



Netfinity[®] FAStT Storage Manager

Subsystem Management Online Help

P/N 19K8481

The following terms are trademarks of International Business Machines Corporation in the United States, other countries, or both:

IBM
Netfinity

Microsoft, Windows, and Windows NT are trademarks of Microsoft Corporation in the United States, other countries, or both.

JavaHelp is a trademark of Sun Microsystems, Inc. in the United States, other countries, or both.

Other company, product, and service names may be trademarks or service marks of others.

First Edition (April 2000)

INTERNATIONAL BUSINESS MACHINES CORPORATION PROVIDES THIS PUBLICATION "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some jurisdictions do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

This publication could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. IBM may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time.

This publication was developed for products and services offered in the United States of America and the United Kingdom. It is possible that this publication may contain references to, or information about, IBM products (machines and programs), programming, or services that are not announced in your country. Such references or information must not be construed to mean that IBM intends to announce such IBM products, programming, or services in your country.

Requests for technical information about IBM products should be made to your IBM reseller or IBM marketing representative.

No part of this publication may be reproduced or distributed in any form or by any means without prior permission in writing from the International Business Machines Corporation.

© **Copyright International Business Machines Corporation 2000. All rights reserved.**

Note to U.S. Government Users — Documentation related to restricted rights — Use, duplication or disclosure is subject to restrictions set forth in GSA ADP Schedule Contract with IBM Corp.

Introduction to the Subsystem Management Window

The Subsystem Management Window is Java-based software that is launched from the [Enterprise Management Window](#) and provides the following primary Storage Subsystem management functions:

- Provides Storage Subsystem options such as [renaming a Storage Subsystem](#), [changing a password](#), or enabling a background [media scan](#)
- Provides the ability to configure logical drives from your Storage Subsystem capacity, define [hosts](#) and [host groups](#), and grant host or host group access to sets of logical drives called [Storage Partitions](#).
- Monitors the health of Storage Subsystem components and reports a detailed status using appropriate icons
- Provides you with the appropriate recovery procedures for a failed logical or physical component
- Presents a view of the Storage Subsystem Event Log
- Presents profile information about physical components such as controllers and drives
- Provides controller management options such as [changing ownership of logical drives](#), setting a controller in [active](#) or [passive](#) mode, or [placing a controller online or offline](#).
- Provides drive management options such as [assignment of hot spares](#), and drive [initialization](#).
- Monitors Storage Subsystem performance

A [Logical View](#) in the left pane presents a picture of the organization of Storage Subsystem capacity into logical entities called [arrays](#) and [logical drives](#). A [Physical View](#) in the right pane presents a picture of the physical devices in the Storage Subsystem, such as controllers, drives, and other components. Selection of a logical drive or other entity in the Logical View shows you which physical components are associated with that entity in the Physical View.

For more information about the features of the Subsystem Management Window, see [Using the Subsystem Management Window](#).

Glossary

[Access Logical Drive](#)

[Active Mode](#)

[Active-Active or Active-Passive Controller Pair](#)

[AL_PA - Arbitrated Loop Physical Address](#)

[Arbitrated Loop](#)

[Array](#)

[Array Sequence Number](#)

[Assigned drive](#)

[Cache memory](#)

[Cache read-ahead](#)

[Cache Without Batteries](#)

[Channel Protection](#)

[Controller](#)

[Current Owner](#)

[Default Host Group](#)

[Defragment](#)

[Degraded Status](#)

[Enterprise Management Window](#)

[Environmental Card Canister](#)

[ESM](#)

[Event Log](#)

[Failed Status](#)

[Fibre Channel](#)

[Firmware](#)

[Free Capacity Node](#)

[Host](#)

[Host Adapter](#)

[Host Group](#)

[Host Interface](#)

[Host Port](#)

[Hot Spare Drive](#)

[I/O Data Connection](#)

[Initialization](#)

[Logical Drive](#)

[Logical Drive-to-LUN Mapping](#)

[Logical Unit Number \(LUN\)](#)

[Logical View](#)

[Management Station](#)

[Media Scan](#)

[Modification Priority](#)

[NL Port](#)

[NVSRAM file](#)

[Operation in Progress Status](#)

[Optimal Status](#)

[Parallel SCSI](#)

[Passive Mode](#)

[Physical View](#)

[Polling Interval](#)

[Port](#)

[Preferred Address - AL_PA](#)

[Preferred Loop ID](#)

[Preferred Owner](#)

[Premium Feature](#)

[RAID](#)

[RAID 0](#)

[RAID 1](#)

[RAID 3](#)

[RAID 5](#)

[RDAC Multi-path Driver](#)

[Reconstruction](#)

[Redundancy](#)

[Replaced Drive Status](#)

[SCSI Channel](#)

[SCSI ID](#)

[Segment Size](#)

[Storage Partition](#)

[Storage Subsystem](#)

[Subsystem Management Window](#)

[Topology](#)

[Unassigned drive](#)

[Unconfigured Capacity Node](#)

[Uninterruptible Power Supply \(UPS\)](#)

[Unresponsive Storage Subsystem](#)

[World Wide Name](#)


[World Wide Node Name](#)

[World Wide Port Name](#)

[Write Cache Mirroring](#)

Using the Subsystem Management Window Help System

The Subsystem Management Window online help is a JavaHelp™ system, featuring a window with a toolbar and two panes:




- **Toolbar** - Features a back arrow, to return to the previously viewed topic, and a forward arrow, to return to the topic that was displayed prior to going back. These are unavailable when you first open the Help Window.

Note: You can also use the back and forward arrows to return to the previously viewed section of a topic, if you are navigating within a topic using hypertext links.

- **Navigation Pane** - The left-hand pane that contains three navigation views: The Table of Contents, Index, and Full Text Search views. Switch between views by



selecting the appropriate tab. If you want to close the Navigation

Pane, click on the top arrow  on the splitter bar. This will close the pane and move the arrows to the left-hand side of the Help Window. Click the bottom arrow to reopen the Navigation Pane.

For more information, see [Tips on Using the Table of Contents](#), [Tips on Using the Index](#), and [Tips on Performing a Full Text Search](#).

- **Content Pane** - The right-hand pane that displays the help topics. If you want to close the Content Pane, click the bottom arrow on the splitter bar. This will close the Content Pane and move the arrows to the right-hand side of the Help Window. Click the top arrow to reopen the Content Pane.

Tips on Using the Table of Contents

- Double-click on a folder to expand or collapse the display of topics it contains.
- Highlight a topic to display its associated help information in the Content Pane.

Tips on Using the Index

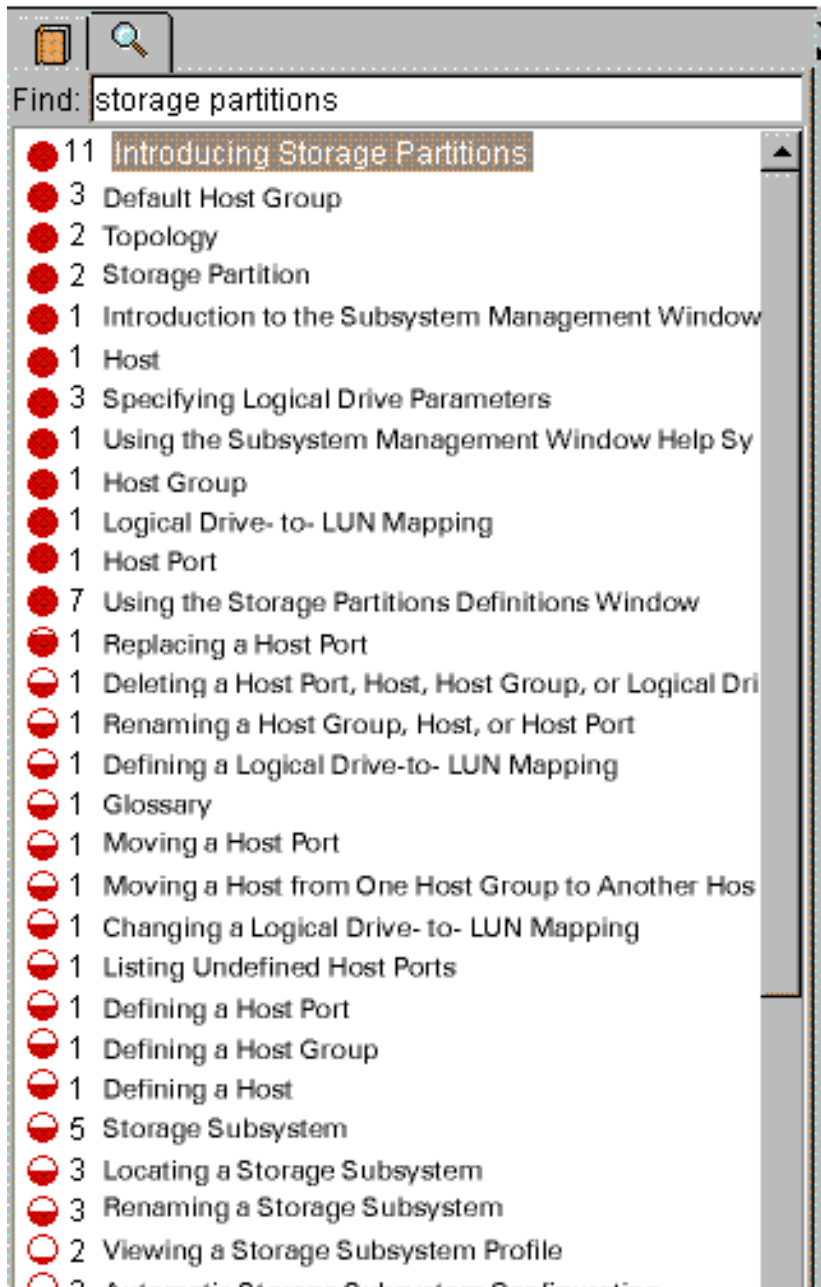
- Double-click on an entry to expand or collapse the display of topics under it.
- Highlight an entry to display its associated help information in the Content Pane.
- To search for an index entry term, type a word in the **Find:** textbox and press **Enter**. This will highlight the first index entry that matches your query. If the

index entry is directly associated with a topic, the topic is displayed in the Content Pane. To see the next match for your query, press **Enter**. Continue pressing **Enter** to see every match returned from your query, one at a time.

Tips on Performing a Full Text Search

To search for a natural language phrase, select the Search tab, type the phrase in the **Find:** text box, then press **Enter**.

Example: Typing in the phrase "storage partitions" produces the following results in the Full Text Search pane:



There are three columns returned in the results.

Column	Description
Ranking	The red circle indicates how well the topic ranked. The more complete the circle is filled in, the higher the ranking. The best match is shown at the top of the pane, followed by other matches in descending rank order.
	For more information, see How are

	the Topics Ranked? .
Number of Matches	The second column contains a number that indicates how many times the query was matched in the listed topic.
Topic title	The selected topic in the navigation pane displays in the topic pane with matched passages highlighted in the help text in an alternate color.

How are the Topics Ranked?

The JavaHelp search engine uses the following two techniques to score passages of text in the help topics as answers to your query.

Technique	Description
Relaxation Ranking	<p>Your query is compared with occurrences of the same or related terms in the help topics. The search engine attempts to find passages in the help topics in which as many as possible of the query terms occur in the same form and the same order.</p> <p>The search engine relaxes the query constraints to identify the passages in which:</p> <ul style="list-style-type: none"> ● Not all of the terms occur ● The terms occur in different forms ● The terms occur in a different order ● The terms occur with intervening words <p>The search engine assigns lower rankings to the passages depending on how they deviate from your query.</p> <p>Tip: To improve the ranking process, include as much information in your query as possible.</p>
Morphing	<p>The JavaHelp search engine finds words with common roots. For example, if you type the word "manage" in the Find: textbox, matches that contain "managed", "management", "manages", and "managing" are also returned.</p> <p>Note: In version 1.0 of JavaHelp, morphing is used only when the help system language is English.</p>

Opening the Help System

To open the help system from...	Procedure
The Subsystem Management Window	<p>Press F1, or select either the Help >> Quick Help or the Help >> Contents menu option.</p> <p>Result: The help system opens. The Quick Help option initially displays the topic "Using the Enterprise Management Window" in the content pane, and the Contents option initially displays the "Introduction to the Enterprise Management Software" in the content pane.</p>


The Storage Partitions Definitions Window	<p>Press F1, or select either the Help >> Quick Help or the Help >> Contents menu option.</p> <p>Result: The help system opens. The Quick Help option initially displays the topic "Using the Storage Partitions Definitions Window" in the content pane, and the Contents option initially displays the "Introduction to Storage Partitions" in the content pane.</p>
A dialog with a help button	<p>Select the help button on the dialog or press F1.</p> <p>Result: The help system opens and a context-sensitive help topic is displayed in the content pane.</p>

Printing the Help System

Printing is not supported by JavaHelp at the time of the help system development and release. However, the **AMW_help.pdf** file on your installation CD contains a copy of the Subsystem Management Window help system that can be printed using Adobe Acrobat Reader™.

Troubleshooting

There are a few known problems in the standard underlying Java components on which the JavaHelp system is based. The following table discusses known problems that you may encounter when using the help system.

If...	Then...
<p>Images are rendered incorrectly in the Content Pane, as shown in the example below:</p> 	<p>Display another help topic, then redisplay the topic with the incorrectly rendered graphic.</p> <p>Result: The graphic is correctly rendered.</p>
<p>You receive a "Null Pointer Exception" Java error while using the help system.</p>	<p>In most cases, you can ignore this error and re-attempt to access the particular topic. If the problem continues to occur, contact your Customer Support representative.</p>

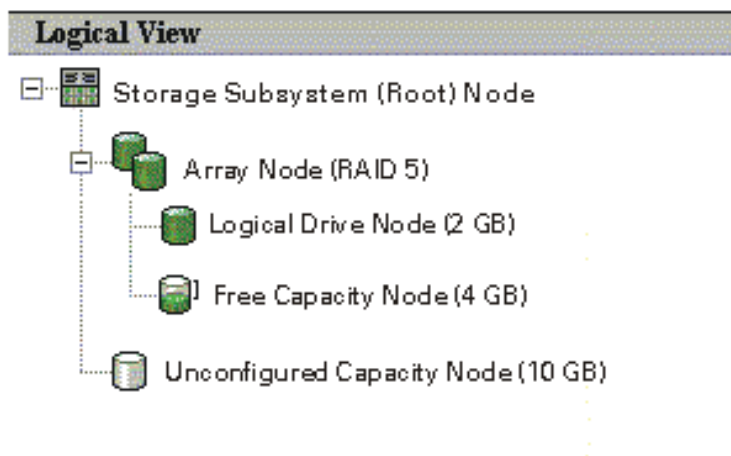
<p>When following a hypertext link in a help topic to another header within the topic, the referenced information displays in the center of the content pane instead of at the top.</p>	<p>This is a known problem, and there is no solution for the current release.</p>
<p>You receive a "Helpset File Error: Helpset is Corrupt or Not Found" error when accessing the help system.</p>	<p>Follow the instructions on the dialog to recover from the problem. If the problem continues to occur, contact your Customer Support representative.</p>







Using the Subsystem Management Window






The Subsystem Management Window has two views, a Logical View and a Physical View. To repartition the window, drag the splitter bar that divides the two views in the desired direction.

About the Logical View

The Logical View has five node types, as illustrated in the following graphic:



Node	Description
Storage Subsystem	<p>This is the root node of the Logical View tree. The icon shown reflects the current Storage Subsystem status, as follows:</p> <p>  Optimal  Needs Attention  Fixing  Unresponsive </p>
Array	<p>This node represents a defined array in the Storage Subsystem. Child nodes include logical drive nodes and free capacity nodes. The icon shown reflects the status (online or offline) of the array, as follows:</p> <p>  Online  Offline </p>

Logical Drive	<p>This node represents a defined logical drive on the array represented by the parent array node. The icon shown reflects the current logical drive status.</p> <p>  Optimal  Degraded Operation  in Progress  Failed  Offline </p>
Free Capacity	<p>This node represents a contiguous chunk of capacity that is present on drives included in an array, but has not been used to create logical drives. The capacity shown is RAID-factored to the defined RAID level of the array. There are no status icons associated with this node.</p> <p>Note: If there is non-contiguous free capacity on the array, then there will be multiple free capacity nodes.</p>
Unconfigured Capacity	<p>This node represents Storage Subsystem capacity that has not been used to create arrays and logical drives. The capacity shown reflects the raw capacity of all unassigned drives. There are no status icons associated with this node.</p>


Important:

- Double click on a node in the Logical View to expand or collapse a view of its child nodes.
- You can only select one node at a time in the Logical View.

About the Physical View

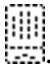
The Physical View displays four enclosure types, two are illustrated in the following graphic.

Physical View



Note: The orientation (vertical or horizontal) of the Physical View is determined by the controller.


The Physical View displays four enclosure types:

Enclosure Type	Description
Controller Enclosure	<p>Displays either one or two controller slots and a controller enclosure components button.</p> <p>Note: You can select each controller in the controller enclosure to perform controller management operations.</p>
Drive Enclosure	<p>Displays the drives or drive slots in the enclosure and one components button. An empty slot is shown with this icon .</p> <p>Note: You can select each drive in the drive enclosure to perform drive management</p>

	operations. Multiple drives can be simultaneously selected in the Physical View.
Combination Enclosure (not shown in the example at left)	Contains a representation of both controllers and drives in the same enclosure, and an enclosure components button.
Generic Enclosure (not shown in the example at left)	Shown as a number of drive slots, when information about the enclosure type is not received from the controller.

To change the enclosure order in the Physical View, use the **Storage Subsystem >> Change Enclosure Order** option. For more information, see [Changing the Enclosure Order](#).

About the Association Between the Logical and Physical View

Selection of a logical component causes an association icon  to be displayed by the selected logical object and the appropriate physical components.

Locating a Storage Subsystem

Use this option to physically locate and identify a Storage Subsystem. An indicator light is activated on each drive in the Storage Subsystem. You can use this option, for example, to identify the particular Storage Subsystem currently displayed in the Subsystem Management Window so that it can be renamed and labeled.

How to Locate a Storage Subsystem

- 1 To physically locate a Storage Subsystem, either select the **Storage Subsystem >> Locate** pull-down menu option, or use the right mouse button to select the Storage Subsystem in the Logical View and select **Locate** from the pop-up menu.

Result: The **Locate** dialog is displayed. The indicator lights on the drives in the Storage Subsystem flash.

- 2 When you have located the Storage Subsystem, select **OK**.

Result: The lights will stop flashing. If there are any other **Locate** operations (**Locate Array**, **Locate Drive**) currently being invoked from another management station, these lights will also stop flashing.

Note: If the lights from the **Locate Storage Subsystem** operation do not stop flashing, use the [Stop All Indications](#) option to stop the process manually.

Stopping All Indications

Use this option in the rare case when a **Locate** operation is not properly closed and the controller continues to flash the indicator lights on drives in the Storage Subsystem. Normally, the indicator lights stop flashing when you select **OK** in the **Locate Storage Subsystem**, **Locate Array**, or **Locate Drive** dialogs.

How to Turn Off Storage Subsystem Indicator Lights

- 1 To turn off indicator lights that are flashing in the Storage Subsystem, either select the **Storage Subsystem >> Stop All Indications** pull-down menu option, or use the right mouse button to select the Storage Subsystem in the Logical View and select **Stop All Indications** from the pop-up menu.

Result: All indicator lights that were flashing because of any **Locate** operation will stop.

- 2 If indications are successfully stopped, a confirmation message is displayed. Select **OK** on this message box to close it.

Viewing a Storage Subsystem Profile

Use this option to:

- View a description of all components and properties of the Storage Subsystem
- Save Storage Subsystem information to a text file

Note: You may want to use the Storage Subsystem Profile as an aid during recovery or as a snapshot of the Storage Subsystem's current configuration.

How to View a Storage Subsystem Profile

- 1 Either select the **Storage Subsystem >> Profile** pull-down menu option or use the right mouse button to select the Storage Subsystem in the Logical View, then select **Profile** from the pop-up menu.

Result: The **Storage Subsystem Profile** dialog is displayed.

- 2 Use the horizontal and vertical scroll bars to view the profile information.
- 3 Select **Close** to close the **Storage Subsystem Profile** dialog.

How to Interpret Storage Subsystem Profile Data

There are seven information areas in the **Storage Subsystem Profile** list. Much of this information is described in other places in the Help system. The information that is unique to the Storage Subsystem Profile display is described in the following table, and links to other help system topics are provided for more information. If the data requires no additional detail, it is not described.

Summary Area

The Access logical drive: field displays the LUN being used for the [access logical drive](#), if present.

For more information on the Start cache flushing at (in percentage), Stop cache flushing at (in percentage), and Cache block size (in KB) fields, see [Specifying Storage Subsystem Cache Settings](#).

For more information on the Media scan duration (in days), see [Changing the Media Scan Duration](#).

Controller Information

There are three data fields, [Topology](#), [World-wide port name](#), and [World-wide node name](#) that only apply to fibre channel host interfaces. These three terms are not displayed on the **Controller Properties** dialog.

For information about any other data field in the Controller Information area, see [Viewing Controller Properties](#).

Array Information

The [array sequence number](#) is displayed after each Array heading.

For more information about free capacity on an array, see [Adding Free Capacity to an Array](#).

Logical Drive Information

For more information about logical drive RAID level, RAID factored capacity of the logical drive, or channel protection, see [Creating a Logical Drive](#).

For more information about segment sizes, see [Creating a Logical Drive](#) and [Changing Segment Size](#).

For more information about logical drive cache settings or logical drive media scan settings, see [Specifying Logical Drive Properties](#).

Drive Information

For more information, see [Viewing Drive Properties](#).

Enclosure Information

The enclosure information is displayed, starting with the controller enclosure information and followed by the information about each drive enclosure in ascending enclosure ID order. For more information, see [Viewing the Status of Controller Enclosure Components](#) and [Viewing the Status of Drive Enclosure Components](#).

Storage Partitions Information

The Topology section describes the Storage Subsystem [Topology](#) definitions that have been configured for the Storage Subsystem.

Note: If no definitions have been configured, this section will list **Default Host Group**.

The [Logical Drive-to-LUN Mapping](#) section shows the mapping data for each logical drive in the Storage Subsystem.

Note: If no specific mappings have been configured, this section will list **Default Host Group** after each logical drive and its LUN number.

For more information, see [Introducing Storage Partitions](#).

How to Save a Storage Subsystem Profile to a Text File

- 1 Select **File**, then **Save As**.

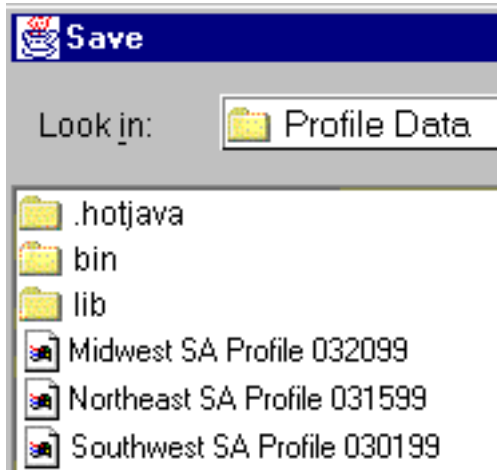
Result: The **Save Profile** dialog is displayed.

- 2 Choose an appropriate directory.
- 3 Type the file name of your choice in the **File Name** text box. You do not need to specify a file extension, but to associate the file with a particular software application that will open it, an extension such as .txt is helpful.

Example: Name the file Southwest SA Profile 030199. This name identifies the Storage Subsystem (Southwest) and the date of the profile (030199).

4 Select **Save**.

Result: A file containing the selected Profile data is saved to the Profile Data directory.



Setting or Changing a Storage Subsystem Password

Use the **Change Password** option to establish a new or change an existing Storage Subsystem password.

Unless you specify a password using this option, ALL options are available within this storage management software. If you specify a password, then any option that is destructive will be password protected. The controller firmware determines which options are destructive.



Caution: Implementing destructive commands on a Storage Subsystem can cause serious damage, including data loss. Use a long password with a minimum of 15 alphanumeric characters to increase security.

Important:

- If no password has previously been set, no current password is required to establish a new password.
- The password is stored on the Storage Subsystem. Therefore, if desired, you will need to provide a password for each Storage Subsystem.
- You cannot change a Storage Subsystem password unless the current password is supplied first. If you have forgotten the password, contact your customer support representative (the Storage Subsystem password must be cleared through a specific controller shell command).
- The maximum password length is 30 characters.
- Passwords are case sensitive. Remember your use of uppercase and lowercase letters when you change a password.
- Trailing white spaces are not stripped from passwords. Be careful not to include trailing spaces in the new password, because they can be difficult to enter accurately in the future.
- If you no longer want to have the destructive operations password protected, enter the current password, then leave the **New password** and **Confirm password** text boxes blank.

How to Set or Change a Storage Subsystem Password

- 1 From the Subsystem Management Window, either select the **Storage Subsystem>>Change Password** pull-down menu option or use the right mouse button to select the Storage Subsystem in the Logical View and select **Change Password** from the pop-up menu.

Result: The **Change Password** dialog is displayed.

Note: If you are setting the password for the first time, you don't need to enter the current password. Skip step 2 and go to step 3.

- 2 Enter the current password in the **Current password:** text box. Then, press the Tab key to move to the **New password:** text box.
- 3 Type the new password in the **New password:** text box. Then, press the Tab key on your keyboard to move to the **Confirm password:** text box.
- 4 Type the new password in the **Confirm password:** text box.

Important: The password in the **Confirm password:** text box must match the password in the **New password:** text box.

- 5 Select **OK** or press Enter to change the Storage Subsystem password.

Result: The **Password Successfully Changed** confirmation dialog is displayed.

- 6 Select **OK** to close the **Password Successfully Changed** dialog.

Note: While you are in the current management session, you will not be asked for a password when you attempt a destructive operation because you have just set the password from this management station. However, the next time you launch the Subsystem Management Window, you will be prompted for a password the first time you attempt a destructive operation.

Error Conditions

<p>Current Password Not Accepted</p>	<p>The password you typed in the Current password: text box did not match the Storage Subsystem password exactly. For example, there could be case inconsistencies or misspellings. Retry the operation using the correct current password.</p>
<p>New and Confirm Password Mismatch</p>	<p>The passwords entered in the New password: text box and the Confirm password: text box did not match. Retry the operation and type the same password in both text boxes.</p>

Entering a Storage Subsystem Password

The **Enter Current Password** dialog is displayed when:

- A Storage Subsystem password has been specified (as shown in [Setting or Changing a Storage Subsystem Password](#))
- -AND-
- You are attempting a password protected option

The options that are password protected are determined by the controller and are generally those that cause destructive Storage Subsystem operations.

Important:

- Only asterisks are displayed when you type a password.
- Passwords are case sensitive.
- Since trailing white spaces are not stripped from passwords when they are set, be careful to include white spaces if they were included as part of the password.
- There is a 30 character maximum password length.
- You will only be asked for the password once during a single management session. However, if a user managing the same Storage Subsystem from another Management Station changes the password while your session is in progress, you will be prompted for a password the next time you attempt a destructive operation.
- The first time a destructive operation is attempted during the management session, the **Enter Current Password** dialog is displayed containing the initial instructions. If an invalid password is entered, the **Enter Current Password** dialog is displayed again with instructions indicating that an invalid password was supplied. The **Enter Current Password** dialog will repeatedly display until either the correct password is supplied or you select **Cancel**.
- There is no limit to the number of times a password can be tried.

How to Enter a Storage Subsystem Password

- 1 In the **Enter Current Password** dialog, type the Storage Subsystem password.
- 2 Select **OK** to enter the password.

Result: The **Enter Current Password** dialog is closed, and the attempted operation completes.

Note: If you don't know the password, select **Cancel**. An error message is displayed indicating that the operation cannot be completed because a password has not been entered. Select **OK** to close it.

Using Passwords in the Script Editor

If destructive commands are contained in a script, a command line password must be entered in the script using the `Set` command. See the online help in the Enterprise Management Window for more information on the Script Editor.

Downloading Firmware or NVSRAM

Use this option to transfer a downloadable package ([firmware](#) or [NVSRAM](#)) to all Storage Subsystem controllers. Executing this option from the management station causes a package file transfer to a controller. After receiving the package, the controller's flash memory is updated with the new package and it is rebooted. If you are downloading a firmware package, the storage management software periodically polls the controller to see if the firmware version has changed to the new version. Next, the first controller in the download sequence passes the target package to the second controller. The second controller's flash memory is updated with the new package, and it is rebooted. If you are downloading a firmware package, the storage management software then polls the second controller periodically until the firmware version has changed to the new version.



Caution:

- If you do not have a multi-path driver installed on your application host operating system, stop I/O to the Storage Subsystem prior to the download to prevent application errors.
- The operating status of the Storage Subsystem controllers is checked when you perform a download operation, and an error message dialog is displayed if any controllers are not optimal. You are given the opportunity to stop or continue the download. Be sure to correct the non-optimal condition before continuing. However, if you choose to proceed with the download, the code on the non-optimal controller will be synchronized once the non-optimal condition has been corrected.
- You may experience an erroneous failure message, stating that the download could not complete, if a network problem occurs while the storage management software is polling a controller in the download sequence to verify that the firmware version has changed. It could be that the download did successfully complete, but the storage management software was unable to verify the operation success. Turn the power to the controller enclosure off and then on. Wait until the Storage Subsystem has finished initializing (this may take several minutes). Then, use the **Storage Subsystem >> Profile** option in the Subsystem Management Window to obtain the current firmware version. If the firmware listed does not match the version you tried to download, retry the download.

How to Download a Firmware or NVSRAM File

- 1 From the Subsystem Management Window, either select the **Storage Subsystem >> Download >> Firmware** or the **Storage Subsystem >> Download >> NVSRAM** pull-down menu option, or use the right mouse button to select the Storage Subsystem in the Logical View. Then, select **Download >> Firmware** or **Download >> NVSRAM** from the pop-up menu.

Result: The **Firmware Download** or **NVSRAM Download** Main Screen is displayed.

The **Current configuration** area lists the current firmware and NVSRAM versions in use by the Storage Subsystem controllers.

The **File Selection** area allows you to select a local or mapped directory and then displays downloadable package files compatible with the current Storage Subsystem configuration stored in that directory. You can then select a file from the displayed list and choose the one you wish to download.

- 2 In the **File Selection** area, change from the home directory that is initially displayed to the directory in which the file to download resides, by either clicking on the folder to open it, or typing the folder name in the **File name:** text box. Then, select the file as described above.

- 3 Select the package file to download from the **File Selection** area by one of the following:
 - Double-click on the file
 - Single-click on the file, then select **OK**
 - Enter the name of the file in the **File name:** text box, then select **OK**

Notes:

- Only downloadable packages that are compatible with the current Storage Subsystem configuration are displayed.
 - When you select a file in the **File Selection** area of the dialog, applicable attributes (if any) of the file are displayed in the **File Information** area.
- 4 Select **OK**.

Result: The **Firmware** or **NVSRAM Download** main screen is closed. The **Confirm Download** dialog is displayed.

Note: If the file selected was not valid or compatible with the current Storage Subsystem configuration, the **File Selection Error** dialog is displayed. Select **OK** on this dialog to close it, then choose a compatible firmware or NVSRAM file.

- 5 Select **Yes** on the **Confirm Download** dialog.

Result: The **Confirm Download** dialog is closed. The download begins, and the **Download Progress** dialog is displayed.

- 6 Monitor the progress and completion status of the download to the controllers. First, the file is transferred (downloaded) to the Storage Subsystem, then each controller is updated with this new file, one at a time. If the update succeeds to the first controller, the second controller then receives the update. The status of the file transfer and the update to each participating controller is displayed on the dialog, as described below.

Note: The progress and status of optimal controllers participating in the download is displayed. Other controllers in the Storage Subsystem are not represented on the dialog.

If the file transfer...

Then the following icon and text are shown

Is in progress



Transferring firmware (or NVSRAM)

Successfully completes



Firmware (or NVSRAM) transferred

Does not successfully complete




Firmware (or NVSRAM) transfer failed

If the update to a controller...

Then the following icon and text are shown

Is in progress




Updating firmware (or NVSRAM)

Is pending




Controller update pending

Successfully completes



Controller firmware (or NVSRAM) updated

Does not successfully
complete



Controller firmware (or NVSRAM)
update failed

If the download...	Result
Succeeded	The Download Progress dialog closes. The controllers are now operating with the new firmware or NVSRAM. Verify the new versions of these files using the Storage Subsystem >> Profile option.
Failed	The Download Failure Information dialog is displayed. Read the information displayed on this dialog. Run the Recovery Guru and follow any recommended recovery procedures.

Viewing Events with the Event Log

What is the Event Log?

The Event Log is a detailed record of events that occur on the Storage Subsystem. You can use the Event Log as a supplementary diagnostic tool to Recovery Guru for tracing Storage Subsystem events. However, always refer to Recovery Guru first when recovering from Storage Subsystem component failures.

Examples of events logged to the Event Log include:

- Configuration events, such as logical drive creation or the assignment of a drive as a hot spare
- Failures of Storage Subsystem components


The Event Log is stored in reserved areas on the disks in the Storage Subsystem.

The storage management software provides a way to view and sort the data in the Event Log, refresh the display to retrieve any new events, display detailed information about a selected event, and save selected Event Log data to a file.

The Event Log contains two levels of events: Critical and Informational. Critical events are sent to any Network Management Station (through [SNMP](#) traps) or e-mail recipient you configured with the **Edit >> Alert** option in the Enterprise Management Window. For more information about SNMP notification, see the online help for the [Enterprise Management Window](#).

How to View the Event Log

To view the Event Log from the [Subsystem Management Window](#), do one of the following:

- Select the **Storage Subsystem >>View Event Log** pull-down menu option.
- Use the right mouse button to select the Storage Subsystem in the Logical View, then select **View Event Log** from the pop-up menu.
- Press the **View Event Log** button  on the Subsystem Management Window toolbar.

Result: The **Event Log** window is displayed.

About the Event Log Window

Event Log Window Views

There are two possible event views in the window:

Summary View - Displays an event summary in a table form. See *Viewing and Interpreting Event Summary Data*, below.

Detail View - Displays details about a selected event. See *Viewing and Interpreting Event Details*, below.

When the **View Details** checkbox is selected, the window divides into the two views. To resize the panes, select the splitter bar between the two views and move it up or down.

How to Display Events

Note: When you open the Event Log Window, the default number of events displayed is 200.

- 1 Specify the number of events to retrieve in the **Display most recent events** spinner box.

The numbers increment by 25, or you can type in the number you want.

- 2 Select **Update** to retrieve new events from the Storage Subsystem for display.



The **Update** button:

- Refreshes the window to display the most recent events
- Retrieves a new number of events if the **Display most recent events** box is changed.
- Resorts the data back to the original order (by sequence number).

Viewing and Interpreting Event Summary Data

- All columns can be sorted by clicking on the column heading. A second click sorts the column in reverse order.
- Column widths can be changed by dragging the edge of the column header.

There are four columns in the Event Summary table:

Column	Description
Date/Time	<p>The date and time stamp of the error, according to the controller clock.</p> <p>Note: The Event Log initially sorts events based on sequence number (shown in the event details). Normally, this will correspond to the date/time. However, the two controller clocks in the Storage Subsystem could be unsynchronized. In this case, some perceived inconsistencies could be displayed in the Event Log relative to events and the date/time shown. Use the Storage Subsystem >> Set Controller Clocks option to synchronize Storage Subsystem controller clocks with the clock on your management station.</p>
Priority	<p>There are two priority values:</p> <ul style="list-style-type: none"> ● Critical  - There is a problem with the Storage Subsystem. However, immediate action by the user may prevent losing access to data. Critical events are used for SNMP trap notifications. All critical events are automatically sent to any Network Management Station (through SNMP traps) or e-mail recipient you configured with the Edit >> Alert option in the Enterprise Management Window. ● Information  - Non-critical information related to the Storage Subsystem.

Category	<table border="1"> <thead> <tr> <th data-bbox="412 195 646 289">Category</th> <th data-bbox="646 195 1425 289">Description</th> </tr> </thead> <tbody> <tr> <td data-bbox="412 289 646 443">Failure</td> <td data-bbox="646 289 1425 443"> <p>Some component on the Storage Subsystem has failed.</p> <p>Examples: Drive failure, battery failure.</p> </td> </tr> <tr> <td data-bbox="412 443 646 642">State Change</td> <td data-bbox="646 443 1425 642"> <p>An element of the Storage Subsystem has changed state.</p> <p>Examples: A logical drive transitioned to optimal; a controller transitioned to offline.</p> </td> </tr> <tr> <td data-bbox="412 642 646 842">Notification</td> <td data-bbox="646 642 1425 842"> <p>Internal controller operations that do not require user action.</p> <p>Example: The controller has completed start-of-day.</p> </td> </tr> <tr> <td data-bbox="412 842 646 995">Command</td> <td data-bbox="646 842 1425 995"> <p>A command has been issued to the Storage Subsystem.</p> <p>Example: A hot spare has been assigned.</p> </td> </tr> <tr> <td data-bbox="412 995 646 1110">Host entry</td> <td data-bbox="646 995 1425 1110"> <p>Host software has posted the entry into the Event Log.</p> </td> </tr> <tr> <td data-bbox="412 1110 646 1352">Error</td> <td data-bbox="646 1110 1425 1352"> <p>An error condition has been detected on the Storage Subsystem.</p> <p>Examples: A controller is unable to sync and purge cache; a redundancy error is detected on the Storage Subsystem.</p> </td> </tr> <tr> <td data-bbox="412 1352 646 1478">General</td> <td data-bbox="646 1352 1425 1478"> <p>Any events that do not fit well into any other category.</p> </td> </tr> </tbody> </table>	Category	Description	Failure	<p>Some component on the Storage Subsystem has failed.</p> <p>Examples: Drive failure, battery failure.</p>	State Change	<p>An element of the Storage Subsystem has changed state.</p> <p>Examples: A logical drive transitioned to optimal; a controller transitioned to offline.</p>	Notification	<p>Internal controller operations that do not require user action.</p> <p>Example: The controller has completed start-of-day.</p>	Command	<p>A command has been issued to the Storage Subsystem.</p> <p>Example: A hot spare has been assigned.</p>	Host entry	<p>Host software has posted the entry into the Event Log.</p>	Error	<p>An error condition has been detected on the Storage Subsystem.</p> <p>Examples: A controller is unable to sync and purge cache; a redundancy error is detected on the Storage Subsystem.</p>	General	<p>Any events that do not fit well into any other category.</p>
Category	Description																
Failure	<p>Some component on the Storage Subsystem has failed.</p> <p>Examples: Drive failure, battery failure.</p>																
State Change	<p>An element of the Storage Subsystem has changed state.</p> <p>Examples: A logical drive transitioned to optimal; a controller transitioned to offline.</p>																
Notification	<p>Internal controller operations that do not require user action.</p> <p>Example: The controller has completed start-of-day.</p>																
Command	<p>A command has been issued to the Storage Subsystem.</p> <p>Example: A hot spare has been assigned.</p>																
Host entry	<p>Host software has posted the entry into the Event Log.</p>																
Error	<p>An error condition has been detected on the Storage Subsystem.</p> <p>Examples: A controller is unable to sync and purge cache; a redundancy error is detected on the Storage Subsystem.</p>																
General	<p>Any events that do not fit well into any other category.</p>																
Description	<p>A description of the event.</p> <p>Example: Drive write failure - retries exhausted</p>																

Viewing and Interpreting Event Details

Use the **View Details** checkbox to show event details of a single selected Summary event. When you have enabled the **View Details** checkbox, click on a single event in the Summary View to show the details about that event in the **Event Details**

View.

Note: Because some of the fields in the **Event Details** View require no explanation, this table only focuses on those fields which require additional detail.

Selected Field	Description
Sequence Number	A 64-bit number that uniquely identifies a specific log entry for a given Storage Subsystem. This number increments with every new entry into the Event Log. Example: 2402
Event type	A 4-digit number that uniquely corresponds to a particular event priority, category and description. Example: 1006 For more information, see Critical Event Descriptions .
Sense key/ASC/ASCQ	Information from the controller that provides additional information about the event. Example: 6/3F/80
Raw data	Hex dump of the raw data contained in the entire Event Log entry. Example: 0x01 00 30 00 04 ff 00 35....

How to Save Selected Events to a File

To save selected events from the Summary area to a file, perform the following steps:

- 1 Select the events you want to save to a file in the Summary View.
 - Press CTL + click to save multiple events that are not next to each other in the display
 - Press SHIFT + click to save a range of events, clicking on the first and the last event you wish to save
 - Select **Select All** to save every displayed event

- 2 Select **Save As**.

Result: The **Save Events** dialog is displayed.

- 3 Save the selected events by choosing an appropriate directory and filename.

The **Save** dialog can filter on files with a `.log` extension. Therefore, you may want to save the file with a `.log` extension using the filename conventions you used for any previous log files.

- 4 Type the file name of your choice in the **File Name** text box.

Example: Name the file `january13.log`.

- 5 Select **Save**.

Result: A file containing the selected events, with a `.log` extension, is saved to the directory.


Monitoring Performance

Use the Performance Monitor to monitor Storage Subsystem performance in real-time and save performance data to a file for later analysis.

Important:

- The Performance Monitor will not dynamically update its display if any configuration changes (for example, creation of new logical drives, change in logical drive ownership, and so on) occur while the monitor window is open. The Performance Monitor window **MUST** be closed and then re-opened for the changes to appear.
- Using the Performance Monitor to retrieve performance data can affect the normal Storage Subsystem performance depending on the [polling interval](#) that you set. For more information, see [Specifying Performance Monitor Settings](#).
- If the Storage Subsystem you are monitoring begins in or transitions to an Unresponsive state, an informational dialog is displayed, stating that the Performance Monitor cannot poll the Storage Subsystem for performance data.

How to Monitor Performance

- 1 From the Subsystem Management Window, do one of the following:
 - Select the Performance Monitor toolbar button .
 - Select the **Storage Subsystem >> Monitor Performance** pull-down menu option.
 - Select the Storage Subsystem node in the Logical View, then **Monitor Performance** from the right-mouse pop-up menu.

Result: The **Performance Monitor** dialog is displayed.

- 2 Select **Settings** to specify the logical drives and/or controllers to monitor and to specify the polling interval. To change the polling interval, choose a number of seconds in the spin box. Each time the polling interval elapses, the Performance Monitor re-queries the Storage Subsystem and updates the statistics in the table.

3 Select **Start**.

Result:

- The **Start:** field at the bottom of the dialog records the current date and time. This information is displayed on the screen until you select **Stop**. (When you select **Stop**, the **Stop:** field at the bottom of the dialog records the stop date and time).
- The **Time Monitored:** field displays the elapsed time between the present time and the time the **Start** button was selected.
- Values are displayed for the selected devices in the Performance Monitor data table. It is updated at the interval specified in the Polling Interval setting. Press **Update** to force an immediate poll of the Storage Subsystem. For more information, see **Interpreting Storage Subsystem Performance Data**.

4 Select **Stop** to stop monitoring the Storage Subsystem.

Interpreting Storage Subsystem Performance Data

The columns in the Performance Monitor data table are described next. Use this data to make performance tuning decisions for your Storage Subsystem. For more information, see [How to Tune Storage Subsystem Performance](#).

Column	Description
Devices	<p>The Device column can display the following values:</p> <ul style="list-style-type: none">● Controller, shown as Controller in Slot A or B - Represents the totals for all logical drives owned by the controller, not just the logical drives presently selected for monitoring.● Logical Drive, shown as Logical Drive name; where Logical Drive name is the name assigned to the logical drive when it was created or after it was renamed.● Storage Subsystem Totals - represents the totals for both controllers in an active-active controller pair, regardless if one, both, or neither are selected for monitoring.

Total I/Os	The number of total I/Os performed by this device since the Start button was pressed. When this number exceeds 99,999, it is displayed in thousands (K), beginning with 100K until the number reaches 9999K, at which time it is displayed in millions (M). For amounts greater than 9999K but less than 100M, the value is displayed in tenths (for example, 12.3M).
Read Percentage	The percentage of total I/Os (previous column) that are read operations for this device. Write percentage can be calculated as 100 minus this value.
Cache Hit Percentage	The percentage of reads that are fulfilled by data from the cache rather than requiring an actual read from disk.
Current KB/second	The current KB/second, or transfer rate. Current means the number of KB/second since the last time the polling interval elapsed, causing an update to occur, or since the Update button was pressed.
Maximum KB/second	The highest value achieved by the current KB/second column for this polling session (since Start was selected.)
Current IO/second	The current I/Os per second. Current means the number of I/Os per/second since the last time the polling interval elapsed, causing an update to occur; or, since Update was pressed.
Maximum IO/second	The highest value achieved by the current I/Os per/second column for this polling session (since Start was selected.)

How to Save Performance Data to a File

- 1 Select **Save As** on the Performance Monitor main screen to save the currently displayed performance statistics.

Result: The **Save Performance Statistics** dialog is displayed.

- 2 Save the selected events by choosing an appropriate directory and filename.

The **Save** dialog can filter on files with a `.perf` extension. Therefore, you may want to save the file with a `.perf` extension using the filename conventions you used for any previous performance monitor statistics files. The `.perf` extension is the default.

- 3 Type the file name of your choice in the **File name** text box.

Example: Name the file `feb14A.perf`.

- 4 Select the file type of your choice in the **Files of type** text box.

File Type	Select this type when...
Report Format	You want to save the data in a report form for viewing or printing.
Comma Delimited Format	You want to save the data in a form that can be imported into a commercial spreadsheet application for further analysis (most leading commercial spreadsheet applications recognize a comma delimiter and use the delimiter to import the data into spreadsheet cells appropriately.)

- 5 Select **Save**.

Result: A file containing the performance statistics, with a `.perf` extension or another extension you have specified, is saved to the directory.

How to Tune Storage Subsystem Performance

Use Performance Monitor data to make Storage Subsystem tuning decisions, as described in the table below.

Data	Implications for performance tuning
Total I/Os	<p>This data is useful for monitoring the I/O activity of a specific controller and a specific logical drive. This is needed to identify possible high-traffic I/O areas.</p> <p>Identify actual I/O patterns to the individual logical drives and compare those with the expectations based on the application. If a particular logical drive has considerably more I/O activity, consider moving the logical drive to the other controller in the Storage Subsystem.</p> <p>You might notice a disparity in the Total I/Os (workload) of controllers, for example, the workload of one controller is heavy or is increasing over time while the workload of the other controller is lighter or more stable. In this case, you might consider changing the controller ownership of one or more arrays to the controller with the lighter workload. Use the logical drive Total I/O statistics to determine which arrays to move. See Changing Controller Ownership.</p> <p>If you notice the workload across the Storage Subsystem (Storage Subsystem Totals Total I/O statistic) continues to increase over time while application performance decreases, this might indicate the need to add additional Storage Subsystems to your enterprise so that you can continue to meet application needs at an acceptable performance level.</p>
Read Percentage	<p>Use the Read Percentage for a logical drive to determine actual application behavior. If there is a low percentage of read activity relative to write activity, you might consider changing the RAID level of an array from RAID 5 to RAID 1 for faster performance. See Changing the RAID Level of an Array.</p>

Cache Hit Percentage

A higher percentage is desirable for optimal application performance.

The cache hit percentage of all of the logical drives may be low or trending downward. At the Storage Subsystem or controller level, this can indicate the need to install more controller cache memory if you do not have the maximum amount of memory installed.

If an individual logical drive is experiencing a low cache hit percentage, consider enabling cache read ahead for that logical drive. [Cache read ahead](#) can increase the cache hit percentage for a sequential I/O workload.

Determining the effectiveness of a logical drive cache read ahead multiplier

To determine if your I/O has sequential characteristics, try enabling a conservative cache read-ahead multiplier (4, for example) using the **Logical Drive >> Properties** option. Then, examine the logical drive cache hit percentage to see if it has improved. If it has, indicating that your I/O has a sequential pattern, enable a more aggressive cache read-ahead multiplier (8, for example). Continue to customize logical drive cache read-ahead to arrive at the optimal multiplier (in the case of a random I/O pattern, the optimal multiplier is zero). See, [Specifying Logical Drive Properties](#). There is a positive correlation between the cache hit percentage and I/O rates.

<p>Current KB/sec and Maximum KB/sec</p>	<p>The transfer rates of the controller are determined by the application I/O size and the I/O rate. In general, a small application I/O request size results in a lower transfer rate but provides a faster I/O rate and a shorter response time. With larger application I/O request sizes, higher throughput rates are possible. Understanding your typical application I/O patterns can give you an idea of the maximum I/O transfer rates that are possible for a given Storage Subsystem.</p> <p>Consider a Storage Subsystem equipped with controllers and fibre channel interfaces that supports a maximum of 100 MB per second transfer rate, or 100,000 KB per second. You are typically achieving an average transfer rate of 20,000 KB per second on the Storage Subsystem. This KB per second average transfer rate is a function of the typical I/O size for the applications using the Storage Subsystem. (The typical I/O size for your applications is 4K, 5,000 I/Os can be transferred per second to reach an average transfer rate of 20,000 KB.) In this case, I/O size is small and there is system overhead associated with each I/O transferred, so you can never expect to see transfer rates that approach 100,000 KB per second. However, if your typical I/O size is large, a transfer rate within a range of 80,000 - 90,000 KB per second might be achieved.</p> <p>Wide Ultra SCSI supports sustained data transfer rates of up to 40MB/second with large host I/O sizes. The maximum transfer rates with relatively small I/O sizes, 4096 bytes, is approximately 17 MB/second.</p>
--	---

<p>Current I/O per sec and Maximum I/O per sec</p>	<p>Factors that affect I/Os per second statistics include access pattern (random or sequential), I/O size, RAID level, segment size, and number of drives in the arrays or Storage Subsystem. The higher the cache hit rate, the higher I/O rates will be.</p> <p>Performance improvements due to a segment size change can be seen in the I/Os per second statistics for a logical drive. Use I/O patterns to determine the optimal segment size. For more information, see Changing Segment Size.</p> <p>Higher write I/O rates are experienced with write caching enabled compared to disabled. In deciding whether to enable write caching for an individual logical drive, consider the current and maximum I/Os per second. You should expect to see higher rates for sequential I/O patterns than for random I/O patterns. Regardless of your I/O pattern, it is recommended that write caching be enabled to maximize I/O rate and shorten application response time. For more information, see Specifying Logical Drive Properties.</p>
--	--

Setting Storage Subsystem Controller Clocks

Use this option to synchronize the Storage Subsystem controller clocks with the Management Station. This option is useful in ensuring that event timestamps written by controllers to the Event Log match event timestamps written to host log files.

Important: Controllers remain available during synchronization.

How to Synchronize Controller Clocks with a Management Station

- 1 To set controller clocks, either select the **Storage Subsystem >> Set Controller Clocks** pull-down menu option, or use the right mouse button to select the Storage Subsystem in the Logical View, then select **Set Controller Clocks** from the pop-up menu.

Result: The **Set Controller Clocks** dialog is displayed.

- 2 Select **OK**.

Result: The controller clocks are synchronized with the Management Station.

Changing the Storage Subsystem Media Scan Settings

Use this option to enable a Storage Subsystem [media scan](#) and specify the duration period (in days) that you want the media scan to run in. You can set the duration within minimum and maximum values supplied by the Storage Subsystem. When enabled, the media scan runs on all logical drives in the Storage Subsystem that are:

- Optimal
- Have no modification operations in progress
- Have the Media Scan parameter enabled on the **Logical Drive Properties** dialog

For more information, see [Specifying Logical Drive Properties](#).

How to Change the Media Scan Settings

- 1 From the Subsystem Management Window, either select the **Storage Subsystem >> Change Media Scan Settings** pull-down menu option or use the right mouse button to select the Storage Subsystem in the Logical View. Then, select **Change Media Scan Settings** from the pop-up menu.

Result: The **Change Media Scan Settings** dialog is displayed.

- 2 Check the **Enable background media scan** checkbox, if it is not already checked.
- 3 Select a number in the spinner box to specify the duration (in days) of the media scan.

4 Select **OK**.

Result: The Media Scan settings for the Storage Subsystem is changed.

The Media Scan duration specifies the number of days over which the media scan should run on the eligible logical drives. For efficiency, the controllers in the Storage Subsystem complete the media scan at a faster rate during times when I/O is not present or very light. Although the media scan may complete more quickly than you have specified, the Subsystem Management software will not restart a complete media scan until the specified number of days has elapsed.

Specifying Storage Subsystem Cache Settings

Use this dialog box to specify:

- When unwritten cache data should be written to disk (flushed)
- When a cache flush should stop
- Cache block size

When the cache holds the specified start percentage of unwritten data, a flush is triggered. When the cache flushes down to the specified stop percentage, the flush is stopped. For example, you can specify that the controller start flushing the cache when the cache reaches 80% full and stop flushing the cache when the cache reaches 16% full.

Note: Both the **Start flushing** and **Stop flushing** parameters in this version of the storage management software are percentages of the total cache memory. The **Stop flushing** parameter is NOT equal to the **Demand Flush Amount** parameter used in previous versions of this storage management software. (The **Demand Flush Amount** parameter was a specified percentage of the **Demand Flush Threshold**; neither of these parameters are currently used.)

How to Change Cache Settings

- 1 Either select the **Storage Subsystem >> Change Cache Settings** pull-down menu option or use the right mouse button to select the Storage Subsystem in the Logical View. Then, select **Change Cache Settings** from the pop-up menu.

Result: The **Change Cache Settings** dialog is displayed.

- 2 Select the percentage of unwritten data in the cache that will trigger a cache flush, in the **Start Value** spinner box.

Note: A percentage that is too low increases the chance that data needed for a host read will not be in the cache. It also increases the number of disk writes necessary to maintain the cache level, increasing system overhead and decreasing performance.

- 3 Select the percentage of unwritten data held in the cache that will stop a cache flush in progress, in the **Stop Value** spinner box.

Note: The lower the percentage that you specify, the higher the chance that data for a host read would require a disk read rather than being present in the cache.

- 4 Select a cache block size that is appropriate for your applications from the **Cache block size:** drop-down list. Available choices include 4 KB (a good choice for file system or database application use) and 16 KB (a good choice for applications that generate sequential I/O, such as multimedia).
- 5 Select **OK**.

Result: The Storage Subsystem controller cache settings are updated.

Redistributing Arrays

Use this option to move arrays (and their associated logical drives) back to their [preferred owners](#). Typically, multi-path drivers move arrays from their preferred controller owner when there is a failure along the connection between the host and the Storage Subsystem.

Important: This option is not available if all arrays are currently owned by their preferred controllers, or if no logical drives exist on the Storage Subsystem.



Caution: If you use this option without a multi-path driver installed on the hosts, stop I/O to the logical drives while this operation is in progress to prevent application errors. You may also need to make operating system modifications to recognize the new I/O path to the logical drives.

How to Redistribute Arrays

- 1 From the Subsystem Management Window, either select the **Storage Subsystem >> Redistribute Arrays** pull-down menu option or use the right mouse button to select the Storage Subsystem in the Logical View. Then, select **Redistribute Arrays** from the pop-up menu.

Result: The **Redistribute Arrays** confirmation dialog is displayed.

- 2 Select **Yes**.

Result: All Storage Subsystem arrays are redistributed back to their preferred owners. Selecting an array in the Logical View displays an associated controller in the Physical View that is both the current owner and the preferred owner.




Important: Under certain application host operating system environments, a reconfiguration of the host driver might be an additional requirement.

Recovering from Storage Subsystem Problems


Use the Recovery Guru to help pinpoint the cause of a problem with the Storage Subsystem and to obtain instructions on how to correct the problem. Where necessary, use the hardware documentation in conjunction with the recovery steps to replace failed components.

Problem Notification

Indicators of a Storage Subsystem problem include:

- A Needs Attention  Storage Subsystem icon displayed in:
 - The Overall Health Status area, Device Tree View, or Device Table of the Enterprise Management Window
 - The Subsystem Management Window Logical View
- The Recovery Guru Optimal toolbar button  in the Subsystem Management Window changes to a flashing Recovery Guru Needs Attention toolbar button .
- Non-optimal component icons are displayed in the Subsystem Management Window Logical and Physical View
- The receipt of critical SNMP or e-mail error messages
- The display of hardware fault lights

How to Recover from Storage Subsystem Problems

- 1 From the Subsystem Management Window, do one of the following:
 - Select the **Recovery Guru** toolbar button .
 - Select the **Storage Subsystem >> Recovery Guru** pull-down menu option.
 - Use the right mouse button to select the Storage Subsystem in the Logical View. Then, select **Recovery Guru** from the pop-up menu.

Result: The Recovery Guru window is displayed. The Recovery Guru window is divided into three Views, Summary, Detail, and Recovery Steps, described in **Recovery Guru Views** in this help topic.

- 2 Select the first problem shown in the Summary area, then follow the Recovery Steps to correct it. Repeat for each listed problem.


Result: If all problems have been corrected, the Storage Subsystem icon will eventually transition from Needs Attention to Optimal. There are certain problems where a Fixing icon will display while an operation such as reconstruction is in progress.

- 3 Select **Recheck** to verify the success of the completed recovery procedure.

Recovery Guru Views

Summary View

The Summary area, shown below, presents a list of Storage Subsystem problems.



Multiple problems within a Storage Subsystem can be related. And, in this case,

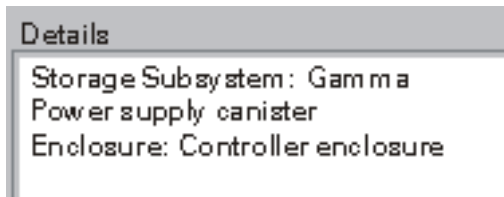
the order in which the problems are corrected can affect the outcome. Select and correct the problems in the order that they are listed in the Summary area.

Note: Multiple power supply canister failures are grouped and listed as one problem in the Summary area. Multiple fan canister failures are also listed as one problem.

Details View

The Details View, shown below, displays information about the selected problem in the Summary area.

Note: You can use standard keyboard cut, copy, and paste commands on text in this area.




Problem Type	Information Shown
All	<ul style="list-style-type: none"> ● Storage Subsystem name
Controller	<ul style="list-style-type: none"> ● Slot identification (where applicable, the appropriate controller slot (A or B) is listed) ● Channel ID, for failed drive SCSI channels ● Serial number of controller
Logical Drive	<ul style="list-style-type: none"> ● Logical Drive group sequence number ● Drive Enclosure Identification ● Affected drive slots (if applicable) ● Logical Drive names ● RAID level ● Status
Drive	<ul style="list-style-type: none"> ● Drive Enclosure Identification ● Affected drive slots, listing slot numbers
Hardware Components	<ul style="list-style-type: none"> ● Component type ● Drive Enclosure Identification

Recovery Steps View

The Recovery Steps View lists the appropriate steps to follow for the selected problem in the Summary area.

Note: You can use standard keyboard cut, copy, and paste commands on text in this area.

There are four headings or statements that may appear in a recovery procedure, described in the table below.

Heading	Description
What Caused the Problem?	Summarizes the likely causes of the problems
 Caution	Lists instructions or information that can prevent data loss, service interruption, or component problems
Important Notes	Lists important information about the recovery process that should be read before proceeding with the recovery steps
Recovery Steps	Lists appropriate steps for a selected problem

Saving Recovery Information to a Text File

- 1 Select **Save As**.

Result: The **Save** dialog is displayed.

- 2 Select the subdirectory location to save the file to, type the name of the file in the appropriate text box, and select **OK**.

Result: The **Save** dialog is closed and the file is saved in HTML format. The file contains all of the information in the Details and Recovery Procedure areas.

Enabling Premium Features

Use this option to enable a premium feature on the Storage Subsystem. To complete this procedure, you must obtain a Feature key file from your storage supplier specific to the premium feature you want to enable.

- 1 Select the **Storage Subsystem >> Premium Features >> Enable** option.

Result: The **Select Feature Key File** dialog is displayed.

- 2 Select the folder in which you placed the generated key file.

Note: The **Select Feature Key File** dialog can filter on files with a `.key` extension.

- 3 Select the appropriate key file, and then select **OK**.

Result: The **Confirm Enable Premium Feature** dialog is displayed.

- 4 Select **Yes**.

Result: The premium feature is enabled.

Note: To check the status of a premium feature, select the **Storage Subsystem >> Premium Features >> List** option. See [Viewing Premium Features](#).

Disabling Premium Features

Use this option to disable a premium feature on the Storage Subsystem.

- 1 Select the **Storage Subsystem >> Premium Features >> Disable** option.

Result: The **Disable Premium Feature** dialog appears, listing any enabled premium features.

- 2 Select one item in the list, and then select **OK**.

Result: The **Confirm Disable Premium Feature** dialog appears.

- 3 Select **Yes**.

Result: The **Working** dialog is displayed while the feature is being disabled. When the feature has been disabled, the **Working** dialog closes.

Notes:

- To check the status of a premium feature, select the **Storage Subsystem >> Premium Features >> List** option. See [Viewing Premium Features](#).
- If you wish to enable the premium feature in the future, you must re-apply the [Feature Key file](#) for that feature.

Viewing Premium Features

Use this option to view a list of premium features on the Storage Subsystem.

- 1 Select the **Storage Subsystem >> Premium Features >> List** option.

Result: The **Premium Feature** dialog is displayed. The dialog lists the following items:

- Premium features enabled on the Storage Subsystem
- [Feature Enable Identifier](#)

- 2 Select **Close** to close the dialog.

Notes:

- To enable a premium feature, see [Enabling Premium Features](#).
- To disable a premium feature, see [Disabling Premium Features](#).
- If you receive a "Premium Features - Out of Compliance" error message during a management session, use the Recovery Guru to resolve the problem.

Renaming a Storage Subsystem

Use this option to set or change the name of a Storage Subsystem.

Important: The storage management software does not check for duplicate names. Check the Enterprise Management Window to verify that the name you have chosen is not already in use by another Storage Subsystem.

Tips on Naming Storage Subsystems

- There is a 30 character limit. All leading and trailing spaces will be deleted from the name.
- Use a unique, meaningful name that is easy to understand and remember.
- Avoid arbitrary names or names that would quickly lose their meaning in the future.
- When you have named the Storage Subsystem, the prefix "Storage Subsystem" is automatically added to the name shown in the Logical View and in the Enterprise Management Window. For example, if you renamed the Storage Subsystem "Engineering," it will display "Storage Subsystem Engineering." You do not need to add the prefix "Storage Subsystem" to the name you specify in this option.

How to Rename a Storage Subsystem

- 1 Use the [Locate Storage Subsystem](#) option to identify the Storage Subsystem you want to rename.
- 2 From the Subsystem Management Window, either select the **Storage Subsystem >> Rename** pull-down menu option or use the right mouse button to select the Storage Subsystem in the Logical View. Then, select **Rename** from the pop-up menu.

Result: The **Rename Storage Subsystem** dialog is displayed. The dialog presents the current name of the Storage Subsystem. In some cases, this is the default name assigned by the subsystem management software.

- 3 Type in a new name.

4 Select **OK**.

Result: The Storage Subsystem is renamed. The name shown in the Logical View and in the Enterprise Management Window will have the prefix "Storage Subsystem" followed by the name. For example, if you renamed the Storage Subsystem "Engineering," it will display "Storage Subsystem Engineering."

5 Prepare a label with the new Storage Subsystem name and place it on the Storage Subsystem.

Automatic Storage Subsystem Configuration

Use this option to automatically configure your Storage Subsystem using all of the remaining unconfigured Storage Subsystem capacity and any free capacity in existing arrays. The result of this option will be a controller-defined set of logical drives, arrays, and hot spare drives. The resulting configuration is normally well suited to the needs of many applications.

Important: Do not use this option if:

- You want to optimize the configuration for a certain application
- You want logical drives with different RAID levels, cache settings, or other parameters
- You have specific hot spare requirements
- You want to keep some free space available for future use

How to Automatically Configure your Storage Subsystem

- 1 From the Subsystem Management Window, select the **Configure >> Automatic Configuration** pull-down menu option.

Result: The **Confirm Automatic Configuration** dialog is displayed.

- 2 Select **Yes**.

Result:

- A controller-defined configuration is applied to the Storage Subsystem.
 - From the initial set of unassigned drives, one is configured as a hot spare
 - From the remaining set of unassigned drives, the maximum number of 5+1 RAID 5 logical drives is configured.
 - If between 1 and 5 drives remain, then they are configured as follows:

Number of remaining drives	Resulting configuration
Between 3 and 5	An N+1 RAID 5 logical drive is created.
2	A 1+1 RAID 1 logical drive is created.
1	A hot spare is created

- The Logical View and the Physical View are updated to reflect the new configuration.

Creating a Logical Drive


A [logical drive](#) is the logical entity you create that will store your data. Use this option to create a single logical drive from either [unconfigured capacity](#) or [free capacity](#). With the advent of larger capacity drives and the ability to share logical drives across multiple hosts, you might want to create more than one logical drive per array to better organize your data. Use the **Create Logical Drive** option for each new logical drive you want to create.

Important:

- For certain operating systems, you can only create a single logical drive on the array and it must use all of the available capacity. The Subsystem Management software enforces these restrictions when they apply to your operating system.
- There can be Host operating system-specific limits on the number of logical drives that the Host can access. Keep this in mind when creating logical drives for use by a particular Host.
- After creating one or more logical drives, run the host-based **SM7devices** utility to correlate logical drive names with host device names. For information on using this utility, see the *Installation and Support Guide*.

Accessing the Option

To create a logical drive, select either the Unconfigured Capacity or Free Capacity node from the Logical View; then do one of the following:

- Select the **Create Logical Drive** button from the toolbar .
- Select **Configure >> Create Logical Drive** from the pull-down menu option.
- Select **Create Logical Drive** from the right-mouse pop-up menu.

Creating a Logical Drive from Unconfigured Capacity

The first step in creating an individual logical drive from unconfigured capacity is to define an array. To define the array, you specify two parameters: RAID level and capacity. The RAID level parameter indicates how the data will be striped across the drives in the array and its level of data protection. The capacity parameter indicates how large you want the array to be (that is, how many drives you want to be included in the array). You use the **Select Array** dialog to specify

these parameters. For more information, see [Selecting an Array](#).

The second step is to specify the parameters for the individual logical drive you plan to create on the array. Use the **Specify Logical Drive Parameters** dialog to specify these parameters. For more information, see [Specifying Logical Drive Parameters](#).

Creating a Logical Drive from Free Capacity

Because the free capacity only resides on an existing array, the overall capacity (that is, the number of drives) and RAID level have already been specified.

Therefore, the **Select Array** dialog is by-passed when you create a logical drive from free capacity. When you select a Free Capacity node, the only step that is required is to specify the parameters for the individual logical drive you plan to create on the array. Use the **Specify Logical Drive Parameters** dialog to specify those parameters. For more information, see [Specifying Logical Drive Parameters](#).

Selecting an Array

The first step in creating an individual logical drive from [unconfigured capacity](#) is to define an array. To define the array, you specify two parameters: RAID level and capacity.

Important: This dialog is available only when you are creating logical drives from the Unconfigured Capacity node. If you are creating a logical drive on an existing array using a Free Capacity node, this dialog is by-passed because the array (and its associated drives and RAID level) have already been specified. For a Free Capacity node, the only dialog that is displayed is the **Specify Logical Drive Parameters** dialog.

- 1 Specify the RAID level that will meet your data storage and protection requirements.

RAID level choices include [RAID 0](#), [RAID 1](#), [RAID 3](#), and [RAID 5](#).

- 2 Select a method for defining the available drives to be used in the array.

Note: By default, the automatic method is selected.

Method



Automatic

Description and Action

Description

Allows you to select from a list of array capacity choices as selected by the controller firmware.

A table displays array capacities, the number of drives that comprise the array, and whether or not the array has [channel protection](#).

Note: Your choice of RAID level affects the capacities of the listed arrays.

Action

1. Select one array and then select **Next** to go to the **Specify Logical Drive Parameters** dialog.



Manual

Description

Important: This option should only be used by experts who understand the concepts of drive redundancy and optimal drive configurations.

Allows you to select which specific drives comprise the array. A table displays drive location information and individual drive capacities.

Note: Your choice of RAID level affects the resulting capacity of the array.

Action

- 1 Select the number of drives you want to include in the array and then select **Apply**.

Result: The area next to the **Apply** button displays the corresponding array capacity and whether the array has [channel protection](#).

Important: If you do not select the appropriate number of drives for the chosen RAID level, an error message is displayed. Change your drive selections and select **Apply** again.

- 2 Select **Next** to go to the **Specify Logical Drive Parameters** dialog.



Specifying Logical Drive Parameters

Use the **Create Logical Drive - Specify Logical Drive Parameters** dialog to specify logical drive name, usage, desired capacity, controller ownership, and storage partition mapping preference.


Important: If you select a Free Capacity node when invoking the Create Logical Drive option and there are no valid logical drive candidates, an error message is displayed and you are returned to the Management Window.

Logical Drive Parameters Summary Area

The information displayed in the Summary area varies according to the choices you made previously, during the logical drive creation process.

If you selected...	The Summary Area displays...
<p>The Unconfigured Capacity node in the Logical View</p> <p>-AND-</p> <p>the Automatic Method in the Select Array dialog</p>	 <ul style="list-style-type: none">● Automatic method icon● RAID level you selected for the array● Maximum capacity (taking into account the RAID level)
<p>The Unconfigured Capacity node in the Logical View</p> <p>-AND-</p> <p>the Manual method in the Select Array dialog</p>	 <ul style="list-style-type: none">● Manual method icon● RAID level you selected for the array● Maximum capacity (taking into account the RAID level)

A Free Capacity node
in the Logical View



- Free Capacity icon
- RAID level of the selected Free Capacity Node
- Maximum capacity (taking into account the RAID level)

How to Specify Logical Drive Parameters

- 1 Specify the capacity units and logical drive capacity.

Use the spinner box to indicate capacity units (in MB or GB) and list box to indicate the specific numeric capacity. The maximum value, shown by default, is the Maximum capacity shown in the Summary Area.

Important:

- You can type up to three decimal places for the capacity.
- If you enter a capacity larger than the maximum, the value will default to the maximum capacity when you move to another area on the dialog or when you select **Finish**.
- If you enter a capacity that is less than the maximum, make sure you leave enough remaining capacity to allow you to create one or more additional logical drives that meet your application needs.

- 2 Specify the logical drive name.

Choose a unique, meaningful name that helps you identify the logical drive. Generally, a meaningful name describes the data that will be stored on the logical drive.

- There is a 30 character maximum
- If you choose the name of a logical drive that duplicates the name of another logical drive in the Storage Subsystem, an error message is displayed and you are prompted to choose another name
- Any leading and trailing spaces in the name are deleted

- 3 Specify the expected usage of the logical drive based on your application needs.

The expected logical drive usage that you specify is used by the controller firmware to indicate an appropriate default logical drive [segment size](#) and [cache read-ahead multiplier](#). The choices available for expected logical drive usage include file system (the default), database, and multimedia.

- 4 Select an appropriate [segment size](#) for your application. You can choose the default (based on expected usage), 8, 16, 32, 64, 128, or 256 K.

Note: You can change the logical drive segment size later using the **Change Segment Size** option. For more information, see [Changing Segment Size](#).

- 5 Specify controller ownership (a preferred controller).

Note: This option is only available if you are using dual active controllers. Choices include either the controller in slot A or the controller in slot B.

If you are creating this logical drive from:	Then...
An Unconfigured Capacity node	<p>The Storage Subsystem controller firmware automatically selects a preferred owner.</p> <p>You can change the preferred owner by selecting the alternate controller radio button.</p>
A Free Capacity node	<p>The preferred owner is set automatically by determining the owner of the other logical drives in the array.</p> <p>Important: You cannot change controller ownership because the same controller MUST own all logical drives in the same array.</p>

- 6 Specify whether the logical drive should have a default mapping or be left unmapped. For more information, see [Introducing Storage Partitions](#).

Mapping	Description
Default Mapping	<p>The logical drive will be available to all host groups or hosts that don't have specific logical drive-to-LUN mappings. (All hosts or host groups in the Default Host Group can access the logical drive.)</p> <p>You can change the default mapping at a later time using the Storage Partitions Definitions Window. For more information, see Changing a Logical Drive-to-LUN Mapping.</p>
No Mapping (map later)	<p>The logical drive cannot be accessed until you make a specific logical drive-to-LUN mapping using the Storage Partitions Definitions window. For more information, see Defining Logical Drive-to-LUN Mappings.</p> <p>Important: This option is not available if you do not have the Storage Partitions premium feature enabled.</p>

- 7 Select **Finish**.

Result: If the logical drive is created successfully, the **Logical Drive Successful** dialog is displayed. If there are any errors, an error message appears.

- 8 Select **OK**.

Result: A logical drive (or array/logical drive) is created with the specified RAID level and parameters. You are returned to the Subsystem Management Window. The Logical View is updated to reflect the addition of the new array and/or logical drive, and the subtraction of capacity from either the Free Capacity or the Unconfigured Capacity node.

Note: To change other logical drive parameters, see [Specifying Logical Drive Properties](#).

- 9 Perform any operating system modifications necessary on the application host so that the applications can use the logical drive. See the *Installation and Support Guide* specific to your operating system.

Introducing Storage Partitions

A Storage Partition is a logical entity consisting of one or more Storage Subsystem logical drives that can be shared among hosts that are part of a host group or accessed by a single host. A Storage Partition is created when you define a collection of hosts (a host group) or a single host and then define a logical drive-to-logical unit number (LUN) mapping. This mapping allows you to define what host group or host will have access to a particular logical drive in your Storage Subsystem.

Storage Partitions are a premium feature of the storage management software and therefore must be enabled either by you or your storage vendor supplier. There is a maximum number of partitions that can be created on the Storage Subsystem depending on the premium feature that has been enabled. Benefits of partitioning include ease of management, amortization of costs, scalability, and flexibility. For a tutorial explaining how to set up Storage Partitions, see [Setting up Storage Partitions, Step-by-Step](#).

There are three main steps to Storage Partitions configuration.

- 1 Create logical drives on your Storage Subsystem. During logical drive creation, you can specify one of two [logical drive-to-LUN mapping](#) settings:

Default or None (map later).

- Default - this setting specifies that a LUN should be automatically assigned to the logical drive. This setting will allow host groups or hosts that do not have specific logical drive-to-LUN mappings (designated by the [default host group](#) node in the Topology View) to have access to the logical drive. If you are not going to be using Storage Partitions, this is the setting you would want to specify.

-OR-

- None (map later) - this setting specifies that a LUN will not be assigned to the logical drive. This setting allows you to define a specific logical drive-to-LUN mapping and create Storage Partitions using the **Configure>>Topology/Storage Partitions** option. If you are going to be using Storage Partitions, this is the setting you would want to specify.

For more information, see [Creating a Logical Drive](#).

- 2 Define the host groups and hosts (and their associated host ports) for which you want to define specific logical drive-to-LUN mappings (your Topology). For more information, see [Defining Storage Partitions Topology](#), below.

- 3 Grant access to one or more logical drives to a defined host group or host. You do this by defining a specific logical drive-to-LUN mapping using the following general steps: (1) Select a defined host group or host in the Topology View, (2) Select the **Define New Mappings** option, (3) Select a logical drive that you want the host group or host to be able to access, and (4) Select the logical unit number (LUN) you want the host group or host to use to access the logical drive.

Each host group or host is granted a unique view of partitioned storage. A defined host group or host can either see:

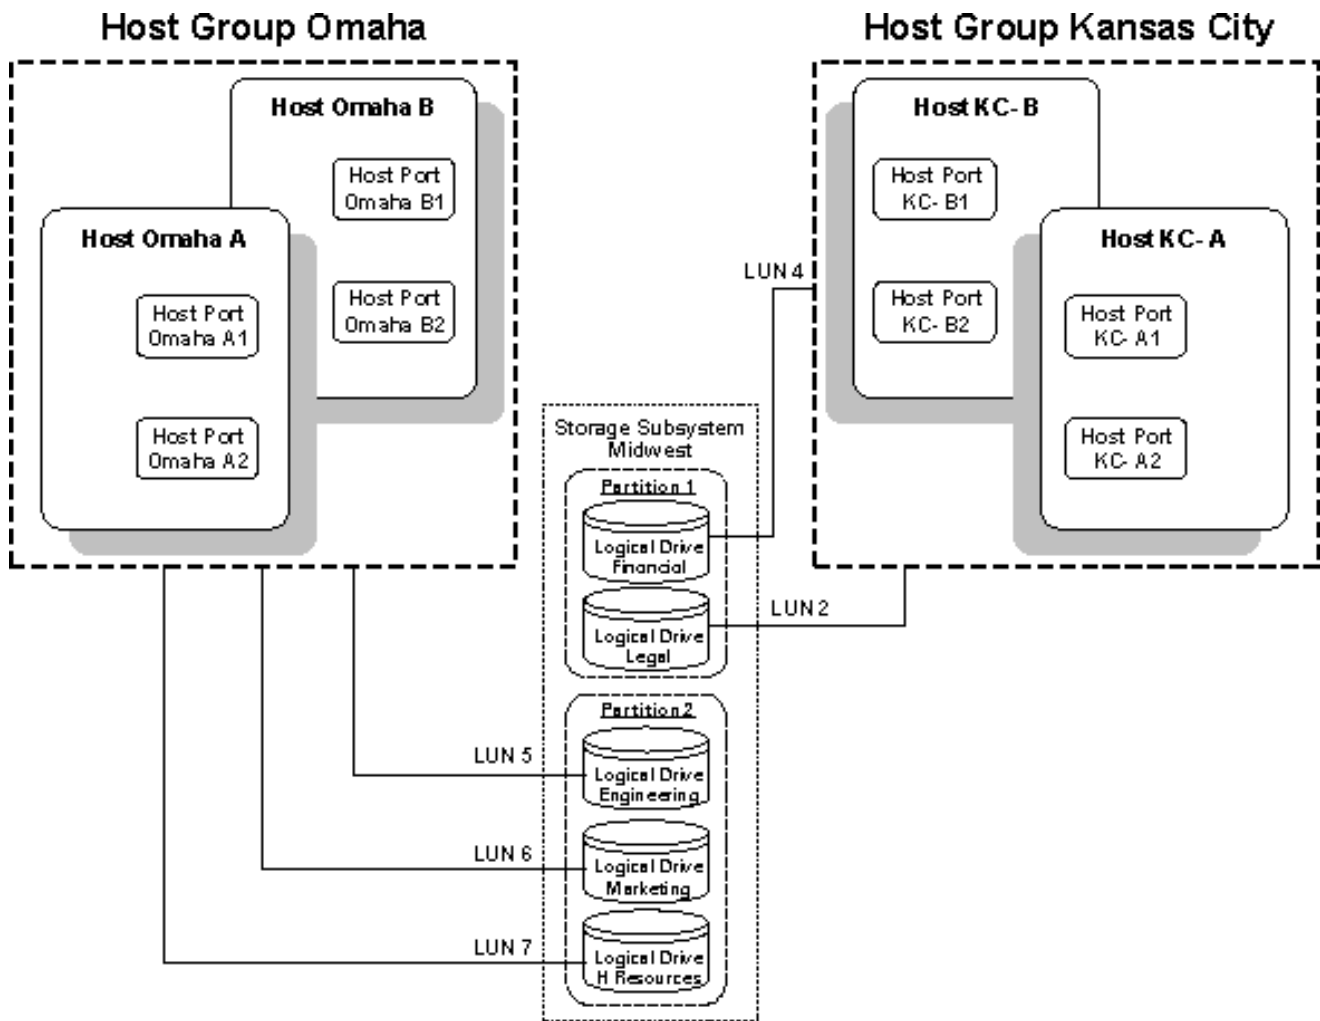
- logical drives with default logical drive-to-LUN mappings. In this case, the host group or host will be part of the [Default Host Group](#).

-OR-

- logical drives to which it has been granted access through a specific logical drive-to-LUN mapping

For more information, see [Defining Logical Drive-to-LUN Mappings](#), below.

The following example shows that two host groups have been defined (host group Omaha and Kansas City) that have access to various logical drives. In this example, two partitions have been used. The first partition consists of logical drives Financial and Legal and it is accessed by the hosts in host group Kansas City using LUNs 4 and 2, respectively. The second partition consists of logical drives Engineering, Marketing, and H Resources and it is accessed by the hosts in host group Omaha using LUNs 5, 6, and 7, respectively.



Storage Partition Topology

A Storage Partition Topology is a collection of topological elements (default host group, host group, host, and host port) shown as nodes in the Topology View of the [Storage Partitions Definitions Window](#). You must define the various topological elements if you want to define specific logical drive-to-LUN mappings for host groups and/or hosts.

Note: Even if you are not planning to use Storage Partitions, you can still define your hosts and host ports that are connected to the particular Storage Subsystem and view the default logical drive-to-LUN mappings to determine which logical unit numbers are being used to access the logical drives in your Storage Subsystem.

Defining Storage Partitions Topology

Default Host Group

A default host group is a standard node in the Topology View that designates all host groups, hosts, and host ports that: (1) do not have any specific logical drive-to-LUN mappings, and (2) share access to any logical drives that were automatically assigned default LUNs by the controller firmware during logical drive creation.

Host Group

A host group is an optional topological element that you define if you want to designate a collection of hosts that will share access to the same logical drives. For more information, see [Defining a Host Group](#).

Host

A host is a computer that is attached to the Storage Subsystem and accesses various logical drives on the Storage Subsystem through its host ports. You can define a host as a separate entity or as part of a host group. You can define specific logical drive-to-LUN mappings to an individual host as well as have the host be part of a host group that shares access to one or more logical drives. For more information, see [Defining a Host](#).

Host Port

A host port is a physical connection on the host adapter that resides within a host. When the host adapter only has one physical connection (host port), the terms host port and host adapter are synonymous.

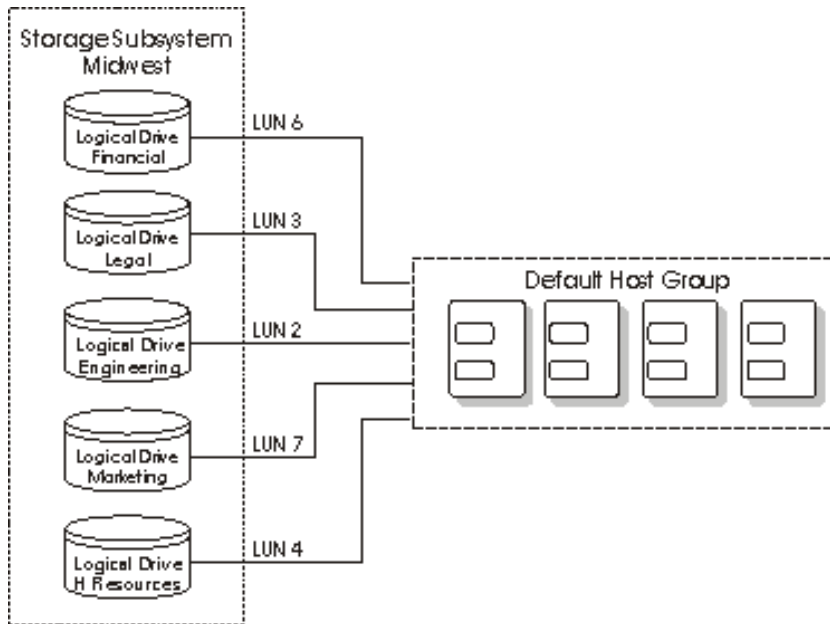
The host ports are automatically discovered by the storage management software. A host port is the actual physical connection that allows a host to gain access to the logical drives in the Storage Subsystem. Therefore, if you want to define specific logical drive-to-LUN mappings for a particular host and create partitions, you must define its associated host ports.

Initially, all discovered host ports belong to the default host group and have access to any logical drives that were automatically assigned default LUNs by the controller firmware during logical drive creation. For more information, see [Defining a Host Port](#).

Any host port in the default host group can automatically access the following logical drives:

- Logical Drives that when created were given default logical drive-to-LUN mappings. For more information, see [Specifying Logical Drive Parameters](#).
- Any logical drives created using previous versions of this storage management software. Such logical drives are automatically given default logical drive-to-LUN mappings. If desired, you can change these mappings using the **Change Mapping** option.

Example: In the following example, the logical drives in the Storage Subsystem Midwest have not been given specific logical drive-to-LUN mappings, therefore any hosts and their associated host ports in the default host group can request data from these logical drives. The LUNs are automatically assigned by the controller firmware.



Reconfiguring Storage Partitions Topology

A Storage Partitions topology is reconfigurable. You can:

[Move a host port](#)

[Replace a host port](#)

[Move a host from one host group into another host group](#)

[Delete a host group, host, or host port](#)

[Rename a host group, host, or host port](#)

Defining Logical Drive-to-LUN Mappings

After you have defined your Storage Partitions topology, you can select a defined host group or host in the Topology View and select the **Define New Mapping** option to assign a specific logical drive and logical unit number (logical drive-to-LUN mapping). This designates that only the selected host group or host has access to the particular logical drive through the assigned LUN. For more information on the **Define New Mapping** option, see [Defining a Logical Drive-to-LUN Mapping](#).

The storage management software manages the logical drive-to-LUN mappings. Host port requests for data issued for specific LUNs are routed to the appropriate logical drive by the controller firmware.

Important:

- Only one logical drive-to-LUN mapping is allowed per logical drive. To grant logical drive access to more than one host, you must first group the hosts together in a logical host group. Then, you can grant logical drive access to that host group; thereby allowing the hosts in the host group to share access to the logical drive.
- Because each host has its own LUN address space, you can use the same LUN in more than one logical drive-to-LUN mapping, as long as that LUN is available for use by each particular host.

participating in the mapping.

- You can only define new logical drive-to-LUN mappings for logical drives that currently have no mappings (this means, they do not have default mappings OR specific mappings).
- If any of the host groups or hosts have specific logical drive-to-LUN mappings defined, then they will no longer appear under the default host group; they will move under the root, Storage Subsystem, node. Also, they will no longer be able to access logical drives with default logical drive-to-LUN mappings. A host group or host can either access logical drives with default mappings or specific mappings, but not both.
- Logical Drive-to-LUN mappings are shared between controllers in the Storage Subsystem.

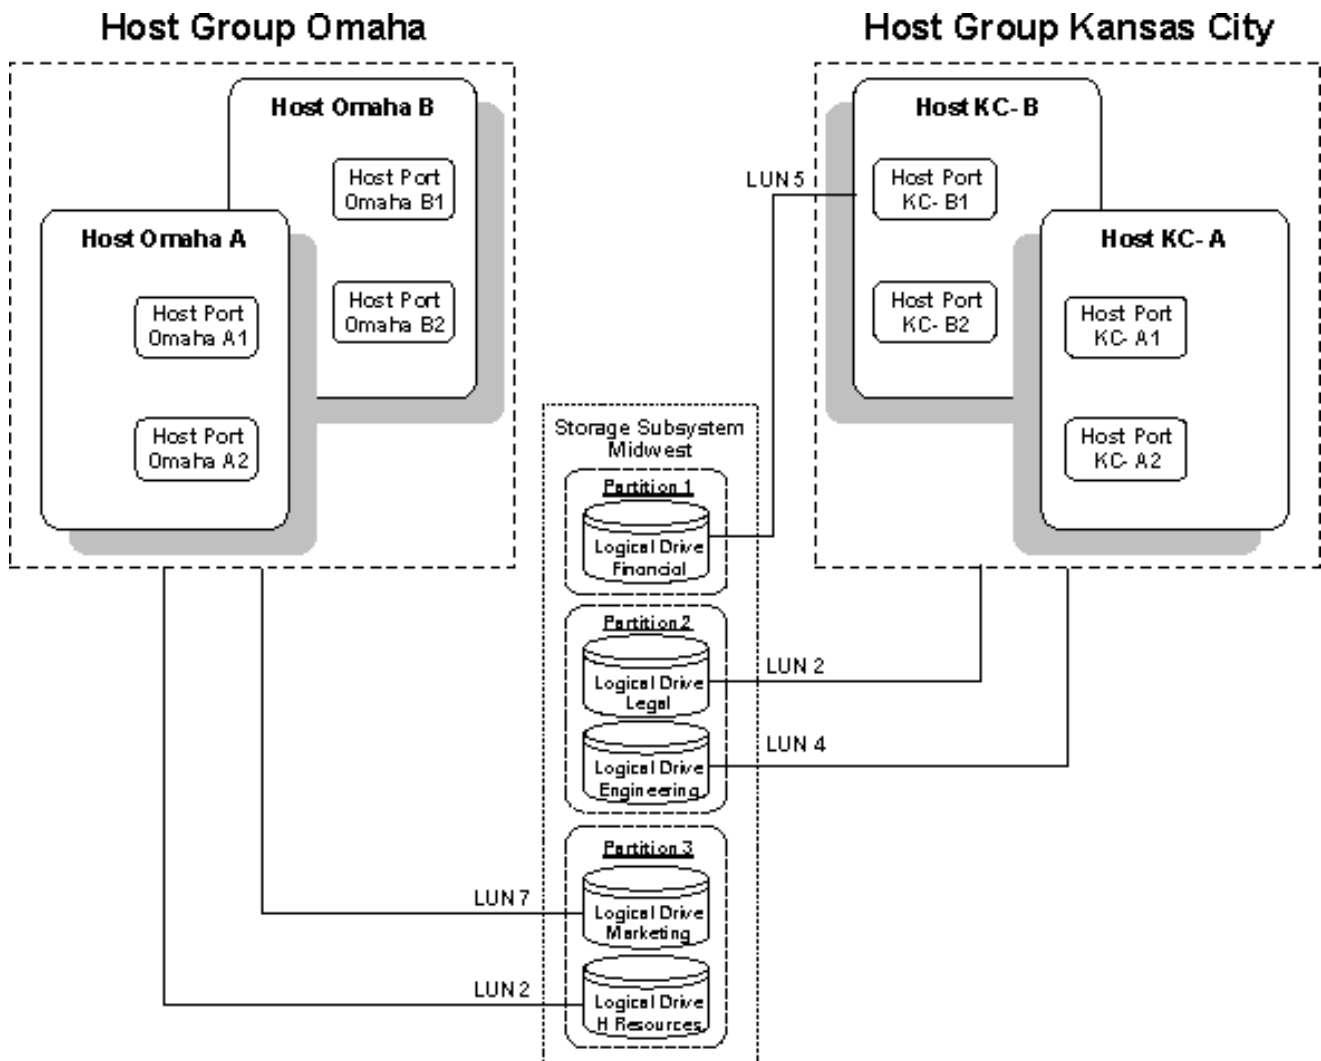
Granting Logical Drive Access to Host Groups

In the following example, three partitions have been configured.

The first partition consists of logical drive Financial. This logical drive is accessed by host KC-B using LUN 5. Even though host KC-B is part of the logical host group Kansas City, host KC-A cannot access this logical drive because the logical drive-to-LUN mapping was created with host KC-B rather than the host group Kansas City.

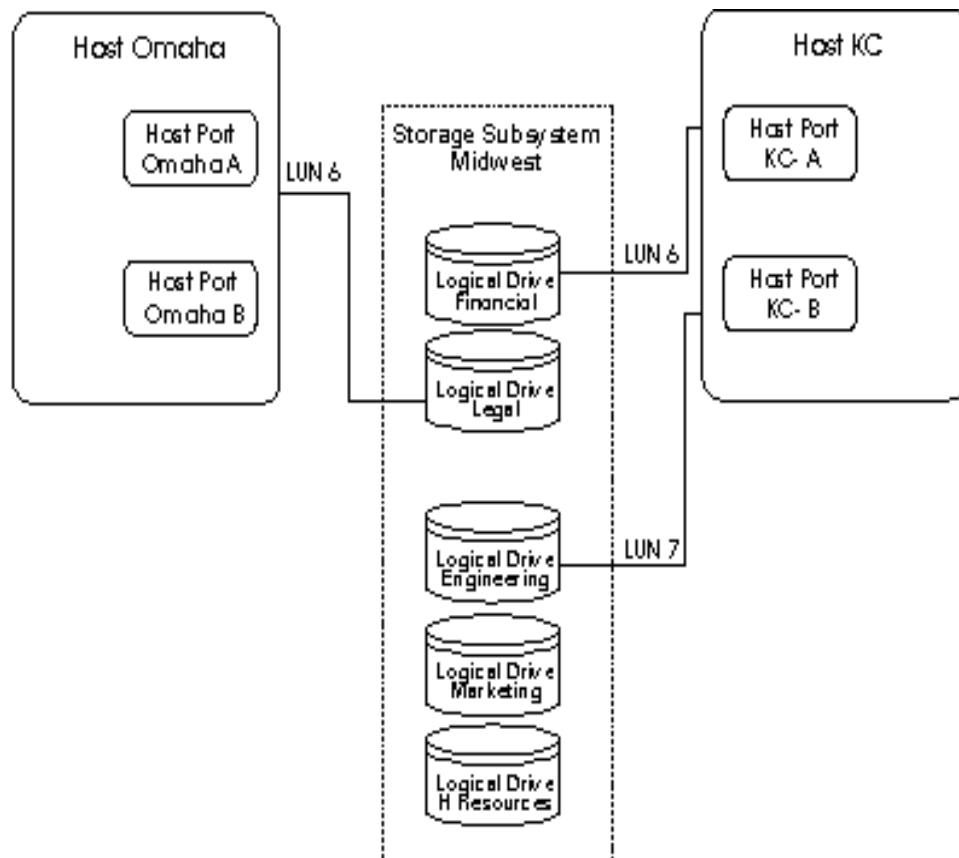
The second partition consists of logical drives Legal and Engineering. These logical drives are accessed by hosts KC-A and KC-B in host group Kansas City using LUNs 2 and 4, respectively.

The third partition consists of logical drives Marketing and H Resources. These logical drives are accessed by hosts Omaha A and B in host group Omaha using LUNs 7 and 2, respectively.



Granting Logical Drive Access to Hosts

In the following example, host Omaha has been granted access to logical drive Legal through LUN 6 through either of its host ports. Host KC has been granted access to logical drive Financial through LUN 6 and logical drive Engineering through LUN 7 through either of its host ports. Note that the same LUN (LUN 6 in this example) can be used for logical drive access for both hosts (Omaha and KC), because each host has its own LUN address space.



Reconfiguring your Logical Drive-to-LUN Mappings

A logical drive-to-LUN mapping is reconfigurable. You can:

- [Change a Logical Drive-to-LUN Mapping](#)
- [Delete a Logical Drive-to-LUN Mapping](#)

Using the Storage Partition Definitions Window

Use the Storage Partition Definitions Window to configure or reconfigure topology definitions and to create or change logical drive-to-LUN mappings.

Note: You must define your host topology if you are going to use storage partitions. However, even if you are not planning to use storage partitions, you can use the topology definition options to define the hosts that are attached to your Storage Subsystem. For more information, see [Introduction to Storage Partitions](#).


The Storage Partition Definitions Window consists of two panes:

- The Topology View (left pane)
- The Mappings View (right pane)

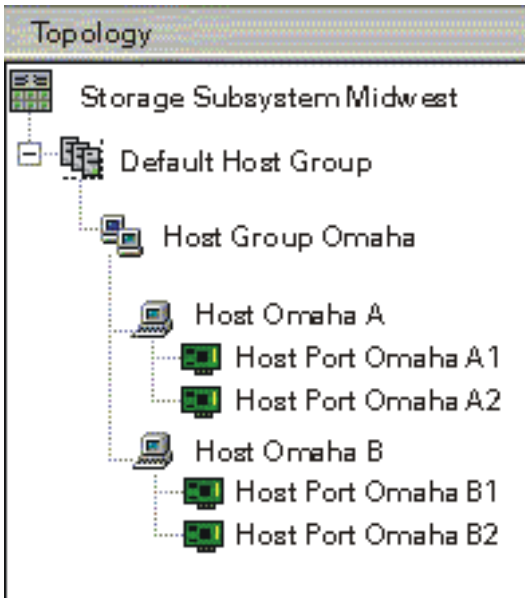
Note: Resize either pane by dragging the splitter bar, located between the two panes, to the right or to the left.

About the Topology View

The Topology View displays the host groups and/or hosts you have defined.

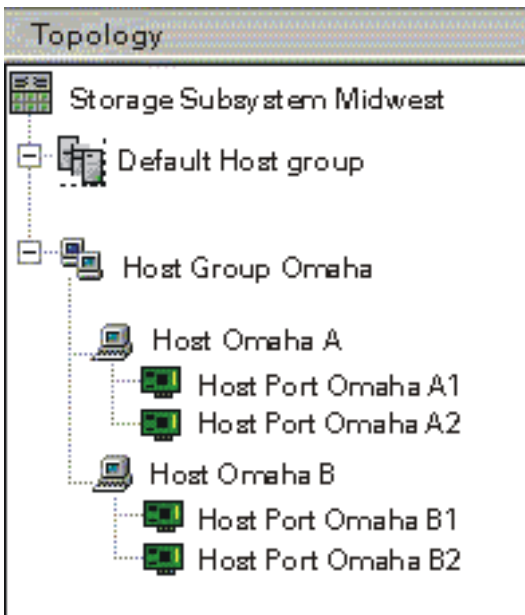


If you have not defined your topology and made specific logical drive-to-LUN mappings, then the Topology View is displayed with the Default Host Group under the Storage Subsystem, as shown in the example to the left.



When you create topology definitions, the defined host groups and hosts initially appear under the Default Host Group.

In the example to the left, the Host Group Omaha and Hosts Omaha A and Omaha B, and Host Ports Omaha A1, Omaha A2, Omaha B1, and Omaha B2 have been defined, but not assigned specific mappings. Therefore, they remain under the Default Host Group.



When you define specific logical drive-to-LUN mappings to host groups and hosts, the host groups and hosts move from under the Default Host Group to under the Storage Subsystem indicating that they have specific logical drive-to-LUN mappings.

In the example to the left, Host Group Omaha and its associated Hosts Omaha A and Omaha B have been given specific logical drive-to-LUN mappings. This action causes the host group and its associated hosts to move out of the Default Host Group. Also, this action causes one storage partition to be used.

Important: After a host group or host has moved out from under the Default Host Group, that host group or host can access ONLY the logical drives that have been mapped to it.

About the Mappings View

The Mappings View displays the default and/or specific logical drive-to-LUN mappings in a table form.

Important: Remember that during logical drive creation, you choose whether to:

- Give the logical drive a default logical drive-to-LUN mapping so that it can be accessed by any discovered hosts attached to the Storage Subsystem (designated by the Default Host Group node).

-OR-

- Leave the logical drive unmapped, so that you can make a specific logical drive-to-LUN mapping and create storage partitions using the options in the Storage Partitions Window. Use this setting if you are going to be using storage partitions.

Mappings View Display Variations

If all logical drives in the Storage Subsystem have been given default logical drive-to-LUN mappings, the mappings are displayed in the Mappings View, as shown in the following example. This is the normal, initial view if you are NOT planning to define storage partitions.

Notes:

- In the Mapping table, the designation of "Default Host Group" under the Scope column indicates default logical drive-to-LUN mappings.
- To display default logical drive-to-LUN mappings, select the Storage Subsystem, the Default Host Group, or any host group or host under the Default Host Group in the Topology View.

Mappings			
Scope	Logical Drive	Unique Logical Drive Identifier	LUN
Default Host Group	Human Resources	1234	0
Default Host Group	Marketing	2345	1
Default Host Group	Engineering	3456	2
Default Host Group	Customer Support	3567	3
Default Host Group	Finance	8769	4
Default Host Group	Accounting	9898	5

If all logical drives in the Storage Subsystem are unmapped, no information is displayed in the Mappings View. The logical drives are available for mapping using the **Define New Mapping** option. This is the normal, initial view if you are planning to define storage partitions.

If some logical drives in the Storage Subsystem have been given default logical drive-to-LUN mappings and other logical drives have been given specific logical drive-to-LUN mappings, both default and specific mappings are displayed in the Mappings View, as shown in the following example. This would normally only occur if you are in the process of defining or changing some default logical drive-to-LUN mappings to specific logical drive-to-LUN mappings for a host group or host. If you are using storage partitions, you will want to map all logical drives to specific logical drive-to-LUN mappings. Therefore, when all mappings have been made, the Default Host Group should not contain any logical drive mappings.

Mappings			
Scope	Logical Drive Name	Unique Logical Drive Identifier	LUN
Default Host Group	Human Resources	1234	0
Default Host Group	Marketing	2345	1
Host Group Omaha	Engineering	3456	2
Host Group Omaha	Customer Support	3567	3
Host Kansas City	Finance	8769	4
Host Kansas City	Accounting	9898	5

Notes:

- In the Mapping table, the designation of "Default Host Group" under the Scope column indicates default logical drive-to-LUN mappings. The designation of "Host Group <host group name>" or "Host <host name>" under the Scope column indicates specific logical drive-to-LUN mappings.

- To display all logical drive-to-LUN mappings, select the Storage Subsystem in the Topology View. To display specific mapping information for a host group or host, select the host group or host in the Topology View.

Using Storage Partition Definitions Pull-down Menus

Menu options **Use this option to...**

File

Exit	Close the Storage Partition Definitions window and return to the Subsystem Management Window.
------	---

Edit

Rename	Rename a host group, host, or host port in your topology. For more information, see Renaming a Host Group, Host, or Host Port .
Delete	Delete a Logical Drive-to-LUN Mapping, or delete a host group, host, or host port topology definition. For more information, see Deleting a Node or a Logical Drive Mapping .

Configure menu options

Topology

Define New Host Group	<p>Define a host group in your topology. Remember that after defining a host group, you must define its associated hosts and host ports.</p> <p>For more information, see Defining a Host Group.</p>
Define New Host	<p>Define a host in your topology. Remember that after defining a host, you must define its associated host port(s).</p> <p>For more information, see Defining a Host.</p>
Define New Host Port	<p>Define a host port in your topology.</p> <p>For more information, see Defining a Host Port.</p>
Move Host	<p>Move a host from one host group to another, or remove a host from a defined host group.</p> <p>For more information, see Moving a Host from One Host Group to Another Host Group.</p>
Move Host Port	<p>Move a host port from one host to another.</p> <p>For more information, see Moving a Host Port.</p>
Replace Host Port	<p>Replace a failed host port.</p> <p>For more information, see Replacing a Host Port.</p>

Logical Drive-to-LUN Mappings

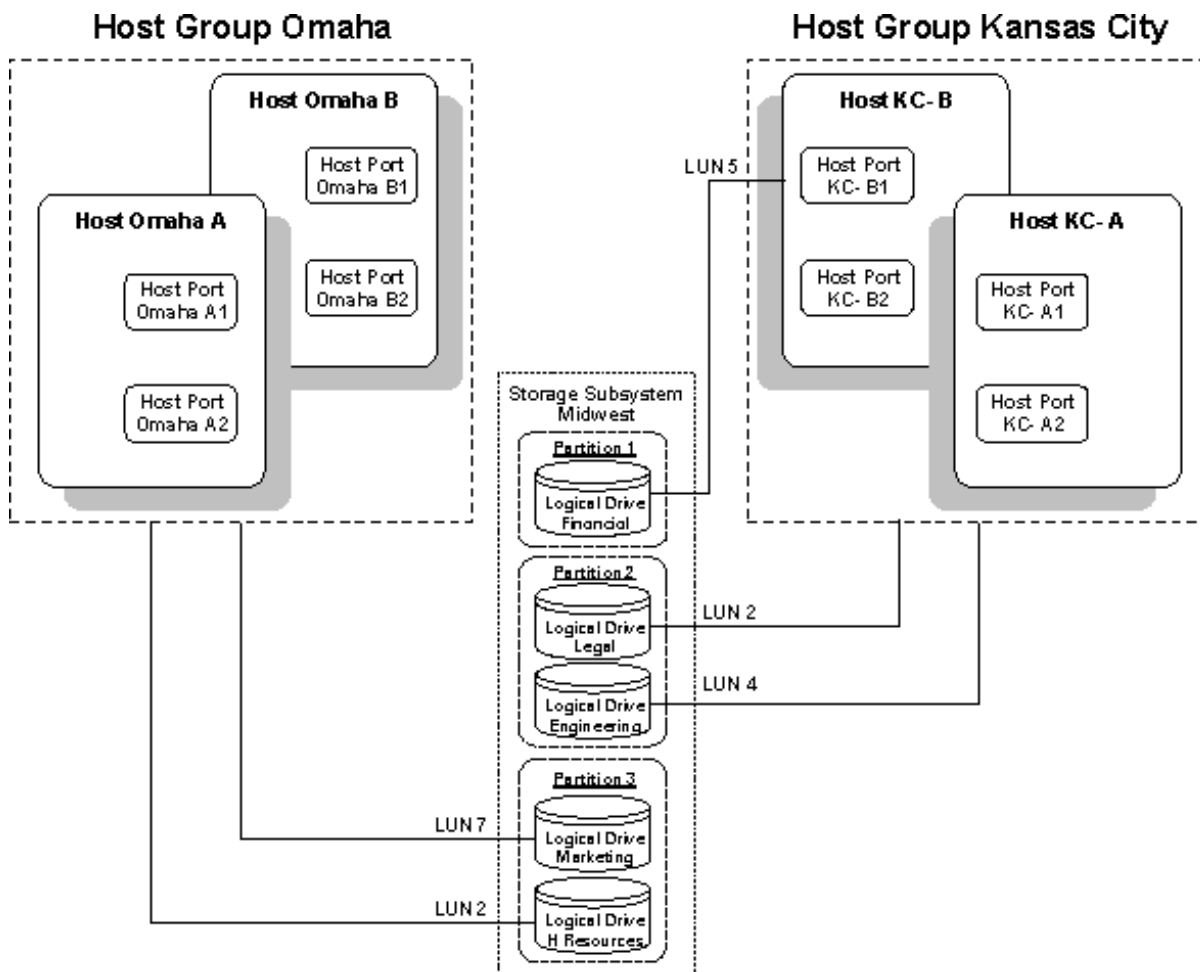
Define New Mapping	<p>Map a logical drive to a LUN and grant access to a defined host group, defined host, or the Default Host Group.</p> <p>For more information, see Defining a Logical Drive-to-LUN Mapping.</p>
Change Mapping	<p>Change the logical drive mapping to a LUN and/or to a defined host group, defined host, or the Default Host Group.</p> <p>For more information, see Changing a Logical Drive-to-LUN Mapping.</p>

Help Menu Options

Invoke the context-sensitive help for the Storage Partition Definitions Window.

Setting up Storage Partitions, Step-by-Step

The following procedures and Online Help topics can be used as a guide to set up Storage Partitions on a Storage Subsystem. Keep in mind that there are other ways to configure a Storage Subsystem; this is presented as one possible scenario. This example explains how to define the topology and logical drive-to-LUN mappings shown here:



Four hosts (Omaha A and B, and KC-A and B) are connected to Storage Subsystem Midwest. The administrator of the Storage Subsystem should make sure that Omaha A and B have access to logical drives Marketing and H Resources and KC-A and KC-B have access to logical drives Legal and Engineering. In addition, the administrator wants host KC-A to have exclusive access to logical drive Financial. To accomplish this, the following steps must be done:

- 1 Create each logical drive (Financial, Legal, Engineering, Marketing, and H Resources) according to the procedures in [Creating a Logical Drive](#).

When entering logical drive-to-LUN mapping settings, make sure "none (map later)" is selected.

- 2 Define the host groups (Omaha and Kansas City) according to [Defining a Host Group](#).

Host groups must be created because multiple hosts must have access to the same logical drives.

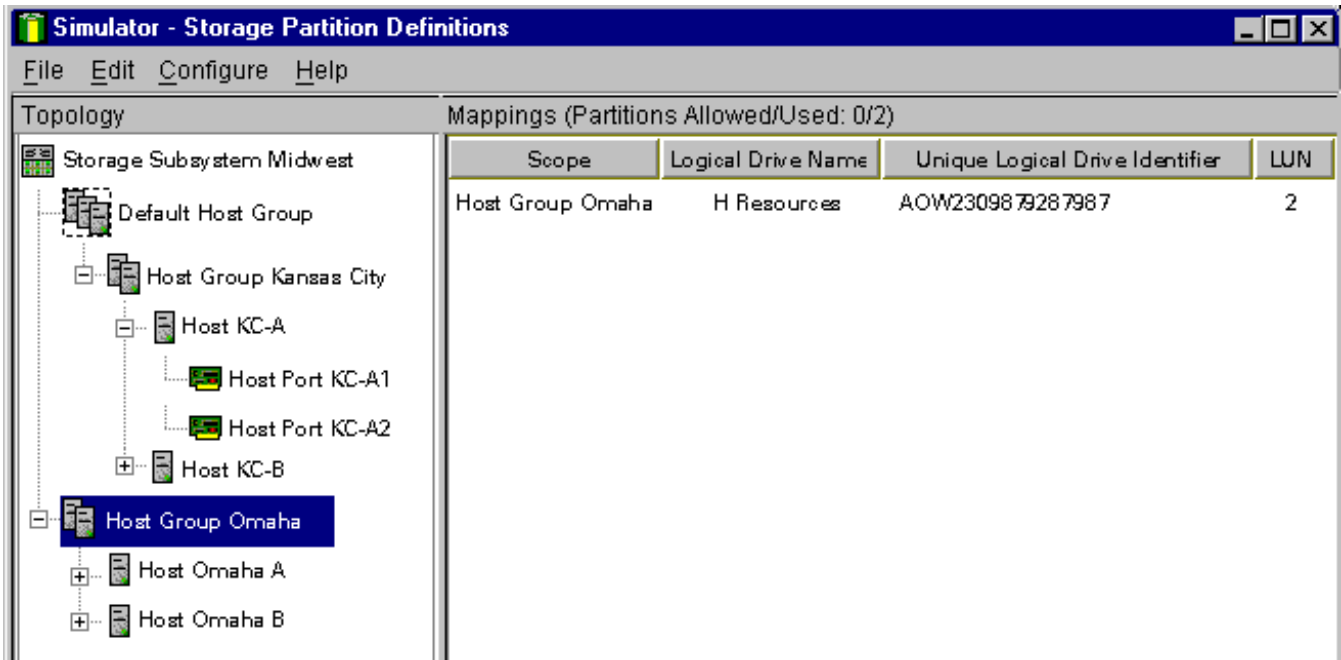
- 3 Define the hosts (Omaha A and B and KC-A and B) under the host groups according to [Defining a Host](#).

- 4 Define the host ports (Omaha A1, A2, B1, and B2 and KC A1, A2, B1, and B2) according to [Defining a Host Port](#).

The host ports are the topological entity that actually allows the host to gain access to the logical drives.

- 5 Assign access for each host group or host according to [Defining a Logical Drive-to-LUN Mapping](#). For example, select host KC-B and define an exclusive logical drive-to-LUN mapping for logical drive Financial (one partition). Then, select host group Kansas City and define logical drive-to-LUN mappings for logical drives Legal and Engineering (a second partition). Finally, select host group Omaha and define logical drive-to-LUN mappings for logical drives Marketing and H Resources (a third partition).

After completing all the steps in the previous procedure, the topology and mappings in the Storage Partition Definitions Window would look like the following:



Defining a Host Group

Use this option to define a new [host group](#).

Important:

- This operation will not succeed if the name you choose duplicates the name of a host group that already exists in the Topology View.
- After defining a host group, you must define its associated hosts and host ports.
- The maximum length of the host group name cannot exceed 30 characters.

How to Define a Host Group

- 1 From the Storage Partition Definitions Window, select the Storage Subsystem or Default Host Group in the Topology View. Then, select either the **Configure >> Topology >> Define New Host Group** pull-down menu option, or **Define New Host Group** from the right-mouse pop-up window.

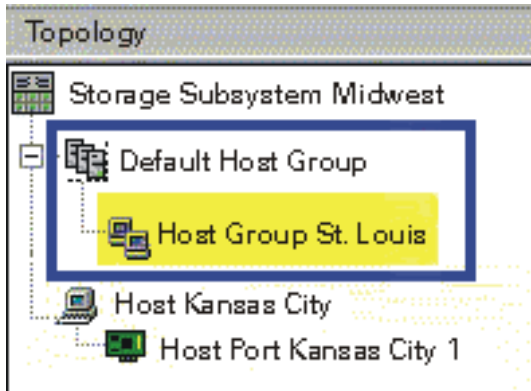
Result: The **Define New Host Group** dialog appears.

- 2 Type the name of the new host group.

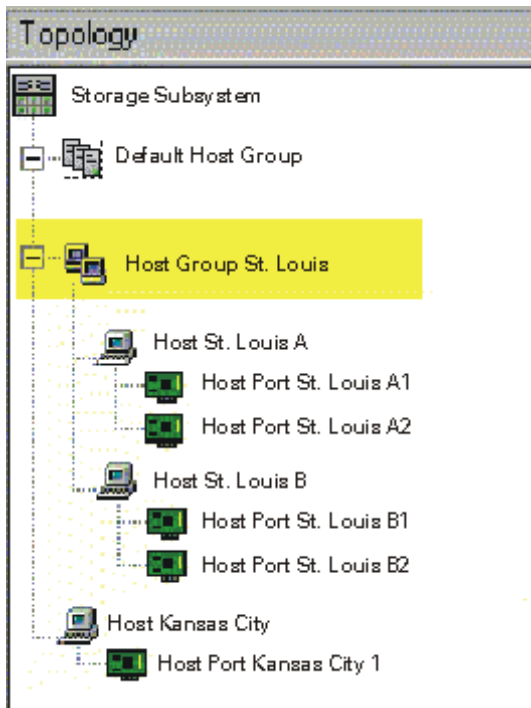
Example: St. Louis

- 3 Select **Apply** to save the newly-defined host group and leave the dialog open to define another host group. Select **OK** to save the newly-defined host group and close the dialog.

Result:



The host group St. Louis is displayed under the Default Host Group in the Topology View. Next, you should define the host group's associated hosts and host ports as shown in the following example.



When you select the St. Louis host group and define a specific logical drive-to-LUN mapping for that host group, the St. Louis host group moves from under the Default Host Group to under the Storage Subsystem in the Topology View, indicating that it has a specific logical drive-to-LUN mapping.

See [Defining Logical Drive-to-LUN Mappings](#).

Defining a Host

Use this option to define a new [host](#).

Important:

- This operation will fail if the name you choose duplicates the name of a host that already exists in the Topology View.
- After defining a host, you must define its associated host ports.
- The maximum length of the host name cannot exceed 30 characters.
- It is highly recommended that you provide the same name as the host's assigned network node name.

How to Define a New Host

- 1 From the Storage Partition Definitions Window, select either the Storage Subsystem, Default Host Group, or a Host Group in the Topology View. Then, select the **Configure >> Topology >> Define New Host** pull-down menu option, or **Define New Host** from the right-mouse pop-up window.

Result: The **Define New Host** dialog appears.

- 2 Type the name of the new host.

Example: St. Louis

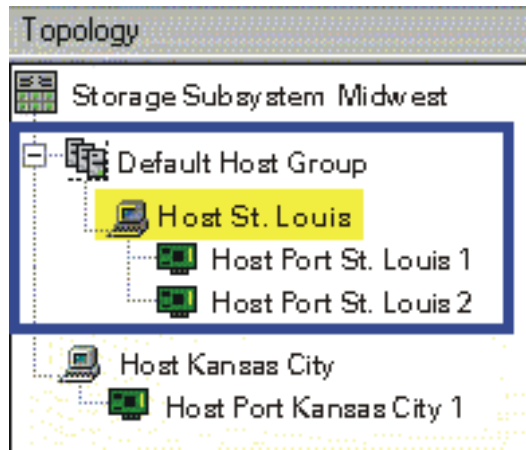
- 3 Select **Apply** to save the newly-defined host and leave the dialog box open to define another host. Select **OK** to save the newly-defined host and close the dialog. Next, define the host's associated host ports.

Result: The St. Louis host appears in one of the following two ways.

If you selected the following node when you defined the host...	Then the host St. Louis appears under...

Storage Subsystem or Default Host Group

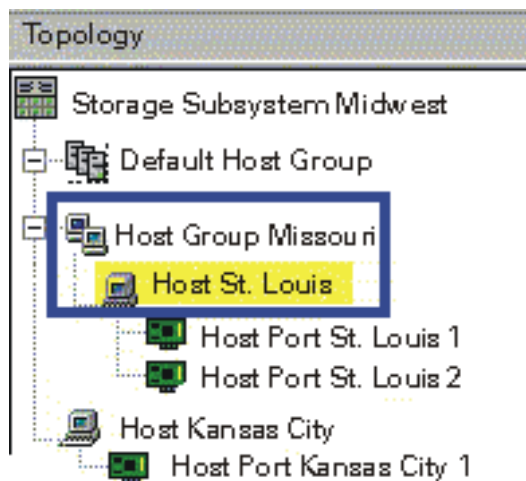
The Default Host Group node in the Topology View.



Note: When you select the St. Louis host and define a specific logical drive-to-LUN mapping for that host, the Host node will move from under the Default Host Group to under the Storage Subsystem node indicating that it has a specific logical drive-to-LUN mapping.

Host Group

The selected Host Group. For example, if you selected the Missouri host group when you defined the St. Louis host, the St. Louis host will appear under the Missouri host group in the Topology View.



Defining a Host Port

Use this option to define a name and host type for a [host port](#).

Important:

- When the host adapter has only one physical connection (host port), the terms host port and host adapter are synonymous.
- You must define all of the host port parameters (the Host port identifier, the Host Port name, and the Host type) during this operation. Leaving one or more of these parameters undefined causes the operation to fail.
- The host port identifier is the world-wide name of the host port. It is also referred to as the adapter node name.
- You can obtain the worldwide-name of the host port by using the host adapter's boot-up utility.

How to Define a Host Port

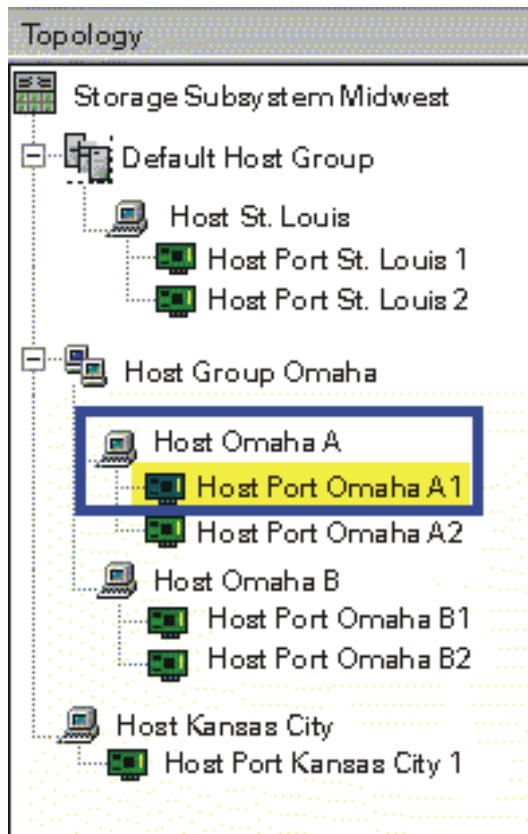
- 1 From the Storage Partition Definitions Window, select a Host in the Topology View; then, select either the **Configure >> Topology >> Define New Host Port** pull-down menu option, or the **Define New Host Port** from the right-mouse pop-up window.

Result: The **Define New Host Port** dialog appears. The **Host** label reflects the name of the Host you selected in the Topology View.

- 2 Select the Host Port parameters from the drop-down list boxes on the **Define New Host Port** dialog.

Host Port Parameter	Description
Host Port Identifier	<p>All host port identifiers that you have not associated with a host (referred to as undefined) are listed in ascending order.</p> <p>It is very important that you choose the correct identifier associated with the Host you selected in the Topology view. Also, you can enter an identifier manually.</p> <p>Note: The maximum length of the host port identifier cannot exceed 16 hexadecimal characters.</p> <p>Example: 200000e08b008250</p>
Host Port Name	<p>Host Port names previously defined using this dialog, display in ascending alphabetical order. However, you must enter a new Host Port name.</p> <p>Notes:</p> <ul style="list-style-type: none">● If no previous Host Port names have been defined, the list is blank.● The maximum length of the host port name cannot exceed 30 characters. <p>Example: For Host Omaha A, name the Host Port Omaha A1.</p>
Host Type	<p>The host type allows you to specify the operating system for this host port. For the first release of this software, the only value allowed is Automatic selection.</p> <p>Notes:</p> <ul style="list-style-type: none">● You cannot enter a value manually.● Multiple Host Types will be available in future releases of this software.

- 3 Select **Apply** to save the newly-defined host port and leave the dialog open to define another host port. Select **OK** to save the newly-defined host port and close the dialog. Next, you should define a specific logical drive-to-LUN mapping for this host.



Result: The Host Port Omaha A1 appears under the Host node Omaha A in the Topology View.

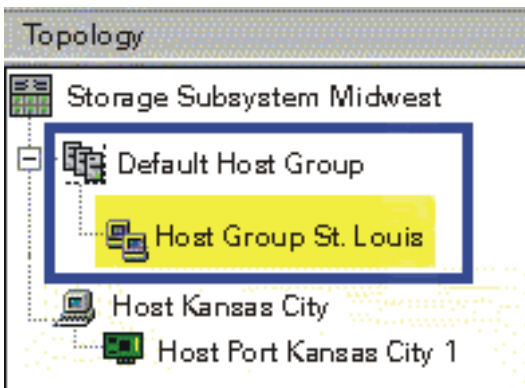
Defining a Logical Drive-to-LUN Mapping

Use this option to grant logical drive access to a particular host group and/or host and to define the logical unit number (LUN) that should be used to access it.

Important:

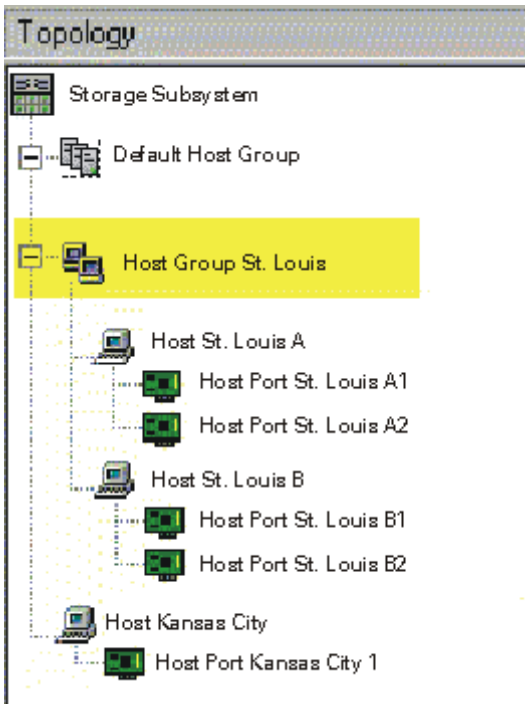
When you have defined a specific logical drive-to-LUN mapping for a host group or host that is currently under the Default Host Group, the host group or host moves out of the Default Host Group and under the Storage Subsystem (root) node indicating that the host group or host has a specific logical drive-to-LUN mapping. Once a specific logical drive-to-LUN mapping is defined for a host group or host, the host group or host no longer has access to the logical drives (if any) in the Default Host Group.

Example:



In the example to the left, the host group St. Louis has been defined, but currently does not have any defined hosts, host ports, or specific logical drive-to-LUN mappings.

Therefore, St. Louis is part of the Default Host Group and would automatically access the logical drives (if any) assigned to the Default Host Group.



In the example to the left, associated hosts and host ports have been defined for the St. Louis host group. Also, a specific logical drive-to-LUN mapping has been defined for the St. Louis host group.

Therefore, the St. Louis host group moves from under the Default Host Group to under the Storage Subsystem in the Topology View.

Now, each host in the St. Louis host group can access **ONLY** the logical drives that have been mapped to it.

Important: This option **FAIL** if:

- All logical drive-to-LUN mappings have been defined (this option will be unavailable).
- The selected host group or host does not have at least one host port defined underneath it.
- You are attempting to create a logical drive-to-LUN mapping for the host group that conflicts with an already established logical drive-to-LUN mapping for a host in the host group.
- You are attempting to create a logical drive-to-LUN mapping for a host in the host group that conflicts with an already established logical drive-to-LUN mapping for the host group.

How to Define a Logical Drive-to-LUN Mapping

- 1 From the Storage Partition Definitions Window, select the Default Host Group, or a host group or host from the Topology View; then, select either the **Configure >> Logical Drive-to-LUN Mappings >> Define New Mappings** pull-down menu option or **Define New Mappings** from the right-mouse pop-up menu.

Result: The **Define New Mappings** dialog is displayed.

The selected node is listed on the dialog in the following form:

If you selected....	Then the selected node label is listed as...
A host group	"Host Group <host group name>"
A host	"Host <host name>"
The Default Host Group	"Default Host Group"

- 2 Select a single logical drive from the logical drive-to-LUN mapping table. The logical drive-to-LUN mapping table lists the names and unique logical drive identifiers (world-wide name) of the logical drives that are available for definition.

Note: Logical Drives can only participate in one logical drive-to-LUN mapping. Logical Drives that have already been given a specific mapping do not appear in the logical drive-to-LUN mapping table. To change a logical drive-to-LUN mapping, use the **Configure >> Logical Drive-to-LUN Mappings >> Change Mapping** pull-down menu option. For more information, see [Changing a Logical Drive-to-LUN Mapping](#).

- 3 Select a single Logical Unit Number (0 through 31) from the LUN drop-down list box.

Important: The full range of logical unit numbers is displayed because the same LUN can potentially participate in multiple mappings, because each host has its own LUN address space. However, a logical drive can only be mapped to a single LUN and a single host group or host.

- 4 Select **Apply** to save the logical drive-to-LUN mapping and leave the dialog open. Select **OK** to save the logical drive-to-LUN mapping and close the dialog.

Result: The logical drive-to-LUN mapping is saved. The Topology View is updated to reflect any movement of host groups or hosts from under the Default Host Group if this was the first specific logical drive-to-LUN mapping defined for the host group or host. The Mappings View is updated to reflect the addition of specific mappings.

Changing a Logical Drive-to-LUN Mapping

Use this option to change either the assigned logical unit number (LUN) for a logical drive, or the logical drive's associated host group or host.

How to Change a Logical Drive-to-LUN Mapping

- 1 From the Storage Partition Definitions Window, select a single logical drive from the Mappings View. Then, select either the **Configure >> Logical Drive-to-LUN Mappings >> Change Mapping** pull-down menu option, or **Change Mapping** from the right-mouse pop-up menu.

Result: The **Change Mapping** dialog is displayed. The name and unique identifier of the selected logical drive are listed on the dialog.

- 2 Change the mapping as follows:

If you want to change the...	Then...
LUN	<p>Select a single Logical Unit Number (0 through 31) from the LUN drop-down list box. By default, the current LUN associated with the selected logical drive is displayed.</p> <p>Important: The full range of logical unit numbers is displayed because the same LUN can potentially participate in multiple mappings, because each host has its own LUN address space. However, a logical drive can only be mapped to a single LUN and a single host group or host.</p>

Host Group or Host	Select a single name from the drop-down list box. The drop-down list box displays the term Default Host Group followed by an alphabetical list of all defined host group or host names in the form Host Group <Host Group Name> or Host <Host Name>. By default, the current host group or host associated with the selected logical drive is displayed.
--------------------	--

- 3 Select **OK** to save the logical drive-to-LUN mapping and close the dialog.

Result: The logical drive-to-LUN mapping is checked for validity and saved. The Mappings View is updated to reflect the new mapping. The Topology View is also updated to reflect any movement of host groups or hosts from under the Default Host Group if the logical drive-to-LUN mapping was changed to a host group or host that did not previously have any specific logical drive-to-LUN mappings.

Important:

- If you change the assignment for a host group or host that only has one logical drive-to-LUN mapping, the host group or host will move back to the Default Host Group and inherit any default mappings.
- If you want to change the last mapping of a host group that also has an associated host with specific mappings, then you must delete or move the host mappings first.

Deleting a Host Port, Host, Host Group, or Logical Drive Mapping

Use this option to:

- Delete a host port, host, or host group from the Topology View.
- Delete a logical drive-to-LUN mapping.

Important:

- You cannot delete the root (Storage Subsystem) node or the Default Host Group node.
- If you are deleting the last logical drive-to-LUN mapping for a host group and the host group also has hosts with specific logical drive-to-LUN mappings, make sure you delete or move those mappings before deleting the last logical drive-to-LUN mapping for the host group.

How to Delete a Node or Logical Drive Mapping

- 1 From the Storage Partition Definitions Window, select a host group, host, or host port from the Topology View, or a single logical drive from the Mappings View; then do one of the following:
 - Select the **Edit >> Delete** pull-down menu option.
 - Select **Delete** from the right-mouse pop-up menu.
 - Press the Delete key.

Result: The **Confirm Delete** dialog is displayed.

- 2 Select **Yes** to delete the node or the logical drive mapping and close the **Confirm Delete** dialog.

Result:

If you selected a... Then...

Host Group, Host, or Host Port under the Default Host Group Node...

<p>Host Group</p>	<p>The host group and its associated hosts and host ports are removed from the Topology View.</p> <p>Any default logical drive-to-LUN mappings associated with the Default Host Group are unaffected and remain in the Mappings View.</p> <p>Any host ports that were associated with the host group become undefined. For information on how to view the undefined host ports, see Showing Undefined Host Ports.</p>
<p>Host</p>	<p>The host and its associated host ports are removed from the Topology View.</p> <p>Any default logical drive-to-LUN mappings associated with the Default Host Group are unaffected and remain in the Mappings View. Also, if the host is part of a host group, the host group is unaffected.</p> <p>Any host ports that were associated with the host become undefined. For information on how to view the undefined host ports, see Showing Undefined Host Ports.</p>

<p>Host port associated with a host group or a host</p>	<p>The host port is removed from the Topology View and is no longer associated with a host group or host. However, any host group or host associated with the host port is unaffected.</p> <p>Any default logical drive-to-LUN mapping associated with the Default Host Group are unaffected.</p> <p>For information on how to view the undefined host ports, see Showing Undefined Host Ports.</p>
---	---

Host Group, Host, or Host Port that has Specific Logical Drive-to-LUN Mappings...

<p>Host Group</p>	<p>The host group and its associated hosts and host ports are removed from the Topology View.</p> <p>Any specific logical drive-to-LUN mappings are deleted from the Mappings View and the associated logical drives are available for a new mapping using the Define New Mapping option.</p> <p>Any host ports that were associated with the host group become undefined. For information on how to view the undefined host ports, see Showing Undefined Host Ports.</p>
-------------------	--

Host	<p>The host and its associated host port(s) are removed from the Topology View.</p> <p>Any specific logical drive-to-LUN mappings are deleted from the Mappings View and the associated logical drives are available for a new mapping using the Define New Mapping option.</p> <p>If the host is part of a host group that has its own specific logical drive-to-LUN mappings, the host group is unaffected. However, if the host is part of a host group that does not have any other mappings, the host group and any other associated hosts/host ports will be moved back to the Default Host Group and inherit any default mappings.</p> <p>Any host ports that were associated with the host become undefined. For information on how to view the undefined host ports, see Showing Undefined Host Ports.</p>
Host Port	<p>The host port is removed from the Topology View and is no longer associated with a host group or host. However, any host group or host associated with the host port is unaffected.</p> <p>Any specific logical drive-to-LUN mappings associated with the host group or host are unaffected. However, if this was the last host port, these logical drive-to-LUN mappings will be inaccessible until you re-define associated host(s) and host port(s).</p> <p>For information on how to view the undefined host ports, see Showing Undefined Host Ports.</p>

A Logical Drive-To-LUN Mapping in the Mappings View...

<p>The only logical drive-to-LUN mapping or the last logical drive-to-LUN mapping associated with a host group or a host</p>	<p>The logical drive-to-LUN mapping is removed from the Mappings View. The associated logical drive is available for a new mapping using the Define New Mapping option.</p> <p>The host group and/or hosts and host ports associated with the mapping are moved back under the Default Host Group in the Topology View and inherit any default mappings.</p> <p>If this mapping is associated with a host group and one or more of its hosts has specific mappings, the deletion will not be allowed until you move or delete the specific mappings of the host(s).</p>
<p>One of many logical drives associated with a host group or a host</p>	<p>The logical drive-to-LUN mapping is removed from the Mappings View and is no longer associated with its respective host group or host. The associated logical drive is available for a new mapping using the Define New Mapping option.</p> <p>The host group or host is unaffected.</p>

Moving a Host Port

Use this option if you are going to move a host adapter ([host port](#)) from one host to another host.

Important:

- If you are replacing a failed host adapter with a new host adapter, then it is not appropriate to use this option. Use the [Replace Host Port](#) option instead.
- You must use the following sequence when moving a host port: (1) Physically remove the host adapter from the host. (2) Use this **Move Host Port** option to move the host adapter (host port) to the desired host in the Topology View. (3) Physically insert the host adapter into the new host. (4) Reboot both hosts so that the changes can be recognized by their operating systems.

How to Move a Host Port

- 1 From the Storage Partition Definitions Window, select a host port in the Topology View; then select either the **Configure >> Topology >> Move Host Port** pull-down menu option, or **Move Host Port** from the right-mouse pop-up menu.

Result: The **Move Host Port** dialog is displayed. The Host Port name and Host port identifier (world-wide name) are listed on the dialog.

Example: You want to move the host port St. Louis A1 from host St Louis A to host St. Louis B. You selected the host port St. Louis A1 under the host St. Louis A.

- 2 Select the desired Host from the **Host name** drop-down list box. The list displays an alphabetical list of all defined hosts.

Example: You select St. Louis B.

3 Select **OK**.

Result: The changed host port-to-host definition is saved and the dialog closes. The host port (in our example, St. Louis A1) is removed from the previous host (St. Louis A) and is now displayed under the new host (St. Louis B). The host port inherits the logical drive-to-LUN mappings of St. Louis B.

Note: If the host port was the only port listed under the previous host, the previous host is still displayed in the Topology View and retains its current logical drive-to-LUN mappings. However, these logical drive-to-LUN mappings will be inaccessible until you re-define at least one associated host port for this host.

Replacing a Host Port

Use this option if you need to replace a failed host adapter in a host.

Important: It is extremely important that you use one of the following sequences when replacing the failed host adapter to ensure that the host will retain its current logical drive-to-LUN mappings. If you do NOT have any default logical drive-to-LUN mappings in the Default Host Group, use sequence 1. If you do have default logical drive-to-LUN mappings in the Default Host Group, use sequence 2.

Replacement Sequence 1: (1) Physically replace the failed host adapter with the new host adapter and reboot the host. Use the host adapter's boot-up utility to obtain the host port identifier (world-wide name) of the new host adapter. (2) Use this **Replace Host Port** option to indicate the new host adapter (host port) that replaced your failed host adapter.

Replacement Sequence 2: (1) Use this **Replace Host Port** option to indicate the new host adapter (host port) that will replace your failed host adapter. You must obtain the host port identifier (world-wide name) of the new host adapter before physically inserting it in the intended host. (2) Physically replace the failed host adapter with the new host adapter. (3) Reboot the host so that the new host adapter is identified by the host operating system.

How to Replace a Host Port

- 1 From the Storage Partition Definitions Window, select the failed host port from the Topology View; then select either the **Configure >> Topology >> Replace Host Port** pull-down menu option, or **Replace Host Port** from the right-mouse pop-up menu.

Result: The **Replace Host Port** dialog is displayed. The current host port identifier (world-wide name) is listed on the dialog.

- 2 Define a replacement host port using the **New identifier** drop-down list. The drop-down list displays all undefined host ports (that is, all host ports that have not been associated with a host) in ascending identifier order. The first undefined host port is displayed by default.

Also, you can enter an identifier manually.

Important: The maximum length of the host port identifier cannot exceed 16 hexadecimal characters.

- 3 Enter a host port name in the **Host port name:** edit box. By default, the edit box displays the host port name associated with the current identifier. Keep the same name or use a similar name.

Important:

- The host port name is a required field.
- If the name of the host port duplicates the name of another host port in the Storage Subsystem, you will receive an error message and a prompt to enter a new host port name.

- 4 Select **OK**.

Result: The changed host port definition is saved and the dialog closes.

Renaming a Host Group, Host, or Host Port

Use this option to change the name of a host group, host, or host port.

Important:

- The maximum length of the name cannot exceed 30 characters.
- The operation will **FAIL** if you choose a name that duplicates the name of another node of the same type. For example, if you were renaming a host that duplicates the name of another host in your Topology View, an error message is displayed and you are prompted to choose another name.
- Any leading and trailing spaces in the name are deleted.

How to Rename a Host Group, Host, or Host Port

- 1 From the Storage Partition Definitions Window, select a host group, host, or host port from the Topology View; then, select either the **Edit >> Rename** pull-down menu option, or **Rename** from the right-mouse pop-up menu.

Result: The **Rename** dialog is displayed. The edit box displays the current name of the selected node.

- 2 Type the new name and select **OK**.

Result: The node is renamed and the dialog closes.

Note: If you delete the text from the edit box and select **OK**, the original name of the node is retained.

Moving a Host from One Host Group to Another Host Group

Use this option to:

- Move a host from one host group to another host group
- Remove a host from a host group

How to Move a Host

- 1 From the Storage Partition Definitions Window, select a host from the Topology View; then, select either the **Configure >> Topology >> Move Host** pull-down menu option, or **Move Host** from the right-mouse pop-up menu.

Result: The **Move Host** dialog is displayed. The Host name is listed on the dialog.

Example: You desire to move the host Omaha A from host group Omaha to host group Kansas City. You selected the host Omaha A under the host group Omaha.

- 2 Select the desired host group from the **Host Group name:** drop-down list. The drop-down list displays a choice of "None", followed by an alphabetical list of all host groups.

Important: The "None" choice is appropriate when you want to remove a host from any defined host group.

Example: You select the host group Kansas City.

3 Select OK.

Result: The changed host-to-host group definition is saved and the dialog closes. There are several result variations:

If you selected	Then
A different host group	<p>The host and its associated host ports are removed from the previous host group and are now displayed under the chosen host group. The host and its associated host ports inherit the new host group's logical drive-to-LUN mappings. If the host also had any of its own specific mappings, these are retained.</p> <p>In our example, Omaha A is removed from the Omaha host group and is now displayed under the Kansas City host group.</p> <p>Condition 1: If the host was the only host under the previous host group, the host group loses its association to the host but retains the host group's current logical drive-to-LUN mappings. However, these logical drive-to-LUN mappings will be inaccessible until you re-define associated host(s) and host port(s) for this host group.</p> <p>Condition 2: If the host was the only host under the previous host group and the host group does not have any logical drive-to-LUN mappings, the host group is moved back under the Default Host Group and inherits any default logical drive-to-LUN mappings.</p>

None	<p>Condition 1: If the host does not have any of its own specific logical drive-to-LUN mappings and only has the mappings associated with the host group, the host (and its associated host ports) are returned to the Default Host Group and inherit any default mappings.</p> <p>Condition 2: If the host has its own specific logical drive-to-LUN mappings in addition to the mappings associated with the host group, the host (and its associated host ports) are moved under the Storage Subsystem indicating that the host still has its own specific logical drive-to-LUN mappings but it is no longer associated with a host group.</p>
------	---

Showing Undefined Host Ports

Use this option to view a list of discovered host ports that have not been currently defined (that is, associated with any hosts). Use the **Define New Host Port** option to associate the undefined host ports with their hosts.

- 1 From the Storage Partition Definitions Window, select the Default Host Group, then select the **Show Undefined Host Ports** option from the right-mouse pop-up menu.

Result: The **Undefined Host Ports** dialog is displayed.

- 2 Select **Close** to close the dialog.

Changing the Enclosure Order

Use this option to change the order of the controller and drive enclosures in the Physical View to match the actual hardware configuration in your Storage Subsystem. The Physical View that is initially displayed is a default view that might not match your Storage Subsystem.

Note: The enclosure order change remains in effect until it is modified again.

How to Change the Enclosure Order

- 1 From the Subsystem Management Window, select the **Configure >> Enclosure Order** pull-down menu option.

Result: The **Change Enclosure Order** dialog box is displayed.

- 2 From the list, select the enclosure you want to move.
- 3 Press the appropriate button on the dialog box to move the enclosure to the new position.
- 4 Select **OK**.

Result: The enclosure order in the Physical View is changed.

Resetting a Storage Subsystem Configuration

Use this option to:

Reclaim disk space from a factory logical drive configuration before manually or automatically configuring the Storage Subsystem.

-OR-

Correct major configuration errors that cannot be corrected using any other method.



Caution: All data currently stored on the Storage Subsystem will be lost. Do not attempt this operation unless all data on the Storage Subsystem has been backed up to tape or to another Storage Subsystem.

Important: Resetting the configuration on this Storage Subsystem will cause the following to occur:

- All logical drives and arrays will be deleted
- The user-supplied name will be deleted
- If you have password protected your Storage Subsystem for destructive operations, that password protection will be lost
- The controller pair will be set to active/passive mode.

How to Reset the Storage Subsystem Configuration

- 1 From the Subsystem Management Window, select the **Configure >> Reset Configuration** pull-down menu option.

Result: The **Confirm Reset Configuration** dialog is displayed.

- 2 Select **Yes**.

Result: Because this is a dangerous operation, a second **Confirm Reset Configuration** dialog is displayed.

3

If...	Then...
You don't want to delete the Storage Subsystem Configuration	Select No . Result: You are returned to the Subsystem Management Window.
You are absolutely sure you want to delete your entire Storage Subsystem configuration.	Select Yes . Result: <ul style="list-style-type: none">● The current configuration is deleted, destroying any existing data on the Storage Subsystem.● A controller-defined configuration is applied (which may include a default logical drive). All other drives are unassigned.● The Logical and Physical Views are updated to reflect the new configuration.

Locating an Array

Use this option to physically locate and identify all drives that comprise a selected array. An indicator light is activated on each drive in the array.

How to Locate an Array

- 1 To physically locate an array, select the array in the Logical View; then, select either the **Array >>Locate** pull-down menu option, or **Locate** from the right-mouse pop-up menu.

Result: The **Locate Array** dialog is displayed. The indicator lights on the drives in the array flash.

- 2 When you have located the array, select **OK**.

Result: The lights will stop flashing. If there are any other **Locate** operations (**Locate Drive**, **Locate Storage Subsystem**) currently running from another Management Station, these lights will also stop flashing.

Note: If the lights from the **Locate Array** operation do not stop flashing, use the [Stop All Indications](#) option to stop the process manually.

Changing the RAID Level of an Array

Use this option to change the RAID level on a selected array. Using this option will change the RAID level of every logical drive that comprises the array. Performance might be slightly affected during the operation.

Important:

- You cannot cancel this operation after it begins.
- Your data remains available during this operation.
- The array must be in an Optimal state before you can perform this operation.
- If you do not have enough capacity in the array to convert to the new RAID level, you will receive an error message and the operation will not continue. If you have unassigned drives, use the **Array >> Add Free Capacity** option to add additional capacity to the array. Then, retry the operation.

RAID Level Drive Number Constraints

There are drive number requirements for choosing a particular RAID level. If you attempt to change to a RAID level for which you do not have the appropriate number of drives in the array, an error message is displayed.

RAID level	Drive Number Requirements
RAID 0	<ul style="list-style-type: none">● You are limited to a maximum of 30 drives in the array
RAID 1	<ul style="list-style-type: none">● You must have an even number of drives in the array. If you do not have an even number of drives and you have some remaining unassigned drives, use the Array>>Add Free Capacity option to add additional drives to the array, then retry the operation.● You are limited to a maximum of 30 drives in the array

RAID 3 or 5	<ul style="list-style-type: none">● You must have a minimum of 3 drives in the array● You are limited to a maximum of 30 drives in the array
-------------	---

How to Change the RAID Level


- 1 Select the array in the Logical View; then, select either the **Array >> Change RAID Level** pull-down menu option, or **Change RAID Level** from the right-mouse pop-up menu.

Result: The RAID level fly-out menu is displayed.

- 2 Select the RAID level ([RAID 0](#), [RAID 1](#), [RAID 3](#), or [RAID 5](#)). The currently selected option is designated with a dot.

Result: The **Change RAID Level Confirmation** dialog is displayed.

- 3 Select **Yes**.

Result: The RAID level operation begins. The logical drive icons under the selected array in the Logical View show a status of Operation in Progress  until the operation is complete. To view the progress or change the priority of the modification operation, select a logical drive in the array, then select **Logical Drive >> Properties**. Note that the priority is set on an individual logical drive basis.

Changing Controller Ownership of an Array

Use this option to change the controller ownership of a selected array.

Important:

- This option is only available when you have an [active/active controller pair](#).
- Changing controller ownership of an array changes the [preferred owner](#) of this array. The preferred owner of an array is listed under this menu option with the designation of "preferred" and normally a dot next to the menu item. In certain situations, the logical drives can be temporarily moved to the other controller. Under these circumstances, the menu option uses the dot to identify the controller that currently owns the logical drives (the [current owner](#)) and the "preferred" designation to show the preferred owner.



Caution: If you do not use a driver that understands multiple I/O paths to a logical drive, shut down any host applications that are currently using this array to prevent application errors when the I/O path changes.

How to Change Controller Ownership

To change controller ownership, select an array in the Logical View; then, either select the **Array >> Change Ownership** pull-down menu option and select the controller from the fly-out menu, or Select Change Ownership from the right-mouse pop-up menu, then select the controller from the fly-out menu.

Result: The ownership of the array is changed. I/O to the logical drives is now directed through this I/O path.

Important: Under certain operating system environments, it might be necessary to reconfigure the host driver before this I/O path can be used.

Adding Free Capacity to an Array

Use this option to expand the capacity of a selected array by adding [unassigned drives](#). The result of the expansion causes additional free capacity to be included in the array. You can use this free capacity to create additional logical drives.

Important:

- You cannot cancel this operation after it begins.
- Your data remains available during this operation.
- The array must be in an Optimal state before you can perform this operation.
- This option will work only if there are unassigned drives in the Storage Subsystem.
- Existing logical drives in the array do not increase in size as a result of this operation. Existing logical drive capacity is redistributed over the larger number of drives.
- There is a controller firmware-determined limit to the number of drives you can pick for a single **Add Free Capacity** operation. Selecting more drives than the allowed number causes an error message to display. However, after you have completed the first **Add Free Capacity** operation, you can repeat the option to add more drives until the desired capacity is reached.
- In a RAID 1 logical drive, you must add a minimum of two drives (You must add an even number of drives, half of which are used for redundancy).
- Whenever possible, select drives that have a capacity equal to the current drive capacities in the array.
- If you must add unassigned drives with a smaller capacity, be aware that the usable capacity of each drive currently in the array will be reduced so that the drive capacity is the same across the array. If there is data on the current drives in the array that could be lost when the usable capacity is reduced, an error message is displayed.
- If you must add unassigned drives with a larger capacity, be aware that the usable capacity of the unassigned drives that you add will be reduced so that they match the current drives capacities in the array.

How to Add Free Capacity to an Array

- 1 Select an array in the Logical View; then, select either the **Array >> Add Free Capacity** pull-down menu option, or **Add Free Capacity** from the right-mouse pop-up menu.

Result: The **Add Free Capacity** dialog is displayed. The dialog presents the RAID level and the existing drive capacities of the selected array. A table is also displayed that contains information about the unassigned drives that are available for selection. For fibre channel drives, the ID is the actual ID rather than the preferred ID.

Important Tips:


- If you are adding SCSI drives, it is best to select drives that are on different channels so that the impact of a channel failure is minimized.
 - Whenever possible, select drives that have a capacity equal to the current drive capacities in the array.
- 2 Select up to the maximum allowable number of drives in the list, using CTRL + CLICK (or SHIFT + CLICK for drives that are next to each other on the list).
 - 3 Select **Add**.

Result: The **Add Free Capacity Confirmation** dialog box is displayed.

- 4 Select **Yes** to add the capacity to the array.



Result: The **Add Free Capacity** operation begins. The Logical View and Physical View changes are listed below.

Logical View Changes

- Each logical drive in the array displays an  Operation in Progress icon while the data blocks are physically reorganized.
- A **Defragment Array** operation automatically takes place as a result of this operation.
- If a Free Capacity node did not previously exist, a new Free Capacity node is added at the bottom of the tree, showing its RAID-factored capacity.
- If one or more Free Capacity nodes previously existed, they are consolidated during the defragment operation to create one Free Capacity node.
- The Unconfigured Capacity node will update to reflect the subtraction of the

usable capacity of the unassigned drives that were added to the array.

Physical View Changes

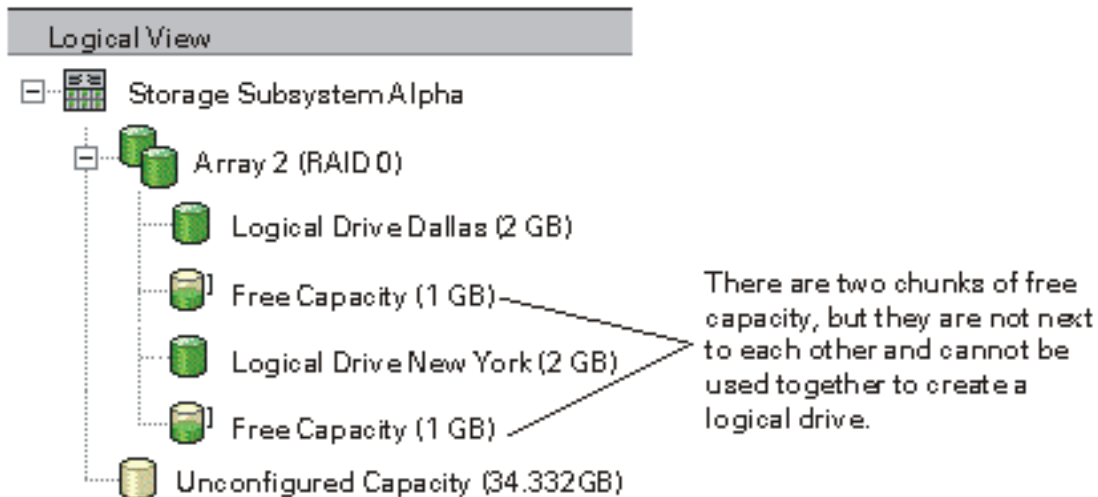
The drives selected for addition to the array were previously represented in the Physical View with an Unassigned drive icon . They are now displayed with an Assigned drive  icon.

Defragmenting an Array

Use this option to consolidate all Free Capacity on a selected Array so that additional logical drives can be created from the maximum amount of free capacity. A fragmented Array can result from the deletion of logical drives or from not using the entire available free capacity in a Free Capacity node during logical drive creation.

Example: You have two logical drives, "New York" and "Dallas" of 2 GB each. You would like to create a new logical drive, "Chicago" of 2 GB in the same array. There is 2 GB of space available in the array, but it is split between two Free Capacity nodes that are not contiguous and cannot be organized into a single logical drive.

In order to create a new logical drive "Chicago" of 2 GB, you must first defragment the array so that the space in the two Free Capacity nodes is contiguous.



Important:

- You cannot cancel the operation once it begins.
- Your data remains accessible during the defragment operation.
- In order to use this option, all logical drives in the array must be online and Optimal, and there can be no logical drive modification operations (such as Change of Segment Size) in progress

How to Defragment an Array

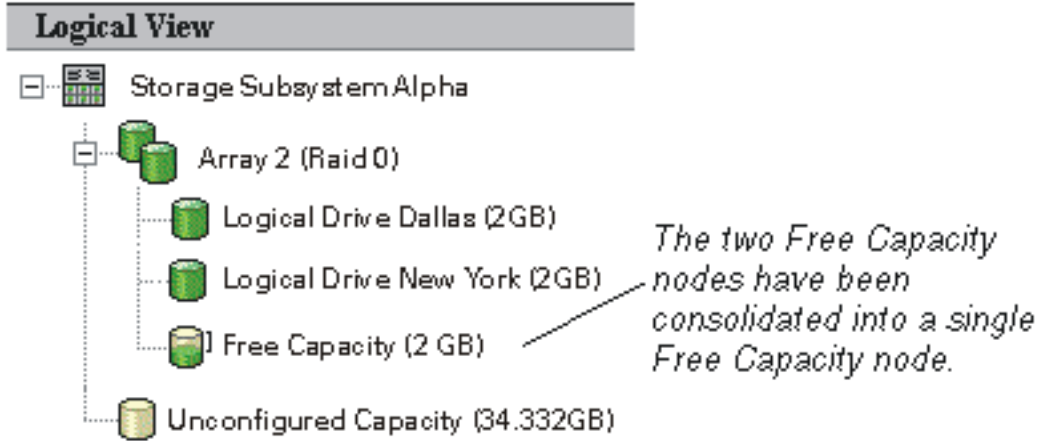
- 1 Select an array in the Logical View, then select either the **Array >> Defragment** pull-down menu option, or **Defragment** from the right-mouse pop-up menu.

Result: The **Confirm Defragment** dialog box is displayed.

- 2 Select **Yes**.

Result: All free data blocks on the Free Capacity nodes are consolidated into one Free Capacity node containing a larger set of contiguous data blocks. The Logical View is updated to display the larger, single Free Capacity node. You can select this Free Capacity node for logical drive creation.

Example: In the previous example, there were two Free Capacity nodes of 1 GB each in Array 2. After a Defragment Array operation, there is a resulting single Free Capacity node of 2 GB.



Placing an Array Offline or Online



Caution: This option is available for use ONLY under the guidance of a customer service representative.

Initializing an Array

Use this option when the Recovery Guru advises the manual initialization of all logical drives in an array.



Caution: All logical drive data on all logical drives in the array will be lost. Do not attempt this operation unless advised to do so by the Recovery Guru.

Important:

- You cannot cancel the operation once it begins.
- Do not use this option if the logical drives in the array are not Optimal.
- Do not use this option if there are any modification operations in progress on the logical drive or the array.
- Do not change cache parameters of the logical drive while the **Initialize** operation is in progress.

How to Initialize a Logical Drive

- 1 Select the array to initialize from the Logical View; then, select either the **Array >> Initialize** pull-down menu option, or **Initialize** from the right-mouse pop-up menu.

Result: The **Confirm Initialization** dialog is displayed.

- 2 Select **Yes**.

Result: The logical drives in the array begin the initialization process. During initialization, each logical drive icon shows an Operation in Progress status

 in the Logical View.

Note: To monitor initialization progress, select the logical drive, then select **Logical Drive>>Properties**.

Reviving an Array

Use this option to revive (attempt to spin up) the drives in a selected array. The Recovery Guru advises using this option when a failed drive SCSI channel or a failed environmental card cause a multi-drive failure. In either case, it is possible that the data on the drives is not corrupt and the logical drives can be successfully revived.



Caution: Do not use this option unless advised to do so by the Recovery Guru. This option is NOT appropriate for certain failure conditions. Using this option without the guidance of a Recovery Guru procedure can result in an optimal logical drive with data corruption.

Important: This operation will fail if one or more of the drives is unusable and cannot be revived.



How to Revive an Array

- 1 Select the array in the Logical View; then, select either the **Array >> Revive** pull-down menu option, or **Revive** from the right-mouse pop-up menu.

Result: The **Confirm Revive** dialog box is displayed.

- 2 Select **Yes**.

Result:

- The logical drives in the array return to an Optimal state  in the Logical View.
- Drives in the Physical View that are associated with the array will return to an optimal state. 

Checking Data Redundancy on an Array

Use this option to check the [redundancy](#) on a selected array only when instructed to do so by the Recovery Guru.

Important:

- This option cannot be used on RAID 0 arrays, which have no redundancy.
- If you use this option on a RAID 1 array, the redundancy check compares the data on the mirrored drives.
- If you perform this operation on a RAID 3 or RAID 5 array, the redundancy check checks the parity information that is striped across the drives.
- To successfully perform this operation:
 - The logical drives in the array must be Optimal.
 - The array must have no logical drive modification operations in progress.
- This option can only be performed on a single array at a time. However, you can alternatively perform a redundancy check on selected logical drives during a [media scan](#) operation. See [Specifying Logical Drive Properties](#) for information on enabling a media scan redundancy check on one or more logical drives in the Storage Subsystem.

How to Perform a Redundancy Check

- 1 Select a single array in the Logical View; then, select either the **Array >> Check Redundancy** pull-down menu option, or **Check Redundancy** from the right-mouse pop-up menu.

Result: The **Check Redundancy** dialog is displayed. The logical drives in the array are displayed in a table column on the dialog. A second column provides redundancy check status for each logical drive.

2 Select **Start**.

Result: The **Check Redundancy** operation begins. The logical drives in the array are sequentially scanned, starting from the top of the table in the logical drive dialog. As each logical drive is scanned:

- The logical drive is highlighted in the logical drive table
- The status of the redundancy check is shown in the associated status column.

Status	This status displays when...
Pending	This is the first logical drive in the array to be scanned, but the Start button has not been pressed to start the redundancy check. -OR- The redundancy check operation is being performed on other logical drives in the array.
Checking	The logical drive is undergoing the redundancy check
Passed	The logical drive passed the redundancy check (there were no inconsistencies in the redundancy information)
Failed	The logical drive failed the redundancy check (there were inconsistencies in the redundancy information)

- The progress bar indicates the completion percentage of the redundancy check for a single logical drive, with 100% being complete. If an operation error occurs before the redundancy check completes, an error message is displayed.

3 Select **Done** after the last logical drive in the array has been checked.

Result: The **Check Redundancy** dialog closes and you are returned to the Subsystem Management Window.

Deleting an Array

Use this option to delete an array and all of the logical drives it contains, resulting in more unconfigured capacity that can be reconfigured to better meet application storage needs.



Caution:

- Deleting an array and its contained logical drives causes the loss of all data contained on those logical drives. Back up the data on all the logical drives in the array and stop all I/O before performing this operation, if necessary.
- If a filesystem is mounted on the logical drive, unmount it before attempting this operation.

Important:

You cannot delete an array that:

- Contains any initializing logical drives
- Contains any reconstructing logical drives
- Is undergoing a copyback operation on one of its assigned drives
- Is undergoing a modification operation such as a change of RAID level or a capacity expansion unless the logical drives in the array are now in a Failed state

How to Delete an Array

- 1 Select an array in the Logical View; then, select either the **Array >> Delete** pull-down menu option, or **Delete** from the right-mouse pop-up menu.

Result: The **Confirm Delete** dialog is displayed.

- 2 Select **Yes**.

Result: All associated assigned drives of the selected array change to the unassigned state in the Physical View. The array and all its associated logical drives are deleted and their icons are removed from the Logical View. The raw capacity from the drives formerly associated with the deleted array is added to an existing Unconfigured Capacity node or to a new Unconfigured Capacity node if one did not previously exist.

Changing Segment Size

Use this option to change the [segment size](#) on a selected logical drive. During this operation, I/O performance is affected but your data remains available.

How Long Does a Change Segment Size Operation Take?

The operation is slower than other modification operations (for example, changing RAID levels or adding free capacity to an array) because of the way the data is physically reorganized and because of the temporary internal backup procedures that occur during the operation.

The specific amount of time a Change Segment Size operation can take depends on many variables, including the I/O load from the hosts, the [modification priority](#) of the logical drive, the number of drives in the array, the number of drive channels, and the processing power of the Storage Subsystem controllers.

If you want this operation to complete faster, you can change the modification priority, although this may decrease system I/O performance. To change the priority, select a logical drive in the array, then select **Logical Drive >> Properties**.

Important:

- You cannot cancel this operation once it begins.
- Do not begin this operation unless the array is Optimal.
- The controller firmware determines the segment size transitions that are allowed. Segment sizes that are inappropriate transitions from the current segment size are unavailable on the menu. Allowed transitions are typically a doubling or halving of current segment size. For example, if the current logical drive segment size is 32K, a new logical drive segment size of either 16K or 64K is allowed.

When to Change Segment Size

A default segment size is set during logical drive creation, based on the array RAID level and the logical drive usage you specify. These two parameters should optimize the segment size appropriately for your environment. However, it is

important to monitor your Storage Subsystem and change segment size when necessary for optimal performance based on the guidelines below.

If your typical I/O size is larger than your segment size, increase your segment size in order to minimize the number of drives needed to satisfy an I/O request. This is especially true in a multi-user, database or filesystem storage environment. Using a single drive for a single request leaves other drives available to simultaneously service other requests.

If you are using the logical drive in a single-user, large I/O environment (such as for multimedia application storage) performance is optimized when a single I/O request can be serviced with a single data stripe (the segment size multiplied by the number of drives in the array that are used for I/O). In this case multiple disks are used for the same request, but each disk is only accessed once.

How to Change Segment Size


- 1 Select a logical drive in the Logical View; then, select either the **Logical Drive >> Change Segment Size** pull-down menu option, or **Change Segment Size** from the right-mouse pop-up menu.

Result: The flyout menu of segment sizes is displayed. The current segment size is marked with a black dot to its left. Any segment sizes that are not appropriate for the array RAID level are not available.

- 2 Select the desired segment size.

Result: The **Change Segment Size Confirmation** dialog box appears.

- 3 Select **Yes**.

Result: The segment size modification operation begins. The logical drive icon in the Logical View will show an Operation in Progress  status while the operation is taking place.

Note: To view the progress or change the priority of the modification operation, select a logical drive in the array, then select **Logical Drive >> Properties**.

Deleting a Logical Drive

Use this option to delete a logical drive that:

- Was created with the wrong parameters or capacity, or
- No longer meets application storage needs.

This option results in an increase of free capacity in the array or additional unconfigured capacity.



Caution:

- Deleting a logical drive causes the loss of all data contained on the logical drive. Back up the data and stop all I/O before performing this operation, if necessary.
- If a filesystem is mounted on the logical drive, unmount it before attempting this operation.

Important:

You cannot delete a logical drive that is:

- Initializing
- Reconstructing
- Part of an array that contains a drive that is undergoing a copyback operation
- Undergoing a modification operation such as a change of segment size unless the logical drive is now in a Failed state

How to Delete a Logical Drive

- 1 Select the logical drive in the Logical View; then, select either the **Logical Drive >> Delete** pull-down menu option, or **Delete** from the right-mouse pop-up menu.

Result: The **Confirm Delete** dialog is displayed.

2 Select **Yes**.

Result: The logical drive is deleted and any data that resided on it is destroyed. In addition:

If the deleted logical drive was...	Then...
The only logical drive in the array	<ul style="list-style-type: none">● The drives associated with the deleted logical drive return to an unassigned state in the Physical View● The array is removed from the Logical View● The raw capacity of the newly unassigned drives is added to the Unconfigured Capacity node in the Logical View
One of multiple logical drives in the array	<ul style="list-style-type: none">● The size of the Free Capacity node is increased if the deleted logical drive was adjacent to a Free Capacity node in the array. Otherwise, a new Free Capacity node is created.

3 If you see system error messages relating to this deleted logical drive, reconfigure or reboot your host system to permanently remove any system information about this logical drive.

Initializing a Logical Drive

A logical drive is automatically initialized when it is first created. However, the Recovery Guru might advise the manual initialization of a logical drive to recover from certain failure conditions.



Caution: All logical drive data will be lost. Do not attempt this operation unless advised to do so by the Recovery Guru.

Important:


- You cannot cancel the operation once it begins.
- Do not use this option if there are any modification operations in progress on the logical drive or the array.
- Do not change cache parameters of the logical drive while the **Initialize** operation is in progress.

How to Initialize a Logical Drive

- 1 Select the logical drive to initialize in the Logical View; then, select either the **Logical Drive >> Initialize** pull-down menu option, or **Initialize** from the right-mouse pop-up menu.

Result: The **Confirm Initialization** dialog is displayed.

- 2 Select **Yes**.

Result: The logical drive begins the initialization process. During initialization, a logical drive icon shows an Operation in Progress status  in the Logical View.

Note: To monitor initialization progress, select the logical drive; then, select **Logical Drive >> Properties**.

Renaming a Logical Drive

Use this option to change the name of a logical drive when the original name specified during its creation is no longer meaningful or appropriate.

Important: If you try to rename a logical drive with a name that is already in use by another logical drive in the Storage Subsystem, an error message is displayed and you are prompted to choose another name for the logical drive.

Tips on naming logical drives

- There is a 30 character limit. All leading and trailing spaces will be deleted from the name.
- Use a unique, meaningful name that is easy to understand and remember. In most cases, the most meaningful name will describe the data that the logical drive stores.
- Avoid arbitrary names or names that would quickly lose their meaning in the future.
- When you have named the logical drive, the prefix "Logical Drive" is automatically added to the logical drive name shown in the Logical View. For example, if you renamed the logical drive "Chicago," the Logical View will display "Logical Drive Chicago." You do not need to add the prefix "Logical Drive" to the name you specify in this option.

How to Rename a Logical Drive

- 1 Select the logical drive in the Logical View, then select either the **Logical Drive >> Rename** pull-down menu option, or **Rename** from the right-mouse pop-up menu.

Result: The **Rename Logical Drive** dialog is displayed, containing the current name of the logical drive.

- 2 Type the new name.
- 3 Select **OK**.

Result: The logical drive is renamed. The name shown in the Logical View is updated to show the prefix Logical Drive followed by the new logical drive name. For example, if you renamed the logical drive "Chicago," the Logical View will display "Logical Drive Chicago."

Specifying Logical Drive Properties

Use the Logical Drive Properties dialog to specify the [modification priority](#), cache properties, and media scan properties of a selected logical drive. Also, you can view the progress of a logical drive modification operation on this dialog.

How to Access the Logical Drive Properties Dialog

To access the **Logical Drive Properties** dialog, select a logical drive in the Logical View; then, select either the **Logical Drive >> Properties** pull-down menu option, or **Properties** from the right-mouse pop-up menu.

Result: The **Logical Drive Properties** dialog is displayed.

Base Properties

Viewing Base Properties Information

The Base tab of the Logical Drive Properties dialog displays the following logical drive data: logical drive name, [world wide name](#), status, capacity, RAID level, and segment size.

Setting Modification Priority

The **modification priority** defines how much processing time is allocated for logical drive modification operations relative to system performance. You can increase the logical drive modification priority, although this may affect system performance.

Operations affected by the modification priority include:

- Copyback
- Reconstruction
- Initialization
- Changing Segment Size
- Defragmentation of an array

- Adding Free Capacity to an array
- Changing RAID Level of an array

Modification Priority rates

The following priority rates are available.

- Lowest
- Low
- Medium
- High
- Highest

Note: The Lowest priority rate favors system performance, and the modification operation will take longer. The Highest priority rate favors the modification operation, and system performance may suffer.

To set the modification priority, click and hold the settings button on the slider bar and move it to the desired priority.

Viewing the Progress of a Modification Operation

The progress bar at the bottom of the **Logical Drive Properties** dialog displays the progress of an operation.

You can view the progress of the following operations:

- Copyback
- Reconstruction
- Initialization
- Change RAID Level*
- Add Capacity*
- Defragment*
- Change Segment Size*

* The Progress Bar label shows "Modification progress" rather than the operation name.

Important: The storage management software cannot obtain progress information from the controllers in the Storage Subsystem if the network management connection to the controllers is down or if the Storage Subsystem is

partially managed. In this case, the progress bar on the dialog displays "zero percent complete" and the text above the progress bar displays "<operation> progress data unavailable." For more information on a partially managed Storage Subsystem or a Unresponsive controller or Storage Subsystem condition, see the Enterprise Management Window help system.

Cache Properties

Cache memory is an area of temporary volatile storage (RAM) on the controller that has a faster access time than the actual drive media. By using cache, you can increase overall application performance because:

- Data blocks for a read from the host might already be in the cache from a previous operation, eliminating the need for drive access
- Write data is written initially to the cache which frees the application to continue

To measure cache effectiveness, examine the **cache hit percentage** in the Performance Monitor. This statistic shows the percentage of data requests that are serviced by the cache and did not require a disk access. For more information, see [Monitoring Performance](#).

Setting Cache Properties

Use the Cache Properties tab to enable or set the following cache parameters.

Note: In certain situations, the controller may temporarily suspend caching operations when caching is not appropriate. This status is called **Enabled but Not Active**.

Cache Parameter	Description
Read Caching	<p>Allows read operations from the host to be stored in controller cache memory.</p> <p>If a host requests data that is not in the cache, the controller reads the needed data blocks from the disk and then places them in the cache. Until the cache is flushed, all other requests for this data are fulfilled with cache data rather than from a physical disk read, increasing throughput.</p> <p>To enable the read caching parameter, select the Read Caching checkbox.</p>
Write Caching	<p>Allows write operations from the host to be stored in cache memory.</p> <p>The logical drive data in the cache is written to disk, or flushed, automatically at the interval you specify in the Flush write caching after parameter.</p> <p>Note: There also are cache settings you can set at the Storage Subsystem level that affect all logical drives. Select Storage Subsystem >> Change Cache Settings. For more information, see Specifying Storage Subsystem Cache Settings.</p> <p>To enable the write caching parameter, select the Write Caching checkbox.</p>
Write cache mirroring	<p>Allows cached data to be mirrored across redundant controllers with the same cache size. Data written to the cache memory of one controller is also written to the cache memory of the other controller. Therefore, if one controller fails, the other can complete all outstanding write operations.</p> <p>Important: To use this option:</p> <ul style="list-style-type: none"> ● You must be using an active/active controller pair ● Both controllers must have the same size cache. To check controller cache size, select the controller in the Physical View; then, select either the Controller >> Properties pull-down menu option, or Properties from the right-mouse pop-up menu.

	<p>To enable the write cache mirroring parameter, select the Write Cache Mirroring checkbox.</p> <p>Note: This option is available only when write caching is also enabled.</p>
<p>Write caching without batteries</p>	<p>Allows write caching to continue, even if the controller batteries are discharged completely, not fully charged, or not present. If you select this parameter without a UPS for back-up power, you could lose data if power fails.</p> <p>Enable the Write Caching Without Batteries parameter by selecting the checkbox.</p> <p>Note: This option is only available if write caching is also enabled.</p>
<p>Cache read-ahead multiplier</p>	<p>Cache read-ahead allows the controller, while it is reading and copying host-requested data blocks from disk into the cache, to copy additional data blocks into the cache. This increases the chance that a future request for data could be fulfilled from the cache. Cache read-ahead is important for multimedia applications that use sequential I/O.</p> <p>The cache read-ahead multiplier values specify the multiplier to use to determine how many additional data blocks are read into cache. Choosing a higher cache read-ahead multiplier may increase the cache hit percentage, however, transfer time might also increase.</p> <p>In the cache read-ahead multiplier spinner box, choose a multiplier.</p> <p>Disabling cache read-ahead</p> <p>If you need to disable cache read-ahead, select 0 in the cache read-ahead multiplier spinner box.</p>

Media Scan Properties

A media scan is a background process that runs on all logical drives in the Storage Subsystem for which it has been enabled, providing error detection on the drive media. Enabling the media scan process allows the process to find media errors before they disrupt normal drive reads and writes. The media scan process scans all logical drive data to verify that it can be accessed, and optionally scans the logical drive redundancy data.

Note: You can set the duration over which the media scan runs by selecting the **Storage Subsystem >> Media Scan Duration** pull-down menu option.

A media scan discovers the following errors and reports them to the Event Log:

Error	Description/Result
Unrecovered media error	<p>The drive could not read the requested data on its first attempt, or on any subsequent retries.</p> <p>Result:</p> <p>Logical Drives with redundancy protection. Data is reconstructed, rewritten to the drive, and verified. The error is reported to the event log.</p> <p>Logical Drives without redundancy protection (RAID 0 logical drives and degraded RAID 1, 3, and 5 logical drives). The error is not corrected but it is reported to the event log.</p>
Recovered media error	<p>The drive could not read the requested data on its first attempt, but succeeded on a subsequent attempt.</p> <p>Result: The data is rewritten to the drive and verified. The error is reported to the event log.</p>

<p>Redundancy mismatches</p>	<p>Redundancy errors are found and a media error is forced on the block stripe so that it is found when the drive is rechecked. If redundancy is repaired, this forced media error is removed.</p> <p>Result: The first 10 redundancy mismatches found on a logical drive are reported to the event log.</p> <p>Note: The media scan checks for redundancy only if the optional redundancy checkbox is enabled. For more information, see <i>Enabling a Redundancy Check</i>, below.</p>
<p>Unfixable error</p>	<p>The data could not be read, and parity or redundancy information could not be used to regenerate it. For example, redundancy information cannot be used to reconstruct data on a degraded logical drive.</p> <p>Result: The error is reported to the event log.</p>

Enabling a Background Media Scan

To enable a background media scan for this logical drive, select the **Enable background scan** check box.

Enabling a Redundancy Check

During a redundancy check, all data blocks in a logical drive are scanned, and:

- In a RAID 3 or 5 logical drive, redundancy is checked and repaired
- In a RAID 1 logical drive, the data is compared between the mirrored drives and data inconsistencies are corrected

Note: RAID 0 logical drives have no data redundancy.

Select the **With redundancy check** radio button to enable a redundancy check during a media scan.

Applying Individual Logical Drive Properties to All Logical Drives

To apply all settings of an individual logical drive (modification priority, cache, and media scan) to all other logical drives in the Storage Subsystem, click the **Apply Settings to All Logical Drives** checkbox. Use this checkbox only if you want all of the logical drives to share the same settings. This option overwrites any current settings on the others logical drives in the Storage Subsystem.

Placing a Controller Offline or Online

Placing a Controller Offline

Use the **Place Offline** option to place a controller in an Offline state.

Note: This option is only available when the controller is online.



Caution:

- Use this option only when instructed to do so by Recovery Guru or a Customer Support representative.
- If you do not use write cache mirroring, data in the cache of the controller you placed offline will be lost.
- If you take a controller offline and you have controller failover protection through a host multi-path driver, the other controller in the pair takes over. Arrays and their associated logical drives that were assigned to the offline controller are automatically reassigned to the remaining good controller. If you do not have a multi-path driver installed on the application host and you take a controller offline while the application is using the associated logical drives, application errors will occur.

How to Place a Controller Offline

- 1 To take a controller offline, select it in the Physical View; then, select either the **Controller >>Place Offline** pull-down menu option, or **Place Offline** from the right-mouse pop-up menu.

Result: The **Confirm Place Offline** dialog is displayed.

- 2 Select **Yes**.

Result: The Place Offline operation begins. The controller icon changes to

Needs Attention, Passive .

- 3 If you do not have a multi-path driver installed on the application host, make the appropriate host operating system modifications so that the applications can recognize the new I/O path.


Placing a Controller Online

Use the **Place Online** option to place the controller back online, making it available for I/O.

Note: This option is only available when the controller is offline.

How to Place a Controller Online

- 1 Select the offline controller in the Physical View; then, select either the **Controller >>Place Online** pull-down menu option, or **Place Online** from the right-mouse pop-up menu.

Result: The **Place Online** option begins. The controller icons changes to Optimal, Passive .

- 2 Use the **Controller >> Change Mode** option to change the controller to Active Mode.

Result: The controller icon changes to Optimal, Active



- 3 Use the **Storage Subsystem >> Redistribute Arrays** to move the formerly owned logical drives back to this controller owner.

Changing a Controller Mode

Use this option to place a controller in [Active](#) or [Passive](#) mode.

Note: The Active option is only available for a selected controller that is currently Passive. Conversely, the Passive option is only available for a selected controller that is currently Active.



Caution: If you do not have a multi-path driver installed on the application host and you change a controller from Active to Passive mode while the application is using the associated logical drives, application errors will occur.

Important: If you change an Active controller to Passive mode and the other controller of the pair is already in Passive mode, an error is displayed and the operation will not complete. At least one controller in the Storage Subsystem must be in Active mode.

How to Change a Controller to Active Mode

- 1 Select the controller in the Physical View; then, select either the **Controller >> Change Mode >> Active** pull-down menu option, or **Change Mode >> Active** from the right-mouse pop-up menu.

Result: The **Confirm Change to Active Mode** dialog is displayed.

- 2 Select **Yes**.

Result: The Controller changes to Active mode. The controller icon in the Physical View changes to Active. 

Important: Arrays and associated logical drives are not moved to the new controller automatically.

- 3 Move one or more arrays and their associated logical drives to the newly active controller.


If...	Then...
The newly active controller previously owned arrays and their associated logical drives (before a previous switch to Passive mode)	Use the Storage Subsystem >> Redistribute Arrays option to move the arrays and logical drives originally owned by this controller back to this controller.
The newly active controller never previously owned arrays and their associated logical drives	Use the Array >> Change Ownership option to move one or more arrays to this controller -OR- Create new arrays and logical drives and specify this controller as the owner.

How to Change a Controller to Passive Mode

- 1 Select the controller in the Physical View; then, select either the **Controller >> Change Mode >> Passive** pull-down menu option, or **Change Mode >> Passive** from the right-mouse pop-up menu.

Result: The **Confirm Change to Passive Mode** dialog is displayed.

- 2 Select **Yes**.

Result: The controller changes to Passive mode. The controller icon in the Physical View changes to Passive . Any arrays and their associated logical drives that were owned by this controller change ownership to the other controller, and all I/O to these logical drives is automatically rerouted by the multi-path driver on the application host.

- 3 If you do not have a multi-path driver installed on the application host, make the appropriate host operating system modifications so that the applications can recognize the new I/O path.

Listing Components Associated with a Controller

This option is used to list the logical and physical components associated with a controller.

How to List Components Associated with a Controller

- 1 Select a controller in the Physical View; then, select either the **Controller >> List Associated Components** pull-down menu option, or **List Associated Components** from the right-mouse pop-up menu.

Result: The **Associated Components** dialog is displayed. Listed information includes the array numbers of associated arrays, logical drive names of associated logical drives, and the enclosure and slot IDs of associated drives.

Notes: The icons on the dialog do not reflect the actual state of the associated arrays, logical drives, and drives. Always use the Physical View to see the current state of associated components.

- 2 Select **Close** to close the dialog.

Viewing Controller Properties

Use this option to view the properties of a selected controller.

How to View Controller Properties

- 1 Select a controller in the Physical View; then, select either the **Controller >>Properties** pull-down menu option, or **Properties** from the right-mouse pop-up menu.

Result: The **Controller Properties** dialog is displayed.

The **Base** tab displays basic controller information, such as the manufacturer and serial number of the controller and the cache/processor size.

The **Interfaces** tab displays information about the drive and host interfaces to this controller, in two tables. For more information, see the sections of this topic called About the Drive Interfaces Table and About the Host Interfaces Table.

- 2 Select **Close** to close the dialog.

About the Drive Interfaces Table

The information in the Drive interfaces table varies, depending on the interface type.

Interface Type	Information displayed
SCSI	<p>Each row in this table describes a controller drive interface, an access path from this controller to drives in the Storage Subsystem. Every drive interface in the Storage Subsystem associated with the selected controller is displayed in this table.</p> <p>There are two columns in the SCSI drive interfaces table:</p> <ul style="list-style-type: none"> ● The Channel column lists the SCSI channel number of the drive interface ● The Current ID column lists the SCSI ID of the controller making the connection to a drive through this channel.
Fibre channel	<p>Each row in this table describes a controller drive interface, an access path from this controller to a single drive port on a Storage Subsystem drive. (There will be two fibre channel drive interfaces to every drive, one to each drive port.) Every drive interface in the Storage Subsystem associated with the selected controller is displayed in this table.</p> <p>There are two columns in the fibre channel drive interfaces table:</p> <ul style="list-style-type: none"> ● The Channel column lists the channel number of the drive interface. There are multiple channels that can be used on a single controller. ● The Current ID column lists the actual Loop ID of the controller port connecting through this channel to a drive port.

About the Host Interfaces Table

The information in the host interfaces table varies depending on the interface type.

Interface Type	Information displayed
SCSI	<p>Each row in this table describes a controller host interface, a connection from a SCSI controller to a host adapter.</p> <p>There are two columns in the SCSI host interfaces table:</p> <ul style="list-style-type: none"> ● The Channel column lists the SCSI channel number of the controller host interface ● The Current ID column lists the SCSI ID of the controller host interface
Fibre channel	<p>Each row in this table describes a controller host interface, a connection from a fibre channel controller port to a port on a host. Each fibre channel controller may have one or more ports that can be used to access ports on a host.</p> <p>There are four columns in the fibre channel host interfaces table:</p> <ul style="list-style-type: none"> ● The Port column lists the controller ports used to connect to ports on a host. ● The Current ID column lists the current fibre channel Loop ID of the controller port associated with the host interface. The Loop ID is in the form Decimal Loop ID / hexadecimal AL_PA. <p>Note: If the Preferred ID is set to 126, indicating that the controller host interface is using a point-to-point connection rather than an arbitrated loop, the value <code>Not Applicable</code> is displayed in this column.</p> <ul style="list-style-type: none"> ● The Preferred ID lists the preferred fibre channel Loop ID of the controller port associated with the host interface. The Loop ID is in the form Decimal Loop ID / hexadecimal AL_PA. You can set the preferred ID through the Change Preferred ID dialog box. On some models of controller, you may have to set the preferred ID using a hardware switch on the controller. ● The NL-Port ID lists the hexadecimal NL-Port Identifier of the controller port associated with this host interface. The NL-Port ID is a 24-bit address used in the S_ID and D_ID fields of a frame.

Changing the Preferred Loop ID of a Controller Host Interface Port

Use this option to change the [preferred loop ID](#) of the controller port used in a host interface connection. This is the only way to set the Preferred Loop ID for most fourth generation controllers that do not contain a hardware switch on the back of the controller enclosure. In addition, you can override or extend the hardware switch settings for fourth generation controllers that do contain a switch.



Important: This option is grayed out until you select a fibre channel host interface in the Host Interfaces table of the [Controller Properties](#) dialog.

How to Change the Preferred Loop ID of a Controller Port

- 1 From the Controller Properties dialog, select a fibre channel host interface in the Host Interfaces table, then select **Change Preferred ID**.

Result: The **Change Preferred Loop ID** dialog is displayed. The Preferred Loop ID currently displayed on the **Controller Properties** dialog for this controller port is shown.

- 2 Select one of the three radio buttons:

<p>Normal (0-125)</p>	<p>Use this option to assign a normal Preferred ID value to the controller port. Select a value from the spinner box.</p> <p>If the controller port currently has a Preferred ID between 0 and 125, this radio button is selected by default, and the current ID is shown in the spinner box.</p> <p> Caution: If your fibre channel host adapter driver manages devices based on Loop-ID (or AL_PA), selecting a Preferred ID value that is already in use by another device might result in data corruption.</p>
<p>N-port only (126)</p>	<p>Use this option only to configure the controller host interface as a point-to-point connection.</p> <p>Important: This forces the controller to operate as an N-port, disabling loop mode.</p> <p>If the controller port currently has a Preferred ID of 126, this radio button is selected by default.</p>
<p>Set during loop initialization (127)</p>	<p>Use this option to assign the controller port to accept any available address during loop initialization.</p> <p> Caution: If your fibre channel host adapter driver manages devices based on Loop-ID (or AL_PA), selecting this option might result in data corruption.</p> <p>If the controller port currently has a Preferred ID of 127, this radio button is selected by default.</p>

3 Select **OK**.

Result: The Preferred Loop ID of the controller port is changed.

Viewing the Status of Controller Enclosure Components

Use this option to:

- View the status of all components in the controller enclosure
- Reset the battery age clock, when you have replaced the battery

How to Check Battery Status

The battery status is shown with an icon and a label that show its status.

Icon and label for an OK battery status



Status: OK

Icon and label for a Failed battery status



Status: Failed

The battery area also shows the current **Age** of the battery and the number of **Days Until Replacement**.

How to Reset the Battery Age Clock



Caution: Only reset the battery age when you are following the direction of Recovery Guru to replace a battery that has neared or reached its expiration date.

- 1 Select the **Controller Enclosure Components** button  in the Physical View.

Result: The **Controller Enclosure Components** dialog is displayed.

- 2 Select **Reset**.


Result: The **Confirm Reset** dialog is displayed.

- 3 Select **Yes**.


Result: The battery age clock is reset to zero.

How to Check the Status of Controller Enclosure Components

- 1 Select the **Controller Enclosure Components** button  in the Physical View.

Note: If one or more components in the controller enclosure are not in an Optimal state, then the **Controller Enclosure Components** button in the Physical View reflects a Needs Attention status. 

Result: The **Controller Enclosure Components** dialog is displayed.

- 2 The status for every controller enclosure component is shown. If you see a Needs Attention  enclosure status in the Physical View, use this dialog display for more information on the potential problem. Then, run Recovery Guru and take the recommended action to correct the problem.

The following are examples of Normal and Needs Attention status icons:

Normal Component Status



Fan
canisters: OK

Needs Attention Component Status



Fan canisters: Failed

- 3 Select **Close**.

Result: The **Controller Enclosure Components** dialog is closed.

Locating a Drive

Use this option to physically locate and identify one or more drives by activating drive indicator lights.

How to Locate a Drive

- 1 To physically locate a drive or drives, select the drives in the Physical View; then, select either the **Drive>>Locate** pull-down menu option, or **Locate** from the right-mouse pop-up menu.

Result: The **Locate Drive** dialog is displayed. The indicator lights on the selected drives flash.

- 2 When you have located the drives, select **OK**.

Result: The lights will stop flashing. If there are any other Locate operations (**Locate Array**, **Locate Storage Subsystem**) currently being invoked from another management station, these lights will also stop flashing.


Note: If the lights from the **Locate Drive** operation do not stop flashing, use the [Stop All Indications](#) option to stop the process manually.

Assigning and Deassigning a Hot Spare

Assigning a Hot Spare

Use this option to assign a selected unassigned drive as a hot spare drive. A hot spare drive contains no data and acts as a standby in case any drive fails in a RAID level 1, 3, or 5 array on the Storage Subsystem. The hot spare drive adds another level of redundancy to the Storage Subsystem.

Important:

- You can only assign Optimal unassigned drives  as hot spare drives.
- The hot spare drive is available to any RAID level 1, 3, or 5 array in the Storage Subsystem. It is not specifically associated with any single array.
- When you assign a drive to be a hot spare, make sure that there is at least one other drive on the Storage Subsystem that is not assigned as a hot spare and that has a capacity equal to or smaller than the drive you assign as a hot spare.
- Select a drive with a capacity equal to or larger than the configured capacity of the drive you want to cover with the hot spare. For example, if you have an 18 GB drive with configured capacity of 8 GB, you could use a 9 GB or larger drive as a hot spare. In general, however, you should not assign a drive to be a hot spare unless its capacity is equal to or greater than the capacity of the largest drive on the Storage Subsystem.
- The maximum number of SCSI hot spare drives you can create is equal to the number of SCSI drive channels supported by the controller. The maximum number of fibre channel hot spare drives you can create is four.

How to Assign a Hot Spare Drive

To assign a hot spare drive, select one or more drives in the Physical View from the available unassigned drives; then, select either the **Drive >> Assign Hot Spare** pull-down menu option, or **Assign Hot Spare** from the right-mouse pop-up menu.


Result: The drive changes from unassigned to standby hot spare.  The drive's

capacity is subtracted from the Unconfigured Capacity node in the Logical View.

Deassigning a Hot Spare


Use this option to return a selected hot spare drive or drives to an unassigned state.

Caution:

- You can only use this option on standby, optimal  hot spares. A hot spare drive cannot be deassigned if it is in use (taking over for a failed drive).
- If a hot spare drive is not in an Optimal state, follow the Recovery Guru procedures for correcting the problem before trying to deassign the drive.

How to Deassign a Hot Spare Drive

To deassign a hot spare drive, select one or more standby hot spare drives in the Physical View; then, select either the **Drive >> Deassign Hot Spare** pull-down menu option, or **Deassign Hot Spare** from the right-mouse pop-up menu.

Result: The drive changes from standby hot spare to unassigned.  The drive's capacity is added to the Unconfigured Capacity node in the Logical View.

Failing a Drive

Use this option to fail a selected drive or drives.



Caution:

- This operation could result in data loss or the loss of data redundancy. Do not use this operation unless advised to do so by the Recovery Guru or a Customer Support Representative.
- This operation will fail when there is a logical drive modification operation in progress.

Conditions in which data loss will occur:

- Failure of a single drive in a RAID 0 logical drive
- Failure of a single drive in a Degraded RAID 1 logical drive, if the drive failed is the mirror drive of another failed drive
- Failure of a single drive in a Degraded RAID 3 or 5 logical drive
- Failure of two or more drives in an Optimal RAID 3 or 5 logical drive

Conditions in which a loss of data redundancy will occur:

- Failure of a single drive in a RAID 1 logical drive
- Failure of a single drive in an Optimal RAID 3 or 5 logical drive

How to Fail a Drive or Drives

- 1 Select one or more drives in the Physical View; then, select either the **Drive >> Fail** pull-down menu option, or **Fail** from the right-mouse pop-up menu.



Result: The **Confirm Fail Drive** dialog is displayed.

2 Select **Yes**.

Result: The drives fail and the dialog closes. The drive in the Physical View is shown in the Needs Attention state  in the selected mode.

Note: If you are going to remove the drive, make sure at least 30 seconds have elapsed.

If this drive was assigned to an array, then the Logical View changes are shown in the table below.

If...	Then...
The logical drive was Optimal before the drive was Failed	The logical drive changes to a Degraded state  in the Logical View.
The logical drive was Degraded before the drive was Failed	The logical drive changes to a Fail state  in the Logical View.

Reviving a Drive

Use this option to revive (attempt to spin up) a selected drive.



Caution: Do not use this option unless you have been advised to do so by the Recovery Guru. This option is NOT appropriate for certain failure conditions.

How to Revive a Drive

- 1 Select one or more drives in the Physical View; then, select either the **Drive >> Revive** pull-down menu option, or **Revive** from the right-mouse pop-up menu.

Result: The **Revive Drive** dialog is displayed.

- 2 Select **Yes**.

Result: The selected drives are revived. The drives return to an Optimal state



in the Physical View.

Reconstructing a Drive

Use this option to manually start reconstruction of a drive ONLY when instructed to do so by Recovery Guru.

Important:




This option is available only when a drive is assigned to a RAID 1, 3, or 5 array with either a [Failed](#) status or a [Replaced](#) status, and reconstruction did not automatically start.

How to Reconstruct Drives Manually

- 1 Select one or more drives in the Physical View; then, select either the **Drive >> Reconstruct** pull-down menu option, or **Reconstruct** from the right-mouse pop-up menu.

Result: The **Confirm Reconstruct** dialog is displayed.

- 2 Select **Yes**.

Result: The **Confirm Reconstruct** dialog is closed and the drive starts to reconstruct. The drive icon in the Physical View changes to Optimal . The associated logical drive icons in the Logical View change to Operation in Progress  during the reconstruction, then to Optimal  when the reconstruction is finished.

To view the reconstruction progress, select a logical drive in the Logical View that is associated with this drive, then select **Logical Drive >> Properties**.

Initializing a Drive

Use this option only when you have moved a sub-set of drives that were previously part of a multi-drive array from one Storage Subsystem to another. Because you did not move the entire set of drives, the logical drive/array information on the drives is incomplete. Therefore, to erase all logical drive/array information on these drives and make them available again for new logical drive creation, you must initialize these drives. This will erase the logical drive/array information and return the selected drives to an unassigned state, adding new or additional unconfigured capacity to the Storage Subsystem. You can use this capacity to create additional logical drives.



Caution: All data on the drive or drives will be lost.


How to Initialize Selected Drives

- 1 From the Physical View, select one or more drives you want to initialize; then, select either the **Drive >> Initialize** pull-down menu option, or **Initialize** from the right-mouse pop-up menu.

Result: The **Initialize Drive** dialog is displayed.

- 2 Select **Yes**.

Result:

- The drive initialization process begins.
- The drive in the Physical View returns to an unassigned state. 
- The Logical View is updated as follows:

If...	Then...
An Unconfigured Capacity node did not exist prior to the drive initialization operation	A new Unconfigured Capacity node is added to the Logical View, containing the raw capacity of the newly initialized drives.
An Unconfigured Capacity node existed prior to the drive initialization operation	The raw capacity of the newly initialized drives is added to the existing Unconfigured Capacity node.

Listing Components Associated with a Drive

Use this option to list the logical and physical components associated with a drive.

How to List Components Associated with a Drive

- 1 Select a drive in the Physical View; then, select either the **Drive >> List Associated Components** pull-down menu option, or **List Associated Components** from the right-mouse pop-up menu.

Result: The **Associated Components** dialog is displayed. Listed information includes the array numbers of associated arrays, logical drive names of associated logical drives, and the slot designation (A or B) of the associated controller.

Notes: The icons on the dialog do not reflect the actual state of the associated arrays, logical drives, and controllers. Always use the Physical View to see the current state of associated components.

- 2 Select **Close** to close the dialog.

Viewing Drive Properties

Use this option to view the properties of a selected drive.

How to View Drive Properties

- 1 Select a drive in the Physical View; then, select either the **Drive >> Properties** pull-down menu option, or **Properties** from the right-mouse pop-up menu.

Result: The **Drive Properties** dialog is displayed.

- 2 Select **Close** to close the dialog.

Interpreting Drive Properties

Because some of the fields in this dialog require no explanation, the table below focuses only on those fields which require additional detail.

Field	Description
Channel/ ID	Displayed for SCSI drives. The SCSI channel number is listed first, then the SCSI ID. Example: Channel 2/ ID 1
Port 1 Channel/ ID	Displayed for fibre channel drives. Port 1: The drive port performs data communications over the fibre channel link. Fibre channel drives have dual ports, for redundancy. Each port is connected to one of the channels. This data field is associated with the first drive port seen by the controller. Channel: The Channel is the controller's drive interface channel number. ID: The Loop ID (actual) of the drive is an ID for the drive unique to the drive port. Each port resides on one


	<p>redundant drive loop. The preferred Loop ID, used to determine the Preferred AL_PA (hard address), might not be available on both loops, so the ID for the drive can vary from Drive Port 1 to Drive Port 2.</p> <p>Example: Port 1 Channel 1/ID 2</p>
Port 2 Channel/ ID	<p>Displayed for fibre channel drives.</p> <p>This is the channel and ID information for the second drive port seen by the controller. See the previous information for Port 1 Channel/ID.</p> <p>Example: Port 2 Channel 2/ID 6</p>
Raw Capacity	<p>This is the full capacity of the drive, including any capacity that is reserved for Storage Subsystem management and control.</p> <p>Example: 8.479 GB</p>
Usable Capacity	<p>This is the capacity of the drive after space reserved for Storage Subsystem management and control is allocated on the drive.</p> <p>Example: 8.44 GB</p>

Viewing the Status of Drive Enclosure Components


Use this option to view the status of all components in the drive enclosure.

How to Check the Status of Drive Enclosure Components

- 1 Select the **Drive Enclosure Components** button  in the Physical View.

Note: If one or more components in the controller enclosure are not in an Optimal state, then the **Drive Enclosure Components** button in the Physical View reflects a Needs Attention status. 

Result: The **Drive Enclosure Components** dialog is displayed.

- 2 The status for every drive enclosure component is shown. If you see a Needs Attention  enclosure status in the Physical View, use this dialog display for more information on the potential problem. Then, run Recovery Guru and take the recommended action to correct the problem.

The following are examples of Normal and Needs Attention status indicators:

Normal Component Status **Needs Attention Component Status**



Fan canisters:
OK



Fan canisters: Failed

- 3 Select **Close**.

Result: The **Drive Enclosure Components** dialog is closed.

Access Logical Drive

An Access Logical Drive is a factory-configured logical drive used for communication between the storage management software and a Host-Agent Managed Storage Subsystem controller.

Important: An Access Logical Drive is NOT needed for Directly Managed Storage Subsystems.

The Access Logical Drive:

- Uses 20 MB of storage space that is not available for application data storage
- Has a unique worldwide name
- Must be registered with the host operating system
- Uses a LUN address

Note: Each controller in the Storage Subsystem uses the same LUN to communicate with its own Access Logical Drive.

The Access Logical Drive and [Logical Unit Number](#) information are shown in the Storage Subsystem Profile as shown in the following examples. For more information, see [Viewing a Storage Subsystem Profile](#).

Example 1: The Storage Subsystem does not have an Access Logical Drive. The Storage Subsystem Profile entry is:

```
Access Logical Drive: None
```

Example 2: The Storage Subsystem controllers reach the Access Logical Drive through LUN 7. The Storage Subsystem Profile entry is:

```
Access Logical Drive: Logical Unit Number 7
```

Active Mode

A mode in which the controller can have specific assigned logical drives and handle I/O requests from the host for those logical drives.

Use the **Controller >> Change Mode** option to set a controller to active or passive mode. For more information, see [Changing a Controller Mode](#).

Active-Active Controller Pair

Also known as dual-active. The status of a controller pair in which both controllers have specific logical drives assigned to them and handle I/O requests from the host for those logical drives.

If one of the controllers fails, and multi-path failover support exists, the logical drives are switched over to the other active controller.

This configuration is recommended over the active-passive configuration for performance reasons.

Active-Passive Controller Pair

In this configuration, one controller (the active controller) handles all of the I/O from the host, and the other controller (the passive controller) is held as a spare, taking over if the active controller fails.

Use the **Controller >> Change Mode** option to set a controller to active or passive mode. For more information, see [Changing a Controller Mode](#).

Arbitrated Loop Physical Address (AL_PA)

The controller ID for Fibre Channel controllers using an Arbitrated Loop topology. This ID is a 1-byte address used to identify each port on an arbitrated loop.

Arbitrated Loop

Fibre Channel topology reporting that the controller is in a loop environment in which two to 126 devices can be interconnected. Communication is managed using an arbitration process in which the lowest port address has the highest priority.

Also known as a Fibre Channel Arbitrated Loop (FC-AL).

Array

A set of drives that the controller logically groups together to provide one or more logical drives to an application host.

Array Sequence Number

A controller-assigned number from 1 to n , where n is the maximum number of logical drives you can create in the Storage Subsystem. (See your Installation and Support Guide for information on this limit.)

These numbers may change as you create and delete logical drives and arrays.

These numbers are displayed in the Storage Subsystem Profile and in the Recovery Guru to identify specific arrays.

Note: This number is not synonymous with Logical Unit Number and is not known or used by the application host. It is needed by the storage management software.

Assigned Drive

An assigned drive is part of an array. In the Physical View of the Subsystem

Management Window, an Optimal assigned drive is represented by this icon .

Cache Memory

Cache memory is an area of temporary volatile storage (RAM) on the controller that has a faster access time than the actual drive media. By using cache, you can increase overall application performance because:

- Data blocks for a read from the host may already be in the cache from a previous operation, eliminating the need for drive access
- Write data is written initially to the cache which frees the application to continue

Use the **Logical Drive>>Properties** dialog to set logical drive cache parameters. See, [Specifying Logical Drive Properties](#).

Note: There also are cache settings you can set at the Storage Subsystem level that affect **all** logical drives. Select **Storage Subsystem >> Change Cache Settings**. For more information, see [Changing Storage Subsystem Cache Settings](#).

Cache Read Ahead

Cache read-ahead, also known as **prefetch**, allows the controller, while it is reading and copying host-requested data blocks from disk into the cache, to copy additional data blocks into the cache. This increases the chance that a future request for data could be fulfilled from the cache. Cache read-ahead is important for multimedia applications that use sequential I/O.

The cache read-ahead multiplier values specify the multiplier to use to determine how many additional data blocks are read into cache. Choosing a higher cache read-ahead multiplier might increase the cache hit percentage, however, transfer time might also increase.

Enable and set cache read-ahead on the **Logical Drive Properties** dialog. For more information, see [Specifying Logical Drive Properties](#).

Cache Without Batteries

This parameter allows write caching to continue even if the controller batteries are discharged completely, not fully charged, or not present. If you select this parameter without a UPS for back-up power, you could lose data if power fails.

Enable the **Write Caching Without Batteries** parameter on the **Logical Drive >> Properties** dialog. For more information, see [Specifying Logical Drive Properties](#).

Channel Protection

When creating logical drives from unconfigured capacity, the storage management software displays either **Yes** ✓ or **No** ✘ in the **Channel Protection** column for each array candidate.

In a SCSI environment, whether or not the array candidate has channel protection depends on the RAID level of the array and how many drives are present on any single drive channel. For example, a RAID 5 array would not have channel protection if more than one drive was present on a single drive channel.

In a Fibre Channel environment, channel protection is usually present for any array candidate because when the Storage Subsystem is properly cabled there are two redundant Fibre Channel Arbitrated Loops for each drive.

Controller

A board and firmware, located in the Storage Subsystem, that control the I/O between the logical drives and the host system.

Current Owner

The *current owner* of an array is the controller of an [active-active](#) pair that currently owns the array. If you select an array in the Subsystem Management Window Logical View, the current controller owner of the array is designated with an association dot in the Physical View.

The current owner may not be the [preferred owner](#). Ownership of the array may shift from the preferred controller when it is being replaced or is undergoing a firmware download. This is considered a routine ownership change and is reported with an informational entry in the event log.

There can also be a forced failover from the preferred controller to the other controller because of I/O path errors. This is reported with a critical entry in the event log, and will be reported by the Enterprise Management software to alert e-mail and SNMP destinations.

Use the [Change Controller Ownership](#) option to change both the current and the preferred owner for a selected array.

To shift logical drives away from their current owners and back to their preferred owners, use the [Redistribute Arrays](#) option.

Default Host Group

A standard node in the Topology View of the Storage Partitions Definitions Window that designates all [host groups](#), [hosts](#), and [host ports](#) that do not have any specific logical drive-to-LUN mappings, and share access to any logical drives that were automatically assigned default LUNs by the controller firmware during logical drive creation.


For more information, see [Introducing Storage Partitions](#) and [Using the Storage Partitions Definitions Window](#).

Defragment

The defragment option consolidates non-contiguous space on an array. Non-contiguous space can be created by deleting logical drives or by not using the entire available free capacity in a Free Capacity node during logical drive creation.

For more information, see [Defragmenting an Array](#).

Degraded Status

A mode of operation for a RAID 1, 3, or 5 logical drive when one drive fails in an array. The controller continues to function in the degraded mode because it uses the data and redundancy on the drives that are still operational. This enables you to continue reading and writing to the logical drive without losing data and to reconstruct the data after you replace the failed drive. The Degraded logical drive icon displayed in the Logical View is .

If a hot spare is configured and it takes over for the failed drive, the logical drive remains in an Optimal status and still shows a single failed drive.

Enterprise Management Window

The Enterprise Management Window is Java-based software that runs on a [Management Station](#) and provides six primary storage management functions:

- Discovers [Hosts](#) and [Storage Subsystems](#) automatically on your local sub-network.
- Provides the ability to manually add and remove Hosts and Storage Subsystems in the management domain.
- Monitors the health of Storage Subsystems and reports a high-level status using appropriate icons.
- Provides the ability to configure alert notifications (e-mail or SNMP) and report critical events to the configured alert destinations.
- Loads the appropriate Subsystem Management Window for a selected Storage Subsystem to allow detailed configuration and management operations.
- Executes a script to perform batch management tasks on a particular Storage Subsystem.

A Device Tree in the left pane presents a picture of the management domain, showing discovered and added devices, their status conditions, and their network management connections. A Device Table in the right pane presents more detailed Storage Subsystem data of selected Storage Subsystem devices.

For more information about the features of the Enterprise Management Window, see the Enterprise Management Window online help system.

Environmental Card Canister

A component in a drive enclosure that monitors the components in that enclosure.
Not all Storage Subsystems have environmental cards.

Also known as ESM.

ESM

See [Environmental Card](#).

Event Log


The Event Log is a detailed log of events that occur on the Storage Subsystem. Examples of events that are logged to the Event Log include:

- Configuration events, such as logical drive creation or the assignment of a drive as a hot spare
- Failures of Storage Subsystem components

The Event Log is stored in reserved areas on the disks in the Storage Subsystem.

For more information, see [Using the Event Log](#).

Failed Status

A logical drive status of Failed occurs when two or more drives have failed in the same array, or one or more drives have failed in a RAID 0 array. All the logical drives on the array are no longer operating. This is the most serious status for a logical drive and in a RAID 1, 3, or 5 array. You have lost data unless you erroneously removed the second drive and it was actually Optimal and NOT failed. (In a RAID 0 array, you have lost data.) Use the Recovery Guru to restore your array to Optimal status. A failed logical drive is represented by this icon  in the Subsystem Management Window Logical View.

Physical components such as controllers, drives, power supplies, and so on can also have Failed status, and icons reflecting a status change to Failed are shown in the Subsystem Management Window Physical View. For more information, see [Using the Subsystem Management Window](#).

Fibre Channel

A host interface that is a channel-network hybrid using an active, intelligent interconnection scheme (topology) to connect devices over a serial bus. The storage management software uses this connection between the host machine (where it is installed) and each controller in the Storage Subsystem to communicate with the controllers.

For more information, see [Topology](#).

Firmware

Software for controller operations. In this version of the storage management software, commands are sent to Storage Subsystem controllers and the controller firmware contains most of the logic needed to implement the storage management tasks.

Free Capacity Node

A Free Capacity Node in the Logical View is a contiguous region of unassigned capacity on a defined array. The space in a Free Capacity node can be used to create logical drives.

Host

A computer that is attached to the Storage Subsystem and accesses various logical drives on the Storage Subsystem through its [host ports](#). You can define a host as a separate entity or as part of a host group using the Storage Partitions Definitions Window. You can define specific logical drive-to-LUN mappings to an individual host as well as have the host be part of a host group that shares access to one or more logical drives.

For more information, see [Defining a Host](#).

Host Adapter

A physical board with one or more host ports that resides in the host and attaches it to the SCSI or fibre channel bus. Also known as a Host Bus Adapter (HBA).

Host Group

An optional topological element that you define using the Storage Partitions Definitions Window if you want to designate a collection of [hosts](#) that will share access to the same logical drives.

For more information, see [Defining a Host Group](#).

Host Interface

The connection between a RAID Controller and the host machine, either [Parallel SCSI](#) or [Fibre Channel](#).

Host Port

A physical connection on the [host adapter](#) that resides within a [host](#). When the host adapter only has one physical connection (host port), the terms host port and host adapter are synonymous.

The host ports are automatically discovered by the storage management software. A host port is the actual physical connection that allows a host to gain access to the logical drives in the Storage Subsystem. Therefore, if you want to define specific logical drive-to-LUN mappings for a particular host and create Storage Partitions, you must define its associated host ports.


For more information, see [Defining a Host Port](#).

Hot Spare Drive


A *hot spare* is a drive containing no data that acts as a standby in case a drive fails in a RAID 1, 3, or 5 Logical Drive. The hot spare drive adds another level of redundancy to your Storage Subsystem. If a drive fails, the hot spare takes over for the failed drive until you replace it.

A hot spare drive is not dedicated to a specific array but instead is **global** (can be used for any failed drive in the Storage Subsystem with the same or smaller capacity).

A **Hot Spare Drive - Standby**, represented by this icon in the Subsystem

Management Window Physical View , is a drive that has been assigned as a hot spare but is currently not taking over for any failed drive.

A **Hot Spare Drive - In Use**, represented by this icon in the Subsystem

Management Window Physical View , is a drive that has been assigned as a hot spare and is currently taking over for a failed drive.

I/O Data Connection

The route that data travels from the system host adapter through the interface cable (either [SCSI](#) or [Fibre Channel](#)) to the controller. This might also be referred to as data path.

Initialization

The erasure of all data, on a drive, logical drive, or array. In previous versions of the storage management software, this was called **format**. For more information, see [Initializing a Logical Drive](#).

Logical Drive

A logical drive is a logical structure you create on a Storage Subsystem for data storage. A logical drive is defined over a set of drives called an [array](#), and has a defined [RAID level](#) and capacity. The drive boundaries of the array are hidden to the host computer.

In order to best suit your application needs, you can configure additional logical drive parameters such as [segment size](#), [modification priority](#), and [cache](#) parameters.

Logical Drive-to-LUN Mapping

An association of a logical drive with a single logical unit number (LUN). Use the Storage Partitions Definitions Window to define this mapping and specify which host group or host has access to the logical drive.

For more information, see [Defining a Logical Drive-to-LUN Mapping](#).

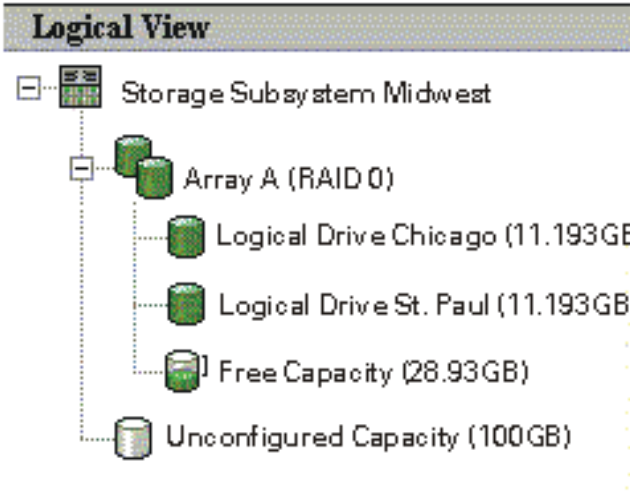
Logical Unit Number (LUN)

The number a host uses to access a logical drive. Each host has its own LUN address space. Therefore, the same LUN may be used by different hosts to access different logical drives. However, a logical drive can only be mapped to a single LUN and a single host group or host.

For example, host 1 may access logical drive Engineering using LUN 5 and host 2 may also use LUN 5 to access logical drive Finance.

Logical View

The Logical View, as shown in the following example, is one of two Subsystem Management Window views. The Logical View provides a tree-structured view of logical nodes. Click on the plus or minus sign beside the node to expand or collapse the view of its child nodes.



The Storage Subsystem, or Root node , has three types of child nodes.

Child Nodes of the Root Node	Description				
Array	An Array Node has two possible types of child nodes, as follows:				
	<table border="1"> <thead> <tr> <th data-bbox="505 1352 824 1472">Child Nodes of the Array Node</th> <th data-bbox="824 1352 1409 1472">Description</th> </tr> </thead> <tbody> <tr> <td data-bbox="505 1472 824 1671">Logical Drive</td> <td data-bbox="824 1472 1409 1671">Represents a configured and defined logical drive. There can be multiple Logical Drive nodes under an Array node.</td> </tr> </tbody> </table>	Child Nodes of the Array Node	Description	Logical Drive	Represents a configured and defined logical drive. There can be multiple Logical Drive nodes under an Array node.
	Child Nodes of the Array Node	Description			
Logical Drive	Represents a configured and defined logical drive. There can be multiple Logical Drive nodes under an Array node.				

	Free Capacity	Represents a region of capacity that can be used for creation of one or more new logical drives within the array. There can be multiple Free Capacity nodes under an Array Node.
<u>Unconfigured Capacity</u>		Represents any Storage Subsystem capacity that has not been configured into an array.

For more information, see [Using the Subsystem Management Window](#).

Management Station

A computer running the storage management software.

Media Scan

A media scan is a background process that runs on all logical drives in the Storage Subsystem for which it has been enabled, providing error detection on the drive media. The advantage to enabling the media scan process is that the process can find media errors before they disrupt normal drive reads and writes. The media scan process scans all logical drive data to verify that it can be accessed, and optionally scans the logical drive [redundancy](#) data, as well.

For more information, see [Changing the Media Scan Duration](#).

Modification Priority

The modification priority defines how much processing time is allocated for logical drive modification operations versus system performance. The higher the priority, the faster logical drive modification operations complete, but the slower system I/O is serviced.

Logical Drive modification operations include reconstruction, copyback, initialization, media scan, defragmentation, change of RAID level and change of segment size.

The modification priority is set for each individual logical drive using a slider bar on the **Logical Drive>>Properties** dialog. There are five relative settings on the reconstruction rate slider bar ranging from Lowest to Highest. The actual speed of each setting is determined by the controller.

Note: The modification priority is changed immediately when you move the slider bar, although you may notice a slight delay if you have many or large logical drives reconstructing. For more information, see [Specifying Logical Drive Properties](#).

NL Port

A Fibre Channel controllers ID in a hexadecimal value; its hex format varies depending on the topology:

- For Private Arbitrated Loop, it is a one-byte Arbitrated Loop Physical Address (AL_PA).
- For all other Arbitrated Loops, it is displayed as a single 24-bit hex number (a triplet of Domain, Area, and AL_PA where each field is one byte).
- For Fabric and Point-to-Point, it is a three-byte hex value used in the D_ID and S_ID (Destination/Source Identifier) fields of Fibre Channel frames.


NVSRAM

NVSRAM (Non-Volatile Static Random Access Memory) is a controller file that specifies default settings for the controller.

Operation in Progress Status

This is a logical drive status that reflects a modification operation in progress. Modification operations include:

- Reconstruction of a drive that is part of the array
- Copyback to a drive that is part of the array
- Initialization of the logical drive
- Change of array RAID Level
- Capacity expansion of the array
- Defragmentation of the array
- Change of segment size

The logical drive is represented with this icon  in the Subsystem Management Window Logical View.

Optimal Status

Status indicating the component is operating in its desired working condition.

Parallel SCSI

A host interface that uses SCSI cables over SCSI buses to connect devices over a parallel bus. The storage management software uses the SCSI connection between the host machine (where it is installed) and each controller in the Storage Subsystem to communicate with the controllers.

For more information, see [SCSI ID](#) and [SCSI Channel](#).

Passive Mode

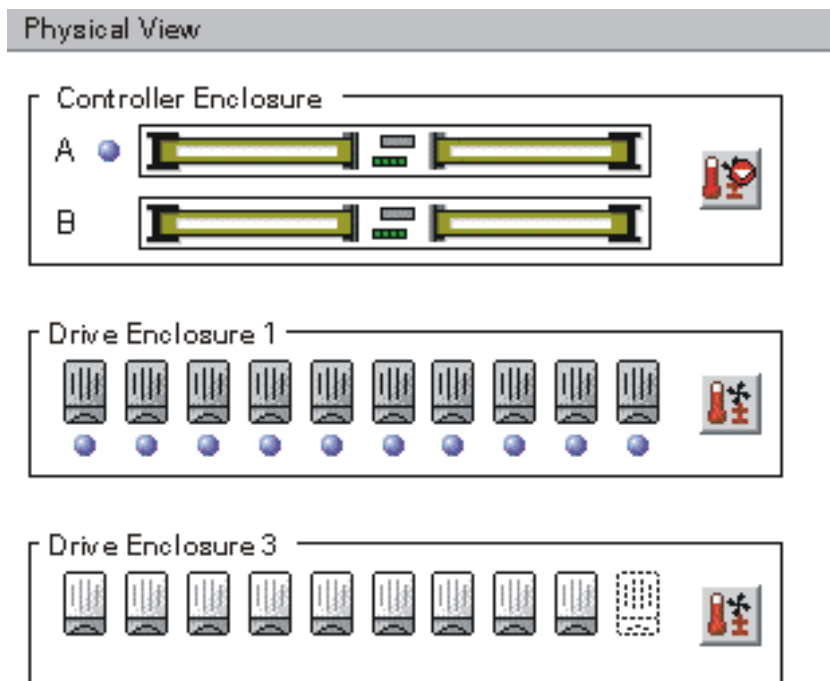
A controller mode indicating that the controller is acting as a spare. If the active controller fails, the passive controller becomes active and takes over the logical drives owned by the failed controller.

For more information, see [Changing a Controller Mode](#).

Physical View

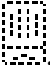
The Physical View is one of two Subsystem Management Window views. The Physical View:

- Provides a view of the physical components in a Storage Subsystem, including their status
- Shows the physical components associated with a selected node in the Logical View



Note: The orientation of the Physical View is determined by actual layout of the controller. For example, if the controller has horizontal drive enclosures, the software will show horizontal drive enclosures in the Physical View.

The Physical View displays four enclosure types:

Enclosure Type	Description
Controller Enclosure (shown above)	<p>Displays either one or two controller slots and a controller enclosure components button.</p> <p>Note: You can select each controller in the controller enclosure to perform controller management operations.</p>
Drive Enclosure (shown above)	<p>Displays the drives or drive slots in the enclosure and one components button. An empty slot is shown with this icon .</p> <p>Note: You can select each drive in the drive enclosure to perform drive management operations.</p>
Combination Enclosure	<p>Contains a representation of both controllers and drives in the same enclosure, and an enclosure components button.</p>
Generic Enclosure	<p>Shown as a number of drive slots, when information about the enclosure type is not received from the controller.</p>

To change the enclosure order in the Physical View, use the **Storage Subsystem >> Change Enclosure Order** option. See [Changing the Enclosure Order](#) for more information.

Association

In the example above:

- The blue association dot shown beside Controller A in the Controller Enclosure identifies the current owner of a selected logical drive in the Logical View.
- The blue association dots underneath the assigned drives in Drive Enclosure 1 identify the drives associated with a selected logical drive in the Logical View.

The Components Button

The Components button on each displayed enclosure shows the status of secondary

components within the enclosure. For more information, see [Viewing the Status of Controller Enclosure Components](#) and [Viewing the Status of Drive Enclosure Components](#).

More Information

For more information, see [Using the Subsystem Management Window](#).

Polling Interval

The polling interval is the time frequency that the Performance Monitor uses to check I/O activity on the Storage Subsystem.

Port

The hardware within a node that performs data communications over the Fibre Channel link.

Preferred Address (AL_PA)

The hard address (AL_PA format) that Fibre Channel controllers attempt to acquire. Displayed only for controllers using the Arbitrated Loop topology. If the hard address is not available, the controller automatically obtains the first address available after the other ports on the loop have obtained their hard addresses.

For more information, see [AL_PA](#) and [Preferred Loop ID](#).

Preferred Loop ID

The Fibre Channel Loop ID in NVSRAM used to determine the Preferred AL_PA (hard address). Displayed only for controllers attached using the Arbitrated Loop topology.

This ID is NOT the AL_PA, but an index into a table of valid AL_PA values (see Annex K of the FC-AL standard). Loop IDs are assigned an arbitration priority with (host-side IDs) 0 being the lowest and 126 being the highest.

Preferred Owner

The *preferred owner* of an array is the controller of an [active-active](#) pair that is designated to own these logical drives. The *current owner* is the controller that currently owns the array.

If the preferred controller is being replaced or undergoing a firmware download, ownership of the logical drives is automatically shifted to the other controller, and that controller becomes the current owner of the array. This is considered a routine ownership change and is reported with an informational entry in the event log.

There can also be a forced failover from the preferred controller to the other controller because of I/O path errors. This is reported with a critical entry in the event log, and will be reported by the Enterprise Management software to e-mail and SNMP alert destinations.

The preferred owner for an array is initially selected by the controller when the array is created. Use the [Change Controller Ownership](#) option to change the preferred owner for a selected array.

To shift logical drives away from their current owners and back to their preferred owners, use the [Redistribute Arrays](#) option.

Premium Feature

A Storage Subsystem feature that may not be available in the standard configuration of the storage management software.

Use the **Storage Subsystem >> Premium Features >> Enable** option to enable a premium feature. See [Enabling Premium Features](#).

RAID

RAID is an acronym for Redundant Array of Independent Disks. It describes a storage solution in which part of the storage capacity is used to store redundant information about user data stored on the remainder of the storage capacity. The redundant information enables regeneration of user data if one of the disk drives in the array fails.

RAID relies on a series of **RAID levels** to determine how drives connect and how the controller reads and writes data and redundancy on the drives. You can create RAID Level 0, 1, 3, and 5 logical drives using this storage management software.

RAID 0

Non-redundant RAID Level where data, without redundancy, is striped across an array/logical drive. All drives are available for storing user data. Data availability is more at risk than with other RAID levels, because any single drive failure causes data loss and a logical drive status of Failed. However, because RAID 0 does not require data redundancy to be written to the media, it provides exceptionally high performance and is sometimes used for performance reasons.

RAID 1

Redundant RAID Level where identical copies of data are maintained on drive pairs, also known as mirrored pairs. Half of the drives are available for storing user data. Drive pair failure causes data loss. The strength of this RAID level is high safety and data availability.

RAID 3

Redundant RAID Level where user data and redundant information (parity) is striped across the drives in an array. The equivalent of one drive's worth of capacity is used for redundant information. This RAID level is best used for applications such as multimedia or medical imaging that write and read large sequential chunks of data. If a single drive fails in a RAID 3 array, all associated logical drives become degraded but the redundant information allows the data to still be accessed. Any two-drive failure in the same array causes associated logical drives to fail and data loss.

RAID 5

Redundant RAID Level where data and redundancy are striped across an array/logical drive. Best used for small/medium, random I/Os. Any two-drive failure in same array causes data loss.

RDAC Multi-path Driver

The RDAC multi-path driver manages the I/O data connection for Storage Subsystems with redundant controllers. If a component (cable, controller, host adapter, and so on.) fails along the I/O data connection, the RDAC Multi-path Driver automatically reroutes all I/O operations to the other controller.

Consult the "Installation and Support Guide" for your operating system for specific information.

Reconstruction

The process of using data and redundancy on all operational drives within an array to regenerate the data to a replacement drive or a hot spare drive. Only data on a RAID 1, 3, or 5 logical drive can be reconstructed.

Redundancy

Redundancy is additional information stored along with the data that allows the controller to reconstruct lost data.

RAID 1 uses mirroring for redundancy. RAID 3 and RAID 5 use redundancy information, sometimes called parity, that is constructed from the data bytes and striped along with the data on each disk.

A redundancy check, performed using the **Array >> Check Redundancy** option:


- Scans the blocks in a RAID 3 or 5 logical drive and checks the redundancy information for each block.

-OR-

- Compares data blocks on RAID 1 mirrored drives.

For more information, see [Performing a Redundancy Check](#).

Replaced Drive Status

A temporary status for a drive that has just replaced a failed drive, but has not yet reconstructed. The drive status is represented in the Physical View by this icon .

SCSI Channel

A number used to identify the path for the transfer of data between a disk drive and a RAID Controller.

SCSI ID

A number used to identify the drive on the SCSI Bus between the controller and drives (Drive-side SCSI Bus).

Segment Size

A segment is the amount of data, in kilobytes, that the controller writes on a single drive in a logical drive before writing data on the next drive. Data blocks store 512 bytes of data and are the smallest units of storage. The size of a segment determines how many data blocks it contains. For example, an 8K segment holds 16 data blocks, and a 64K segment holds 128 data blocks.

Important: The segment size was expressed in number of data blocks in previous versions of this storage management software. It is now expressed in Kilobytes.

When you create a logical drive, the default segment size is a good choice for the expected logical drive usage. The default segment size can be changed using the **Logical Drive>>Change Segment Size** option. For more information, see [Changing Segment Size](#).

Storage Partition

A logical entity consisting of one or more Storage Subsystem logical drives that can be shared among hosts that are part of a [host group](#) or accessed by a single [host](#). A Storage Partition is created when you define a collection of hosts (a host group) or a single host and then define a logical drive-to-logical unit number (LUN) mapping. This mapping allows you to define what host group or host will have access to a particular logical drive in your Storage Subsystem.

For more information, see [Introducing Storage Partitions](#).

Storage Subsystem

A **Storage Subsystem** is a storage entity managed by the storage management software. A Storage Subsystem consists of a collection of both physical components (such as drives, controllers, fans, and power supplies) and logical components (such as arrays and logical drives). A Storage Subsystem can span multiple physical enclosures.

Storage Subsystem replaces the term **RAID Module** that was used in the previous versions of this storage management software.

Subsystem Management Window

From the **Subsystem Management Window**, you can manage a Storage Subsystem in the management domain.

Use this window to:

- Manage physical components, such as controllers and drives
- Configure and manage logical components
- Recover from errors
- Monitor performance

For more information, see [Introduction to the Subsystem Management Window](#) and [Using the Subsystem Management Window](#).

Topology

An interconnected set of nodes.

A logical topology, in the context of Storage Partitions definitions, is a collection of [Default Host Group](#), [Host Group](#), [Host](#) and [Host Port](#) nodes in the Topology View. For more information, see [Introducing Storage Partitions](#).


A physical topology, for Storage Subsystems with fibre channel host interfaces, is the interconnect method between a Storage Subsystem controller and a host. The **Storage Subsystem Profile** option displays one of the following values in the Topology field:

Topology	Description
Point-to-Point	The controller port and host port are directly connected as N-ports (non-loop).
Arbitrated Loop - Public	The controller port is connected to a switch port via a fibre channel arbitrated loop. The host port may be on the local loop or connected to the switch. Devices on the loop and devices attached to the switch can communicate with each other.
Arbitrated Loop - Private	The controller port and host port are connected by a fibre channel arbitrated loop.
Fabric Attach	The controller port is connected as an N-port to the fabric (switch) port. The host port is connected somewhere else in the fabric. The switch facilitates communication between the controller port and the host port.

For more information, see [Viewing a Storage Subsystem Profile](#).

Unassigned Drive

An unassigned drive is not yet used in a Storage Subsystem array, or assigned as a hot spare. An optimal, unassigned drive is represented in the Subsystem

Management Window Physical View with this icon .


Unconfigured Capacity Node

Unconfigured capacity is capacity present in the Storage Subsystem from drives that have not been assigned to an array. The Unconfigured Capacity Node in the Subsystem Management Window Logical View can be used to create new arrays.

Uninterruptible Power Supply (UPS)

A UPS is a battery-operated power supply connected to the Storage Subsystem to keep it running during a power failure.

Unresponsive Storage Subsystem Status

An Unresponsive Storage Subsystem status occurs when the management station cannot communicate with the only controller or both controllers in a Storage Subsystem. An Unresponsive status is represented by this icon  in the Device Tree and Device Table views of the Enterprise Management Window. You will be unable to launch an Subsystem Management Window for a Storage Subsystem that is unresponsive.

For information about correcting this problem, see the Enterprise Management Window online help system.

World Wide Name

A unique hexadecimal identifier for a peripheral device.

World Wide Node Name

An 8-byte hexadecimal value used to uniquely identify a Fibre Channel controller.

World Wide Port Name

An 8-byte hexadecimal value used to uniquely identify a Fibre Channel controller port.


Write Cache Mirroring

Allows cached data to be mirrored across two redundant controllers that have the same cache size. The data written to the cache memory of one controller is also written to the cache memory of the other controller. Therefore, if one controller fails, the other can complete all outstanding write operations.


Enable Write Cache Mirroring for each logical drive using the **Logical Drive >> Properties** dialog. For more information, see [Specifying Logical Drive Properties](#).

Optimal Status

An Optimal Status indicates every component in the Storage Subsystem is in the desired working condition.

A Storage Subsystem has an Optimal status, represented by this icon  in the Subsystem Management Window Logical View and in the Enterprise Management Window, to indicate that the Storage Subsystem is Optimal.


Fixing Status

This temporary Storage Subsystem status, represented by this icon  in the Subsystem Management Window Logical View and the Enterprise Management Window, displays when a Needs Attention condition has been corrected and the Storage Subsystem is currently transitioning to an Optimal state (for example, a reconstruction operation is in progress). A Fixing status requires no action other than if you want to check on the progress of the operation in the Subsystem Management Window.

Note: Some recovery actions cause the Storage Subsystem state to change directly from Needs Attention to Optimal, without an interim state of Fixing.

Needs Attention Status


A Needs Attention Status indicates a problem on a Storage Subsystem that requires your intervention to correct it. To correct the problem, you should launch the Subsystem Management Window for the particular Storage Subsystem and then use the Recovery Guru to pinpoint the cause of the problem and obtain appropriate instructions.

A Storage Subsystem has a Needs Attention status, represented by this icon  in the Subsystem Management Window Logical View and in the Enterprise Management Window, to represent a coarse-level indication that one or more of its components are in a Needs Attention state. A Needs Attention icon is also used in the Physical View of the Subsystem Management Window on controllers, drives, or the components button to pinpoint the exact component having the problem.

Simple Network Management Protocol (SNMP)

A standard network management protocol used by Network Management Stations to manage devices. If you have configured [SNMP trap messages](#) to be sent for Storage Subsystem critical events using the **Alert >> Destinations** option, the Enterprise Management software sends remote notification of these critical events to the network management stations you designated.

Critical Event Descriptions

The following reference provides more information about events with a critical priority, shown with this icon  in the Event Log.

- The Critical Event Type is shown in critical event messages that are forwarded to configured e-mail or SNMP destinations.
- The Critical Event Type and the Sense Key/ASC/ASCQ data are both shown in the Event Log Details.

Critical Event Type and Sense Key/ASC/ASCQ	Description
Event 1001 - Channel failed Sense key/ASC/ASCQ: 6/3F/C3	The controller failed a channel, and will not access drives on this channel any more. The FRU Group Qualifier (byte 26) in the sense data will indicate the 1-relative channel number of the failed channel. This condition is typically caused by a drive ignoring SCSI protocol on one of the controller's destination channels. The controller typically fails a channel if it issued a reset on a channel, and it continued to see drives ignore the SCSI Bus Reset on this channel. Select the Recovery Guru to obtain the recovery procedure.
Event 1010 - Impending drive failure (PFA) detected Sense key/ASC/ASCQ: 6/5D/80	A drive has reported that a failure prediction threshold has been exceeded. This indicates that the drive may fail within 24 hours. Select the Recovery Guru to obtain the recovery procedure.

<p>Event 1015 - Incorrect mode parameters set on drive Sense key/ASC/ASCQ: 6/3F/BD</p>	<p>The controller was unable to query the drive for its current critical mode page settings, or was unable to change these to the correct setting. Currently, this indicates the Qerr bit is set incorrectly on the drive specified in the FRU field of the Request Sense data.</p>
<p>Event 202E - Read drive error during interrupted write Sense key/ASC/ASCQ: 3/11/8A</p>	<p>A media error has occurred on a read operation during interrupted write processing.</p>
<p>Event 2109 - Controller cache not enabled - cache sizes do not match Sense key/ASC/ASCQ: 6/A1/00</p>	<p>The controller will not allow mirroring to be enabled if the alternate controller's cache size is different. Make sure that both controllers have the same cache size.</p>
<p>Event 210B - Cache between controllers not synchronized Sense key/ASC/ASCQ: 6/2A/01</p>	<p>The Mode Select parameters changed because of a cache synchronization error during the processing of the most recent Mode Select request. This normally occurs if the controller's alternate pair has not completed its start-of-day routine. If this message persists, contact your Customer Support Representative.</p>
<p>Event 210C - Controller cache battery failed Sense key/ASC/ASCQ: 6/0C/80</p>	<p>The controller has detected that the battery is (1) not physically present, (2) it is fully discharged, or (3) it has reached its expiration date. Select the Recovery Guru to obtain the recovery procedure.</p>
<p>Event 210E - Controller cache memory recovery failed after power cycle or reset Sense key/ASC/ASCQ: 6/0C/81</p>	<p>Recovery from a Data Cache error was unsuccessful. User data may have been lost.</p>
<p>Event 2110 - Controller cache memory initialization failed Sense key/ASC/ASCQ: 6/40/81</p>	<p>The controller has detected the failure of an internal controller component (RAID Buffer). This failure may have been detected during operation as well as during an on-board diagnostic routine.</p>

<p>Event 2113 - Controller cache battery nearing expiration Sense key/ASC/ASCQ: 6/3F/D9</p>	<p>The cache battery is within the specified number of weeks of failing. Select the Recovery Guru to obtain the recovery procedure.</p>
<p>Event 2229 - Drive failed by controller Sense key/ASC/ASCQ: None</p>	<p>The controller has failed a drive because of a problem with the drive. Select the Recovery Guru to obtain the recovery procedure.</p>
<p>Event 222D - Drive manually failed Sense key/ASC/ASCQ: 6/3F/87</p>	<p>The drive was manually failed by a user. Select the Recovery Guru to obtain the recovery procedure.</p>
<p>Event 2247 - Data lost on logical drive during unrecovered interrupted write Sense key/ASC/ASCQ: 6/3F/EB</p>	<p>An error has occurred during interrupted write processing during the start-of-day routine causing the logical drive to transition to the failed state. Select the Recovery Guru to obtain the recovery procedure.</p>
<p>Event 2248 - Drive failed - write failure Sense key/ASC/ASCQ: 6/3F/80</p>	<p>The drive failed a write command to it. The drive will be marked failed. Select the Recovery Guru to obtain the recovery procedure.</p>
<p>Event 2249 - Drive capacity less than minimum Sense key/ASC/ASCQ: 6/3F/8B</p>	<p>The capacity of the drive is not large enough to support all the logical drives that must be reconstructed on it. You should replace the drive with a larger capacity drive.</p>
<p>Event 224A - Drive has wrong block size Sense key/ASC/ASCQ: 6/3F/8C</p>	<p>The drive's block size does not match that of the other drives in the logical drive. The drive will be marked failed. Select the Recovery Guru to obtain the recovery procedure.</p>
<p>Event 224B - Drive failed - initialization failure Sense key/ASC/ASCQ: 6/3F/86</p>	<p>The drive failed either a Format Unit command, or a Write operation (issued when a logical drive was initialized). The drive will be marked failed. Select the Recovery Guru to obtain the recovery procedure.</p>

<p>Event 224D - Drive failed - no response at start of day Sense key/ASC/ASCQ: 6/3F/85</p>	<p>The drive failed a Read Capacity or Read command during the start-of-day routine. The controller was unable to read the configuration information stored on it. Select the Recovery Guru to obtain the recovery procedure.</p>
<p>Event 224E - Drive failed - initialization/reconstruction failure Sense key/ASC/ASCQ: 6/3F/82</p>	<p>The (previously-failed) drive was marked failed because either (1) the drive failed a Format Unit command issued to it, or (2) the reconstruction on the drive failed due to the controller being unable to restore it (for example, an error occurring on another drive required for reconstruction). Select the Recovery Guru to obtain the recovery procedure.</p>
<p>Event 2250 - Logical Drive failure (3F E0) Sense key/ASC/ASCQ: 6/3F/E0</p>	<p>The controller has marked the logical drive failed. User data and/or redundancy (parity) can no longer be maintained to ensure availability. The most likely cause is the failure of a single drive in non-redundant configurations or a second drive in a configuration protected by one drive. Select the Recovery Guru to obtain the recovery procedure.</p>
<p>Event 2251 - Drive failed - reconstruction failure Sense key/ASC/ASCQ: 6/3F/8E</p>	<p>A drive failed due to a reconstruction failure during the start-of-day routine. Select the Recovery Guru to obtain the recovery procedure.</p>
<p>Event 2252 - Drive marked offline during interrupted write Sense key/ASC/ASCQ: 6/3F/98</p>	<p>An error has occurred during interrupted write processing causing the logical drive to be marked failed. Drives in the array that did not experience the read error will transition to the Offline state and log this error.</p>

<p>Event 2254 - Redundancy (parity) and data mismatch was detected Sense key/ASC/ASCQ: 6/8E/01</p>	<p>The controller detected inconsistent redundancy (parity)/data during a parity verification.</p>
<p>Event 2255 - Logical Drive definition incompatible with ALT mode - ALT disabled Sense key/ASC/ASCQ: 6/91/3B</p>	<p>The controller will operate in normal redundant controller mode without performing Auto-LUN transfers.</p>
<p>Event 2602 - Automatic controller firmware synchronization failed Sense key/ASC/ASCQ: 02/04/81</p>	<p>The versions of firmware on the redundant controllers are not the same because the automatic controller firmware synchronization failed. Controllers with an incompatible version of firmware may cause unexpected results. Try the firmware download again. If the problem persists, contact your Customer Support Representative.</p>
<p>Event 2801 - Storage Subsystem running on UPS battery Sense key/ASC/ASCQ: 6/3F/C8</p>	<p>The Uninterruptible Power Source (UPS) has indicated that AC power is no longer present and the UPS has switched to standby power. While there is no immediate cause for concern, you should save your data frequently, in case the battery is suddenly depleted.</p>
<p>Event 2803 - UPS battery - two minutes to failure Sense key/ASC/ASCQ: 6/3F/C9</p>	<p>The Uninterruptible Power Source (UPS) has indicated that its standby power source is nearing depletion. You should take actions to stop I/O activity to the controller. Normally, the controller will change from a write-back caching mode to a write-through mode.</p>
<p>Event 2804 - UPS battery failed Sense key/ASC/ASCQ: None</p>	<p>The UPS battery has failed.</p>
<p>Event 2807 - Environmental card failed Sense key/ASC/ASCQ: None</p>	<p>An environmental card has failed. Select the Recovery Guru to obtain the recovery procedure.</p>

<p>Event 2808 - Enclosure ID not unique Sense key/ASC/ASCQ: 6/98/01</p>	<p>The controller has determined that there are multiple drive enclosures with the same ID selected. Make sure that each drive enclosure has a unique ID setting.</p>
<p>Event 280A - Controller enclosure component missing Sense key/ASC/ASCQ: 6/3F/C7</p>	<p>A component other than a controller is missing in the controller enclosure (for example, fan, power supply, or battery). The FRU codes indicate the faulty component. Select the Recovery Guru to obtain the recovery procedure.</p>
<p>Event 280B - Controller enclosure component failed Sense key/ASC/ASCQ: 6/3F/C7</p>	<p>A component other than a controller has failed in the controller enclosure (for example, fan, power supply, battery), or an over-temperature condition has occurred. The FRU codes indicate the faulty component. Select the Recovery Guru to obtain the recovery procedure.</p>
<p>Event 280D - Drive enclosure component failed Sense key/ASC/ASCQ: 6/3F/C7</p>	<p>A component other than a drive has failed in the drive enclosure (for example, fan, power supply, battery), or an over-temperature condition has occurred. The FRU codes indicate the faulty component. Select the Recovery Guru to obtain the recovery procedure.</p>
<p>Event 280E - Standby power source not fully charged Sense key/ASC/ASCQ: 6/3F/CA</p>	<p>The Uninterruptible Power Source (UPS) has indicated that its standby power source is not at full capacity.</p>
<p>Event 3019 - Logical Drive ownership changed due to failover Sense key/ASC/ASCQ: None</p>	<p>The multi-path driver software has changed ownership of the logical drives to the other controller because it could not access the logical drives on the particular path. Select the Recovery Guru to obtain the recovery procedure.</p>

Specifying Performance Monitor Settings

Use this option to select logical drives and controllers to monitor or to change the polling interval.

Important: Chosen settings apply to this monitoring session only (that is, the settings will revert back to the default settings if you exit and re-open the Performance Monitor.)

Selecting Logical Drives and Controllers to Monitor

To select logical drives and controllers to monitor, press the SHIFT+CLICK or CTRL+CLICK to select multiple contiguous or non-contiguous devices in the list, respectively. To select every device, press **Select All**.

Note: The **Storage Subsystem Totals** row can be selected or unselected for display on the **Settings** dialog.

Changing a Polling Interval

To change the polling interval, choose a number of seconds in the spin box. Each time the polling interval elapses, the Performance Monitor re-queries the Storage Subsystem and updates the statistics in the table.

If you're monitoring the Storage Subsystem in real time, update the statistics frequently by selecting a short polling interval, for example, 5 seconds.

If you're saving results to a file to look at later, choose a slightly longer interval, for example 30 to 60 seconds, to decrease the system overhead and the performance impact.

Applying Performance Monitor Settings

To apply the new settings, select **OK**.

Feature Key File

The Storage Subsystem controller uses this file to enable an authorized premium feature. The file contains the [Feature Enable Identifier](#) of the Storage Subsystem for which the premium feature is authorized and generated data about the premium feature to be enabled.

Feature Enable Identifier

A unique identifier for the Storage Subsystem, used in the process of generating a [Feature Key file](#).

SNMP Trap Message

If you have configured trap messages to report predefined critical events that have occurred on a Storage Subsystem, the Management Station sends these alerts to specified Host destinations running an SNMP service (usually a Network Management Station).

For information on configuring a host destination to receive SNMP trap messages, see [Configuring SNMP and E-mail Alert Notification Settings](#).