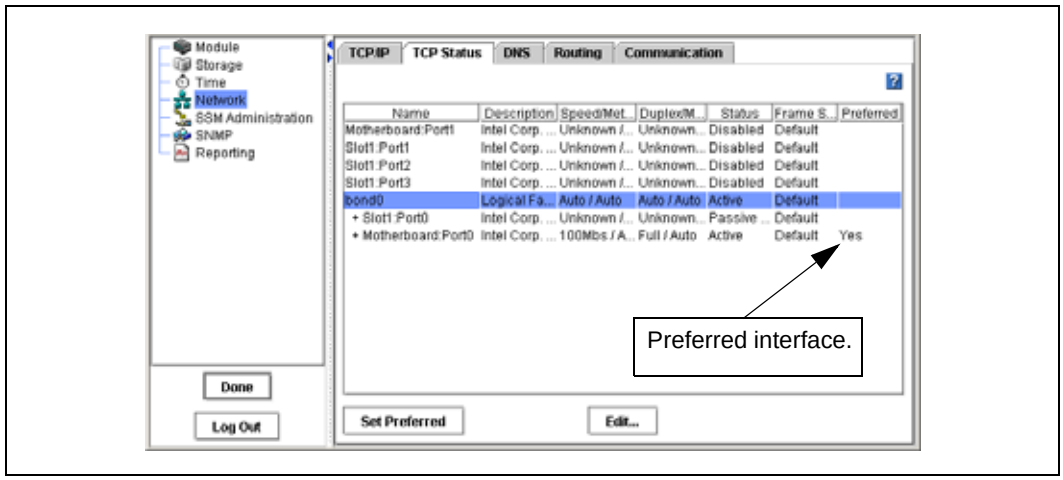

Figure 59. Viewing the Status of a NIC Aggregation Bond

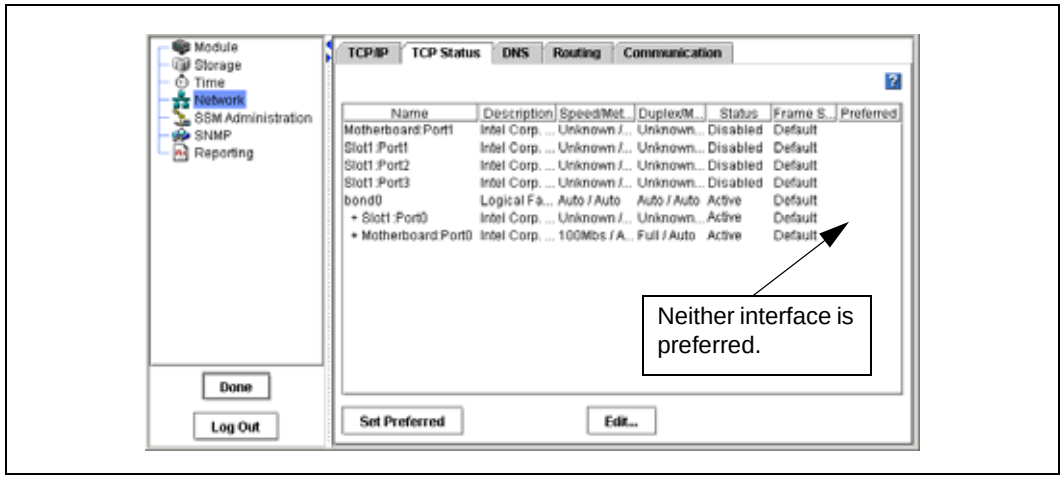

Note: If the bonded NIC experiences rapid, sequential ethernet failures, the Console may display the SSM as failed (flashing red) and access to data on that SSM fails. However, as soon as the ethernet connection is reestablished, the SSM and the Console display the correct information.

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

2. Click Delete Bond.

   Because the IP addresses change, the Search for SSMs window opens. For detailed information, see "Finding Storage Server Modules on the Network" on page 11 in Chapter 1, "Getting Started."

3. Finish searching for the SSM, using the desired method.

   Finding the SSM might take a few minutes. You can exit the search window and use the Find menu at your convenience.

Note: You can also use the Configuration Interface to delete a NIC bond. See 'Using the Configuration Interface," on page 237.

4.6 Disabling a Network Interface

You can disable the network interfaces on the SSM.

- 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 SSM, 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 SSM.
If you disable all the interfaces, you must reconfigure at least one interface using the Configuration Interface to access the SSM. See 'Using the Configuration Interface," on page 237.

4.6.1 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 48.

3. Click Disable Interface.
4. Click OK.

   A confirmation message opens.

5. Click OK.

4.6.1.1 If SSM is in a Management Group

If the SSM 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.

4.6.2 Configuring a Disabled Interface

If one interface is still connected to the SSM but another interface is disconnected, you can reconnect to the second interface using the Storage System Console. See “Configuring the IP Address Manually” on page 59.

If both interfaces to the SSM are disconnected, you must attach a terminal, or PC or laptop to the SSM with a null modem cable and configure at least one interface using the Configuration Interface. See ‘Using the Configuration Interface,” on page 237.

4.7 TCP Status

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

4.7.1 The TCP Status Tab

Review the status of the network interfaces on the TCP Status tab, shown in Figure 60.

Figure 60. Viewing the TCP Status

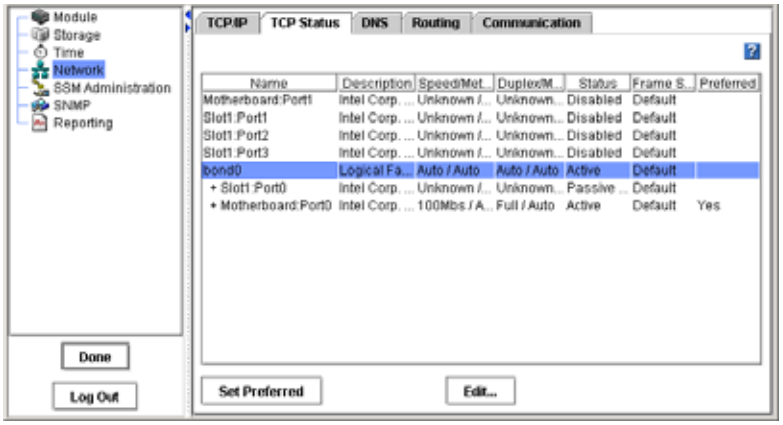

Table 15. Status of and Information about Network Interfaces

Column	Description
Name	Name of the interface. Entries are • Motherboard:Port0 • Motherboard:Port1 • Slot1:Port0 • bond0 - the bonded interface(s) [displays only if SSM 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.

Table 15. Status of and Information about Network Interfaces (Continued)

Column	Description
Status	Describes the state of the interface. See Table 12 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.

4.7.2 Editing the TCP Speed and Duplex

You can change the speed and duplex of the 1000BASE-T TCP interfaces. 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.

Note: 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.

4.7.2.1 Best Practice

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

Table 16. Setting SSM Speed and Duplex Settings

SSM 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 61.

Figure 61. Editing TCP Speed, Duplex, and Frame Size

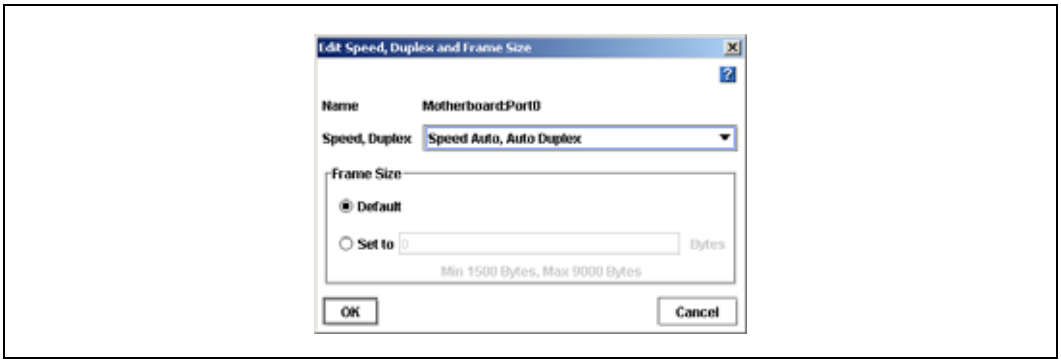

3. Select the combination of speed and duplex that you want.
4. Click OK.
   A series of status messages display. Then the changed setting displays in the TCP status report.

Note: You can also use the Configuration Interface to edit the TCP speed and duplex. See 'Using the Configuration Interface," on page 237.

4.7.3 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.

4.7.3.1 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 SSMs on the network, and set compatible frame sizes on all clients.

The frame size on the SSM should correspond to the frame size on Windows and Linux application servers.Table 17 shows recommended SSM frame sizes and the corresponding frame sizes for Windows and Linux clients.

Table 17. Setting Corresponding Frame Sizes on SSMs and Windows or Linux Clients

SSM Frame Size	Windows Client Frame Size	Linux Client Frame Size
1500 (Default)	1542 (Default)	1500 (Default)
4046	4088	4042
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 SSM 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.

Change the speed and duplex settings while the SSM 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.

4.7.3.2 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 62.

Figure 62. Editing TCP Speed, Duplex, and Frame Size

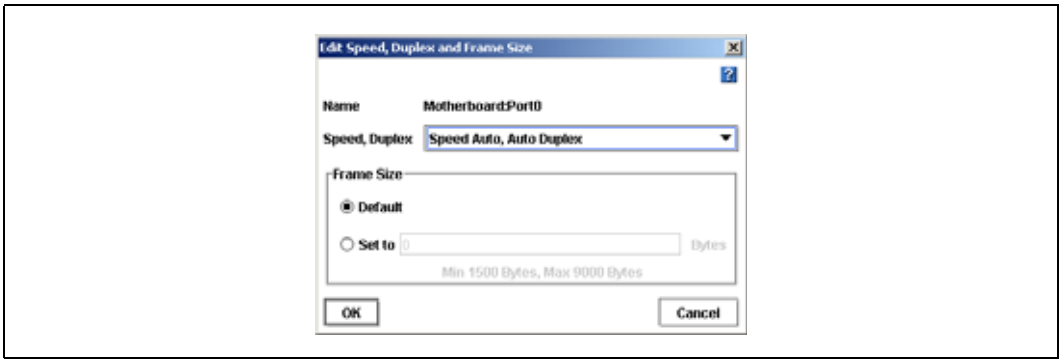

3. Click the Set To radio button.
4. Enter a value between 1500 and 9000 bytes in the Set To field.
5. Click OK.
   A series of status messages display. Then the changed setting displays in the TCP status report.

Note: You can also use the Configuration Interface to edit the frame size. See 'Using the Configuration Interface," on page 237.

4.8 Using a DNS Server

The SSM can use a DNS server to resolve host names. For example, if you enter a host name to specify an NTP time server, the SSM 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.

4.8.1 DNS and DHCP

If you configure the SSM 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 SSM. These DNS servers are listed as IP addresses in the SSM configuration window in the Network category on the DNS tab. You can remove these DNS servers, but the SSM will not be able to resolve host names until you enter a new DNS server.

4.8.2 DNS and Static IP Addresses

If you assigned a static IP address to the SSM and you want the SSM to recognize host names, you must manually add a DNS server to the Network DNS tab.

Note: If you initially set up the SSM 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 SSM and log in, if necessary.
2. The SSM Configuration window opens. Select Network from the SSM configuration categories.
3. Click the DNS tab to bring it to the front, shown in Figure 63.

Figure 63. Adding DNS Servers

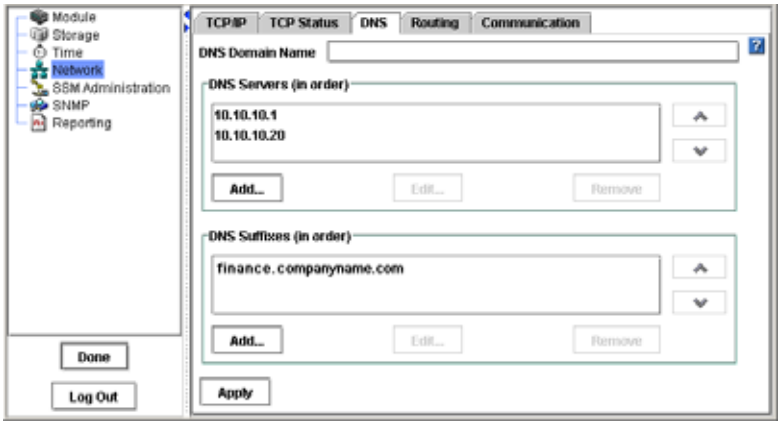

4.8.3 Adding the DNS Domain Name

Add the name of the DNS domain in which the SSM resides.

1. On the DNS tab, type the DNS domain name.
2. Click Apply when you are finished.

4.8.4 Adding a DNS Server

Add up to three DNS servers for use with the SSM.

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.

4.8.5 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 SSM 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.

4.8.6 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.

4.8.7 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.

4.8.8 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.

4.8.9 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.

4.9 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 SSM. To create a route that will survive an SSM reboot or shut down, you must enter a default gateway on the TCP/IP tab. See “Configuring the IP Address Manually” on page 59.

Information for each route listed includes the device, the network, gateway, and mask, and flags.

4.9.1 Adding Routing Information

1. On the Network View, double-click the SSM and log in, if necessary.
2. The SSM Configuration window opens. Select Network from the SSM configuration categories.

3. Click the Routing tab to bring it to the front, shown in Figure 64.

Figure 64. Adding Network Routing Information

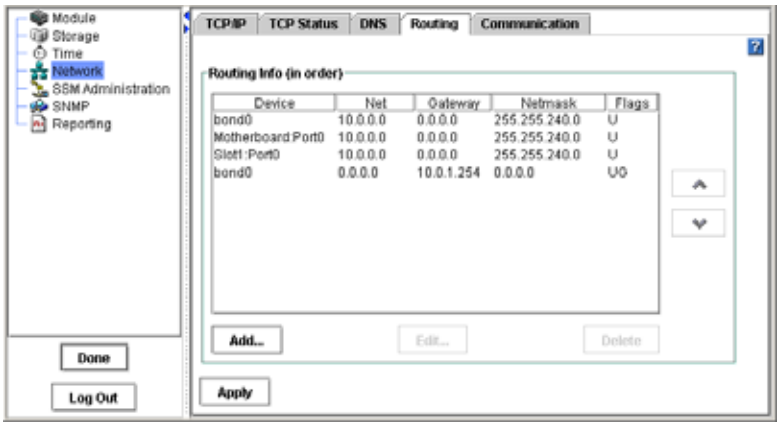

4. Click Add.
   The Add Routing Information dialog opens, shown in Figure 65.

Figure 65. Adding Routing Information

5. Select the port to use for routing in the Device list.
6. Type the IP address portion of the network address in the Net field.
7. Type the IP address of the router in the Gateway field.
8. Select the netmask.
9. Click OK.
10. Use the arrows on the routing table panel to order devices according to the needs of your network.
    The SSM will attempt to use the routes in the order in which they are listed.
11. Click Apply when you are finished.

4.9.2 Editing Routing Information

1. On the routing tab, select the information you want to change.
2. Click Edit.
   The Edit Routing Information dialog opens, shown in Figure 66.

Figure 66. Editing Routing Information

3. Change the relevant information.
4. Click OK.
5. Click Apply.

4.9.3 Deleting Routing Information

1. On the routing tab, select the information you want to delete.
2. Click Delete.
   A confirmation message opens.
3. Click OK.
4. Click Apply when you are finished.

4.10 Configuring a Direct Connection Between the SSM and an EBSD Host

If you want to configure a direct (point-to-point) connection between the SSM and the EBSD host computer, you must specify the route to be used for communication between the SSM and the EBSD host.

1. On the TCP/IP tab, edit the Ethernet port to be used for communication as shown below.

Table 18. SSM Network Interface Settings

SSM Network Interface Setting	Value
IP Address	The IP address of the SSM
Subnet Mask	The same subnet as the EBSD host
Default Gateway	The IP address of the SSM

2. On the Routing tab, add a route for communication with the EBSD host.

Table 19. SSM Route Settings

SSM Route Setting	Value
Device	The network interface you configured in step 1.
Net	The IP address of the EBSD host.
Gateway	The IP address of the SSM
Netmask	255.255.255.255

3. On the EBSD host computer, use the command line to configure the host computer's ethernet interface as follows:

Table 20. EBSD Host Network Interface Settings

EBSD Host Network Interface Setting	Value
IP Address	The IP address of the EBSD host
Subnet Mask	The same subnet as the SSM
Default Gateway	The IP address of the SSM

4. On the EBSD host computer, use the command line to add a route to communicate with the SSM.

Table 21. EBSD Host Route Settings

EBSD Host Route Setting	Value
Device	The network interface you configured in step 3
Net	The IP address of the SSM
Gateway	The IP address of the SSM
Netmask	0.0.0.0

Note: If the network interfaces on the SSM and the EBSD host are both 10/100 NICs, then you must use crossover connection cables.

4.11 Configuring SSM Communication

Use the Communication tab to configure the network interface used by the SSM to communicate with other SSMs on the network and to update the list of managers that the SSM can communicate with.

Figure 67. Selecting the Storage System Software Interface and Updating the List of Managers

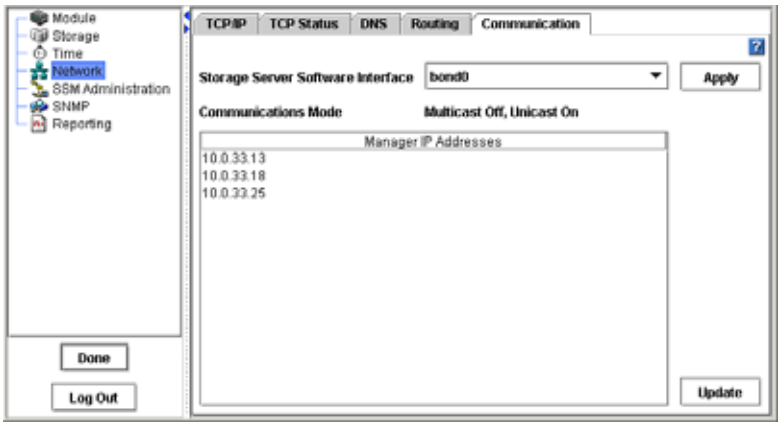

4.11.1 Selecting the Interface Used by the Storage System Software

The Storage System Software uses one network interface for communication with other SSMs on the network. In order for clustering to work correctly, the Storage System Software communication interface must be designated on each SSM. The interface can be

- a single NIC that it 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 an SSM 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 SSM from the management group.

To select a different communication interface:

1. Make sure that the SSM is not in a management group.
2. Select Network from the SSM configuration categories.
3. Click the Communication Mode tab to bring it to the front, shown in Figure 67.
4. Select an interface from the Storage System Software Interface drop-down list.
5. Click Apply.

4.11.2 Updating the List of Manager IP Addresses

Update the list of manager IP addresses to ensure that a manager running on this SSM is communicating correctly with all managers in the management group.

Note: Each time you update the list of managers, you must reconfigure application servers that use the management group to which this SSM 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 SSM.

1. Select Network from the SSM configuration categories.
2. Click the Communication Mode tab to bring it to the front, shown in Figure 67.
3. Click Update.

   The list is updated with the current SSM in the management group and a list of IPs with every manager's enabled network interfaces.

Note: If you are not logged into all the managers in the management group, you will be asked to log in before Update continues.

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 "Communication Mode" on page 126 in Chapter 9, "Working with Management Groups."

Setting the Date and Time 5

5.1 Date and Time Overview

The 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 NTP**
  Configure the SSM to use an external time service (NTP).
- **Setting Date and Time**
  Set the date and time on the SSM if not using an external time service.

5.1.1 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 won't start at the intended time. See "Resetting the Management Group Time" on page 137 in Chapter 9, "Working with Management Groups."

5.1.2 Getting There

1. On the Network View, double-click the SSM and log in, if necessary.
2. The SSM Configuration window opens. Select Time from the SSM configuration categories.
   The Time window opens, shown in Figure 68

Figure 68. Setting the Time Zone and the Date and Time

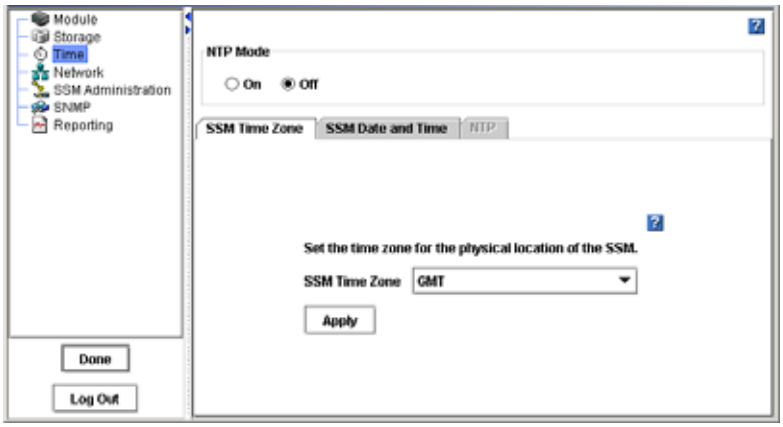

5.2 Setting the SSM Time Zone

You must set the time zone whether you use NTP or not. Set the time zone for the physical location of the SSM. HTTP files display the time stamp according to this local time zone.

1. Click the SSM Time Zone tab to bring it to the front.
2. Choose the time zone for the location of the SSM.
3. Click Apply.

5.3 Setting SSM Date and Time

Set the date and time on the SSM.

If using NTP, the NTP server controls the date and time for the SSM. See “Using NTP” on page 87.

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.

5.3.1 Setting the Date and Time

1. If you are not using an NTP server, make sure NTP mode is set to Off.
   The SSM Date and Time tab is enabled.
2. Click the SSM Date and Time tab to bring it to the front, shown in Figure 69.

Figure 69. Setting the SSM Date and Time

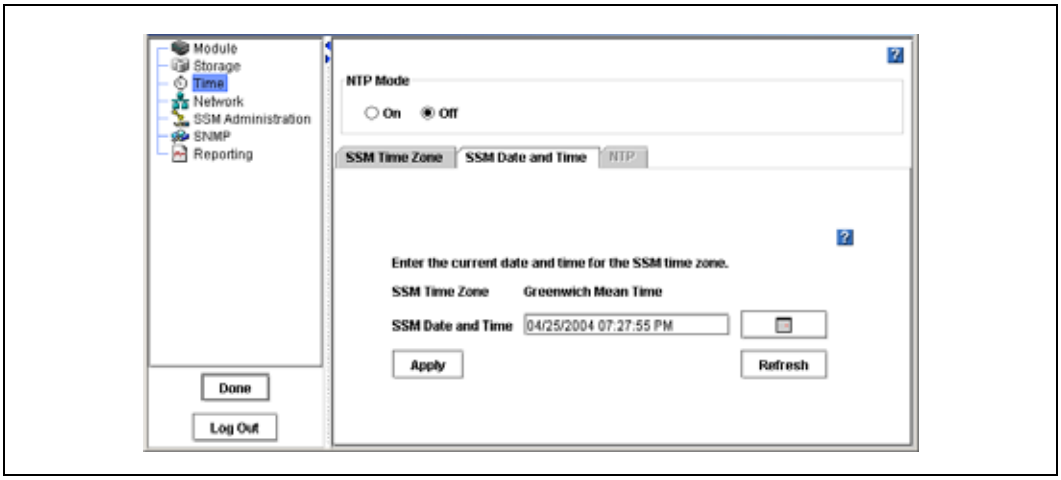

3. Select a time zone.
4. Change the date and time to the correct date and time for that time zone. Type the time directly in the field. Type the date in the field as well, or use the calendar button to select the date.
5. Click Apply.

5.4 Using NTP

You can use Network Time Protocol (NTP) to manage the time for the SSM. NTP updates are fixed at 5 minute intervals. You still must set the time zone for the SSM. See “Setting the SSM Time Zone” on page 86.

1. Select On in the NTP Mode area.
   The Add NTP Server dialog opens, shown in Figure 70.

Figure 70. Adding an NTP Server

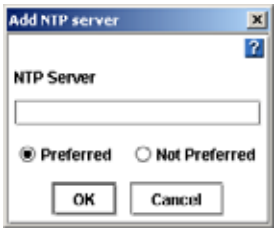

2. Type the IP address of the NTP server you want to use.
3. Click whether you want the NTP server to be designated preferred or not preferred.

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 71.

Figure 71. Viewing the List of NTP Servers

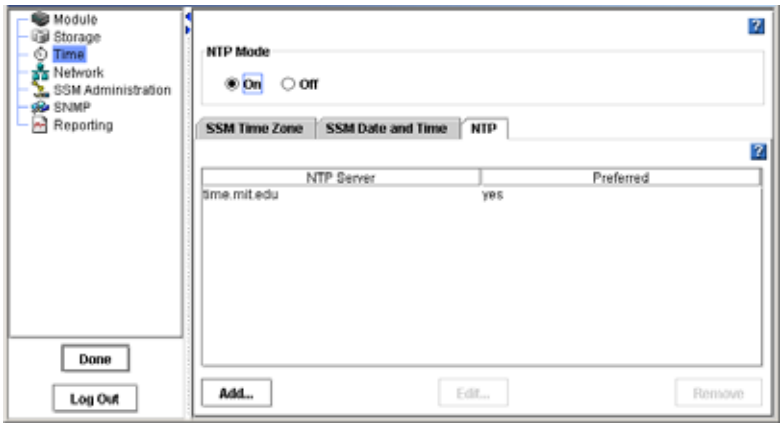

5.4.1 Editing NTP Servers

You can change the 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. Make certain that the NTP Mode is On.
2. On the NTP tab, select the NTP server you want to edit.
3. Click Edit.
   The Edit NTP Server window opens, shown in Figure 72.

Figure 72. Editing an NTP Server

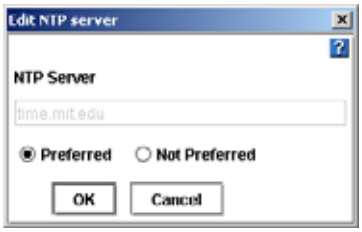

4. Change the preference of the NTP server.
5. Click OK.
   The list of NTP servers displays the changed NTP server in the list.

Note: To change the IP of an NTP server, you must remove the server no longer in use and add a new NTP server.

Administrative Users and Groups 6

6.1 User and Group Administration Overview

The Storage System Software 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.

6.1.1 Getting There

1. On the Network View, double-click the SSM and log in, if necessary.
   The Module Configuration window opens.
2. Select SSM Administration from the configuration categories.
   The Groups tab opens, as shown in Figure 73.

Figure 73. Viewing the SSM Administration Groups Tab

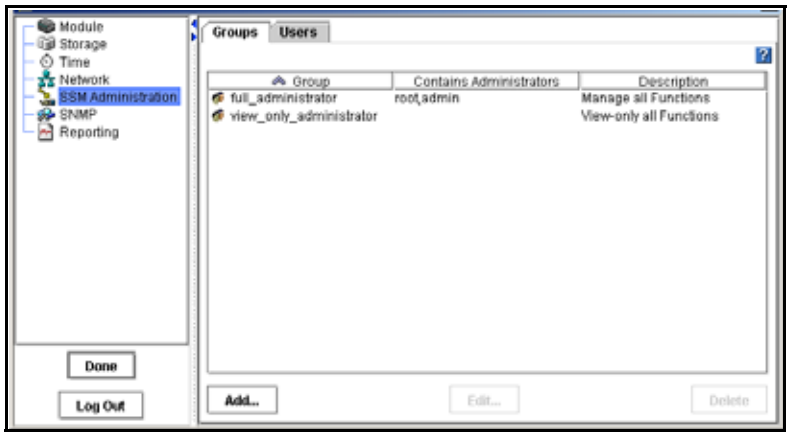

6.2 Managing Administrative Groups

The SSM comes configured with a set of default administrative groups. Use these groups or create new ones.

6.2.1 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 22. 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)

6.2.2 Adding Administrative Groups

Administrative groups are listed on the SSM Administration window on the Groups tab, shown in Figure 73.

6.2.2.1 Adding a Group

1. Select SSM Administration from the configuration categories.
2. Click the Groups tab to bring it to the front.
3. Click Add on the Groups tab.
   The Create Administrative Group window opens, shown in Figure 74.

Figure 74. Adding an Administrative Group

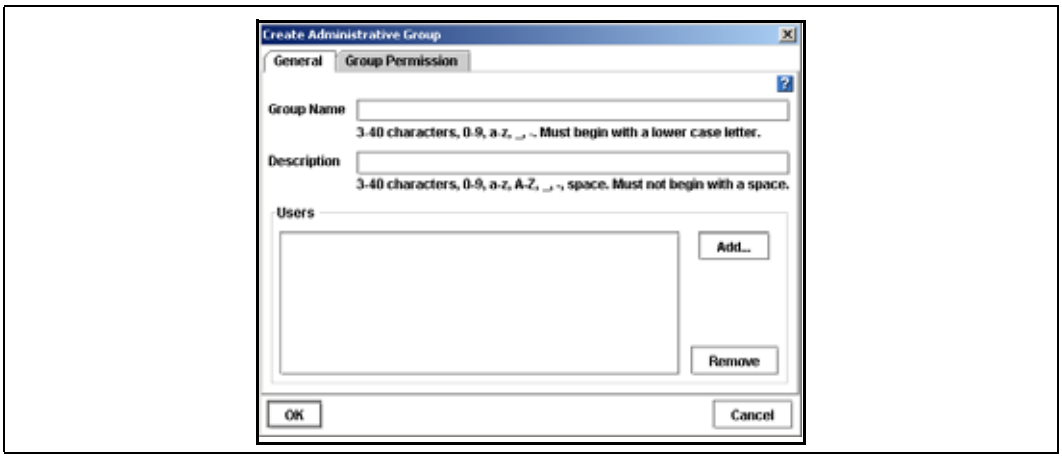

4. Type a Group Name and Description. Both are required.

Table 23. Administrative Group Name Requirements

Group Name Requirements	Example
• 3 to 40 characters • start with a letter • using the letters a-z, A-Z, numbers 0-9, _, -	• jSoftware_Admins • Region11_Managers

6.2.2.2 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 75.

Figure 75. Adding an Administrative User to a Group

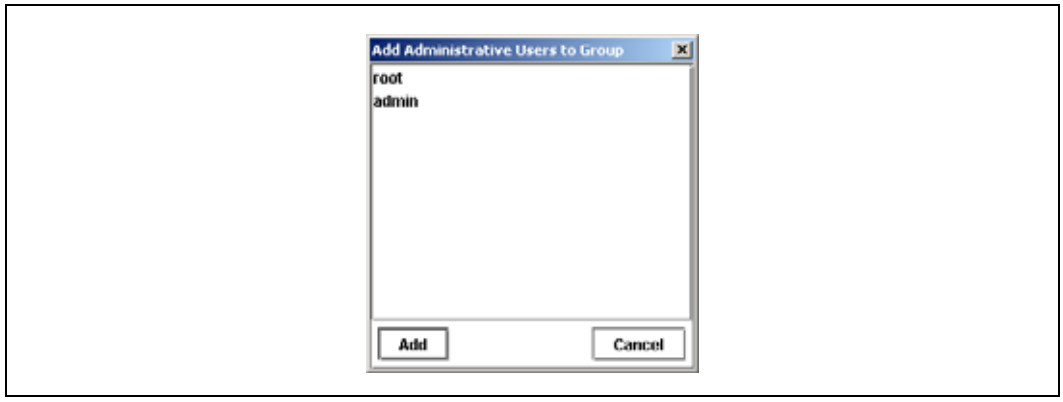

2. Select one or more users you want to add to the group.
3. Click Add.

6.2.3 Adding Administrative Group Permissions

Administrative groups can have

- different levels of access to the SSM, such as read/write, and
- 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 76.

Figure 76. Adding Permissions to Administrative Groups

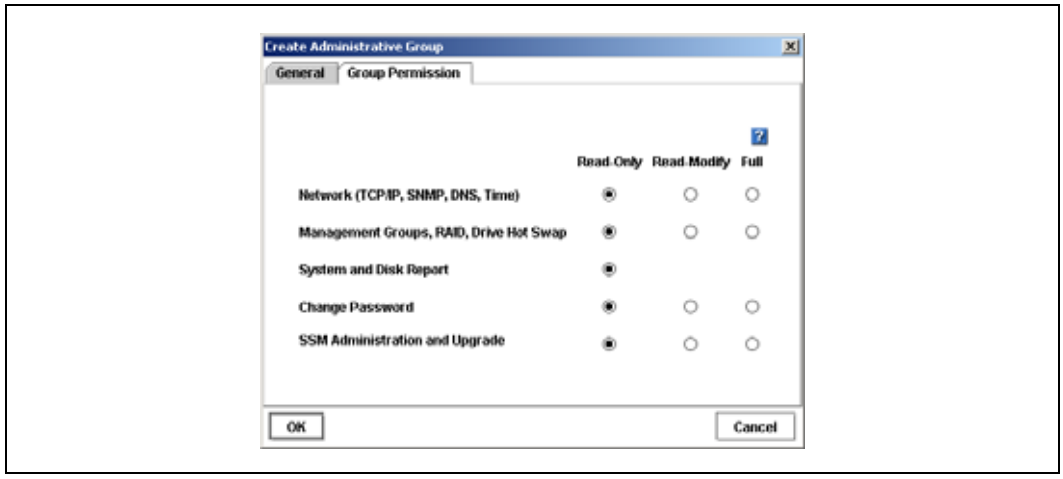

2. Click the permission level for each function for the group you are creating.
3. Click the General tab and complete the rest of the information if you have not already done so.
4. Click OK to finish adding the group.

   The SSM Administration window opens with the Groups tab in front. The new group is added to the list.

6.2.4 Description of Administrative Group Permissions

Table 24. 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 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.

6.2.4.1 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).

Figure 77. Sorting Administrative Groups

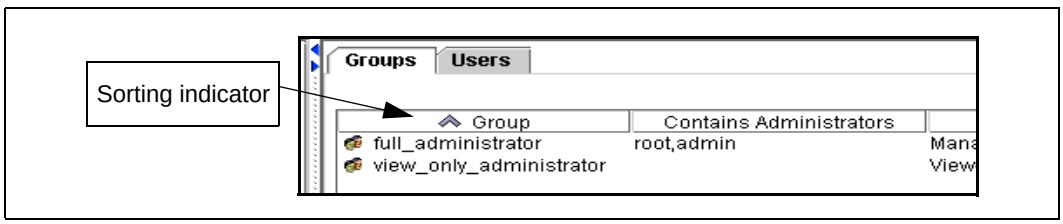

6.2.5 Editing Administrative Groups

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 78.

Figure 78. Editing an Administrative Group

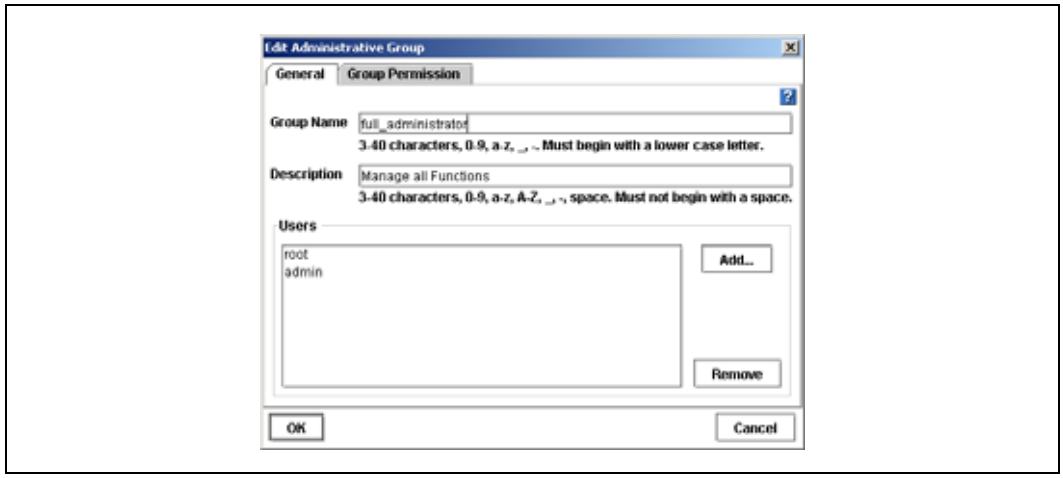

4. Change the name and description as necessary.

6.2.5.1 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.

6.2.5.2 Changing Administrative Group Permissions

Change the management capabilities available to members of a group. The default setting is Read Only for each category.

Changing Permissions for a Group

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.

6.2.6 Deleting Administrative Groups

Delete administrative groups from the SSM.

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.

6.3 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.

6.3.1 Adding Administrative Users

Administrative users are listed on the SSM Administration window on the Users tab along with their group membership and a description.

Figure 79. Adding Administrative Users

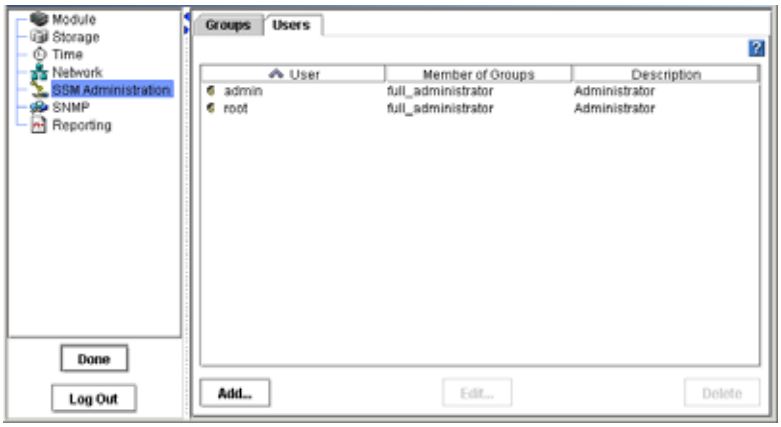

6.3.1.1 Adding an Administrative User

1. Select SSM Administration from the configuration categories.
2. Click the Users tab to bring it to the front, shown in Figure 79.
3. Click Add.
   The Create Administrative User window opens, shown in Figure 80.

Figure 80. Adding an Administrative User

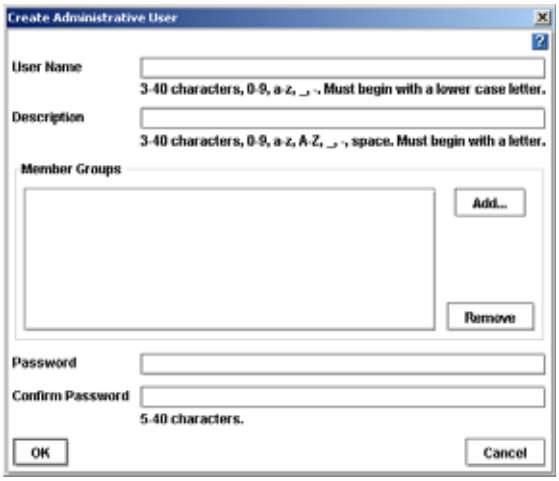

4. Type a User Name and Description.
5. Type a password and confirm that password.

6.3.1.2 Adding a Member Group

1. Click Add in the Member Groups section.

The Add Administration Groups window opens, shown in Figure 81.

Figure 81. Adding a Group to an Administrative User

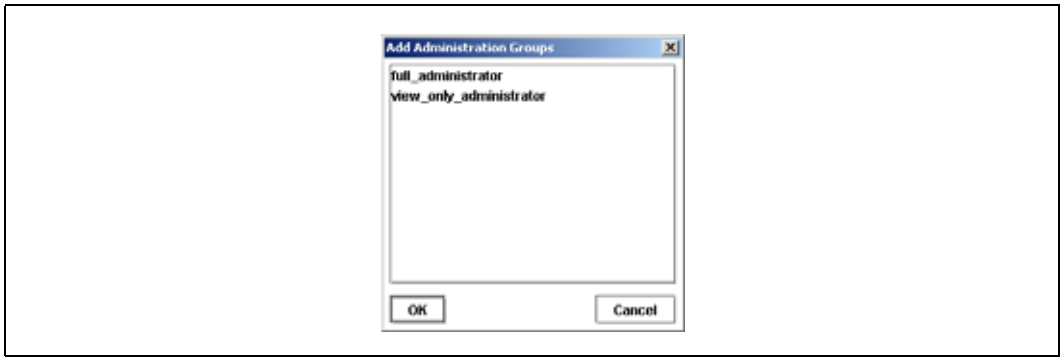

2. Select one or more groups you want to add.
3. Click OK.
   The Create Administrative User window opens.
4. Click OK to finish adding the administrative user.

6.3.1.3 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).

Figure 82. Sorting Administrative Users

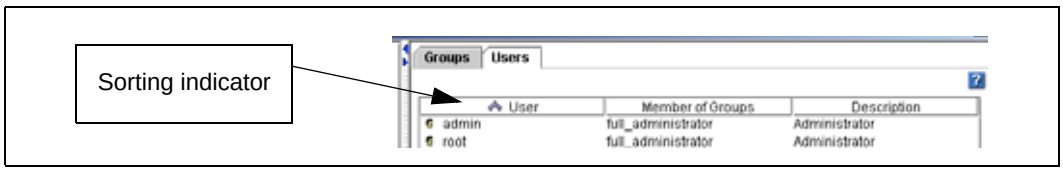

6.3.2 Editing Administrative Users

Change information about 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 83.

Figure 83. Editing an Administrative User

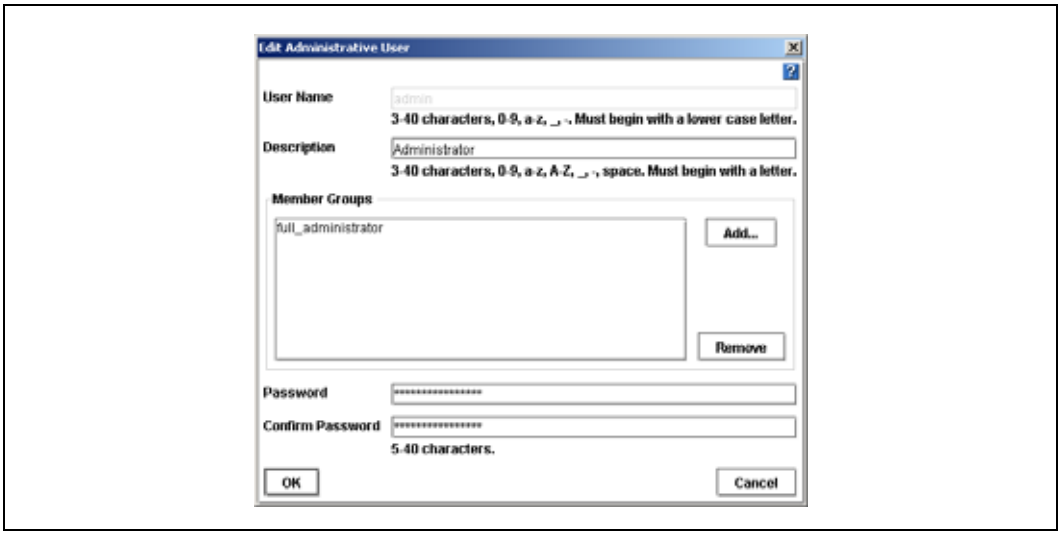

5. Change the necessary information.
6. Click OK.

6.3.3 Deleting Administrative Users

Delete administrative users from the SSM.

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.

Using SNMP 7

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 “Locating the Storage System MIB” on page 105 for a list of Storage System MIBs.

7.1 Getting There

1. On the Network View, double-click the SSM and log in, if necessary.
   The SSM Configuration window opens.
2. Select SNMP from the configuration categories.
   The SNMP General tab opens, shown in Figure 84.

Figure 84. Using SNMP

7.2 Configuring the SNMP User

The SNMP user is a proxy user that allows SNMP access to the SSM. The SNMP user name and password are required to use the SNMP Agent with the SSM.

Use this option to:

- Select an existing administrative user as SNMP user
- Add a new SNMP user
- Change the password for an existing SNMP user

Note: You can only have one SNMP proxy user.

1. Select SNMP from the configuration categories.

The SNMP General window opens, shown in Figure 84.

2. Click the SNMP User tab.
   The SNMP User window opens, shown in Figure 85.
   If an SNMP User has been designated, that name will show in the SNMP User field.

Figure 85. Managing SNMP User

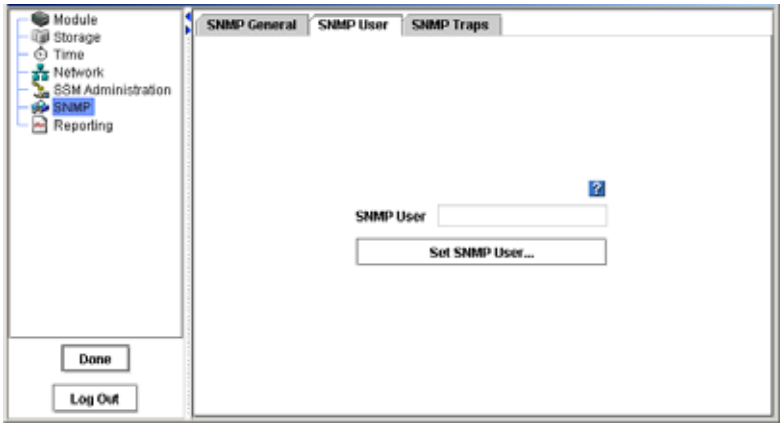

7.3 Selecting an Existing Administrative User as SNMP User

1. Click Set SNMP User.
   The Set SNMP User window opens, shown in Figure 86.

Figure 86. Set SNMP User

2. Select Existing User.
   The window displays the required fields for selecting an existing user, shown in Figure 87.

Figure 87. Selecting an Existing User

3. Select the name of the existing administrative user in the User Name field.
4. Type the user password in the Password field.

Note: You can change the password for the selected user by selecting Change Password and then entering the new password.

5. Click OK.

7.4 Adding New SNMP User

Add a new SNMP User.

1. Click Set SNMP User.

   The Set SNMP User window opens, shown in Figure 86. If the SNMP user already exists, the name will be displayed in the User Name field.

2. Select New User.

   The window displays the required fields for creating a new user, shown in Figure 88.

Figure 88. Adding a new SNMP User

3. Type in the User Name.
4. Enter the password for the new user.
5. Confirm the password.
6. Click OK.

If there is an existing SNMP user, the Delete User window opens, asking whether you want to delete existing users.

7. Click Yes to delete the existing SNMP user and replace it with the new one.

   or

   Click No to replace the SNMP User with the new user without deleting the old user.

   The new user will be the SNMP proxy user and the old user will be inactive, but will be listed in the SSM Administration configuration category as an administrative user.

7.5 Changing the Password for Existing SNMP User

1. Click Set SNMP User.

   The Set SNMP User window opens, shown in Figure 87. The name of the existing SNMP user appears in the User Name field.

2. Type the user's current password.
3. Select Change Password.
4. Type the new password.
5. Confirm the new password.
6. Click OK.

7.6 Enabling the SNMP Agent

In order to enable the SNMP Agent, you must add an SNMP User first. If no SNMP User is designated, a message opens notifying you to add an SNMP user.

1. Add an SNMP User.

   See "Configuring the SNMP User" on page 99.

2. Click Enable SNMP Agent.

   The Enable Agent fields become activated, shown in Figure 89.

Figure 89. Enabling the SNMP Agent

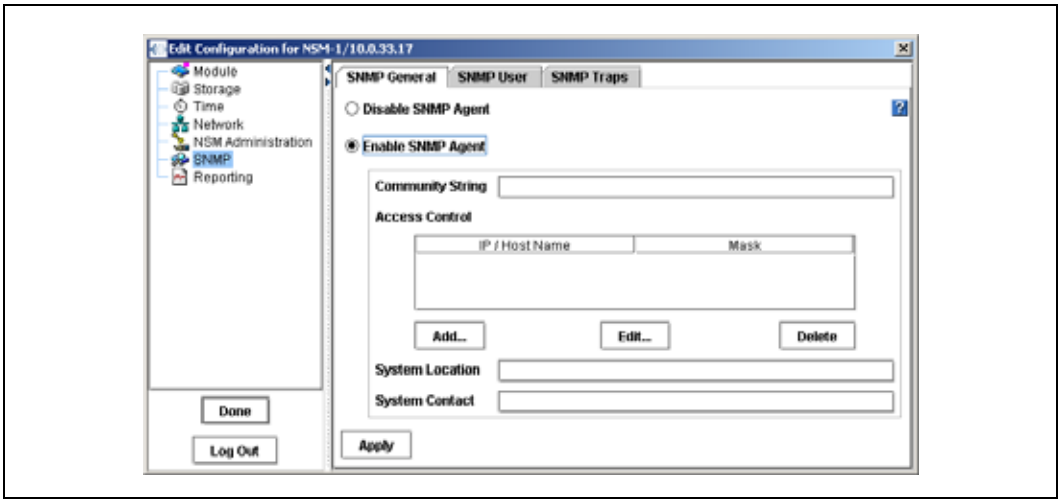

3. Type the Community String.

Note: The community string identifies a group of hosts that are allowed read-only access to the SNMP data. The community "public" is typically used to denote a read access community. This string is entered into the SNMP Management tool (not included) when attempting to access the system.

7.7 Choosing Access Control

1. Click Add to add an SNMP client.

   The Add SNMP Client window opens, shown in Figure 90. You can add SNMP Client either by specifying IP addresses or host names.

Figure 90. Adding an SNMP Client

7.7.0.1 By Address

1. Click By Address.
2. Type the IP Network Address.
3. Select an IP Netmask from the list. Select Single Host if the SNMP Client is a single computer.
4. Click OK.

   The IP address and netmask appear in the Access Control list.

Note: 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 that match that IP and netmask combination can access SNMP.

7.7.0.2 By Name

1. Click By Name.
2. 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. See "Using a DNS Server" on page 77.
3. Click OK.

   The host name appears in the Access Control list.

7.8 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 91.

Figure 91. Editing a Host in the Access Control List

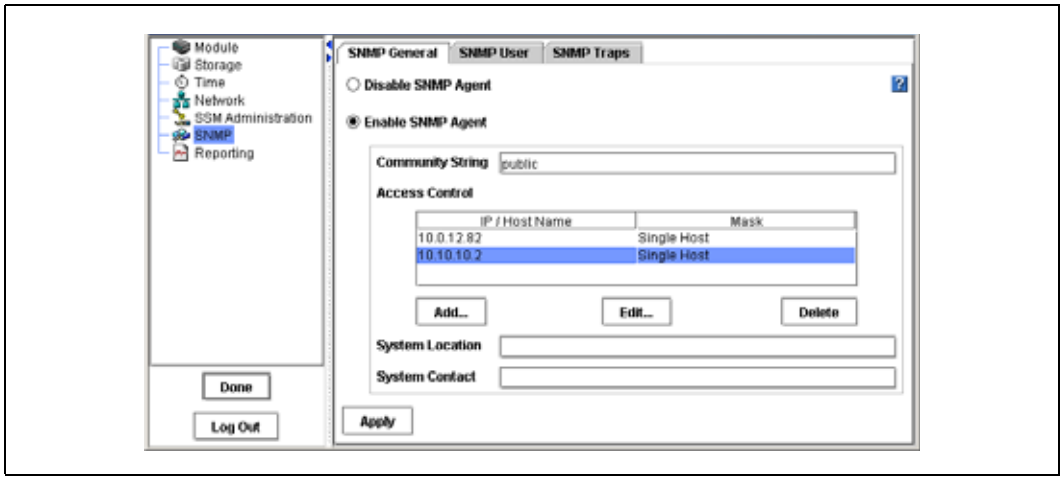

2. Click Edit.

   The Edit SNMP Client window opens, shown in Figure 92.

Figure 92. Editing SNMP Client from the Access Control List

3. Change the appropriate information.

4. Click OK.

7.9 Deleting Access Control Entries

Delete an SNMP client from the list.

1. Select a host listed in the Access Control list, shown in Figure 91.
2. Click Delete.
   A confirmation message opens.
3. Click OK.

7.10 Entering System Information (Optional)

1. Enter System Location information such as address, building name, room number, etc. Normally this will be network administrator information — the person you would contact if you could not connect to SNMP clients.
2. Enter System Contact information such as name, telephone, e-mail, etc.

7.11 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.

7.12 Locating the Storage System MIB

The Storage System MIB files are located on the Storage System Console CD under the MIBS folder. 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 CD. Load these MIB files in the Network Management System if required.

7.13 Disabling the SNMP Agent

Disable the SNMP Agent if you do not plan to use SNMP applications to monitor your network of SSMs.

1. On the Network View, double-click the SSM and log in, if necessary.
   The SSM Configuration window opens.
2. Select SNMP from the configuration categories.
   The SNMP General window opens, shown in Figure 84.

7.14 Disabling SNMP

1. On the SNMP General window, select Disable SNMP Agent.
2. Click Apply.

7.15 Enabling SNMP Traps

Enable SNMP Traps if you plan to use an SNMP tool to notify you when a monitoring threshold is reached.

1. On the Network View, double-click the SSM and log in, if necessary.
   The SSM Configuration window opens.
2. Select SNMP from the configuration categories.
   The SNMP General window opens, shown in Figure 84 on page 99.
3. Select the SNMP Traps tab.
   The SNMP Traps window opens, shown in Figure 93.

Figure 93. Enabling SNMP Traps

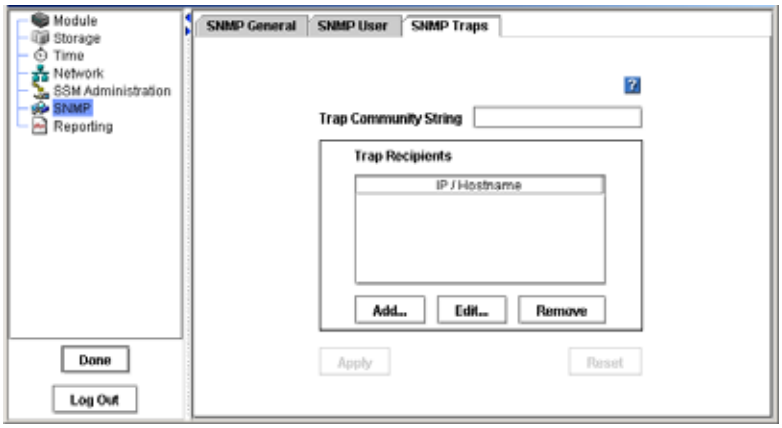

7.16 Enabling SNMP Traps

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 94. Adding an SNMP trap Recipient

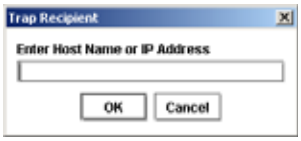

3. Enter the host name or IP address for the SNMP client that is receiving the traps.
4. Click OK.
5. Repeat steps 2 through 4 for each host in the trap community.
6. Click Apply when you are finished adding hosts.

7.16.0.1 Editing the Trap Recipient

1. Select the host you want to change from the list of Trap Recipients.
2. Click Edit.
   The Trap Recipient window opens.
3. Change the host name or IP address.

4. Click OK.
5. Click Apply when you are finished editing trap recipients.

7.16.0.2 Removing the Trap Recipient

1. Select the host you want to remove from the list of Trap Recipients.
2. Click Remove.
   A confirmation window opens.
3. Click OK to remove the trap recipient.
   The host is removed from the list.
4. Click Apply when you are finished removing trap recipients.

7.17 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.

Reporting 8

8.1 Reporting Overview

The SSM offers multiple reporting capabilities. Use reporting to:

- View real-time statistical information about the SSM
- View and save log files
- Set up active monitoring of selected variables
- Set up e-mail notification
- View alerts
- Run hardware diagnostics

When you select Reporting from the SSM configuration category list, the Reporting window opens, shown in Figure 95.

Figure 95. Viewing the Reporting Window

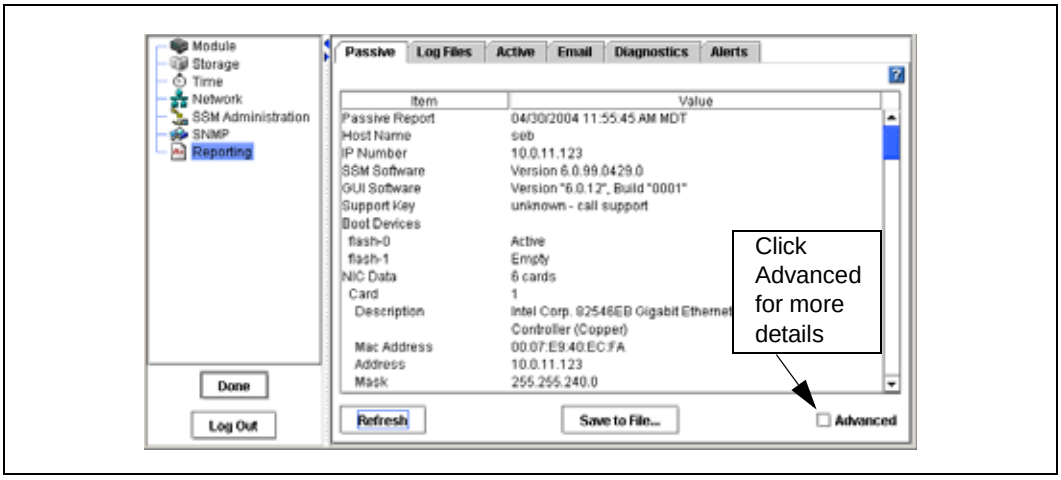

8.2 Using Passive Reports

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 Passive tab. Select Advanced to see additional statistics.

1. Select Reporting from the configuration categories.
   The Reporting window opens.
2. Click Refresh to display statistics on the Passive tab.
3. [Optional] To view extended statistics about the SSM, click Advanced at the bottom of the window.

8.2.1 Saving the Report to a File

1. On the Passive Tab, click 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 .doc extension. It is a text file and will open with Word in Windows, or any text editor in Linux or UNIX.

8.2.2 Passive Reporting Detail

This list details selected information available on the Passive Reporting window. Not all items are listed here.

Table 25. Selected Details of the Passive Report

This Term	Means This
Passive Report on	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 Storage System Console.
Boot Devices	Status information about the compact flash card(s) used to boot the SSM.
NIC Data	Information about NICs in the SSM.
Card	Indicates which NIC in the list is being described.
Description	Card name/manufacturer and capable speed of the NIC.
MAC Address	Physical address of the NIC. Each card has a unique MAC (media access control) address.
Address	IP address of the NIC.
Mask	Network mask for NIC.
Gateway	Gateway that the SSM is using.
Mode	Shows manual/auto/disabled. Manual equals a static IP, auto equals DHCP, disabled means the interface is disabled.
DNS Data	Information about DNS, if a DNS server is being used.
Server 1, Server 2	IP address of the DNS servers.
Memory	Information about memory in the SSM.
Total	Total amount of memory in KB.
Free	Total amount of free memory in KB.
CPU	Information about the CPU.
Speed	Clock speed of the microprocessor.
Load Average	Information about the average load on the system.
RAID	Information about RAID.

Table 25. Selected Details of the Passive Report (Continued)

This Term	Means This
Minimum Rebuild Speed	Minimum amount of data in KB/sec that will be transferred during a RAID rebuild. The higher this number, the less bandwidth available for users because the system will not transfer at a rate lower than what is set.
Maximum Rebuild Speed	The maximum amount of data in KB/sec that will be transferred during a RAID rebuild.
Statistics	Information about the RAID for the SSM.
Unit #	Identifies disks that make up the RAID configuration, their RAID level, chunk size, and device name.
Drive Status	Information about the drives in the SSM.
Drive #	Drive 1 through 16. Indicates a specific drive in the SSM.
Capacity	Size of the drive.
Drive Interface Card	Lists RAID card model numbers and version numbers.

8.3 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 112.

1. Select Reporting from the configuration categories.
   The Reporting Window opens.
2. Select the Log Files tab.
   The Log Files window opens, shown in Figure 96.

Figure 96. Saving Log Files to a Local Machine

3. Scroll down the Choose Logs to Save list.
4. Select the file or files you want to save.
   To select multiple files, use the Ctrl key.
5. Click Save Files.
   The Save dialog opens.
6. Select a location for the file or files.
7. Click Save.
   The file or files are saved to the designated location.

8.3.1 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.

8.3.1.1 Adding a Remote Log

1. Select Reporting from the configuration categories.
   The Reporting Window opens.
2. Select the Log Files tab.
   The Log Files window opens, shown in Figure 96.
3. Click Add below the list of remote logs.
   The Add Remote Log Target window opens, shown in Figure 97.

Figure 97. Adding a Remote Log

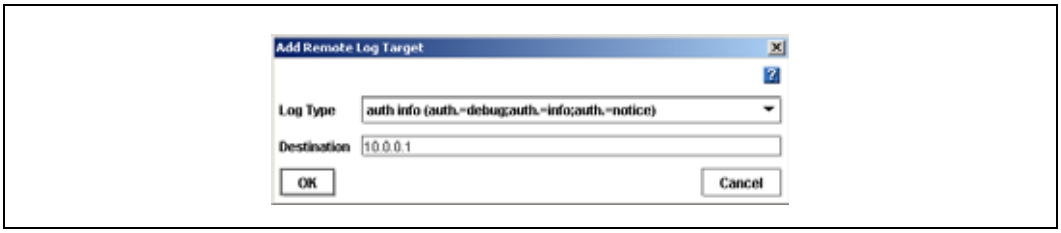

4. In the Log Type list, select the log that you want to direct to a remote computer.
   The Log Type list only contains logs that support syslog.
5. In the Destination field, type the IP address or host name of the computer that will receive the logs.
6. Click OK.
   The remote log displays in the Remote logs list on the Log Files tab.

8.3.1.2 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.”

8.3.1.3 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.

8.3.1.4 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.

8.4 Using Active Monitoring

Use active monitoring to track the health of the SSM. Active monitoring allows you to set up notification through e-mails, alerts in the Storage System Console, and SNMP traps. You can choose which variables to monitor and choose the notification methods for alerts related to the monitored variables. For a detailed list of monitored variables, see "List of Monitored Variables" on page 118.

Note: Critical variables, such as the CPU temperature and motherboard temperature, have thresholds that trigger a shutdown of the SSM.

1. On the Network View, double-click the SSM and log in, if necessary.
   The SSM Configuration window opens.
2. Select Reporting from the configuration categories.
   The Reporting Window opens.
3. Select the Active tab.
   The Active Reporting window opens, shown in Figure 98.

Figure 98. Setting Active Monitoring Variables

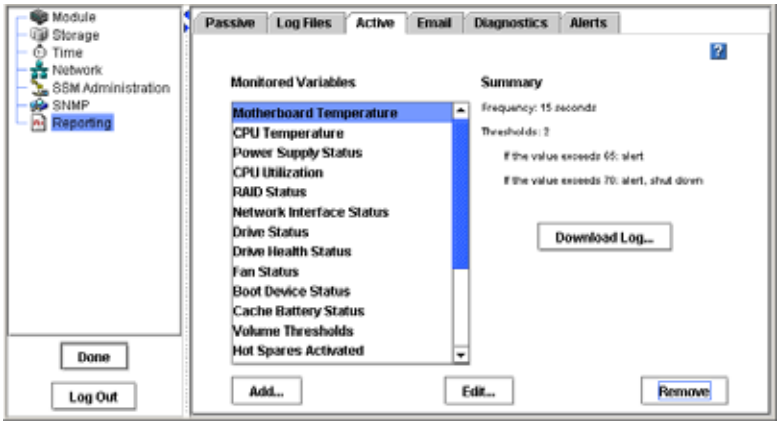

8.4.1 Adding Variables to Monitor

The variables that the SSM is currently monitoring are listed in the box. All variables are configured and set for Console alerts. You can only add variables that have been previously removed.

1. Click Add.
   The Configure Variable wizard opens to Step 1, shown in Figure 99.

Figure 99. Adding a Variable, Step 1

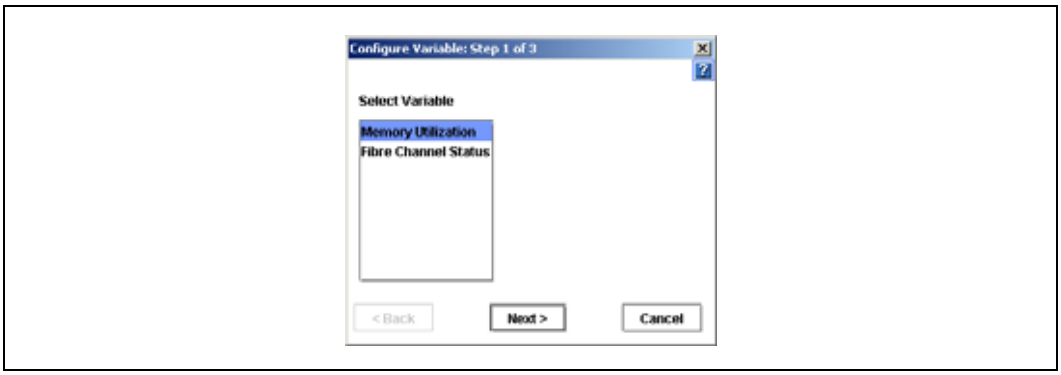

2. Select the variable that you want to monitor and click Next.
   The Configure Variable wizard, Step 2, opens, shown in Figure 100.

Figure 100. Adding a Variable, Step 2

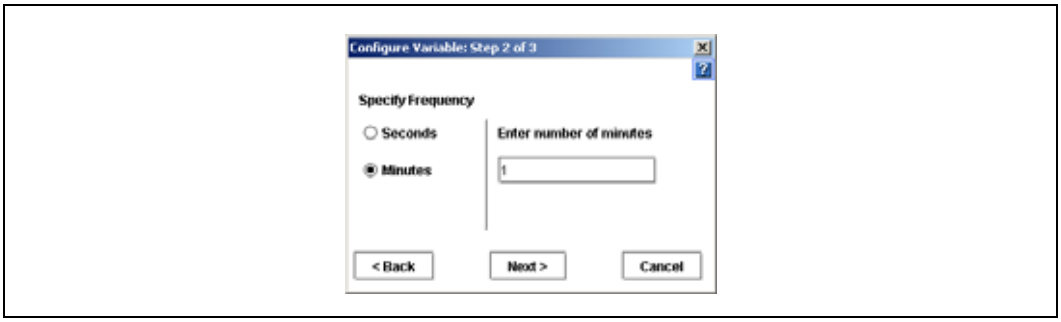

3. Specify the frequency for monitoring the variable and click Next.
   The Configure Variable wizard, Step 3, opens.

Note: If you change the frequency for a variable, the log file for that variable is **recreated**. If you want to save the history, download a copy of the log file before you change the frequency.

Figure 101. Setting Alerts for Monitored Variables

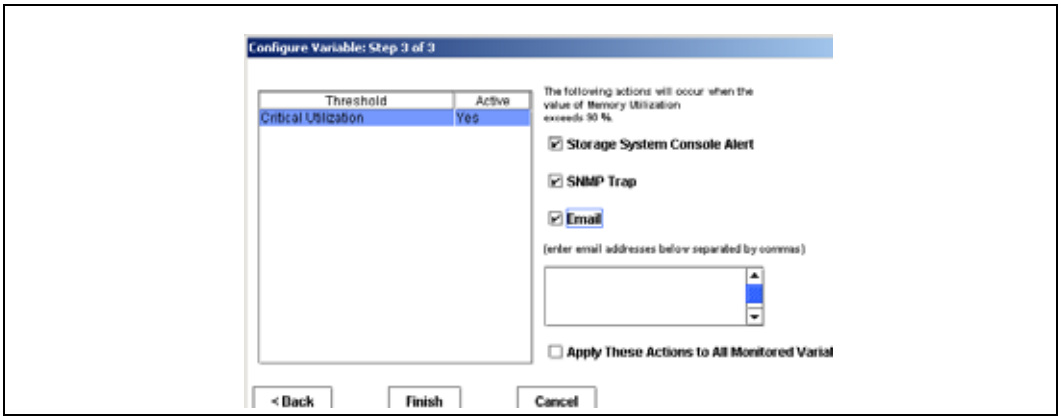

4. For each threshold listed, select the type of alert you want to receive.

Types of Alerts Available for Active Monitoring

Type of Alert	Where Alerts Are Sent	For More Information
Console alerts	To the Alert Message area of the Storage System Console.	See "Alert Messages View" on page 9 in Chapter 1, "Getting Started."
SNMP traps	To the SNMP trap community managers. You must have configured the SSM to use SNMP.	See "Enabling SNMP Traps" on page 106 in Chapter 7, "Using SNMP."
Email	To specified email addresses. Type the email addresses to receive the notification, separated by commas. Then configure Email Notification.	See "List of Monitored Variables" on page 118.

5. [Optional] To apply the alert actions (including the e-mail 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.

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 e-mail 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.

8.4.2 Downloading a Variable Log File

If you change the frequency of a monitored variable, the log file is recreated. To save the history of a variable, download a copy of the log file before you change the frequency specification of a variable.

1. In the list of monitored variables, click the variable for which you want to save the log file.
2. Click Download Log 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.

8.4.3 Editing a Variable

Depending upon the variable, you can edit the frequency and the notifications. See the variable specifications in "List of Monitored Variables" on page 118.

1. Select from the list the variable you want to edit.
2. Click Edit.
   The Configure Variable wizard opens to Step 2, seen in Figure 100 on page 115.

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.

Note: If you change the frequency for a variable, the log file for that variable is recreated. If you want to save the history, download a copy of the log file before you change the frequency.

4. [Optional] Change the alert notification method.
5. [Optional] To apply the alert actions (including the e-mail 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.

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 e-mail 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.

8.4.4 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.

Figure 102. Viewing the Monitoring Variable Summary on the Active Window

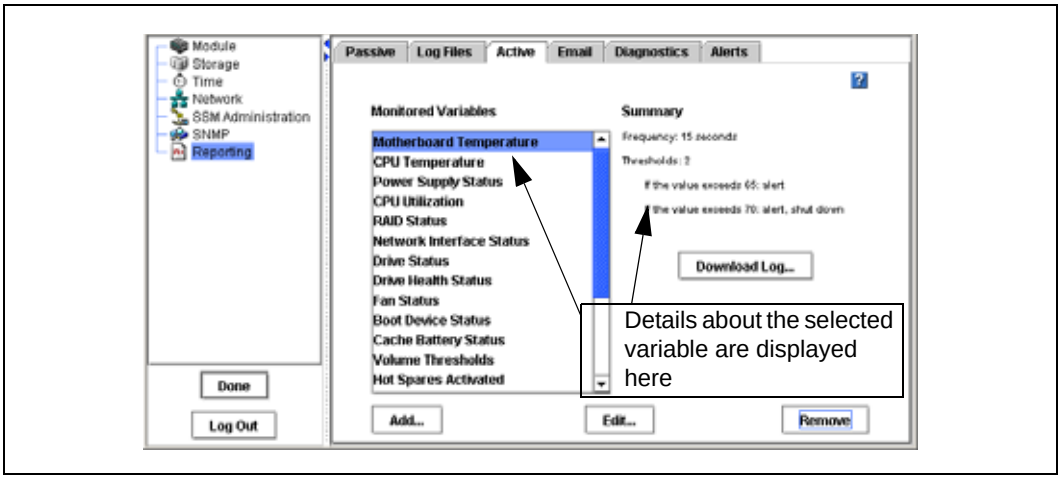

8.4.5 Removing a Variable from Active Monitoring

Some variables can be removed from active monitoring. Permanent variables, such as motherboard temperature, cannot be removed. See “List of Monitored Variables” on page 118.

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.

8.4.6 List of Monitored Variables

Table 26 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.
- The preset threshold for the variable.

Table 26. List of Variables Available for Active Monitoring

Variable Name	Units	Perm. Variable	Specify Freq.	Default Freq.	Default Action/ Threshold	Preset Thresholds
CPU Temperature	Celsius	Yes	No	15 seconds	Shutdown, Console alert at > 80º	--
Motherboard Temperature	Celsius	Yes	No	15 seconds	Shutdown, Console alert at > 60º	--
Power Supplies (Primary and Secondary)	--	No	Yes	1 minute	Console alert at Faulty	--
Memory Utilization	Percent	No	Yes	1 minute	Console alert at > 90%	Warning at 80% Critical at 90%
CPU Utilization	Percent	No	Yes	1 minute	Console alert at > 90%	Warning at 80% Critical at 90%
Drive Status (1 through 16)	--	No	Yes	1 minute	Console alert if changes	--
RAID Status	--	Yes	Yes	15 seconds	Console alert if changes	--
Volume Thresholds	--	No	Yes	15 minutes	Console alert if threshold exceeded for any volume or snapshot in the mgt. group	User sets in Storage System Console

Table 26. List of Variables Available for Active Monitoring (Continued)

Variable Name	Units	Perm. Variable	Specify Freq.	Default Freq.	Default Action/ Threshold	Preset Thresholds
Hot Spares Activated	-	No	Yes	15 minutes	Console alert if a spare is activated	--
Volume Status	-	No	Yes	15 minutes	Console alert if volume status changes	--
Network Interface Status	-	No	Yes	1 minute	Console alert if NIC status changes	--
Snapshot Status	-	No	Yes	1 minute	Console alert if snapshot status changes	--
Remote Copy Status	-	No	Yes	15 minutes	Console alert if fails	--
Remote Copy Complete	-	No	Yes	15 minutes	Console alert if true	--
Remote Copy Failovers	-	No	Yes	15 minutes	Console alert if true	--
Remote Management Group Status	-	No	Yes	1 minute	Console alert if changes	--
Volume Restripe Complete	-	No	Yes	1 minute	Console alert if completed	-
Drive Health Status	Normal, Marginal, Faulty	Yes	Yes	1 minute	Console alert if not normal	-
Fan Status	Normal, Faulty	No	Yes	1 minute	Console alert if changes	-
Boot Devices Status	Active, Inactive, Failed, Empty, Unformatted, Not recognized, Unsupported	No	Yes	30 seconds	Console alert if not normal	-
Fibre Channel Status	Waiting for firmware, Active, Enabled, Not ready, Initialized	No	Yes	1 minute	Console alert if changes	-
Cache Battery Status	Normal, Missing, Faulty	No	Yes	1 minute	Console alert if changes	-

8.5 Setting E-mail Notification

Use the e-mail tab to configure the SMTP settings if you request e-mail notification on the Active tab.

To request e-mail notification when alerts occur:

1. On the Network View, double-click the SSM and log in, if necessary.
   The SSM Configuration window opens.
2. Select Reporting from the configuration categories.
   The Reporting Window opens.
3. Select the E-mail tab.
   The E-mail window opens, shown in Figure 103.

Figure 103. Configuring E-mail Settings for E-mail Alert Notifications

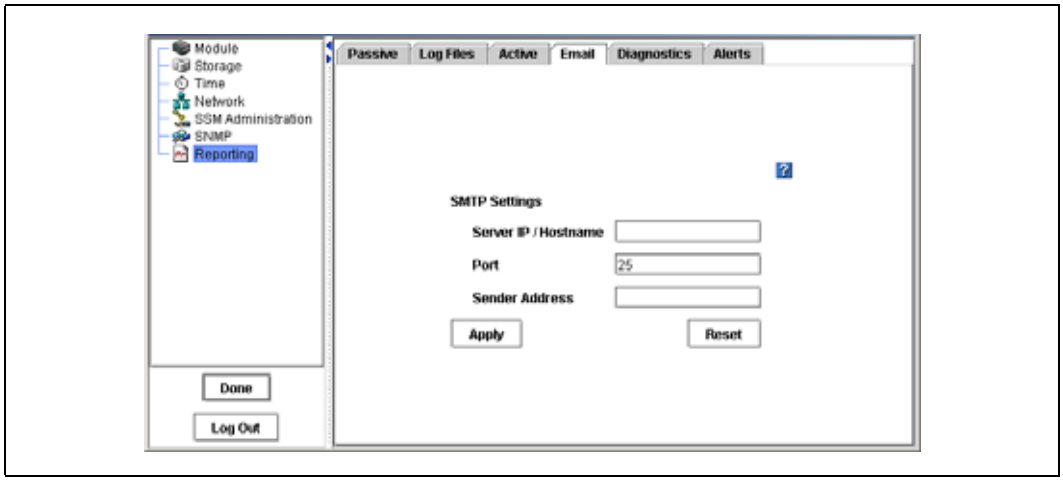

4. Enter the IP or host name of the e-mail server.
5. Enter the port.
   The standard port is 25.
6. (Optional) If your e-mail server is selective about valid sender addresses on incoming e-mails, enter a sender address.
   If you do not enter a sender address, e-mail notifications will display “root@hostname,” where hostname is the name of the SSM, in the e-mail From field.
7. Click Apply.

Note: If you are requesting e-mail notification, be sure to set up the SMTP settings on the E-mail tab.

8.6 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. For help resolving hardware problems, call your Technical Support representative.

To run diagnostic tests:

1. On the Network View, double-click the SSM and log in, if necessary.
   The SSM Configuration window opens.
2. Select Reporting from the configuration categories.
   The Reporting Window opens.
3. Select the Diagnostics tab.
   The Diagnostics window opens, shown in Figure 104.

Figure 104. Viewing the List of Diagnostics

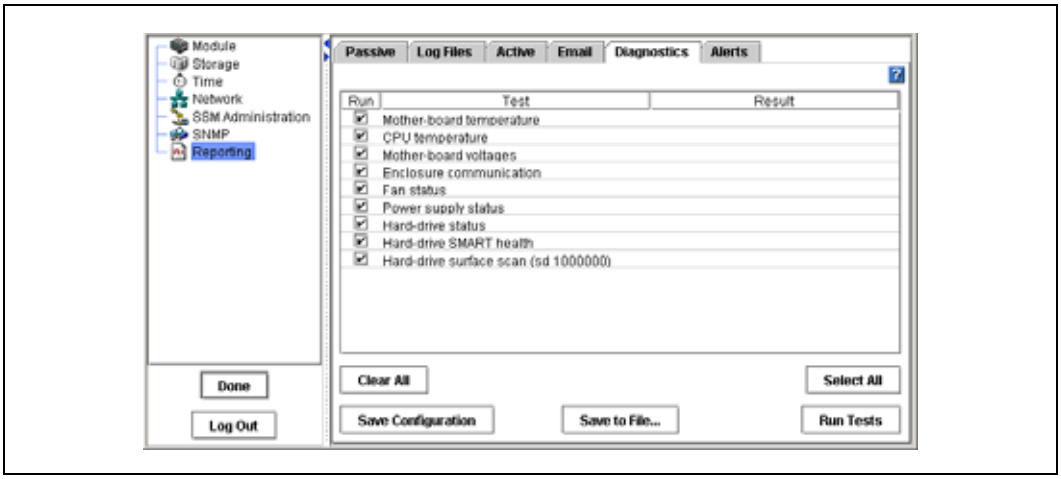

4. Select the diagnostic tests that you want to run.
   The default setting is to run all tests. Click to deselect any tests that you do not want to run. To clear all selections, click Clear All.

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.

5. Click Run Tests.
   A progress message displays. When the tests complete, the results of each test display in the Result column.
6. [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 ".doc" file in the designated location. It is a text file and will open with Word in Windows, or any text editor in Linux or UNIX.
7. [Optional] To save the list of diagnostics that you selected so that next time you open the Diagnostics window it will be preconfigured with your selections, click Save Configuration.

8.6.1 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 (.doc) file.
3. Open the report file.

8.6.2 List of Diagnostic Tests

Table 27 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.

Table 27. List of Hardware Diagnostic Tests and Pass/Fail Criteria

Diagnostic Test	Description	Pass Criteria	Fail Criteria
Mother board 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 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
Hard drive status	Checks the status of each of the eight drives.	All drives are “On and Secured”	One or more drives not “On and Secured”
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
Hard drive surface scan test	Performs a partial read-scan across the surface of all disk drives (SDA-SDP).	No read errors found	Read errors found on one or more disks

8.7 Viewing Alerts

Any time that an actively monitored variable causes an alert, the alert is logged by the SSM. If the Storage System Console is open, alerts display in the Alerts Message area on the Console main window. If the Console is not open, these alerts are still logged, and you can view them next time you open the Console.

Note: The Alerts tab 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.

1. On the Network View, double-click the SSM and log in, if necessary.
   The SSM Configuration window opens.
2. Select Reporting from the configuration categories.
   The Reporting Window opens.
3. Select the Alerts tab.
   The Alerts window opens, shown in Figure 105.

Figure 105. Viewing Alerts

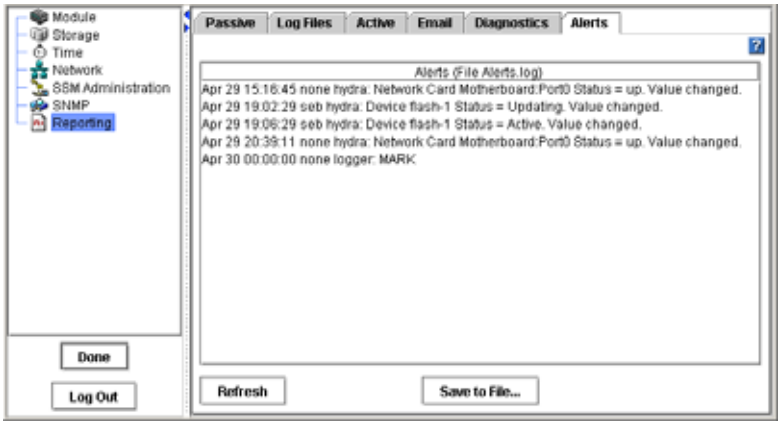

4. To refresh the list of alerts, click Refresh.
5. [Optional] To save the list of alerts, click Save to File. Then select a location for the file.

Working with Management Groups 9

9.1 Management Group Overview

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.

9.1.1 Topics Covered in This Chapter

- Managers
- Quorum
- Setting the management group time
- Setting the local bandwidth
- Backing up the management group configuration

9.1.2 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.

9.1.2.1 Functions of Managers

Managers control data replication, keep track of system status, coordinate reconfigurations as SSMs are brought up and taken offline, and re-synchronize data when SSMs fail and recover.

9.1.2.2 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 Engine to function. Therefore, for optimal fault tolerance, you should have 3 or 5 managers in your management group. Three or five managers provide the best balance between fault tolerance and performance.

Table 28. Managers and Fault Tolerance Management Groups

Number of Managers	Number for a Quorum	Management Fault Tolerance	Explanation
1	1	None	If the manager fails, no data control takes place.
2	2	None	If one manager fails, there is not a quorum. Not Recommended
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	If one manager fails, 3 remain, so there is still a quorum.
5	3	High	If one or two managers fail, 3 remain so there is still a quorum.

9.1.3 Communication Mode

The Storage Server Console and Storage System Engine support unicast communication among SSMs and application servers.

9.1.3.1 Unicast Communication

Unicast is communication between a single sender and a single receiver over a network. Unicast communication allows application servers to direct messages to SSM managers which are located in different subnets. When you configure application servers to access Storage System Engine volumes, you must use the IP addresses of the SSMs that are running managers.

9.1.3.2 Adding or Removing Managers

Any time you add or remove managers in a management group, a window opens which displays all the IP addresses of those managers along with a reminder to reconfigure the application servers that are affected by the change.

Note: Unicast requires that the SSMs running managers have static IP addresses (or reserved IP addresses if using DHCP). See "Configuring the IP Address Manually" on page 59 in Chapter 4, "Managing the Network."

9.2 Requirements for Creating Management Groups

When creating a management group, you must configure the following parameters.

Management Group Requirement	What it means
Configure SSMs	Before you create a management group, you should configure all the SSMs for that management group. When planning your storage, remember that all SSMs in a **cluster** must be configured alike. Refer to “Configuration Tasks” on page 1 in Chapter 1, “Getting Started.”
Log in to SSMs	You must be logged in to the SSM 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 SSM added to the group has the manager started automatically. You can add managers to other SSMs later.
Assigning manager IP addresses	The SSMs that are running managers must have static IP addresses (or reserved IP addresses if using DHCP). See “Configuring the IP Address Manually” on page 59 in Chapter 4, “Managing the Network.”

9.3 Creating a Management Group

Create a management group as 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
- Setting the local bandwidth

9.3.1 Getting There

1. Open the Console.

   If you have not created a management group, but you have some SSMs on the network, the Console displays those SSMs as available. See Figure 106.

Figure 106. Viewing SSMs Before Creating a Management Group

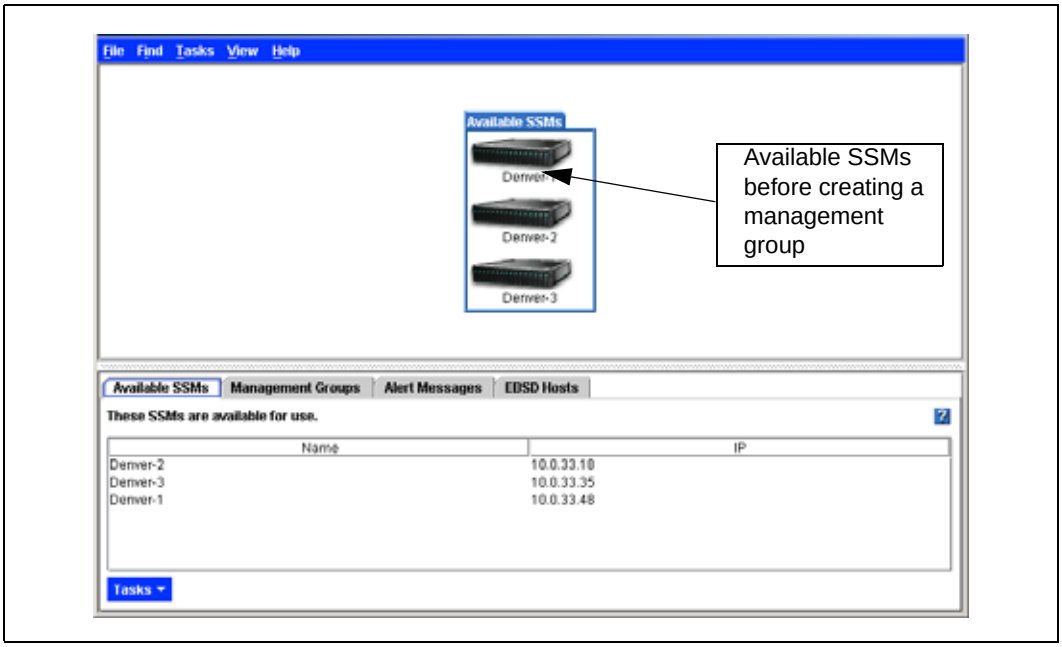

2. Log in to one or more of the SSMs you want to add to the new management group.
3. Click Done from the SSM Edit Configuration window to return to the main Console window with the SSM tab view, shown in Figure 107.

Figure 107. The SSM Tab View

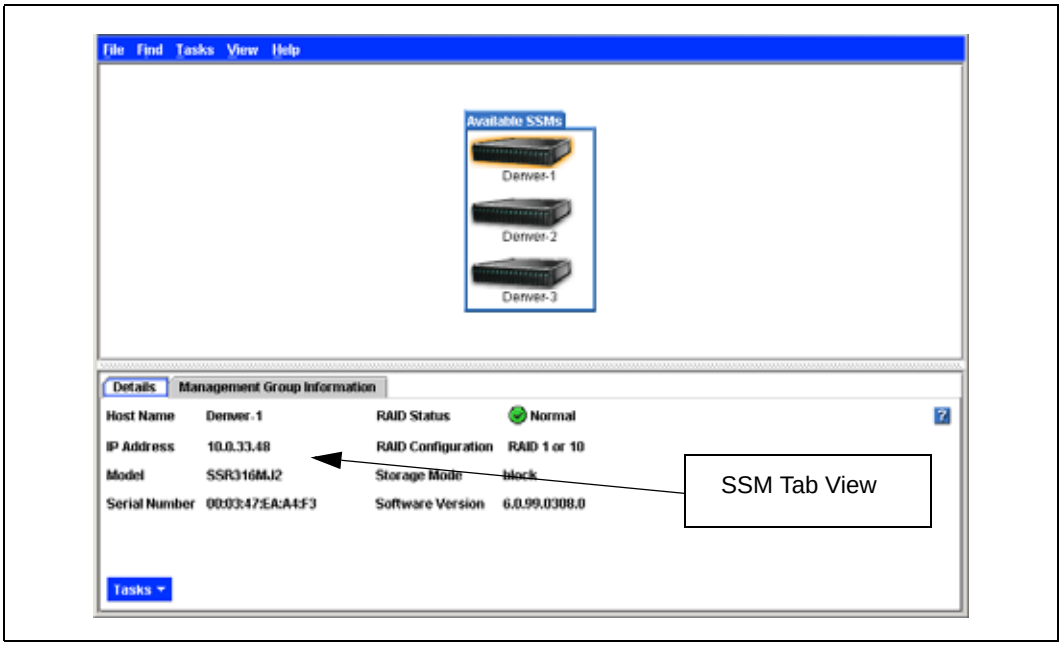

9.3.2 Adding the First SSM to Create a New Management Group

1. Select the first SSM to include in the management group.
2. Click the Management Group Information tab in the tab view, shown in Figure 108.

Figure 108. Management Group Information Tab

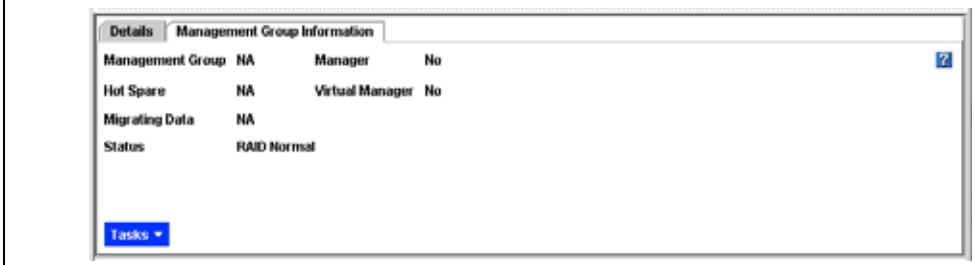

3. Click the Tasks button and select Add to new or current Management Group.
   The Add to or create a Management Group window opens, shown in Figure 109.

Figure 109. Creating a New Management Group

4. Select New Management Group and type a name for the management group.
5. Click Add.
   A message opens, shown in Figure 110.

Figure 110. Starting Manager Message for SSM Joining a Management Group

6. Click OK.
   The Managers IP Addresses window opens, shown in Figure 111.

Figure 111. List of Manager IP Addresses for Management Group

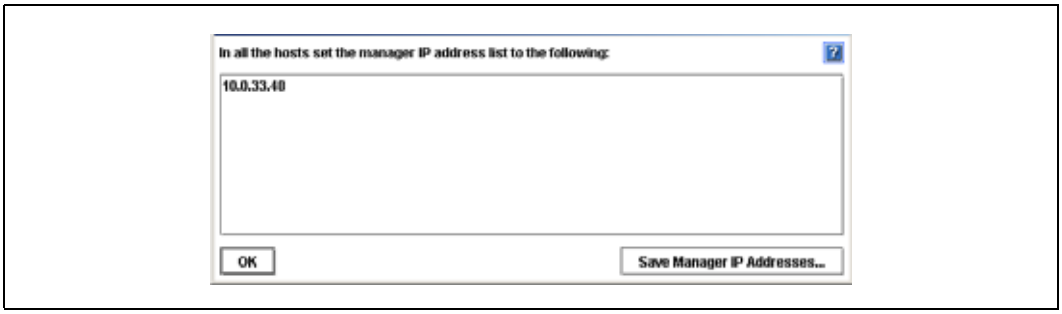

7. Click OK.

   The SSM joins the management group and starts the manager. The Console displays the newly created management group, shown in Figure 112.

Figure 112. New Management Group with One SSM

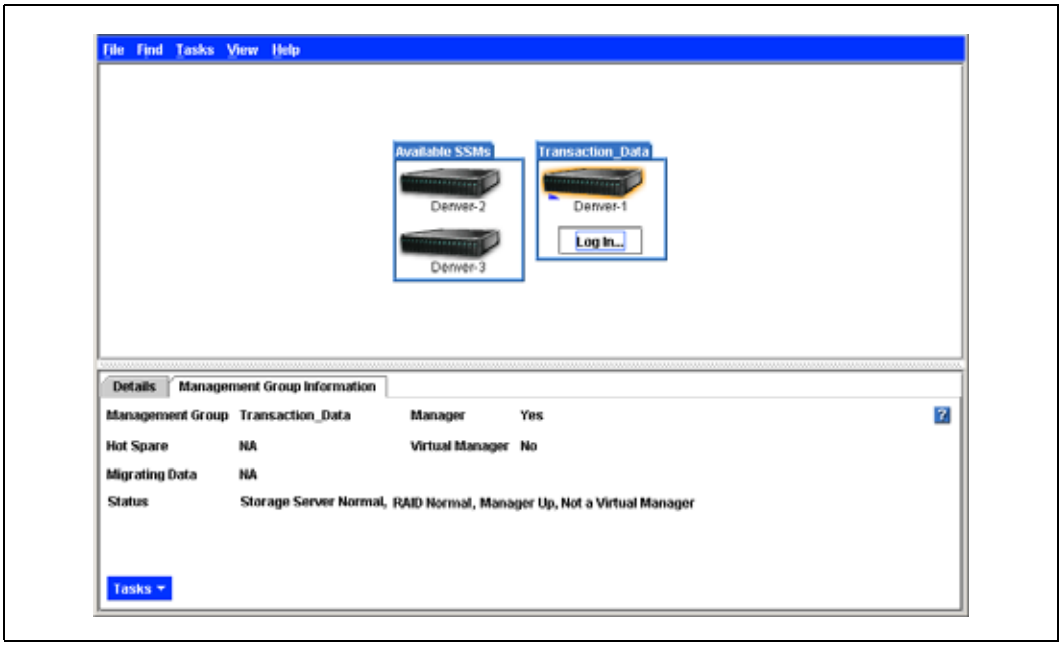

8. [Optional] Select the next SSM that you want to add.
9. Click Add to Management Group on the Management Group Information tab.

   The Add to or create a Management Group window opens with the existing management group selected.
10. Click Add.

Note: If you are not logged in to the SSM, you are prompted to do so now.

The SSM is added to the specified management group.

11. Repeat steps 8 through 10 to add additional SSMs.

9.3.3 Adding Managers to the Management Group

After adding the SSMs to the management group, you can start managers on the additional SSMs in the management group. The number of managers you start depends upon the overall design of your storage system. See “Managers Overview” on page 125 for more information about how many managers to add.

1. Select an SSM in the management group on which to start a manager.

   The SSM tab view opens.

2. On the Management Group Information tab, select the Tasks menu and click Start Manager, as shown in Figure 113.

Figure 113. Starting a Manager

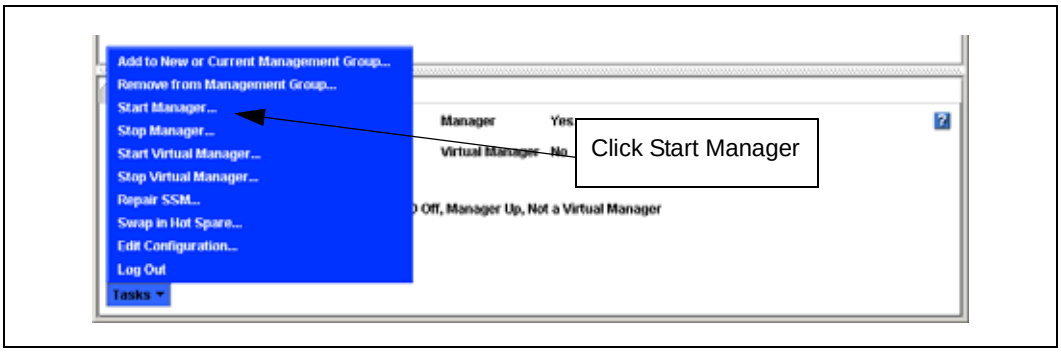

3. Repeat steps 1. and 2 to start managers on additional SSMs.

9.3.4 Logging In to a Management Group

You must log in to a management group to administer the functions of that group.

Note: Log in to a management group by logging in to an SSM that is designated as a manager for that management group.

1. Click Log In on the management group in the Network View, shown in Figure 114.

Figure 114. Logging in to a Management Group

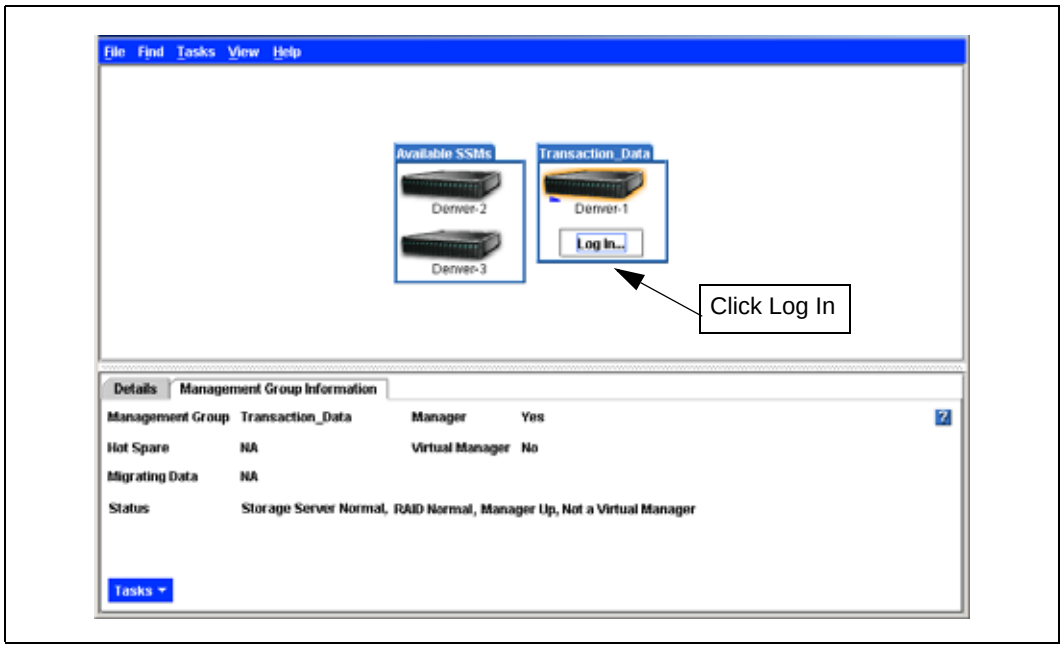

The Manager Log In List window opens, shown in Figure 115. Any SSMs to which you are already logged in display Yes in the Logged In column.

Figure 115. List of SSMs Running Managers

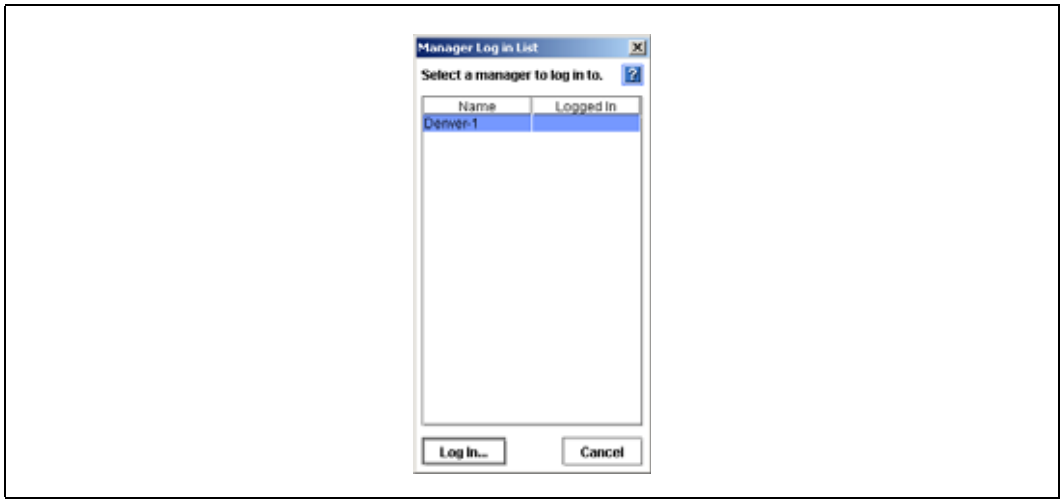

2. Select an SSM and click Log In.

   Whatever view of the Console is displayed when you log in to a management group, that is the view that returns after logging in to that management group.

Note: If you are not already logged in to the SSM, you are prompted to do so now.

Figure 116. Viewing a Management Group in the Console

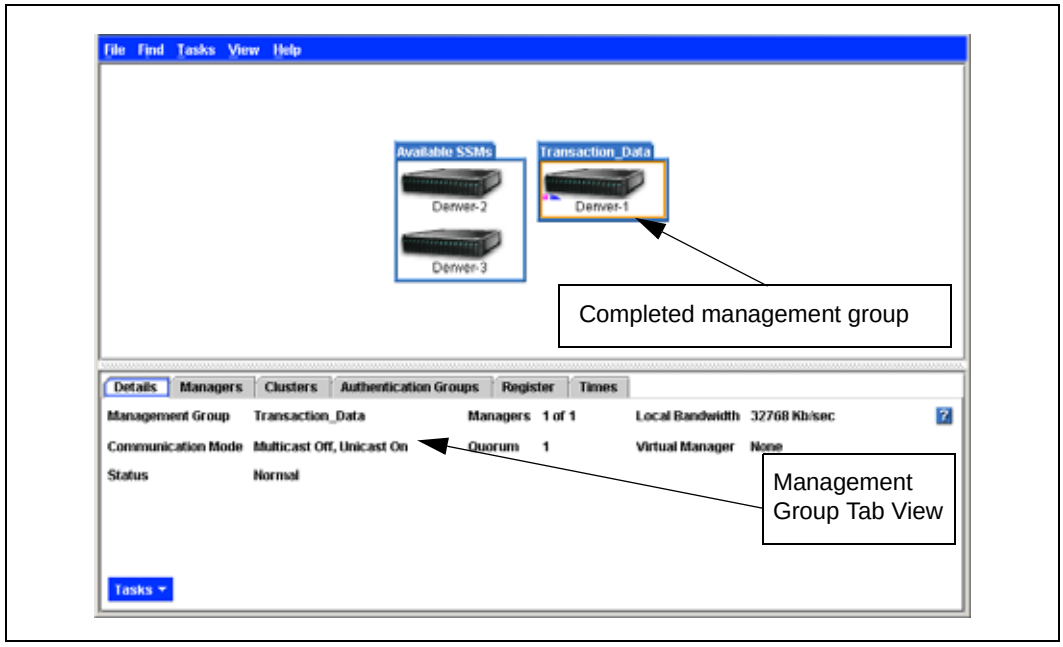

9.3.5 Management Group Tab View

When you have logged in to and selected a management group, the management group tab view opens. The tabs provide access to management group information and features, such as viewing management group properties, registering features, and creating clusters and authentication groups.

9.3.5.1 Details Tab

Details about the management group are listed along with the Task menu for taking action on the management group.

Information provided on the Details tab includes:

- The name of the management group
- The number of managers operating within the group
- How many of those managers are required for a quorum and are operational (see “Managers Overview” on page 125 for information about quorums)
- The synchronization bandwidth that is set when you edit the management group
- The communication mode of the management group

9.3.5.2 Managers Tab

All the SSMs included in the management group are listed on the managers tab. SSMs that are running managers display Yes in the Manager column.

Information provided on the Managers tab includes:

- Host name of the SSM
- Whether that SSM is running a manager
- Whether that SSM is running a virtual manager

9.3.5.3 Clusters Tab

All the clusters created within the management group are listed on the clusters tab. For more information, see Chapter 11, "Working with Clusters."

Information provided on the Clusters tab includes:

- Name of the cluster

9.3.5.4 Authentication Groups Tab

All the authentication groups that are associated with volumes in a management group are created and managed from the management group. For more information about working with authentication groups, see "Working with Authentication Groups" on page 215.

Information provided on the authentication groups tab

- Name of the authentication group
- Authentication mode of the group
- The subnet/mask of the authentication group, if the authentication is set up for subnet and mask
- Volumes associated with the group

9.3.5.5 Register Tab

Register to use add-on features available for specialized storage features. Information available on the Register tab includes:

- The number of licenses, if any, for the upgrade features.
- Version information about the software components of the system. The version information is provided for customer support should you ever have a support issue.

9.3.5.6 Times Tab

Resynchronize the management group time any time you change the time on an SSM in the management group that is running a manager.

Note: Use NTP to ensure closely synchronized times on the SSMs in the management group.

Information available on the Times tab includes:

- Current time setting of the management group
- Current time setting of each SSM in the management group

9.4 Registering Add-on Modules

Add-on modules are available for specialized storage features. Intel sales issues registration numbers to activate add-on modules. The add-on modules available include

- Scalability Pak
- Configurable Snapshot Pak
- Remote Data Protection Pak
- Client Server Clustering Pak

For more information about add-on modules and feature registration, see Chapter 16, “Feature Registration.”

9.5 Editing a Management Group

When editing a management group you can change the local bandwidth. The local bandwidth setting controls the copy rate within the local management group. Therefore it sets the data restripe rate for the management group. If you use Remote Copy between two clusters within one management group, local bandwidth will also control the remote copy rate. For more information about setting the bandwidth for Remote Copy, see the Remote Copy User Manual.

Note: When Remote Copy is used to copy a snapshot from one management group to another, the remote bandwidth setting of the management group containing the remote volume determines the rate per second that the manager will devote to copying data.

9.5.1 Setting or Changing the Local Bandwidth

After a management group has been created, you can edit the management group to change the local bandwidth. This is the maximum rate per second that a manager will devote to non-application processing, such as moving data and synchronizing hot spare SSMs. The default rate is 32768 Kb (4 MB) per second. You cannot set the range below 2048 Kb (256 KB).

The bandwidth setting is in Kb (kilobytes) per second. The industry standard for networking bandwidth is in bits per second (bps). Use the following table to convert megabits to kilobytes for setting the local bandwidth.

Table 29. Typical network types with speeds in bps and KB

Network Type	Bits Per Second (bps)	Kilobytes Per Second (KB)
Fractional T1	256 Kbps	32
Fractional T1 (1/2)	768 Kbps	96
T1	1.544 Mbps	197
Bonded T1 (2)	3.088 Mbps	395
Bonded T1 (4)	6.176 Mbps	790
Ethernet, 10Base-T	10 Mbps	1280
T3	44.736 Mbps	5726

Table 29. Typical network types with speeds in bps and KB (Continued)

Network Type	Bits Per Second (bps)	Kilobytes Per Second (KB)
Ethernet, 100Base-T	100 Mbps	12,800
OC-3	155 Mbps	19,840
OC-12	622 Mbps	79,616
Ethernet, 1000Base-T	1 Gbps	128,000
OC-192	10 Gbps	1,280,000

1. Log in to the management group.
2. Click Edit Management Group.
   The Edit Management Group window opens, shown in Figure 117.

Figure 117. Editing a Management Group

3. Change the local bandwidth.
4. Click OK.
   The new rate displays on the Details tab in the management group tab view.

9.5.2 Logging Out of a Management Group

Logging out of a management group prevents unauthorized access to that management group and the SSMs in that group.

1. Select the management group to log out of.
2. From the Tasks menu on the Details tab, select Log Out of Management Group.

9.6 Adding a SSM to an Existing Management Group

SSMs can be added to management groups at any time. Adding an SSM to a management group increases the storage space available to the group. The newly added SSM can also be used as a hot spare for a cluster within the management group.

Note: All SSMs in a cluster must be configured alike. See "Configuring Storage Server Modules" on page 20 of Chapter 1, "Getting Started."

1. Select the SSM that you want to add to a management group.
   The SSM tab view opens with the Details tab on top.
2. From the Tasks menu on the Details tab, select Add to New or Current Management Group.

The Add to or create a Management Group window opens, shown in Figure 118.

Figure 118. Adding an SSM to Existing Management Group

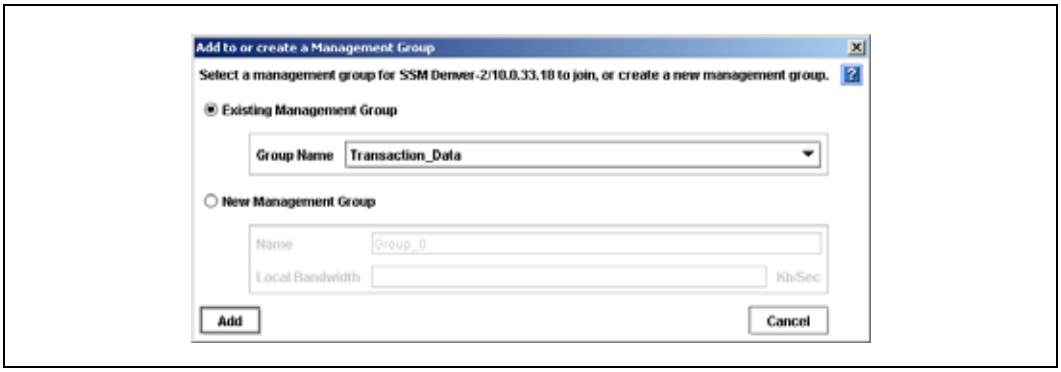

3. Select the correct management group from the list of existing management groups.
4. Click Add.

Note: If you are not logged in to the SSM that you are adding to the management group, you are prompted to log in now.

5. [Optional] If you want the SSM to run a manager, select the SSM in the management group, right-click and select Start Manager.

Adding Manager IP Addresses to Application Servers

When you add a manager to a management group, a window opens which displays all the IP addresses of the managers in that management group and a reminder to reconfigure the application servers that are affected by the change.

- Click Save Manager IP Addresses to save this list.

  You can print the list and use it as a reference when reconfiguring application servers.

9.7 Resetting the Management Group Time

Any time you change the time setting of an SSM that is running a manager, you must reset the time of the management group as well. If the manager SSM time is different from the management group time, then

- file creation times on volumes and snapshots might be affected and
- scheduled snapshots might not kick off at the intended time

Note: Use NTP to ensure closely synchronized times on the SSMs in the management group.

9.7.1 Reset Management Group Time

First verify the time settings of the SSMs running managers. If necessary, change time settings to ensure all the manager SSMs have the same time. For information about setting the time on the SSM, see Chapter 5, “Setting the Date and Time.” Then refresh the management group time.

1. Log in to the management group.
2. Select the Times tab.
3. From the Tasks menu, click Refresh All.
   Verify the time settings on the SSMs running managers.
4. Click Reset Management Group Time.
   A confirmation message opens.
5. Click OK.
   All the times listed on the Times tab should be the same.

9.8 Starting and Stopping Managers

Start or stop managers on SSMs already in a management group.

1. Log in to the management group.
2. Click the Managers tab in the tab view.
   The Managers tab displays, shown in Figure 119.

Figure 119. Starting a Manager

3. Select from the list the SSM on which you want to start a manager.
4. From the Tasks menu, click Start Manager.
5. Click OK.

Note: If you are not logged in to the SSM, you are prompted to log in now.

Adding Manager IP Addresses to Application Servers

When you add a manager to a management group, a window opens, shown in Figure 120, which displays all the IP addresses of the managers in that management group and a reminder to reconfigure the application servers that are affected by the change.

Figure 120. Adding Manager IP Addresses to Application Servers

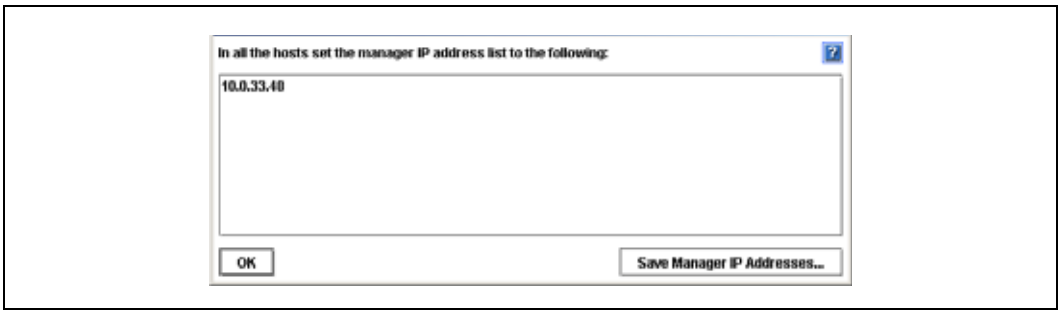

- Click Save Manager IP Addresses to save this list.
  You can print the list and use it as a reference when reconfiguring application servers.

9.8.1 Stopping Managers

Under normal circumstances, you stop a manager when you are removing an SSM from a management group.

Implications of stopping managers

- Quorum may be decreased
- Fewer copies of configuration data are available
- Fault tolerance may be lost
- Data integrity and security may be compromised

Warning: Stopping a manager can result in the loss of fault tolerance. If all managers are stopped, access to the data in that management group is lost.

1. Select the management group in the Network View.
2. Log in to the management group.
3. Select the Managers tab in the Tab View.
   The Managers tab displays.
4. Select the SSM on which you want to stop the manager.
5. Right-click and select Stop Manager.

Note: If you are not logged in to the SSM, you are prompted to log in now.
   A confirmation message opens

6. Click OK to confirm stopping the manager.

Removing Manager IP Addresses from Application Servers

When you remove a manager from a management group, a window opens which displays all the IP addresses of the managers in that management group and a reminder to reconfigure the application servers that are affected by the change.

- Click Save Manager IP Addresses to save this list.

You can print the list and use it as a reference when reconfiguring application servers.

9.9 Removing an SSM from a Management Group

Remove an SSM from an existing management group.

Prerequisites

Stopping or removing the SSM from data storage activities.

- Remove all snapshots and volumes from the cluster containing the SSM. See “Deleting a Snapshot” on page 208 and “Deleting a Volume” on page 186.
- Remove the SSM from any cluster to which it belongs. See “Removing a SSM from a Cluster” on page 164.
- Stop the manager on the SSM, if it is running a manager. See “Stopping Managers” on page 139.

9.9.1 Removing the SSM

1. Log in to the management group from which you want to remove an SSM.
2. Select the SSM to remove.
3. Right-click and select Remove from Management Group.
4. A confirmation message opens.
5. Click OK.
   The SSM is removed from the management group, and moved to Available status.

9.10 Backing Up Management Group Configuration

Use Backup Configuration of Management Group to save one or both of the following on your local machine

- a binary file of the management group configuration from which you can restore the management group, and/or
- 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 SSMs in that management group. To back up SSM configurations, see "Backing Up the Storage System Module Configuration File" on page 29.

1. Log in to the management group.
2. From the Tasks menu on the Details tab, select Back up Configuration of Management Group.
   The Back up Configuration of Management Group window opens, shown in Figure 121.

Figure 121. Backing up the Management Group Configuration

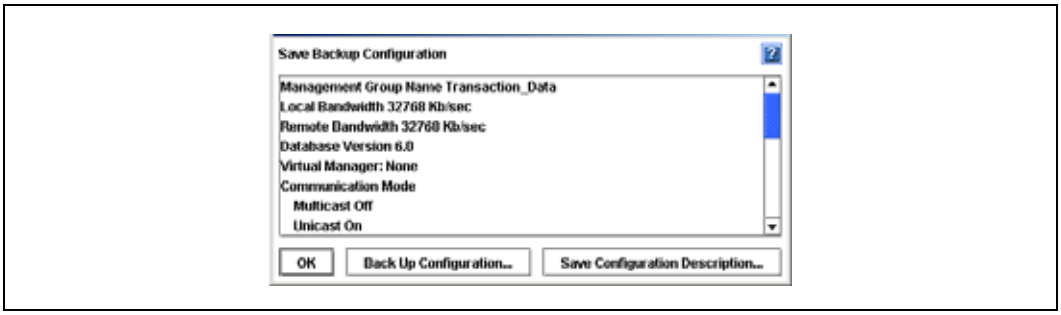

9.10.1 Backing Up Management Group Configuration

1. Click Back Up Configuration.
   The Save window opens, shown in Figure 122.

Figure 122. Save Window for Backing up the Management Group Configuration

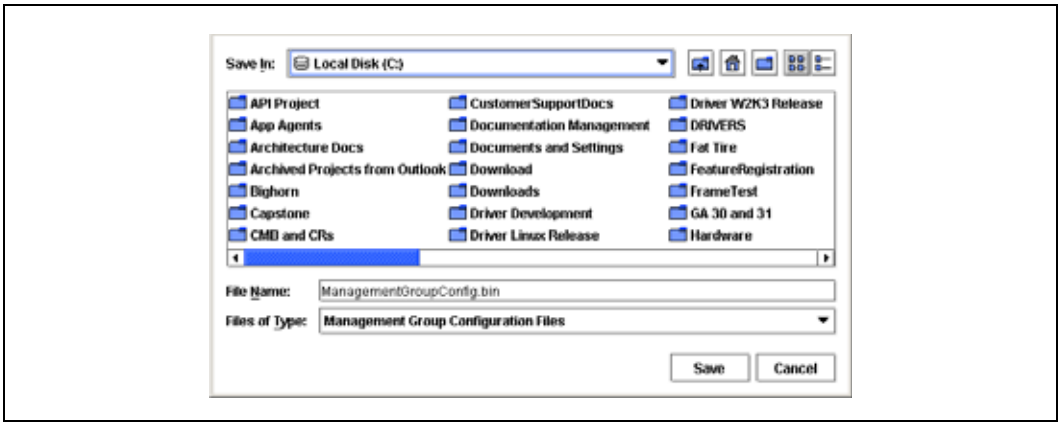

2. Navigate to a folder on the system running the Console in which to store the management group configuration binary file.
3. Use the default name or type a new name for the file.
4. Click Save.
   The configuration file is saved as a binary file in the folder you selected.
5. Click OK to close the Backup Configuration window

9.10.2 Saving the Management Group Configuration Description

1. Click Save Configuration Description.
   The Save window opens.
2. Navigate to a folder on the system running the Console in which to store the management group 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.

9.11 Restoring a Management Group

For disaster recovery, you can use the management group binary file to recreate a management group. The restore procedure restores everything except snapshots, since the data stored in volumes and snapshots is gone.

9.11.1 Requirements for Restoring a Management Group

- **Hardware** - You must have the same number of SSMs available that are the same capacity or greater.
- **IP Addresses** - You must use the same IP addresses for the replacement SSMs that were assigned to the original SSMs. 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 SSMs you are using to restore your management group are in the Available pool in the Console.
2. Right-click in the Network view and select Restore Management Group.
   A standard Open window opens, shown in Figure 123.

Figure 123. Opening the Configuration Binary File

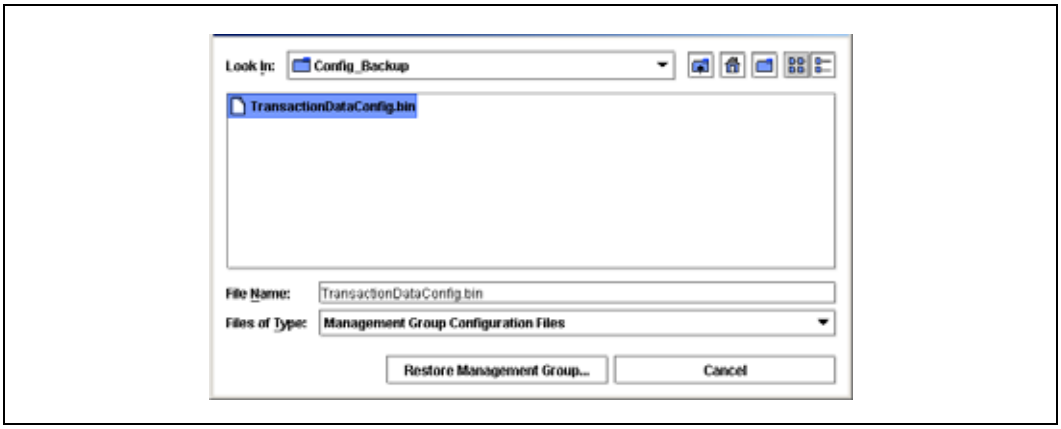

3. Navigate to the location of the configuration binary file.
4. Select the file and click Restore Management Group.
   The Verify Management Group Configuration window opens, shown in Figure 124.

Figure 124. Verifying the Management Group Configuration

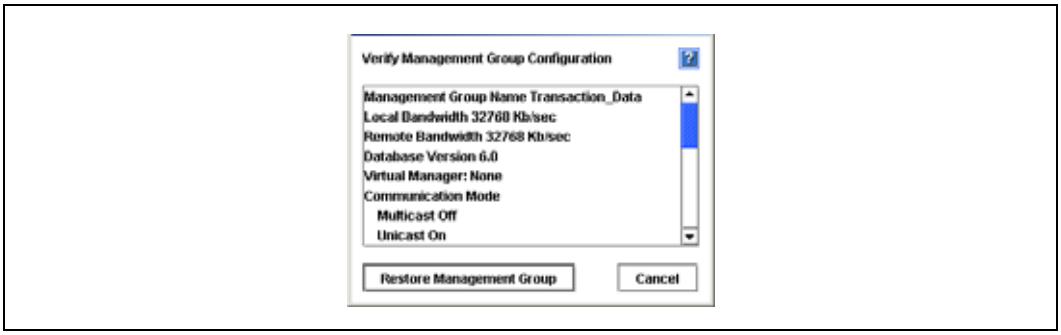

5. After you have reviewed the configuration parameters, click Restore Management Group.

Note: If you are not logged in to the SSMs that will be restored into the management group, you are prompted to log in now. You will also be prompted to log in to a remote management group if it was part of your original configuration.

6. The management group is restored.

9.12 Deleting a Management Group

Delete a management group when you are completely reconfiguring your network storage.

Warning: When a management group is deleted, all data stored on SSMs in that management group are lost.

Prerequisites

- Logging in to each SSM in the management group,
- Stopping the virtual manager and managers on the individual SSMs

9.12.1 Deleting a Management Group

1. Log in to the management group in the Network View.
   The management group tab view opens.
2. From the Tasks menu on the Details tab, select Delete Management Group.
3. A confirmation message opens.

Note: If you are not logged in to all the SSMs in the management group, you are prompted to log in now. If you have not stopped managers on any SSMs, you are prompted to stop them now.

4. Click OK.
5. When the management group is deleted, the SSMs return to available status in the Network View.

Disaster Recovery Using A Virtual Manager 10

10.1 Virtual Manager Overview

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.

See "Managers and Quorum" on page 126 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.

10.1.1 When to Use a Virtual Manager

Use a virtual manager in the following configurations:

- A management group across two sites.
- A management group in a single location with two SSMs.

10.1.1.1 Management Group Across Two Locations With Shared Data

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.

10.1.1.2 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 126 of Chapter 9, "Working with Management Groups."

Running two managers and adding a virtual manager to this management group provides the capability of regaining quorum if one manager goes down.

10.1.2 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.

10.1.2.1 Definitions

Here 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 125
- **Manager** - either of these managers.

Figure 125. Correct Two-site Failure Scenarios Using Virtual Managers

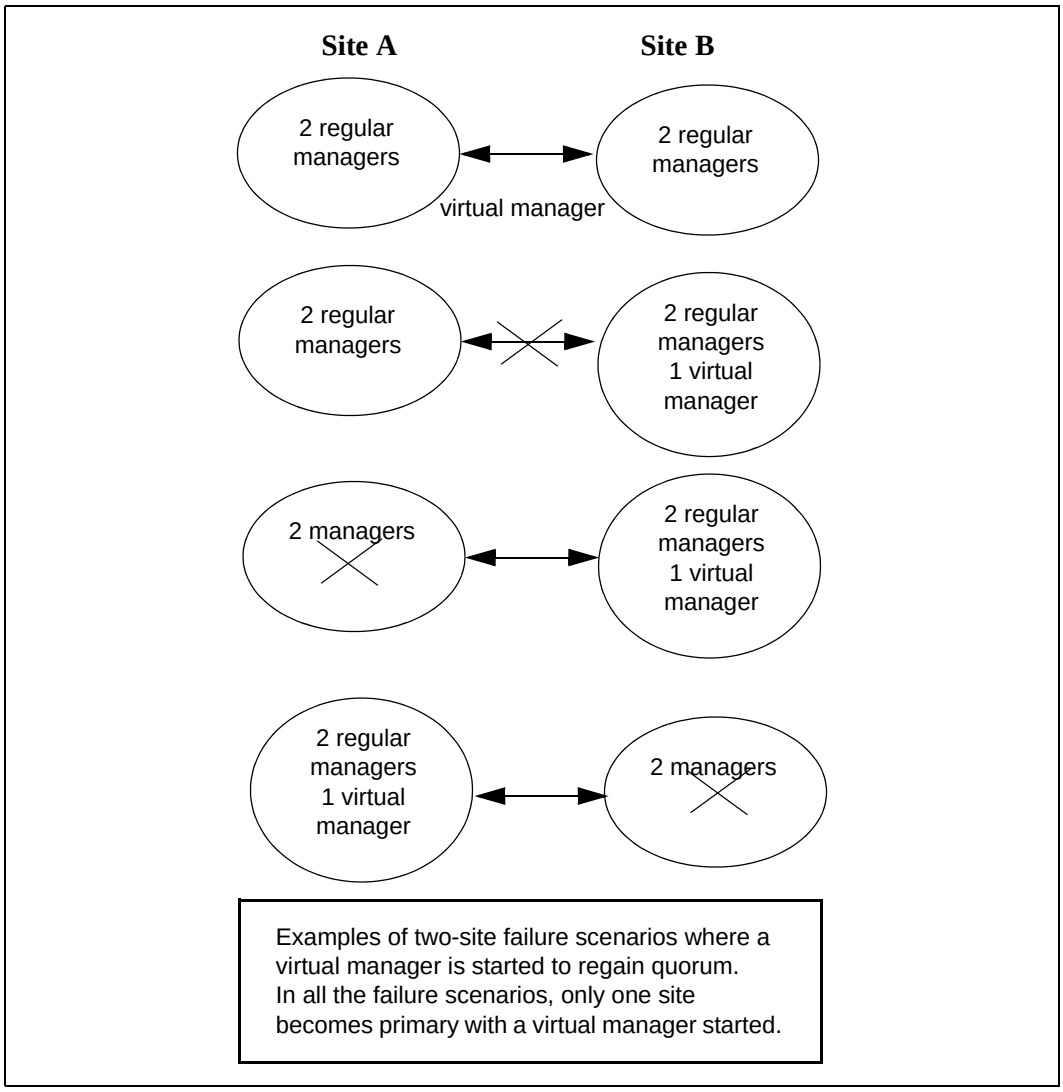

Figure 126. Incorrect Uses of Virtual Manager to Regain Quorum

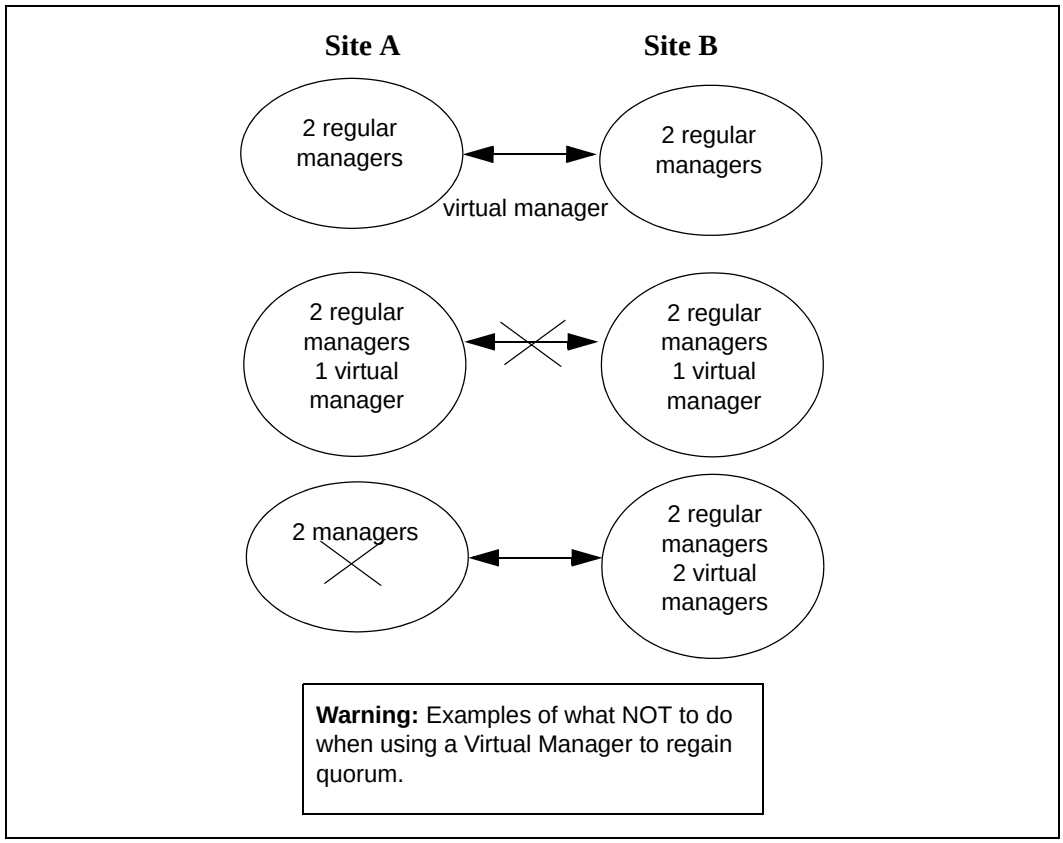

10.1.3 Requirements for Using a Virtual Manager

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.

Table 30. Requirements for Using a Virtual Manager

<table>
<tr><th>Requirement</th><th>What it means</th></tr>
<tr><td>Use a Virtual Manager with an Even Number of Regular Managers Running on SSMs</td><td>
<table>
<tr><th>Disaster Recovery Scenario</th><th># of SSMs Running Regular Managers</th><th>Total # of Managers Including the Virtual Manager</th></tr>
<tr><td>2 sites with shared data</td><td>4</td><td>5</td></tr>
<tr><td>2 SSMs in Management Group</td><td>2</td><td>3</td></tr>
</table>
</td></tr>
<tr><td>Add a Virtual Manager When Creating Management Group</td><td>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.</td></tr>
<tr><td>A Virtual Manager Must Only Be Started Once, and Run Only Until the Site is Restored or Communication is Restored</td><td>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.</td></tr>
</table>

10.2 Configuring a Cluster for Disaster Recovery

In addition to using a virtual manager, you must configure your cluster and volume correctly for disaster recovery. This section describes how to configure your system, including the virtual manager.

10.2.1 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 that spans both sites. That volume must contain all the data in each site.

10.2.2 Configuration Steps

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.

Management Group Name - Transaction_Data

SSM names

- Denver-1
- Boulder-1
- Denver-2
- Boulder-2

For more information, see “Changing the Storage System Module Host Name” on page 25 of Chapter 2, “Working with Storage System Modules.”

Create Management Group - Plan 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.

For more information, see “Managers Overview” on page 125 in Chapter 9, “Working with Management Groups.”

Add SSMs to Cluster in Alternating Order

Create the cluster. When adding SSMs to the cluster, add them in alternating order, as shown in Figure 127. 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.

Cluster Name - Card_Payments

Add SSMs to cluster in the following order

- **1st SSM** - Denver-1
- **2nd SSM** - Boulder-1
- **3rd SSM** - Denver-2
- **4th SSM** - Boulder-2

For more information, see “Creating a Cluster” on page 157 in Chapter 11, “Working with Clusters.”

Warning: If SSMs are added to the cluster in any order other than alternating order by site, you will not have a complete copy of data on each site.

Figure 127. Adding SSMs to Cluster in Alternating Site Order

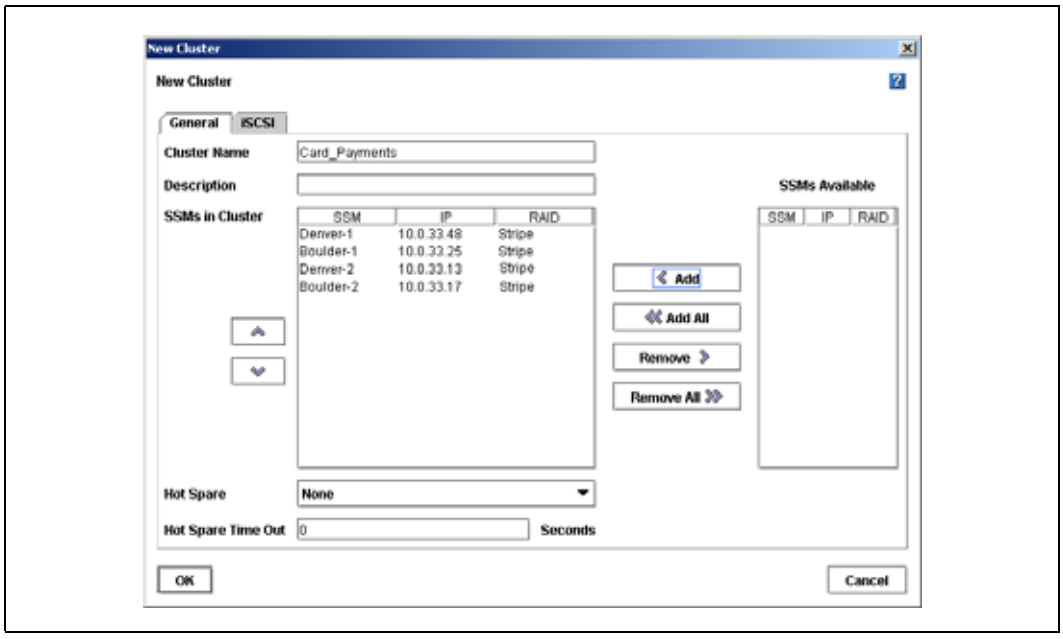

Create the Volume with 2-way Replication

Create the volume on the cluster with 2-way replication. Two way replication ensures that two copies of the data are 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.

For more information, see "Planning Data Replication" on page 173 in Chapter 12, "Working with Volumes."

Figure 128. Cluster with SSMs Added in Alternating Order

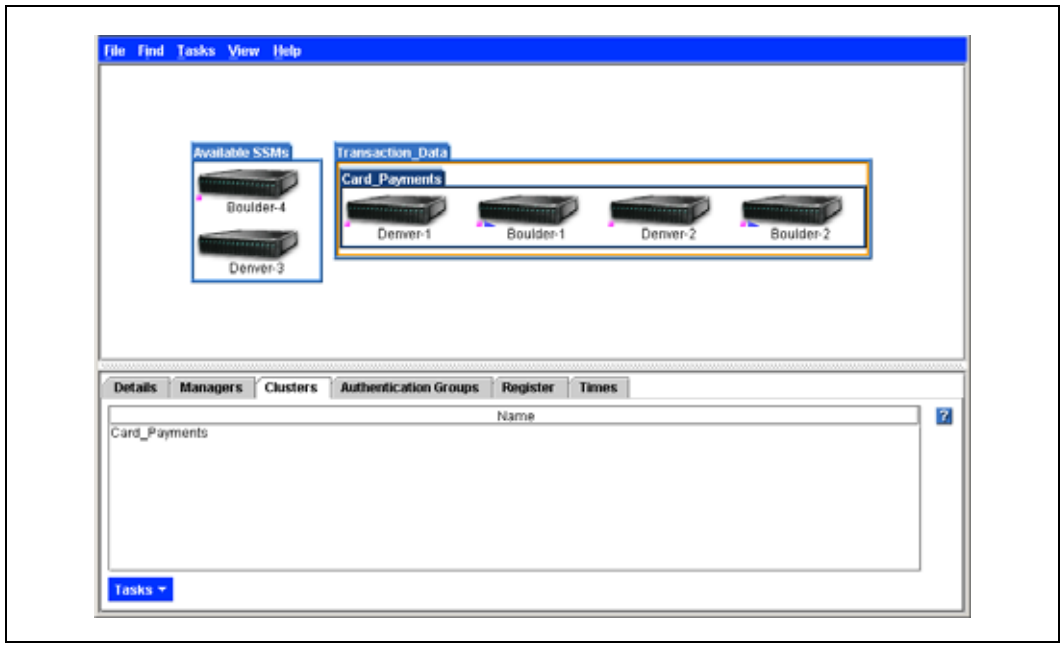

10.3 Configuring a Virtual Manager

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.”

10.3.1 Adding a Virtual Manager

Add a virtual manager to a management group.

1. Right-click on the management group and select Add or Delete Virtual Manager.

   A confirmation dialog opens.

2. 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.

Figure 129. Management Group with Virtual Manager Added

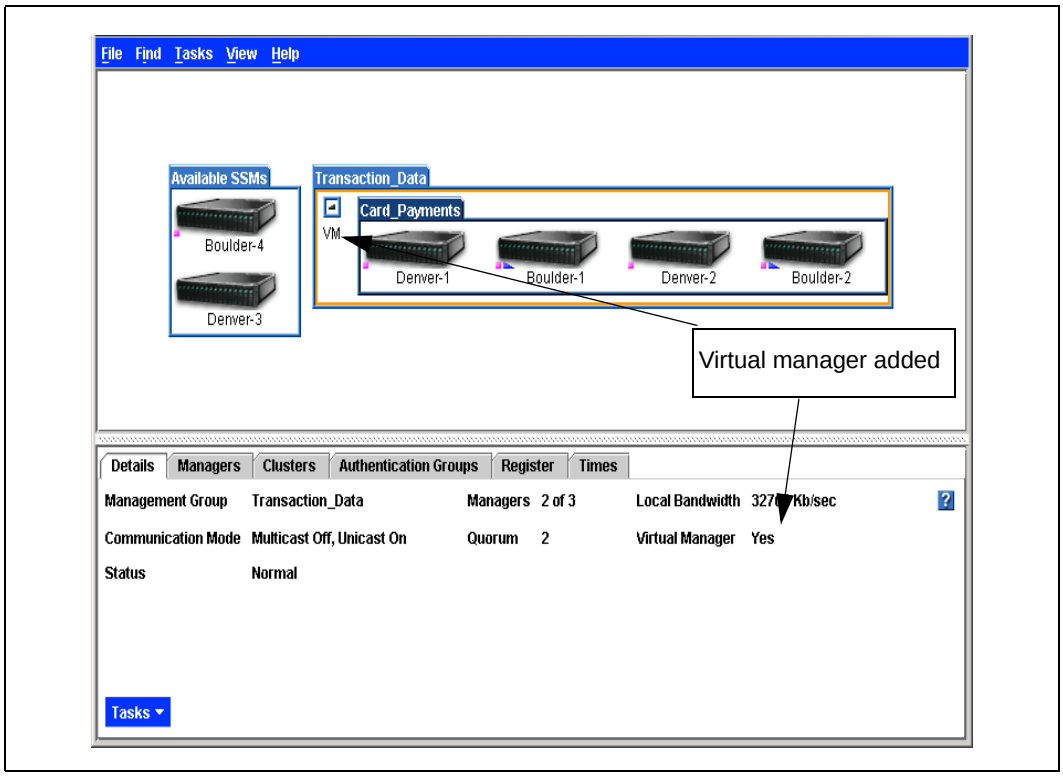

The virtual manager remains added to the management group until needed.

10.4 Starting a Virtual Manager to Regain Quorum

Only start a virtual manager when it is needed to regain quorum in a management group.

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.

10.4.1 Starting a Virtual Manager

A virtual manager must be started on an SSM, ideally one that isn't already running a manager. However, if necessary, you can start a virtual manager on an SSM that is already running a manager. Figure 130 shows a management group with a down manager.

1. Click the SSM on which you want to start the virtual manager.
2. From the Tasks menu on the Details tab, select Start Virtual Manager, shown in Figure 130.

Figure 130. Starting a Virtual Manager

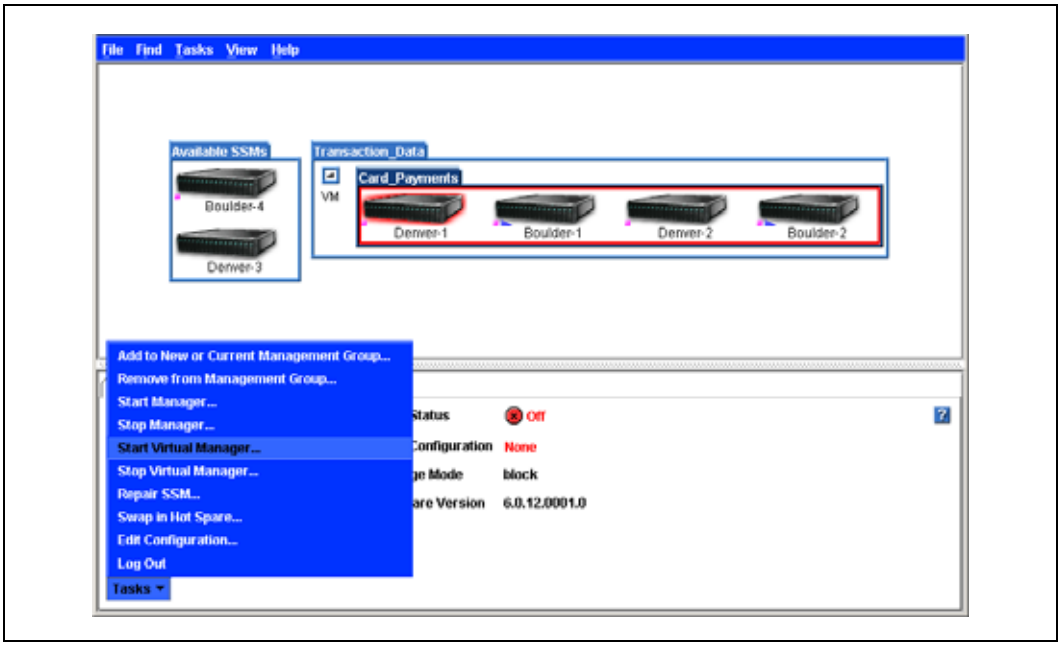

The virtual manager starts on that SSM, and the black triangle—the graphic indicator of the virtual manager—appears under the SSM, shown in Figure 131. See "Icons Used in the Storage Server Console" on page 4 for a key to all the graphic indicators.

Figure 131. Indicator of the Virtual Manager

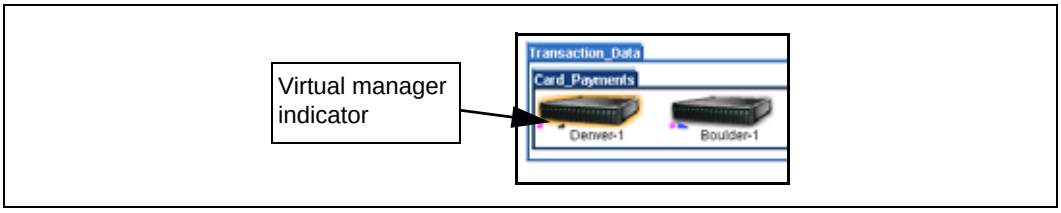

Note: If you attempt to start a virtual manager on an SSM that appears to be up in the Console, and you receive a message that the SSM is down, start the virtual manager on a different SSM. This situation can occur when quorum is lost because the Console displays the SSM in a normal state, even though the SSM is down.

10.5 Stopping a Virtual Manager

When the situation requiring the virtual manager is resolved—either the site recovers or the communication link is restored—you stop the virtual manager. Stopping the virtual manager returns the management group to a fault tolerant configuration.

1. Select the SSM that is running the virtual manager.
2. Click the Management Group Information tab to bring it to the front.
3. Click Stop Virtual Manager.
   A confirmation message appears.
4. Click OK.
   A window listing manager IP addresses opens.
5. Click OK.
   The virtual manager is stopped. However, it remains part of the management group and part of the quorum.

10.5.1 Removing a Virtual Manager

You can remove the virtual manager from the management group altogether.

1. Select the management group that has the virtual manager.
2. From the Tasks menu on the Details tab, select Add or Delete Virtual Manager.
   A confirmation message 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.

Working with Clusters 11

11.1 Clusters Overview

Within a management group you create sub-groups of SSMs called clusters. A cluster is a grouping of SSMs from which you create volumes.

Think of a cluster as a pool of storage. You add storage to the pool by adding SSMs. You then carve volumes out of the pool. Volumes seamlessly span the SSMs in the cluster.

11.1.1 Mixing SSMs of Different Capacities in Clusters

Clusters can contain SSMs with different capacities. However, all SSMs in a cluster will operate at a capacity equal to that of the smallest capacity SSM.

Prerequisites

All the SSMs in a cluster must be configured alike.

Before you create a cluster, you must have created a management group.

11.1.2 Hot Spares Overview

A cluster of SSMs can contain a hot spare SSM. A hot spare is an SSM that is not used for data storage, but stands by in case an SSM in the cluster goes down. A hot spare SSM is designated in the Console by the icon show below.

Figure 132. Hot Spare SSM Icon

In order to have more than one SSM in a cluster beyond the 30-day evaluation period, you must purchase the Scalability Pak. See

11.1.2.1 Requirements for Hot Spares

To designate a hot spare SSM for a cluster, the following requirements apply.

Table 31. Hot Spare Requirements

Hot Spare Requirements
A cluster must contain at least 3 SSMs to have one SSM designated as a hot spare.
At most, one hot spare SSM can be designated per cluster. However, a cluster does not require a hot spare.
The hot spare SSM must be equal to or greater in size than the other SSMs in the cluster.

11.1.2.2 Using Hot Spares

If an SSM in a cluster goes down, and a hot spare is designated for that cluster, then the spare is automatically activated and the data starts to migrate to the new SSM. At this point the cluster no longer contains a hot spare. When the down SSM comes back up, it becomes the hot spare.

When a hot spare is activated, it is not configured as a manager. If you want to designate the activated hot spare as a manager, you must start the manager.

11.1.2.3 Setting the Hot Spare Time Out

The hot spare time out designates the amount of time before a hot spare is activated in the cluster. When a hot spare is activated the system will migrate data onto the new SSM. This data migration may take some time. Setting the hot spare time out allows you to control for situations in which you don't want the hot spare activated, for example, if your network has high latency.

The time out begins counting from the time that the SSM begins blinking in the Console. The default time is set to 0 seconds so that the hot spare takes over as soon as the system detects that the SSM is unavailable.

For example, if you set the time out to 60 seconds, then the hot spare is activated 1 minute after the system detects that the SSM is unavailable.

11.1.2.4 Swap in Hot Spare

You can manually force a SSM designated as a hot spare to activate in the cluster, if an SSM in that cluster is not available and the cluster is blinking red in the Console. Swapping in a hot spare overrides the hot spare time out setting. However, the setting remains intact in the cluster and continues to apply once the cluster configuration has returned to normal.

11.1.3 Clusters and iSCSI

If you plan to use iSCSI with the Storage System Engine, 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.

- iSCSI Failover - If you are using an initiator that does not support multiple addresses per target, such as the Microsoft* iSCSI Initiator, to ensure iSCSI failover you must configure a virtual IP for the SSMs in a cluster.
- iSNS Server - If you use an iSNS server, configure your cluster to register the iSCSI target with the iSNS server.

11.1.3.1 iSCSI Failover and Virtual IP

A virtual IP address ensures that if a SSM in a cluster becomes unavailable, clients using an initiator that does not support multiple addresses per target, such as the Microsoft iSCSI Initiator, can still access the volume through the other SSMs in the cluster. If the initiator you are using does support multiple addresses, you may not want to use a virtual IP.

Table 32. Requirements for a Virtual IP

Requirements for a Virtual IP
SSMs must be in same subnet address range as the virtual IP
The virtual IP must be routable regardless of which SSM it is assigned to.
Microsoft iSCSI clients must be able to ping the virtual IP
Must be unique to all SSMs on the network.
Must be a specific IP reserved for this purpose. If you use DHCP, you must use a static IP.
All Microsoft iSCSI initiators must be configured to connect to this IP for failover.

11.1.3.2 Using an iSNS Server

An iSNS server simplifies the discovery of iSCSI targets on multiple clusters on a network. You can have up to 3 iSNS servers.

11.2 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 a two clusters, each with one SSM, the most reliable configuration is to create two management groups with one SSM in each group.

1. Log in to the management group for which you want to create a cluster.
   The management group tab view opens.
2. Click the Clusters tab.
   The Clusters tab opens.

Figure 133. Viewing the Clusters Tab

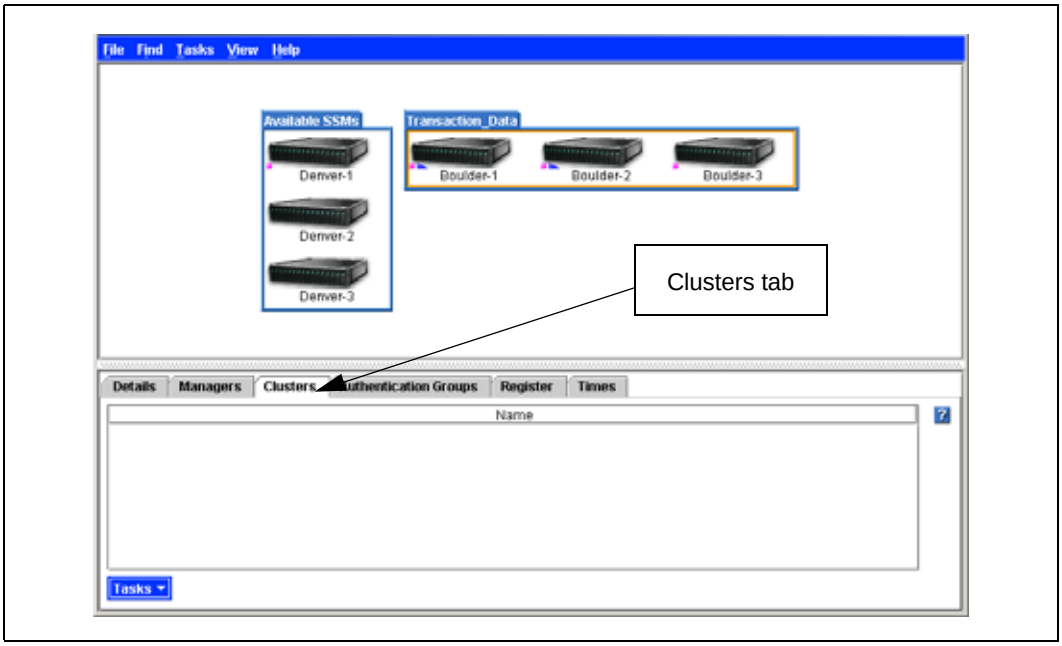

3. From the Tasks menu, click New Cluster.
   The New Cluster window opens with the General tab on top.

Figure 134. Creating a New Cluster

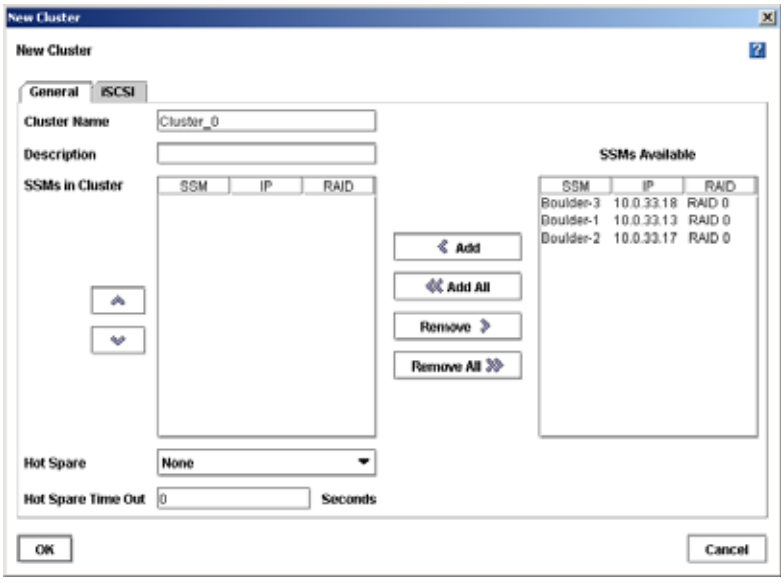

4. Type a meaningful name for the cluster.
   A cluster name is case sensitive and must be from 1 to 127 characters.
5. [Optional] Type a description of the cluster.

6. Select one or more SSMs from the SSMs Available list.

Note: The SSMs in the list are all those included in the management group that are not already in a cluster.

7. Click Add.

   The selected SSMs move to the SSMs in Cluster list.

   or

   Click Add All to move all the SSMs from the Available list to the SSMs in Cluster list.

11.2.1 Designating a Hot Spare

You must purchase the Scalability Pak to use the hot spare feature beyond the 30-day evaluation period.

8. [Optional] Click the Hot Spare drop down list to designate a hot spare.

   Only SSMs in the cluster are displayed in the Hot Spare list. Hot spares cannot be used for storage—that is, you cannot create volumes on them.

9. [Optional] If you designate a hot spare you can set the hot spare time out.

11.2.2 Configure Virtual IP and iSNS for iSCSI

[Optional] To configure iSCSI failover for initiators that do not support multiple addresses per target, such as the Microsoft Initiator, add a virtual IP for the cluster.

10. Click the iSCSI tab to bring it to the front.

Figure 135. Configuring a Virtual IP for iSCSI

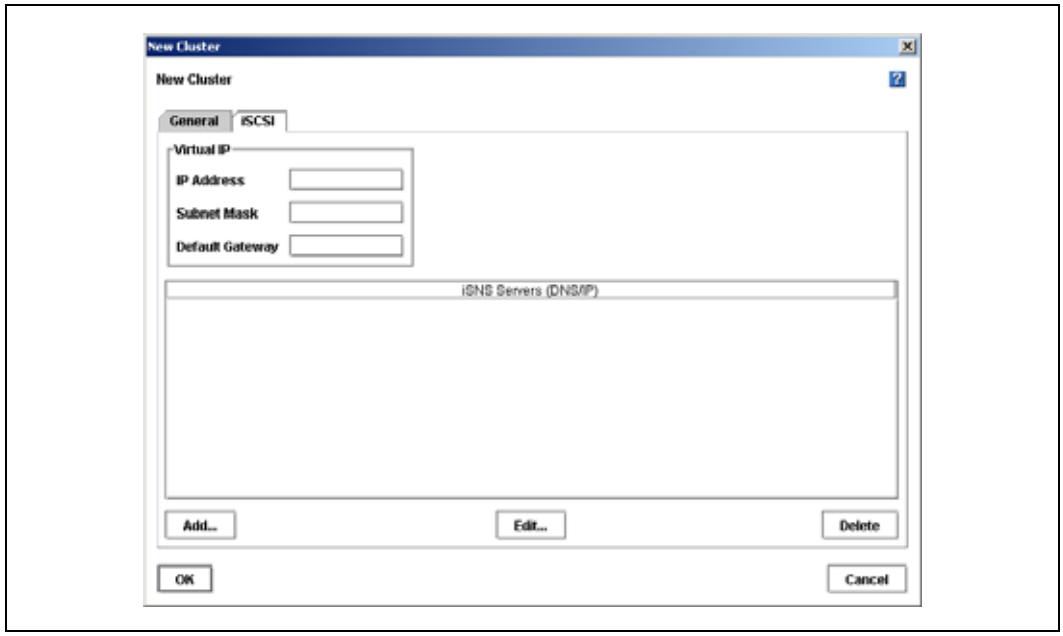

11. Add the IP address, subnet mask and default gateway if required.

11.2.3 Adding an iSNS Server

[Optional] Add an iSNS server.

Note: If you use an iSNS server, do not add Target Portals in the Microsoft iSCSI Initiator.

12. Click Add.
    The Add iSNS Server window opens.

Figure 136. Adding an iSNS Server

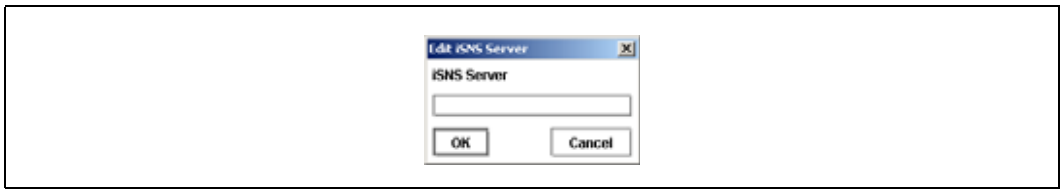

13. Type the IP address of the iSNS server.
14. Click OK.
    The server is added to the list.

Figure 137. Viewing the List of iSNS Servers

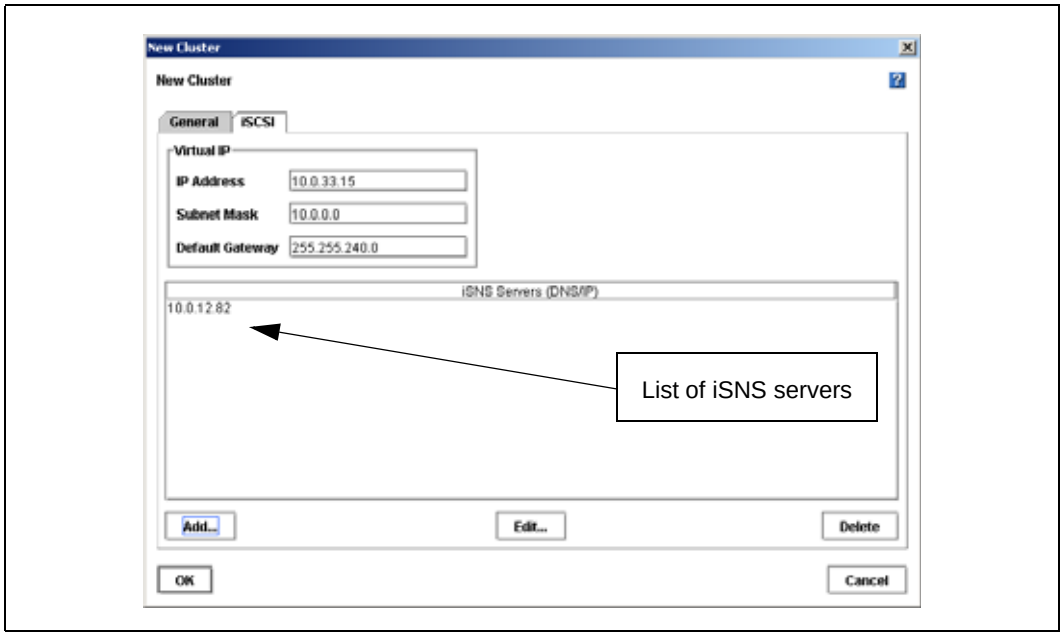

15. Click OK when you have finished.
    The cluster is created and displayed inside the management group.
16. Select the cluster to open the clusters tab view.

Figure 138. Viewing a Cluster and the Cluster Tab View

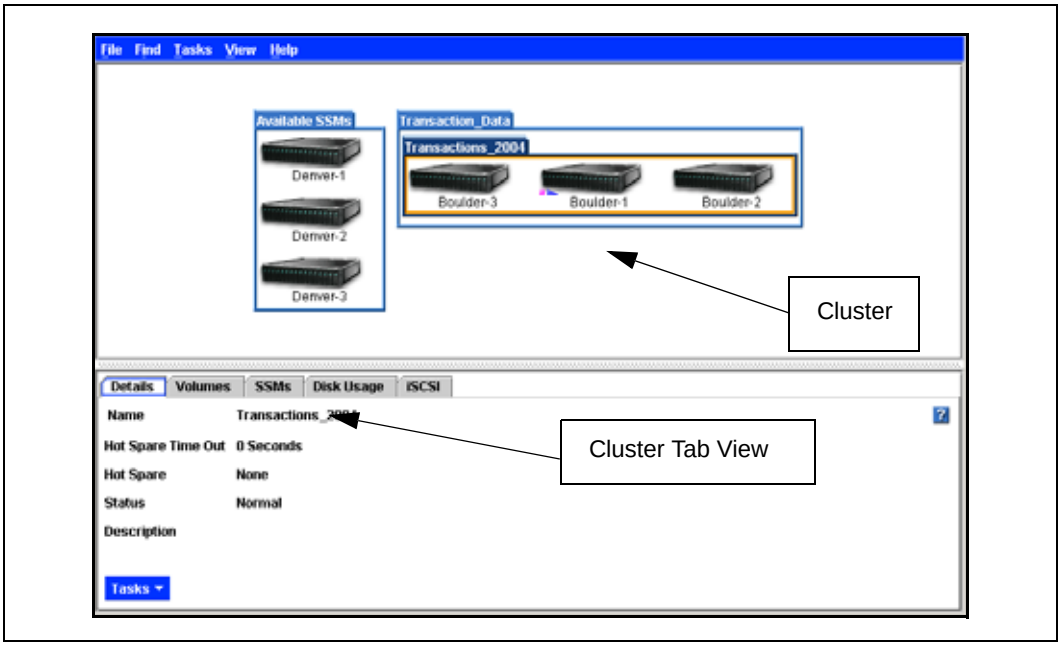

11.2.4 The Cluster Tab View

The tabs provide access to cluster information and features, such as editing or deleting the cluster, creating volumes, and monitoring disk usage for the cluster.

11.2.4.1 Details Tab

Includes name, status, description, and the buttons to edit or delete the cluster.

11.2.4.2 Volumes Tab

Includes the name, replication level, size, hard threshold and soft threshold, and the buttons to create, edit, and delete a volume.

11.2.4.3 SSMs Tab

For each SSM in the cluster, the SSMs tab includes the host name, IP address, and whether the module is a hot spare for the cluster.

11.2.4.4 Disk Usage Tab

Displays usage statistics for the cluster and the modules, volumes, and snapshots contained in the cluster. The usage table provides the following information:

For the cluster - Total space in the cluster, percent allocated to volumes and snapshots.

For each SSM in the cluster - Total space on the SSM, the percent of SSM space to which data has been written, and whether the SSM is a hot spare.

For each volume in the cluster - The volume hard threshold and percent of hard threshold space that data has been written to, and the replication level.

For each snapshot of a volume - The snapshot hard threshold and percent of hard threshold space that data has been written to, and the replication level.

11.2.4.5 iSCSI Tab

Displays the virtual IP address if there is one, and lists any iSNS servers configured for the cluster.

11.3 Editing a Cluster

When editing a cluster, you can change the description, add or remove SSMs, and change the hot spare designation of an SSM. 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.

11.3.1 Getting There

1. Select the cluster you want to edit.
2. From the Tasks menu on the Details tab, select Edit Cluster.
   The Edit Cluster window opens.

Figure 139. Editing a Cluster

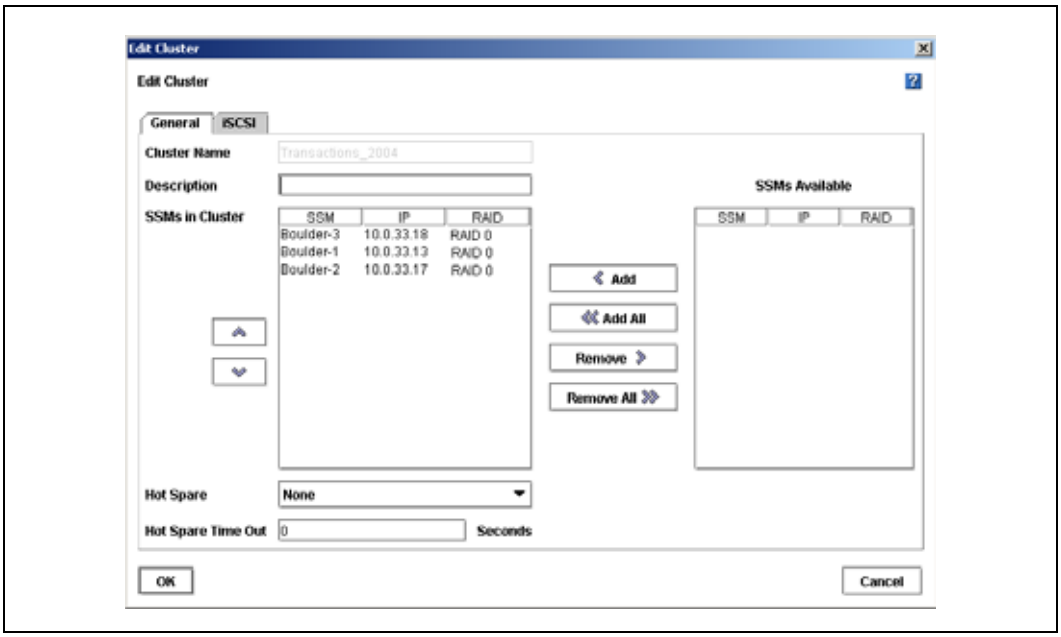

11.3.2 Adding a SSM to an Existing Cluster

Add a SSM to an exiting cluster to expand the storage for that cluster or to designate the SSM as a hot spare.

Prerequisites

- Configure the SSM to match the SSMs already in the cluster.
- Add the SSM to the management group that contains the cluster.

Note: If you mix SSMs with different capacities in a cluster, all SSMs in the cluster will operate at a capacity equal to that of the smallest capacity SSM.

1. Select an SSM from the SSMs Available list.
2. Click Add.
   The SSM moves to the SSMs in Cluster list.
   or
   Click Add All to move all the SSMs from the SSMs Available list to the SSMs in Cluster list.
3. Click OK when you are finished.

11.3.3 Changing the Hot Spare Designation

You can add a hot spare or remove a hot spare from a cluster as long as the volumes in that cluster have a replication priority of availability.

Note: A hot spare cannot reside in a cluster which contains volumes that have a replication priority of redundancy.

11.3.3.1 Adding a Hot Spare

To add a hot spare, the cluster must contain sufficient SSMs to handle the volumes and snapshots that currently exist in that cluster.

1. Click the Hot Spare drop down list and select the SSM to designate as the hot spare.
2. Click OK when you are finished.

Note: A cluster must contain at least 3 SSMs to have one SSM designated as a hot spare.

11.3.3.2 Removing a Hot Spare

To remove a hot spare, simply change the designation in the list to none. The hot spare then becomes an SSM in the cluster, adding more space for storage.

1. Click the Hot Spare drop down list and select None from the list.
2. Click OK when you are finished.
   The SSM returns to the cluster as available storage.

11.3.4 Changing the Hot Spare Time Out

The hot spare time out designates the amount of time before a hot spare is activated in the cluster. You can change the value at any time. The results of changing the time out value are listed below.

- Cluster operating normally – changing hot spare time out has an effect only if an SSM in the cluster becomes unavailable.
- Cluster with unavailable SSM – reducing the time out value will activate the hot spare earlier.
  For example, an SSM is not available and the hot spare time out is configured for 6 hours. After 3 hours you reduce the time out to 1 hour, thinking that will activate the hot spare in 60 minutes. However, the hot spare activates immediately. This is because the clock that is tracking the time out started when the SSM became unavailable and it considers the 1 hour interval to have passed already.
- Cluster with unavailable SSM – increasing the time out value will activate the hot spare later.
  For example, an SSM is not available and the hot spare time out is configured for 6 hours. Four hours have passed. You increase the time out to 8 hours, adding an additional 2 hours to the time out interval, before the hot spare activates

To change the hot spare time out

1. Change the value for the hot spare time out.
2. Click OK.

11.3.5 Removing a SSM from a Cluster

You can remove an SSM from a cluster only if the cluster contains sufficient modules to maintain the existing volumes and replication level.

1. Select an SSM from the SSMs in Cluster list.
2. Click Remove.
   The SSM moves to the SSMs Available list.
3. Click OK when you are finished.

11.3.6 Changing or Removing the Virtual IP

Anytime you change or remove the virtual IP address you are changing the configuration that clients are using. Therefore it is important to disconnect any clients before making this change.

11.3.6.1 Preparing Clients

- Quiesce any applications that are accessing volumes in the cluster.
- Log off the active sessions in the initiator for those volumes.

11.3.6.2 Changing or Removing the Virtual IP Address

1. In the Edit Cluster window, click the iSCSI tab.
2. Change or delete the entries in the IP Address, Subnet Mask and Default Gateway fields.

11.3.6.3 Finishing Up

1. Click OK when you are finished changing or removing the virtual IP.
2. Reconfigure the iSCSI initiator with the changes.
3. Reconnect to the volumes.
4. Restart the applications that use the volumes.

11.3.7 Changing or Removing an iSNS Server

If you change the IP 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.

11.3.7.1 Preparing Clients

- Quiesce any applications that are accessing volumes in the cluster.
- Log off the active sessions in the initiator for those volumes.

11.3.7.2 Changing an iSNS Server

1. Select the iSNS server to change.
2. Click Edit.
   The Edit iSNS Server window opens.
3. Change the IP address.

4. Click OK.

11.3.7.3 Deleting an iSNS Server

1. Select the iSNS server to delete.
2. Click Delete
   A confirmation message opens.
3. Click OK.

11.3.7.4 Finishing Up

1. Click OK when you are finished changing or removing an iSNS server.
2. Reconfigure the iSCSI initiator with the changes.
3. Reconnect to the volumes.
4. Restart the applications that use the volumes.

11.4 Swapping in a Hot Spare

You can manually swap in a hot spare if an SSM in the cluster is not available and is blinking red in the Console.

1. Right-click on the SSM designated as the hot spare for the cluster.
   The hot spare SSM has the hot spare icon next to it. 
   A confirmation message opens.
2. Click OK.
   The SSM begins the process of data migration.

11.5 Repairing a SSM

Repairing an SSM allows you to replace a failed disk in an SSM that is in a cluster configured for 2-way or 3-way replication and only trigger one resync of the data stored on SSMs in that cluster, rather than restriping. Resyncing the data is a shorter operation than a restripe.

11.5.1 Prerequisites for Using Repair SSM

- Volume must have 2-way or 3-way replication.
- SSM must be blinking red in the Console.
- If the SSM is running a manager, stopping that manager must not break quorum.

11.5.2 How Repair SSM Works

Replacing a failed disk requires removing the SSM from the cluster and management group, replacing the disk, and returning the SSM to the cluster. Because of the replication level, removing and returning the SSM to the cluster would normally cause the remaining SSMs in the cluster to restripe the data twice—once when the SSM is removed from the cluster and once when it is returned. Repairing the SSM creates a placeholder in the cluster, in the form of a "ghost" SSM. This ghost SSM keeps the cluster intact while you remove the SSM, replace the disk, configure RAID, and return the SSM to the cluster. The returned SSM only has to resynchronize with the other 2 SSMs in the cluster.

11.5.2.1 Repairing a SSM

When an SSM in a cluster has a disk failure, the Network View displays the SSM and the cluster as blinking red and needing attention.

Figure 140. SSM with Failed Disk

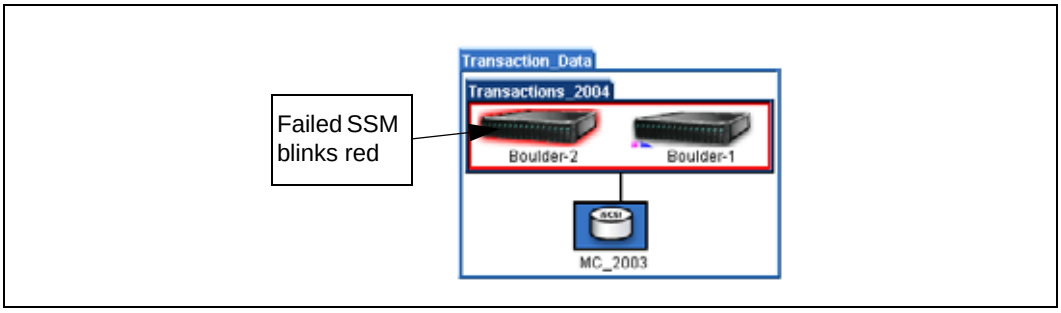

1. If the SSM is running a manager, stop the manager.
2. Select the SSM in the Network View.
3. Right-click and select Repair SSM.
   A confirmation message opens.
4. Click OK.
   The SSM leaves the management group and moves to the Available group. A placeholder, or "ghost" SSM remains in the cluster. It is labeled with an IP address instead of a host name.

Figure 141. Viewing the Ghost SSM

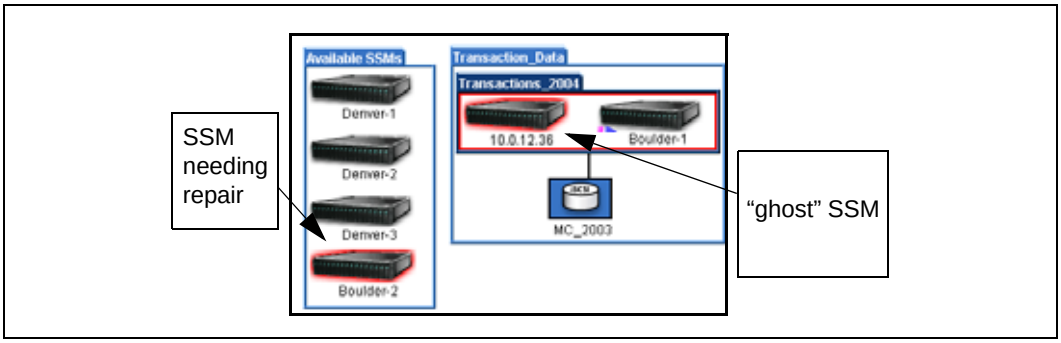

5. Replace the disk in the SSM.
   After you replace the disk you must power the disk on and reconfigure RAID.

6. Add the repaired SSM to the management group.
   The SSM returns to the management group and the ghost SSM is still in the cluster.

Figure 142. Returning the SSM to the Management Group

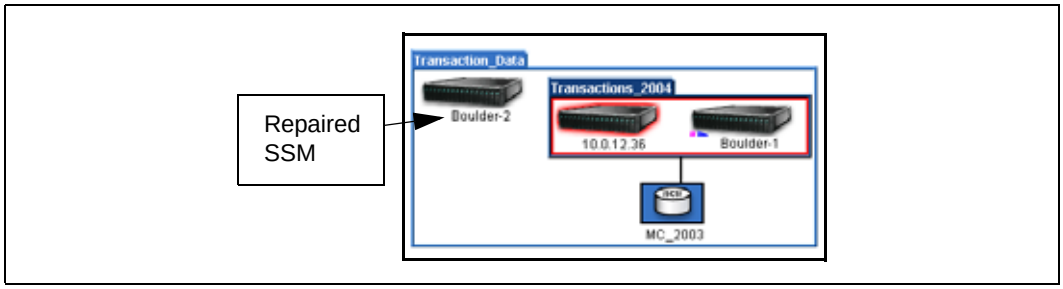

7. Edit the cluster and add the repaired SSM to the cluster.

Note: The repaired SSM 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 SSM to the cluster in the original order

8. In the Edit Cluster window remove any SSMs in the list that are **below** the ghost SSM.
   The removed SSMs return to the SSMs Available column.

Figure 143. Returning the Repaired SSM to the Cluster

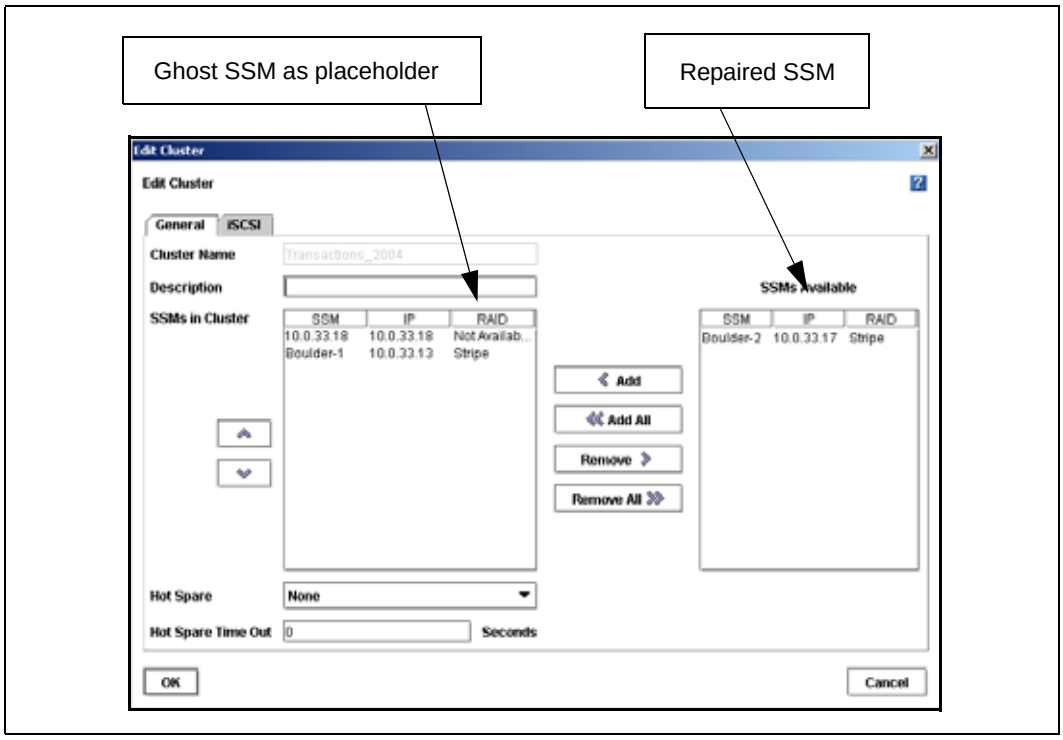

9. Remove the ghost SSM.
10. Select the repaired SSM and add it to the cluster.
    It will be in the place reserved by the ghost SSM.

11. Add any remaining SSMs to the cluster.
12. Click OK.
    The SSMs are in the cluster in their original order. The ghost SSM is removed from the cluster.
13. Select the ghost SSM and remove it from the management group.
    A confirmation message opens, warning that the SSM cannot be found on the network.
14. Click OK to confirm removing SSM from the management group.
    Another confirmation message opens.
15. ClickOK.
    The ghost SSM disappears from the Console.

11.6 Deleting a Cluster

Volumes and snapshots must be deleted from a cluster before you can delete the cluster.

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. Select the cluster you want to delete.
   The Cluster tab view opens.
3. From the Tasks menu on the Details tab, 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 SSMs return to the management group as available.

Working with Volumes 12

12.1 Volume Overview

A volume is a logical entity that is made up of storage on one or more SSMs. 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 SSMs.

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,” on page 125 and Chapter 11, “Working with Clusters,” on page 155.

12.1.1 Topics Covered in This Chapter

- Planning volume size and thresholds
- Planning data replication and data priority
- Choosing access modes for volumes
- Using iSCSI and volumes
- Creating and managing volumes

12.2 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?
- What applications will access the volume?
- How will those applications access the volume - EBSD driver, iSCSI initiator, or Fibre Channel host?
- Do you plan to grow the volume or to keep it the same size?
- What clients are going to access the volume and how will you configure permissions for those clients?

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.

12.2.1 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. See the Remote Copy User's Manual for detailed information about remote volumes.

12.2.2 Planning Volume Size

Volume size is the size of the virtual device communicated to the operating system and the applications. Volume size falls into one of three categories

- Volumes that are smaller than the storage capacity of the cluster
- Volumes that are equal in size to the storage capacity of the cluster
- Volumes that are larger than the storage capacity. Creating larger volumes makes it easy to add additional storage resources to the cluster at a later date.

How you plan to use the volume is one factor in setting the size. Other factors in planning size are calculating the hard threshold and whether you plan to use snapshots.

12.2.3 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 changed when using snapshots.

12.2.3.1 Best Practice if Not Using Snapshots

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 189.

12.2.3.2 Best Practice if Using Snapshots

For volumes that will be used with snapshots, set the hard threshold size less than the volume size. Next, set the soft threshold less than the hard threshold.

12.2.4 Planning Snapshots

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.

Note: Volume size, volume thresholds, and using snapshots should be planned in conjunction. If you intend to use snapshots, review Chapter 13, "Working with Snapshots."

12.2.5 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 178 for strategies to manage volume growth.

12.2.5.1 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.

12.2.5.2 Best Practice If Using Snapshots

If the hard threshold is less than the volume size, set the soft threshold to a percentage of the hard threshold. When a soft threshold alert is received,

- provision more storage for the cluster (if required),
- increase the hard threshold, and
- re-adjust the soft threshold to be a percentage of the new hard threshold.

12.2.6 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 SSMs, replication levels are tied to the number of available SSMs in a cluster. (Hot spare SSMs are not available for data storage, and therefore not available when calculating replication levels.)

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.

12.2.6.1 Replication Level

Three replication levels are available depending upon the number of available (non-hot spare) SSMs in the cluster. The level of replication you choose also affects the Replication Priority you can set.

Table 33. Setting a Replication Level for a Volume

With This Number of Available SSMs 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 (not a recommended configuration for high availability)	• None • 2-Way	• 1 copy of data in the cluster, no replication. • 2 copies of data in the cluster. One replica is created.
3 or more	• None • 2-Way • 3-Way	• 1 copy of data in the cluster (no replication). • 2 copies of data in the cluster (one replica). • 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.

12.2.6.2 How Replication Levels Work

When you choose 2-way or 3-way replication, data is written to either 2 or 3 consecutive SSMs in the cluster. For example:

2-Way Replication

A cluster with 3 SSMs, configured for 2-way replication. There have been 5 writes to the cluster. Figure 144 illustrates the write patterns on the 3 SSMs.

Figure 144. Write Patterns in 2-Way Replication

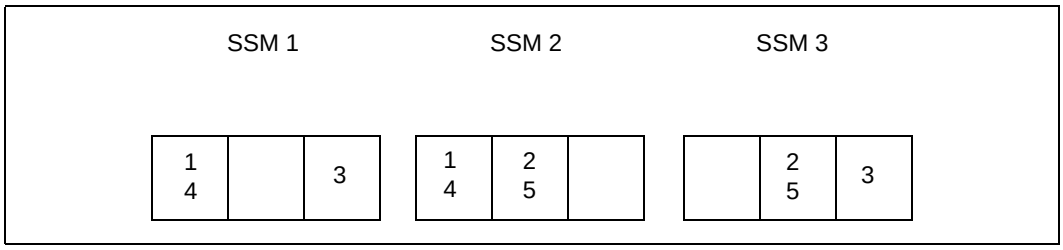

12.2.6.3 Replication Priority

Set the replication priority according to whether data availability or data redundancy is the goal for the volume.

Availability Priority

If data availability is your priority, you can set any replication level.

Redundancy Priority

If redundancy is the priority, you must select either 2-way or 3-way replication.

Note: When you have volumes with a priority of redundancy, you cannot use a hot spare in the cluster.

Table 34. Replication Levels, Priority Settings, and Volume Availability

Volume is available to a client with a replication level of:			
And a priority setting of:	**None**	**2-way**	**3-way**
Availability	All SSMs must be up	1 of every 2 consecutive SSMs must be up	1 of every 3 consecutive SSMs must be up
Redundancy	N/A	All SSMs must be up	2 of every 3 consecutive SSMs must be up

Warning: A management group with 3 SSMs is the minimum configuration for fault tolerant operation. Although the system allows you to configure 2-way replication on 2 SSMs, this does not guarantee data availability in the event that 1 SSM becomes unavailable. See "Managers Overview" on page 125.

12.2.6.4 Best Practice

If your volumes contain critical data, configure them for 2-way replication and a priority of redundancy.

12.2.7 Planning Access to Volumes

Will client applications access volumes using the EBSD driver, an iSCSI initiator, or a Fibre Channel host? All access modes are supported such that

- You can toggle between access modes for both volumes and snapshots. All data in the volume or snapshot is retained when switching between modes.
- iSCSI, EBSD and Fibre Channel volumes and snapshots can reside in a single cluster.

12.2.8 Planning Volumes and iSCSI

Volumes can be configured for iSCSI. If you configure volumes for iSCSI, you can also configure CHAP for those volumes.

CHAP (Challenge-Handshake Authentication Protocol) is a standard authentication protocol. Storage System Software supports no CHAP, 1-way CHAP, or 2-way CHAP.

- No CHAP - authorized clients can access the volume without logging in.
- 1-way CHAP - clients must log in with a target secret to access the volume. 1-way CHAP is set at the volume level.
- 2-way CHAP - clients must log in with a target secret, and volume identity must be confirmed for the client. 2-way CHAP is set at the authentication group level. See "Authentication Groups and iSCSI" on page 215.

CHAP is optional. However, if you configure CHAP parameters for the volume (1-way CHAP) or for the authentication group (2-way CHAP), you must remember to configure the Storage System Software authentication group and iSCSI initiator with the appropriate parameters. Table 35 lists the requirements for configuring CHAP.

12.2.8.1 Requirements for Configuring CHAP

Table 35. Entering Information to Configure iSCSI CHAP

CHAP Level	What to Configure in the Console	What to Configure in the iSCSI Initiator
No CHAP	• Volume - target type iSCSI. • Authentication Group - Initiator Name only if required for selected level of access	• No configuration requirements
1-way CHAP	• Volume - target type iSCSI, create a target secret.	• Enter the target secret when logging on to available target.
2-way CHAP	• Authentication Group - enter Initiator Name (assigned by the Initiator) and Initiator Secret (make up your own). • Volume - select iSCSI as target type, create a target secret.	• Create the initiator secret. • Enter the target secret when logging on to the available target

Note: A volume must have an authentication group associated with it in order to mount it on a server.

12.2.9 Planning Volumes and Fibre Channel

Storage volumes can also be configured for Fibre Channel access. Designate the volume as Fibre Channel when you create it. Then configure an authentication group to associate a logical unit number (LUN) with the volume. "Creating an Authentication Group" on page 217.

12.3 Requirements for Volumes

When creating a volume, you define the following parameters.

Table 36. Parameters for Volumes

Volume Parameter	Configurable in Volume Type	What it Means
Type	Any	Whether the volume is primary or remote. • Primary volumes are used for data storage. • Remote volumes are used for configuring Remote Copy for business continuance, backup and recovery, or data mining/ migration. **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 Storage System 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.

Table 36. Parameters for Volumes (Continued)

Volume Parameter	Configurable in Volume Type	What it Means
Replication Level	Any	The number of copies of the data to create on SSMs in the cluster. The replication level must be at most the number of SSMs in the cluster or 3, whichever is smaller. See “Planning Data Replication” on page 173.
Replication Priority	Any	• Availability - Default setting. These volumes will remain available as long as at least one SSM out of every n (n = replication level) remains active. When the unavailable SSM returns to active status in the cluster, than the volume resynchronizes across the replicas. • 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 an SSM in the cluster becomes unavailable, thereby preventing 2-way replication, the volume goes offline until the SSM is again available.
Size	Primary	The logical 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 SSMs 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. The default value in the size field is equal to the available space on the cluster.
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 must be less than or equal to the volume size. Remote volumes contain no data and do not have a size. Therefore, 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. Because remote volumes have no size, and cannot have a hard threshold, they also cannot have a soft threshold.
Checksum	Any	Whether to use checksumming to verify data transmission. Volume checksumming is in addition to standard IP and ethernet checksumming. Enabling checksumming for a volume increases data integrity at some cost to system performance.
Target Type	Any	Access mode for the volume. Choices are • the EBSD driver, • an iSCSI initiator, or • Fibre Channel. Note: be certain to plan how the authentication group for this volume is configured if you plan to switch access modes. See Chapter 15, “Working with Authentication Groups.”
iSCSI Target Secret	iSCSI	[Optional for iSCSI] The target secret is a password that is associated with a volume and that must be known by an initiator that wants to use the volume. Use the target secret if you configure 1-way or 2-way CHAP.

12.4 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.

Note: Volume size, replication level, and snapshots should be planned in conjunction. if you intend to use Snapshots, review Chapter 13, "Working with Snapshots."

12.4.1 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 490 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.

12.4.1.1 Managing the Volume Growth Capacity

When you receive the alert that the soft threshold has been reached, you take the following actions.

- Increase the volume size. For our example above, you increase the volume size by 20 percent to 600 GB.
- Increase the hard threshold by about 20 percent, to 590 GB.
- Increase the soft threshold to 585 GB.

See "Editing a Volume" on page 183 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 SSMs.

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 SSM volume.

12.5 Creating a Volume

A volume resides on the SSM(s) contained in a cluster.

1. Log in to the management group for which you want to create a volume.

The management group tab view opens.

2. Select the cluster on which you want to create a volume.
   The cluster tab view opens.
3. Click the Volumes tab.
   The Volumes tab opens, shown in Figure 145.

Figure 145. Viewing the Volumes Tab

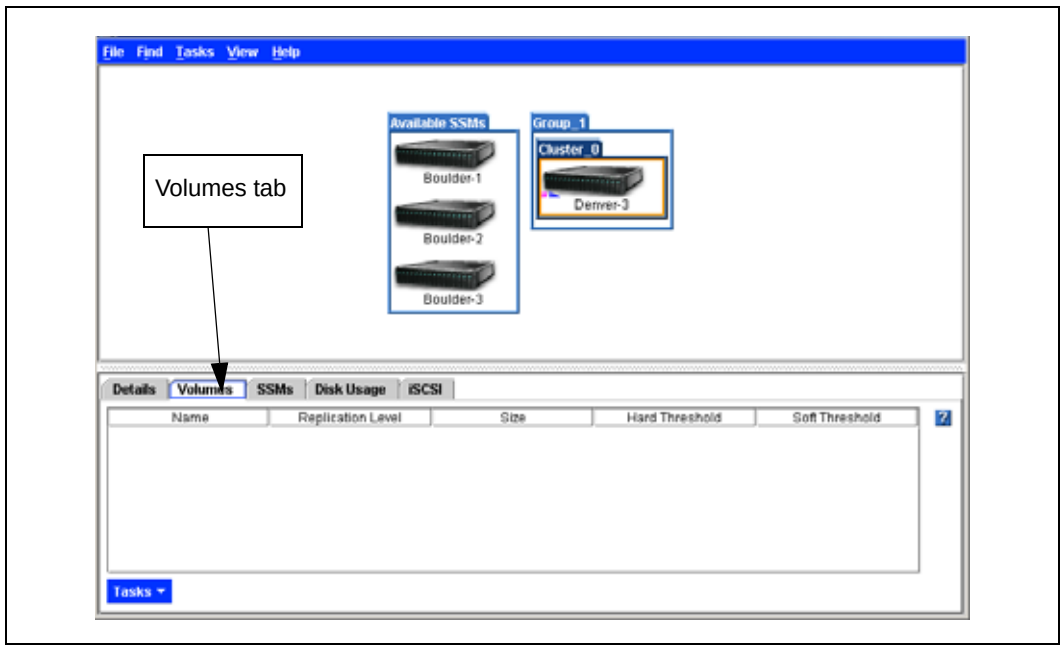

4. From the Tasks menu, click New Volume.
   The New Volume window opens, shown in Figure 146.

Figure 146. Creating a New Primary Volume

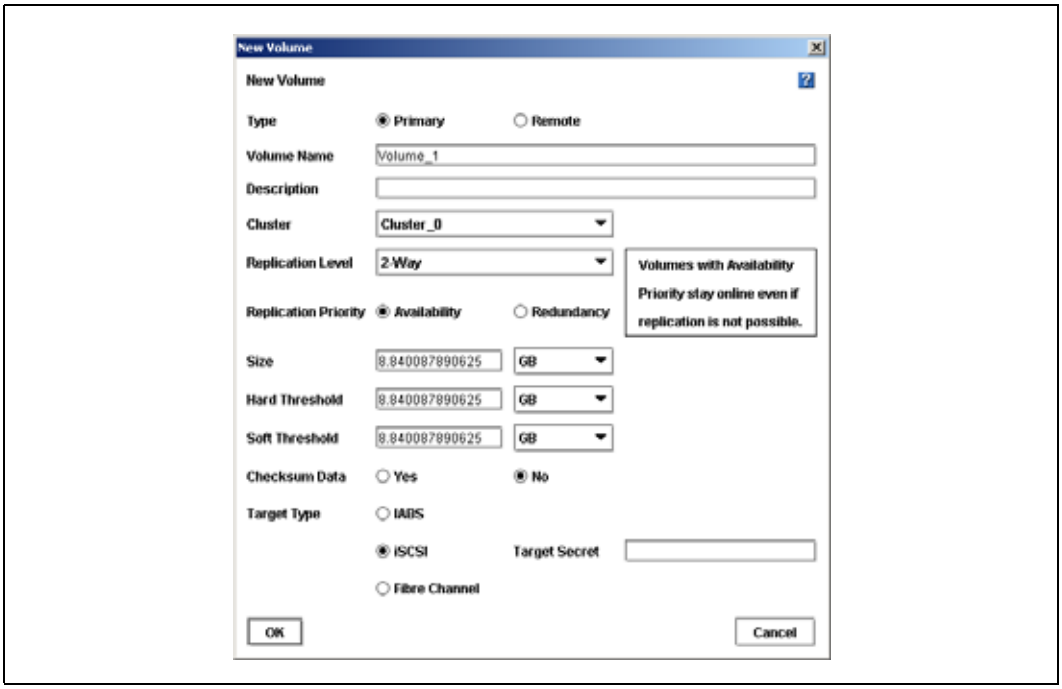

5. Select primary as the volume type.

   The window for a new primary volume is shown in Figure 146. For information about creating a remote volume, see the Remote Copy User's Manual.

6. Type a name for the volume.
7. [Optional] Type a description of the volume.
8. Select a replication level.

   You must purchase the Scalability Pak to use the N-way replication feature beyond the 30-day evaluation period.

9. Select a replication priority.

   If you select a replication level of None, the replication priority must be Availability. See Figure 147.

Figure 147. Setting Replication to None

10. Type a size and select the units.
11. Type a hard threshold and select the units.
12. Type a soft threshold and select the units.

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.

13. Select the cluster you want to contain the volume.
14. Select whether you want to enable checksumming.
15. Select the Target Type for the volume.
16. [Optional for iSCSI only] For 1-way or 2-way CHAP, type a target secret.
17. Click OK.

    The Storage System Software creates the volume and it is attached to the cluster, shown in Figure 148.
18. Select the new volume in the network view.

    The volume tab view opens, also shown in Figure 148.

Figure 148. Viewing a Volume in a Cluster

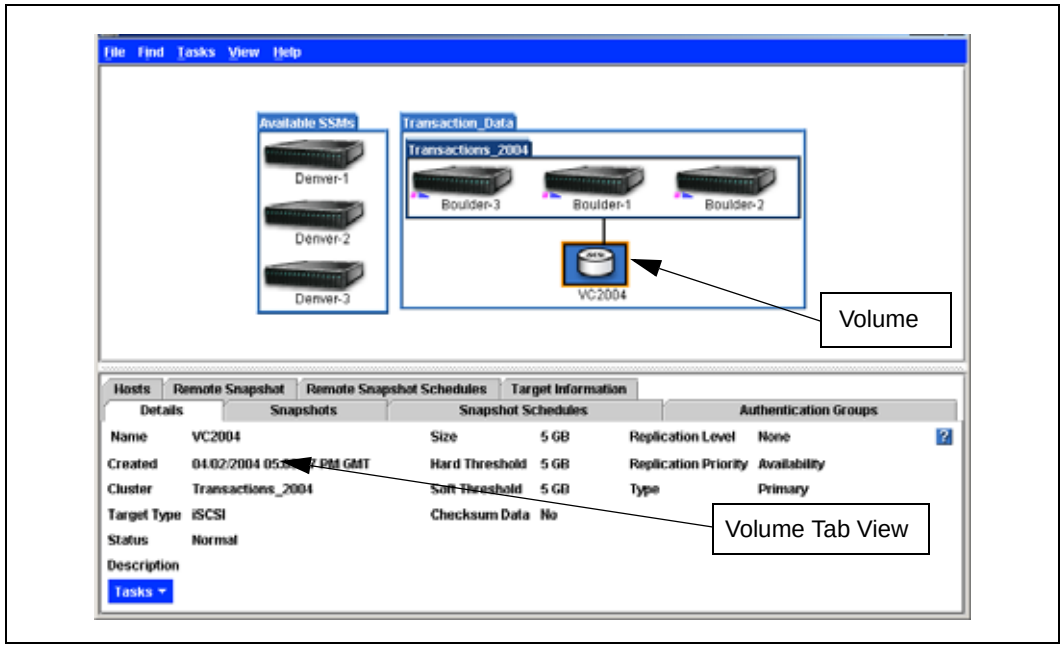

12.5.1 The Volume Tab View

The tabs provide access to volume information and features, such as creating snapshots, and associating authentication groups with the volume.

12.5.1.1 Details Tab

Displays information about the selected volume. You can also edit and delete volumes from this tab.

12.5.1.2 Snapshots Tab

Displays information about the existing snapshots. You can also create, edit, or delete snapshots from this tab. See Chapter 13, “Working with Snapshots,” on page 187 for more information about snapshots.

12.5.1.3 Snapshot Schedules Tab

Displays the name of the snapshot schedule, the hard and soft thresholds set for the snapshots, the actual schedule and any error that prevented a scheduled snapshot from taking place.

12.5.1.4 Authentication Groups Tab

Displays information about authentication groups associated with this volume and enables you to associate groups, edit permissions, and disassociate groups.

12.5.1.5 Hosts Tab

Lists all EBSD hosts that are associated to EBSD volumes. For EBSD hosts, to retrieve the IP, Mode (read, read/write, or none), Type (Windows 2000, 2003, or Linux), and driver version for the volume, put the IP address of the host in Find By Module IP or Host Name.

No host information is displayed for iSCSI and Fibre Channel volumes. Fibre Channel volumes may display "FC Internal" or the display may be blank. iSCSI volumes display the SSM name that the iSCSI initiator is logged into.

Note: Authentication groups must be created at the management group level. See “Working with Authentication Groups” on page 215.

12.5.1.6 Remote Snapshots Tab

Displays the names of the primary and remote snapshots, the management groups they reside in, and the status of the copying from primary to remote. Create a remote snapshot here, or cancel one that is in progress.

12.5.1.7 Remote Snapshot Schedules Tab

Displays the name of the remote snapshot schedule, the hard and soft thresholds set for the remote snapshots, the actual schedule and any error that prevented a scheduled remote snapshot from taking place.

12.5.1.8 Target Information Tab

For iSCSI volumes, displays the virtual IP address if configured and lists any iSNS servers.

For Fibre Channel volumes, displays the Fibre Channel Port World Wide Name (WWPN) for that volume.

12.6 Editing a Volume

When editing a volume, you can change the description, replication level, replication priority, size, hard and soft thresholds, the cluster that contains the volume, whether checksumming is enabled, and the target type of the volume.

Table 37. Requirements for Changing Volumes

Item	Requirements for Changing
Description	Must be from 0 to 127 characters.
Cluster	The target cluster must • Reside in the same management group. • Have sufficient unallocated space for the hard threshold and replication level of the volume being moved. When moving a volume to a different cluster, that volume will temporarily exist on both clusters.
Replication Level	The cluster must have sufficient SSMs 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 SSMs are in the cluster to make the volume available. For a detailed explanation, see Table 35 on page 176. For example, if you have 2-way replication with 3 SSMs in the cluster, you can change from Availability to Redundancy if all the SSMs in the cluster are available and have enough space for replicating the data.
Size	To increase the size of the volume: • There must be sufficient unallocated space in the cluster, or • You can move the volume to a cluster that has enough unallocated space, or • You can add an SSM to the cluster. To decrease the size of the volume: • The size entered must be greater than the hard threshold. You cannot decrease the volume size to a value less than the hard threshold. • You also cannot decrease the size of the volume below the size needed for data currently stored on the volume.
Hard Threshold	• 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. • To decrease the hard threshold, first decrease the size of the volume and then decrease the hard threshold to the same value as the size.
Soft Threshold	• The soft threshold must be equal to or less than the hard threshold. • To decrease the soft threshold, first decrease the hard threshold, and decrease the soft threshold to a value less than the hard threshold.
Target Type	You can change the target type between iSCSI, EBSD and Fibre Channel**.** If you change target type, be sure that you have configured the appropriate settings in the iSCSI initiator, the EBSD driver, or the Fibre Channel host.
Target Secret	If using 1-way or 2-way CHAP, you can edit the target secret. Be certain that you edit the target secret in the iSCSI initiator as well.

Warning: Decreasing the volume size or hard threshold is not recommended without careful planning. See "Changing the Size" on page 185 and "Changing the Hard Threshold" on page 185.

12.6.1 Getting There

1. Select the volume you want to edit in the Network View.
   The volume tab view opens.
2. Click Edit Volume.
   The Edit Volume window opens, shown in Figure 149.

Figure 149. Editing a Volume

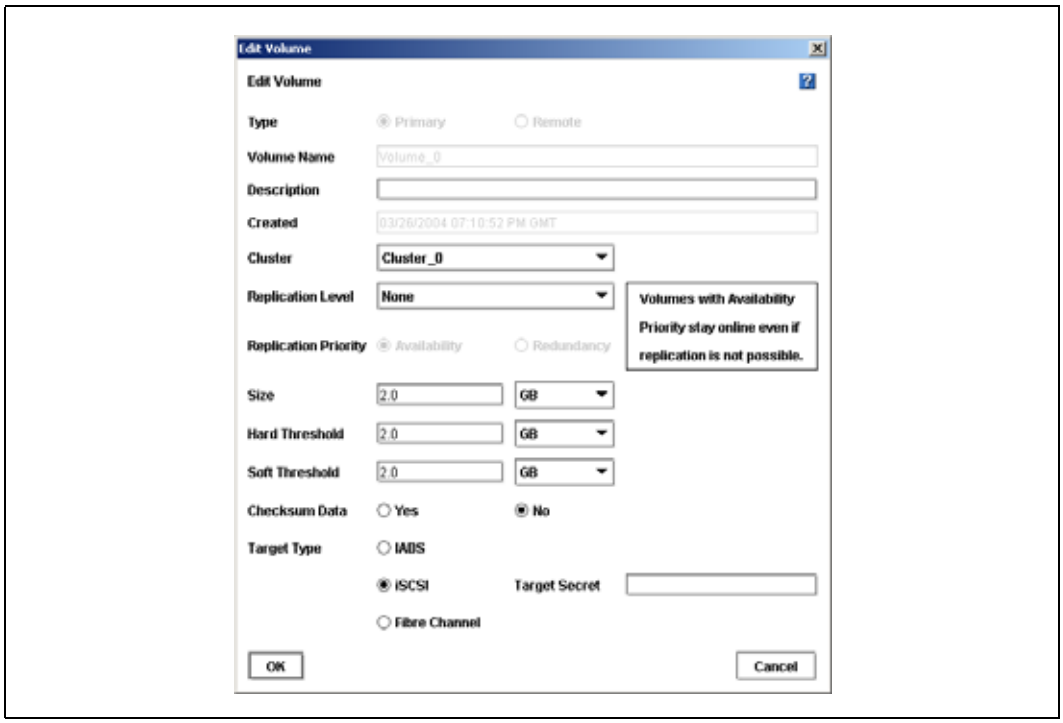

12.6.2 Changing the Volume Description

1. In the Description field, change the description.
2. Click OK when you are finished.

12.6.3 Changing the Cluster

1. In the cluster list, select the cluster to which you want to move the volume.
2. Click OK when you are finished.
   The volume will reside on both clusters until all of the data is moved to the new cluster.

12.6.4 Changing the Replication Level

1. In the Replication Level drop down, select the level of replication you want.
2. Click OK when you are finished.

12.6.5 Changing the Replication Priority

1. Select the replication priority you want.
2. Click OK when you are finished.

12.6.6 Changing the Size

1. In the size field, change the number and change the units if necessary.
2. Click OK when you are finished.

12.6.7 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.

12.6.8 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.

12.6.9 Changing the Target Type

1. Stop any clients from accessing the volume.
2. Select the desired access mode for the volume.
3. [Optional] Type a target secret if you have changed the mode from EBSD to iSCSI and you want to use 1-way or 2-way CHAP.
4. Click OK when you are finished.

12.6.10 Changing the Target Secret

1. Stop any clients from accessing the volume.
2. Type in the new target secret.
3. Click OK when you are finished.

12.7 Fixing a Replica-Challenged Redundant Volume

If a SSM goes offline and needs to be repaired or replaced, and a replicated volume 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 Storage System Console.
3. Right-click and select Edit Volume.
4. Change the data priority from data redundancy to data availability.
5. Remove the SSM from the cluster.
   Repair or replace the module.
6. [optional] Add the new or repaired SSM 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.

12.8 Deleting a Volume

Delete a volume to remove that volume's data from the SSM and make that space available. When deleting volumes, you must delete all snapshots of that volume before you can delete the volume itself.

Warning: Deleting a volume removes that volume's data permanently from the SSM.

Warning: Deleting a Fibre Channel volume requires a reboot of Windows 2000 and Windows 2003 systems.

Prerequisites

- Delete all snapshots of the volume that you want to delete.
- Stop applications from accessing the volume.
- Disable the drives on the host.

See the EBSD Driver User Manual for detailed instructions about working with the client and the disks on the client.

1. Select the volume you want to delete.
   The volume tab view opens.
2. Click Delete Volume.
   A confirmation window opens.
3. Click OK.
   The volume is removed from the cluster.

Working with Snapshots 13

13.1 Snapshots Overview

Snapshots provide a fixed version of a volume for read-only access.

13.1.1 Snapshots vs. Backups

Unlike backups, which 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, but not device or storage media failure.

Prerequisites

Before you create a snapshot, you must have created

- a management group
- a cluster, and
- a volume.

13.1.2 Topics Covered in This Chapter

- Single snapshots and scheduled snapshots
- Managing capacity using volume and snapshot thresholds
- Creating snapshot schedules

13.2 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, all the snapshots created after that snapshot are deleted. Also, using a third-party utility, you can copy a snapshot to a different server and open the snapshot as a volume on that server.

Snapshots can be used for

- 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

13.2.1 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.

13.3 Requirements for Snapshots

Review in Chapter 12, “Working with Volumes” the section “Planning Volumes” on page 171 to ensure that you configure snapshots correctly. When creating a snapshot, you define the following parameters.

Table 38. Snapshot Parameters

Snapshot Parameter	What it means
Snapshot Name	The name of the snapshot that is displayed in the Storage System Console. A snapshot name must be from 1 to 127 characters and is case sensitive.
Description	[Optional] A description of the snapshot.
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.
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.
Target Type	Access mode for the snapshot. Choices are the EBSD driver, an iSCSI initiator, or a Fibre Channel host. Note: be certain to plan how the authentication group for this snapshot is configured if you plan to switch access modes. See Chapter 15, “Working with Authentication Groups.”
iSCSI Target Secret	[Optional for iSCSI] The target secret is a password that is associated with a snapshot and that must be known by an initiator that wants to use the snapshot. Use the target secret if you configure 1-way or 2-way CHAP. See Chapter 12, “Working with Volumes” “Requirements for Configuring CHAP” on page 176.

13.4 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 IXA SDKs 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 172.

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 after it. (Hard thresholds of the volume or snapshot directly after the deleted snapshot will increase up to the size of the volume, and soft thresholds will increase up to the hard threshold.) 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, you must calculate the effect of adding the hard thresholds back into the volume.

13.4.1 Easiest Method for Planning Capacity

Make the snapshot hard threshold equal to the volume size, and the soft threshold equal to the hard threshold.

13.4.2 Most Flexible 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.

The tables illustrate the effects of decreasing or not decreasing the hard threshold to reduce the space on the cluster occupied by snapshots.

Table 39. Space Used by Snapshots when Hard Threshold not Reduced

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 40. Space Used by Snapshots when Hard Threshold is Reduced

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

Note: Deleting files on a file system does not create space on the volume. For file level capacity management, use application or file system-level tools.

13.5 Planning Snapshots

When planning to use snapshots, take the purpose and size considerations into account.

Note: When considering the size of snapshots in the cluster, remember that the replication level of the volume is duplicated in the snapshot.

13.5.1 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 will never exceed the volume size.

13.5.2 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 will never exceed the volume size.

13.5.3 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 deleted?

13.6 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.

   The management group tab view opens.

2. Select the volume on which you want to create a snapshot.

   The volume tab view opens, shown in Figure 150.

Figure 150. Volume Tab View

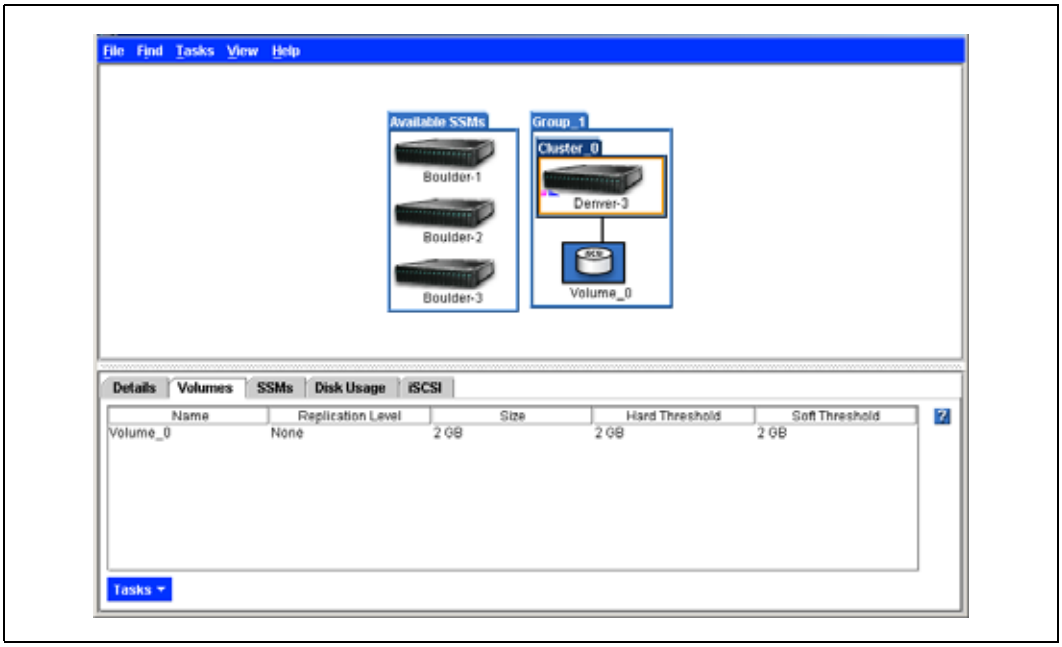

3. Click the Snapshots tab to bring it to the front, as shown in Figure 152.
4. From the Tasks menu, select New Snapshot.
   The New Snapshot window opens, shown in Figure 151.

Figure 151. Creating a new Snapshot

5. Type a name for the snapshot.

Names are case sensitive. They cannot be changed after the snapshot is created.

6. [Optional] Type in a description of the snapshot.
7. [Optional] Change the hard and soft thresholds for the snapshot.

   You must purchase the Configurable Snapshot Pak to change the hard threshold of a snapshot beyond the 30-day evaluation period.

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 189.

8. [Optional] Select a target type for the snapshot if you want it to be different than the target type for the volume.
9. Click OK when you are finished.

   The Snapshots tab opens with the new snapshot listed. The new snapshot also displays in the Network view, as shown in Figure 152.

Note: Snapshots are listed below the volume in descending date order - from newest to oldest.

Figure 152. Viewing the new Snapshot

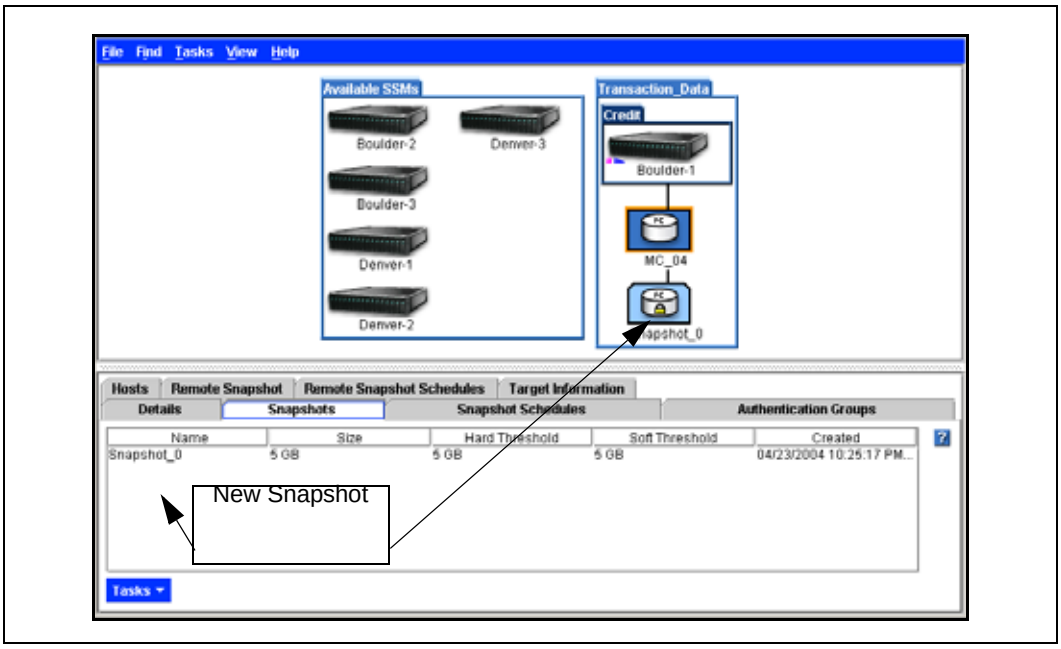

13.6.1 The Snapshot Tab View

Clicking on the snapshot itself opens the snapshot tab view, shown in Figure 153.

Figure 153. Snapshot Tab View

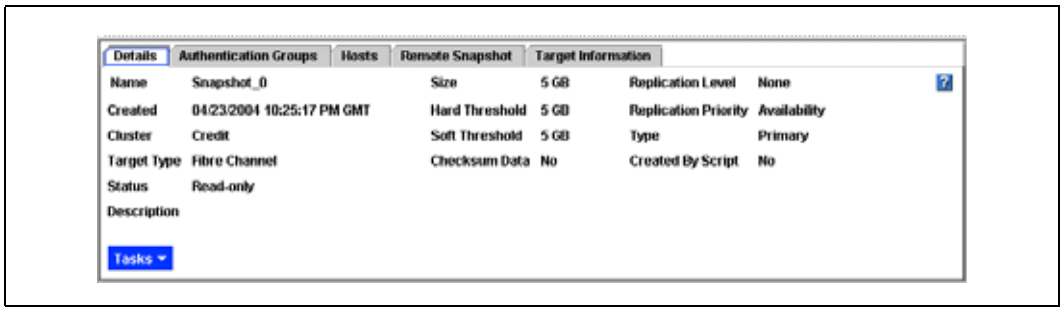

The tabs provide access to snapshot information and features, such as editing snapshots, rolling back a volume, and associating authentication groups with the snapshot. The Tasks button on each tab provides access to the actions you can take related to that tab.

13.6.1.1 Details Tab

Displays information about the selected snapshot. Use the Tasks menu to edit and delete snapshots and roll back volumes from this tab.

13.6.1.2 Authentication Groups Tab

Displays information about authentication group associations, which are inherited from the parent volume, except for permissions which are read-only for snapshots. Additionally, Fibre Channel snapshots will not inherit the parent volume LUN numbers. Use the Tasks menu to associate additional authentication groups, edit permissions, and disassociate groups. For information about authentication groups, and associating them to snapshots, see Chapter 15, “Working with Authentication Groups.”

13.6.1.3 Hosts Tab

Lists all EBSD hosts that are associated to the snapshot. To retrieve the IP, Mode, Type, and driver version for the snapshot, put the IP address of the host in Find By Module IP or Host Name.

13.6.1.4 Remote Snapshot Tab

Lists remote snapshots associated with a snapshot. Buttons include creating a remote snapshot and canceling a remote snapshot that is in progress.

13.6.1.5 iSCSI

Displays the Initiator name assigned by the Microsoft iSCSI Initiator and the masked Target Secret.

13.7 Mounting a Snapshot

A snapshot is a read-only volume. Read-only volumes cannot be mounted on Windows 2000. You can mount read-only volumes with EBSD in Windows 2003 and on Linux.

Warning: Windows 2000 allows buffered writes to read-only volumes. Once Windows 2000 attempts to write the data from the buffer to the read-only volume, the write will fail.

13.8 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 191.

1. Log into the management group that contains the snapshot that you want to edit.
2. Select the snapshot you want to edit.
   The snapshot tab view opens.
3. From the Tasks menu, select Edit Snapshot.
   The Edit Snapshot window opens, shown in Figure 154.

Figure 154. Editing a Snapshot

4. Navigate to the field you want to change and change the information.

Table 41. 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.

Table 41. Data Requirements for Editing a Snapshot

Item	Requirements for Changing
Soft Threshold	Soft threshold size must be equal to or less than the hard threshold size.
Target Type	Ensure that the corresponding driver, initiator, or host is configured for the snapshot.
Target Secret	If using 1-way or 2-way CHAP, you can edit the target secret. Be certain that you edit the target secret in the iSCSI initiator as well.

5. Click OK when you are finished.
   The snapshot tab view opens, shown in Figure 152.

13.9 Manually Copying a Volume from a Snapshot

Once you have mounted the snapshot on a host you can do the following:

- Copy the snapshot 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 217.
2. Associate the authentication group to the snapshot for read-only access. Set that group’s permissions to read-only.
   See “Creating an Authentication Group Association” on page 223.
3. Configure client read-only access to the snapshot volume according to the instructions for your operating system in the EBSD Driver User Manual.
   Now you can access the snapshot
   - — 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

13.10 Creating Snapshot Schedules

Using the Storage System Console you can schedule recurring snapshots. Recurring snapshots can be scheduled in a variety of frequencies and with a variety of retention policies.

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.

Scripted snapshots is an add-on feature. You must purchase the Configurable Snapshot Pak to use snapshot scripting beyond the 30-day evaluation period.

13.10.1 Requirements for Scheduling Snapshots

Scheduled snapshots require particular attention to capacity management. Additionally, you must ensure that the time settings on the IXA SDKs running managers and the time setting of the management group are synchronized.

Note: Use NTP for ensuring that all the IXA SDKs in the management group have synchronized time settings.

Table 42. 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 189. 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>. It is possible for the new snapshot and the one to be deleted to coexist in the cluster for some period of time. 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 IXA SDK times with management group time	The time setting on the IXA SDKs running managers and the time setting of the management group must be synchronized. If they are not synchronized, then the snapshot schedule might run incorrectly. Be sure to configure the correct time on the IXA SDK. See “Date and Time Overview” on page 85. Also, see “Resetting the Management Group Time” on page 137.

13.10.2 Creating Snapshot Schedules

You can create one or more snapshot schedules 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. Select the volume for which you want to schedule snapshots.
   The volume tab view opens.
2. Click the Snapshot Schedules tab to bring it to the front.
3. From the Tasks menu, select New Schedule.
   The New Snapshot Schedule window opens, shown in Figure 155.

Figure 155. Creating a Snapshot Schedule

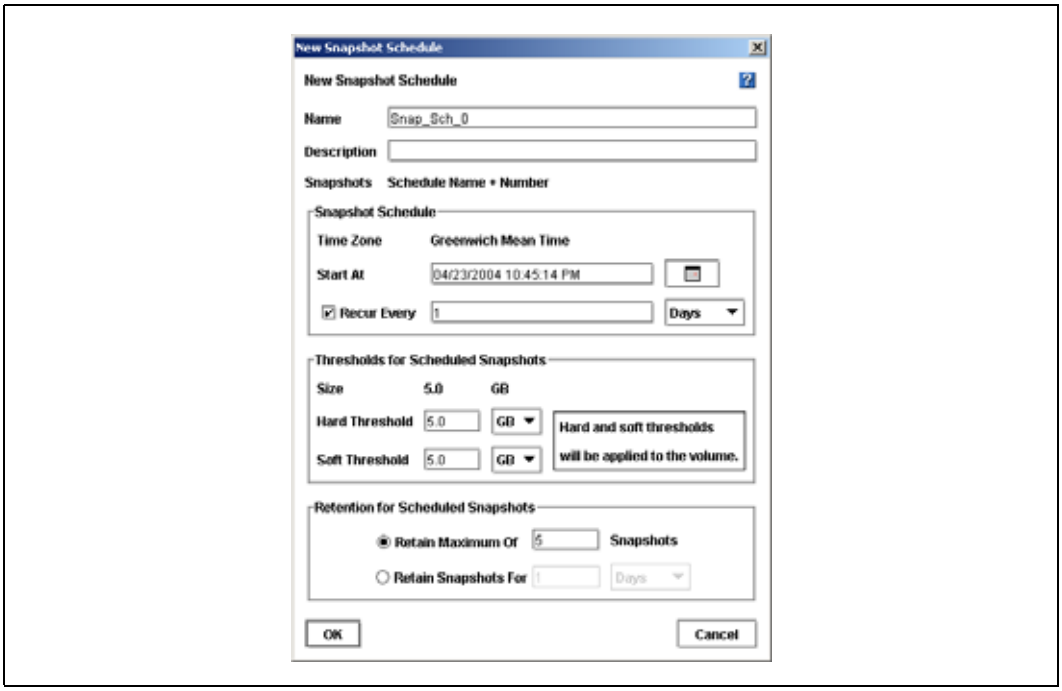

4. Type a name for the snapshots.

   The name will be used with sequential numbering. For example, if the snapshot name is Backup, the list of scheduled snapshots will be named Backup1, Backup2, Backup3.

5. [Optional] Enter a snapshot description.

6. [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 189.

7. Enter a start date and time.

   The date and time must be valid, but they can occur in the past.

8. Select a recurrence schedule.

   The recurrence schedule can be in minutes, hours, days or weeks.

9. Set a retention schedule.

   The retention schedule can be for specified number of snapshots, or for a designated period of time.

10. Click OK.

    The New Snapshot Schedule window closes and the new snapshot schedule appears on the tab, shown in Figure 156.

Figure 156. List of Scheduled Snapshots

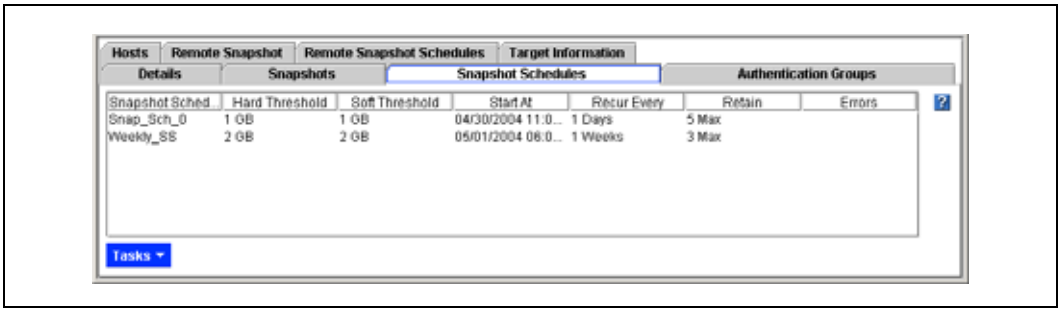

13.10.3 Editing Snapshot Schedules

You can edit everything in the snapshot schedule except for the name.

1. Select the volume for which you want to edit the snapshot schedule.
   The volume tab view opens.
2. Click the Snapshot Schedules tab to bring it to the front.
3. Select the schedule you want to edit.
4. From the Tasks menu, select Edit Schedule.
   The Edit Snapshot Schedule window opens, shown in Figure 157.

Figure 157. Editing a Snapshot Schedule

5. Change the desired information.

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 189.

6. Click OK.

13.10.4 Deleting Snapshot Schedules

1. Select the volume for which you want to delete the snapshot schedule.
   The volume tab view opens.
2. Click the Snapshot Schedule tab to bring it to the front.
3. Select the schedule you want to delete.
4. From the Tasks menu, select Delete Schedule.

13.11 Scripting Snapshots

Application-based scripting is available for taking snapshots. Using application-based scripts allows automatic snapshots of a volume. For detailed information, see Chapter 14, "Working with Scripting."

Check your vendor's web site for specific applications for which sample scripts have been developed.

13.12 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.

13.12.1 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.

Table 43. Requirements for Rolling Back a Volume

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 189.
New Soft Threshold	Soft threshold size must be equal to or less than the hard threshold size.

Table 43. Requirements for Rolling Back a Volume

Parameter	Requirements for Changing
Authentication Groups	You must associate authentication groups to the new volume. See "Associating Authentication Groups Overview" on page 222.
Fibre Channel LUN #	You must add a new LUN # for the Fibre Channel volume.
Hosts	You must reconfigure hosts to connect to the new volume.

Warning: After rolling back a volume to a snapshot, you lose all data stored after the rolled back snapshot.

Prerequisites

- Stop applications from accessing the volume.

13.12.1.1 Rolling Back the Volume

1. Log in to the management group that contains the volume that you want to roll back.
2. 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. From the Tasks menu on the Details tab, select Roll Back Volume.
   The Roll Back Volume window opens, shown in Figure 158.

Figure 158. Rolling Back a Volume

5. Type a new name for the rolled back volume.
   You can also change the hard threshold and soft threshold if necessary.
6. [Optional] You can change the target type for the rolled back volume.

7. Click OK.
   The Roll Back Volume confirmation message, shown in Figure 159, explains that the original volume and all newer snapshots will be deleted.

Figure 159. Verifying the Volume Roll Back

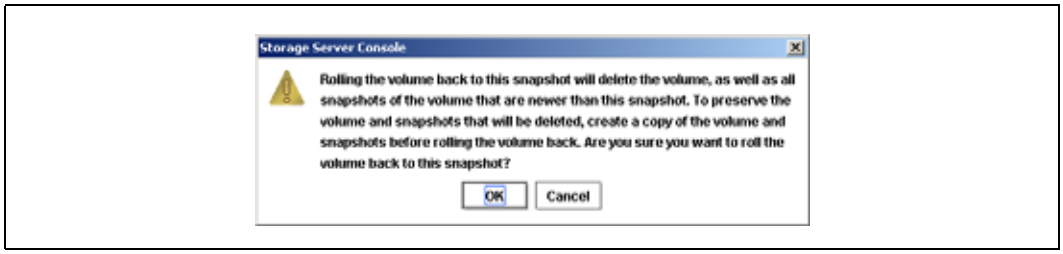

8. Click OK.
   The snapshot version of the volume is restored as a read/write volume.
9. Reassociate authentication groups to the restored volume.
10. For Fibre Channel volumes, add a new LUN number.
11. Reconfigure hosts to access the new volume.

Warning: All snapshots between the current date and the roll back are deleted. The original volume is also deleted.

13.13 Deleting a Snapshot

Deleting a snapshot removes that snapshot's data from the IXA SDK and removes the snapshot from the Network View. The space used by the deleted snapshot is added to the snapshot or volume directly after it.

Prerequisites

- Stop applications from accessing the snapshot.
- Disable the snapshot if it is mounted by a host.

Note: These prerequisites do not apply to Fibre Channel snapshots.

13.13.0.1 Deleting a Snapshot

1. Log into the management group that contains the snapshot that you want to delete.
   The management group tab view opens.
2. Select the snapshot that you want to delete.
3. Review the Details tab to ensure you have selected the correct snapshot.
4. From the Tasks menu on the Details tab, click Delete Snapshot.
   A confirmation message opens.
5. Click OK.

Warning: Deleting a snapshot causes that snapshot's data to be unavailable from the IXA SDK.

Working with Scripting 14

14.1 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 Chapter 13, “Working with Snapshots.” for detailed information about snapshot requirements. Information about Remote Copy can be found in Appendix F, “Using Remote Copy”.

Sample scripts for specific applications may be available from your vendor.

14.1.1 Overview of Scripting

The tasks supported by scripting includes

- Taking a snapshot of the volume
- Mounting the snapshot
- [Optional] Unmounting or deleting the snapshot
- Increasing volume thresholds

Two tools, named **`java.commandline.CommandLine`** and **`aebsvm`** are provided to access the Storage System Console functionality. The tools are described in the following sections and sample scripts may be available which illustrate correct use of these tools.

14.2 Tools for Scripting

Two software tools are available to use in scripts. The first one, `java commandline.CommandLine`, is used to create and delete snapshots, and to automatically increase volume thresholds. The second one, `aebsvm`, is used to mount the snapshot.

14.2.1 Java commandline. CommandLine

`Java commandline.CommandLine` is the program that actually invokes the snapshot function in the Storage System Console for creating and deleting snapshots. In addition, the program can respond when a soft threshold is reached on a volume and automatically increase the hard and soft thresholds on that volume.

- First, set the environment

Table 44. Setting the Environment for Using Scripting Tools

Operating System	Syntax	Example
Windows	set CLASSPATH <full path to Console.jar>	`set CLASSPATH C:\Program Files\Storage_System\Storage_ System_Software\ 6.0\Console\ Console.jar`
Unix (C Shell type)	setenv CLASSPATH <full path to Console.jar>	`setenv CLASSPATH /opt/ Storage_System/ Storage_System_Software/ 6.0/ Console/ Console.jar`
Unix (Bourne or Kshell or Bash)	export CLASSPATH=<full path to Console.jar>	`export CLASSPATH=/opt/ Storage_System/ Storage_System_Software/ 6.0/ Console/ Console.jar`

- Then, run the tool
  `java 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 45. Parameters for `java commandline.CommandLine`

Parameter	What It Is
admin name	Value = text Name of the administrator with full administrative privileges. Can be either the primary or remote administrator, if they are different.
admin password	Value = text The administrator's Storage System Console password. Can be either primary or remote password, if they are different.
manager ip	Value = IP address IP address of an SSM running a manager in the management group containing either the source volume or the remote volume.
volume name	Value = text Name of the volume to snapshot. May be the primary volume if using Remote Copy.
snapshot name	Value = text Name of the snapshot to create.
primary volume name	Value = text Name of the primary volume to make remote.
remote volume name	Value = text Name of the remote volume created in the Console.
remote snapshot name	Value = text Name of the remote snapshot.

Table 45. Parameters for `java commandline.CommandLine`

Parameter	What It Is
remote snapshot description	Value = text Description of the remote snapshot.
soft threshold*	Value = number 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 threshold*	Value = number 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 [Optional] Description associated with the snapshot.
failure timeout seconds	Value = number The number of seconds to wait until exiting with a failure.
grow size	Value = number The size in MegaBytes by which to increase the volume thresholds.
volume_snapshot	Use this value as written. (This is verbatim.)
volume_delete	Use this value as written. (This is verbatim.)
volume_remote_snapshot	Use this value as written. (This is verbatim.)
volume_make_primary	Use this value as written. (This is verbatim.)
volume_make_remote	Use this value as written. (This is verbatim.)
volume_autogrow_set	Use this value as written. (This is verbatim.)
volume_autogrow_get	Use this value as written. (This is verbatim.)

NOTE: * You must provide either all three items, or none of them. For example, you cannot provide only a soft threshold value.

14.2.2 aebsvm

`aebsvm` is the program that mounts the snapshot or volume. Table 46 lists the parameters available in `aebsvm`.

Getting Help

1. Type `aebsvm help` and press Enter.

Table 46. Parameters for `aebsvm`

Parameter	What It Is
mgmt group name	Value = text Name of the management group containing the source volume.
snapshot name	Value = text Name of the snapshot or volume created using java commandline.CommandLine.
auth group name	Value = text Name of the authentication group associated with the volume in the Storage System Console.

Table 46. Parameters for `aebsvm` (Continued)

Parameter	What It Is
local ip	Value = IP address The IP address of the machine the script is running on.
number of managers	Value = number The number of managers in the management group that contains the source volume.
managers' ip	Value = IP address Separator = space The IP addresses of the SSM managers in the management group.
lock_mode	Value = ro, rw The attribute of the volume - read-only or read-write.

14.3 Scripted Commands for Volumes and Snapshots

Below are examples of the Storage System Software functions that can be accomplished using application-based scripts.

14.3.1 Creating a Snapshot

Create a snapshot using

`java commandline.CommandLine`

java 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

14.3.2 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

14.3.3 Mounting a Snapshot

Below is an example of mounting the snapshot using **`aebsvm.`**

aebsvm <mgmt group name> <snapshot name> <auth group name> <local ip> <number of managers> <each managers's ip> <lock_mode>

Example

Joe Jones plans to mount his XRayReview snapshot and mount it on another server where the group named adminusers (the administrators of the orthopedic section) can access the images for filing the patient database.

aebsvm Images XRayReview adminusers 10.0.20.212 3 10.0.13.79 10.0.33.47 10.0.33.87 ro

14.3.4 Increasing Volume Hard and Soft Thresholds

You can create a script that will automatically increase 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 accomodate 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 SSMs or to increase the volume length, follow instructions as described in Chapter 11, “Working with Clusters” or Chapter 12, “Working with Volumes.”

14.3.4.1 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 auto_grow_set X-Rays 512 600

14.3.4.2 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 auto_grow_get X-Rays 60

14.4 Scripted Commands for Remote Copy

Scripting operations for Remote Copy uses the same tools that are available for scripting snapshots, with the addition of parameters specific to Remote Copy. Using the command line parameters allows you to create scripts for

- creating a primary snapshot
- creating a remote snapshot
- making a primary volume into a remote volume
- failing over to a remote snapshot

14.4.1 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 (Mega-bytes)> <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 looks 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 saxaphone 10.10.45.72 volume_remote_snapshot `RemVolX_Rays RSS2_xrays` “second remote snapshot” jjones trumpet 10.0.111.212 PSS2_xrays 15

14.4.2 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 admin 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

14.4.3 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.

14.4.3.1 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>]

14.4.3.2 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 saxaphone 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

14.4.4 Scripting Failover

Scripting failover uses a **`java commandline.CommandLine`** script along with the **`aebsvm`** script for mounting a snapshot.

14.4.4.1 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>]

14.4.4.2 Mount New Primary Volume

aebsvm <remote mgmt group name> <remote volume name> <auth group name> <local ip> <number of managers> <each managers's ip> <lock_mode>

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 saxaphone 10.10.45.72 volume_make_primary `RemVolX_Rays 512000 512000 30`

aebsvm Remote_ImagesRemVolX_Rays adminusers 10.3.11.19 3 10.3.11.27 10.3.11.31 10.3.11.12 ro

Working with Authentication Groups 15

Authentication groups are sets of application servers or other machines that access a volume.

You first create an authentication group at the management group level, and then, when you create a volume, you associate the authentication group to the specific volume it needs to access.

For related information about associating authentication groups with volumes, see “Associating Authentication Groups Overview” on page 222.

15.0.1 Topics Covered in This Chapter

- Authentication groups and access modes
- Creating and managing authentication groups
- Associating volumes with authentication groups

15.1 Types of Volume Access

The Storage System Software supports three modes of access to volumes - through an iSCSI initiator, the EBSD driver, or a Fibre Channel host. Volumes can be configured for use with the EBSD driver, an iSCSI initiator or a Fibre Channel host, and can be changed between those types. You can configure an authentication group for one, two or all three modes if desired.

When creating an authentication group, you choose the level of host access you want. Choices are

- All hosts on the network - Authenticate all hosts, that is, allow all to have access. This choice is not available for Fibre Channel.
- Hosts with
  - a specific initiator name and [optional] initiator secret (iSCSI),
  - specific subnet and mask (EBSD), or
  - specific World Wide Name (Fibre Channel)

  Authenticate all hosts with the specified parameters.
- No hosts - Do not authenticate any hosts, that is, let no one have access.

15.2 Authentication Groups and iSCSI

If you use 2-way iSCSI CHAP (Challenge-Handshake Authentication Protocol), you must configure a secret for each initiator (application server) listed in the authentication group. The initiator secret is created in the iSCSI initiator and is then entered along with the initiator name when you configure the authentication group. See the section on “Planning Volumes and iSCSI” on page 175 for a detailed description of configuring CHAP.

15.3 Authentication Groups and Fibre Channel

When configuring authentication groups for Fibre Channel, you must create an authentication group for each Fibre Channel volume. Then, when you assign the authentication group to the volume, you assign a LUN number for that volume.

15.4 Assigning LUN Numbers to Volumes

Some requirements for authentication group access and LUN numbers include the following:

- LUN numbers must be assigned to volumes via authentication groups before the hosts can access them.
- LUN numbers must be unique per host. For example, one host can only access one LUN # 0. If you plan to have one host accessing two LUNs, each authentication group association must have a different LUN number.

15.4.0.1 Best Practice – Hosts with Separately Numbered LUNs

Figure 160 shows a typical best practice configuration for assigning LUN numbers and associating LUNs to hosts. Each host is accessing two LUNs. Each of the two LUN associations has a different LUN number.

Figure 160. Example Best Practice Configuration for Assigning LUN Numbers

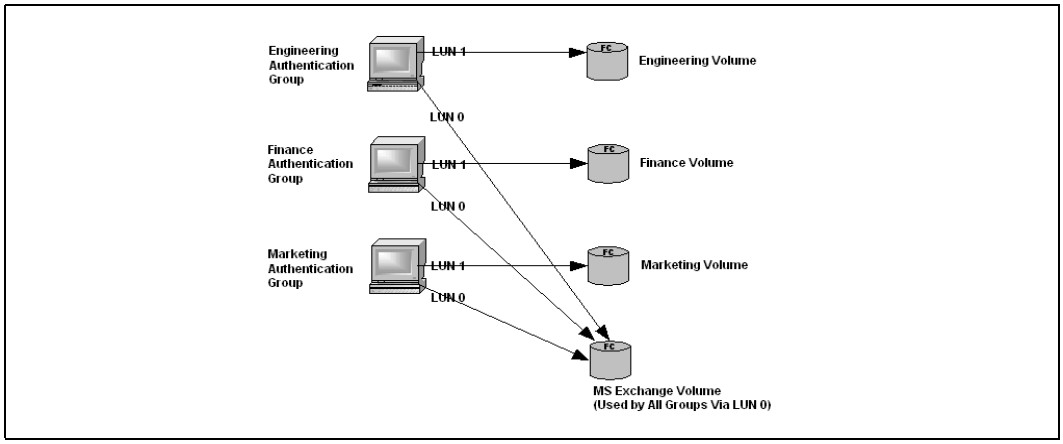

15.4.0.2 Prohibited – Host with Duplicate Numbered LUN

Figure 161 illustrates a prohibited LUN numbering configuration.

Figure 161. LUN Numbering Configuration that is NOT Allowed

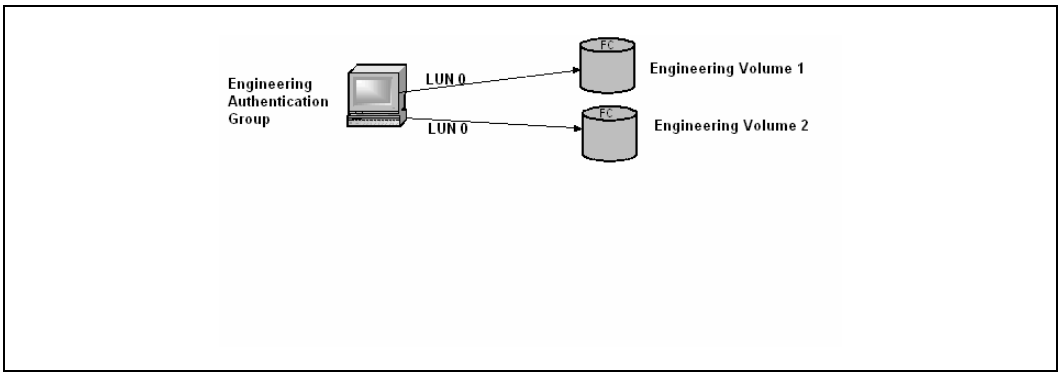

15.4.0.3 Possible – Hosts with a Shared LUN

It is possible to associate multiple hosts to one LUN, in which case you should make one host association read/write and the other host associations read only, as shown in Figure 162.

Figure 162. LUN Numbering Configuration with one LUN Shared Among Three Hosts

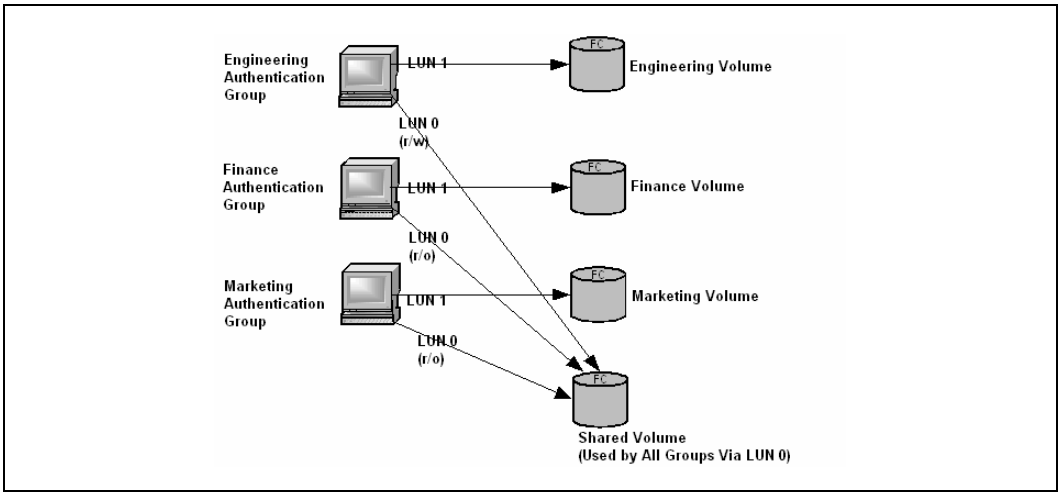

15.5 Creating an Authentication Group

1. Log into the management group and select that management group in the network view. The management group tab view opens.
2. Click the Authentication Groups tab to bring it to the front.
3. From the Tasks menu, select New Authentication Group.
   The New Authentication Group window opens, shown in Figure 163.

Figure 163. Creating a New Authentication Group

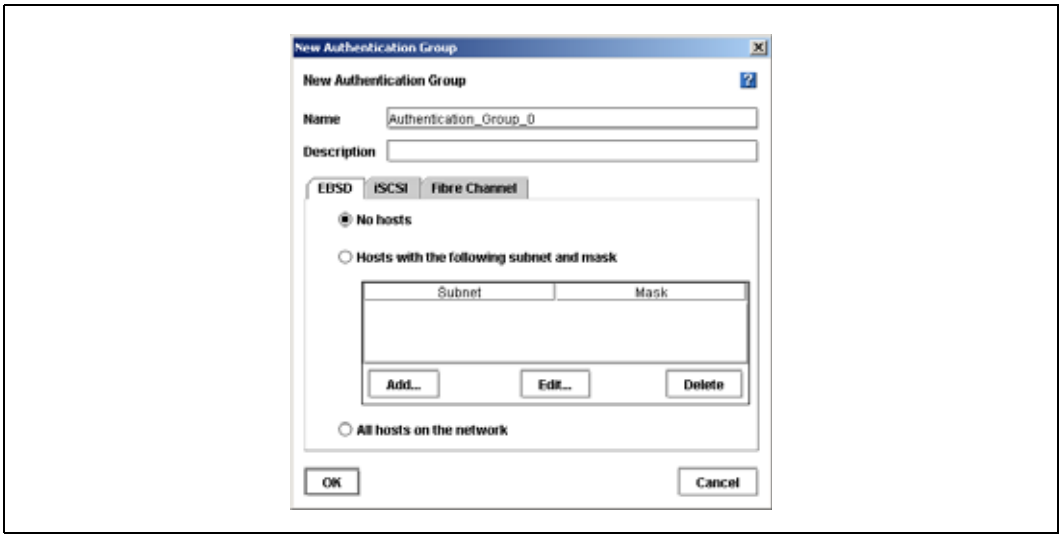

4. Type a name and description for the authentication group.
   The authentication group name is case sensitive.
5. Select the tab for the appropriate type of host access.

15.5.0.1 Configuring EBSD

1. On the EBSD tab, shown in Figure 163, select the level of access for hosts using the EBSD driver.

Table 47. Choosing the Level of Access for Hosts Using the EBSD Driver

Authentication Method	What Happens
No hosts	No application server gains access.
Hosts with the following subnets and masks	Only hosts on the designated subnet and mask gain access. If selecting this method, 1. Click Add. 2. Enter a subnet and mask. 3. Click OK.
All hosts on the network	All hosts gain access.

15.5.0.2 Configuring iSCSI

1. On the iSCSI tab, shown in Figure 164, select the level of access for hosts using an iSCSI initiator.

Warning: Allowing more than one iSCSI application server to connect to a volume could result in data corruption.

Figure 164. Creating iSCSI Access in New Authentication Group

Table 48. Choosing the Level of Access for Hosts Using an iSCSI Initiator

Authentication Method	What Happens
No hosts	No application server gains access.
Hosts with the following initiator name and, if using 2-way CHAP, initiator secret.	Only initiators with the supplied name can access the volume. If selecting this method, 1. Click Add. 2. Enter initiator name that is created by the initiator. 3. [Optional] Enter the initiator secret that you configured in the iSCSI initiator. 4. Click OK.
All hosts on the network	All hosts gain access. **Warning:** This could cause data corruption.

15.5.0.3 Configuring Fibre Channel

1. On the Fibre Channel tab, shown in Figure 165, select the level of access for Fibre Channel hosts.

Figure 165. Creating Fibre Channel Access in New Authentication Group

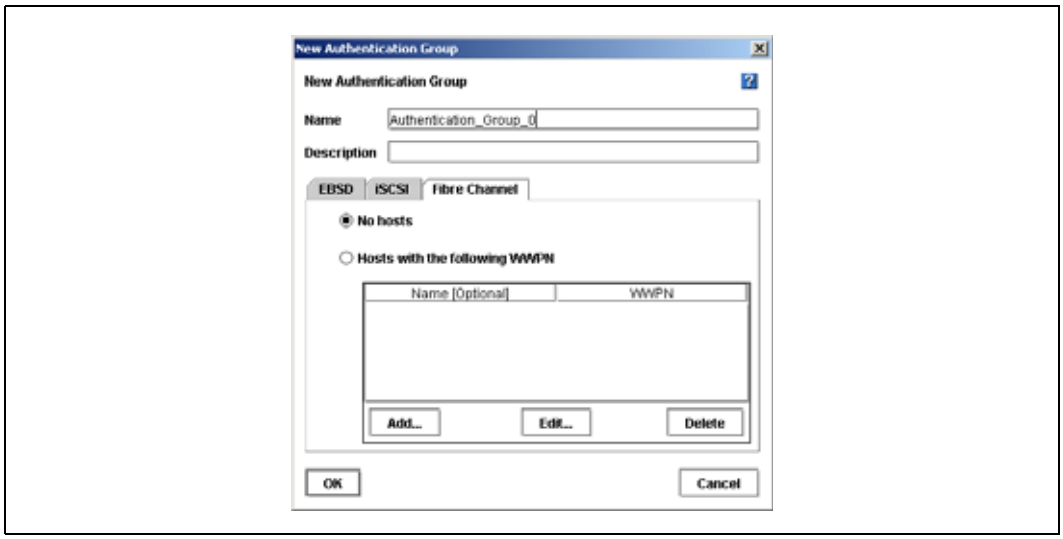

Table 49. Choosing the Level of Access for Fibre Channel Hosts

Authentication Method	What Happens
No hosts	No application server gains access.
Hosts with the following Port World Wide Name (WWPN).	Only hosts with the supplied WWPN can access the volume. If selecting this method, 1. Click Add. 2. [Optional] Enter the host name. 3. Enter the host WWPN. This name must be in hexadecimal (0-9, a-f). 4. Click OK. **NOTE:** A WWPN may only be in one authentication group.

15.5.0.4 Finishing Up

1. Click OK when you are finished.

   The Authentication Group tab opens, with the new group displayed in the list, shown in Figure 166. You can now associate that authentication group with any of the volumes created in the management group. See “Creating an Authentication Group Association” on page 223.

Figure 166. Viewing the Authentication Group Tab

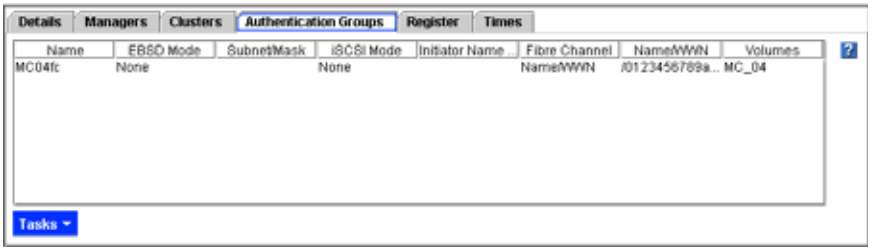

15.6 Editing an Authentication Group

You can edit authentication groups, including the description and authentication level. You cannot change the name of an existing authentication group.

See "Authentication Groups and iSCSI" on page 215 and "Authentication Groups and Fibre Channel" on page 216 before changing iSCSI or Fibre Channel authentication group parameters.

1. Log into the management group and select that management group in the network view.
   The management group tab view opens.
2. Click the Authentication Groups tab to bring it to the front.
3. Select from the list the group you want to edit.
4. From the Tasks menu, select Edit Authentication Group.
   The Edit Authentication Group window opens, shown in Figure 167.

Figure 167. Editing an Authentication Group

5. Select the appropriate tab for the group you are editing.
6. Change the appropriate information.

7. Click OK when you are finished.

15.7 Deleting an Authentication Group

You must delete the group's volume associations before you can delete the group itself. See "Deleting an Authentication Group Association" on page 225.

1. Log into the management group and select that management group in the network view.
   The management group tab view opens.
2. Click the Authentication Groups tab to bring it to the front.
3. Select from the list the group you want to delete.
4. From the Tasks menu, select Delete Authentication Group.
   A confirmation window opens.
5. Click OK to delete the group.

15.8 Associating Authentication Groups Overview

In order for an authentication group to access a volume or snapshot, the group must first be associated to that volume. First you create the authentication groups in the management group, then you associate those groups with specific volumes.

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.
- At least one authentication group has been created for the management group.

Warning: When associating or deleting associations for Fibre Channel LUNs, the host server's Disk Management (or equivalent) window must be closed.

15.9 Requirements for Authentication Group Associations

- Only one authentication group can have read/write access to a volume.
- Additional authentication groups with read-only access with can be associated with the same volume or snapshot.
- For Fibre Channel volumes or snapshots, you must assign LUN numbers as part of the authentication group association.

15.10 Creating an Authentication Group Association

Associating an authentication group to a volume or snapshot controls that group's access to the volume.

1. Log into the appropriate management group.
2. Select the volume or snapshot to which you want to associate an authentication group.
   The volumes or snapshot tab view opens.
3. Click the Authentication Groups tab to bring it to the front.
4. From the Tasks menu, select Associate Group.
   The New Volume Authentication Group Association window opens, shown in Figure 168.

Figure 168. Creating a New Group Association

5. Select the authentication group you want to associate with the volume or snapshot.
6. Select the permission level for the association.

Table 50. Characteristics of Permission Levels

Type of Access	Allows This
No Access	Prevents the group from accessing the volume or snapshot.
Read Access	Restricts the group to read-only access to the data on the volume or snapshot.
Read/Write Access (not available for snapshots)	Allows the group read and write permissions to the volume.

7. For a Fibre Channel volume, enter the LUN number.

Warning: Before you associate a LUN to an authentication group, the host server's Disk Management (or equivalent) window must be closed.

8. Click OK when you are finished.

The authentication group tab opens with the new association listed, shown in Figure 169.

Figure 169. Viewing New Authentication Group Association

15.11 Editing Permissions

Editing permissions involves changing the access rights for the group.

1. Log into the appropriate management group.
2. Select the volume or snapshot for which you want to edit the permissions.
   The volume or snapshot tab view opens.
3. Click the Authentication Groups tab to bring it to the front.
4. Select the group association to edit.
5. From the Tasks menu, select Edit Permissions.
   The Edit Permissions window opens, shown in Figure 170.

Figure 170. Editing Authentication Group Permissions on a Volume or Snapshot

6. Select the level of access you want for the volume or snapshot (read/write access is not available for snapshots).

Note: For Fibre Channel associations, you can also change the LUN number, if necessary.

7. Click OK when you are finished.
   The Authentication Groups tab opens with the edited permissions displayed, shown in Figure 171.

Figure 171. Viewing the Edited Authentication Group Permissions

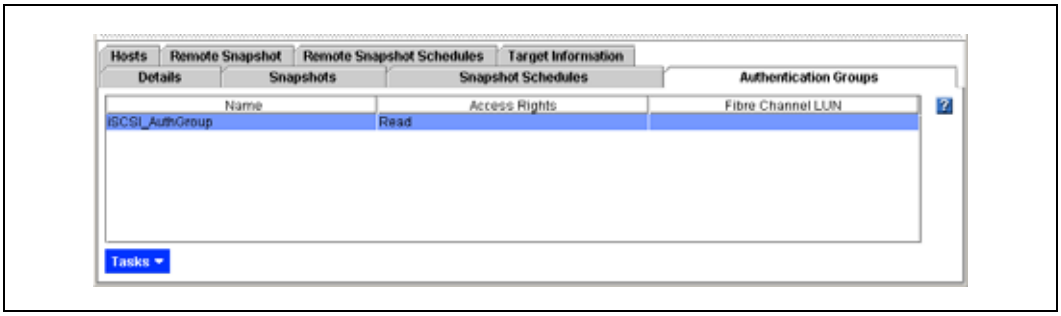

15.12 Deleting an Authentication Group Association

Deleting an authentication group association removes that association from the volume or snapshot and prevents the group from accessing the volume.

Warning: Before you delete an association to a LUN, the host server's Disk Management (or equivalent) window must be closed.

Note: To prevent a group from accessing a volume without deleting the association, change the permissions to "No Access."

1. Log into the appropriate management group.
2. Select the volume or snapshot for which you want to delete the authentication group association.

   The volume or snapshot tab view opens.
3. Click the Authentication Groups tab to bring it to the front.
4. Select the group you want to disassociate.
5. From the Tasks menu, select Disassociate Group.

   A confirmation window opens.
6. Click OK to confirm deleting the group association.

   The Authentication Groups tab opens with the association gone.

Feature Registration 16

16.1 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

All add-on features and applications are available when you begin using the Storage System Software. When you begin using any add-on feature or application, 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.

16.2 Evaluating Features

Add-on features and applications are active and available when you install and configure your system.

16.2.1 30-Day Evaluation Period

When you use any feature that requires registration, a message opens asking you to verify that you want to enter a 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 inaccessible 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 accessibility 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.

16.2.2 Tracking the Time Remaining in the Evaluation Period

Track the time left on your 30-day evaluation period by the management group Register tab, shown in Figure 172 or by the reminder notices that open periodically, as shown in Figure 173.

Figure 172. Evaluation Period Countdown on Register Tab

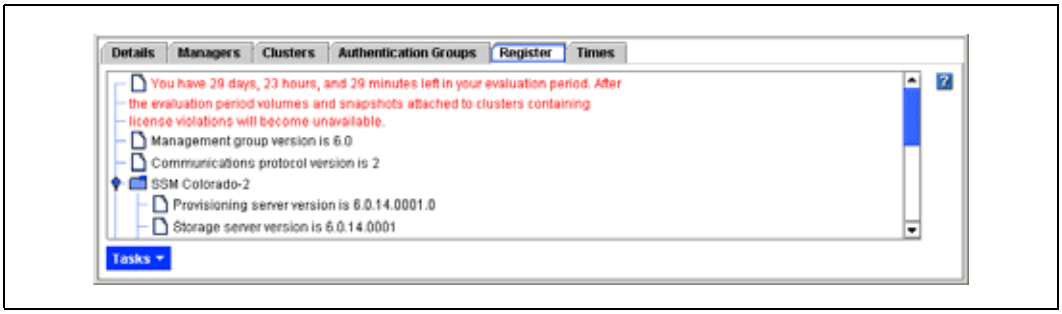

Figure 173. Evaluation Period Countdown Message

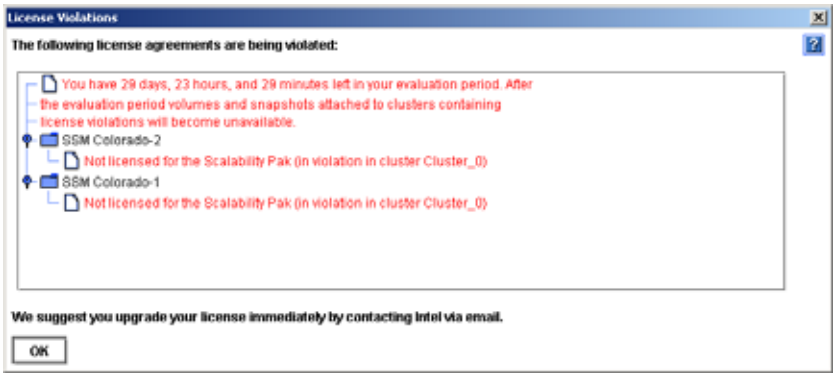

16.3 Evaluating the Scalability Pak

The Scalability Pak includes multi-node virtualization and clustering. Features included are

- multiple nodes in a cluster
- hot spares
- N-way replication
- virtual manager

16.3.1 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, add a hot spare to the cluster, or configure a virtual manager. Please read “Hot Spares Overview” on page 155, “Planning Data Replication” on page 173 and Chapter 10, “Disaster Recovery Using A Virtual Manager,” before working with these features.

16.3.2 Backing Out of the License Evaluation Period

If you decide not to purchase the Scalability Pak, you must remove the additional SSM(s) from the evaluation cluster. 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. First back up any volumes you plan to retain.
   Table 51 describes additional steps to safely back out of the Scalability Pak evaluation.

Table 51. 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 of 1 SSM • Delete the volume • Move the volume to another cluster, or • Add storage to the SSM
2- or 3-way replication	Decrease replication level
Virtual manager	Stop virtual manager

2. Remove the extra SSMs from the cluster.

16.3.3 Evaluating the Configurable Snapshot Pak

The Configurable Snapshot Pak includes programmable snapshots. Features included are

- variable capacity allocation
- scheduled snapshots
- scripting for snapshots

16.3.4 Starting the License Evaluation Period

The Configurable Snapshot Pak 30-day evaluation period begins if you

- Create a snapshot schedule
- Set the snapshot hard or soft threshold to a value that is different than the volume size

Please read “Managing Capacity Using Volume and Snapshot Thresholds” on page 189, and “Creating Snapshot Schedules” on page 196 before working with these features.

16.3.5 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, or reset the hard and soft thresholds. 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. First back up any volumes you plan to retain.
   Table 52 describes additional steps to safely back out of the Configurable Snapshot Pak evaluation.

Table 52. Safely Backing Out of Configurable Snapshot Pak Evaluation

Feature Being Evaluated	Steps to Back Out
Scheduled snapshots	• Delete the snapshot schedule
Variable capacity allocation	• One snapshot of the volume, edit the thresholds on the volume (NOT the snapshot) and set the hard and soft thresholds equal to the volume size. • Two or more snapshots, set the thresholds of the volume and the thresholds of the most recent snapshot to be equal to the volume size.

16.4 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

16.4.1 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 Copy User Manual* before working with these features.

16.4.2 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. First back up any volumes you plan to retain.

   Table 53 describes additional steps to safely back out of the Remote Data Protection Pak evaluation.

Table 53. 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 • Delete the remote volume
Remote snapshots - retaining the data on the remote target	• Make the remote volume into a primary volume • Disassociate the primary and remote management groups, if the remote copy was between management groups.

16.5 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.

Note: Turning off the scripting evaluation ensures that no scripts will continue to run the 30-day evaluation clock unintentionally.

16.5.1 Turn On Scripting Evaluation

To use scripting while evaluating add-on features or applications, enable the scripting evaluation period.

1. Select the management group.
2. Select the Register tab.
3. From the Tasks menu, select Feature Registration.
4. Select the Scripting Evaluation tab, shown in Figure 174.

Figure 174. Enabling Scripting Evaluation

5. Check the box to enable the use of scripts during a license evaluation period.

For more information about scripting, see Chapter 14, “Working with Scripting.”

16.5.2 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. From the Tasks menu, select Feature Registration.
4. Select the Scripting Evaluation tab, shown in Figure 174.
5. Clear the check box.
6. Click OK.

Table 54 describes additional steps to safely back out of the scripting evaluation.

Table 54. 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 Scheduled snapshots Snapshots with hard and soft thresholds different than volume size Remote copy volumes and snapshots Automatic threshold increases	• Back out of any configurable snapshots, scheduled snapshots, or remote copying • Delete any scripts • 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 run the 30-day evaluation clock unintentionally.

16.6 Registering Features and Applications

When registering IXA SDKs 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).

16.6.1 Using License Keys

License keys are assigned to individual SSMs. One license key is issued per SSM and that key licenses all the features requested for that module. 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 the remote copy functionality, 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.

16.6.1.1 Submitting IXA SDK Serial Numbers

First you must submit the serial numbers of all the SSMs that you want to register.

1. Select the management group for which you want to register features or applications.
2. Select the Register tab, shown in Figure 175.

   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.

Figure 175. Registering Features and Applications

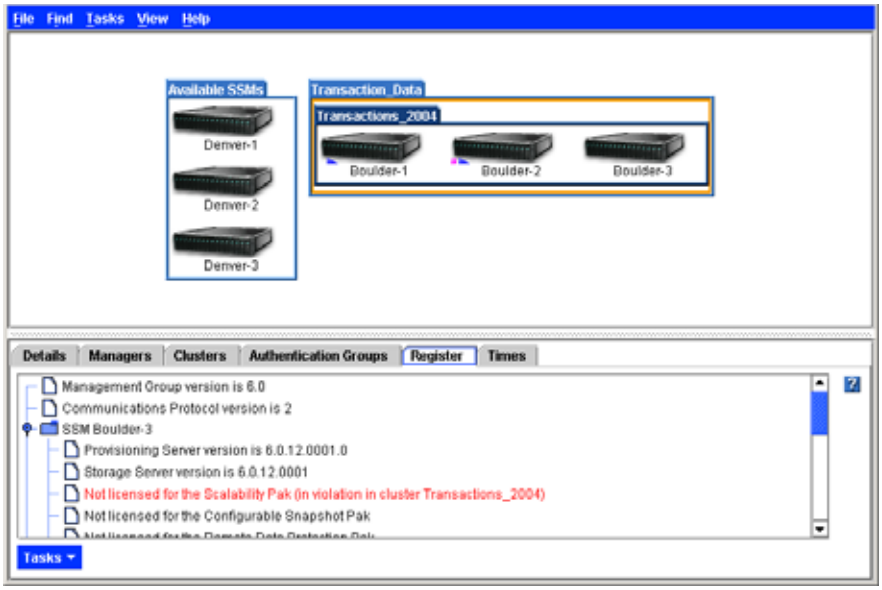

3. From the Tasks menu, select Feature Registration.

   The Feature Registration window opens, shown in Figure 176. Listed are all the SSMs in that management group.

Figure 176. Opening the Feature Registration Window

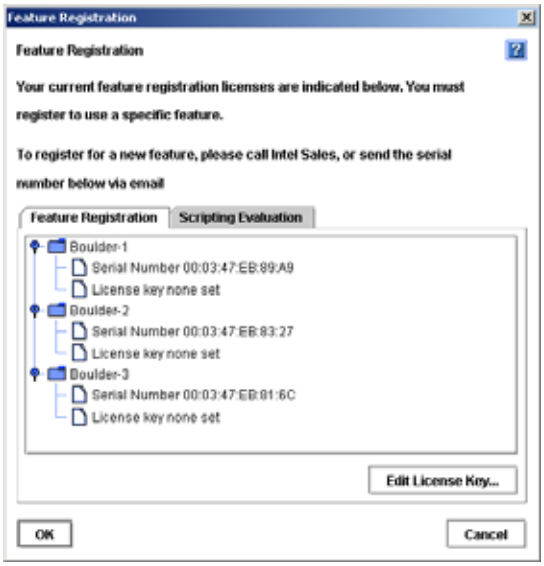

4. For each SSM listed in the window that you want to register, submit the serial number as instructed in the Feature Registration window.

Control + C copies the serial number so that you can paste it into an application such as Notepad or Word.

Note: Record the host name or IP address of the module along with the serial number. This record will make it easier to add the license key to the correct module when you receive it.

16.6.1.2 Entering License Keys

When you receive the license keys add them to the SSMs in the Feature Registration window.

1. Select the management group.
2. Select the Register tab.
3. From the Tasks menu, select Feature Registration.
4. Select an SSM and click Edit License Key.
   The Edit Feature Registration window opens, shown in Figure 177.

Figure 177. Entering License Key

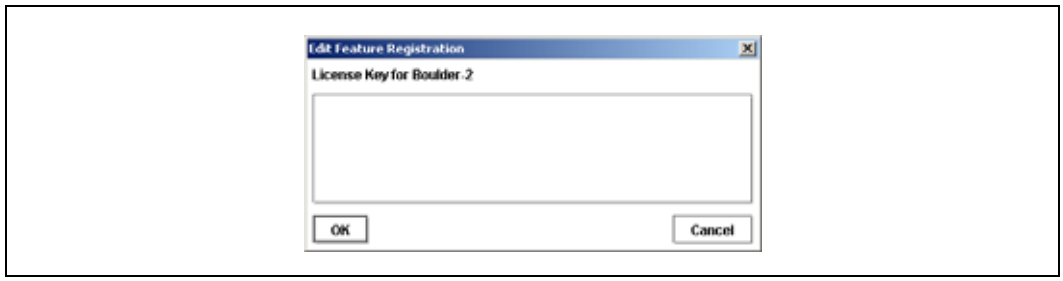

5. Copy and paste the appropriate license key for that SSM into the window.
6. Click OK.
   The license key information is updated in the Feature Registration window, as shown in Figure 178.

Figure 178. Viewing License Keys

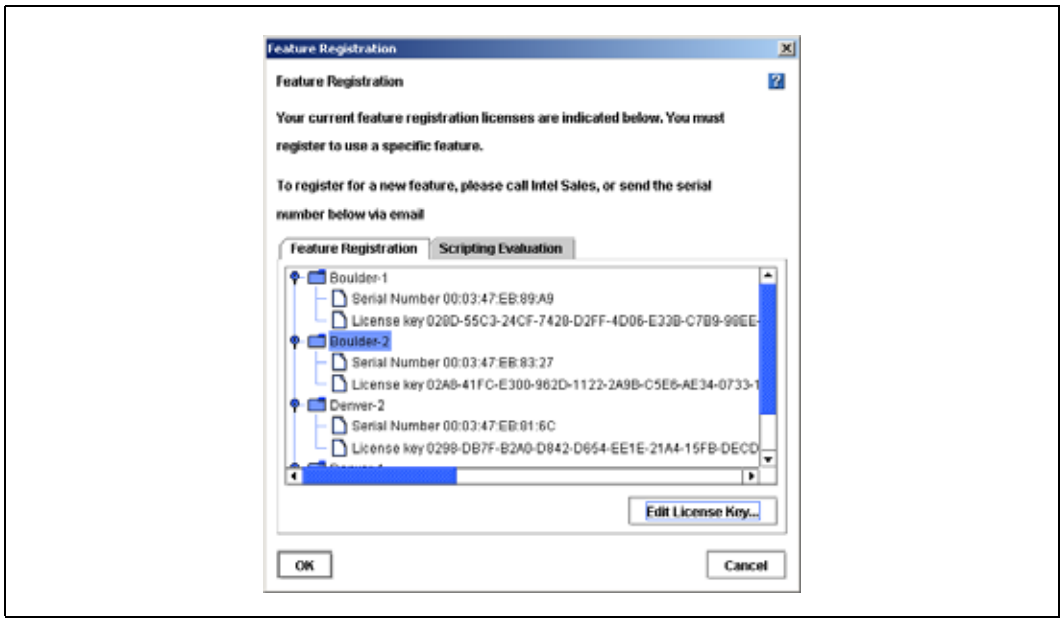

Using the Configuration Interface A

The Configuration Interface is the command line interface that uses a direct connection to configure 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 cluster configuration
- reset the SSM configuration to factory defaults

A.1 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.

A.1.1 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 179.

Figure 179. Opening the Configuration Interface

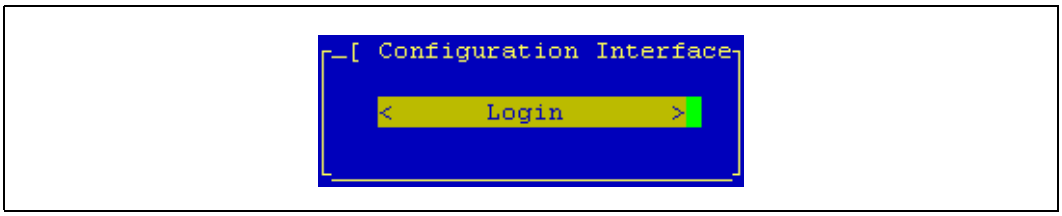

A.1.2 Connecting to the Configuration Interface with Linux/UNIX

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 SSM

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.

6. When the session is connected to the SSM, the Configuration Interface window opens, shown in Figure 179.

A.2 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 Login window opens, shown in Figure 180.

Figure 180. Enter User Name and Password

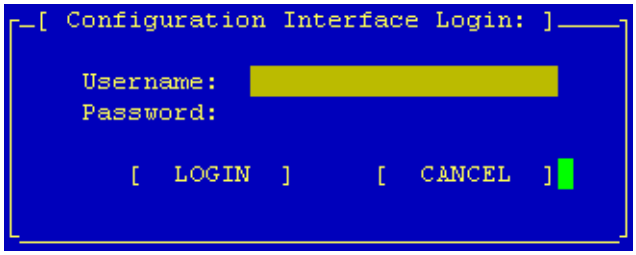

2. Type the user name and password of the administrative user established when the SSM was first configured.

Note: This user is viewable in the Storage System Console under SSM Administration. Click Users and find the admin user on the list.

3. Tab to Login and press Enter.
   The Configuration Interface main menu opens, shown in Figure 181.

Figure 181. Configuration Interface Main Menu

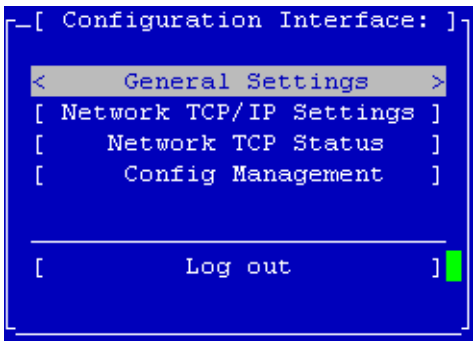

A.3 Configuring Administrative Users

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 182.

Figure 182. General Settings Window

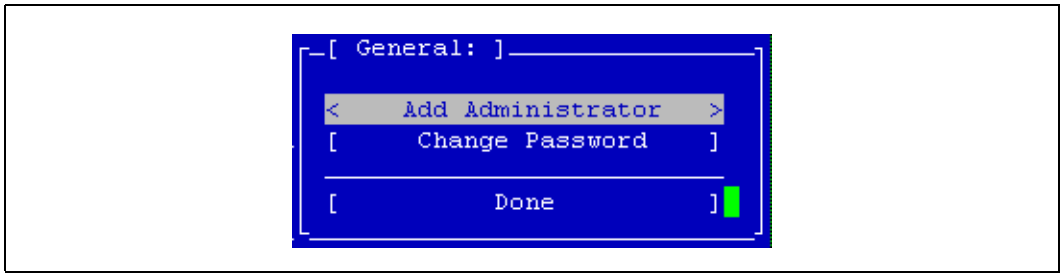

2. To add an administrative user, tab to Add Administrator and press Enter. Then enter the new user's name and password.
3. To change the password for the user that you are currently logged in as, tab to Change Password and press Enter. Then enter the new password.
4. On the General window, tab to Done and press Enter.

A.4 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.

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 183.

Figure 183. Selecting an Interface to Configure

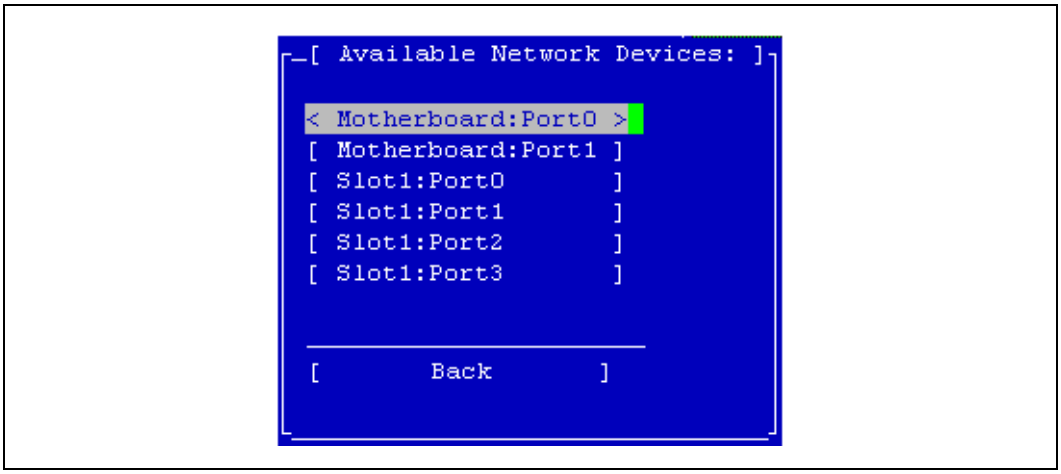

2. Tab to select the network interface that you want to configure and press Enter.

   The Network Settings window opens.

   If the interface you selected is a bond, then the Logical Interface Device window displays first. Click Change Settings, shown in Figure 184 to open the Network Settings window for the bond.

Figure 184. Entering the Host Name and Settings for an Interface

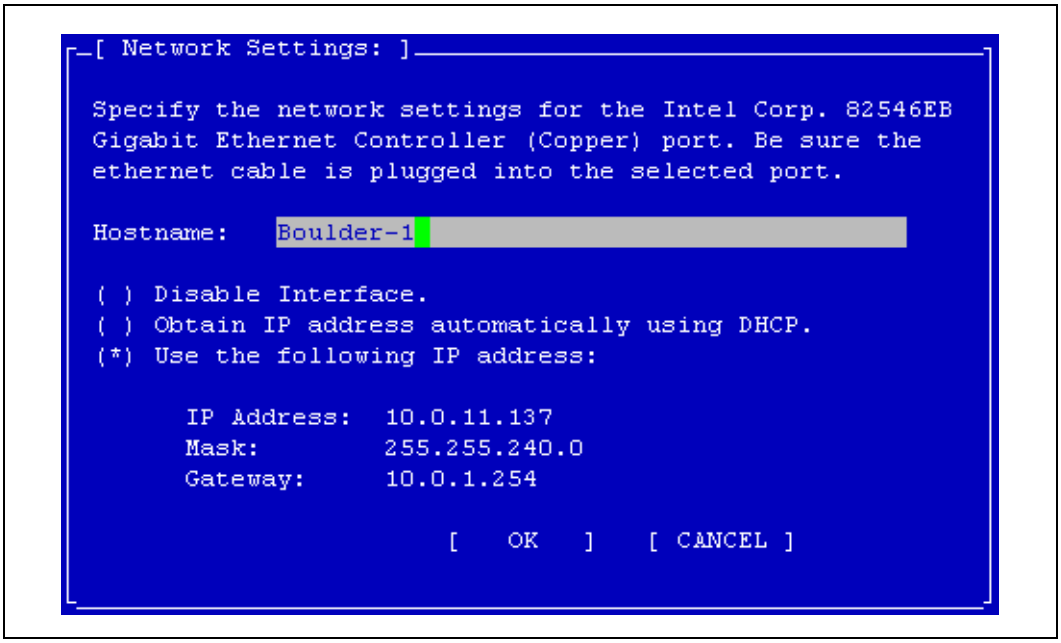

3. Enter the host name and tab to the next section to configure the network settings.

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.

A.5 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 60.

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 185. The logical bond is listed in the window.

Figure 185. Selecting a Bonded Interface in the Available Network Devices Window

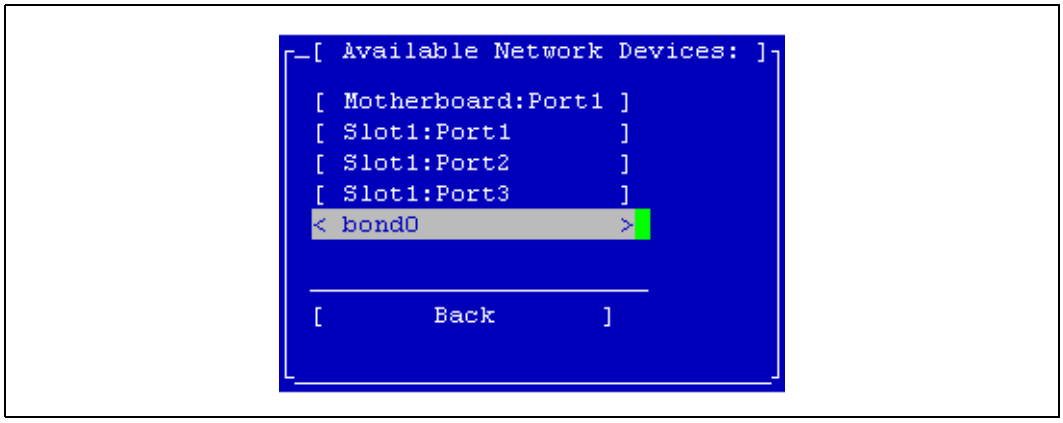

2. Tab to select the bond and press Enter.

   The Logical Failover Device window opens, shown in Figure 186.

Figure 186. Deleting a NIC Bond

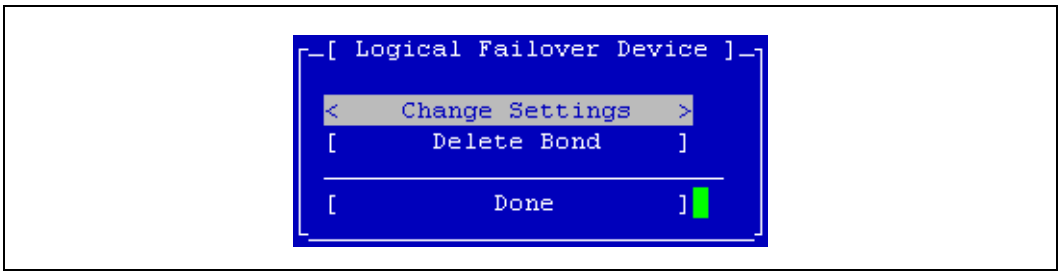

3. Tab to Delete Bond and press Enter.
   A window opens, asking you to confirm the changes.
4. Press Enter.
5. On the Available Network Devices window, tab to Back and press Enter.

A.6 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 74.

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 75.

1. On the Configuration Interface main menu, tab to Network TCP Status and press Enter.
   The Available Network Devices window opens, shown in Figure 187.

Figure 187. Available Network Devices Window

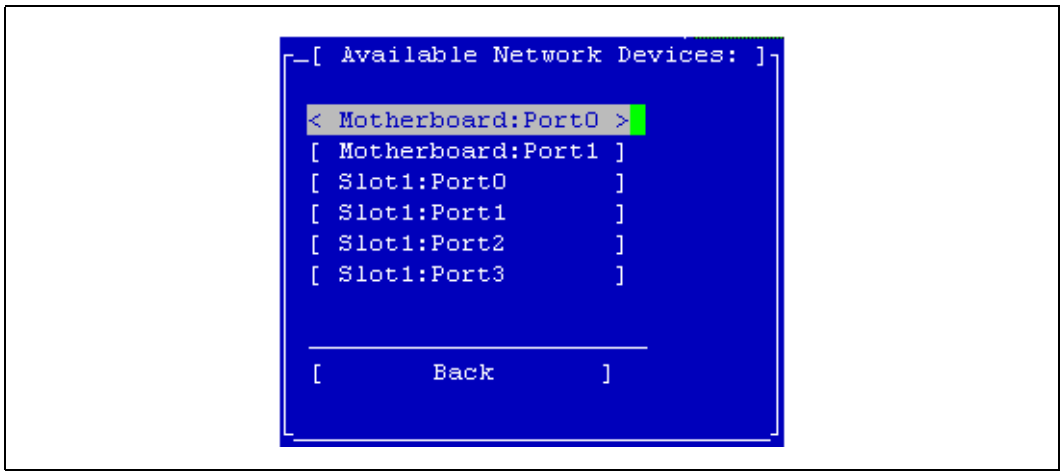

2. Tab to select the network interface for which you want to set the TCP speed and duplex and press Enter.

   The Network TCP Status window opens, shown in Figure 188.

Figure 188. Setting the Speed, Duplex, and Frame Size

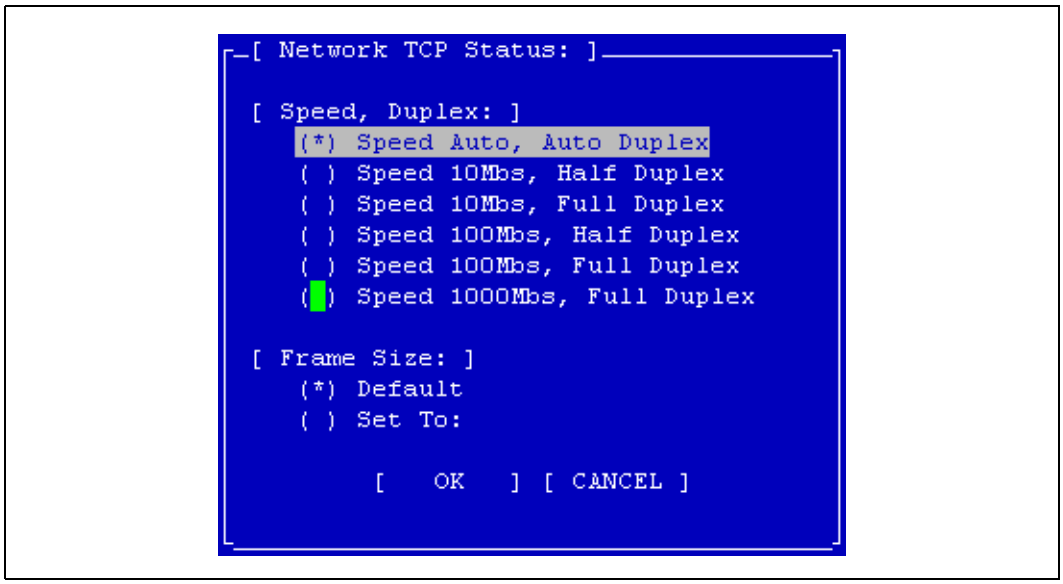

3. To change the speed and duplex of an interface, tab to a setting in the Speed / Duplex list.
4. To change the frame size, select Set To in the Frame Size list. Then tab to the field to the right of Set To and type a frame size.

   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.

A.7 Removing the SSM from a Management Group

If an SSM is part of a management group, you can use the Configuration Interface to delete all data on the SSM and remove it from the management group in one step.

Warning: Removing the 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 General window opens, shown in Figure 189.

Figure 189. Removing the SSM from a Management Group

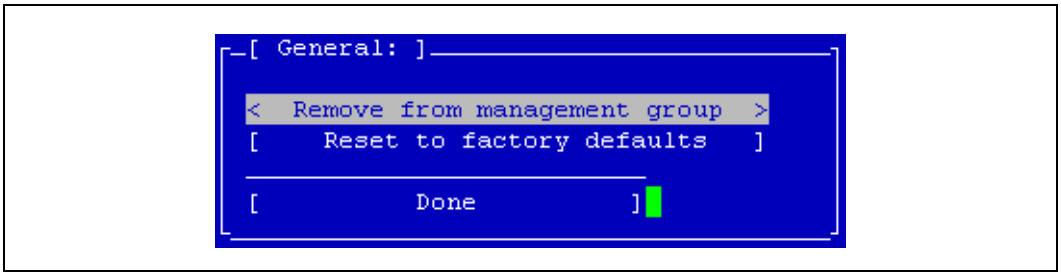

2. Tab to Remove from Management Group and press Enter.

   A window opens, warning you that removing the module from a management group will delete all data on the SSM.

3. Press Enter.
4. On the General window, tab to Done and press Enter.

A.8 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 General window opens, shown in Figure 190.

Figure 190. Resetting the SSM to Factory Defaults

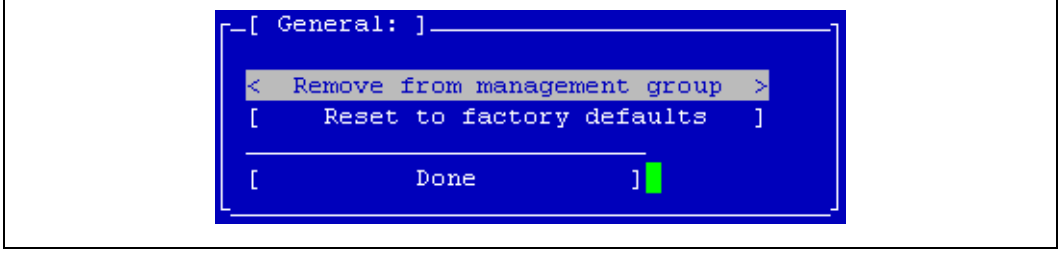

2. Tab to Reset to Factory Defaults and press Enter.

   A window opens, warning you that resetting the SSM configuration will delete all data on the SSM.

3. Press Enter.
4. On the General window, tab to Done and press Enter.

SNMP MIB Information B

B.1 SNMP Agent

The SNMP Agent resides in the SSM. 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.

B.2 The Supported MIBs

- MIB II
- Host Resources MIB
- UCD Extensions MIB
- SNMPv3 MIB

B.2.1 Exceptions

B.2.1.1 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)

B.2.1.2 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)

B.2.1.3 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)

B.2.1.4 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)

Using the EBSD Driver for Windows 2000 C

C.1 Recommended Configuration

- Windows 2000 with Service Pack 2 or above
- Pentium III processor or greater
- 1 GB minimum RAM
- 10 MB minimum free hard disk space

C.2 Overview of EBSD Driver for Windows 2000

Install and configure the EBSD driver for Windows® 2000 on any computer that accesses volumes on a cluster of SSMs.

To configure a client to access a volume, you must install and configure the driver and configure the volume to be accessed. This manual describes all of the driver configuration tasks. Volume configuration tasks are described in the User Manual in the chapter entitled Working with Volumes. The table below lists the required tasks and where to find information about them.

Table 55. EBSD Driver Configuration Tasks

Configuration Task	Configuration Tool	Instructions
Create a volume and associate the volume with an authentication group.	Storage System Console	In the User Manual • Working with Volumes chapter, see “Creating a Volume” • Working with Authentication Groups chapter, see “Creating an Authentication Group” and “Creating an Authentication Group Association”
Install the EBSD Driver on a Windows computer	EBSD Driver	“Installing the EBSD Driver”
Create a disk on the EBSD client computer.	EBSD Driver	“Adding EBSD Disks to Your System”
Write the disk signature.	Windows® Disk Manager on client computer	“Enabling Write Cache on Volumes”
Partition a basic disk.	Windows Disk Manager on client computer	“Partitioning Basic EBSD Disks”

C.3 Installing or Updating the EBSD Driver

C.3.1 Installation Overview

The EBSD for Windows installation CD provides two installation options for the EBSD driver.

- Install a driver on a local machine directly from the CD. If you have an earlier version of the driver already installed on the local machine, the installation wizard directs you to the location where you can update the driver. See "Updating the EBSD Driver".
- Copy the EBSD driver files to a local machine or a network share. Then install the driver individually on local machines using the Windows Add/Remove Hardware wizard.

Note: You need administrative privileges during installation and configuration.

Note: Throughout the procedures in this guide, the term "EBSD client" refers to any computer, such as an application server, that accesses volumes on a cluster of SSMs. The term "disk" refers to the EBSD disks that you create on the EBSD client computer, while the term "volume" refers to the volumes created on Storage System Software using the Storage System Console and to any volumes that are created on EBSD disks using the Windows Disk Management tool.

C.4 Copying the EBSD Driver Files [Optional]

Skip this section if you only want to install or update the EBSD driver from the CD. Go directly to "Installing the EBSD Driver", or "Updating the EBSD Driver".

1. Insert the resource CD that came with your system into the CD drive of the EBSD client PC.

   The installation wizard should automatically open. If the installation wizard does not open, navigate to the CD drive, the InstData\VM\ folder, and run EBSD**60_setup.exe**.

2. On the Choose Product window of the installation wizard, select EBSD Driver Bundles and click Next.

   The Choose Install Folder window opens.

3. Choose to accept the default directory (C:\ProgramFiles\Storage_System\Storage_System_Software\6.0\Drivers) or browse to the directory where you want the EBSD folder to be copied. This location can be another location on the network, such as a file server.

4. Click Next.

   Review the Pre-installation Summary window.

5. Click Install.

   The folder, EBSD, is copied into the directory you specify.

C.5 Installing the EBSD Driver

You can install the EBSD driver locally using the CD or using the folder from a network share. See "Copying the EBSD Driver Files [Optional]" on page 252.

The installation of the EBSD driver includes the installation of the Microsoft® Windows 2000 DiskPart utility which is used by the driver for volume management. The DiskPart utility has its own installation wizard which runs as part of the EBSD driver installation.

C.5.1 Beginning the Driver Installation

Table C.1.

Installing from the CD	Installing from a Network Share
1. Insert the EBSD for Windows CD into the CD drive of the EBSD client PC. The installation wizard should automatically start. If not, run EBSD**60_setup.exe** from the InstData\VM\ folder on the CD. 2. On the Choose Product Component window, select EBSD Driver. 3. After reviewing the Pre-installation Summary window and clicking Next, the DiskPart.exe installation begins. 4. After completing the installation of the DiskPart Utility, a message opens describing how to locate the driver using the Add/Remove Hardware wizard. 5. Click OK. The Add/Remove Hardware wizard opens. 6. Continue with Step 1 below.	1. Copy the EBSD folder to the computer on which you want to install the driver. 2. From the EBSD folder, run **DiskPart_setup.exe**. 3. Open the Add/Remove Hardware wizard from the Control Panel. 4. Continue with Step 1 below.
Note: The DiskPart Setup wizard copies files to the folder you specify and the Driver installation wizard also copies the file diskpart.exe in the system PATH.	

Figure 191. Opening the Add/Remove Hardware Wizard

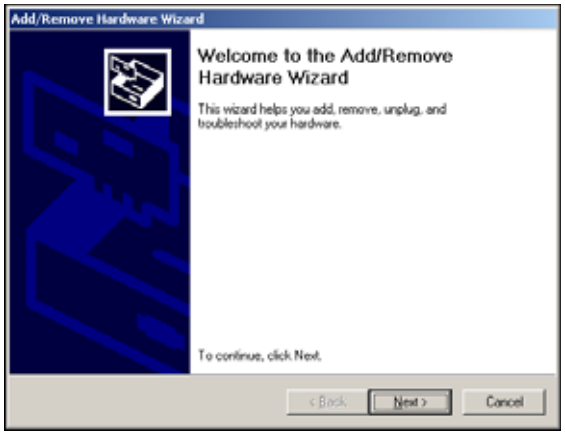

1. On the Welcome to the Add/Remove Hardware wizard window, click Next.
   The Choose a Hardware Task window opens, shown in Figure 192.

Figure 192. Choosing the Hardware Task

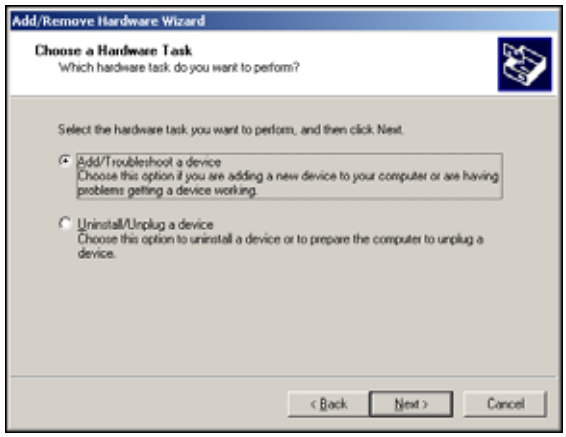

2. Select Add/Troubleshoot a Device and click Next.

   The wizard searches for new Plug and Play hardware. It should find nothing to install and then open the Choose a Hardware Device list, shown in Figure 193.

Figure 193. Adding New Device

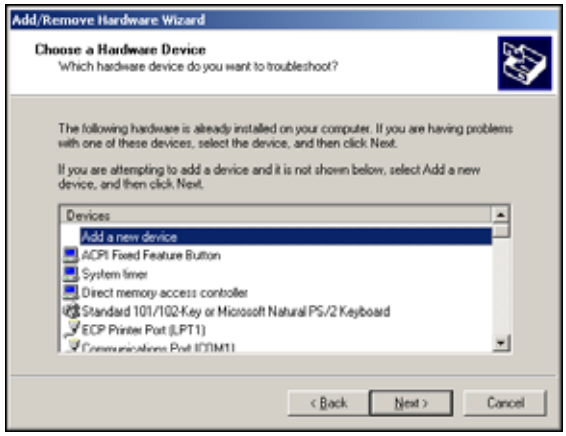

3. Select Add a new device and click Next.

   The Find New Hardware window opens, shown in Figure 194.

Figure 194. Choosing "Select Hardware from List"

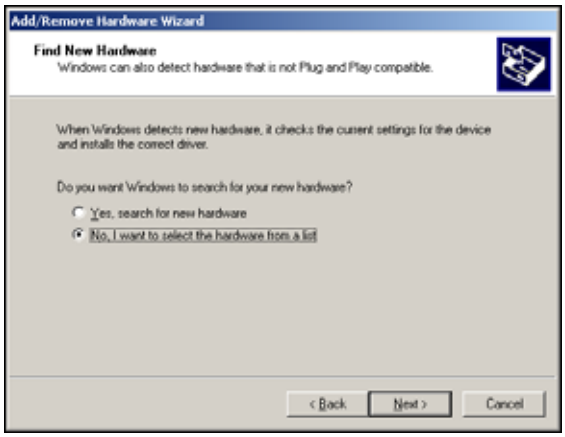

4. Select "No, I want to select the hardware from a list," and click Next.
   The Hardware Type window opens, shown in Figure 195.

Figure 195. Choosing "Other Devices"

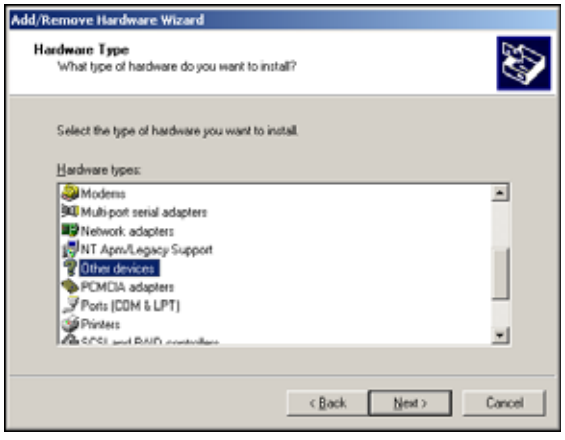

5. Scroll down and select "Other devices" from the list, shown in Figure 195, and click Next.
   The Select Device Driver window opens, shown in Figure 196.

Figure 196. Selecting a Device Driver

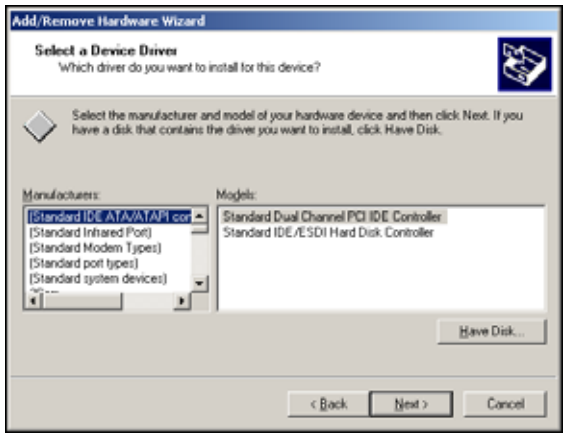

6. Click Have Disk.

   The Install From Disk window opens, shown in Figure 197.

Figure 197. Install from Disk Window

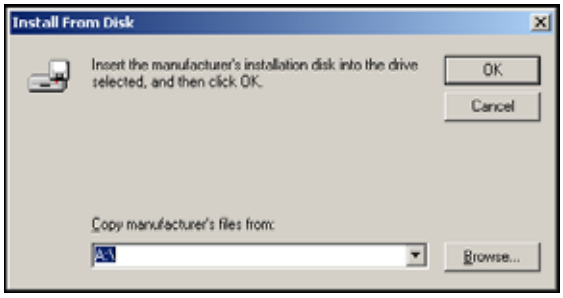

C.5.2 Locating the EBSD Driver Files

1. Browse to the C:\Program Files\Storage_System\Storage_System_Software\6.0\Drivers\EBSD folder where the **aebs.inf** file is located.

   or

   You can install the driver from any directory into which you have copied the files.

2. Select the **aebs.inf** file and click Open.

   Focus returns to the Install From Disk window with the directory containing the selected file displayed.

Figure 198. Install from Disk Window

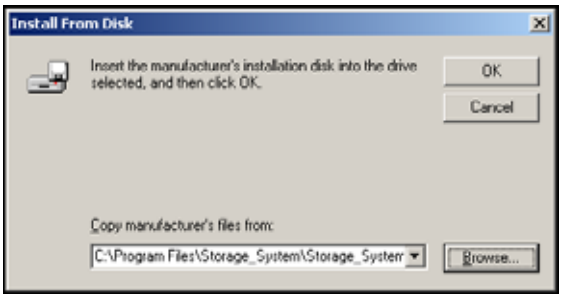

3. Click OK.
   Focus returns to the Select a Device Driver window with the driver displayed.

Figure 199. Selecting the Driver

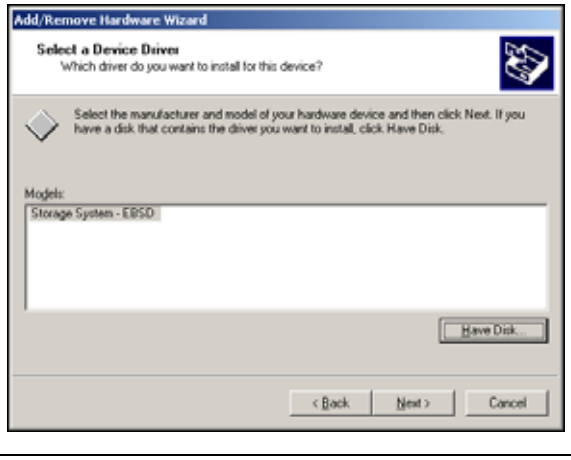

4. Select the EBSD driver and click Next.
   The Start Hardware Installation window opens, shown in Figure 200.

Figure 200. Verifying the Driver

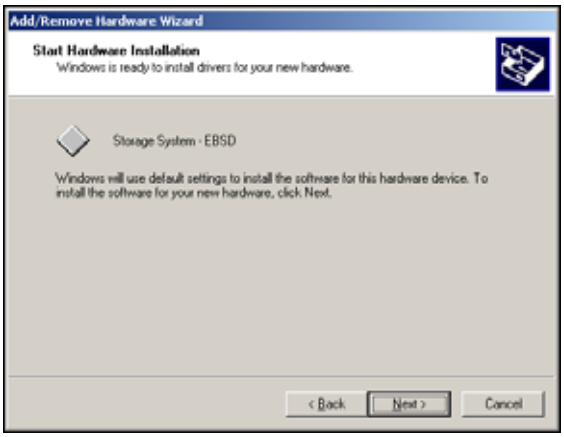

5. Verify that the EBSD driver appears in the window and click Next.

Note: You may see a message regarding digital signatures not found. Click Yes to continue the installation.

The EBSD System Configuration window opens, shown in Figure 201.

Figure 201. Configuring File Server, SQL Server, and Exchange Services to Come Online after a Reboot

6. [Optional] If you want to automatically configure services and applications to come online after a reboot, select the appropriate boxes.

   You can change these settings later in the EBSD driver advanced settings. See “Configuring Services for File Server, MS SQL and MS Exchange” on page 283.

   The Finish EBSD installation window opens. Instructions for modifying the driver configuration are on the window.

Figure 202. Finishing Installation and More Instructions

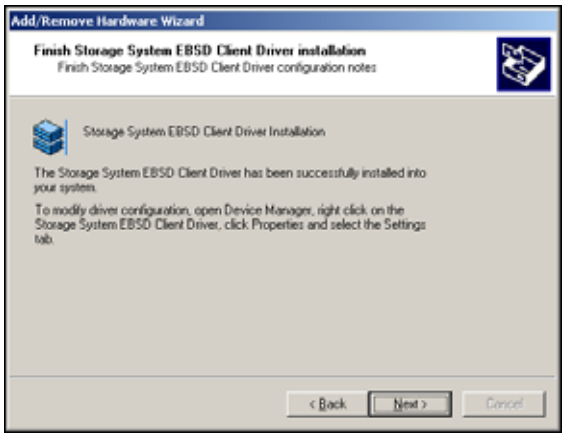

7. Click Next.

   The Completing the Add/Remove Hardware window opens. Review the window to verify that it is the driver that is being installed.

Figure 203. Completing the Add/Remove Hardware Wizard

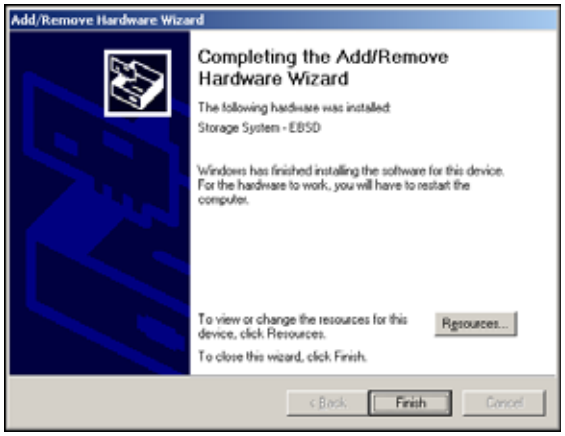

C.5.3 Restarting Windows to Apply Settings

1. Click Finish.

   A message opens notifying you to restart your computer in order for the settings to take effect.

2. Click No.
3. Complete the steps in the Driver Installation wizard.
4. Manually reboot your computer to apply the driver settings.

   See “Configuration Overview” for information about configuring the driver.

C.6 Updating the EBSD Driver

If you have an earlier version of the EBSD driver installed on your system, the installation wizard directs you to the Windows Device Manager where you can update the driver.

Table C.1.

Updating from the EBSD for Windows CD	**Updating from your vendors web or FTP site (if applicable).**
1. Insert the EBSD for Windows® into the CD drive of the EBSD client PC. Your browser should automatically open and start the installation wizard. If not, run EBSD**60_setup.exe** from the InstData\VM folder on the CD. 2. On the Choose Product Component window, select EBSD Driver. 3. Click Next. Review the Pre-installation Summary window. 4. Click Install. The files are installed to the location you specified and a message opens. 5. Review the message and click OK. The message displays the location of the driver update files. Those files are in the C:\Program Files\Storage_System\ Storage_System_Software\6.0\Drivers\ EBSD folder. The Update Installation wizard then opens the Windows Device Manager. 6. Continue with step 1 below.	1. Using a web browser, open your vendors web or FTP site where the files are stored. 2. Open the folder for the release version that you are upgrading to. 3. Copy install.htm and the InstData folder to your hard drive. 4. Navigate to install.htm on your hard drive. 5. Double-click install.htm. 6. Complete the installation wizard. 7. Continue with step 1 below.

Note: If the DiskPart utility has not been installed in the system PATH, run **DiskPart_setup.exe**. See "Installing the EBSD Driver" on page 252.

Note: Install the update in the same directory as the original EBSD driver installation. The default directory is C:\Program Files\Storage_System\Storage_System_Software\6.0\Drivers\EBSD.

C.6.1 Updating the Device Driver in the Windows 2000 Device Manager

1. Expand the Network Storage Devices and select the EBSD driver, shown in Figure 204.

Figure 204. Selecting the EBSD Driver

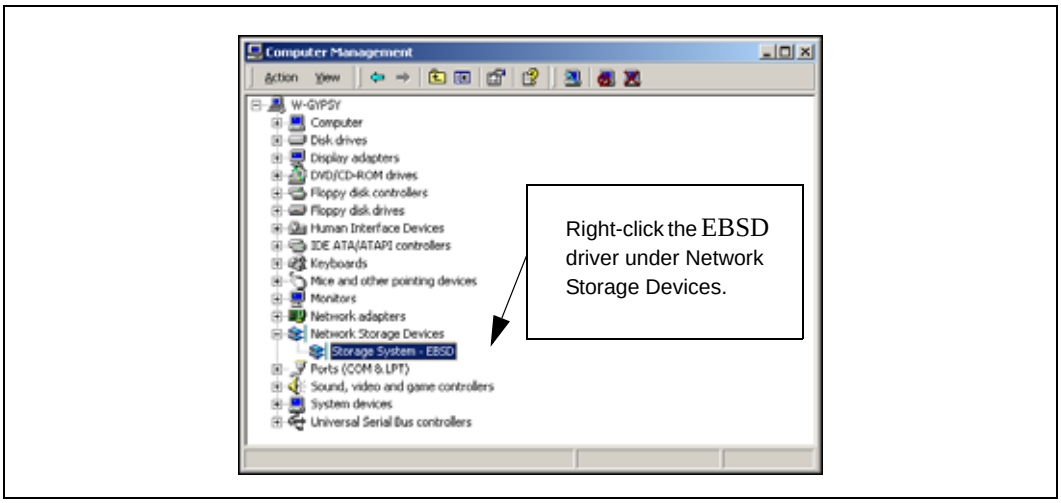

2. Right-click on the driver and select Properties.
   The EBSD Properties window opens.
3. Select the Driver tab to bring it to the front.

Figure 205. Updating the EBSD Driver

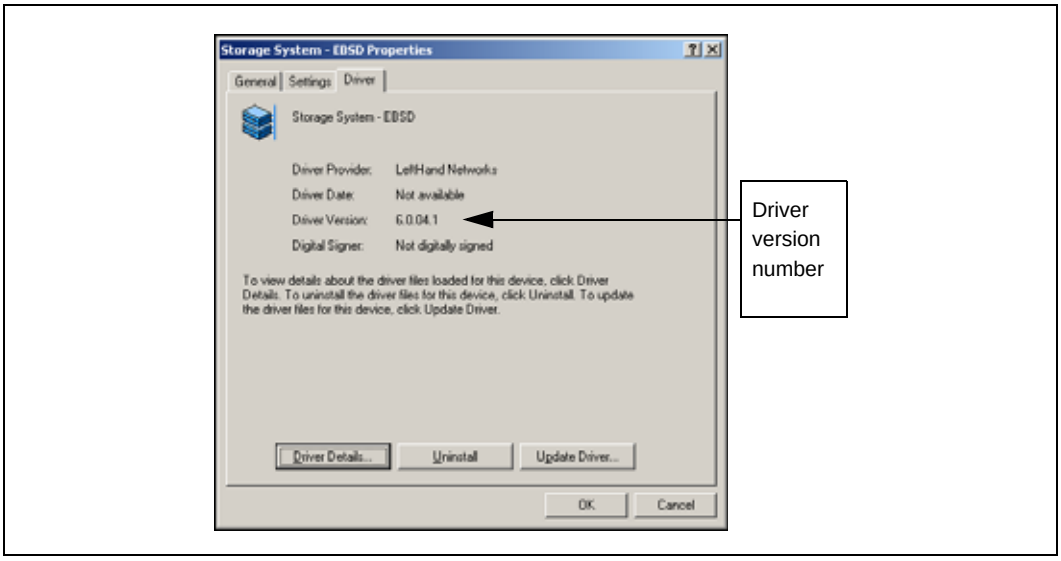

Note: The driver version number is displayed on the Driver tab.

4. Click Update Driver.
   The Upgrade Device Driver wizard opens.

Figure 206. Updating the Driver

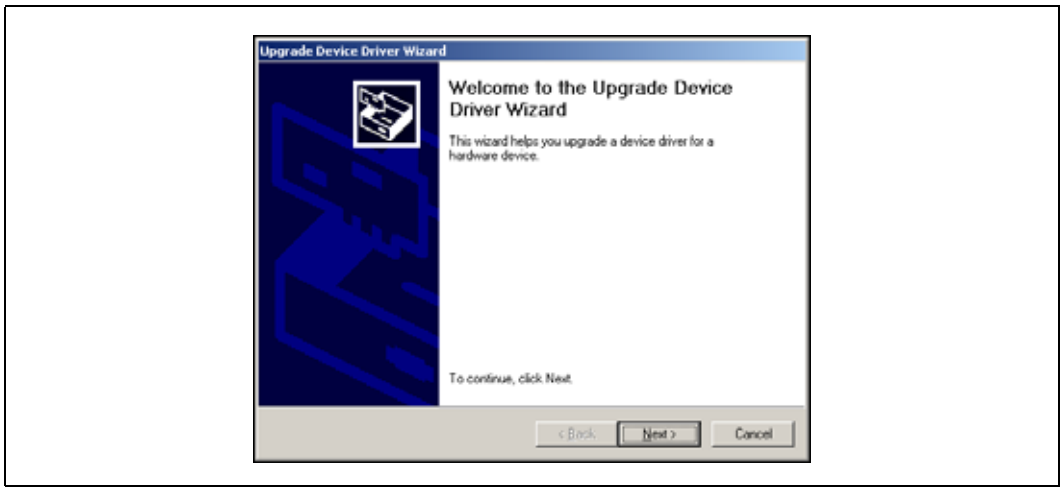

5. Click Next.
   The Install Hardware Device Drivers window opens.

Figure 207. Installing Hardware Device Drivers

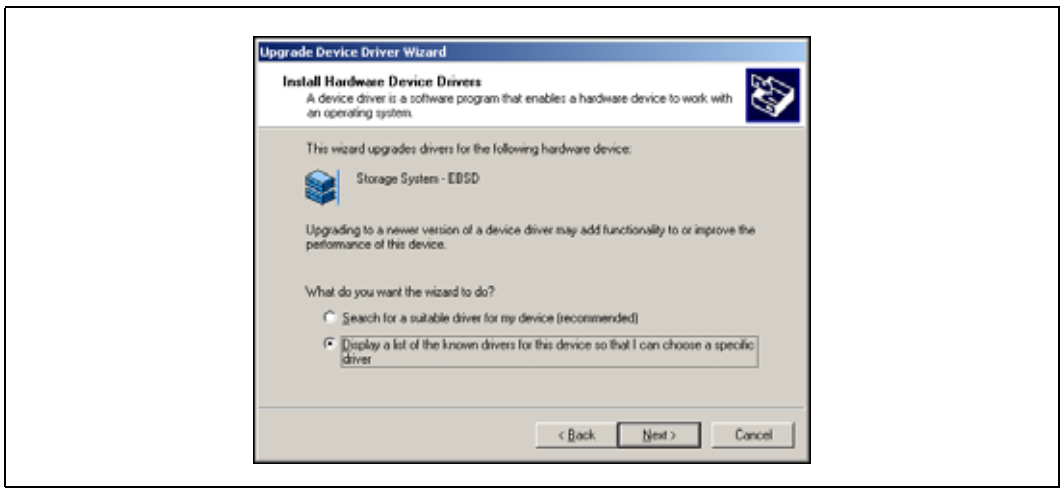

6. Select "Display list of known drivers. . ." and click Next.
   The Select a Device Driver window opens with the EBSD Driver displayed.

Figure 208. Selecting the Device Driver

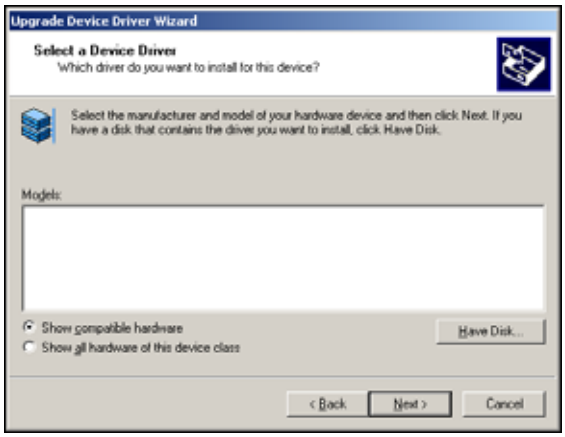

7. Click Have Disk.
   A browse window opens where you can specify the location of the driver files.

Figure 209. Browsing for the Driver Files

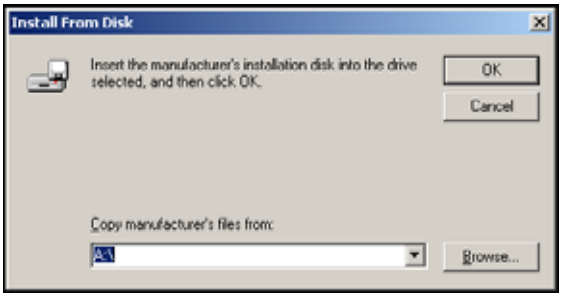

8. Click Browse to navigate to the location where you installed the EBSD driver update files.
   The files are in
   C:\Program Files\Storage_System\Storage_System_Software\6.0\Drivers\EBSD.
9. Select the **aebs.inf** file and click Open.
   Focus returns to the Install from Disk window.

Figure 210. aebs.inf file Selected

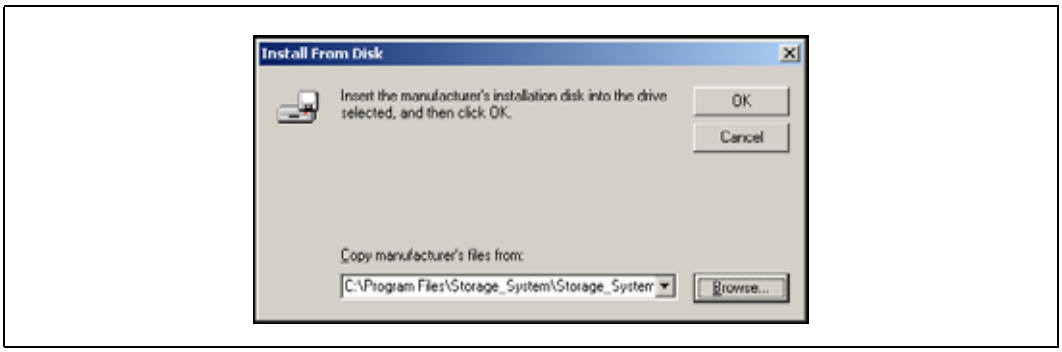

10. Click OK.
    Focus returns to the Select a Device Driver window.

Figure 211. Selecting the Device Driver

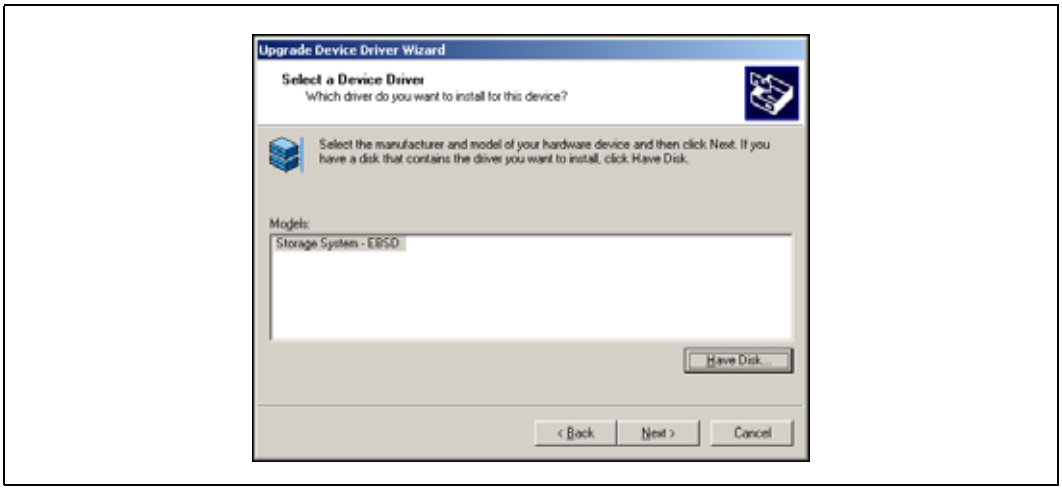

11. Select EBSD and click Next.
    The Start Device Driver Installation window opens.

Figure 212. Starting the Update Installation

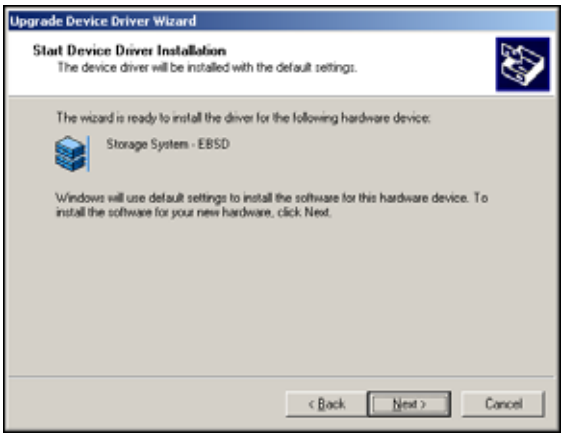

12. Verify that the EBSD driver appears in the window and click Next.

Note: You may see a message regarding digital signatures not found. Click Yes to continue the installation.

The EBSD System Configuration window opens, shown in Figure 201.

Figure 213. Starting File Server, SQL Server, and Exchange Services

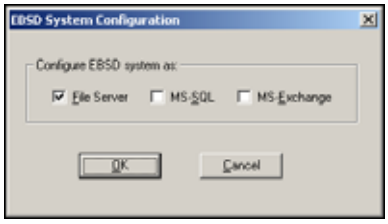

13. [Optional] If you want to automatically configure services and applications to come online after a reboot, select the appropriate boxes.
14. You can change these settings later in the EBSD driver advanced settings. See “Configuring Services for File Server, MS SQL and MS Exchange” on page 283.

    The Finish Device Driver Installation window opens.

Figure 214. Finishing the Update Installation

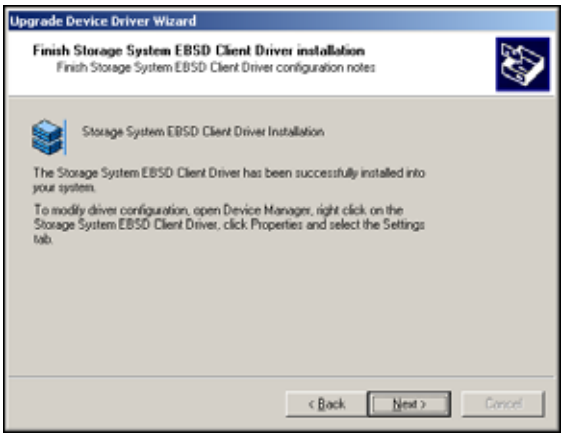

15. Click Next.

    The Completing the Upgrade Device wizard window opens.

Figure 215. Completing the Upgrade Wizard

16. Verify that it lists EBSD and click Finish.

    Focus returns to the EBSD Properties window.

Figure 216. Closing the EBSD Driver

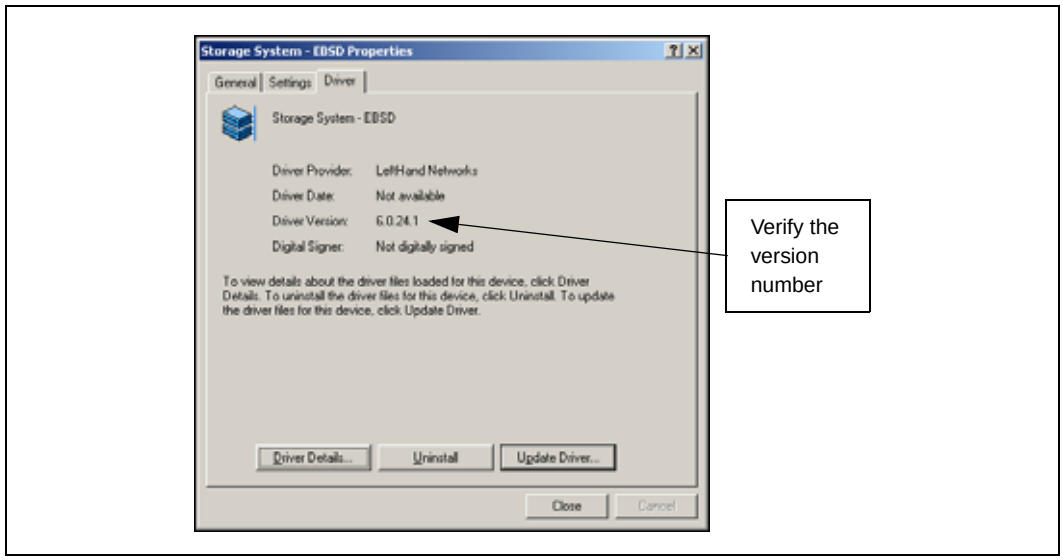

17. Verify that the version number has changed to reflect the upgrade version.
18. Click OK to close the EBSD driver.
    A message opens, prompting you to reboot the computer.
19. Click Finish.
    A message opens notifying you to restart your computer in order for the settings to take effect.
20. Click No.
21. Complete the steps in the Driver Installation wizard.
22. Manually reboot your computer to apply the driver settings.
    See "Configuration Overview" for information about configuring the client drive.

C.7 Configuring the EBSD Driver

C.7.1 Configuration Overview

Once the EBSD driver is installed, it must be configured.

Note: You need administrative privileges during installation and configuration.

Note: Configure volumes and associate them with authentication groups in the Storage System Console before configuring the EBSD driver. You use information about the volumes, including the

management group configuration, volume name, and authentication group name, to configure EBSD disks.

C.8 Opening the EBSD Driver

1. Open Windows Device Manager.
   A list of devices opens in the right hand pane, shown in Figure 217.

Figure 217. Selecting the EBSD Driver

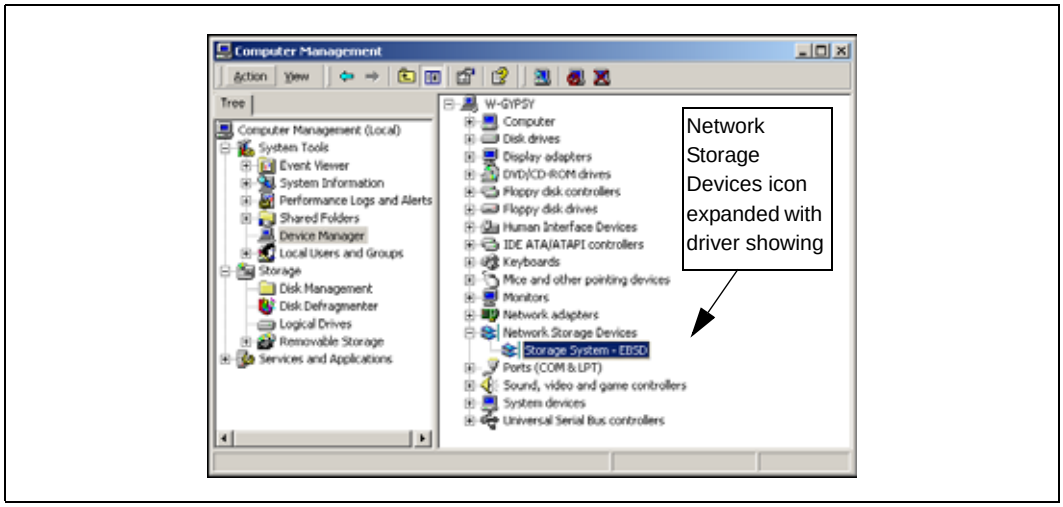

2. Expand the Network Storage Devices list and select the EBSD driver.
3. Double-click on the driver or click the Action menu and select Properties.
   The EBSD Properties window opens, shown in Figure 218.

Figure 218. EBSD Properties Dialog

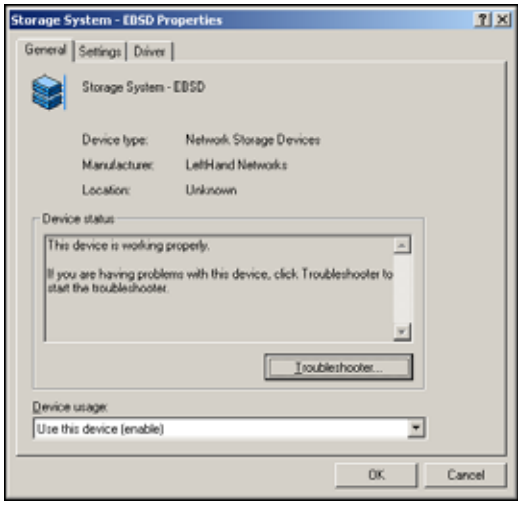

C.9 Adding EBSD Disks to Your System

An EBSD disk maps to a volume on the SSM. Before you begin adding EBSD disks, use the Storage System Console to locate each volume the EBSD client will access and write down the following information.

- Management group name
- IP addresses of all managers in the management group
- Volume name
- Name of the authentication group that is associated with the volume

When you create a volume on the EBSD client, you will need to enter this information exactly as it appears in the Console.

1. Click the Settings tab to bring it to the front, shown in Figure 219.

Figure 219. EBSD Driver Settings

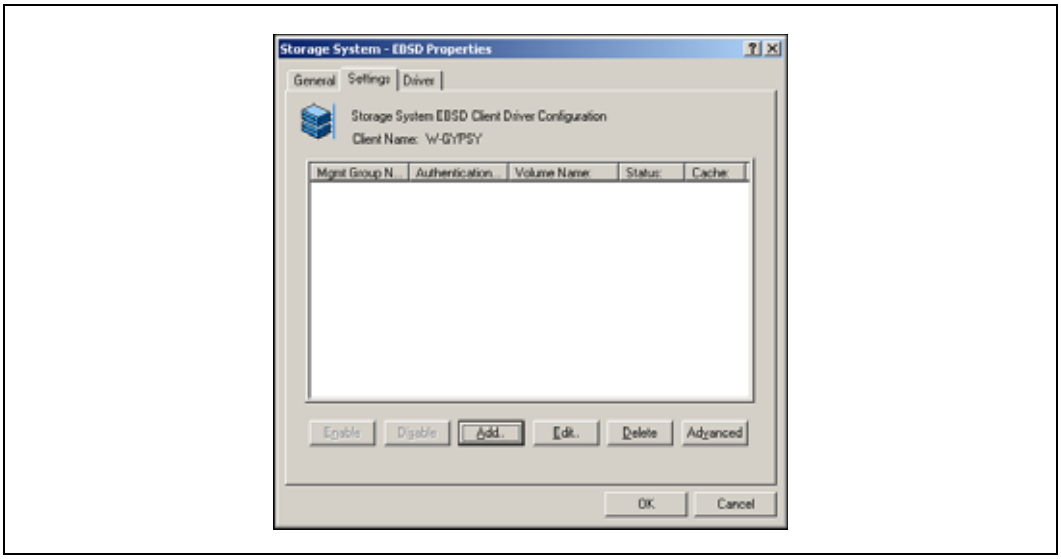

2. Click Add to add EBSD disks.
   The Add EBSD Volume window opens, shown in Figure 220.

Figure 220. Adding an EBSD Disk

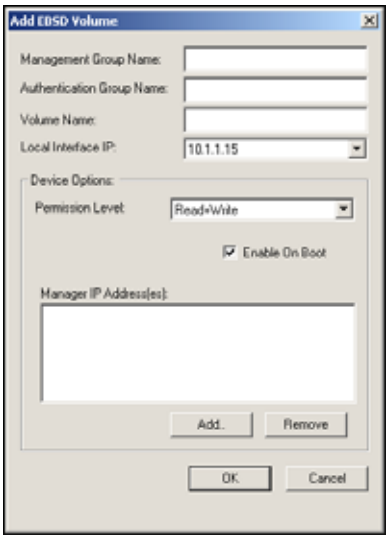

Use the Add EBSD Volume window to create a disk that corresponds to a volume that exists on a cluster of SSMs.

3. Complete the fields in the Add EBSD Volume window.

   SeeTable 56 for a list of field descriptions and requirements for completing the Add EBSD Volume window

Note: Be sure to type the names of the management group, authentication group, and volume exactly as they appear in the Storage System Console.

Table 56. Requirements for Adding an EBSD Disk

Add EBSD Disk Field	Description and Requirements
Management Group Name	Type the name of the management group that contains the volume.
Authentication Group Name	Type the name of the authentication group that is associated with the volume in the Storage System Console.
Volume Name	Type the name of the volume exactly as it appears in the Storage System Console.
Local Interface IP	Select from the list the IP address of the computer that is running the EBSD driver. If the client computer has more than one NIC, select the IP address of the NIC that you want the client to use to access the volume.

Table 56. Requirements for Adding an EBSD Disk (Continued)

Add EBSD Disk Field	Description and Requirements
Permission Level	Default = Read+Write Select the permission level of the EBSD disk. The permission level for the disk cannot be greater than the permission level of the authentication group associated with the volume in the Storage System Console. For example, if the authentication group has read only permissions, you cannot give the EBSD disk read and write permission.
Enable on Boot	Default = Enabled (checked) Ensures that volumes come online after a reboot. ***Note:*** Change this to Disabled only if you are using 3rd party clustering software such as Veritas Cluster Server or Microsoft Cluster Server.
Manager IP Addresses	Enter the IP address of at least one of the managers in the management group containing the volume. ***Note:*** To ensure volume availability, enter the IP addresses of all managers in the management group. If the manager the EBSD client is using to access the volume becomes unavailable, the EBSD client can use any of the other managers to access the volume.

4. Click OK when you have finished completing the EBSD disk information.

   The new EBSD disk appears in the list on the Settings tab, shown in Figure 221. The Status column displays Starting, and then changes to Online when the new disk is ready.

Figure 221. Listing of EBSD Disks

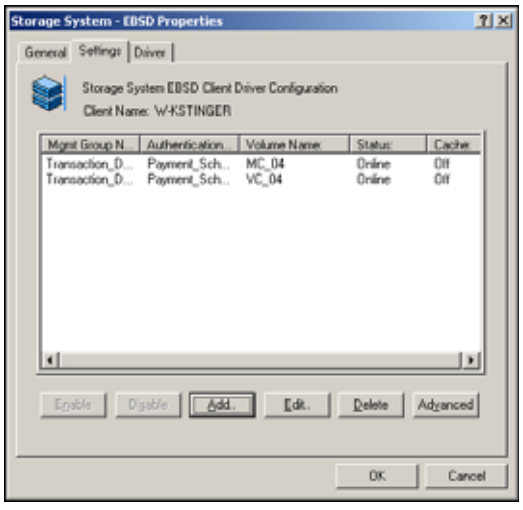

5. Repeat steps 2 through 4 for each EBSD disk you want to add.
6. Click OK when you are finished.

The EBSD Properties window closes.

C.10 Enabling Write Cache on Volumes

Write cache is disabled by default on EBSD volumes. The cache status is shown in the Driver Settings window as “on” or “off.”

Write cache can be enabled on read+write volumes. It cannot be enabled on read only volumes or snapshots. Also, if the system does not have enough memory resources, write cache cannot be enabled and an event log message will be recorded.

C.10.0.1 The Write Through Command

The EBSD driver automatically supports the Write Through command when write cache is enabled. This command, which is set by applications (for example, SQL Server), provides an additional level of safety by causing certain critical packets of data to be written immediately to the disk, bypassing the cache.

C.10.1 Requirements for Changing Write Cache

- All data transfers to or from the volume must be completed before enabling or disabling write cache. You cannot change the write cache setting when data is being written.
- The system should have at least 256MB free RAM when you enable write cache. Verify the amount of available RAM in Windows Task Manager on the Performance tab.

Note: When enabling write cache, standard UPS power protection to prevent possible data loss is recommended.

C.10.2 Enabling Write Cache

1. Open Windows Device Manager.
2. Expand the Disk Drives list on the right.

   The EBSD disks are listed as shown in Figure 222. See “Identifying the Storage System Software Volume That Corresponds to an EBSD Disk” if you need to verify which disk you are working with.

Figure 222. Viewing EBSD Disks

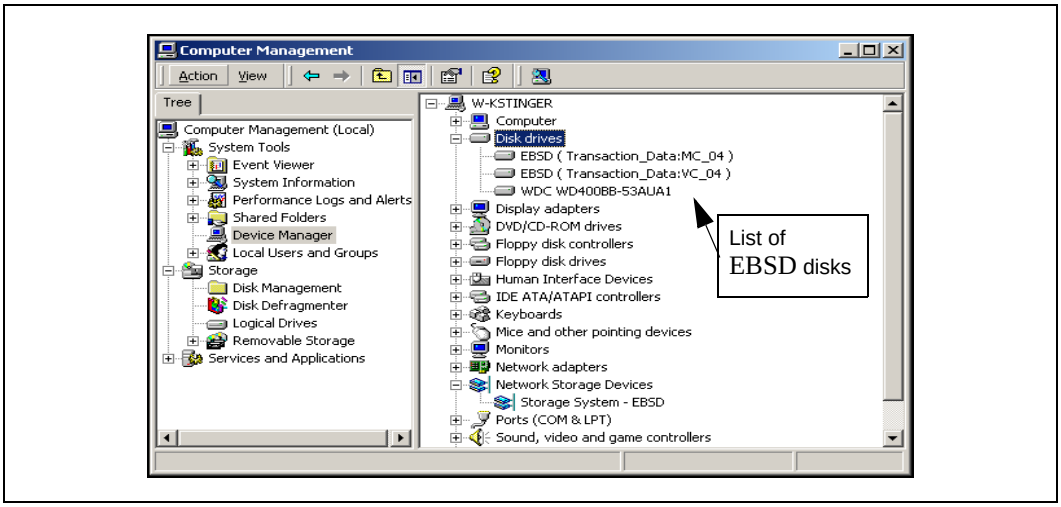

3. Select the disk for which you want to enable write cache.
4. Double-click on the disk or right-click and select Properties from the menu.
   The Disk Properties window opens, shown in Figure 223.

Figure 223. Opening the Disk Properties from the Device Manager

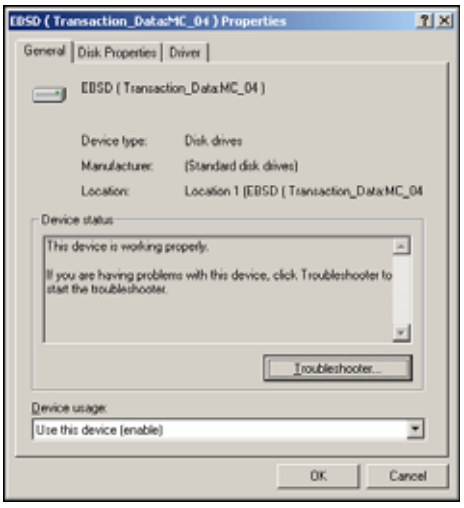

5. Click the Disk Properties tab to bring it to front, shown in Figure 224.

Figure 224. Opening the Disk Properties Tab

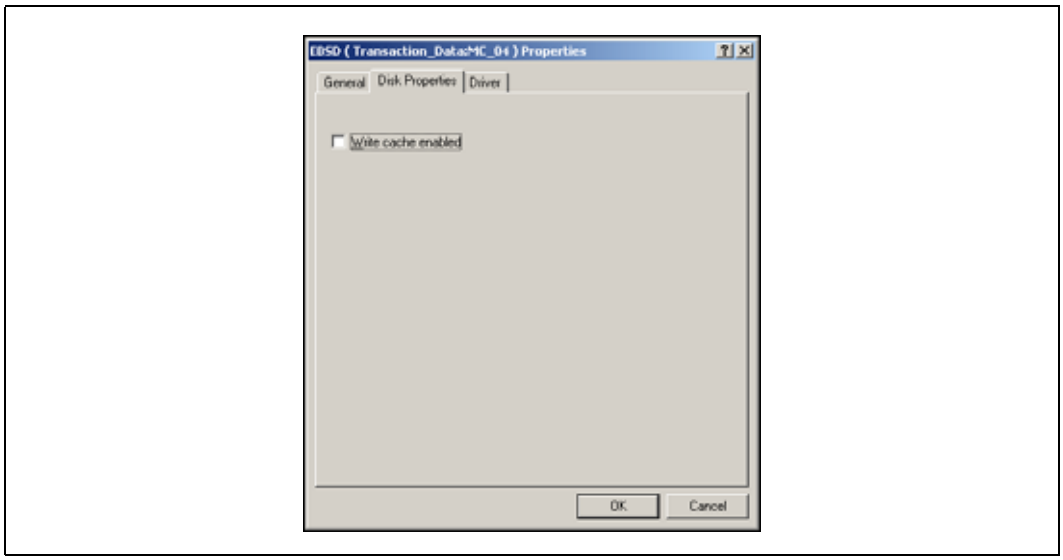

6. Select the check box to enable write cache.
   A warning message opens, describing the risk of losing data.
7. Click OK.
   Focus returns to the Disk Properties tab.
8. Click OK to close the Disk Properties window.

C.10.2.1 Verifying Write Cache Status

Verify that write cache is enabled after changing the setting.

1. Open the EBSD driver.
2. Click the Settings tab to bring it to the front.

Figure 225. Viewing the Status of Write Cache on the EBSD Volumes

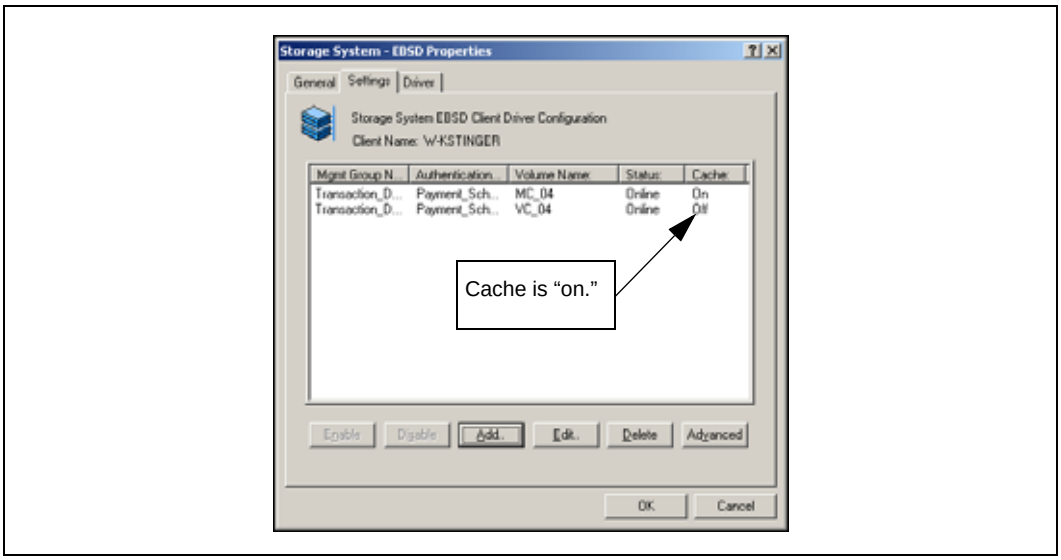

3. Verify that cache is on for those volumes on which you enabled it.

C.10.3 Disabling Write Cache on Volumes

Make certain that all data transfers are complete before disabling write cache.

1. Open Windows Device Manager.
2. Expand the Disk Drives list on the right.
   The EBSD disks are listed as shown in Figure 222. See “Identifying the Storage System Software Volume That Corresponds to an EBSD Disk” to verify which disk you are working with.
3. Select the disk for which you want to disable write cache.
4. Double-click on the disk or right-click and select Properties from the menu.
   The Disk Properties window opens, shown in Figure 223.
5. Clear the check box to disable write cache.
6. Click OK to close the Disk Properties window.

C.11 Writing the Disk Signature

Next you write disk signatures on the new EBSD disks as if they are locally attached hard drives.

1. Open Windows Disk Management.
   The Disk Management window opens and begins to scan for new devices.
   Next, the Write Signature and Upgrade Disk wizard opens, shown in Figure 226.

Figure 226. Opening the Write Signature and Upgrade Disk Wizard

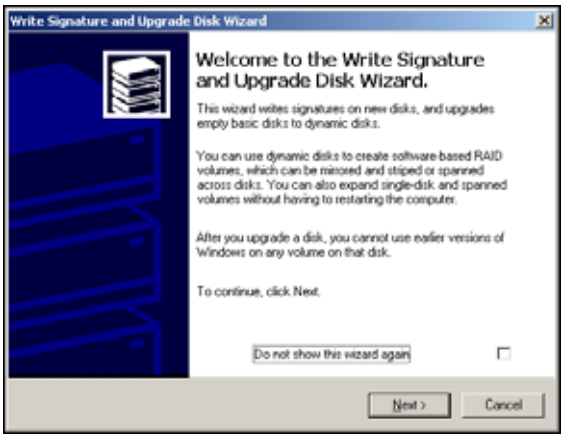

Note: If the Write Signature and Upgrade Disk wizard does not open, you can manually create signatures on the disks.
- Right click on the disk in the Disk Management window.
- Select Write Signature from the menu.

2. Click Next to continue the wizard.
   The Select Disk to Write Signature window opens, shown in Figure 227.

Figure 227. Selecting Disks to Write Signatures

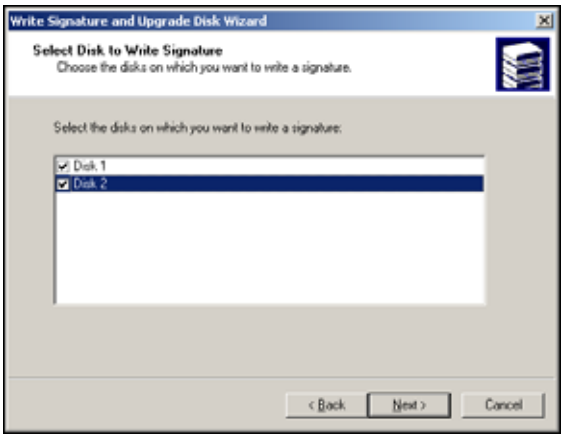

3. Select the disks for which you want to write signatures and click Next.
   The Select Disks to Upgrade window opens, shown in Figure 228.

Note: It is recommended that you use only basic disks with the Storage System Software system. For a detailed discussion of basic disks, disk management programs, and expanding basic partitions, visit the support section of the Company Name web site.

Figure 228. Clearing Disk Selections

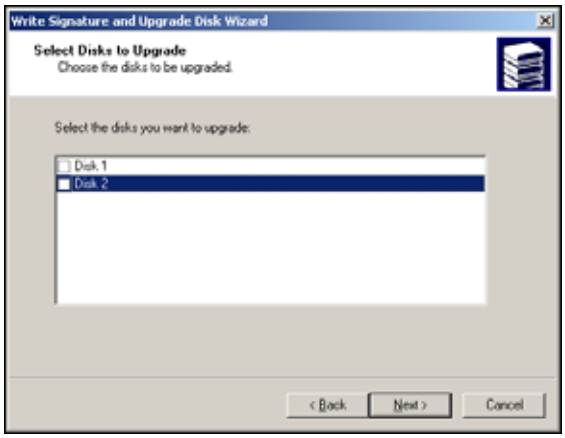

4. Clear the check box(es) and click Next.

   The Completing the Write Signature and Upgrade Disk wizard window opens, shown in Figure 229.

Figure 229. Completing the Write Signature and Upgrade Disk Wizard

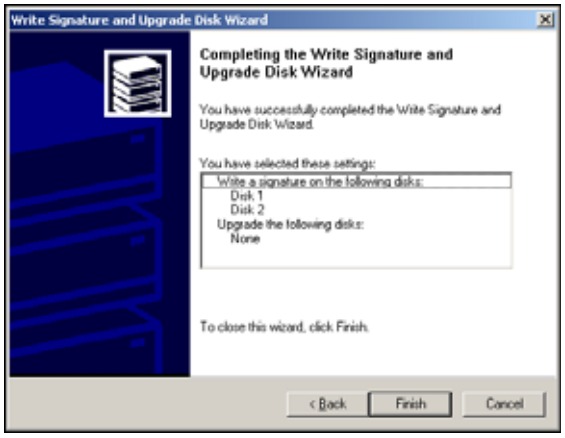

5. Review the selections and click Finish.

   The disks are ready to partition and format. They display in the Disk Management window with the disk area showing Unallocated, shown in Figure 230.

Figure 230. Viewing Unpartitioned EBSD Disks in the Disk Management Window

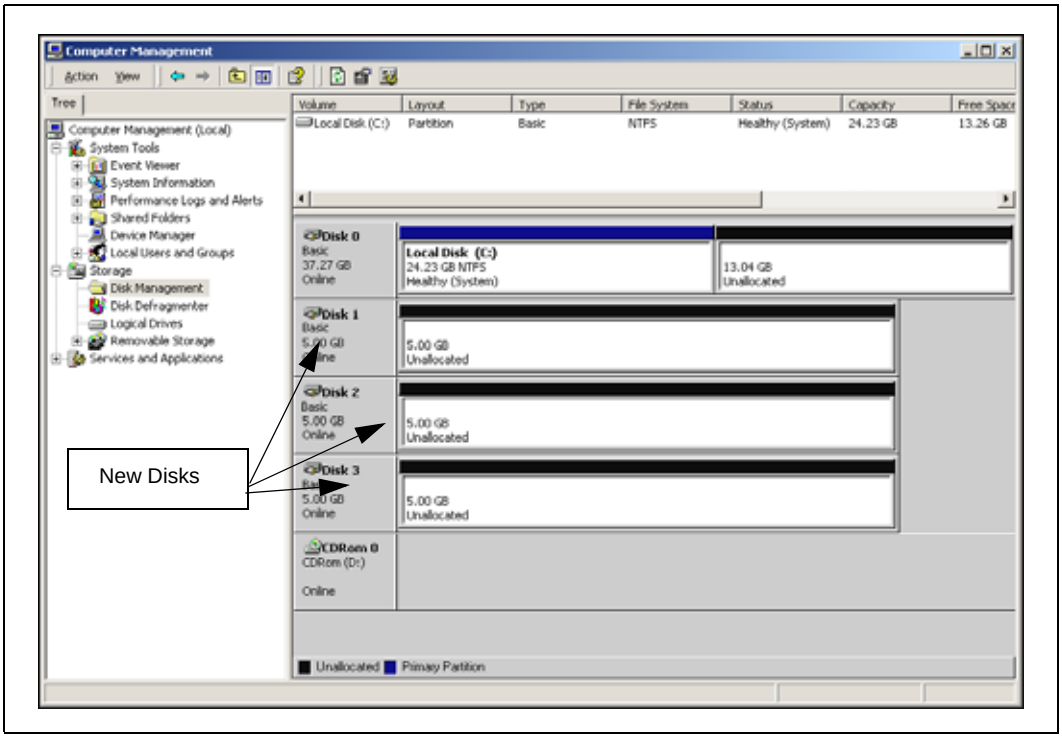

C.12 Partitioning Basic EBSD Disks

Partition an EBSD disk.

1. In the Disk Management window, right-click on the partition area labeled Unallocated with the size.

   For example, in Figure 230, Disk 3 shows 5 GB Unallocated.

2. Select Create Partition from the menu.

   The Create Partition wizard opens.

3. Click Next.

   The Select Partition Type window opens, shown in Figure 231.

Figure 231. Selecting the Type of Partition

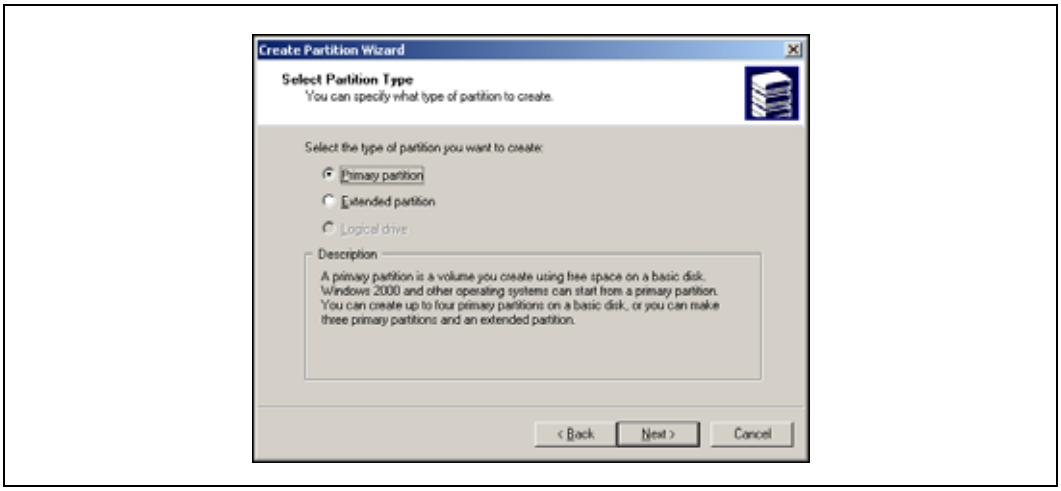

4. Select Primary Partition and click Next.
   The Specify Partition Size window opens, shown in Figure 232.

Figure 232. Selecting the Partition Size

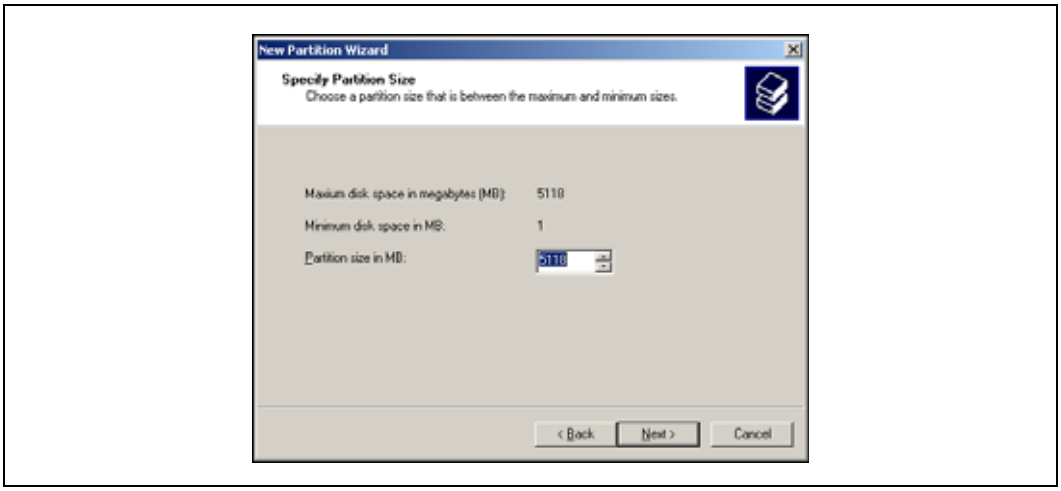

5. Select a size for the partition and click Next.

Note: The standard Windows disk driver is limited to 2TB devices at the block level. The EBSD driver uses this standard disk driver, so native volumes are limited to 2TB.
For basic disks, the 2TB limit is a hard limit with no known work-around.

C.12.1 Assigning Drive Letters and Formatting Partitions

1. The Assign Drive Letter or Path window opens, shown in Figure 233.

Figure 233. Assigning a Drive Letter to a Basic Disk

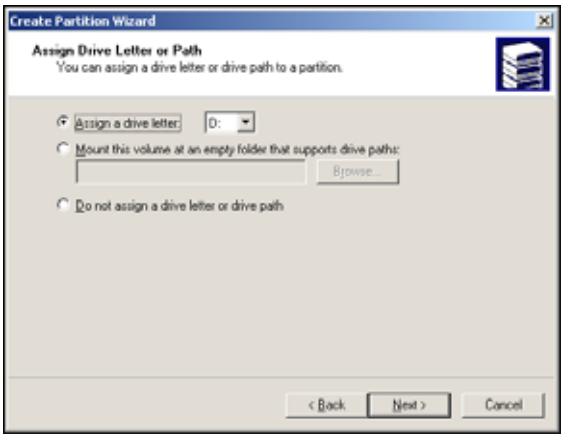

2. Select a drive letter and click Next.
   The Format Partition window opens, shown in Figure 234.

Figure 234. Formatting the Partition

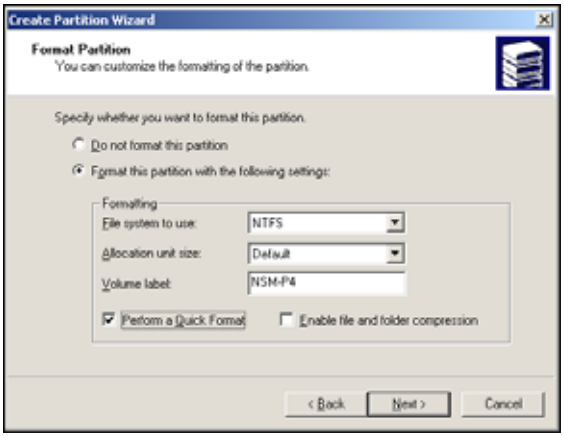

3. Complete the Formatting information and click Next.
   The final Create Partition wizard window opens showing a summary of the partition specifications.
4. Review the specifications and then click Finish.
   The partition is created and formatted and appears in the Disk Management window.

C.13 Configuring Applications and Services to Come Online After A Reboot

To ensure that applications and services using EBSD volumes come online after a reboot of the system, add the appropriate services using the Advanced Settings in the EBSD driver. Adding services changes the Startup type of those services to Manual. Then the EBSD driver starts those services either when

- all the EBSD volumes arrive after a reboot of the system or
- after 8 minutes has passed. If 8 minutes passes but all volumes have not arrived, an event log message is recorded and a message notifies the administrator that all EBSD volumes have not arrived.

Removing the services in the Advanced Settings changes the Startup type of those services to back to their original settings.

Note: The EBSD driver adds the 'ArrivalDelay' registry entry to the aebsagent registry group. This entry delays invocation of services until the EBSD volume is online, or for 8 minutes (480 seconds), whichever occurs first. You can change the default delay at
HKEY_LOCAL_MACHINE\SYSTEM\
CurrentControlSet\Services\aebscheck\
ArrivalDelay:REG_DWORD:0x1e0
The default value for ArrivalDelay is 480 (seconds).

C.13.1 Configuring File Services, SQL Server, and Exchange

File server, SQL Server and Exchange services and their dependent services can be quickly preconfigured in the Advanced Settings window.

C.13.2 Configuring Other Applications with User Services

All other applications require that you identify the Service name before going to Advanced Settings. In the Advanced Settings window you use the User Services button to configure those other applications.

C.13.2.1 Identifying the Service Name for Applications

For applications that are not file server applications, SQL Server, or Exchange, you must identify the application's service name when configuring the application in Advanced Settings. To obtain the application's service name, go to that service's Properties window. Note that the service name is not the same name as the name listed in the Services window.

For example, configure Veritas Backup Exec to come online after a reboot. Veritas Backup Exec is controlled by the Backup Exec Device & Media Service.

1. Click Start > Settings > Control Panel.
2. Select Administrative Tools > Services.
   The Services window opens.

Figure 235. Opening the Services Window and Selecting the Backup Exec Device and Media Service

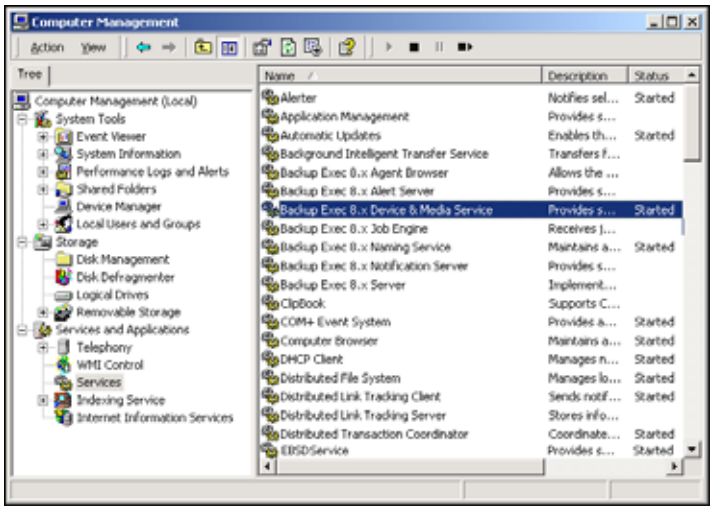

3. Locate the service you want to configure and right-click to open the menu.
   Using our example, select Backup Exec Device & Media Service.
4. Right-click and select Properties.
   The Backup Exec Device & Media Service. Properties dialog opens, shown in Figure 236.

Figure 236. Opening the Service Properties Dialog

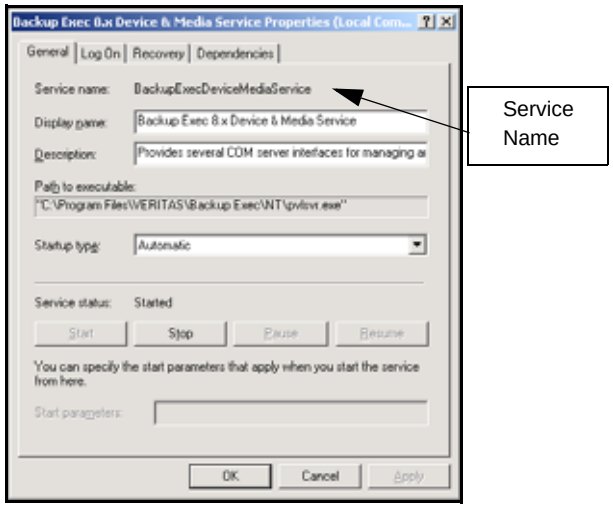

5. Make a note of the Service name listed.
   For example, the name for Backup Exec Device & Media Service is BackupExecDeviceMediaService.
6. Repeat steps 3 through 5 for each service you want to configure to come online after reboot.

7. Close the Services window when you are finished.

C.13.3 Configuring Services In The EBSD Driver

Advanced Settings offers three predefined applications—file server, MS SQL, and MS Exchange—to configure for coming online after a reboot. All other applications are configured individually.

1. Open the EBSD driver.
2. Click the Settings tab to bring it to the front, as shown in Figure 237.

Figure 237. Settings Tab with the Advanced Button

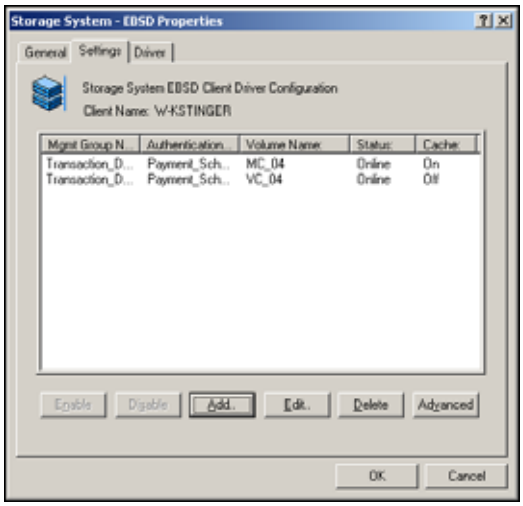

3. Click Advanced to open the EBSD Advanced Settings window, shown in Figure 238.

Figure 238. Advanced Settings Window

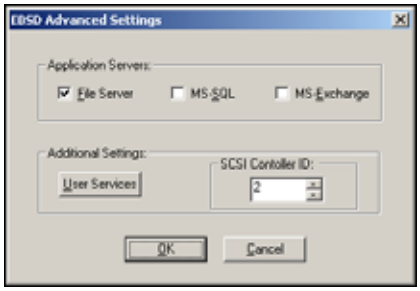

C.13.3.1 Configuring Services for File Server, MS SQL and MS Exchange

If you selected File Server, SQL Server, or MS Exchange when you completed the Add Hardware wizard, your settings display in the Advanced Settings window. You can change these selections.

1. Select the application for which you are configuring services.

2. Click OK.
   A confirmation message opens.
3. Click OK.

C.13.3.2 Configuring Other Applications with User Services

For all other applications, use the Additional Settings section.

1. On the Advanced Settings window, click User Services.
   The EBSD Service Settings window opens, shown in Figure 239.

Figure 239. Configuring Other Application Services with Their Dependencies

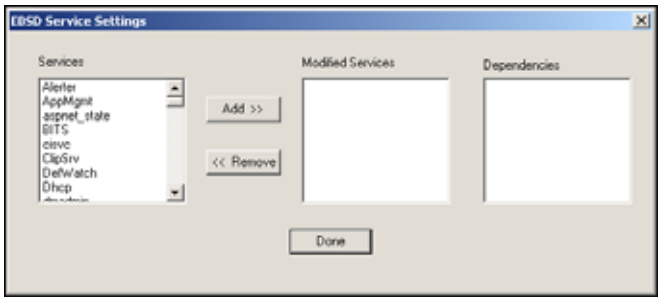

2. Using the list of service names you identified from the Services properties windows, select the services you want to configure for coming online after reboot.
   For example: Select BackupExecDeviceMediaService.to configure Backup Exec Device & Media Service come online after reboot.
3. Click Add.
   The service you selected displays in the Modified Services list. Any dependent services display in the Dependencies List. Those dependent services are automatically included in this configuration.
   For example, one dependent service of **BackupExecDeviceMediaService** is **Backup Exec Job Engine**. Depending upon the system, there may be other dependent services under Backup Exec Device & Media Service.
4. Repeat steps 2 and 3 for all the services you want to configure.
5. Click Done when you are finished.
   Focus returns to the Advanced Settings window.
6. Click OK.
   Focus returns to the Settings tab.
7. Click OK to close the EBSD driver.

C.13.4 Resetting Services

You can return the application services to their original settings.

1. Open the EBSD driver.
2. Click the Settings tab to bring it to the front.

3. Click Advanced to open the EBSD Advanced Settings window, shown in Figure 238.

C.13.4.1 Resetting File Server, MS SQL, MS Exchange

Reset File Server, MS SQL or MS Exchange to their original setting.

1. Clear the desired application in the Advanced Settings window.
2. Click OK to close the Advanced Settings window.

C.13.4.2 Resetting Other Applications

Reset any other applications by using the User Services.

1. Click User Services to open the EBSD Service Settings window, shown in Figure 240.
   A list of services is listed in the Modified Services pane on the right.

Figure 240. Modified Services in Advanced Settings

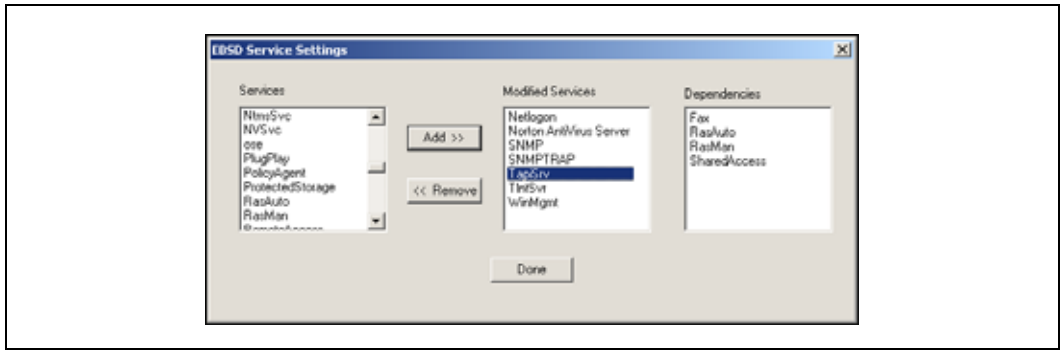

2. Select the service you want to reset and click Remove.
3. Click OK.
   The service name appears in the Services list on the left.
4. Repeat steps 2 and 3 for each service you want to reset.
5. Click Done when you are finished.
   Focus returns to the Settings tab.
6. Click OK to close the EBSD driver.

C.14 Changing the SCSI Controller ID

The default SCSI controller ID for the EBSD driver is 2. In a clustered application server environment, you may have other SCSI controllers with an ID of 2. If other controllers have an ID of 2, a conflict occurs and EBSD volumes will not come online.

To bring the EBSD volumes online, change the SCSI controller ID on the EBSD driver.

1. Disable all volumes on the EBSD driver.
   See “Disabling and Re-enabling EBSD Disks”.
2. Open the EBSD driver.

3. Click the Settings tab to bring it to the front, as shown in Figure 241.

Figure 241. Settings Tab with the Advanced Button

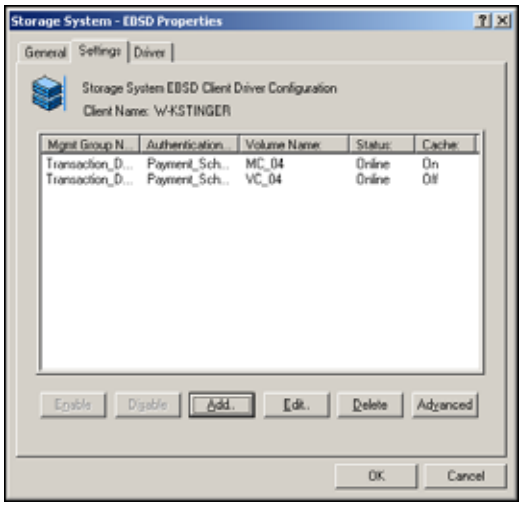

4. Click Advanced to open the EBSD Advanced Settings window, shown in Figure 242.

Figure 242. Advanced Settings Window

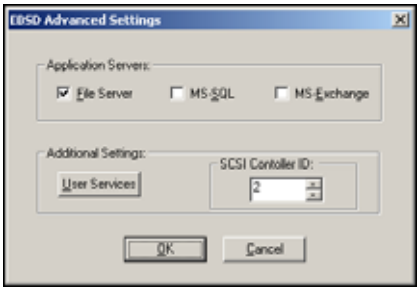

5. Change the value in the SCSI Controller ID box.
   The ID must be a number between 0 and 7. Make sure the port is not already in use.
6. Click OK.
7. Reenable the volumes.

C.15 Managing EBSD Disks

C.15.1 Overview of Managing EBSD Disks

After you have installed and configured the EBSD driver and configured EBSD disks, you will continue to manage those disks and their corresponding volumes on the SSMs. The table below lists various management tasks and where to find information about them.

Table 57. EBSD Driver Management Tasks

Management Task	Management Tool	Instructions
Identifying the Storage System Console volume that corresponds to an EBSD disk.	• Windows Disk Management	"Identifying the Storage System Software Volume That Corresponds to an EBSD Disk"
Accessing read only volumes and snapshots from an EBSD client.	• EBSD driver • Storage System Console	"Accessing Read Only Volumes and Snapshots from an EBSD Client" In the User Manual • Working with Authentication Groups chapter, see "Editing Permissions" and "Creating an Authentication Group Association" • Working with Snapshots chapter, see "Creating Snapshots"
Expanding volumes	• Storage System Console • Windows Disk Management	In the User Manual • Working with Volumes chapter, "Editing a Volume" "Expanding Volumes"
Disabling and re-enabling EBSD disks	• Unplug or Eject Hardware • Windows Device Manager • EBSD driver	"Identifying the Storage System Software Volume That Corresponds to an EBSD Disk" "Unplugging or Ejecting the Hardware" "Disabling the EBSD Disk" "Enabling EBSD Disks"
Moving EBSD disks and preserving data on the SSM	• Unplug or Eject Hardware • Windows Device Manager • EBSD driver • Storage System Console • Windows Disk Management	"Identifying the Storage System Software Volume That Corresponds to an EBSD Disk" "Unplugging or Ejecting the Hardware" "Deleting the EBSD Disks from the Client" "Preparing a New Client" "Adding EBSD Disks to the New Client" In the User Manual • Working with Authentication Groups chapter, "Authentication Groups Overview"
Deleting EBSD disks and removing data from the SSM	• Windows Disk Management • Unplug or Eject Hardware • Windows Device Manager • EBSD driver • Storage System Console	"Identifying the Storage System Software Volume That Corresponds to an EBSD Disk" "Deleting Partitions or Volumes from the Client" "Unplugging or Ejecting the Hardware" "Deleting the EBSD Disks" In the User Manual • Working with Volumes chapter, "Deleting a Volume"
Uninstalling the EBSD driver	• Windows Device Manager • EBSD driver	"Overview of Uninstalling the Driver" "Uninstalling the EBSD Driver"

C.16 Identifying the Storage System Software Volume That Corresponds to an EBSD Disk

When managing EBSD disks in Windows 2000, you must identify which EBSD disk corresponds to a specific volume in the Storage System Console.

1. Open Windows Disk Management, shown in Figure 243.

Figure 243. Identifying the Storage System Software Volume that Corresponds to an EBSD Disk

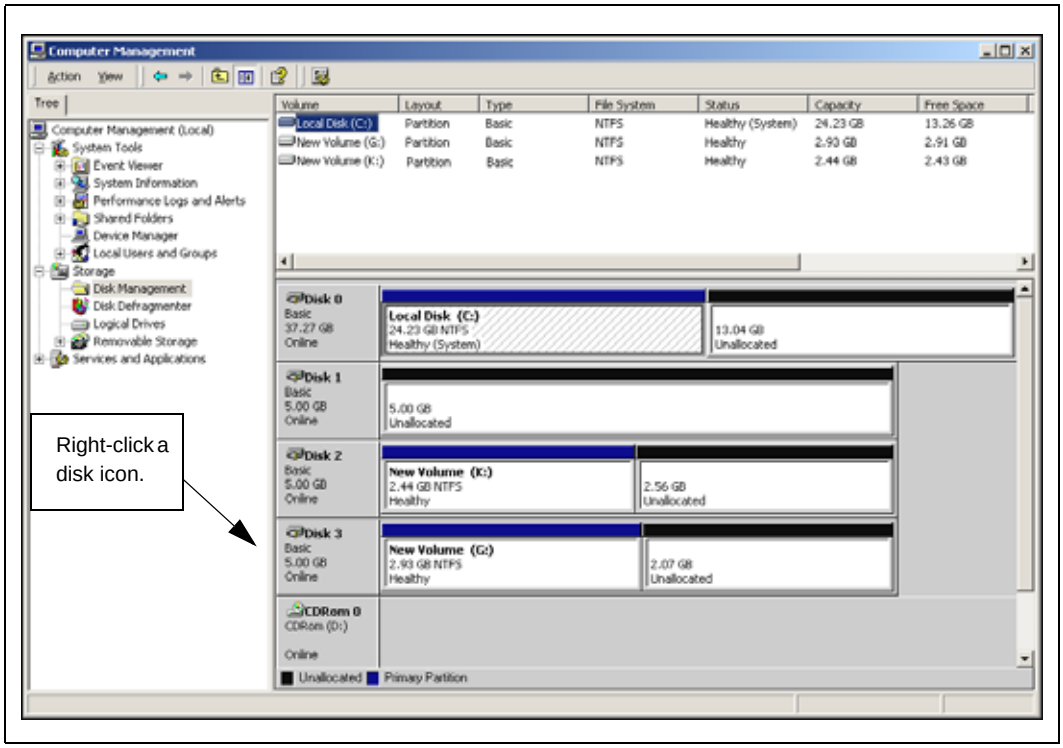

2. Right-click on a disk icon, as shown in Figure 243, and select Properties.
   The Disk Properties window, shown in Figure 244, opens for that disk.

Figure 244. Viewing Disk Properties

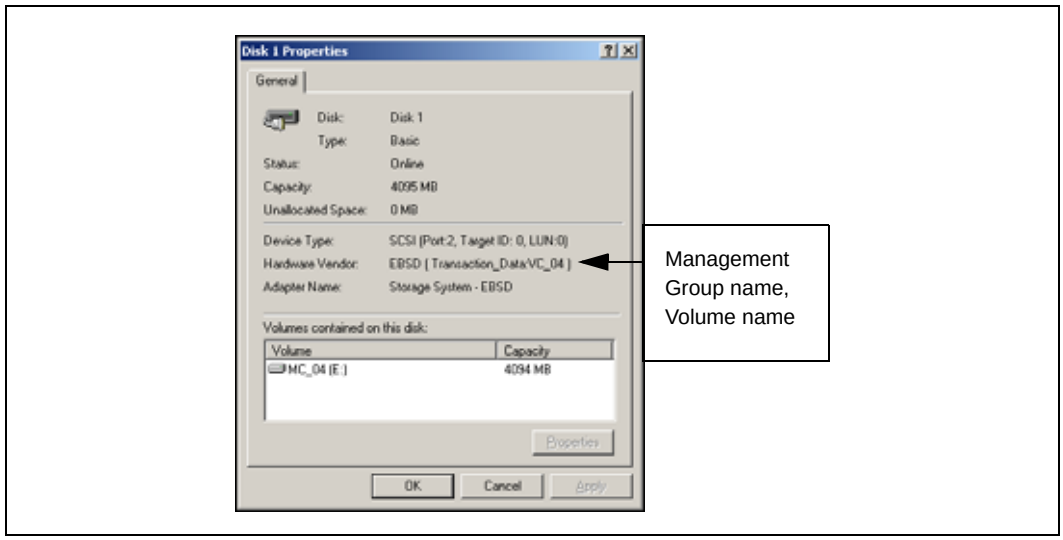

3. On the Disk Properties window, determine the name of the volume by reading the data in the Hardware Vendor field: EBSD (Management Group Name: Volume Name).

   For example, in Figure 244, the EBSD disk corresponds to the MC03 volume in the Transaction_Data management group.

   The management group name and volume name in this line are the names as designated in the Storage System Console.

4. Click OK when you are finished.

C.17 Editing EBSD Volumes

You may encounter circumstances in which you want to edit an EBSD volume. For example, if you work with snapshots or read only volumes, you may have to change some settings in an EBSD volume.

Note: Editing an EBSD volume requires that you disable the volume before you edit it.

Note: Stop any applications from accessing the disk you are going to edit. Applications cannot write to a disk that is disabled.

C.17.1 Unplug/Eject Hardware

First, unplug/eject the hardware. See “Unplugging or Ejecting the Hardware”.

C.17.2 Disable the Disk

Next, disable the EBSD volume so that no activity can take place on that volume. See “Disabling the EBSD Disk”.

C.17.3 Open the EBSD Driver and Edit Disk

Open the EBSD driver. See "Adding EBSD Disks to Your System".

1. Select the Settings tab.
2. Select the volume for which you want to change permissions and click Edit.
   The Add EBSD Volume window opens.
3. Make the desired changes.
4. Click OK.

C.17.4 Re-enable the Disk

Re-enable the EBSD volume. See "Enabling EBSD Disks".

C.18 Accessing Read Only Volumes and Snapshots from an EBSD Client

You can access snapshots and volumes configured with read only permissions. Mount these snapshots just as you would a regular volume. However, when these snapshots are mounted, there is no indication that they are read only snapshots of a regular volume. However, there are limitations to be aware of.

- Read only volumes and snapshots appear to be read/write, but changes do not get committed. What looks like a change to a snapshot is actually volatile. When the system reboots, or the snapshot is disabled and re-enabled, the changes are not saved.
- Snapshots must have an authentication group with read only access associated with them before mounting.
- The permission level of read only volumes and snapshots must be read only. If the permission level is not read only, the snapshot or volume will not come online and it will remain in the "starting" state.

C.18.1 Changing Volumes to Read Only

You can change a volume configured with read/write access to read only access.

C.18.1.1 Unplug/Eject Hardware

- First, unplug/eject the hardware. See "Unplugging or Ejecting the Hardware".

C.18.1.2 In the EBSD Driver

- Disable the EBSD disk.
  See "Disabling the EBSD Disk".

C.18.1.3 In the Storage System Console

- Change permissions for the volume.

  See “Editing Permissions” in the Storage System Console Online Help, or in the User Manual in the chapter entitled Working with Authentication Groups.

C.18.1.4 In the EBSD Driver

- Change permissions in the Settings tab.

1. Open the EBSD driver.

   See “Adding EBSD Disks to Your System”.
2. Select the Settings tab.
3. Select the disk for which you want to change permissions and click Edit.

   The Add EBSD Volume window opens.
4. Change the Permissions Level to Read Only.
5. Click OK.

- Re-enable the EBSD disk. See “Enabling EBSD Disks”.

C.18.2 Moving Read Only Volumes to a Different Client

Follow the instructions in “Overview of Deleting or Moving EBSD Disks”.

C.18.3 Mounting Snapshots of Basic EBSD Disks

When mounted, snapshots of basic EBSD volumes appear as basic disks to Windows Disk Manager. When mounting snapshots of basic disks, there are no limitations. You can mount multiple snapshots of the same read/write volume on the same EBSD client server.

C.19 Expanding Volumes

To increase the size of a volume, increase its size in the Storage System Console and then expand the volume or create new volumes/partitions in the Windows 2000 Disk Management window.

1. Increase the volume size, hard threshold, and soft threshold in the Storage System Console.

   Wait for the volume status to change from “Re-striping” to “Normal.”

   See “Editing a Volume” in the User Manual in the chapter entitled Working With Volumes.
2. Open the Disk Management window.

   The newly added storage space appears as Unallocated, shown in Disk 6 in Figure 245.

   — If the newly added storage space does not appear, click the Action menu and select Rescan Disks.
3. Increase the volume size or create new partitions as follows:

   — Create a new partition on the unallocated space.

   — Expand the existing volume using diskpart.exe.

Figure 245. Increasing the Size of Basic Volumes

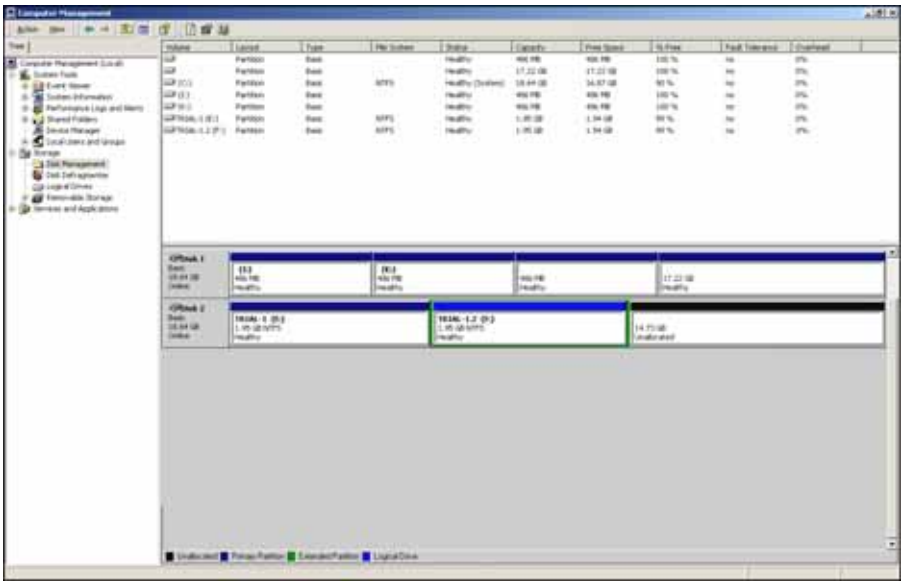

Note: You can increase or decrease the size of volumes in the Storage System Console. However, there are limitations to using the new size with a Windows 2000 EBSD client:
-- Limitations for NTFS and Basic Disks. You may be able to use 3rd party tools, such as Disk Doubler™, Norton Utilities™, and Partition Magic®, to increase and decrease the size of basic Microsoft disks. However, the 2TB limit still applies.

C.20 Disabling and Re-enabling EBSD Disks

C.20.1 Overview of Disabling and Re-enabling EBSD Disks

Disabling EBSD disks stops activity between the EBSD client and the cluster containing the volume. The partition and drive mapping remain intact and the disk can be re-enabled when appropriate.

Note: Before you begin, be sure that the EBSD disks you plan to disable are not in use.

C.21 Disabling and Re-enabling EBSD Disks

Review the information in "Identifying the Storage System Software Volume That Corresponds to an EBSD Disk" before beginning this process. You must know the relationship of the volume name to the EBSD disk name before beginning.

C.21.1 Unplugging or Ejecting the Hardware

1. Halt all applications accessing the EBSD volume about to be disabled.
2. Double-click the Unplug or Eject Hardware icon from the Windows taskbar, shown in Figure 246.

Figure 246. Unplug or Eject Hardware Icon

The Unplug or Eject Hardware window opens, shown in Figure 247.

Figure 247. Viewing the Unplug or Eject Hardware Window

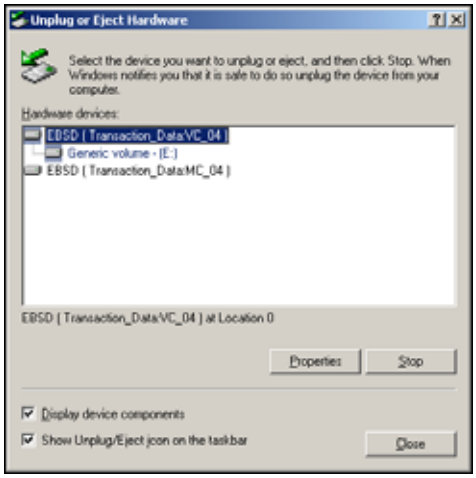

3. Make sure the check boxes at the bottom of the window are checked.

   Display device components expands the EBSD disks to show the volumes associated with each disk.

   Show Unplug/Eject icon on the taskbar ensures that the icon remains displayed on the Windows taskbar.
4. Select from the list the device you want to disable.
5. Click Stop.

   A confirmation window opens, shown in Figure 248.

Figure 248. Confirming Devices to be Stopped

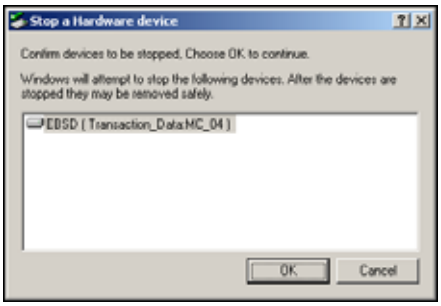

6. Verify the device and click OK.
   A message opens, verifying that the device can now be safely removed from the system.
7. Click OK.
   The Unplug or Eject Hardware window returns and the selected device is no longer in the list.
8. Repeat steps 4 through 7 for each device that you want to disable.
9. Click Close.

C.21.2 Disabling the EBSD Disk

1. Open Windows Device Manager.
2. Expand the Disk Drives list the right.
   The disks you unplugged display an exclamation point icon, as shown in Figure 249.

Figure 249. Viewing the Unplugged EBSD Disk under Expanded Disk Drives

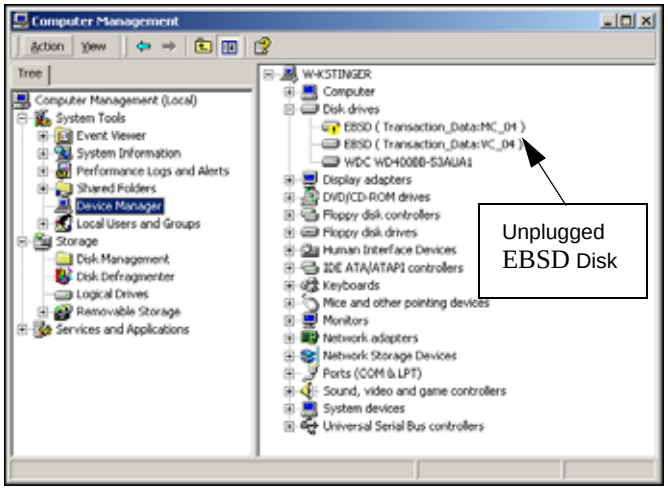

3. Expand the Network Storage Devices list and select the EBSD driver.
4. Double-click on the driver or click the Action menu and select Properties.

The EBSD Properties window opens, shown in Figure 250.

Figure 250. Disabling an EBSD Disk

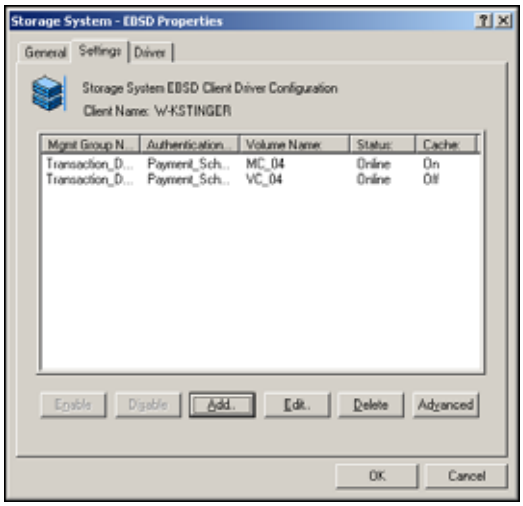

5. Click the Settings tab to bring it to the front.
6. Select from the list the volume that you want to disable.
7. Click Disable.

   A message opens, warning that the data on the volume will become unavailable, and to make certain that all applications are stopped.
8. Click Yes.
9. Repeat steps 6 through 8 for each volume that you want to disable.

Note: Make sure that these disks are already stopped through the Unplug/Eject Hardware option, described in "Unplugging or Ejecting the Hardware".

10. Click OK when you are finished.

    The EBSD Properties window closes. The Disk Drive list no longer displays the disks.

C.21.3 Enabling EBSD Disks

You can enable EBSD disks that have been disabled, but not deleted, as long as you have not changed the corresponding volume configuration on the SSM(s) in the Storage System Console.

1. Open Windows Device Manager.
2. Expand the Network Storage Devices list and select the EBSD driver.
3. Double-click on the driver or click the Action menu and select Properties.

   The EBSD Properties window opens.
4. Click the Settings tab to bring it to the front, shown in Figure 251.

Figure 251. Enabling EBSD Disks

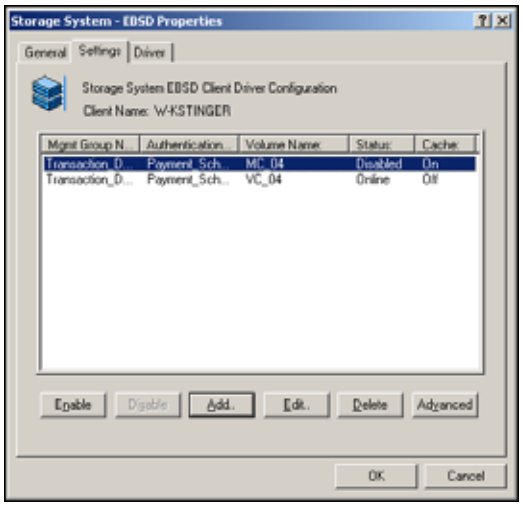

5. Select from the list the volume that you want to enable.
   In this example, we are enabling the MC03 volume.
6. Click Enable.
   The Status column displays “Starting,” shown in Figure 252. Then the status changes to Online.

Figure 252. “Starting” Status of an EBSD Disk

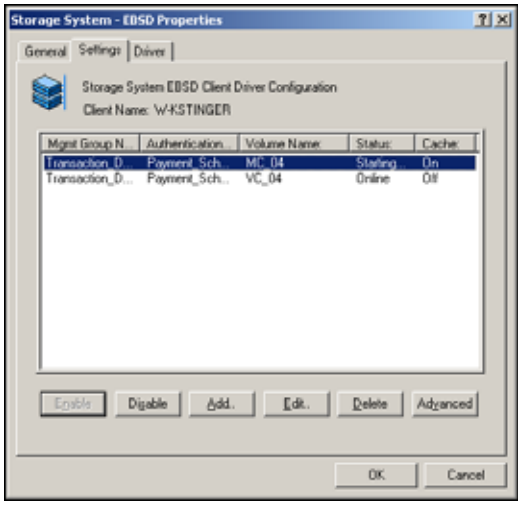

7. Repeat steps 5 and 6 for each volume that you want to enable.
8. Click OK when you are finished.
   The EBSD Properties window closes.

Verifying That the Disks are Enabled

The Disk Drive list now displays the disks. When you open the Disk Management window, the disks and their volumes are again visible.

C.22 Deleting or Moving EBSD Disks

C.22.1 Overview of Deleting or Moving EBSD Disks

You can delete EBSD disks from a client while preserving data on the volume in the cluster. You can also delete EBSD disks and delete all the data on the SSM.

Deleting EBSD disks requires the following tasks:

- Unplugging the hardware
- Deleting EBSD disks from the EBSD driver

You can then reconnect those EBSD disks on a different server to access the same volumes (and data) on the cluster. The partitions are preserved, though the drive letters might change, depending on whether you change the client machine.

If you reconnect EBSD disks on a client that does not belong to the authentication group associated with the volume in the Storage System Console, you must create a new authentication group for the client and associate the new group with the volume.

Note: Reassigning EBSD disks to a different client preserves the data that is stored on the cluster. The reassigned EBSD disks are connected to the same volumes on the cluster.

C.23 Deleting or Moving EBSD Disks While Preserving Data

Before you begin, review the information in "Identifying the Storage System Software Volume That Corresponds to an EBSD Disk". You must know the names of the EBSD disks you are reassigning to ensure that you are working with the correct volumes in the Storage System Console.

Note: Before you begin, be sure that the EBSD disks you plan to remove are not in use.

C.23.1 Unplugging or Ejecting the Hardware

The first step in deleting or moving an EBSD disk while preserving the data on the SSM is to unplug or eject the EBSD disk from Windows using the Unplug or Eject Hardware icon on the Windows taskbar, shown in Figure 253.

Figure 253. Unplug or Eject Hardware Icon

Repeat this process for each disk that you want to delete or move. For more information about unplugging or ejecting hardware, see "Unplugging or Ejecting the Hardware".

C.23.2 Deleting the EBSD Disks from the Client

The second step in deleting or moving an EBSD disk while preserving the data on the SSM is to delete the EBSD disk from the EBSD client.

1. Open Windows Device Manager.
2. Expand Disk Drives from the list on the right.

   The disks you unplugged display an exclamation point icon, as shown in Figure 254.

Figure 254. Viewing the Unplugged EBSD Disk under Expanded Disk Drives

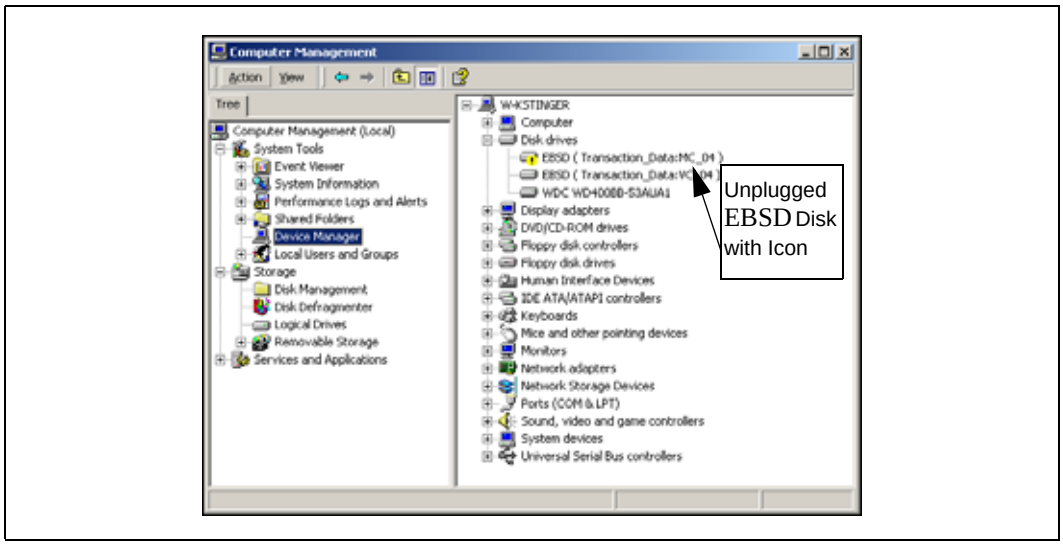

3. Expand Network Storage Devices and select the EBSD driver.

Figure 255. Selecting the EBSD Driver

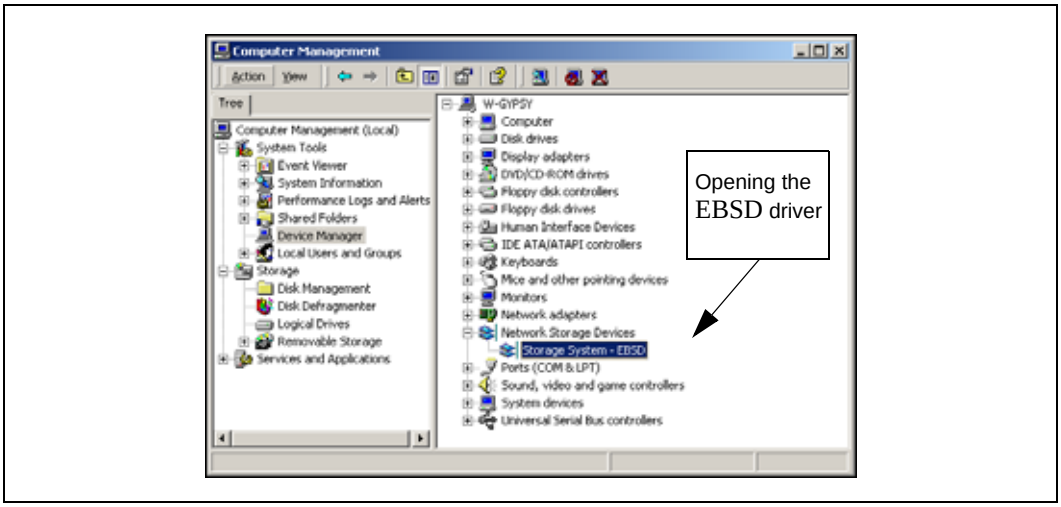

4. Double-click the driver, or right-click and select Properties.
   The EBSD Properties window opens.
5. Click the Settings tab to bring it to the front, as shown in Figure 256.

Figure 256. Deleting an EBSD Disk

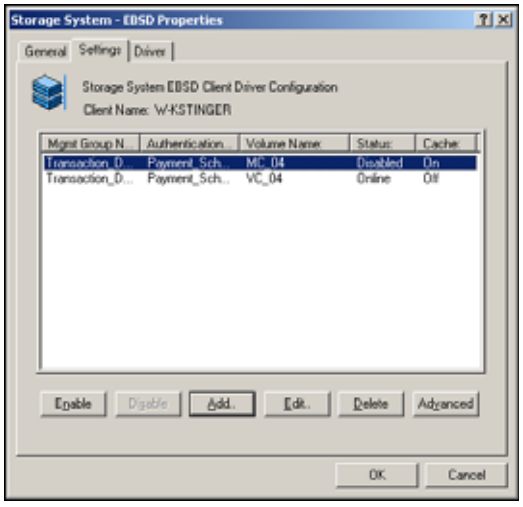

6. Select from the list the disk you want to delete.
7. Click Delete.
   A warning message opens, shown in Figure 257, warning that you cannot access data on this disk once it is deleted.

Figure 257. Warning Message before Deleting EBSD Disk

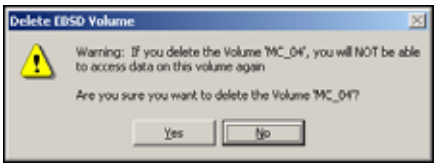

8. Click Yes to confirm deleting the disk.
   The message window closes.
9. Repeat steps 6 through 8 for each disk you want to delete.
10. Click OK to close the EBSD driver.

C.23.3 Preparing a New Client

Before you reconnect the EBSD disks on a new client, you may need to reconfigure the volumes in the Storage System Console.

C.23.3.1 Associate New Client with an Authentication Group

If the new client does not belong to the authentication group associated with the volume in the Storage System Console, you must create a new authentication group for the client and associate the new group with the volume. See “Authentication Groups Overview” in the User Manual in the chapter entitled Working with Authentication Groups.

C.23.3.2 Install EBSD Driver

If the EBSD driver has not been installed on the new client machine, you must install it before adding the EBSD disks. For detailed instructions, see “Installing the EBSD Driver” and “Configuration Overview”.

C.23.4 Adding EBSD Disks to the New Client

Once you have prepared the new client, add the EBSD disks to the new client.

To add the EBSD disk to the new client, open the EBSD driver and click Add on the Settings tab. For detailed instructions on adding an EBSD disk, see “Adding EBSD Disks to Your System”.

C.23.5 Finishing Up

After you reconnect the EBSD disks in the EBSD driver, go to Disk Management. Follow any system instructions that appear. For example, you may be prompted to assign new drive letters to the partitions or volumes.

C.24 Deleting EBSD Disks and Removing Data from the SSM

Removing EBSD disks and ensuring that data is removed from the volume on the cluster requires the following tasks on the client:

- Deleting partitions
- Unplugging the hardware
- Deleting EBSD disks

and the following task in the Storage System Console:

- Deleting volumes from the cluster. See “Deleting a Volume” in the User Manual in the chapter entitled Working with Volumes.

Note: Make sure the EBSD disks you plan to delete are NOT in use. Deleting EBSD disks from the client and deleting volumes from the cluster removes all data stored in those volumes. Once removed, that data cannot be retrieved.

Before you begin, review the information in “Identifying the Storage System Software Volume That Corresponds to an EBSD Disk”. You must know the names of the EBSD disks you are removing to ensure that you are removing the correct volumes in the Storage System Console.

C.24.1 Deleting Partitions or Volumes from the Client

Before you delete an EBSD disk, you must first remove the partitions or volumes from the disk.

1. Open Windows Disk Manager.
2. Delete the partitions on the disk that you want to delete.
   — Right-click the partition and select Delete Partition.
   A confirmation message opens.
3. Click Yes.
   The partition is deleted and the disk becomes unallocated.
4. Repeat steps 2 and 3 for all the disks you want to delete.

C.24.2 Unplugging or Ejecting the Hardware

The second step in deleting an EBSD disk and deleting the data on the SSM is to unplug or eject the EBSD disk from Windows using the Unplug or Eject Hardware icon on the Windows taskbar, shown in Figure 258.

Figure 258. Unplug or Eject Hardware Icon

Repeat this process for each disk that you want to delete. For more information about unplugging or ejecting hardware, see “Unplugging or Ejecting the Hardware”.

C.24.3 Deleting the EBSD Disks

1. Open Windows Device Manager.
2. Expand the Disk Drives list on the right.
   The disks you unplugged display an exclamation point icon, as shown in Figure 259.

Figure 259. Viewing the Unplugged EBSD Disk under Expanded Disk Drives

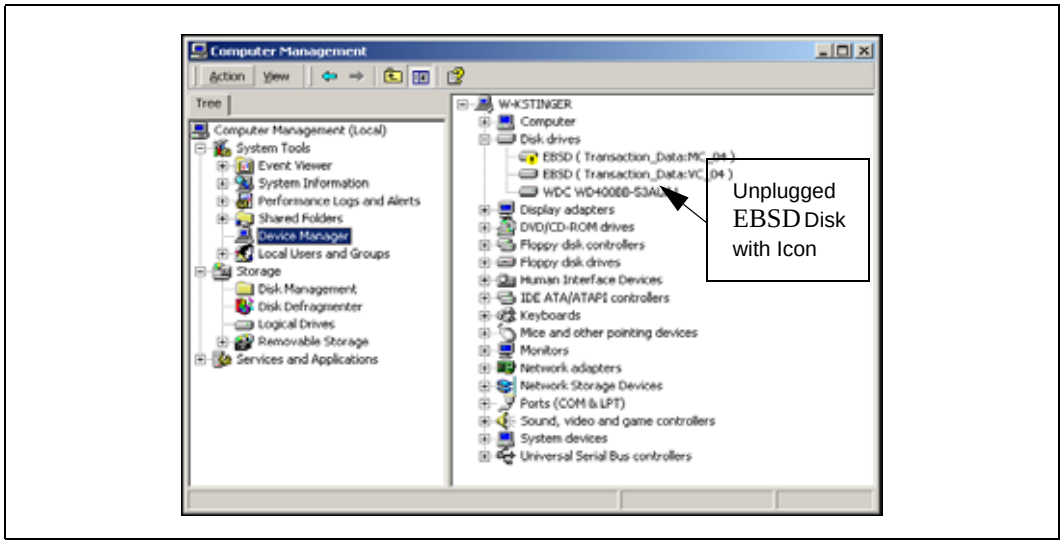

3. Expand Network Storage Devices and select the EBSD driver.
4. Double-click on the driver or click the Action menu and select Properties.
   The EBSD Properties window opens.
5. Click the Settings tab to bring it to the front, as shown in Figure 260.

Figure 260. Deleting an EBSD Disk

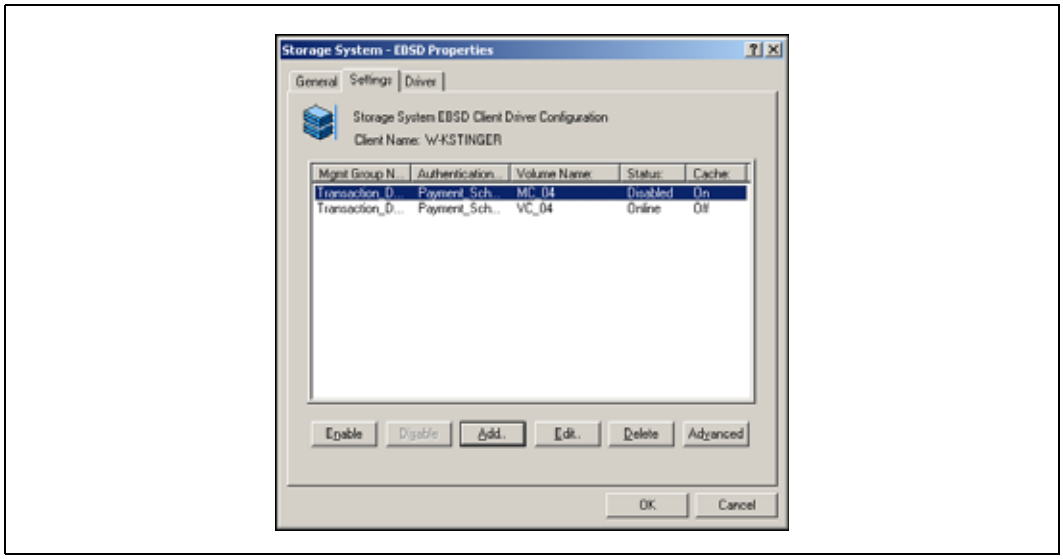

6. Select from the list the disk you want to delete.

7. Click Delete.

   A warning message opens, shown in Figure 261, warning that you cannot access data on this volume once it is deleted.

Figure 261. Warning Message before Deleting EBSD Disk

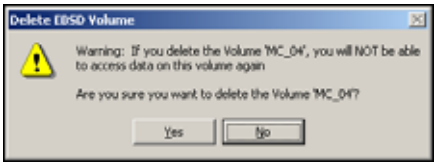

8. Click Yes to confirm deleting the disk.

   The message window closes.

9. Repeat steps 6 through 8 for each disk you want to delete.

10. Click OK to close the EBSD driver.

The final step in deleting the EBSD disk is to delete the volume in the Storage System Console. See "Deleting a Volume" in the User Manual in the chapter entitled Working with Volumes.

C.25 Uninstalling the EBSD Driver

C.25.1 Overview of Uninstalling the Driver

Before uninstalling the EBSD driver for Windows 2000, you should remove the EBSD disks. Review the procedures for removing disks in one of the following two sections, depending upon your goals.

- See "Overview of Deleting or Moving EBSD Disks" if you want to recreate the EBSD disks later, or on another client, or otherwise preserve the data.
- See "Deleting EBSD Disks and Removing Data from the SSM" if your goal is to completely remove EBSD from the client and scrub the SSM of all the data stored in the corresponding volume.

C.26 Uninstalling the EBSD Driver

1. Open Windows Device Manager.
2. Expand the Network Storage Devices list and select the EBSD driver.
3. Double-click on the driver or click the Action menu and select Properties.

   The EBSD Properties window opens.

4. Click the Driver tab to bring it to the front, as shown in Figure 262.

Figure 262. Uninstalling the EBSD Driver

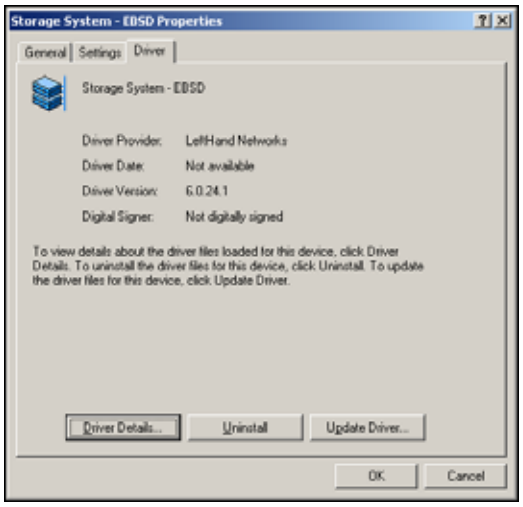

5. Click Uninstall.

   A warning message opens, shown in Figure 263.

Figure 263. Warning Before Uninstalling

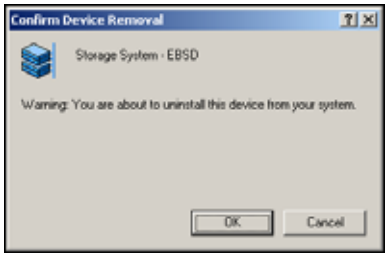

6. Click OK.

   Another message opens, telling you to delete or disable any EBSD volumes that are online.

7. Click OK.

   The driver is uninstalled. A message opens telling you to reboot your computer to complete the uninstall.

8. Restart your computer to complete the uninstall.

Using the EBSD Driver for Linux D

D.1 EBSD Driver for Linux Overview

Install and configure the EBSD Driver for Linux on the computer that accesses the SSM.

Note: You will need root privileges during installation and configuration. Use the X11R6 window environment.

D.1.1 Copying Driver Bundle to a Network Share (Optional)

Copying the driver bundle results in separate tar.gz files plus an install.sh file copied to the location you specify. The tar.gz files correspond to the versions of Linux that are currently supported by the EBSD driver. Running the install.sh file from the directory containing all the tar.gz files installs the appropriate driver for the Linux version on that system.

Skip this section if you only want to install the EBSD from the CD. Go directly to "Installing the EBSD Driver for Linux".

1. Start the Windows Manager of your choice, such as KDE or GNOME.
2. Insert the driver CD into the CD drive of the EBSD host server.
   Autorun should automatically start. A folder directory window also opens, displaying the contents of the CDROM.
3. The autorun window opens, asking you to verify that you want to run Autorun.
   The autorun window may be behind the CD folder directory.
4. Click Yes to have the automatic install process run.
   The automatic installation starts and steps you through the install process.
5. On the Choose Product Component window of the automatic installation, select Install EBSD Bundles.
6. Click Next. The Choose Install Folder window opens.
7. Choose to accept the default directory (/opt/Storage_System/Storage_System_Software/6.0/Drivers) or browse to the directory where you want the EBSD bundle installed. This location can be another location on the network, such as a file server.
8. Click Next. Review the Pre-installation Summary window.
9. Click Install. The EBSD driver bundle is copied into the directory you specified.
   The EBSD driver bundle contains tar.gz files for all the versions of Linux that are currently supported plus an install.sh file to install the driver.

D.2 Installing the EBSD Driver for Linux

You can install the EBSD driver locally using the CD or using driver bundle from a network share.

See the instructions in the table below:

- "Installing from the CD

  or

- "Installing from a Network Share

See "Copying Driver Bundle to a Network Share (Optional)" on page 305.

Table 58. Installing the EBSD Driver

Installing from the CD	Installing from a Network Share
1. Start the Windows Manager of your choice, such as KDE or GNOME. 2. Insert the driver CD into the CD drive of the EBSD client PC. Autorun should automatically start. The autorun window opens, asking you to verify that you want to run Autorun. The autorun window may be behind the CD folder directory, which also opens. 3. Click Yes to have the automatic install process run. The automatic installation starts and steps you through the install process. 4. On the Choose Product Component window, select EBSD Driver. 5. Click Next. The Choose Install Folder window opens. 6. Choose to accept the default directory (/opt/Storage_System/Storage_System_Software/6.0/Drivers). 7. Review the Pre-installation Summary window and click Install. 8. Click Done on the Congratulations window. 9. Continue with Step 1 below.	1. [Optional] Copy one or all tar.gz file(s) and the install.sh file to the computer on which you want to install the driver. See "Copying Driver Bundle to a Network Share (Optional)" for information about the tar.gz files. 2. Run the script install.sh. 3. Continue with Step 1 below.

See "Configuring the EBSD Driver for Linux" for information about configuring the EBSD driver.

D.2.1 What the Install Script Does

The install script installs the EBSD driver into the following locations:

- /usr/local/sbin/ for aebsvm
- /opt/Storage_System/Storage_System_Software/6.0/Drivers/$(uname -r)/etc/ for ebsd.conf.sample
- /etc/init.d/ for ebsd
- Script also installs run level links in /etc/rc?.d

D.3 Upgrading the EBSD Driver for Linux

1. Stop all operations to EBSD devices (i.e., unmount /ebsddisk).
2. Install driver from Installation CD or the driver bundle.
3. Run /etc/init.d/ebsd restart, or service ebsd restart

4. Cat /proc/ebsd/client to verify driver version and devices online.

D.4 Configuring the EBSD Driver for Linux

Once the EBSD driver is installed on Linux, it must be configured and started.

D.4.1 Creating ebsd.conf

1. Copy /opt/Storage_System/Storage_System_Software/6.0/Drivers/$(uname -r)/etc/ebsd.conf sample /etc/ebsd.conf
2. Modify the /etc/ebsd.conf file.
   — Add a device entry in ebsd.conf for each volume or snapshot that has been configured on the SSM. Required parameters are listed in Table 59.

Table 59. Parameters in ebsd.conf

Parameter	What It Is
[device#]	This is the device section identifier. It must be named device# where # is the device number. Corresponding block device is created as /dev/ebsd/disk#. Valid device numbers = 0 to 63.
type = volume | snapshot	EBSD device types is a volume or snapshot.
client_name = %s	The EBSD driver client's hostname.
ip_bind = x.x.x.x	The IP address of the client that you want this driver to bind to. This address identifies the interface over which the driver will communicate to this volume in a multihomed system.
management_group = %s	The name of the management group that contains the volume.
auth_group = %s	The authentication group assigned to this volume.
volume_name = %s	The volume name that is used to create the local EBSD disk.
access_mode = r | ro | rw	r or ro= read only rw = read+write The access mode the driver should use to access this volume.
use_unicast = true | false	Whether the driver should use unicast discovery.
unicast_list = x.x.x.x, x.x.x.x	Used if the use_unicast flag is set to true. Coma separated list of ip addresses to be used for unicast discovery.
*use_multicast = true | false	Whether the driver should use multicast discovery.
* Use either unicast or multicast. Do not use both together.	

D.4.1.1 Sample Device Entry in /etc/ebsd.conf

```
########################
# Sample device entries:
[device0]
type = volume
client_name = myclient
ip_bind = 10.0.1.63
management_group = my_mgtgroup
```

```
auth_group = public
volume_name = my_volume_0
access_mode = rw
use_unicast = true
unicast_list = 10.0.0.12, 10.0.0.13
use_multicast = false
########################
[device1]
type = volume
client_name = myclient
ip_bind = 10.0.1.63
management_group = my_mgtgroup
auth_group = public
volume_name = my_volume_1
access_mode = rw
use_unicast = true
unicast_list = 10.0.0.12, 10.0.0.13
use_multicast = false
```

D.4.2 Connecting the EBSD Driver to the SSM EBSD Server

1. Run the startup script which will spawn a child process for the new device.

`/etc/init.d/ebsd start`

   loads the driver.

2. Start and wait until all devices are online.

   There is a timeout of 2 minutes.

D.4.3 Verifying EBSD Devices

The EBSD driver creates a block device for each volume.

1. Check the current status of the device entries. Use

`cat /proc/ebsd/client`

   It should say "online."

```
suzy1:~ # cat /proc/ebsd/client
Version:04/04/03,4.1.14.0004
Majors:176:177:178:
Device: 0 ( my_mgtgroup_0:my_volume_0 )
    Status:   Online ( Active )
    Read:     0 B (Requests: 0 )
    Write:    0 B (Requests: 0 )
    Ops:      0 ( sync = 0 )
    Cycles:   147
    BSize:    512
    Capacity: 52428800 kb
Device: 1 ( my_mgtgroup_0:my_volume_1 )
    Status:   Online ( Active )
    Read:     0 B (Requests: 0 )
    Write:    0 B (Requests: 0 )
    Ops:      0 ( sync = 0 )
    Cycles:   147
    BSize:    512
    Capacity: 419430400 kb
```

Figure 264. Sample cat /proc/ebsd/client

Table 60. Parameters for /proc/ebsd/client

Parameter	What It Is
Status	Status of the device • Starting • Deleted • Online - may be either Active or Lost Manager
Read	Amount of data read and the number of read requests in bytes (KiB, MiB, or GiB)*
Write	Amount of data written and the number of write requests in bytes
Ops	Combined total of read requests and write requests
Cycles	Internal - the number of times the EBSD task looped
BSize	Block size in bytes
Capacity	Size of the attached volume in kilobytes
* KiB, MiB, and GiB are calculated in increments of 1024 bytes. For example, 1 KiB = 1024 bytes; 1 MiB = 1024 KiB; 1 GiB = 1024 MiB	

2. To verify that the block devices were created, use
`ls -la /dev/ebsd/`

```
root@lnx-demo /root]# ls -la /dev/ebsd
total 0
crwxr-xr-x   1 root    root    176,  0 Apr  4 14:39 ebsdctrl
brwxr-xr-x   1 root    root    177,  0 Apr 10 17:53 disk0
brwxr-xr-x   1 root    root    177,  1 Apr 10 17:53 disk1
```

Figure 265. Sample ls -la /dev/ebsd/

- You can format the block devices by using any of the OS filesystem utilities.

  For example, you can use

`mkfs -t ext2 /dev/ebsd/disk0`

D.4.4 Mounting the Block Device EBSD Disk

Once the ext2 filesystem is created the disk can be mounted. For example:

1. Make a mount point for the disk:

   For example,

`mkdir /mnt/ebsd0`

2. Mount the EBSD disk.

   For example,

`mount /dev/ebsd/disk0 /mnt/ebsd0`

   At this point you can treat the mounted disk like any other OS file directory. You can copy files, add and delete files, and perform other file functions there.

D.5 Adding an EBSD Disk at Runtime

1. Modify ebsd.conf.
2. Run aebsvm --add-all.

   The new disk is added to /dev/ebsd/. Use it as a raw or block device. See "Verifying EBSD Devices".

D.6 Starting the EBSD Driver

If enabled for the current run level, the EBSD service is started when the operating system is booted.

You can also start the EBSD service manually using the following command(s):

```
service ebsd start
```

or

```
/etc/init.d/ebsd start
```

The EBSD service reads the file /etc/sysconfig/ebsd to initialize the EBSD volume manager tool (aebsvm) and the EBSD configuration file (ebsd.conf).

The file contains the following default information:

```
aebstool=/usr/local/sbin/aebsvm
ebsdconf=/etc/ebsd.conf
```

You can modify this file for different directories and names.

The EBSD start service first checks for the EBSD configuration file (ebsdconf environment variable), and the EBSD volume manager tool (ebsdtool environment variable). If the configuration file and the EBSD volume manager tool exist, then the EBSD service:

- Loads the EBSD driver if needed.
- Starts all the devices listed in the EBSD configuration file ($ebsdconf).
- Waits until all the devices become online or exceed the timeout value.
- Mount all the EBSD devices listed in /etc/fstab and not marked as noauto.

D.7 Stopping the EBSD Driver

The EBSD stop service first checks for the EBSD configuration file (aebsconf environment variable), and the EBSD volume manager tool (aebstool environment variable). If the configuration file and the EBSD volume manager tool exist, then the EBSD service:

- Reports a system hang warning if a mounted device has lost connection.
- Stops all processes using mounted EBSD devices.
- Unmounts all mounted EBSD devices.
- Stops and removes all the EBSD devices.
- Unloads the EBSD driver.

D.8 Status of the EBSD Driver and Devices

To display the status of the EBSD driver and the associated devices, enter the following:

```
service ebsd status
```

or

```
/etc/init.d/ebsd status
```

This status command displays the information in the file /proc/ebsd/client, shown in Figure 264.

D.9 Disconnecting an EBSD Device

Disconnecting an EBSD device stops activity on the SSM but preserves the data in the volume.

- Prerequisite

Stop all applications using the EBSD device.

Note: If an application attempts to write to a disconnected or disabled EBSD device (raw or file), the application will hang.

D.9.1 Unmounting the EBSD Disk

Unmount the disk.

1. Execute the command
umount
for example:
`umount /mnt/ebsd0`
2. Run
ebsd`vm --remove #`
where # = the device you want to disconnect. (Remember, device numbers can be from 0 - 63.) This disconnects the SSM volume from the host device entry.

D.10 Deleting an EBSD Device

Deleting the EBSD device erases the volume's data from the SSM.

Note: Make sure the EBSD disks you plan to delete are NOT in use. Deleting EBSD disks from the host and deleting volumes from the cluster removes all data stored in those volumes. Once removed, that data cannot be retrieved.

Note: Be sure to disconnect the device before deleting it.

1. Modify the /etc/ebsd.conf file.
   — Remove the appropriate device entry in ebsd.conf.

D.11 Uninstalling the EBSD Driver for Linux

1. Navigate to the following directory
/opt/Storage_System/Storage_System_Software/6.0/Drivers
   — Copy the file install.sh to your client system.
2. Run
`./install.sh -u`
3. Type cd.. and press Enter.
4. Run
`./Uninstall_EBSD_Driver`
The Install wizard opens.
5. Click Uninstall.

6. Click Done when the wizard is finished uninstalling the driver.
7. Type cd.. and press Enter.

D.11.1 Finishing Up

Finally, remember to remove the EBSD volumes from the Storage System Console.

D.11.2 Troubleshooting

D.11.2.1 Error: Could not Load the EBSD Driver on your System

1. Run uname -r to determine what version of Linux you are running.
2. Determine whether that version is supported by the EBSD driver.

D.11.2.2 Driver Successfully Loaded but Adding Device Returns Failed (i.e. aebsvm --add 0 returns “failed”)

1. Check the error message and correct the problem.

D.11.2.3 Driver Successfully Loaded, Adding a Device Appears Successful, but when you Check the config file, the Device was not Added

1. Check cat /proc/ebsd/client.
2. If device does not exist, add proper entry into the EBSD config file.

   or

   If device status stuck in starting mode, check the following:

   — Verify all entries in the EBSD config file.

   — Verify the network connection to the SSMs.
3. Re-check cat /proc/ebsd/client after making corrections to any issues found.

D.11.2.4 During Unmounting

If during unmounting you get "device/filesystem busy" you probably have a process accessing /dev/ebsd/disk0.

Using the EBSD Driver for Windows 2003 E

E.1 Recommended Configuration

- Windows 2003
- Pentium III processor or greater
- 1 GB minimum RAM
- 10 MB minimum free hard disk space

E.2 Overview of EBSD Driver for Windows 2003

Install and configure the EBSD driver for Windows Server 2003 on any computer that accesses volumes on a cluster of SSMs.

To configure a client to access a volume, you must install and configure the driver and configure the volume to be accessed. This manual describes all of the driver configuration tasks. Volume configuration tasks are described in the **User Manual** in the chapter entitled Working with Volumes. The table below lists the required tasks and where to find information about them.

Table 61. EBSD Driver Configuration Tasks

Configuration Task	Configuration Tool	Instructions
Create a volume and associate the volume with an authentication group.	Storage System Console	In the Users Manual • Working with Volumes chapter, see “Creating a Volume” • Working with Authentication Groups chapter, see “Creating an Authentication Group” and “Creating an Authentication Group Association”
Install the EBSD driver on a Windows computer	EBSD driver	“Installing the EBSD Driver”
Create a disk on the EBSD client computer.	EBSD driver	“Adding EBSD Disks to Your System”
Write the disk signature.	Windows Disk Manager on client computer	“Initializing New Disks”
Partition a basic disk.	Windows Disk Manager on client computer	“Partitioning Basic EBSD Disks”

E.3 Installing or Updating the EBSD Driver

E.3.1 Installation Overview

The EBSD for Windows installation CD provides two installation options for the EBSD driver.

- Install a driver on a local machine directly from the CD. If you have an earlier version of the EBSD driver already installed on the local machine, the installation wizard directs you to the location where you can update the driver. See "Updating the EBSD Driver".
- Copy the EBSD driver files to a local machine or a network share. Then install the driver individually on local machines using the Windows Add Hardware wizard.

Note: You need administrative privileges during installation and configuration.

Note: Throughout the procedures in this guide, the term "EBSD client" refers to any computer, such as an application server, that accesses volumes on a cluster of SSMs.
The term "disk" refers to the EBSD disks that you create on the EBSD client computer, while the term "volume" refers to the volumes created on Storage System Software and the Storage System Console, and to any volumes that are created on EBSD disks using the Windows Disk Management tool.

E.4 Copying the EBSD Driver Files [Optional]

Skip this section if you only want to install or update the EBSD driver from the CD. Go directly to "Installing the EBSD Driver", or "Updating the EBSD Driver".

1. Insert the resource CD that came with your system into the CD drive of the EBSD client PC.
   The installation wizard should automatically open. If the installation wizard does not open, navigate to the CD drive, the InstData\VM\ folder, and click **EBSD60_setup.exe**.
2. On the Choose Product window of the installation wizard, select EBSD Driver Bundles and click Next.
   The Choose Install Folder window opens.
3. Choose to accept the default directory (C:\ProgramFiles\Storage_System\Storage_System_Software\6.0\Drivers) or browse to the directory where you want the EBSD folder to be copied. This location can be another location on the network, such as a file server.
4. Click Next.
   Review the Pre-installation Summary window.
5. Click Install.
   The folder, EBSD, is copied into the directory you specify.

E.5 Installing the EBSD Driver

You can install the EBSD driver locally using the CD or using the folder from a network share. See "Copying the EBSD Driver Files [Optional]" on page 316.

The EBSD driver uses the Microsoft® Windows® 2003 DiskPart utility for volume management.

E.5.1 Beginning the Driver Installation

Table 62. Driver Installation

Installing from the CD	Installing from a Network Share
1. Insert the EBSD for Windows CD into the CD drive of the EBSD client PC. The installation wizard should automatically start. If not, run **EBSD60_setup.exe** from the InstData\VM\ folder on the CD. 2. On the Choose Product Component window, select EBSD Driver. 3. After completing the installation a message opens describing how to locate the driver using the Add Hardware wizard. 4. Click OK. The Add Hardware wizard opens. 5. Continue with Step 1 below.	1. Copy the EBSD folder to the computer on which you want to install the driver. 2. Open the Add Hardware wizard from the Control Panel. 3. Continue with Step 1 below.

Figure 266. Opening the Add Hardware Wizard

1. On the Welcome to the Add Hardware wizard window, click Next.
   The wizard asks if you have connected the hardware, shown in Figure 267.

Figure 267. Is the Hardware Connected?

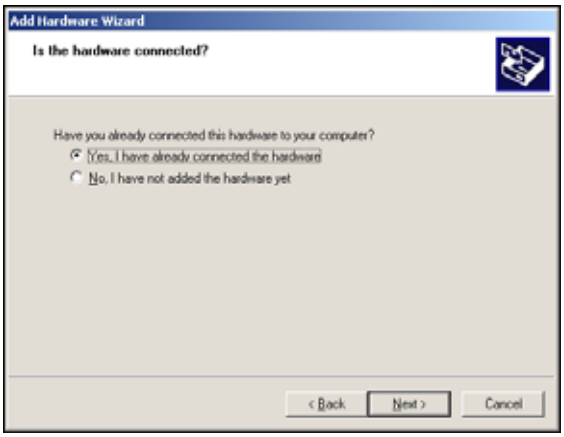

2. Click Yes, the hardware is connected, and click Next.
   The list of installed hardware opens, shown in Figure 268.

Figure 268. Selecting Add new Device from the Hardware List

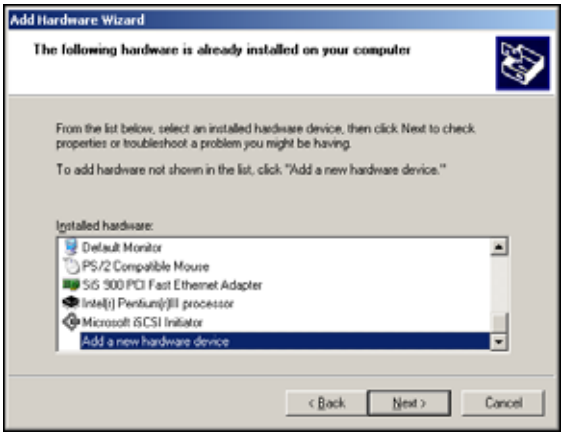

3. Scroll to the bottom of the list and select “Add a new hardware device,” and click Next.
   The Search for hardware window opens, shown in Figure 269.

Figure 269. Choosing the Advanced Option to Manually Select the Hardware from a List

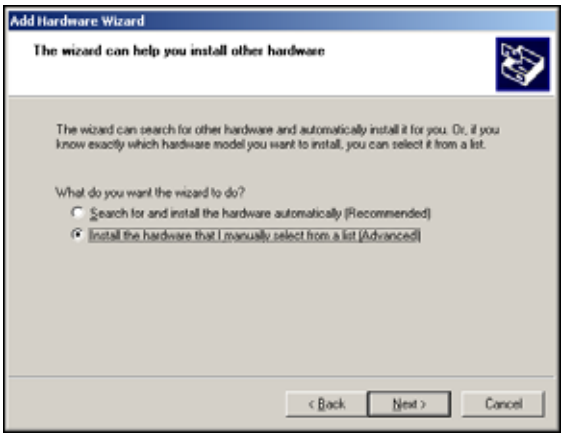

4. Select “Install the hardware that I manually select from a list (Advanced)” and click Next.
   The Common hardware types list window opens, shown in Figure 270.

Figure 270. Selecting Show All Devices

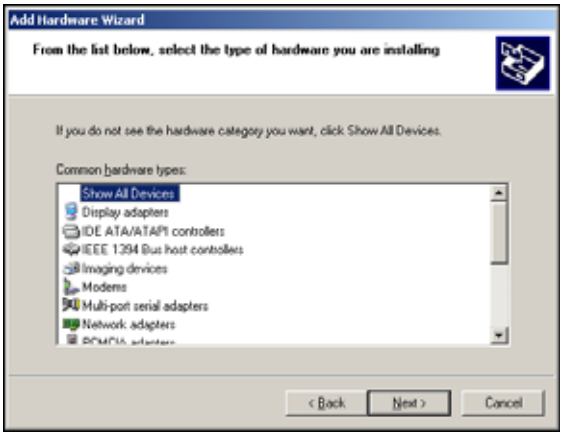

5. Select Show All Devices and click Next.
   The Select Device Driver window opens, shown in Figure 271.

Figure 271. Selecting a Device Driver

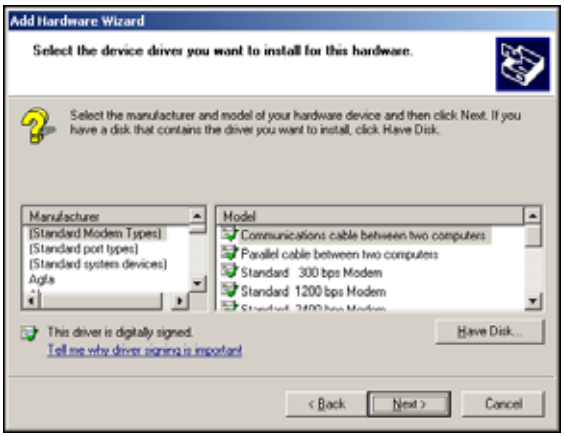

6. Click Have Disk.

   The Install From Disk window opens, shown in Figure 272.

Figure 272. Installing from Disk

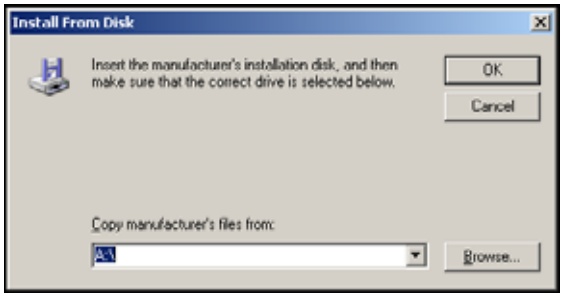

E.5.2 Locating the EBSD Driver Files

1. Browse to the C:\ProgramFiles\Storage_System\ Storage_System_Software\6.0\Drivers folder where the **aebs.inf** file is located.

   You can install the driver from any directory into which you have copied the files.

2. Select the **aebs.inf** file and click Open.

   Focus returns to the Install From Disk window with the directory containing the selected file displayed, shown in Figure 273.

Figure 273. Install from Disk Window

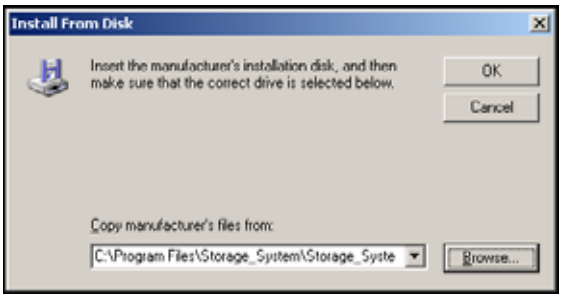

3. Click OK.

   Focus returns to the Select a Device Driver window with the EBSD driver displayed, shown in Figure 274.

Figure 274. Selecting the EBSD Driver

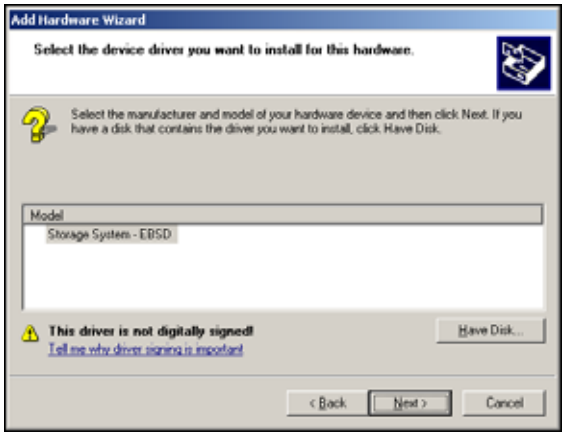

4. Select the EBSD Driver and click Next.

   The Start Device Driver Installation window opens, shown in Figure 275.

Figure 275. Verifying the Driver

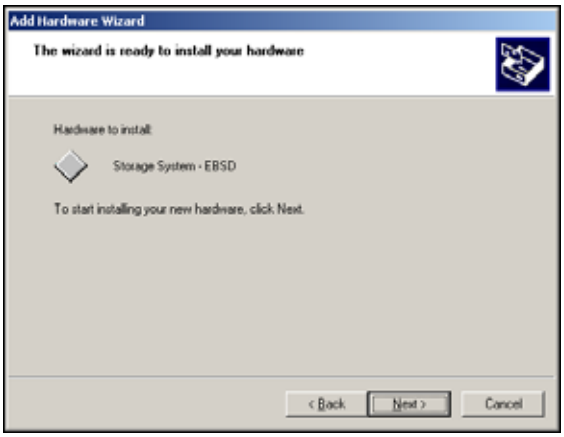

5. Verify that the EBSD Driver appears in the window and click Next.

Note: You may see a Security Alert message asking you to verify that you want to install the driver software. Click Yes to continue installing the driver.

The EBSD System Configuration window opens, shown in Figure 276.

Figure 276. Configuring file Server, SQL Server, and Exchange Services to come Online after a Reboot

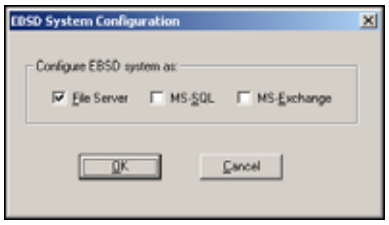

6. [Optional] If you want to automatically configure services and applications to come online after a reboot, select the appropriate boxes.

   You can change these settings later in the EBSD driver advanced settings. See “Configuring File Services, SQL Server, and Exchange” on page 346.

   The Finish EBSD installation window opens, shown in Figure 277. Instructions for modifying the driver configuration are on the window.

Figure 277. Finishing Installation and more Instructions

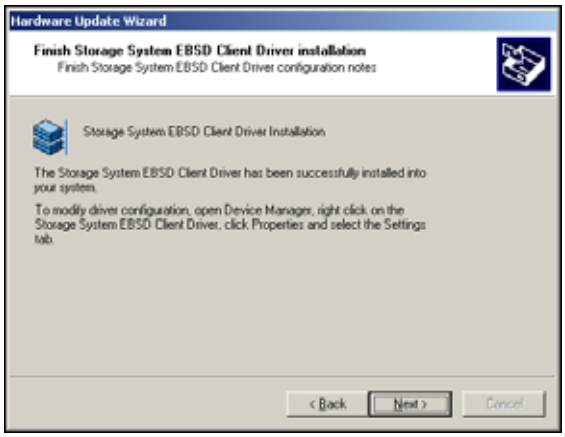

7. Click Next.

   The Completing the Add Hardware window opens, shown in Figure 278. Review the window to verify that it is the EBSD driver that is being installed.

Figure 278. Completing the Add Hardware Wizard

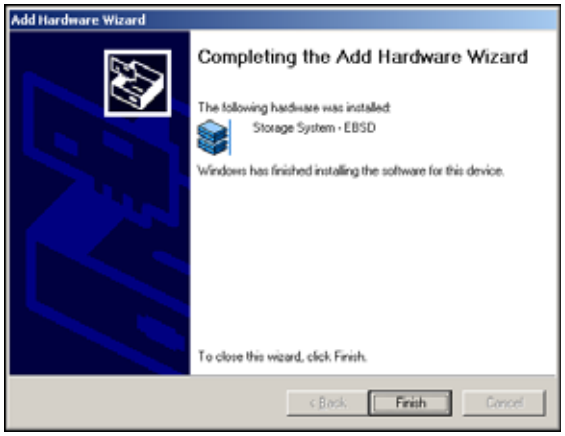

8. Click Finish to complete the driver installation.

E.6 Updating the EBSD Driver

When you are updating the EBSD driver, the installation wizard directs you to the Windows Device Manager.

Table 63. Updating the EBSD Driver

Updating from the EBSD for Windows CD	Updating from your vendors web or FTP site (if applicable).
1. Insert the EBSD for Windows CD into the CD drive of the EBSD client PC. Your browser should automatically open and start the installation wizard. If not, run **EBSD60_setup.exe** from the InstData\VM\ folder on the CD. 2. On the Choose Product Component window, select EBSD Driver. 3. Click Next. Review the Pre-installation Summary window. 4. Click Install. The files are installed to the location you specified and a message opens. 5. Review the message and click OK. The message displays the location of the driver update files. Those files are in the C:\ProgramFiles\Storage_System\ Storage_System_Software\6.0\Drivers folder. The Update Installation wizard then opens the Windows Device Manager. 6. Continue with step 1 below.	1. Using a web browser, open your vendors web or FTP site where the files are stored. 2. Open the folder for the release version that you are upgrading to. 3. Copy install.htm and the InstData folder to your hard drive. 4. Navigate to install.htm on your hard drive. 5. Double-click install.htm. 6. Complete the installation wizard. 7. Continue with step 1 below.

Note: Install the update in the same directory as the original EBSD driver installation. The default directory is C:\ProgramFiles\Storage_System\ Storage_System_Software\6.0\Drivers.

E.6.1 Updating the Device Driver in the Windows 2003 Device Manager

1. Expand the Network Storage Devices and select the EBSD driver, shown in Figure 279.

Figure 279. Selecting the EBSD Driver

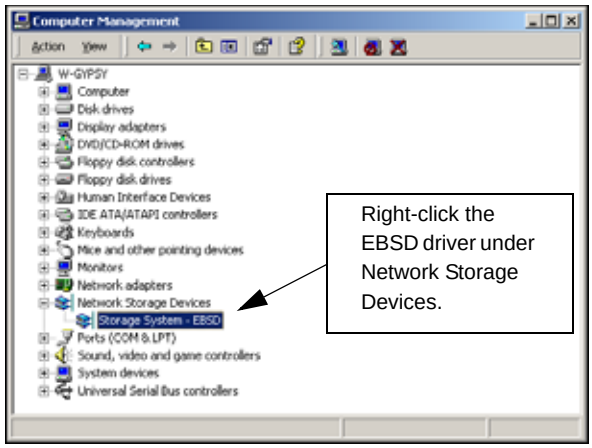

2. Right-click on the driver and select Properties.
   The EBSD Properties window opens, shown in Figure 280.

Figure 280. Updating the EBSD Driver

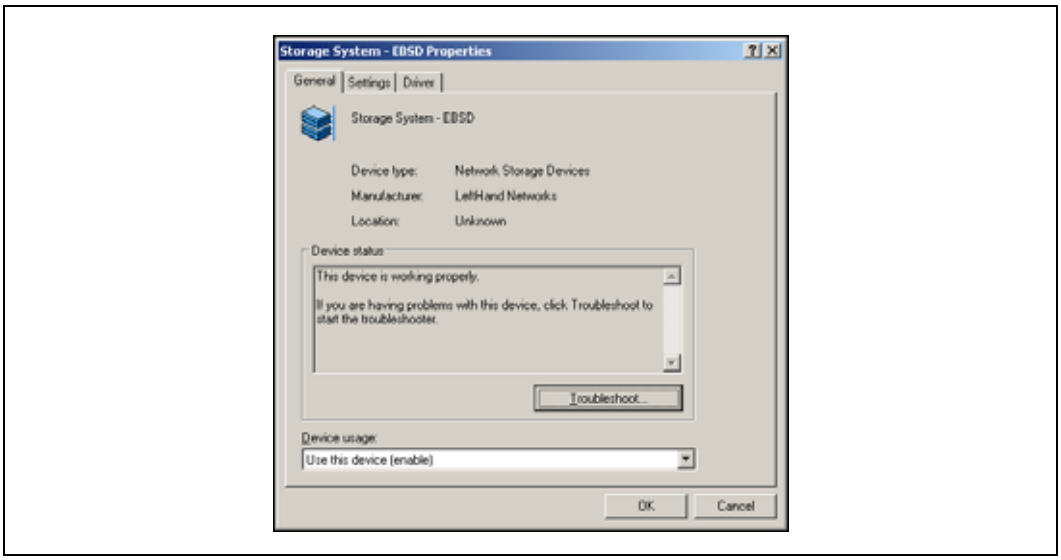

3. Select the Driver tab, shown in Figure 281.

Figure 281. Checking the Driver Version Number

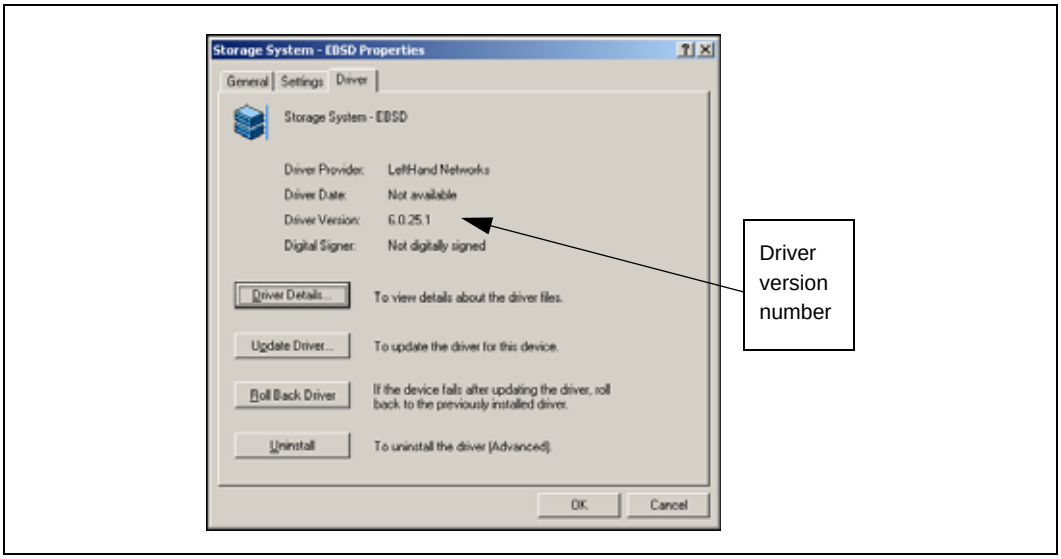

Note: The driver version number is displayed on the Driver tab.

4. Click Update Driver.
   The Hardware Update wizard opens, shown in Figure 282.

Figure 282. Updating the Driver

5. Select “Install from a list...” and click Next.
   The Choose search and installation options window opens, shown in Figure 283.

Figure 283. Choosing Search and Installation Options

6. Select “Don’t search. I will choose the driver to install.” and click Next.
   The Select a device driver window opens with the EBSD Driver displayed.

Figure 284. Selecting the Device Driver

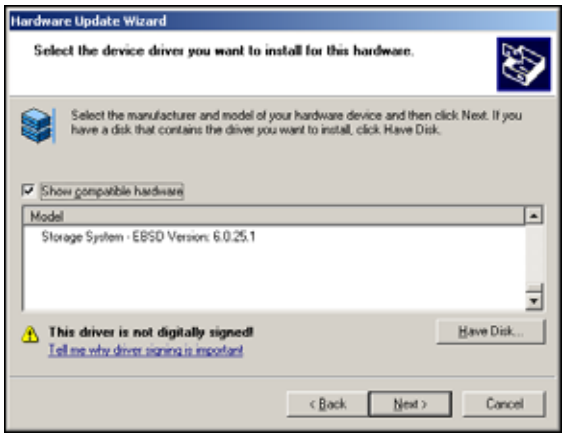

7. Click Have Disk.
   A browse window opens where you can specify the location of the driver files.

Figure 285. Browsing for the Driver Files

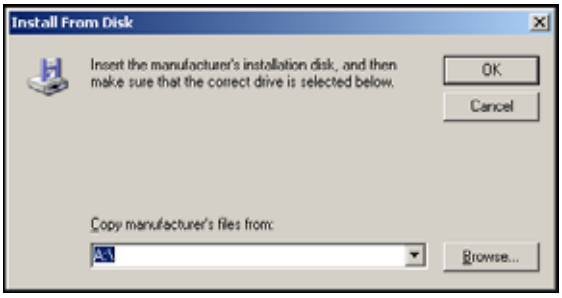

8. Click Browse to navigate to the location where you installed the EBSD driver update files.
   The files are in C:\ProgramFiles\ Storage_System\Storage_System_Software\6.0\Drivers.
9. Select the **aebs.inf** file and click Open.
   Focus returns to the Install from Disk window.

Figure 286. aebs.inf File Selected

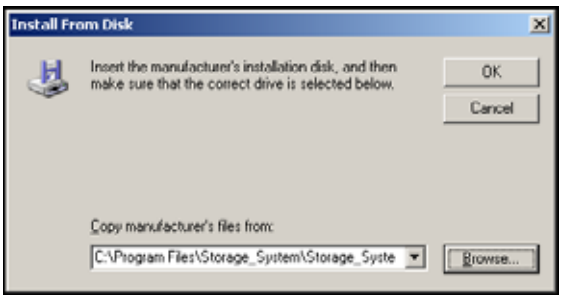

10. Click OK.

    Focus returns to the Select a Device Driver window.

Figure 287. Selecting the Device Driver

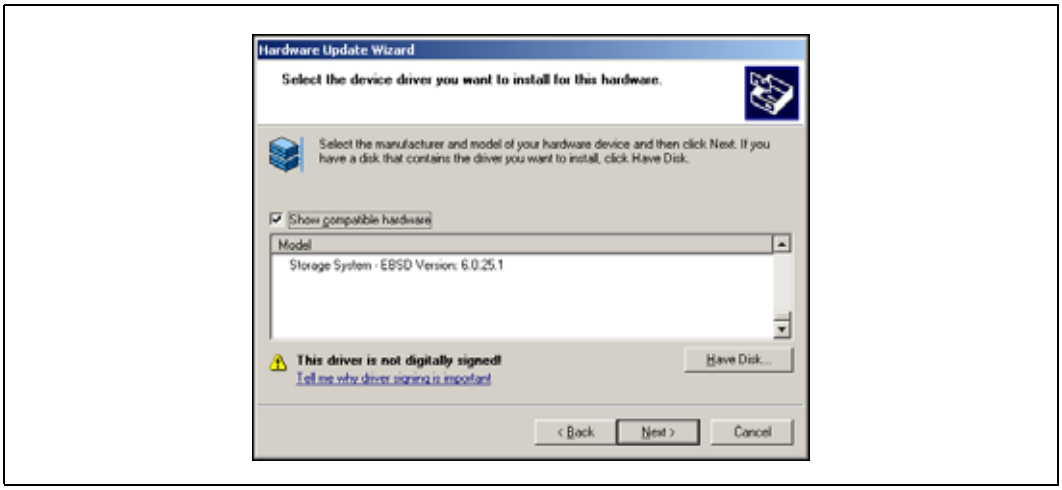

11. Select EBSD and click Next.

    The Start Device Driver Installation window opens.

Figure 288. Starting the Update Installation

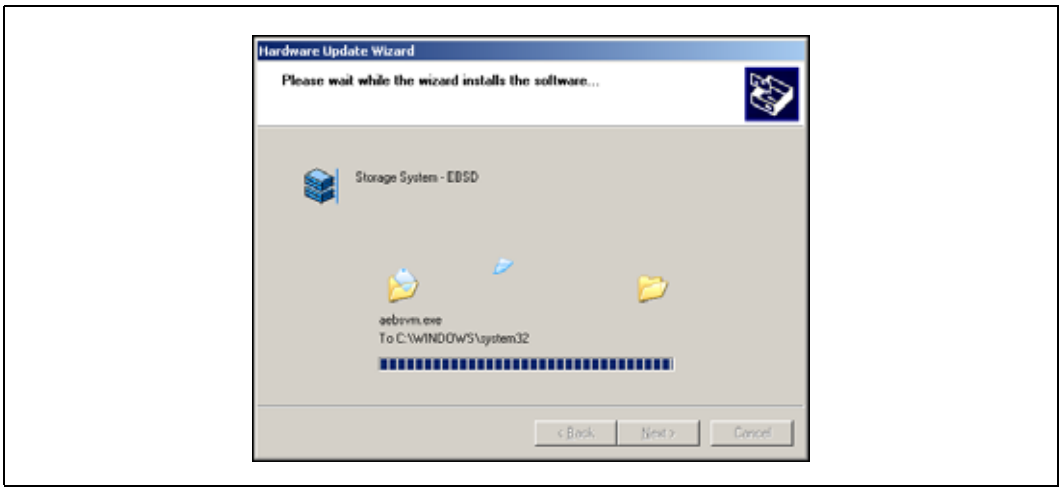

12. Verify that the EBSD driver appears in the window and click Next.

Note: You may see a Security Alert message asking you to verify that you want to install the driver software. Click Yes to continue installing the driver.

The EBSD System Configuration window opens, shown in Figure 289.

Figure 289. Starting File Server, SQL Server, and Exchange Services

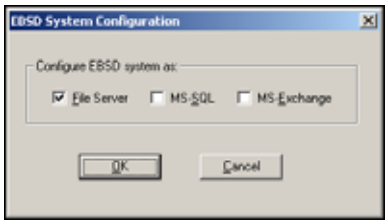

13. [Optional] If you want to automatically configure services and applications to come online after a reboot, select the appropriate boxes.

    You can change these settings later in the EBSD driver advanced settings. See "Configuring File Services, SQL Server, and Exchange" on page 346.

    The Finish Device Driver Installation window opens.

Figure 290. Finishing the Update Installation

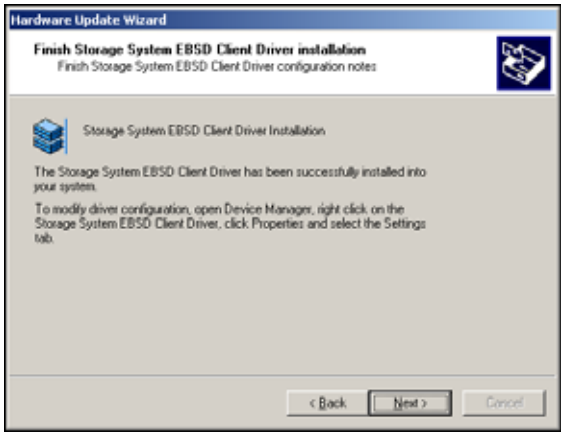

14. Click Next.

    The Completing the Upgrade Device wizard window opens.

Figure 291. Completing the Upgrade Wizard

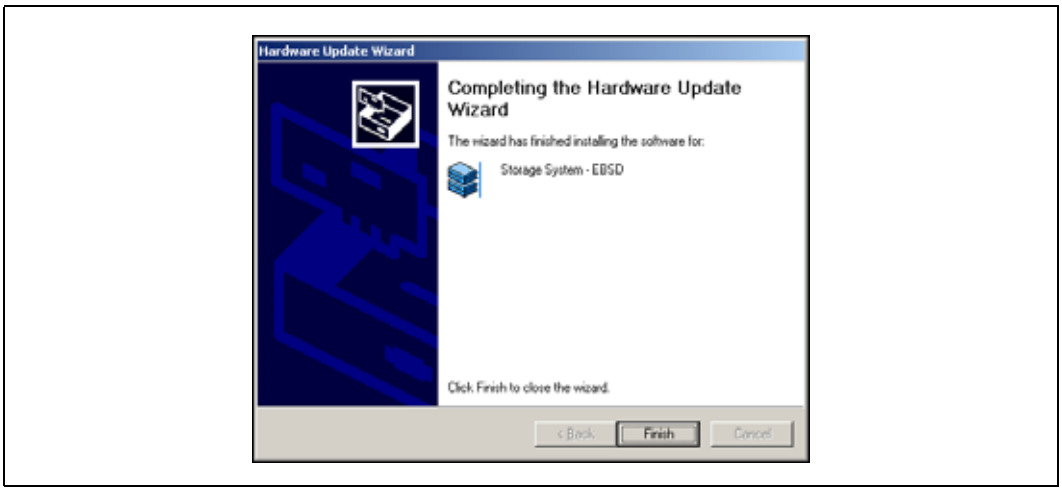

15. Verify that it lists EBSD and click Finish.
    Focus returns to the EBSD Properties window.

Figure 292. Closing the EBSD Driver

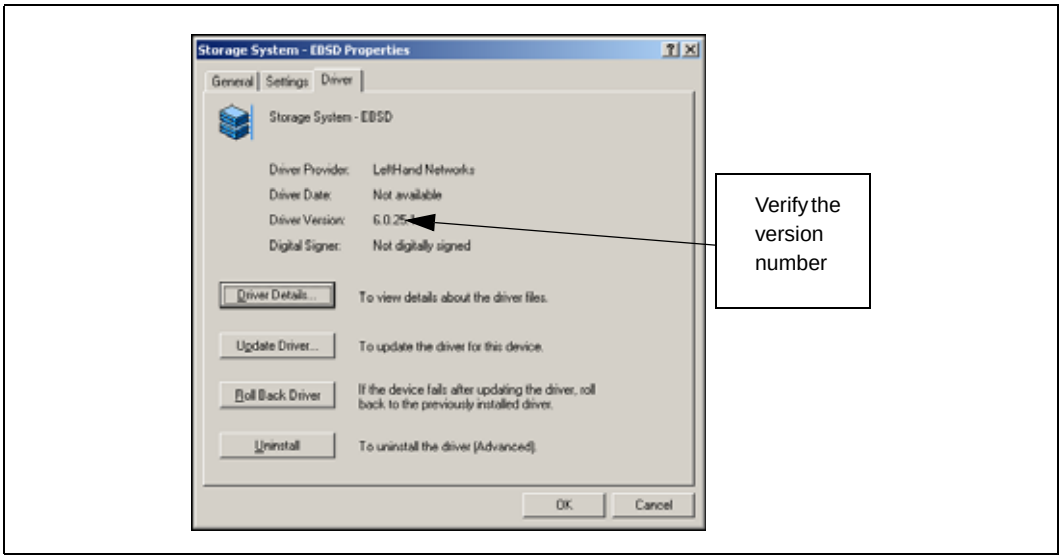

16. Verify that the version number has changed to reflect the upgrade version.
17. Click Close to close the EBSD driver.
    A message opens, prompting you to reboot the computer.
18. Click Finish.
    A message opens notifying you to restart your computer in order for the settings to take effect.
19. Click No.
20. Complete the steps in the Driver Installation wizard.
21. Manually reboot your computer to apply the driver settings.

See "Configuration Overview" for information about configuring the client drive.

E.6.2 Rolling Back the Driver Update

To return to the previous version of the driver, use the Roll Back feature found in Windows 2003.

1. Open the EBSD Properties window, shown in Figure 292.
2. Click Roll Back.
   A confirmation message opens, verifying that you want to roll back to the previous driver.
3. Click Yes.
   A warning opens, stating that the driver does not have a signature.
4. Click Yes to continue.
   A number of confirmation messages may open.
5. Follow any instructions in the message boxes.
6. Reboot to complete the rolling back.
7. Verify that the correct driver version number is shown in the EBSD Properties window, shown in Figure 292.

E.7 Configuring the EBSD Driver

E.7.1 Configuration Overview

Once the EBSD driver is installed, it must be configured.

Note: You need administrative privileges during installation and configuration.

Note: Configure volumes and associate them with authentication groups in the Storage System Console before configuring the EBSD driver. You use information about the volumes, including the management group configuration, volume name, and authentication group name, to configure EBSD disks.

E.8 Opening the EBSD Driver

1. Open Windows Device Manager, shown in Figure 293.

Figure 293. Selecting the EBSD Driver

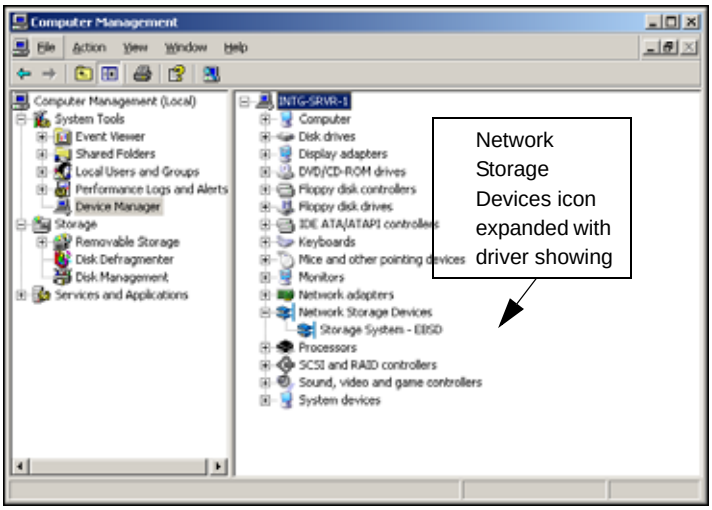

2. Expand the Network Storage Devices list and select the EBSD driver.
3. Double-click on the driver or click the Action menu and select Properties.
   The EBSD Properties window opens, shown in Figure 294.

Figure 294. EBSD Properties Dialog

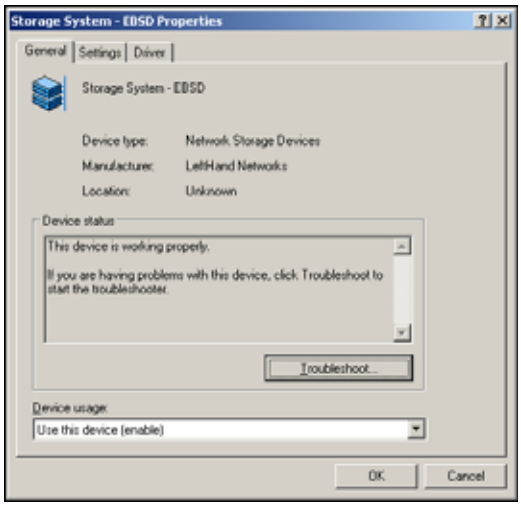

E.9 Adding EBSD Disks to Your System

An EBSD disk maps to a volume on the IXA SDK. Before you begin adding EBSD disks, use the Storage System Console to locate each volume the EBSD client will access and write down the following information.

- Management group name
- IP addresses of all managers in the management group
- Volume name
- Name of the authentication group that is associated with the volume

When you create a volume on the EBSD client, you will need to enter this information exactly as it appears in the Console.

1. Click the Settings tab to bring it to the front, shown in Figure 295.

Figure 295. EBSD Driver Settings

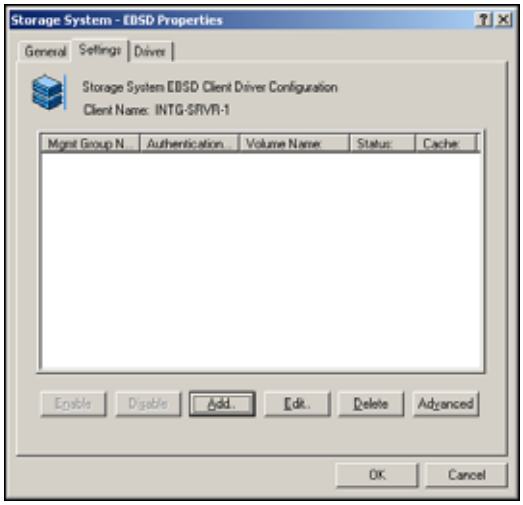

2. Click Add to add EBSD disks.
   The Add EBSD Volume window opens, shown in Figure 296.

Figure 296. Adding an EBSD Disk

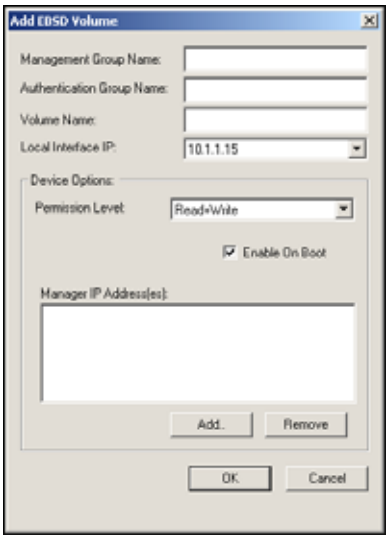

Use the Add EBSD Volume window to create a disk that corresponds to a volume that exists on a cluster of SSMs.

3. Complete the fields in the Add EBSD Volume window.

   SeeTable 64 for a list of field descriptions and requirements for completing the Add EBSD Disk window

Note: Be sure to type the names of the management group, authentication group, and volume exactly as they appear in the Storage System Console.

Table 64. Requirements for Adding an EBSD Disk

Add EBSD Disk Field	Description and Requirements
Management Group Name	Type the name of the management group that contains the volume.
Authentication Group Name	Type the name of the authentication group that is associated with the volume in the Storage System Console.
Volume Name	Type the name of the volume exactly as it appears in the Storage System Console.
Local Interface IP	Select from the list the IP address of the computer that is running the EBSD driver. If the client computer has more than one NIC, select the IP address of the NIC that you want the client to use to access the volume.

Table 64. Requirements for Adding an EBSD Disk

Add EBSD Disk Field	Description and Requirements
Permission Level	Default = Read+Write Select the permission level of the EBSD disk The permission level for the disk cannot be greater than the permission level of the authentication group associated with the volume in the Storage System Console. For example, if the authentication group has read only permissions, you cannot give the EBSD disk read and write permission.
Enable on Boot	Default = Enabled (checked) Ensures that volumes come online after a reboot. ***Note:*** Change this to Disabled only if you are using 3rd party clustering software such as Veritas Cluster Server or Microsoft Cluster Server.
Manager IP Addresses	Enter the IP address of at least one of the managers in the management group containing the volume. ***Note:*** To ensure volume availability, enter the IP addresses of all managers in the management group. If the manager the EBSD client is using to access the volume becomes unavailable, the EBSD client can use any of the other managers to access the volume.

4. Click OK when you have finished completing the EBSD disk information.

   The new EBSD disk appears in the list on the Settings tab, shown in Figure 297. The Status column displays Starting, and then changes to Online when the new disk is ready.

Figure 297. Listing of EBSD Disks

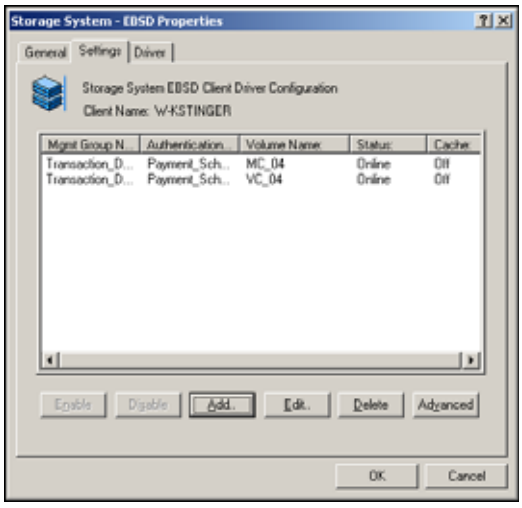

5. Repeat steps 2 through 4 for each EBSD disk you want to add.
6. Click OK when you are finished.

The EBSD Properties window closes.

E.10 Enabling Write Cache on Volumes

Write cache is disabled by default on EBSD volumes. The cache status is shown in the Driver Settings window as “on” or “off.”

Write cache can be enabled on read+write volumes. It cannot be enabled on read only volumes or snapshots. Also, if the system does not have enough memory resources, write cache cannot be enabled and an event log message will be recorded.

E.10.0.1 The Write Through Command

The EBSD driver automatically supports the Write Through command when write cache is enabled. This command, which is set by applications (for example, SQL Server), provides an additional level of safety by causing certain critical packets of data to be written immediately to the disk, bypassing the cache.

E.10.1 Requirements for Changing Write Cache

- All data transfers to or from the volume must be completed before enabling or disabling write cache. You cannot change the write cache setting when data is being written.
- To enable write cache, the system should have at least 256MB free RAM. Verify the amount of available RAM in Windows Task Manager on the Performance tab.

Note: When enabling write cache, using standard UPS power protection to prevent possible data loss is recommended.

E.10.2 Enabling Write Cache

1. Open the Windows Device Manager.
2. Expand the Disk Drives list.

   The EBSD disks are listed as shown in Figure 298. See “Identifying the Storage System Software Volume That Corresponds to an EBSD Disk” if you need to verify which disk you are working with.

Figure 298. Viewing EBSD Disks

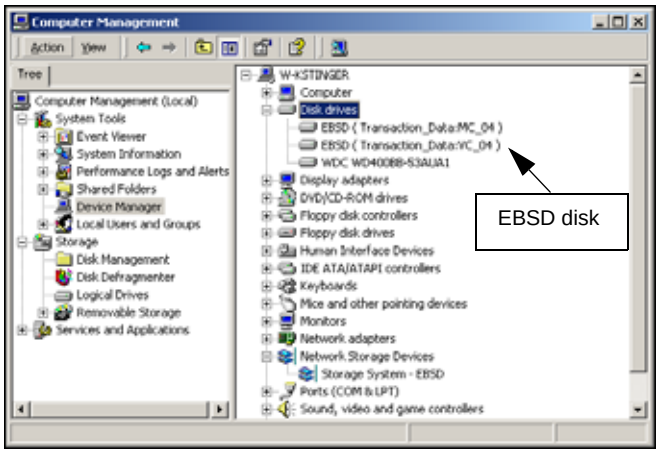

3. Select the disk for which you want to enable write cache.
4. Double-click on the disk or right-click and select Properties from the menu. The Disk Properties window opens, shown in Figure 299.

Figure 299. Opening the Disk Properties from the Device Manager

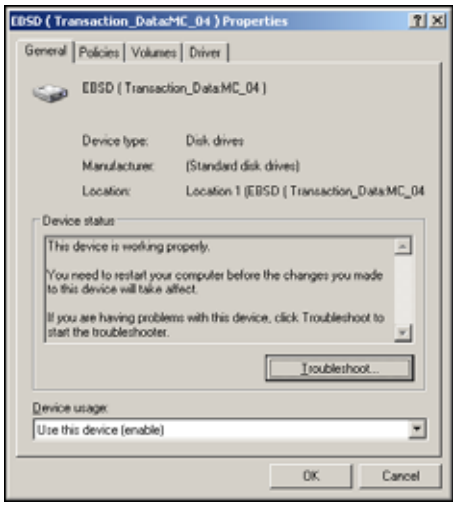

5. Click the Policies tab to bring it to front, shown in Figure 300.

Figure 300. Opening the Policies Tab

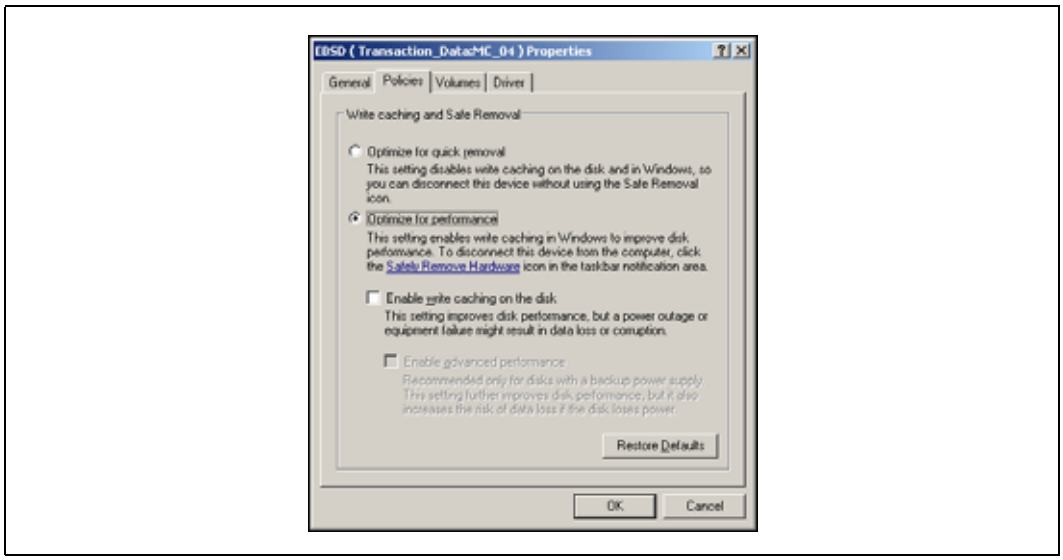

6. Select Optimize for performance.
7. Select the check box to enable write cache.
8. Click OK.
   Focus returns to the Policies tab.
9. Click OK to close the Disk Properties window.

E.10.2.1 Verifying Write Cache Status

Verify that the write cache is enabled after changing the setting.

1. Open the EBSD driver.
2. Click the Settings tab to bring it to the front.

Figure 301. Viewing the Status of Write Cache on the EBSD Volumes

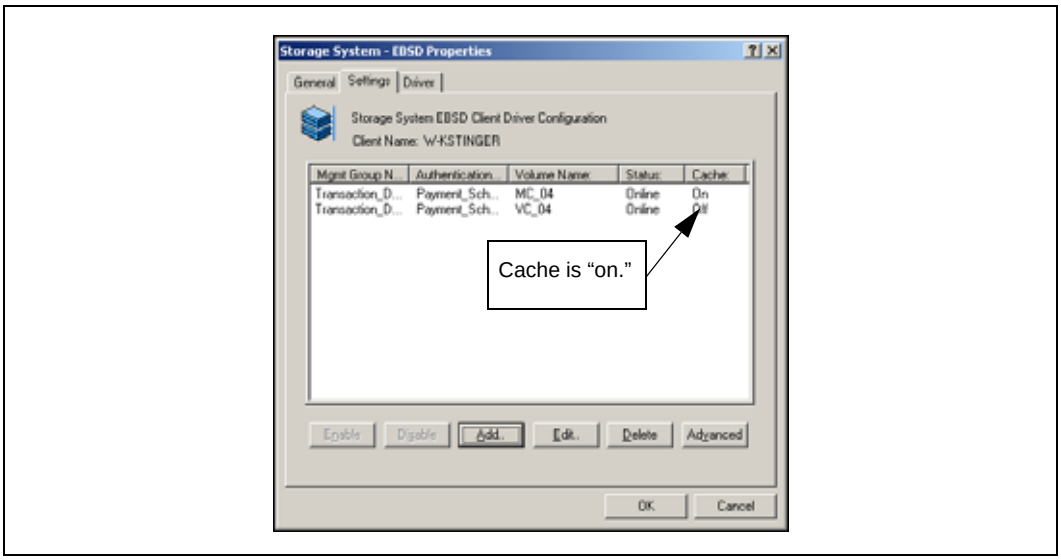

3. Verify that cache is on for those volumes on which you enabled it.

   The “Optimize for performance” setting on the Policies tab of the Disk Properties window will now match the EBSD Properties setting.

E.10.3 Disabling Write Cache on Volumes

Make certain that all data transfers are complete before disabling write cache.

1. Open the Windows Device Manager.
2. Expand the Disk Drives list.

   The EBSD disks are listed as shown in Figure 298. See “Identifying the Storage System Software Volume That Corresponds to an EBSD Disk” to verify which disk you are working with.
3. Select the disk for which you want to disable write cache.
4. Double-click on the disk or right-click and select Properties from the menu.

   The Disk Properties window opens, shown in Figure 299.
5. Select the Policies tab.
6. Clear the check box to disable write cache.
7. Click OK to close the Disk Properties window.

E.11 Initializing New Disks

Next you initialize the new EBSD disks.

1. Open Disk Management.

   The Initialize and Convert Disk wizard opens, shown in Figure 302.

Note: If the wizard does not open automatically, right-click the disk labeled "Not Initialized" and select Initialize Disk.

Figure 302. Opening the Initialize and Convert Disk Wizard

2. Click Next to continue the wizard.
   The Select disk to initialize window opens, shown in Figure 303.

Figure 303. Selecting Disks to Initialize

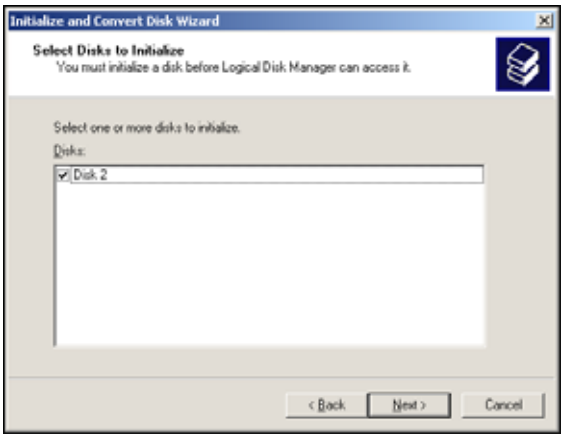

3. Select the disks that you want to initialize and click Next.
   The Select disks to convert window opens, shown in Figure 304.

Note: It is recommended that you use only basic disks with the Storage System Software system. For a detailed discussion of basic disks, disk management programs, and expanding basic partitions, visit the support section of the Company Name web site.

Figure 304. Clearing Disk Selection

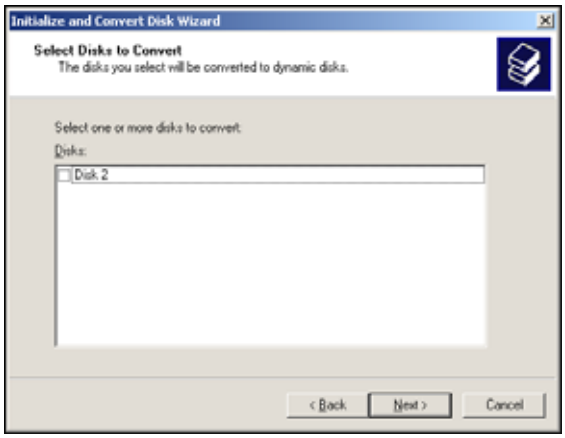

4. Make sure the check box(es) are cleared and click Next.

   The Completing the Initialize and Convert Disk wizard window opens, shown in Figure 305.

Figure 305. Completing the Initialize and Convert Disk Wizard

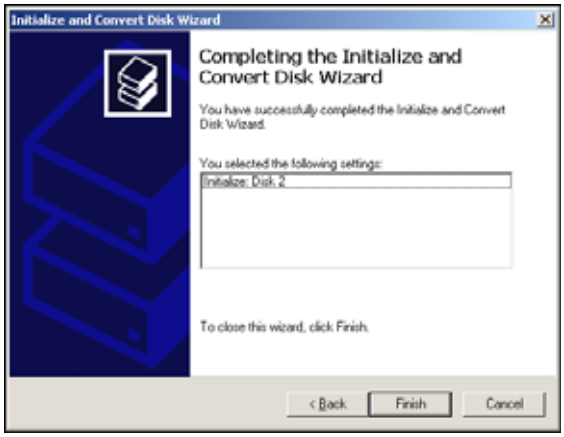

5. Review the selections and click Finish.

   The disks are ready to partition and format. They display in the Disk Management window with the disk area showing Unallocated, shown in Figure 306.

Figure 306. Viewing Unpartitioned EBSD Disks in the Disk Management Window

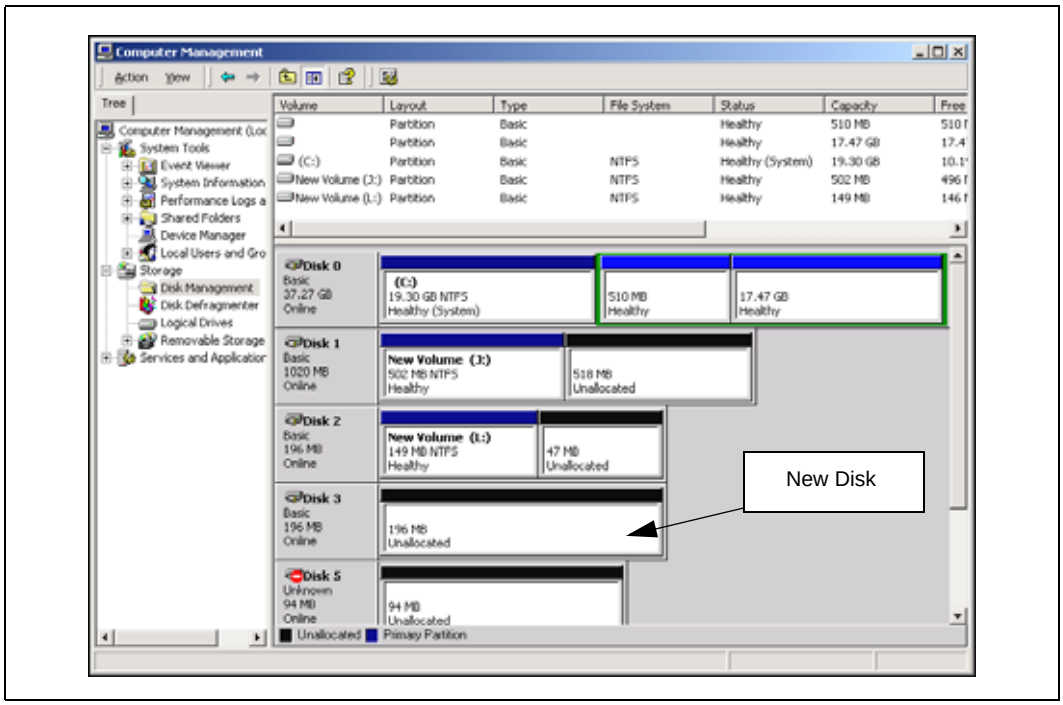

E.12 Partitioning Basic EBSD Disks

Partition an EBSD disk.

1. In the Disk Management window, right-click on the partition area labeled Unallocated with the size.

   For example, in Figure 306, Disk 3 shows 196 MB Unallocated.

2. Select New Partition from the menu.

   The New Partition wizard opens.

3. Click Next.

   The Select Partition Type window opens, shown in Figure 307.

Figure 307. Selecting the Type of Partition

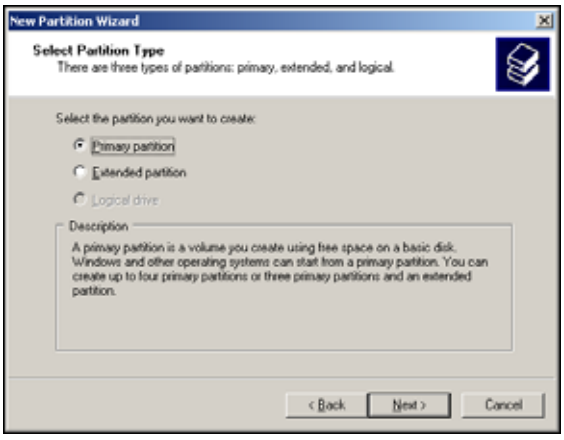

4. Select Primary Partition and click Next.
   The Specify Partition Size window opens, shown in Figure 308.

Figure 308. Selecting the Partition Size

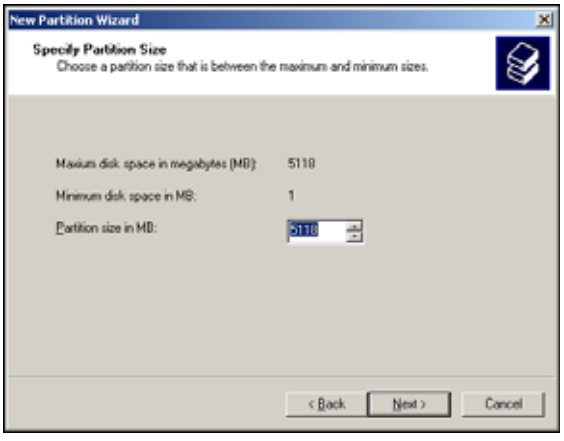

5. Select a size for the partition and click Next.

Note: The standard Windows disk driver is limited to 2TB devices at the block level. The EBSD driver uses this standard disk driver, so native volumes are limited to 2TB.
For basic disks, the 2TB limit is a hard limit with no known work-around.

E.12.1 Assigning Drive Letters and Formatting Partitions

1. The Assign Drive Letter or Path window opens, shown in Figure 309.

Figure 309. Assigning a Drive Letter to a Partition

2. Select a drive letter and click Next.
   The Format Partition window opens, shown in Figure 310.

Figure 310. Formatting the Partition

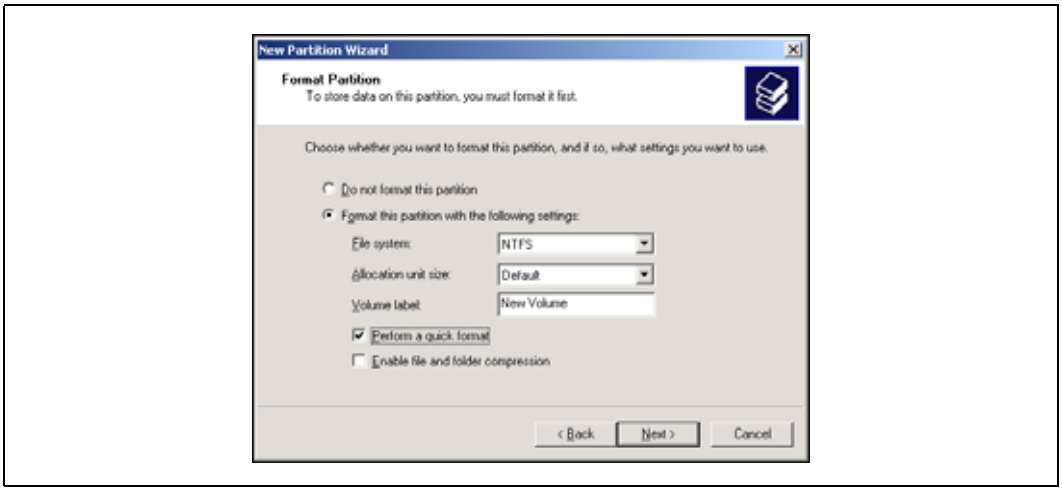

3. Complete the Formatting information and click Next.
   The Completing the New Partition wizard window opens showing a summary of the partition specifications, as shown in Figure 311.

Figure 311. Completing the New Partition Wizard

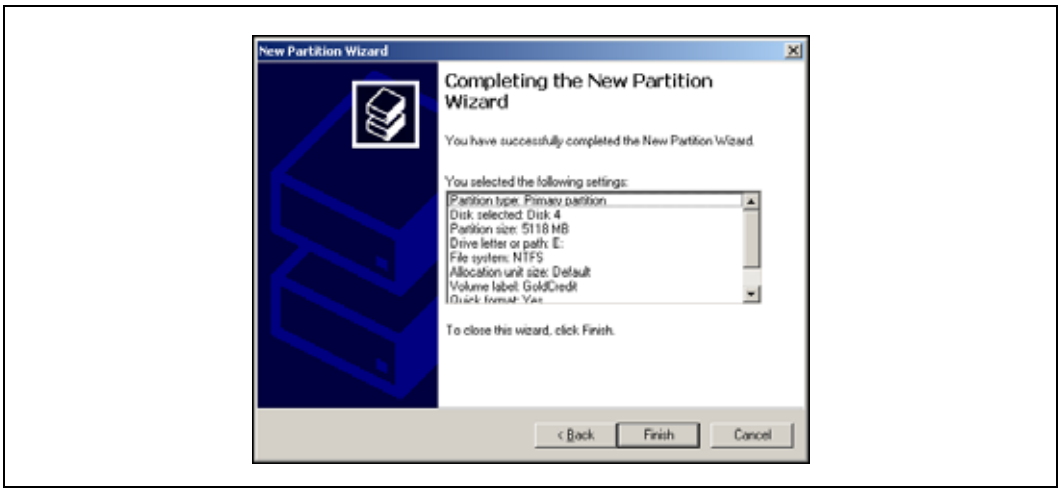

4. Review the specifications and then click Finish.
   The partition is created and formatted and appears in the Disk Management window.

Note: Accessing an EBSD disk at the raw / block level will corrupt the disk database which resides at the end of the disk. Accessing EBSD disks at the block level is not recommended.

E.13 Configuring Applications and Services to Come Online After A Reboot

To ensure that applications and services using EBSD volumes come online after a reboot of the system, add the appropriate services using the Advanced Settings in the EBSD driver. Adding services changes the Startup type of those services to Manual. Then the EBSD driver starts those services either when

- all the EBSD volumes arrive after a reboot of the system or
- after 8 minutes has passed. If 8 minutes passes but all volumes have not arrived, an event log message is recorded and a message notifies the administrator that all EBSD volumes have not arrived.

Removing the services in the Advanced Settings changes the Startup type of those services to back to their original settings.

Note: The EBSD driver adds the 'ArrivalDelay' registry entry to the aebsagent registry group. This entry delays invocation of services until the EBSD volume is online, or for 8 minutes (480 seconds), whichever occurs first. You can change the default delay at HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\aebscheck\

ArrivalDelay:REG_DWORD:0x1e0
The default value for ArrivalDelay is 480 (seconds).

E.13.1 Configuring File Services, SQL Server, and Exchange

File server, SQL Server and Exchange services and their dependent services can be quickly preconfigured in the Advanced Settings window.

E.13.2 Configuring Other Applications with User Services

All other applications require that you identify the Service name before going to Advanced Settings. In the Advanced Settings window you use the User Services button to configure those other applications.

E.13.2.1 Identifying the Service Name for Applications

For applications that are not file server applications, SQL Server, or Exchange, you must identify the application's service name when configuring the application in Advanced Settings. To obtain the application's service name, go to that service's Properties window. Note that the service name is not the same name as the name listed in the Services window.

For example, configure Veritas Backup Exec to come online after a reboot. Veritas Backup Exec is controlled by the Backup Exec Device & Media Service.

1. Click Start > Settings > Control Panel.
2. Select Administrative Tools > Services.
   The Services window opens, shown in Figure 312.

Figure 312. Opening the Services Window and Selecting the Backup Exec Device and Media Service Service

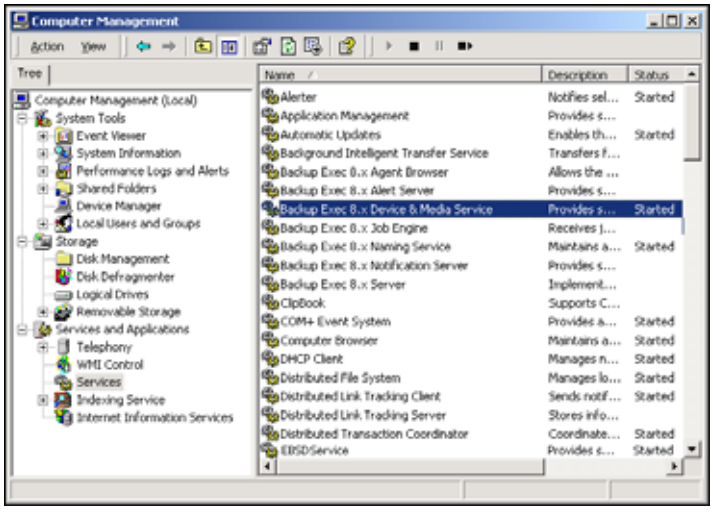

3. Locate the service you want to configure and right-click to open the menu.
   Using our example, select Backup Exec Device & Media Service.

4. Right-click and select Properties.
   The Backup Exec Device & Media Service Properties dialog opens, shown in Figure 313.

Figure 313. Opening the Service Properties Dialog

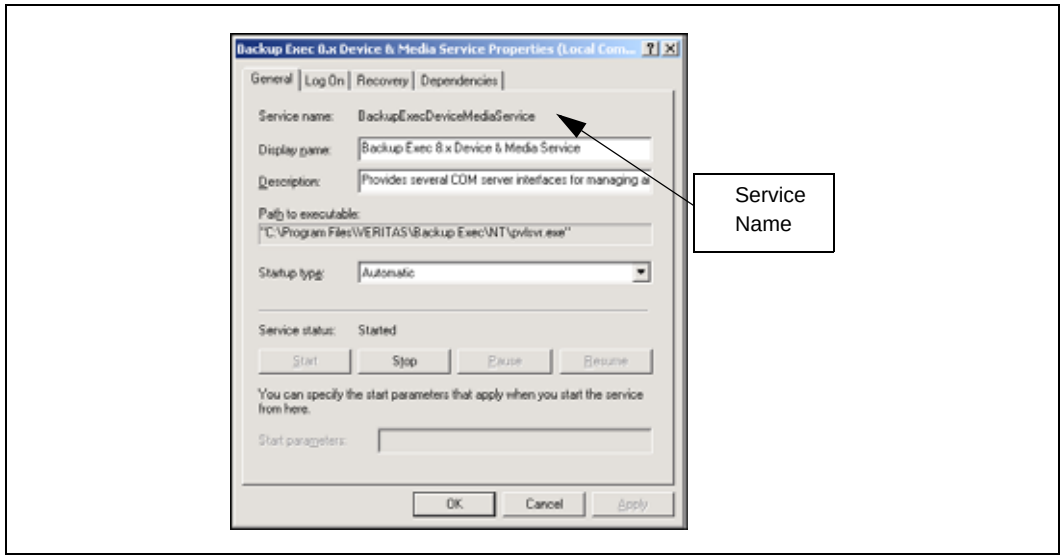

5. Make a note of the Service name listed.
   For example, the name for Backup Exec Device & Media Service is BackupExecDeviceMediaService.
6. Repeat steps 3 through 5 for each service you want to configure to come online after reboot.
7. Close the Services window when you are finished.

E.13.3 Configuring Services In The EBSD Driver

Advanced Settings offers three predefined applications—file Server, MS SQL, and MS Exchange—to configure for coming online after a reboot. All other applications are configured individually.

1. Open the EBSD driver.
2. Click the Settings tab to bring it to the front, as shown in Figure 314.

Figure 314. Settings Tab with the Advanced Button

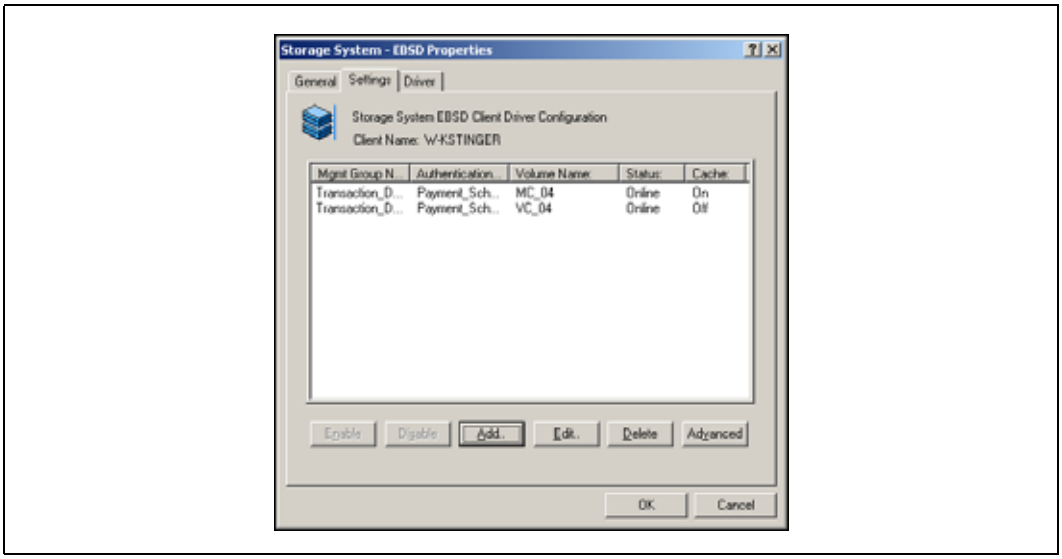

3. Click Advanced to open the EBSD Advanced Settings window, shown in Figure 315.

Figure 315. Advanced Settings Window

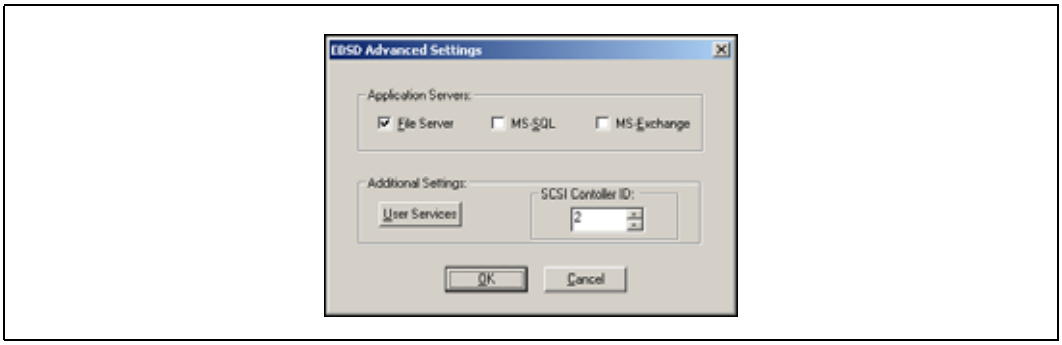

E.13.3.1 Configuring Services for File Server, MS SQL and MS Exchange

If you selected File Server, SQL Server, or MS Exchange when you completed the Add Hardware wizard, your settings display in the Advanced Settings window. You can change these selections.

1. Select the application for which you are configuring services.
2. Click OK.
   A confirmation message opens.
3. Click OK.

E.13.3.2 Configuring Other Applications with User Services

For all other applications, use the Additional Settings section.

1. On the Advanced Settings window, click User Services.
   The EBSD Service Settings window opens, shown in Figure 316.

Figure 316. Configuring other Application Services with their Dependencies

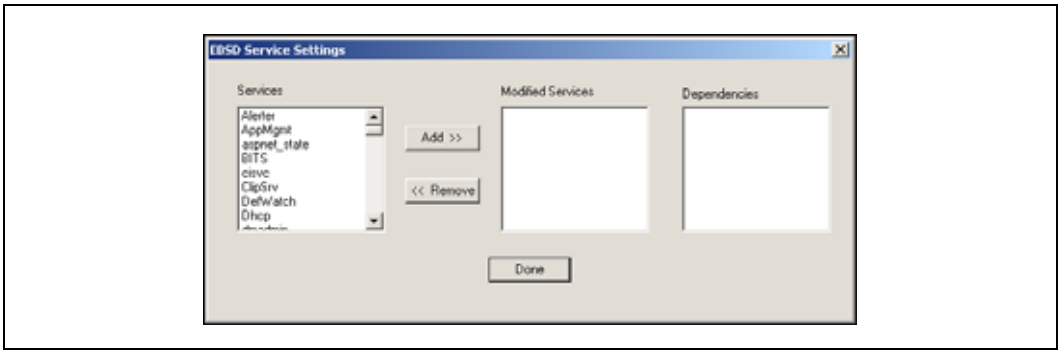

2. Using the list of service names you identified from the Services properties windows, select the services you want to configure for coming online after reboot.

   For example: Select BackupExecDeviceMediaService.to configure Backup Exec Device & Media Service come online after reboot.

3. Click Add.

   The service you selected displays in the Modified Services list. Any dependent services display in the Dependencies List. Those dependent services are automatically included in this configuration.

   For example, one dependent service of **BackupExecDeviceMediaService** is **Backup Exec Job Engine**. Depending upon the system, there may be other dependent services under Backup Exec Device & Media Service.

4. Repeat steps 2 and 3 for all the services you want to configure.

5. Click Done when you are finished.

   Focus returns to the Advanced Settings Window.

6. Click OK.

   Focus returns to the Settings tab.

7. Click OK to close the EBSD driver.

E.13.4 Resetting Services

You can return the application services to their original settings.

1. Open the EBSD driver.
2. Click the Settings tab to bring it to the front.
3. Click Advanced to open the EBSD Advanced Settings window, shown in Figure 315.

E.13.4.1 Resetting File Server, MS SQL, MS Exchange

Reset File Server, MS SQL or MS Exchange to their original setting.

1. Clear the desired application in the Advanced Settings window.
2. Click OK to close the Advanced Settings window.

E.13.4.2 Resetting Other Applications

Reset any other applications by using the User Services.

1. Click User Services to open the EBSD Service Settings window, shown in Figure 317.
   A list of services is listed in the Modified Services pane on the right.

Figure 317. Modified services in Advanced Settings

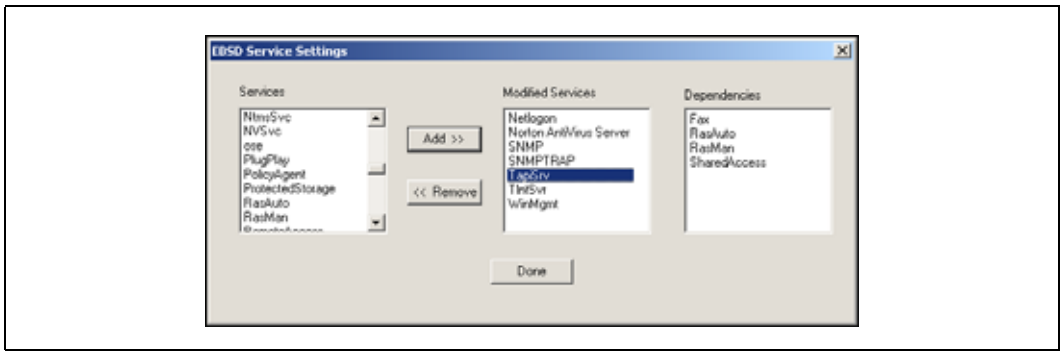

2. Select the service you want to reset and click Remove.
3. Click OK.
   The service name appears in the Services list on the left.
4. Repeat steps 2 and 3 for each service you want to reset.
5. Click Done when you are finished.
   Focus returns to the Settings tab.
6. Click OK to close the EBSD driver.

E.14 Changing the SCSI Controller ID

The default SCSI controller ID for the EBSD driver is 2. In a clustered application server environment, you may have other SCSI controllers with an ID of 2. If other controllers have an ID of 2, a conflict occurs and EBSD volumes will not come online.

To bring the EBSD volumes online, change the SCSI controller ID on the EBSD driver.

1. Disable all volumes on the EBSD driver. See “Disabling and Re-enabling EBSD Disks”.
2. Open the EBSD driver.
3. Click the Settings tab to bring it to the front, as shown in Figure 318.

Figure 318. Settings Tab with the Advanced Button

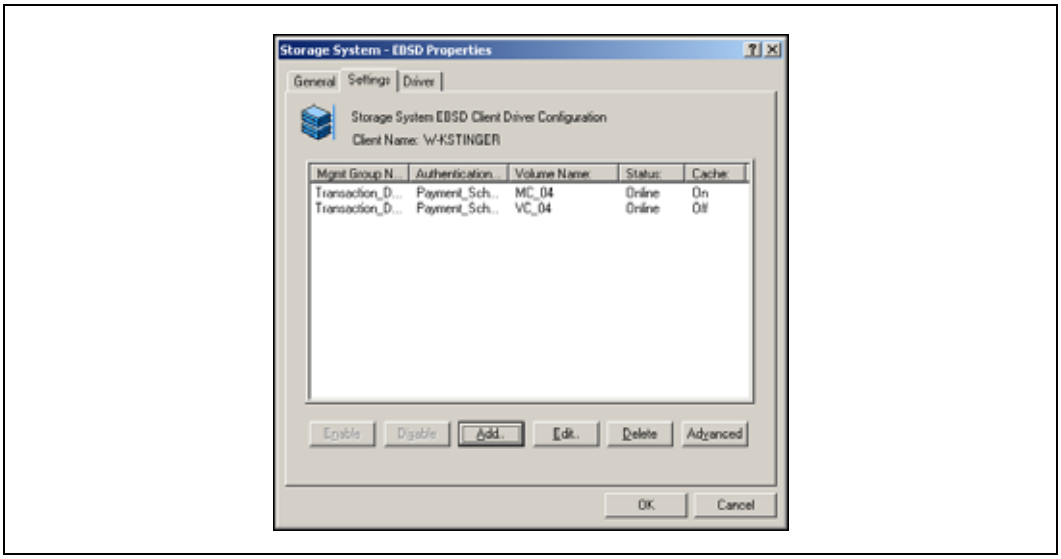

4. Click Advanced to open the EBSD Advanced Settings window, shown in Figure 319.

Figure 319. Advanced Settings Window

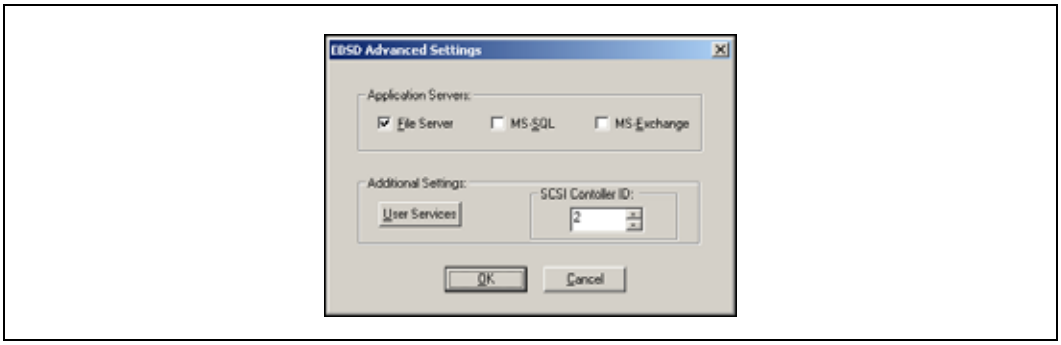

5. Change the value in the SCSI Controller ID box.
   The ID must be a number between 0 and 7. Make sure the port is not already in use.
6. Click OK.
7. Reenable the volumes.

E.15 Managing EBSD Disks

E.15.1 Overview of Managing EBSD Disks

After you have installed and configured the EBSD driver and configured EBSD disks, you will continue to manage those disks and their corresponding volumes on the SSMs. The table below lists various management tasks and where to find information about them.

Table 65. EBSD Driver Management Tasks

Management Task	Management Tool	Instructions
Identifying the Storage System Console volume that corresponds to an EBSD disk.	• Windows Disk Management	"Identifying the Storage System Software Volume That Corresponds to an EBSD Disk"
Accessing read only volumes and snapshots from an EBSD client.	• EBSD driver • Storage System Console	"Accessing Read Only Volumes and Snapshots from an EBSD Client" In the SAN User Manual • Working with Authentication Groups chapter, see "Editing Permissions" and "Creating an Authentication Group Association" • Working with Snapshots chapter, see "Creating Snapshots"
Expanding volumes	• Storage System Console • Windows Disk Management	In the SAN User Manual • Working with Volumes chapter, "Editing a Volume" "Expanding Volumes"
Disabling and re-enabling EBSD disks	• Safely Remove Hardware • Windows Device Manager • EBSD driver	"Identifying the Storage System Software Volume That Corresponds to an EBSD Disk" "Safely Removing the Hardware" "Disabling the EBSD Disk" "Enabling EBSD Disks"
Moving EBSD disks and preserving data on the SSM	• Safely Remove Hardware • Windows Device Manager • EBSD driver • Storage System Console • Windows Disk Management	"Identifying the Storage System Software Volume That Corresponds to an EBSD Disk" "Safely Removing the Hardware" "Deleting the EBSD Disks from the Client" "Preparing a New Client" "Adding EBSD Disks to the New Client" In the SAN User Manual • Working with Authentication Groups chapter, "Authentication Groups Overview"
Deleting EBSD disks and removing data from the SSM	• Windows Disk Management • Safely Remove Hardware • Windows Device Manager • EBSD driver • Storage System Console	"Identifying the Storage System Software Volume That Corresponds to an EBSD Disk" "Deleting Partitions or Volumes from the Client" "Safely Removing the Hardware" "Deleting the EBSD Disks" In the SAN User Manual • Working with Volumes chapter, "Deleting a Volume"
Uninstalling the EBSD driver	• Windows Device Manager • EBSD driver	"Overview of Uninstalling the Driver" "Uninstalling the EBSD Driver"

E.16 Identifying the Storage System Software Volume That Corresponds to an EBSD Disk

When managing EBSD disks in Windows 2003, you must identify which EBSD disk corresponds to a specific volume in the Storage System Console.

1. Open Windows Disk Management, shown in Figure 320.

Figure 320. Identifying the Storage System Software Volume that Corresponds to an EBSD Disk

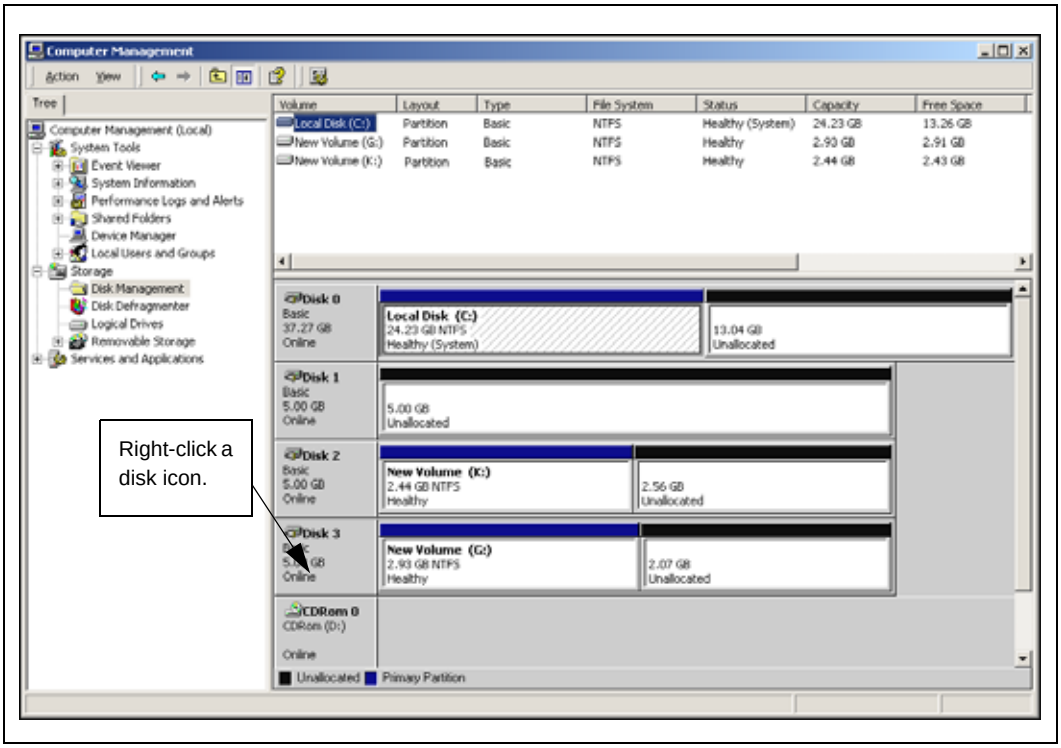

2. Right-click on a disk icon, as shown in Figure 320, and select Properties.
   The Disk Properties window, shown in Figure 321, opens for that disk.

Figure 321. Viewing Disk Properties

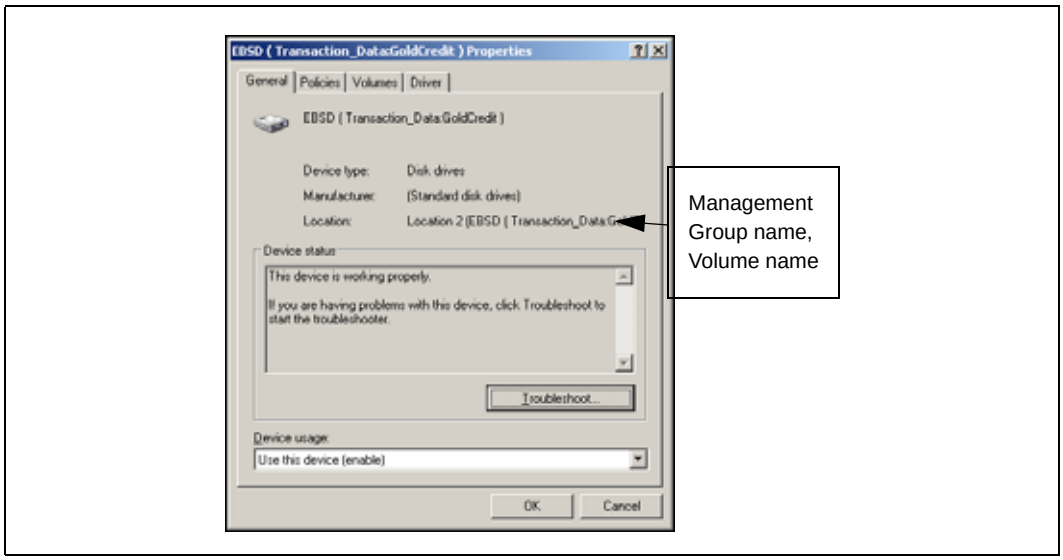

3. On the Disk Properties window, determine the name of the volume by reading the data in the Location field: EBSD (Management Group Name: Volume Name).

   For example, in Figure 321, the EBSD disk corresponds to the MC03 volume in the Transaction_Data management group.

   The management group name and volume name in this line are the names as designated in the Storage System Console.

4. Click OK when you are finished.

E.17 Editing EBSD Volumes

You may encounter circumstances in which you want to edit an EBSD volume. For example, if you work with snapshots or read only volumes, you may have to change some settings in an EBSD volume.

Note: Editing an EBSD volume requires that you disable the volume before you edit it.

Note: Stop any applications from accessing the disk you are going to edit. Applications cannot write to a disk that is disabled.

E.17.1 Safely Remove Hardware

First, remove the hardware. See “Safely Removing the Hardware”.

E.17.2 Disable the Disk

Next, disable the EBSD volume so that no activity can take place on that volume. See “Disabling the EBSD Disk”.

E.17.3 Open the EBSD Driver and Edit Disk

Open the EBSD driver. See “Adding EBSD Disks to Your System”.

1. Select the Settings tab.
2. Select the volume that you want to change and click Edit.
   The Add EBSD Volume window opens, as shown in Figure 322. You can change everything except the Management Group name and the volume name.

Figure 322. Editing an EBSD Volume

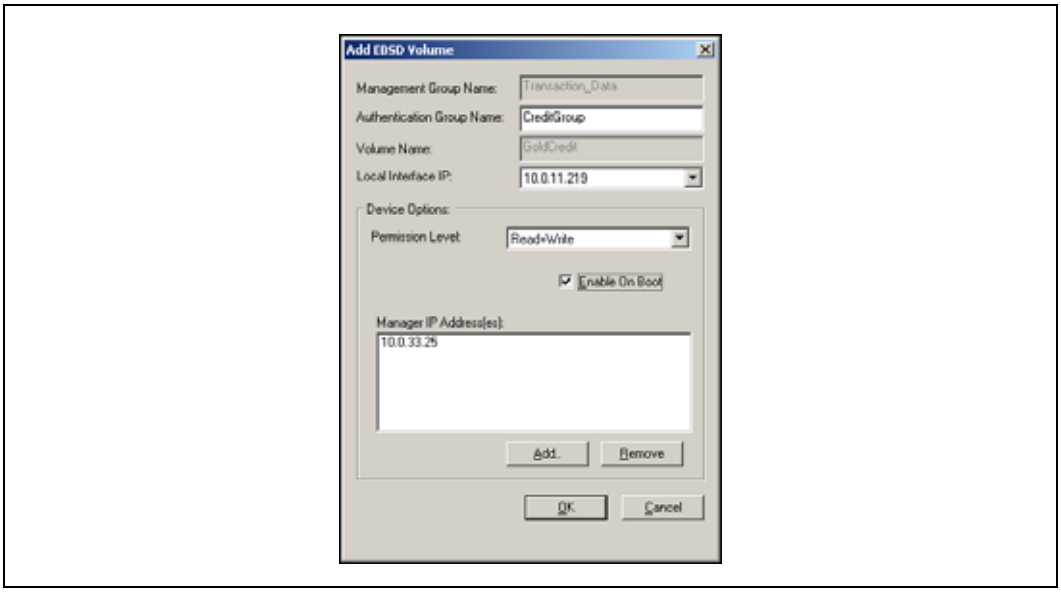

3. Make the desired changes.
4. Click OK.

E.17.4 Re-enable the Disk

Re-enable the EBSD volume. See “Enabling EBSD Disks”.

E.18 Accessing Read Only Volumes and Snapshots from an EBSD Client

You can access snapshots and volumes configured with read only permissions. Mount these snapshots just as you would a regular volume. However, when these snapshots are mounted, there is no indication that they are read only snapshots of a regular volume. However, there are limitations to be aware of.

- Read only volumes and snapshots appear to be read/write, but changes do not get committed. What looks like a change to a snapshot is actually volatile. When the system reboots, or the snapshot is disabled and re-enabled, the changes are not saved.

- Snapshots must have an authentication group with read only access associated with them before mounting.
- The permission level of read only volumes and snapshots must be read only. If the permission level is not read only, the snapshot or volume will not come online and it will remain in the "starting" state.

E.18.1 Changing Volumes to Read Only

You can change a volume configured with read/write access to read only access.

E.18.1.1 Safely Remove Hardware

- First, remove the hardware. See "Safely Removing the Hardware".

E.18.1.2 In the EBSD Driver

- Disable the EBSD disk.
  See "Disabling the EBSD Disk".

E.18.1.3 In the Storage System Console

- Change permissions for the volume.
  See "Editing Permissions" in the Storage System Console Online Help, or in the User Manual in the chapter entitled Working with Authentication Groups.

E.18.1.4 In the EBSD Driver

- Change permissions in the Settings tab.

1. Open the EBSD driver.
2. Select the Settings tab.
3. Select the disk for which you want to change permissions and click Edit.
   The Add EBSD Volume window opens.
4. Change the Permissions Level to Read Only.
5. Click OK.

- Re-enable the EBSD disk. See "Enabling EBSD Disks".

E.18.2 Moving Read Only Volumes to a Different Client

Follow the instructions in "Overview of Deleting or Moving EBSD Disks".

E.18.3 Mounting Snapshots of Basic EBSD Disks

When mounted, snapshots of basic EBSD volumes appear as basic disks to Windows Disk Manager. When mounting snapshots of basic disks, there are no limitations. You can mount multiple snapshots of the same read/write volume on the same EBSD client server.

Note: When you mount a snapshot, it is assigned the same drive letter as its original volume. If this drive letter is already assigned to another disk in the system, then the snapshot will not be assigned a drive letter. You must manually assign a drive letter to the snapshot.

E.19 Expanding Volumes

To increase the size of a volume, increase its size in the Storage System Console and then expand the volume or create new partitions in the Windows 2003 Disk Management window.

1. Increase the volume size, hard threshold, and soft threshold in the Storage System Console.

   Wait for the volume status to change from “Re-striping” to “Normal.”

   See “Editing a Volume” in the SAN User Manual in the chapter entitled Working With Volumes.

2. Open the Disk Management window.

   The newly added storage space appears as Unallocated, shown in Disk 2 in Figure 323.

   — If the newly added storage space does not appear, click the Action menu and select Rescan Disks.

3. Increase the volume size or create new partitions as follows:

   — Create a new partition on the unallocated space.

   — Expand the existing volume using diskpart.exe.

Figure 323. Increasing the Size of Basic Volumes

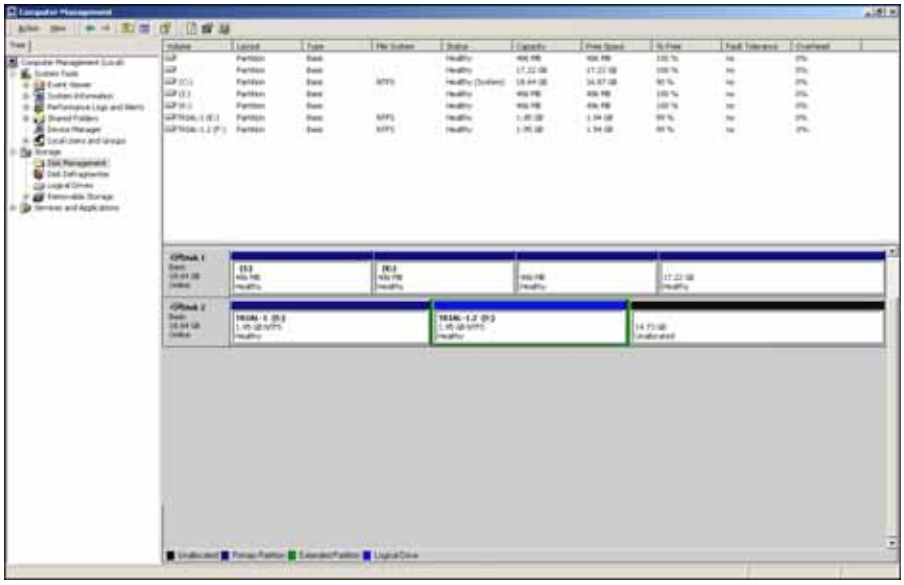

Note: You can increase or decrease the size of EBSD volumes in the Storage System Console. However, there are limitations to using the new size with a Windows 2003 EBSD client:
- Limitations for NTFS and Basic Disks. You may be able to use 3rd party tools, such as Disk

Doubler™, Norton Utilities™, and Partition Magic®, to increase and decrease the size of basic Microsoft disks. However, the 2TB limit still applies.

E.20 Disabling and Re-enabling EBSD Disks

E.20.1 Overview of Disabling and Re-enabling EBSD Disks

Disabling EBSD disks stops activity between the EBSD client and the cluster containing the volume. The partition and drive mapping remain intact and the disk can be re-enabled when appropriate.

Note: Before you begin, be sure that the EBSD disks you plan to disable are not in use.

E.21 Disabling and Re-enabling EBSD Disks

Review the information in "Identifying the Storage System Software Volume That Corresponds to an EBSD Disk" before beginning this process. You must know the relationship of the volume name to the EBSD disk name before beginning.

E.21.1 Safely Removing the Hardware

1. Halt all applications accessing the EBSD volume about to be disabled.
2. Double-click the Safely Remove Hardware icon from the Windows taskbar, shown in Figure 324.

Figure 324. Safely Remove Hardware Icon

The Safely Remove Hardware window opens, shown in Figure 325.

Figure 325. Viewing the Safely Remove Hardware Window

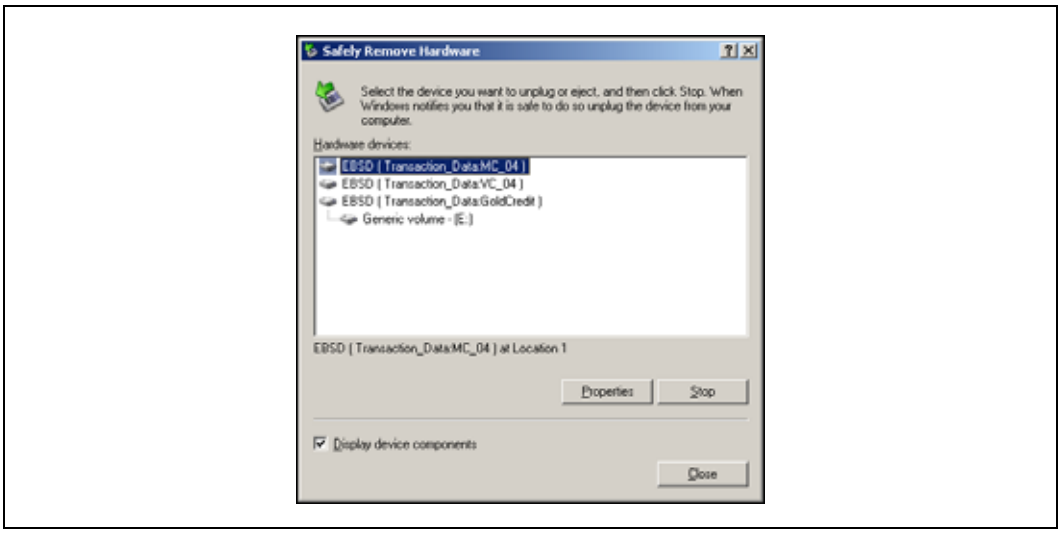

3. Select from the list the device you want to disable.
4. Click Stop.
   A confirmation window opens, shown in Figure 326.

Figure 326. Confirming Devices to be Stopped

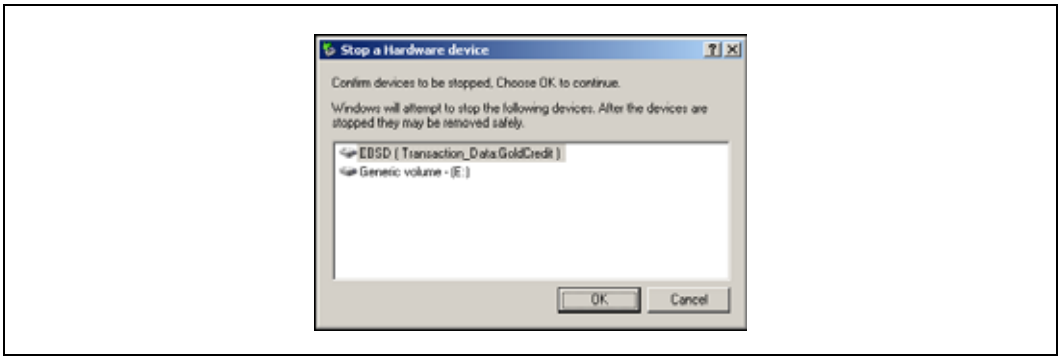

5. Verify the device and click OK.
   A message opens, verifying that the device can now be safely removed from the system.
   The Safely Remove Hardware window returns and the selected device is no longer in the list.
6. Click OK.
7. Repeat steps 3 through 6 for each device that you want to disable.
8. Click Close.

E.21.2 Disabling the EBSD Disk

1. Open Windows Device Manager.
2. Expand the Disk Drives list on the right.
   The disks you removed display an exclamation point icon, as shown in Figure 327.

Figure 327. Viewing the Removed EBSD Disk under Expanded Disk Drives List

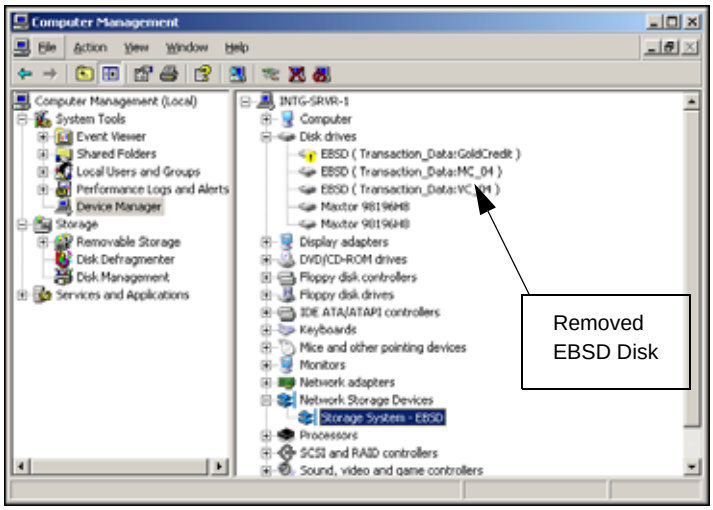

3. Expand the Network Storage Devices list and select the EBSD driver.
4. Double-click on the driver or click the Action menu and select Properties.
   The EBSD Properties window opens, shown in Figure 328.

Figure 328. Disabling an EBSD Disk

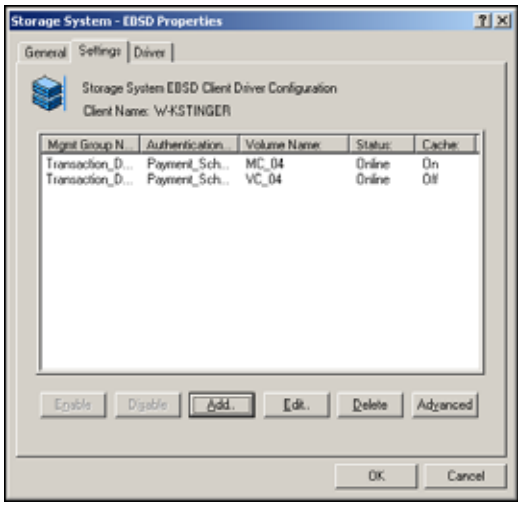

5. Click the Settings tab to bring it to the front.
6. Select from the list the volume that you want to disable.
7. Click Disable.
   A message opens, warning that the data on the volume will become unavailable, and to make certain that all applications are stopped.
8. Click Yes.

9. Repeat steps 6 through 8 for each volume that you want to disable.

Note: Make sure that these disks are already stopped through the Safely Remove Hardware option, described in "Safely Removing the Hardware".

10. Click OK when you are finished.
    The EBSD Properties window closes. The Disk Drive list no longer displays the disks.

E.21.3 Enabling EBSD Disks

You can enable EBSD disks that have been disabled, but not deleted, as long as you have not changed the corresponding volume configuration on the SSM(s) in the Storage System Console.

1. Open Windows Device Manager.
2. Expand the Network Storage Devices list and select the EBSD driver.
3. Double-click on the driver or click the Action menu and select Properties.
   The EBSD Properties window opens.
4. Click the Settings tab to bring it to the front, shown in Figure 329.

Figure 329. Enabling EBSD Disks

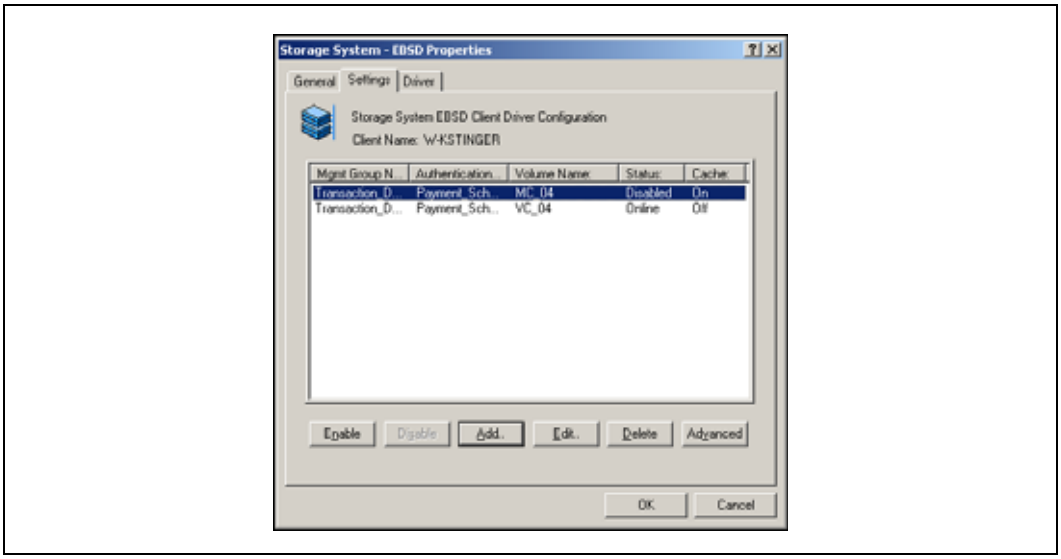

5. Select from the list the volume that you want to enable.
   In this example, we are enabling the VC03 volume.
6. Click Enable.
   The Status column displays "Starting," shown in Figure 330. Then the status changes to Online.

Figure 330. “Starting” Status of an EBSD Disk

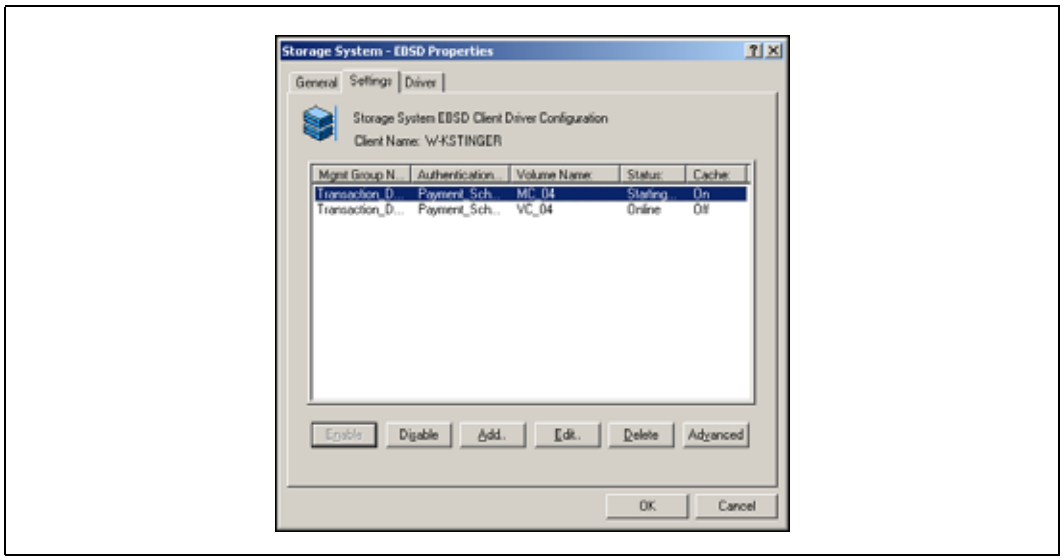

7. Repeat steps 5 and 6 for each volume that you want to enable.
8. Click OK when you are finished.
   The EBSD Properties window closes.

Verifying That the Disks are Enabled

The Disk Drive list now displays the disks. When you open the Disk Management window, the disks and their volumes are again visible.

E.22 Deleting or Moving EBSD Disks

E.22.1 Overview of Deleting or Moving EBSD Disks

You can delete EBSD disks from a client while preserving data on the volume in the cluster. You can also delete EBSD disks and delete all the data on the SSM.

Deleting EBSD disks requires the following tasks:

- Removing the hardware
- Deleting EBSD disks from the EBSD driver

You can then reconnect those EBSD disks on a different server to access the same volumes (and data) on the cluster. The partitions are preserved, though the drive letters might change, depending on whether you change the client machine.

If you reconnect EBSD disks on a client that does not belong to the authentication group associated with the volume in the Storage System Console, you must create a new authentication group for the client and associate the new group with the volume.

Note: Reassigning EBSD disks to a different client preserves the data that is stored on the cluster. The reassigned EBSD disks are connected to the same volumes on the cluster.

E.23 Deleting or Moving EBSD Disks While Preserving Data

Before you begin, review the information in "Identifying the Storage System Software Volume That Corresponds to an EBSD Disk". You must know the names of the EBSD disks you are reassigning to ensure that you are working with the correct volumes in the Storage System Console.

Note: Before you begin, be sure that the EBSD disks you plan to remove are not in use.

E.23.1 Safely Removing the Hardware

The first step in deleting or moving an EBSD disk while preserving the data on the SSM is to remove the EBSD disk from Windows using the Safely Remove Hardware icon on the Windows taskbar, shown in Figure 331.

Figure 331. Safely Remove Hardware Icon

Repeat this process for each disk that you want to delete or move. For more information about removing hardware, see "Safely Removing the Hardware".

E.23.2 Deleting the EBSD Disks from the Client

The second step in deleting or moving an EBSD disk while preserving the data on the SSM is to delete the EBSD disk from the EBSD client.

1. Open Windows Device Manager.
2. Expand the Disk Drives list on the right.
   The disks you removed display an exclamation point icon, as shown in Figure 332.

Figure 332. Viewing the Removed EBSD Disk Under Expanded Disk Drives List

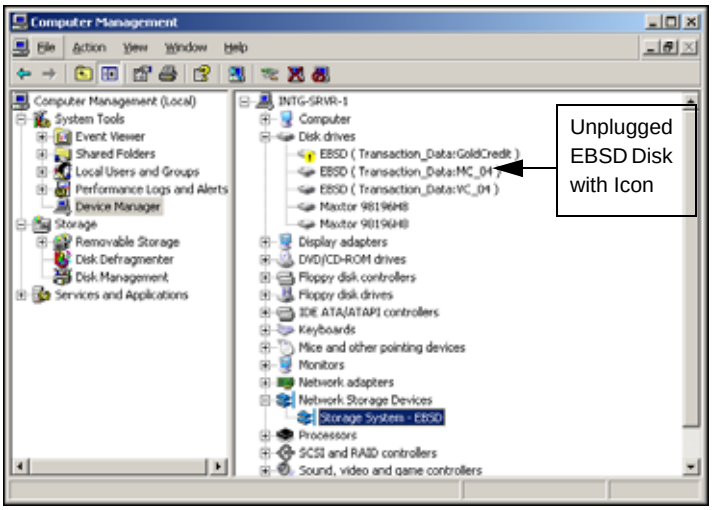

3. Expand Network Storage Devices and select the EBSD driver.

Figure 333. Selecting the EBSD Driver

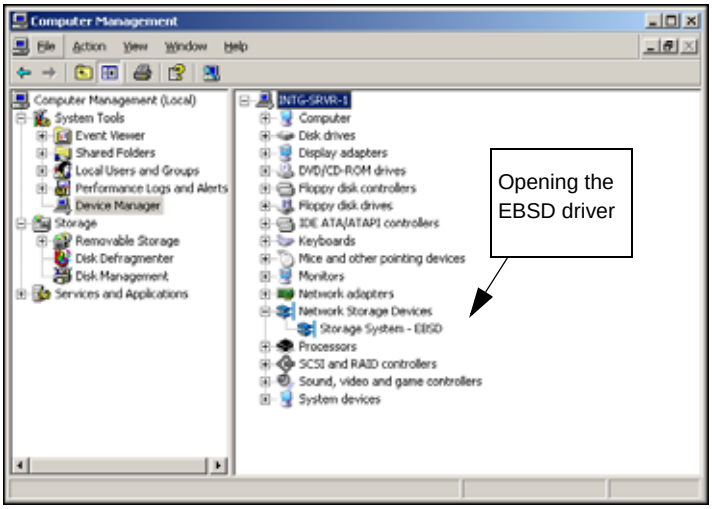

4. Double-click the driver, or right-click and select Properties.
   The EBSD Properties window opens.
5. Click the Settings tab, shown in Figure 334.

Figure 334. Deleting an EBSD Disk

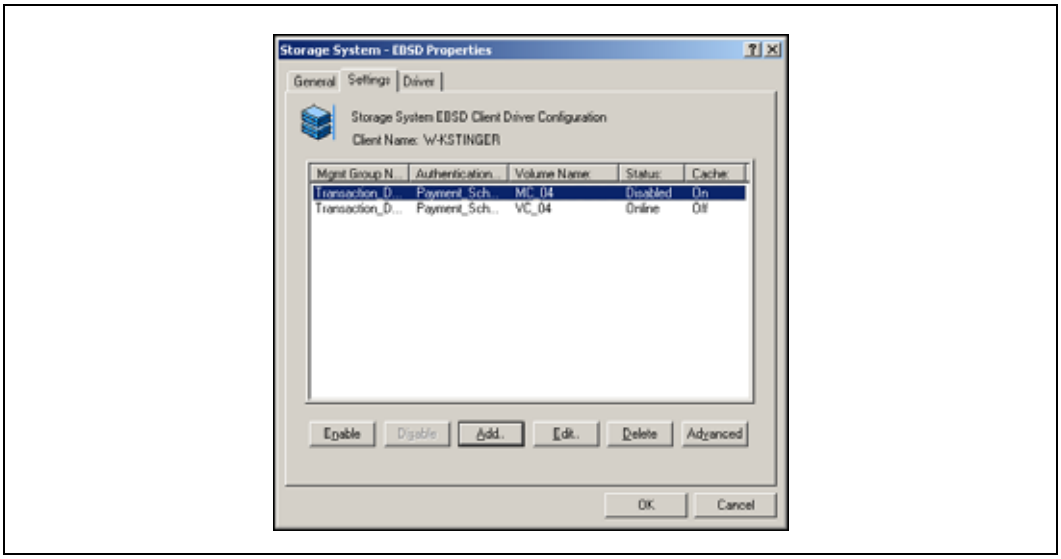

6. Select from the list the disk you want to delete.
7. Click Delete.

   A warning message opens, shown in Figure 335, warning that you cannot access data on this disk once it is deleted.

Figure 335. Warning Message Before Deleting EBSD Disk

8. Click Yes to confirm deleting the disk.

   The message window closes.
9. Repeat steps 6 through 8 for each disk you want to delete.
10. Click OK to close the EBSD driver.

E.23.3 Preparing a New Client

Before you reconnect the EBSD disks on a new client, you may need to reconfigure the volumes in the Storage System Console.

E.23.3.1 Associate New Client with an Authentication Group

If the new client does not belong to the authentication group associated with the volume in the Storage System Console, you must create a new authentication group for the client and associate the new group with the volume. See “Authentication Groups Overview” in the User Manual in the chapter entitled Working with Authentication Groups.

E.23.3.2 Install EBSD Driver

If the EBSD driver has not been installed on the new client machine, you must install it before adding the EBSD disks. For detailed instructions, see "Installing the EBSD Driver" and "Configuration Overview".

E.23.4 Adding EBSD Disks to the New Client

Once you have prepared the new client, add the EBSD disks to the new client.

To add the EBSD disk to the new client, open the EBSD driver and click Add on the Settings tab. For detailed instructions on adding an EBSD disk, see "Adding EBSD Disks to Your System".

E.23.5 Finishing Up

After you reconnect the EBSD disks in the EBSD driver, go to Disk Management. Follow any system instructions that appear. For example, you may be prompted to assign new drive letters to the partitions or volumes.

E.24 Deleting EBSD Disks and Removing Data from the SSM

Removing EBSD disks and ensuring that data is removed from the volume on the cluster requires the following tasks on the client:

- Deleting partitions
- Removing the hardware
- Deleting EBSD disks

and the following task in the Storage System Console:

- Deleting volumes from the cluster. See "Deleting a Volume" in the SAN User Manual in the chapter entitled Working with Volumes.

Note: Make sure the EBSD disks you plan to delete are NOT in use. Deleting EBSD disks from the client and deleting volumes from the cluster removes all data stored in those volumes. Once removed, that data cannot be retrieved.

Before you begin, review the information in "Identifying the Storage System Software Volume That Corresponds to an EBSD Disk". You must know the names of the EBSD disks you are removing to ensure that you are removing the correct volumes in the Storage System Console.

E.24.1 Deleting Partitions or Volumes from the Client

Before you delete an EBSD disk, you must first remove the partitions or volumes from the disk.

1. Open Windows Disk Management.
2. Delete the partitions on the disk that you want to delete.

— Right-click the partition and select Delete Partition.

A confirmation message opens.

3. Click Yes.

   The partition is deleted and the disk becomes unallocated.

4. Repeat steps 2 and 3 for all the disks you want to delete.

E.24.2 Safely Removing the Hardware

The second step in deleting an EBSD disk and deleting the data on the SSM is to remove the EBSD disk from Windows using the Safely Remove Hardware icon on the Windows taskbar, shown in Figure 336.

Figure 336. Safely Remove Hardware Icon

Repeat this process for each disk that you want to delete. For more information about safely removing hardware, see "Safely Removing the Hardware".

E.24.3 Deleting the EBSD Disks

1. Open Windows Device Manager.
2. Expand the Disk Drives list on the tree.

   The disks you removed display an exclamation point icon, as shown in Figure 337.

Figure 337. Viewing the Removed EBSD Disk under Expanded Disk Drives List

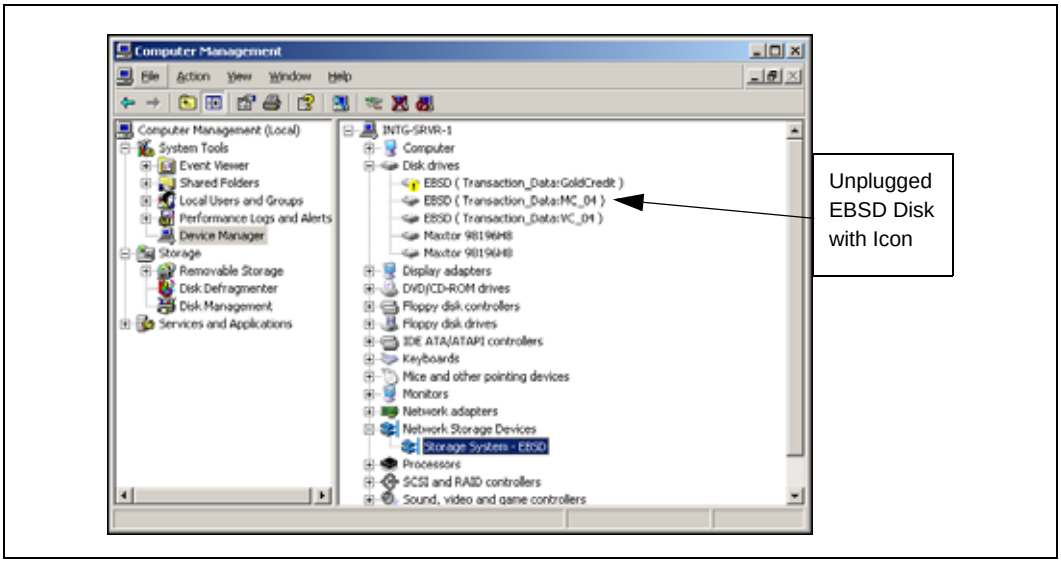

3. Expand Network Storage Devices and select the EBSD driver.
4. Double-click on the driver or click the Action menu and select Properties.

   The EBSD Properties window opens.

5. Click the Settings tab to bring it to the front, as shown in Figure 338.

Figure 338. Deleting an EBSD Disk

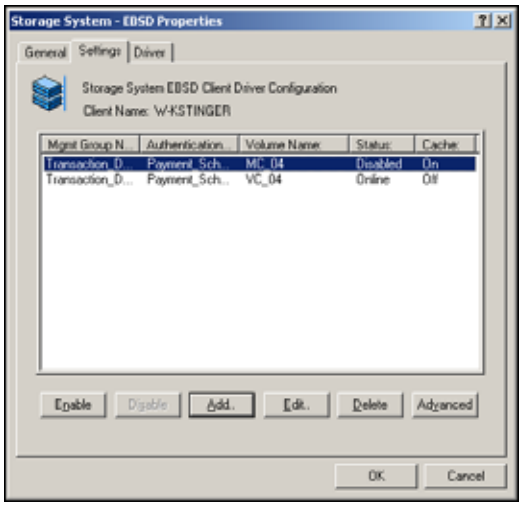

6. Select from the list the disk you want to delete.
7. Click Delete.

   A warning message opens, shown in Figure 339, warning that you cannot access data on this volume once it is deleted.

Figure 339. Warning Message Before Deleting EBSD Disk

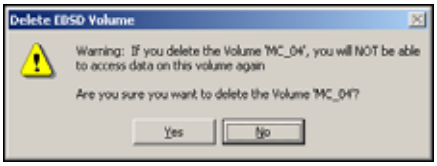

8. Click Yes to confirm deleting the disk.

   The message window closes.
9. Repeat steps 6 through 8 for each disk you want to delete.
10. Click OK to close the EBSD driver.

The final step in deleting the EBSD disk is to delete the volume in the Storage System Console. See "Deleting a Volume" in the SAN User Manual in the chapter entitled Working with Volumes.

E.25 Uninstalling the EBSD Driver

E.25.1 Overview of Uninstalling the Driver

Before uninstalling the EBSD driver for Windows 2003, you should remove the EBSD disks. Review the procedures for removing disks in one of the following two sections, depending upon your goals.

- See "Overview of Deleting or Moving EBSD Disks" if you want to recreate the EBSD disks later, or on another client, or otherwise preserve the data.
- See "Deleting EBSD Disks and Removing Data from the SSM" if your goal is to completely remove EBSD from the client and scrub the SSM of all the data stored in the corresponding volume.

E.26 Uninstalling the EBSD Driver

1. Open Windows Device Manager.
2. Expand the Network Storage Devices list and select the EBSD driver.
3. Double-click on the driver or click the Action menu and select Properties.
   The EBSD Properties window opens.
4. Click the Driver tab to bring it to the front, as shown in Figure 340.

Figure 340. Uninstalling the EBSD Driver

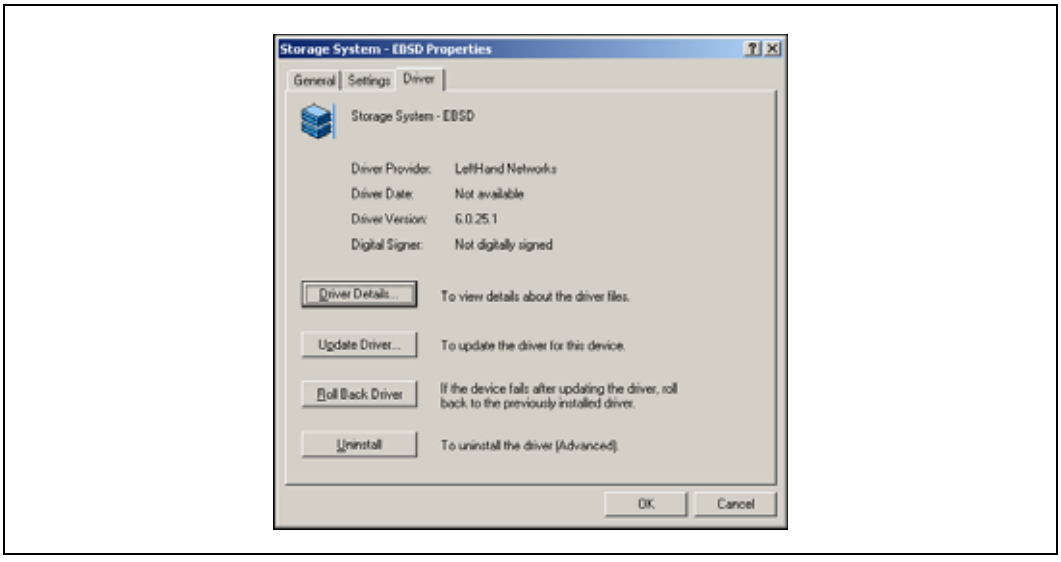

5. Click Uninstall.
   A warning message opens, shown in Figure 341.

Figure 341. Warning before Uninstalling

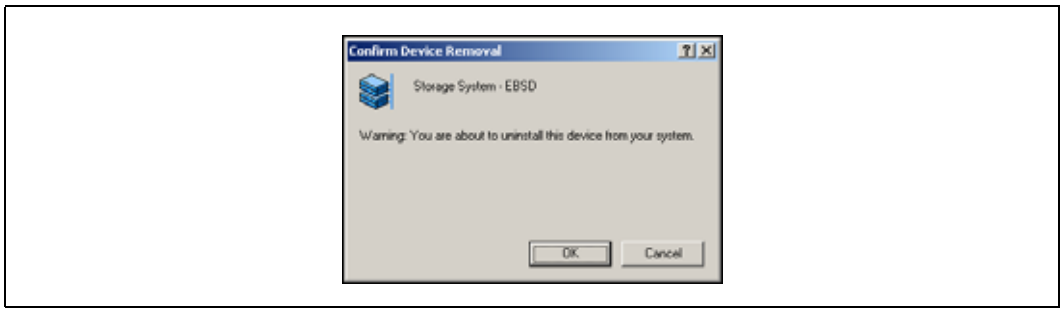

6. Click OK.

   Another message opens, telling you to delete or disable any EBSD volumes that are online.

7. Click OK.

   The driver is uninstalled. A message opens telling you to reboot your computer to complete the uninstall.

8. Restart your computer to complete the uninstall.

Using Remote Copy F

F.1 Remote Copy Overview

Remote Copy provides a powerful and flexible method for replicating data and keeping that replicated data available for business continuance, backup and recovery, data migration, and data mining.

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 SSM in a cluster that will contain a primary volume or a remote volume. For information about registering Remote Copy licenses, see Chapter 16, “Feature Registration” in the Storage System Software User’s Guide.

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.

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.

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.

This chapter provides instructions for registering, configuring, and using Remote Copy for business continuance, backup and recovery, and failover.

For information about how Remote Copy works and how to plan capacity for Remote Copy, see Chapter 1, “Understanding and Planning Remote Copy.”

F.1.1 Glossary for Remote Copy

The following terminology is used in describing the components and processes involved in Remote Copy.

Table 66. 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.
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. It 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.

F.1.2 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.

Figure 342. Basic Flow of Remote Copy

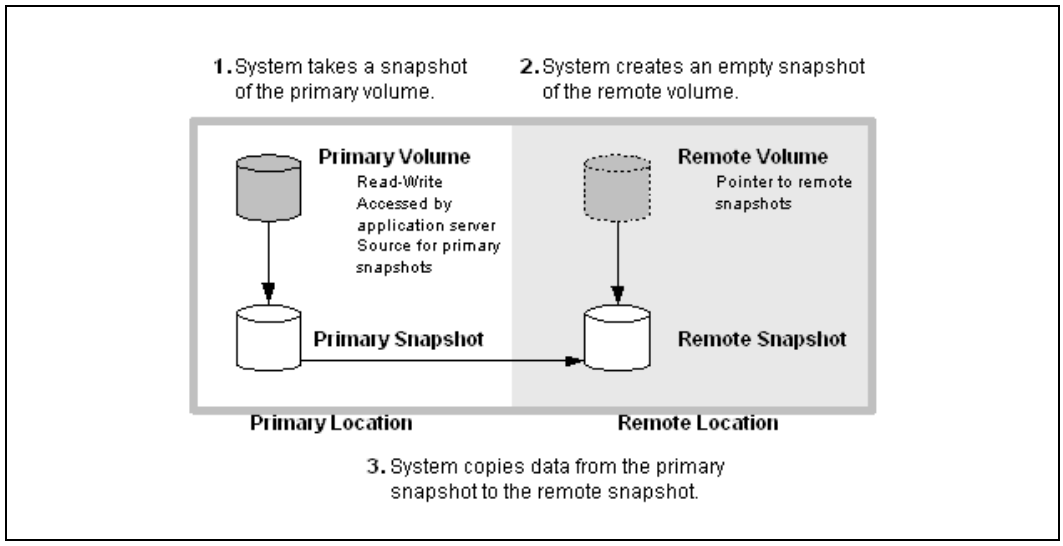

Note: Both primary and completed remote snapshots are the same as regular snapshots. See the chapter "Working with Snapshots" in the Storage System Software User's Manual.

Note: Remote Copy can be used on the same site, even in the same management group and cluster.

F.1.3 Graphical Representations of Remote Copy

The icons depicted in Figure 343 show special graphical representations of Remote Copy.

F.1.3.1 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 343. The pages move in the direction of the data flow from primary to remote snapshot.

Figure 343. Icons Depicting the Primary Snapshot Copying to the Remote Snapshot

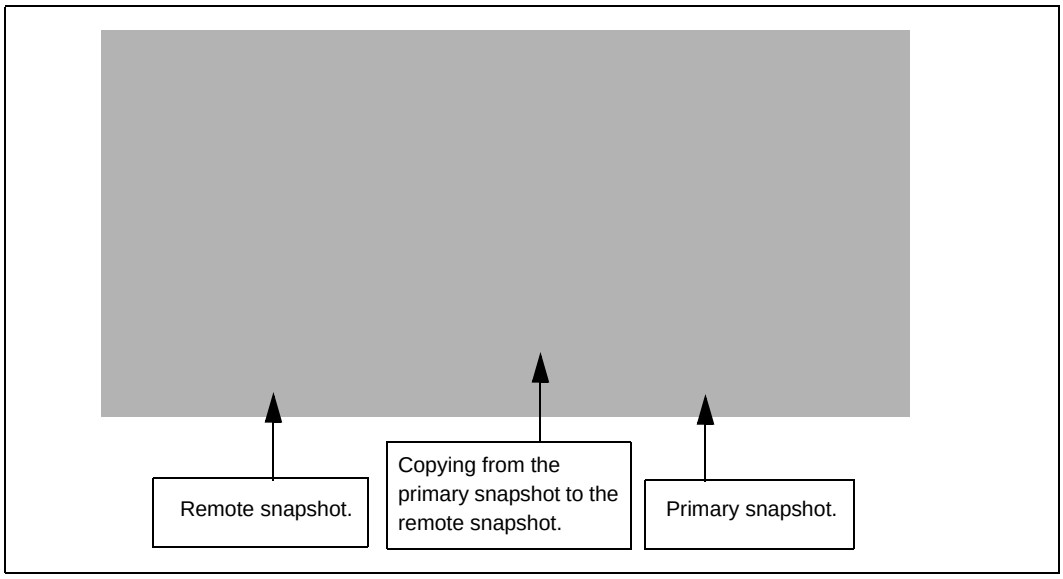

F.1.3.2 Graphical Legend for Remote Copy Icons

The graphical legend available from the Help menu depicts the icons associated with Remote Copy. Figure 344 displays the Remote Copy states icons from the graphical legend.

Figure 344. Icons for Remote Copy as Displayed in the Graphical Legends Window

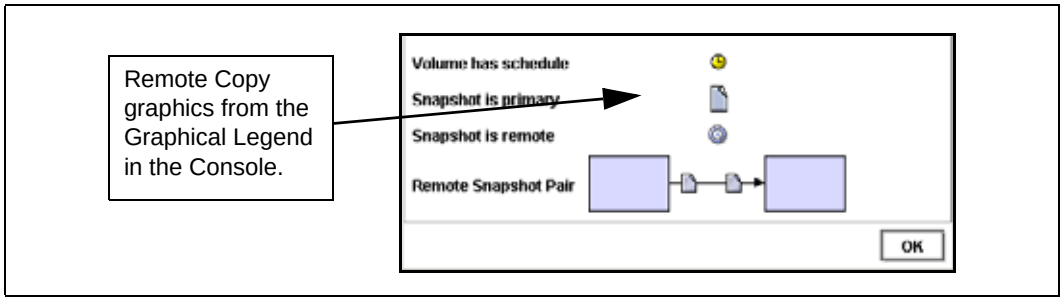

F.1.4 Remote Copy and Volume Replication

Remote Copy is asynchronous replication of data. Volume replication is synchronous replication. Volume replication is described in detail in the Storage System Software User Manual in the chapter, “Working with Volumes.” Using synchronous volume replication on multiple SSMs within a cluster in combination with asynchronous Remote Copy on a different cluster of SSMs creates a robust high-availability configuration.

F.1.5 Uses for Remote Copy

Table 67. Uses for Remote Copy

Use Remote Copy for	How It Works
• Business continuance/ disaster recovery	Using Remote Copy, store remote snapshots off-site. The remote snapshots remain continuously available in the event of a site or system failure.
• 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.

F.1.6 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.

F.2 Planning for Remote Copy

Remote Copy works at the management group, cluster, volume, snapshot, and SSM level.

Table 68. Remote Copy and Management Groups, Clusters, Volumes, Snapshots, and SSMs

Storage System Level	Remote Copy Configuration
Management Groups	• Remote snapshots can be created in the same management group or in a different management group than the primary volume. • 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	• Remote snapshots can be created in the same cluster or in a different cluster than the primary volume.

Table 68. Remote Copy and Management Groups, Clusters, Volumes, Snapshots, and SSMs

Storage System Level	Remote Copy Configuration
Volumes	• Primary volumes contain the data to be copied to the remote snapshot. • Data is copied to the remote snapshot via the remote volume. • The remote volume is a pointer to the remote snapshot. The remote volume has a size of 0.
Snapshots	• Once data is copied from the primary snapshot to the remote snapshot, the remote snapshot behaves as a regular snapshot.
SSM	• Active monitoring of each SSM 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.

F.2.1 Planning the Remote Snapshot

In order to create a remote snapshot:

- you must be logged in to both the management group that contains the primary volume and the management group containing the target cluster where the remote snapshot will be created.
- you must designate or create a remote volume in that remote management group.
- you must have enough space on the target cluster for the remote snapshot.

F.2.1.1 Logging in to the Management Group

Log in to both management groups before you begin. 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.

F.2.1.2 Designating or Creating the Remote Volume

You can 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 panel and then select the Remote option on the New Volume window.

For more information about the three methods of creating remote volumes, see “Creating a Remote Volume”.

F.3 Using Schedules for Remote Copy

Scheduled remote snapshots provide high availability for business continuance/disaster recovery and provide a consistent, predictable update of data for remote backup and recovery.

F.3.1 Planning the Remote Snapshot Schedule

When creating a remote snapshot schedule, a number of considerations are important to plan. All of these issues impact the amount of storage available in the system.

- **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 two remote snapshots of the primary volume. Since the first remote snapshot copies the entire volume, it will take longer to copy. The second remote snapshot copies only **changes** made to the volume since the first remote snapshot. Since 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 management group containing the remote volume, since that 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? See the chapter on snapshots in the Storage System Software User's Manual for information about managing capacity using volume and snapshot thresholds.

  If the cluster does not have sufficient space available, the remote snapshot will appear in the Console and it will flash red. On the Details tab of the remote snapshot, the status says "Read only, not enough space in cluster to start copy."

- **Retention** - how long do you want to retain the primary snapshots? The remote snapshots? You can 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. A remote snapshot that is in the process of being copied is not counted in the retention policy.

Note: If you retain more remote snapshots than primary snapshots, the remote snapshots become regular snapshots when their corresponding primary snapshots are deleted. You can identify them as remote snapshots by their names, since the naming convention is established as part of creating the remote snapshot schedule.

F.3.2 Best Practices

- Retain at least two primary snapshots to ensure that only incremental copying is required for primary snapshots.
- Do not delete a primary snapshot before copying to the remote snapshot completes. An incomplete remote copy snapshot cannot be accessed.

Use the checklist in Table 69 to help plan scheduled remote snapshots.

F.3.2.1 Scheduled Remote Copy Planning Checklist

Table 69. Scheduled Remote Copy Planning Checklist

Configuration Category	Parameters
Snapshot Schedule	
Start Time	• Start date (mm/dd/yyyy) and • Start time (mm:hh:ss) for the schedule to begin
Recurrence	• Recurrence (✓). Recurrence is a yes/no choice. You can schedule a remote snapshot to occur one time in the future and not have it recur. • Frequency (minutes, hours, days or weeks)
Primary Setup	
Hard Threshold Soft Threshold	Set the hard threshold and soft threshold for the primary snapshot.
Retention	Retain either • Maximum number of snapshots (#) • 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 • Maximum number of snapshots (#). This number equals completed snapshots only. In-progress snapshots take additional space on the cluster while they are being copied. • Set period of time (minutes, hours, days or weeks)

F.4 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 Chapter 16, “Feature Registration” in the Storage System Software User’s Guide.

F.4.1 Number of Remote Copy Licenses Required

Register Remote Copy on each management group that contains SSMs that will participate in Remote Copy. If there are SSMs 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 SSMs that will contain a remote volume, and another cluster of 3 SSMs that will not use Remote Copy, you only need 2 Remote Copy licenses.

F.4.2 Registering Remote Copy

For information about starting the 30-day evaluation period and about registering Remote Data Protection Pak, Chapter 16, “Feature Registration.” in the Storage System Software User Guide.

F.5 Creating the Remote Snapshot

Creating a remote snapshot is the main task in Remote Copy. You can 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 4 main steps:

- First, log in to the management groups that will contain primary and remote volumes.
- Second, create a primary snapshot on the primary volume.
- Third, create a remote volume or select an existing remote volume.
- Fourth, specify the settings for the remote snapshot.

F.5.1 Getting There

1. Log in to the management group that contains the primary volume for which you are creating the remote snapshot.
2. Log in to the management group that will contain the remote volume and 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.
3. Right-click the primary volume and select Remote Copy > New Remote Snapshot.
   The New Remote Snapshot window opens, shown in Figure 345.

Figure 345. Creating a New Remote Snapshot

F.5.2 Creating the Primary Snapshot

1. In the Primary section of the New Remote Snapshot window, click New Snapshot.

The New Snapshot window opens, shown in Figure 346.

Figure 346. Creating a New Primary Snapshot

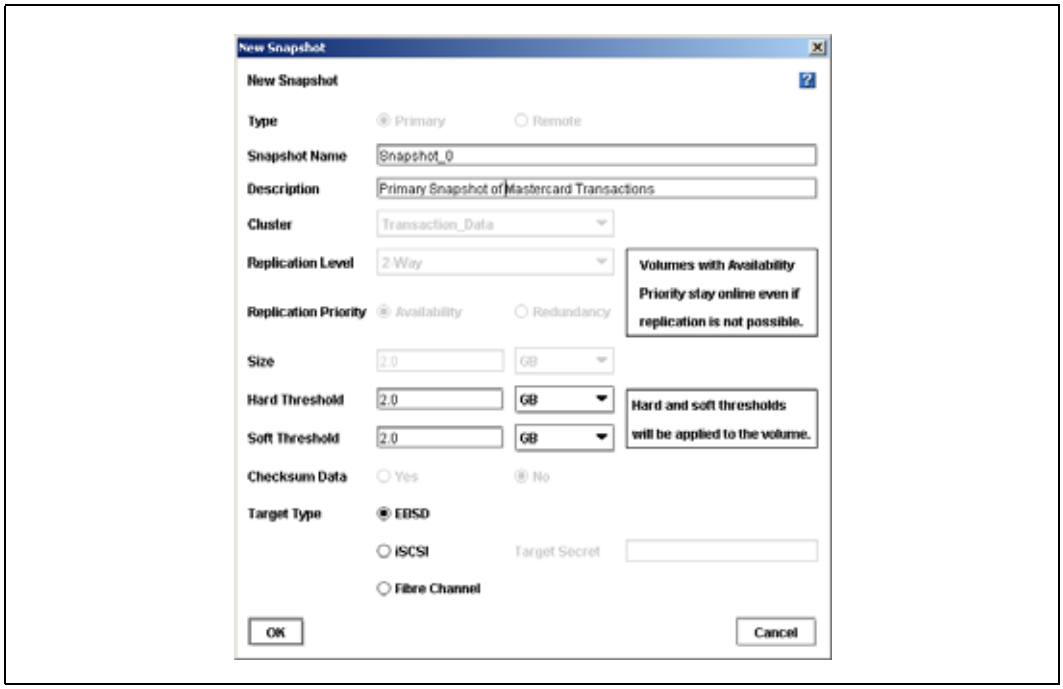

2. Type a name for the primary snapshot.

   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.” The Console displays volume and snapshot names under the icons. If a name is longer than the width of the icon, the end of the name is cut off (however, the full name does show on the corresponding Details tab and on other relevant tab views).

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 347. At this point the primary snapshot has been created.

Figure 347. New Primary Snapshot Created

F.5.3 Creating a Remote Volume

If you have already created the remote volume, select the management group and existing remote volume in the Remote section of the New Remote Snapshot window. Then go to "Completing the Remote Snapshot".

You can 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 panel and then select the Remote option on the New Volume window.

F.5.3.1 Making an Existing Volume Into a Remote Volume

Selecting an existing volume to become a remote volume will cause

1. a snapshot of all existing data to be created for that volume and then
2. all the data in that volume will be deleted so that the remote volume will have zero length and zero hard and soft thresholds.

F.5.3.2 Creating a New Remote Volume.

When you create the remote snapshot, use the Remote Snapshot window, shown in Figure 345, to create the volume. Alternately, you can create a new volume from the cluster details panel and select the Remote option in the New Volume Window.

Note: The fastest way to create a remote volume is to create the volume in the process of creating the remote snapshot, using the Remote Snapshot window.

To create the remote volume from the New Remote Snapshot window:

1. In the Remote section, select the Management Group to contain the remote snapshot.

You must be logged into the management group to continue.

2. To create a new remote volume, click New Volume.
   The Cluster List window opens, shown in Figure 348.

Figure 348. Selecting a Cluster for the Remote Volume

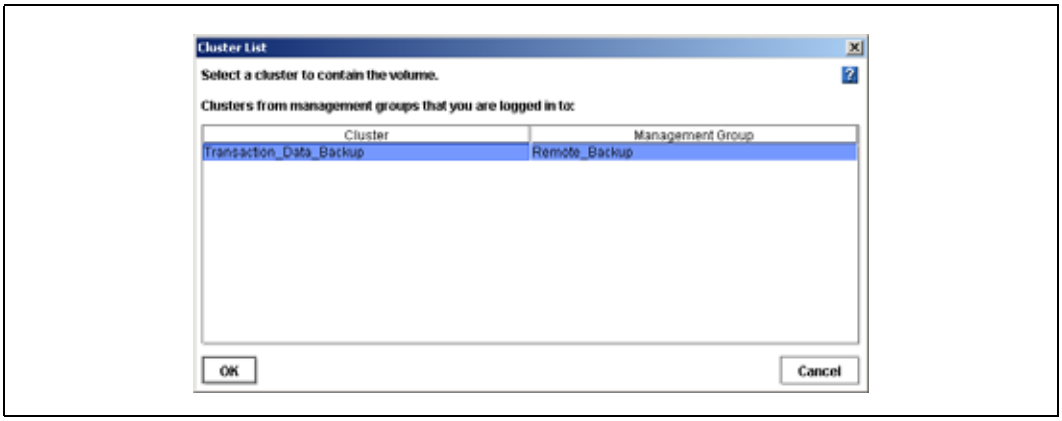

3. Select a cluster for the remote volume and click OK.
   The New Volume window opens, shown in Figure 349. See the chapter on volumes in the Storage System Software User's Manual for detailed information about creating volumes.

Figure 349. Creating a New Remote Volume

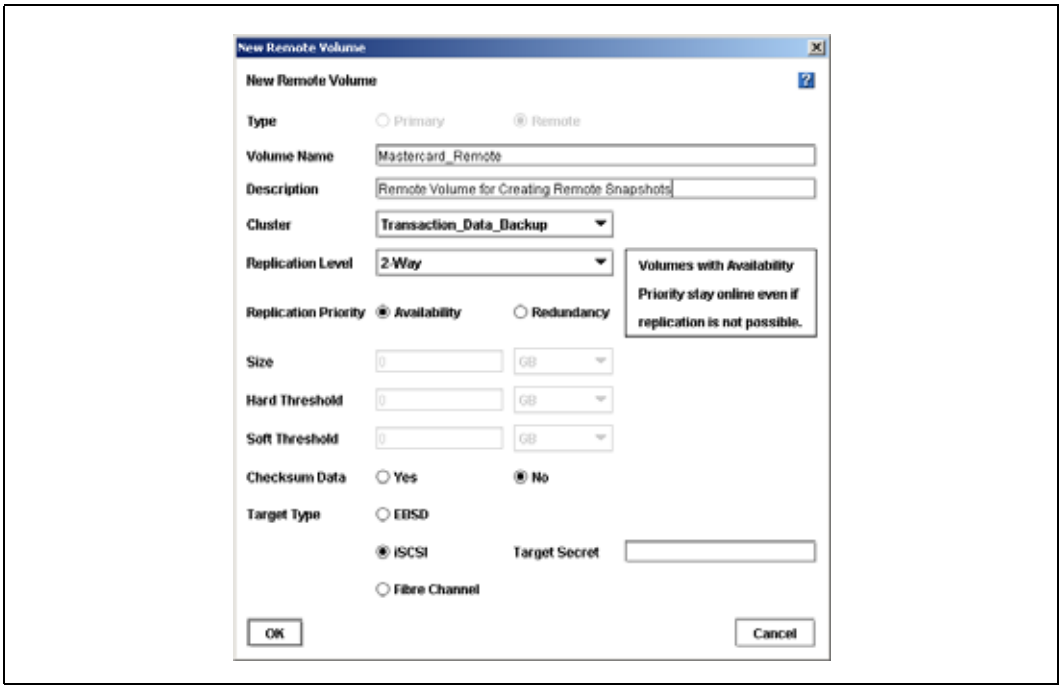

4. Type a name for the volume.
   A volume name must be from 1 to 127 characters and is case sensitive.
5. [Optional] Type a description of the volume.
6. Select the replication level.

You can set different replication levels for the remote volume and the primary volume.

Note: You cannot set the size or thresholds for the remote volume. Those values are 0, since the remote volume is a placeholder for data.

7. Select a replication priority.
   If you select a replication level of None, you cannot set a replication priority. See the chapter on volumes in the Storage System Software User's Manual for detailed information about creating volumes.
8. Select the Target Type for the volume.
9. [Optional] If the volume is an iSCSI target and you want to use 1-way or 2-way CHAP, type a target secret.
10. Click OK to return to the New Remote Snapshot window.
    The new remote volume has been created at this point.

F.5.4 Completing the Remote Snapshot

1. Type a name for the remote snapshot.
2. [Optional] Type a description for the snapshot.
   The completed window is shown in Figure 350.
3. Click OK.

Figure 350. Completing the New Remote Snapshot Dialog

F.5.4.1 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, as shown in Figure 351.

Note: If you create a remote snapshot of a volume with a remote snapshot still in progress, the second remote snapshot will not begin copying until the first remote snapshot is complete.

Figure 351. Viewing the Remote Snapshot

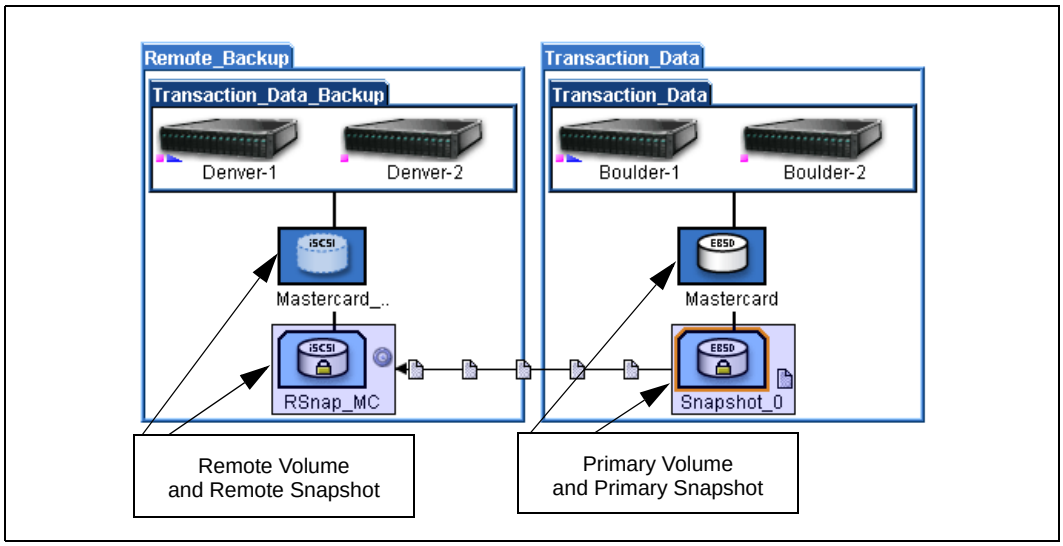

F.6 Canceling a Remote Snapshot

When you cancel a remote snapshot that is in progress, the remote snapshot is deleted and the primary snapshot remains.

To cancel a remote snapshot that is in progress

1. Click the primary or remote snapshot.
   The snapshot tab view opens.
2. Click the Remote Snapshot tab.
3. Select from the list the remote snapshot you want to cancel.
4. Click Cancel Remote Snapshot.
   A confirmation message opens.
5. Click OK.

F.7 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.
2. Right-click the remote snapshot and select Edit Snapshot from the menu.
   The Edit Snapshot window opens, shown in Figure 352.

Figure 352. Editing a Remote Snapshot

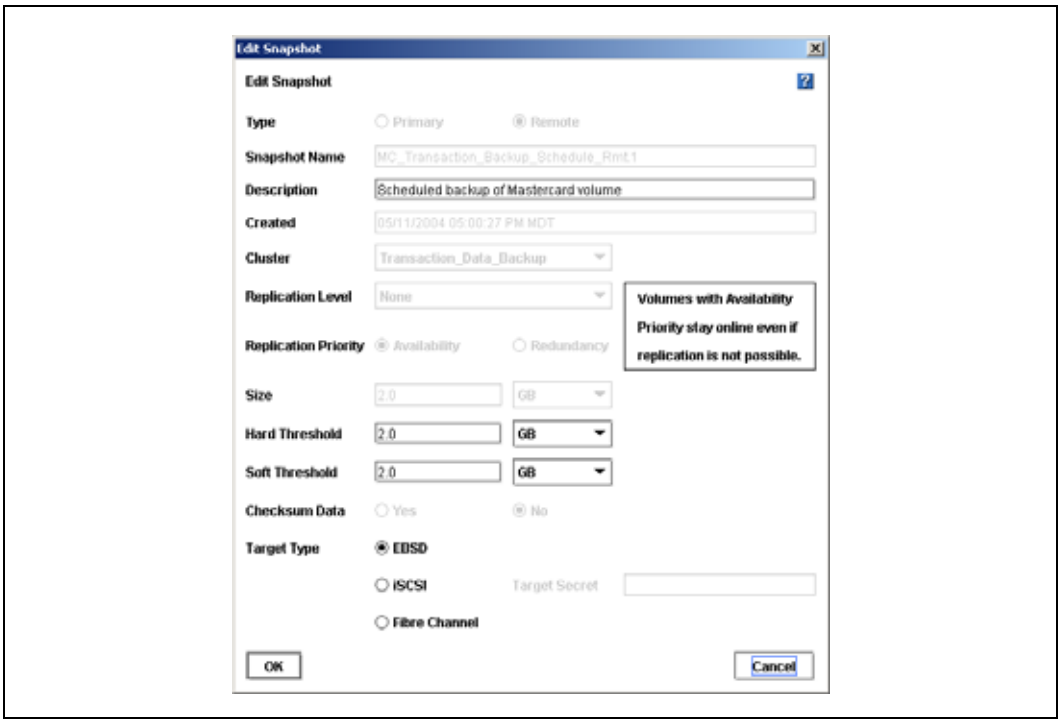

3. Change the desired information and click OK.

F.8 Deleting a Remote Snapshot

1. Log in to the management group that contains the remote snapshot.
2. Right-click the remote snapshot and select Delete Snapshot from the menu.

F.9 Viewing a List of Remote Snapshots

You can view the list of remote snapshots associated with a primary volume or a primary snapshot.

1. Click the primary volume or the primary snapshot for which you want to view the list of remote snapshots.
2. Click the Remote Snapshot tab.

   The tab view opens, shown in Figure 353. The report on the tab lists the primary management group and snapshots along with the remote management group and snapshots. The status column lists the percent of copying completed from the primary to the remote snapshot.

Figure 353. Viewing the List of Remote Snapshots

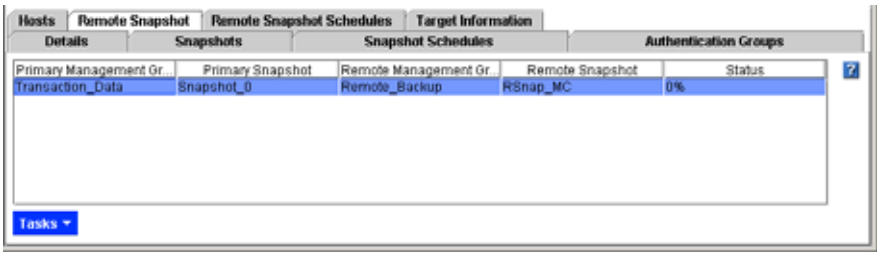

F.10 Setting the Remote Bandwidth

The remote bandwidth sets the maximum rate for data transfer between management groups. The remote bandwidth setting is the upper limit of the range of data transfer—that is, the copy rate will be equal to, or less than, the rate set.

The remote bandwidth specifies the speed at which data is received from another management group. This means that 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.

1. Right-click the remote management group and select Edit Remote Management Groups.
   The Edit Remote Management Groups window opens, shown in Figure 354.

Figure 354. Editing a Remote Management Group

2. Select the management group for which you want to change the remote bandwidth.

3. Click Edit Bandwidth.
   The Edit Remote Bandwidth window displays.

Figure 355. Editing the Remote Bandwidth

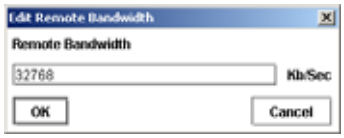

4. Change the bandwidth setting as desired.
   For example, change the value to 93 KB to use no more than about one-half the capacity of a T1 line.

Note: Both bandwidth settings are configured in kilobytes. Be careful when configuring this parameter as you may be used to using bits for networking settings.

F.11 Setting the Monitoring Variables for Remote Copy

There are four variables for monitoring Remote Copy. Notification for these variables comes as an alert message in the Console. You can configure Active Monitoring to receive email notification or for SNMP traps. The Remote Copy variables that are monitored include

- 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)

For detailed information about configuring Active Monitoring, see the Reporting chapter of the Storage System Software User's Manual.

F.12 Creating a Remote Snapshot Schedule

Scheduled remote snapshots provide high availability for business continuance/disaster recovery and provide a consistent, predictable update of data for remote backup and recovery.

The first step in using Remote Copy scheduling is to plan for scheduled creation and deletion of primary and remote snapshots. See "Planning for Remote Copy".

Once you have defined your plan, you are ready to create the remote snapshot schedule.

- First create the schedule
- Second, configure the primary volume and snapshot, and
- Third, create the remote volume and configure remote snapshots.

F.12.1 Creating the Schedule

1. Right-click the volume for which you want to create the remote snapshot schedule and then select Remote Copy > New Remote Snapshot Schedule.
2. From the Tasks menu, select New Schedule.
   The New Remote Snapshot Schedule window opens, shown in Figure 356.
3. Type a name for the schedule.
4. [Optional] Type a description for the schedule.

Figure 356. Creating a New Remote Snapshot Schedule

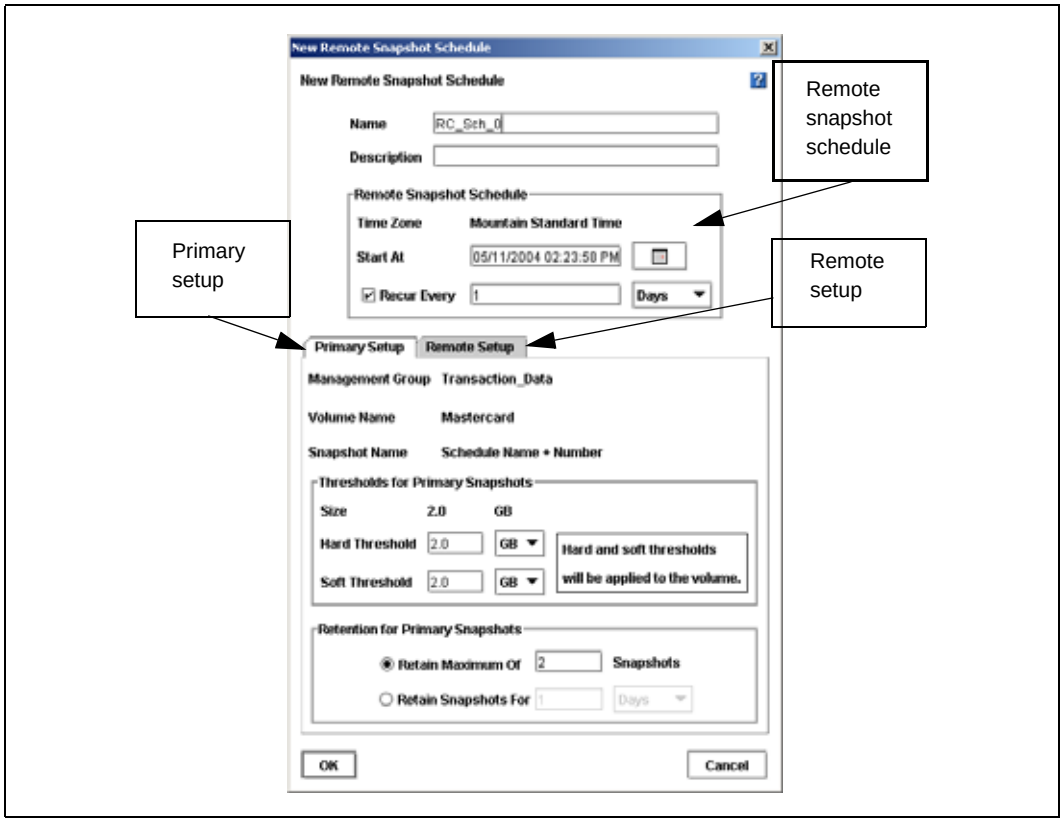

F.12.1.1 Remote Snapshot Schedule

The time zone displayed in the Remote Snapshot Schedule area is the time zone set on the SSM through which you are logged in to the management group.

F.12.1.2 Best Practice

Set all SSMs in the management group to the same time zone. Reset the management group time before creating a remote snapshot schedule. For detailed information, see “Resetting the Management Group Time” in the chapter “Working with Management Groups” in the Storage System Software User’s Manual.

1. Select a start date and time for the schedule.

2. [Optional] Select a recurrence interval for the schedule.

F.12.1.3 Configuring the Primary Volume and Snapshots

1. On the Primary Setup tab, specify the hard threshold and the soft threshold for the primary snapshots.
2. Specify the retention policy for the primary snapshots.

F.12.1.4 Configuring the Remote Volume and Snapshots

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

Figure 357. The Remote Setup Tab

2. Select the management group to contain the remote volume and remote snapshots.
3. Click New Volume to create the remote volume.

   You can use an existing volume as the remote volume. See “Making a Volume Into a Remote Volume”.
4. Specify a retention policy for the remote snapshots.
5. Click OK.

F.12.1.5 What the System Does

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.

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 will recur according to the schedule.

F.12.2 Editing a Remote Snapshot Schedule

When editing a remote snapshot schedule, 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

Note: Be certain to plan threshold changes carefully. See the chapter on snapshots in the Storage System Software User's Manual for detailed information about threshold requirements.

1. Select the primary volume that has the schedule you want to edit.
2. Click the Remote Snapshot Schedules tab.
3. Select from the list the schedule to edit.
4. From the Tasks menu, select Edit Schedule.
   The Edit Remote Snapshot Schedule window opens, shown in Figure 358.
5. Change the desired information.
6. Click OK.

Figure 358. Editing a Remote Snapshot Schedule

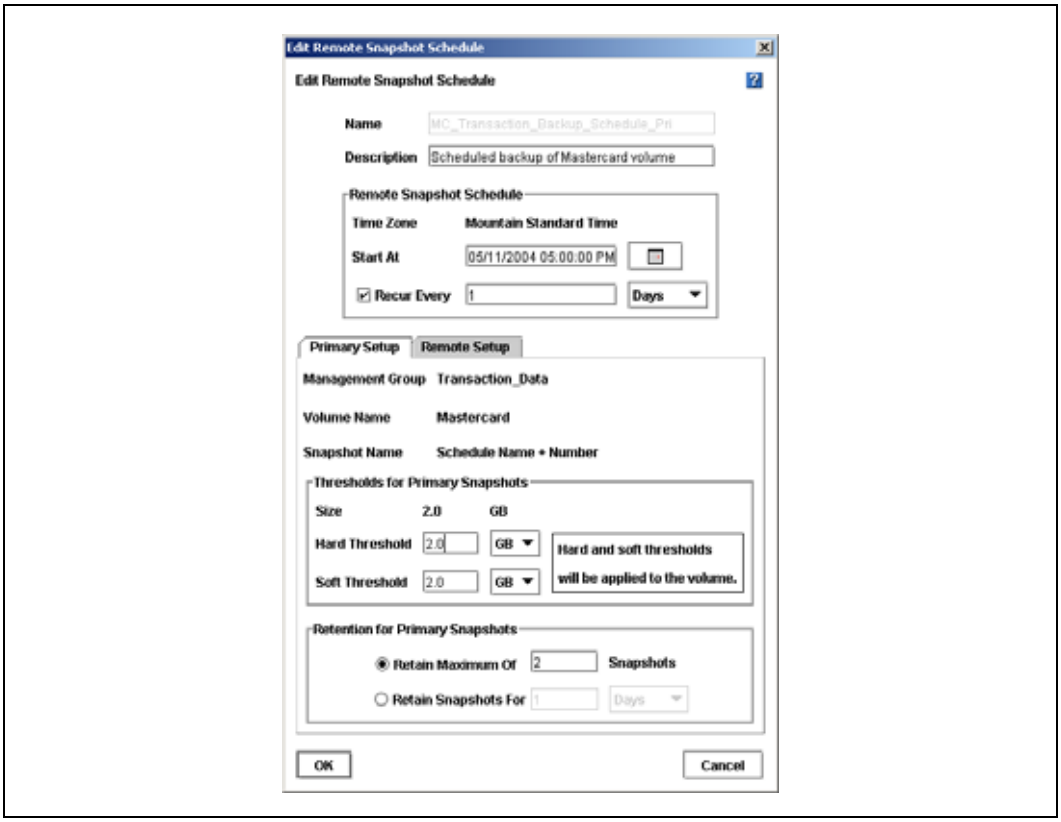

F.12.3 Deleting a Remote Snapshot Schedule

1. Select the volume for which you want to delete the remote snapshot schedule.
   The volume tab view opens.
2. Click the Remote Snapshot Schedule tab to bring it to the front.
3. Select the schedule you want to delete.
4. From the Tasks menu, select Delete Schedule.
   A confirmation message opens.
5. Click OK.

F.13 Changing the Roles of Primary and Remote Volumes

Changing the roles of primary and remote volumes comes into play during failover recovery. You use these procedures when you are resynchronizing data between the acting primary volume and the recovered or newly configured production site primary volume.

F.13.1 Making a Volume Into a Remote Volume

You can make any volume into a remote volume. First the system takes a snapshot of the volume to preserve the existing data that is 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.

1. Log in to the management group containing the volume that you want to convert.
2. Right-click the volume in the network view and select Remote Copy > Make Remote.
   The Make Volume Remote window opens, shown in Figure 359.

Figure 359. Making a Volume Into a Remote Volume

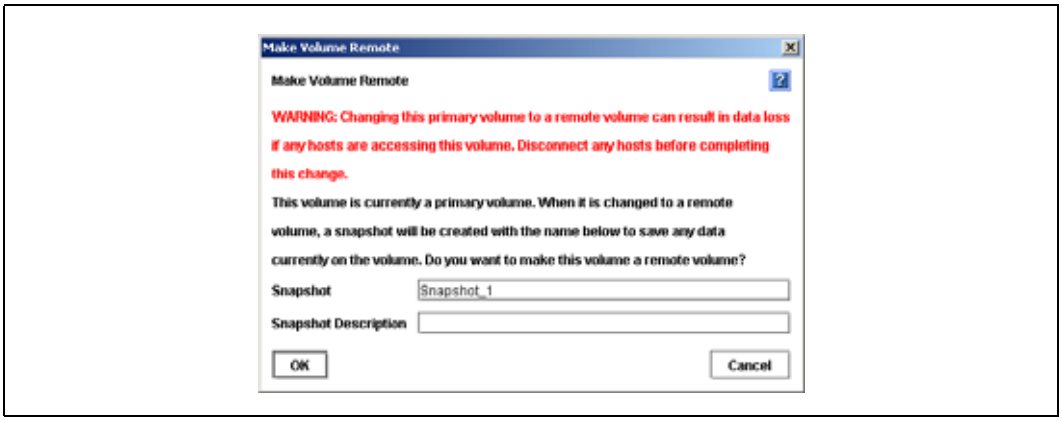

3. Type a name for the snapshot that will be created.
   This snapshot preserves any existing data on the volume.
4. [Optional] Type a description for the snapshot.
5. Click OK.
   The snapshot is created and the volume becomes a remote volume.

F.13.2 Making a Remote Volume Into a Primary Volume

You can 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.

1. Log in to the management group containing the remote volume that you want to convert.
2. Right-click the remote volume in the network view and select Remote Copy > Make Primary.
   The Make Primary window opens, shown in Figure 360.

Figure 360. Making a Remote Volume into a Primary Volume

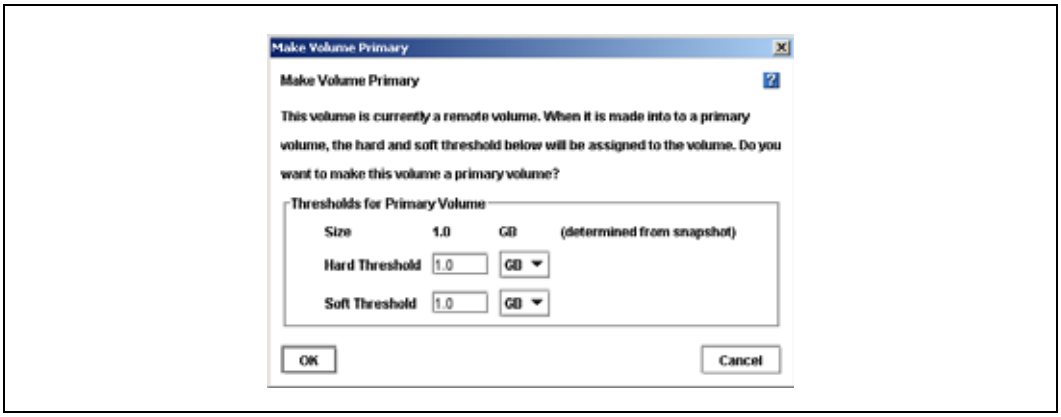

3. Set the hard and soft thresholds if required.
4. Click OK.

F.14 Creating Split Mirrors

Creating split mirrors is the process you use for data mining and data migration. A split mirror is a remote snapshot whose relationship to the primary volume has been severed. Split mirrors are usually created for one-time use and then discarded.

F.14.1 Creating a Read/Write Split Mirror

To create a read/write split mirror involves making a remote volume into a primary volume. Then you configure an EBSD client to access the new primary volume.

F.14.2 Creating a Read Only Split Mirror

To create a read only split mirror, simply configure an EBSD client to access a remote snapshot. The snapshot is read only, so it can act as a read only split mirror. See the EBSD Driver for Windows User's Manual for instructions about configuring EBSD clients.

F.15 Configuring Failover

Configuring Remote Copy for failover provides for business continuance and high availability. When configuring failover you take into consideration both the failover path and the recovery from failover.

F.15.1 Planning Failover

To achieve failover you plan the following parameters:

- 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]

F.15.2 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 snapshot schedules can be created and managed using scripts. Detailed information about Snapshot Scripting can be found in Chapter 15, “Working with Scripting” in the Storage System Software User’s Manual.

F.16 Resuming Production After Failover

After failover occurs, three scenarios exist for resuming production.

- Failback Recovery - return 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 are different for each scenario.

F.16.1 Synchronizing Data After Failover

After a failover, there will usually be two snapshots or volumes that have conflicting data. Recovering and synchronizing such data depends on multiple factors, including the application involved.

F.16.1.1 Example Scenario

The following example illustrates only one process for synchronizing data. Remember that such synchronization is optional.

Table 70. Example Scenario

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

F.16.1.2 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.

F.16.2 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.

F.16.2.1 Synchronizing the Data Between the Acting Primary Volume and the Original Primary Volume

F.16.2.1.1 Create Snapshots of Data

First you create snapshots that contain the data that you need to synchronize. The steps to create those snapshots are described inTable 71.

Table 71. Creating Snapshots of Data to Synchronize

Action/Activity	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 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 ActPrimarySS_0	Returns the remote management group to its original configuration.

F.16.2.1.2 Synchronize the Data

Synchronize the snapshots OrigPrimarySS_0 and ActPrimarySS_0 created in Steps 2 and 3 of Table 71 as appropriate for the application.

F.16.2.2 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 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 snapshot schedule on the new primary volume.
6. Reconfigure scripts for failover on the application servers.

F.16.3 Setting Up a New Production Site

Setting up a new production site involves creating a new primary volume and syncing up 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.

F.16.4 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 remote snapshot schedule on the acting primary volume.
2. Make a new remote volume on the new backup site as part of creating that remote snapshot or remote snapshot schedule.
3. Reconfigure scripts for failover on the application servers.

F.17 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 destroys any snapshots that were created after the snapshot that is rolled back to.

F.17.1 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.

F.17.1.1 Prerequisites

- Stop applications from accessing the volume. See the EBSD User Manual for detailed information about the steps involved.

Warning: After rolling back a volume to a snapshot, you lose all data that was stored since the rolled back snapshot was created.

Warning: Any uncompleted remote copy snapshot that is newer than the snapshot that you are rolling back to will be cancelled.

1. Log in to the management group that contains the primary volume that you want to roll back.
2. Select the snapshot that you want to roll back to.
3. Review the snapshot Details tab to ensure you have selected the correct snapshot.
4. From the Tasks menu, select Roll Back Volume.
   The Roll Back Volume window opens, shown in Figure 361.

Figure 361. Rolling Back a Primary Volume

5. Type a new name for the rolled back primary volume.
6. Select a volume type.
7. [Optional] If the volume is an iSCSI target and you want to use 1-way or 2-way CHAP, type a target secret. See the chapter on volumes in the Storage System Software User's Manual for detailed information about creating volumes.
   You can also change the hard threshold and soft threshold if necessary.

Table 72. Requirements for Rolling Back a Primary Volume

Item	Requirements for Changing
New Primary Volume Name	Must be from 1 to 127 characters. Names are case sensitive.
Hard Threshold	Hard threshold size must be equal to or less than the size of the volume.
Soft Threshold	Soft threshold size must be equal to or less than the hard threshold size.

8. Click OK.

   The Roll Back Volume confirmation message, shown in Figure 362, explains that the original primary volume and all newer primary snapshots will be deleted.

Figure 362. Verifying the Primary Volume Roll Back

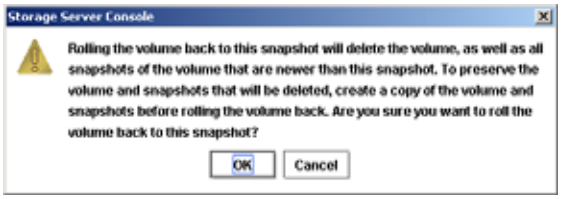

9. Click OK.

   The primary snapshot version of the primary volume is restored as a read/write volume.

10. Reconfigure application servers to access the new volume.

Warning: All primary snapshots between the current date and the roll back are deleted. The original primary volume is also deleted. The remote volume and remote snapshots remain intact.

F.17.2 Rolling Back a Remote Volume

A remote volume cannot be rolled back. In order to roll back a remote volume, you must make the remote volume into a primary volume.

F.18 Disassociate Remote Management Groups

Management groups become associated when linked by remote snapshots or remote snapshot schedules. When you have management groups that no longer share remote snapshots or remote snapshot schedules with each other, you can disassociate those management groups. Disassociating management groups destroys all the shared knowledge between those groups.

1. Log in to both management groups that you want to disassociate.
2. Right-click the remote management group and select Edit Remote Management Group.

   The Edit Remote Management Groups window opens, shown in Figure 363.

Figure 363. Editing a Remote Management Group

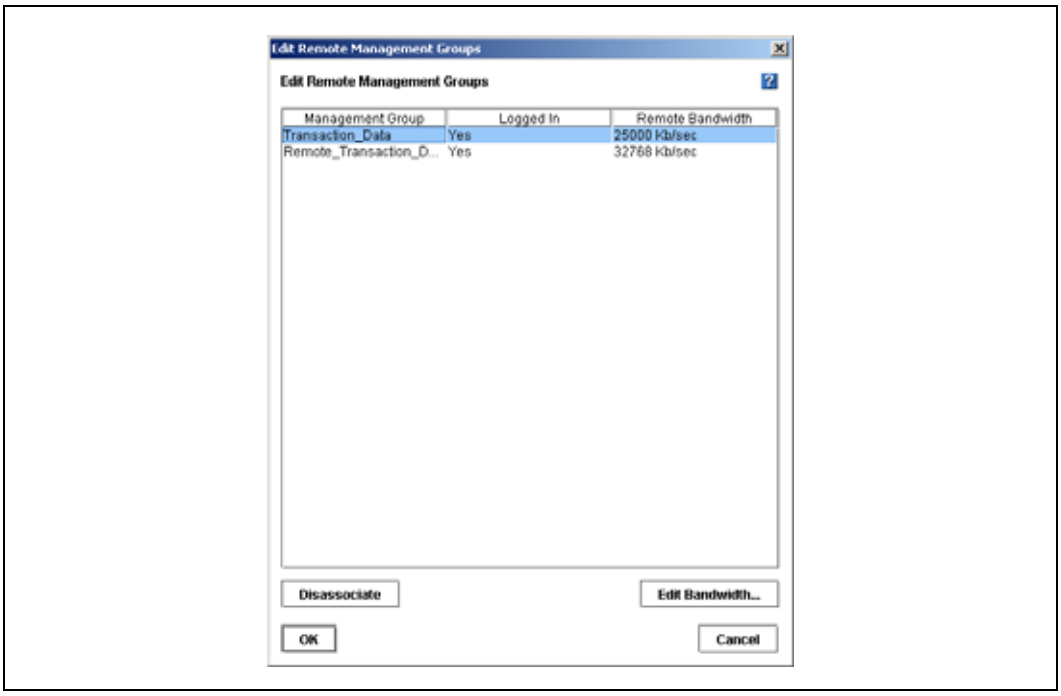

3. Select the management group or groups you want to disassociate.
4. Click Disassociate.

   A confirmation message opens, describing the results of disassociating the management groups.

Warning: Disassociating the management groups
- cancels any in-progress remote snapshots and
- deletes all snapshot schedules that are shared between the selected management groups

5. Click OK.

F.19 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.

F.19.1 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.

F.19.2 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 364 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.

F.19.2.1 Configuration Diagram

Figure 364. High Availability Example Configuration

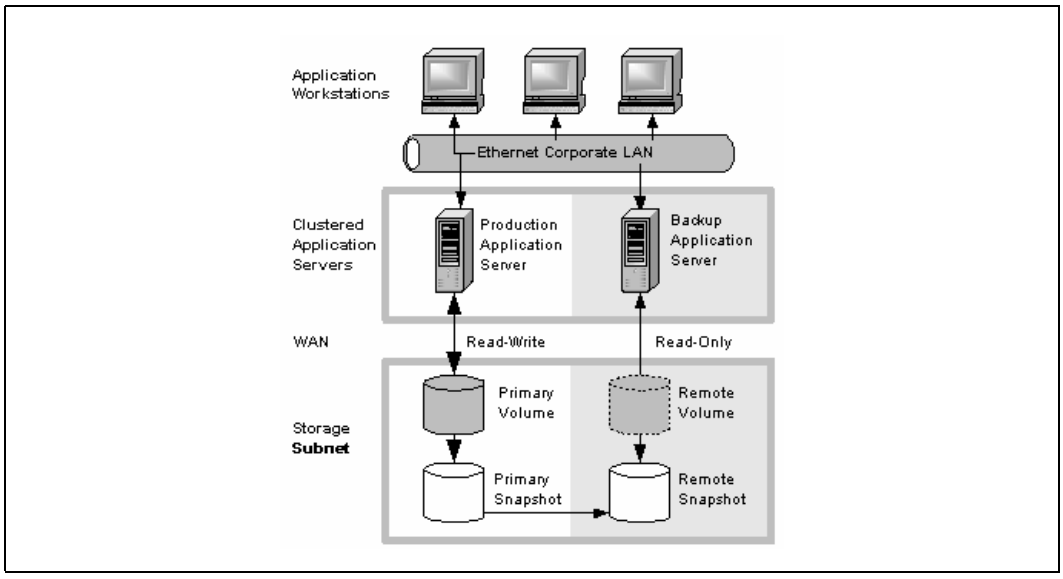

F.19.3 How This Configuration Works for High Availability

If the production application server or volumes become unavailable, application processing fails over to the backup application server. As shown in Figure 365, 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 365. High Availability Configuration During Failover

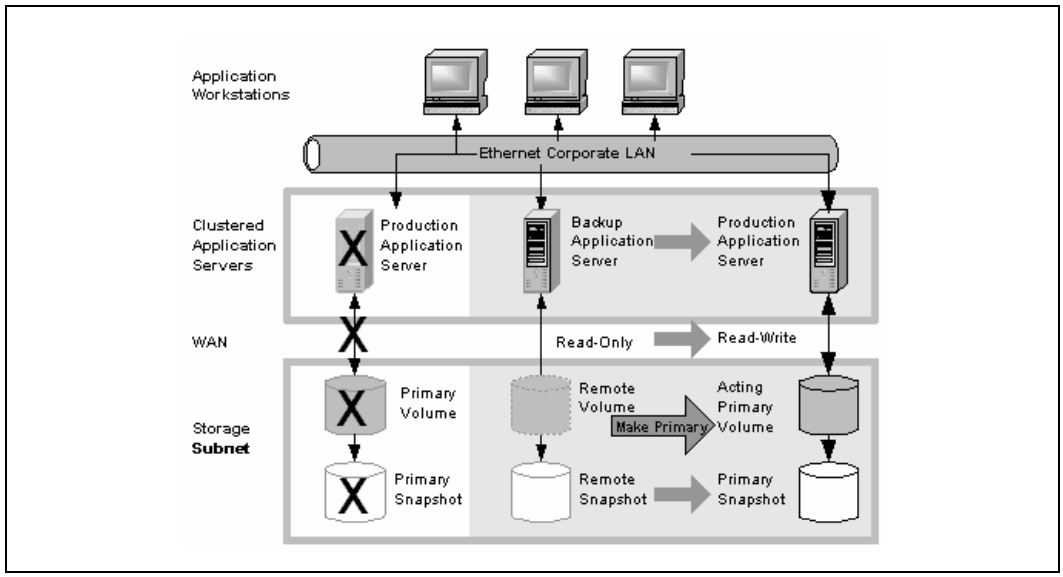

F.19.3.1 Data Availability If the Primary Volume or Production Application Server Fails

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.

F.19.3.2 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.

F.19.3.3 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 366. 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).

F.19.3.4 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. See “Synchronizing the Data Between the Acting Primary Volume and the Original Primary Volume” on page 395.

Figure 366. High Availability Configuration During Failback

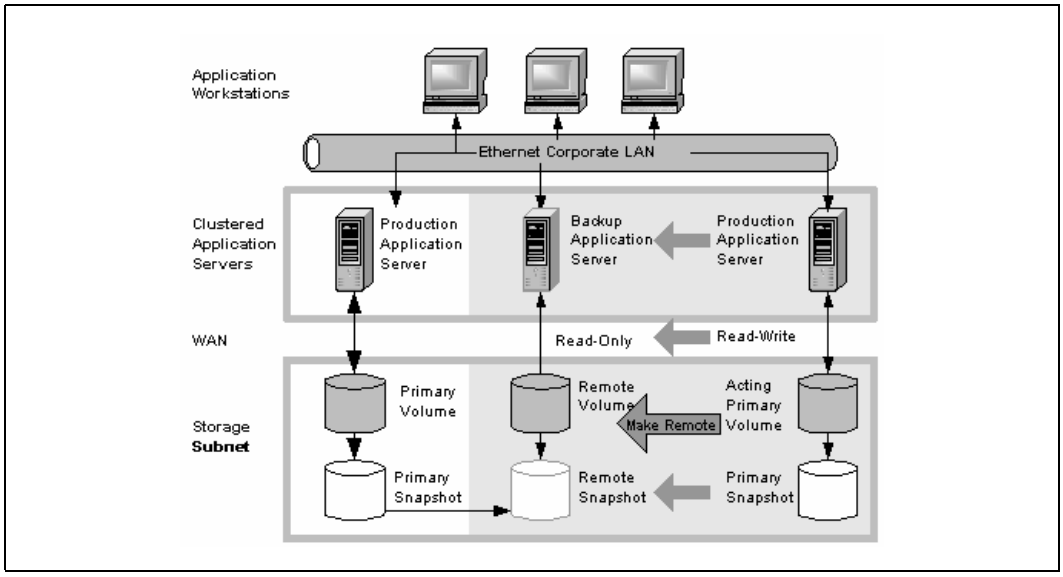

F.19.4 Best Practices

F.19.4.1 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, see the chapter on volumes in the Storage System Software User’s Manual for details.

F.19.4.2 Example Configuration

This example, illustrated in Figure 367, uses 3 SSMs per cluster. However, this scenario can use any number of SSMs. Information about creating clusters and volumes can be found in the Storage System Software User’s Manual.

- 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.
  See the EBSD User Manual for instructions about configuring EBSD clients.
- 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. See “Creating a Remote Snapshot Schedule”.

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 367. High Availability During Failover - Example Configuration

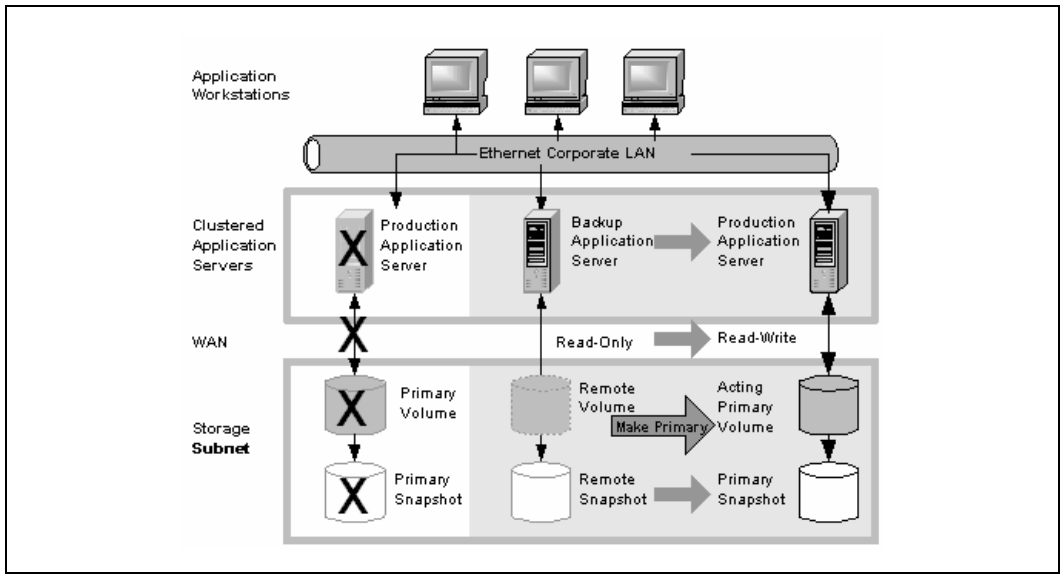

F.19.5 Achieving Affordable Disaster Recovery

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.

F.19.6 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.

F.19.6.1 Configuration Diagram

Figure 368. Affordable Disaster Recovery Example Configuration

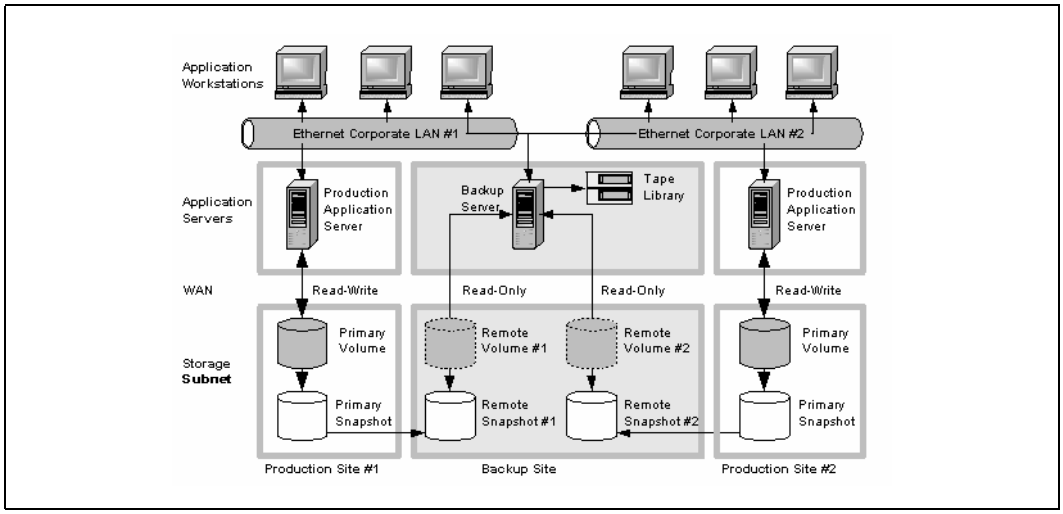

F.19.7 How this Works for Affordable Disaster Recovery

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 369. If you created tape backups, you can recover data from tape backups, as shown in Figure 370.
- 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 369. Restoring from a Remote Volume

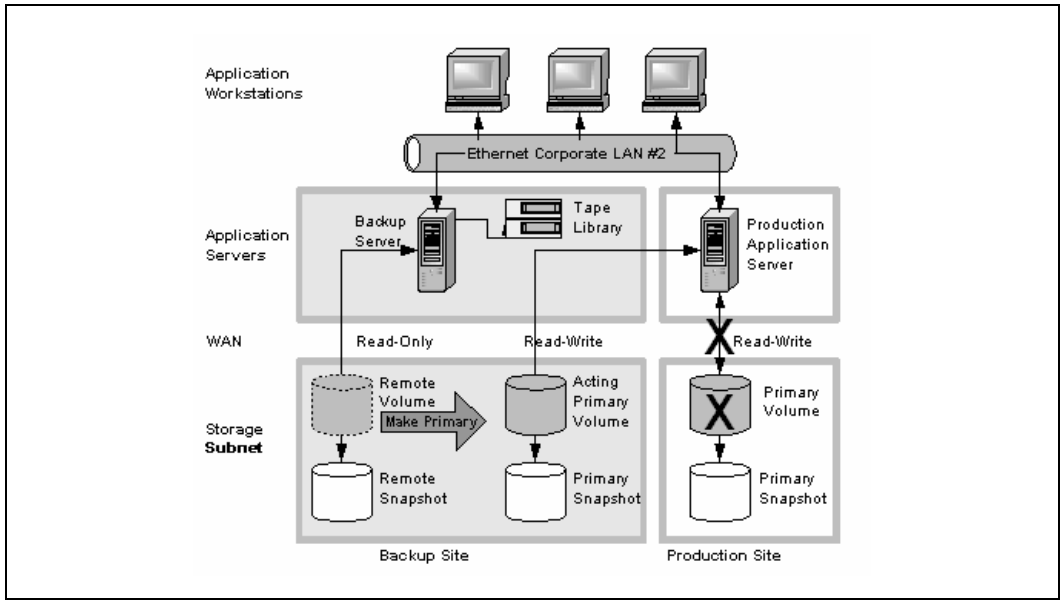

Figure 370. Restoring from Tape Backup

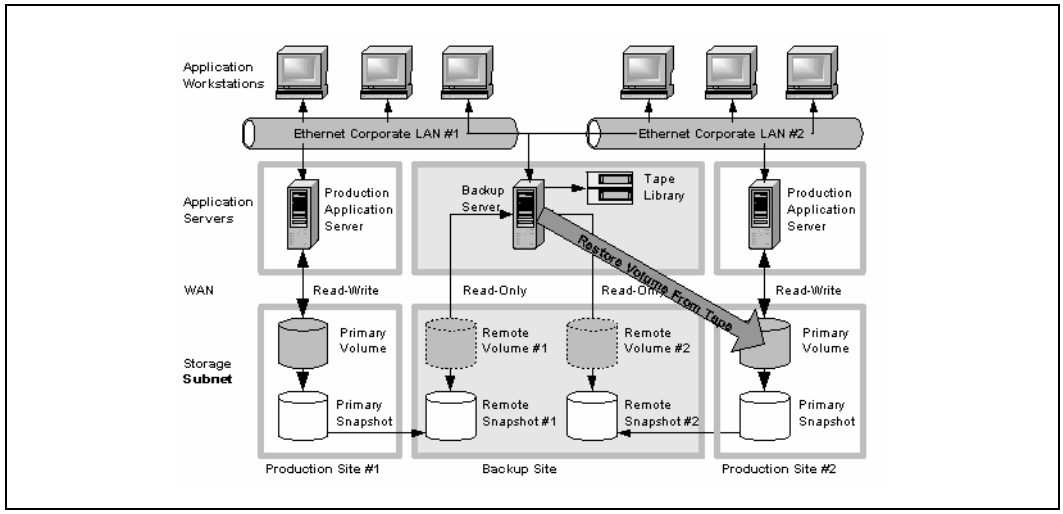

F.19.8 Best Practices

F.19.8.1 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.

F.19.8.2 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.

F.19.8.3 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.

F.20 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 scripting, configure the EBSD driver 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.

F.20.1 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.

F.20.2 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.
- Create routine tape backups of the remote snapshots.

See the example configuration illustrated in Figure 371.

F.20.2.1 Configuration Diagram

Figure 371. Off-site Backup and Recovery Example Configuration

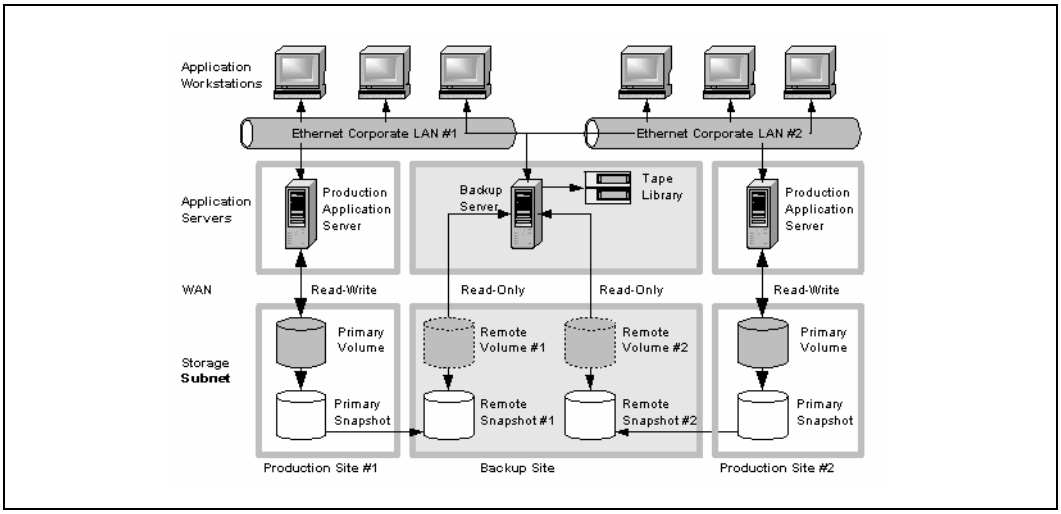

F.20.3 How This Configuration Works for Off-site Tape Backup

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.

F.20.4 Best Practices

F.20.4.1 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.

F.20.4.2 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.

F.20.5 Achieving Non-Destructive Rollback

As discussed in “Rolling Back a Primary Volume”, 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.

F.20.6 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 372 for an illustration of this configuration.

F.20.6.1 Configuration Diagram

Figure 372. Non-destructive Rollback Example

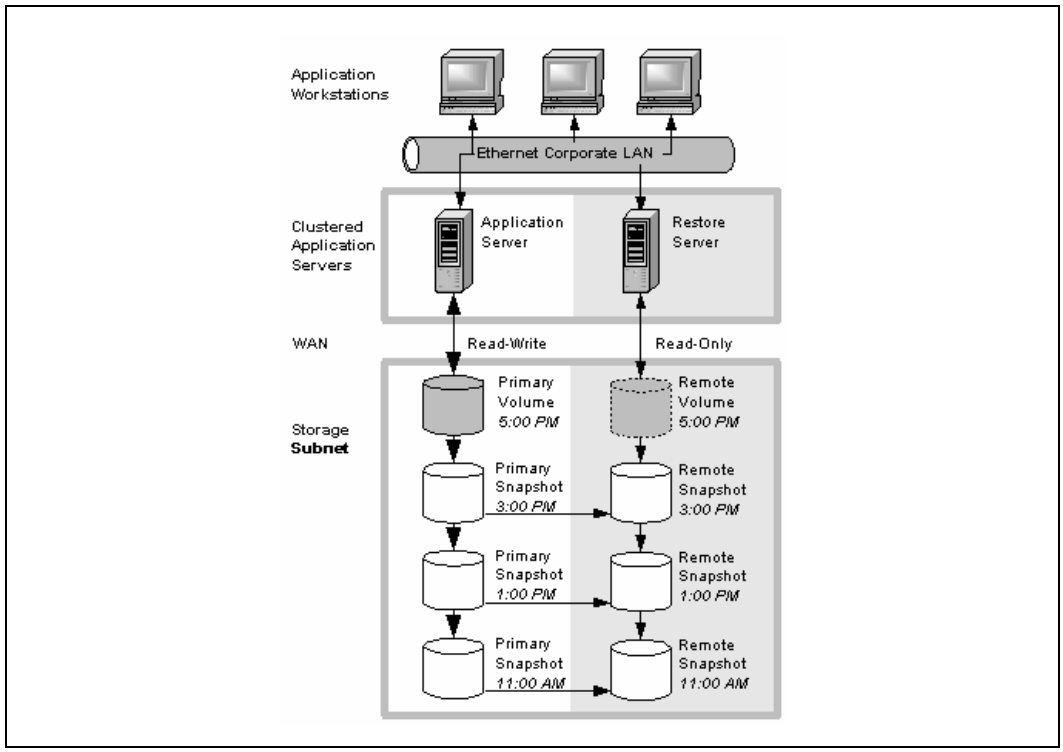

F.20.7 How This Configuration Works for Non-Destructive Rollback

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 373 shows rollback of the primary snapshot while leaving the remote snapshots intact.

Figure 373. Non-destructive Rollback from the Primary Snapshot

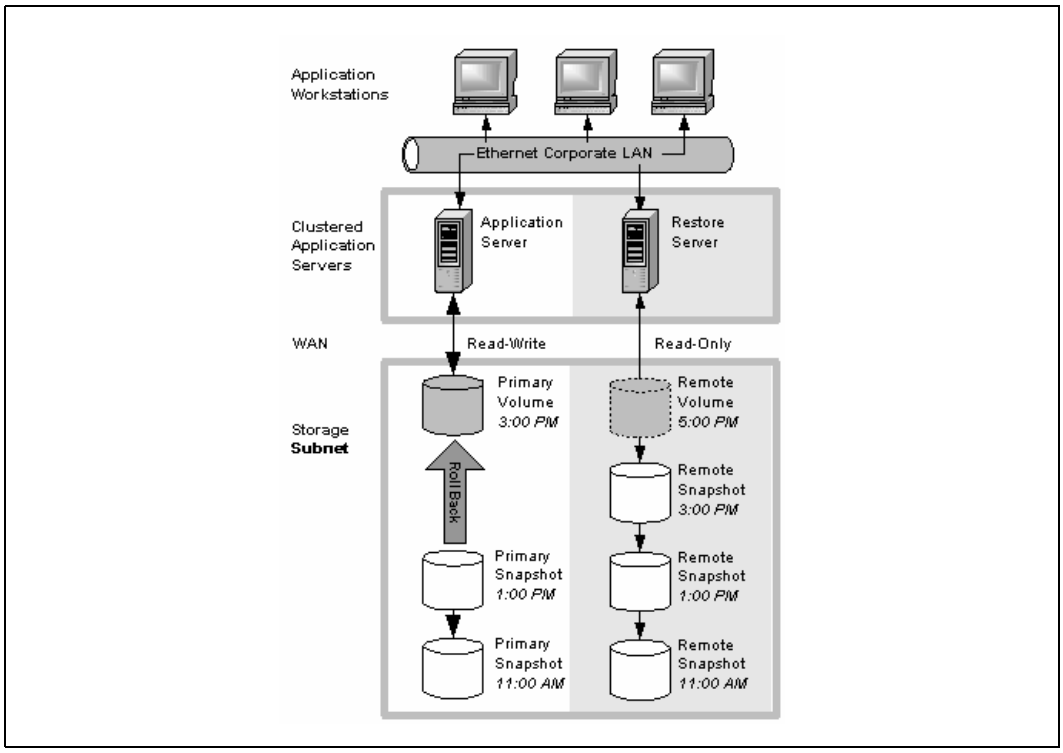

- To roll back the remote snapshot, you must first make the remote volume into a primary volume. This will stop scheduled creation of remote snapshots, which may jeopardize your high availability, disaster recovery, or routine backup strategies. Figure 374 shows rollback of the remote snapshot.

Figure 374. Non-destructive Rollback from the Remote Snapshot

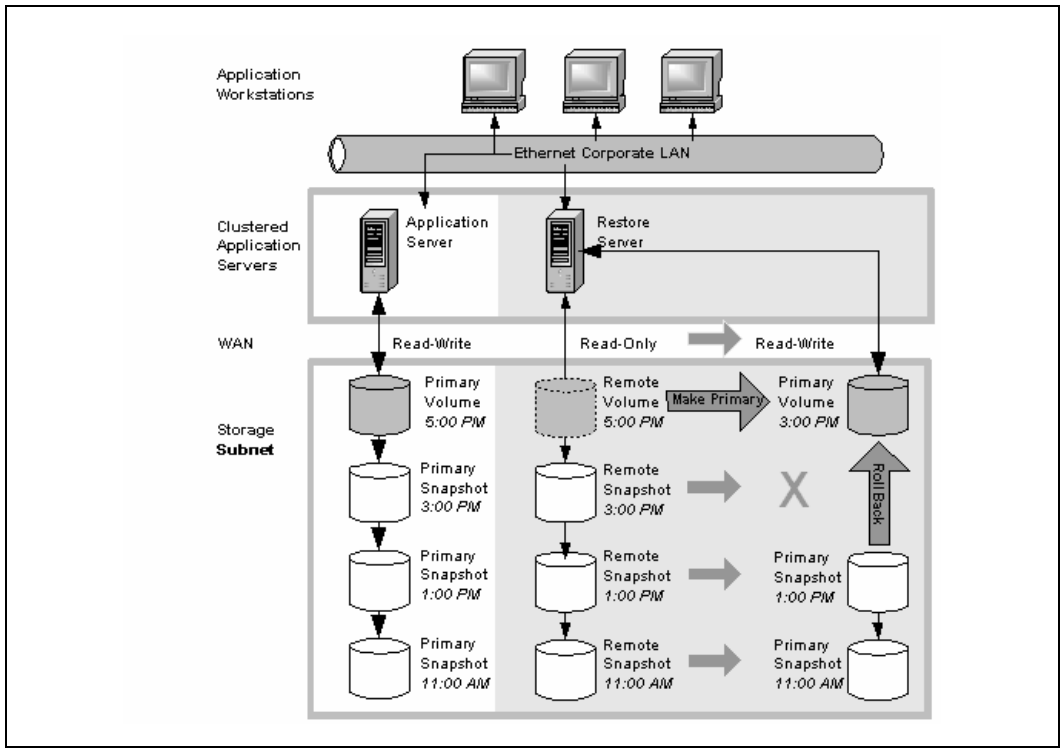

F.20.8 Best Practices

F.20.8.1 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.

   Remote snapshots remain intact.
3. After the primary volume is rolled back, scheduled creation of remote IP copies will continue.

F.21 Using Remote Copy for Data Migration

Remote Copy allows a one-time 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.

F.21.1 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.

F.21.2 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 375.

F.21.2.1 Configuration Diagram

Figure 375. Data Migration Example Configuration

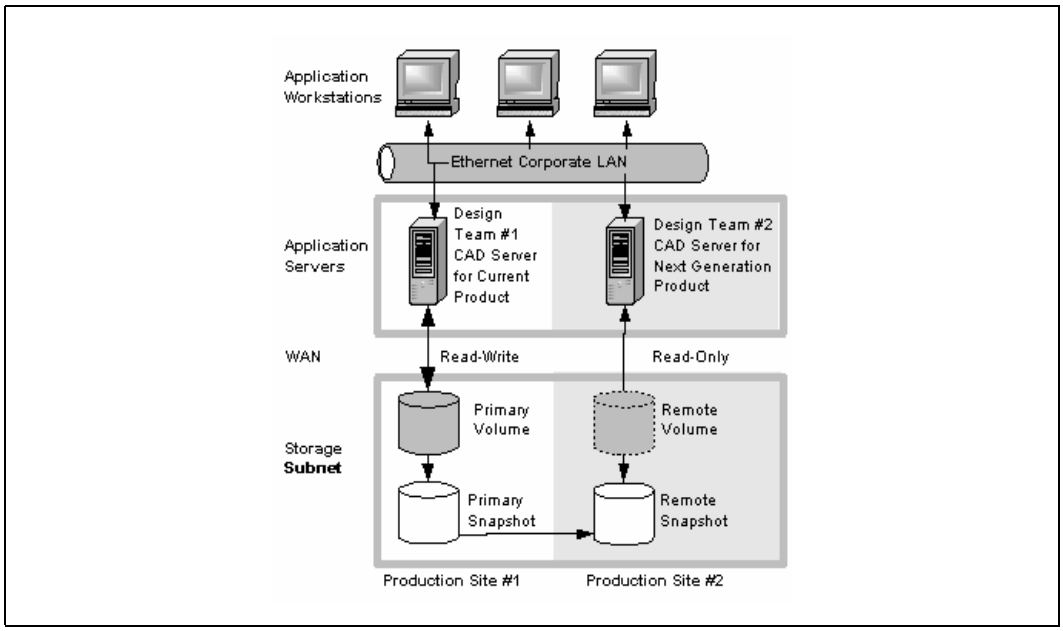

F.21.3 How This Configuration Works for Data Migration

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 376 shows migration of data by making a remote volume into a primary volume.

Figure 376. Configuration after Data Migration