Document Name:	Spheras Storage Director Installation and User Guide

Part Number	MAN-00005-UG
Revision	1.0

Revision History

Rev	Approved	Date	Change Description	Reviewed By
1.0	ECO-3679	Sep., 2003	Released with SPHSSD 2.1	CCB

Table 1 Revision History

Contents

Figures

Preface

Proprietary Rights Notice

This document contains proprietary information of Adaptec, Inc. The information contained herein is not to be used by or disclosed to third parties without the prior express written permission of an officer of the Adaptec, Inc. Any product or products described herein is/are a licensed product of Adaptec, Inc.

Document Description

Document GEN-00005-UG, First Edition, September 2003

This document describes the SANbloc 2500, Spheras Storage Director. This document remains the official reference source for all revisions/releases of this product until rescinded by an update.

Disclaimer

It is the policy of Adaptec, Inc. to improve products as new technology, components, software and firmware become available. Adaptec, Inc. reserves the right to make changes to any products herein at any time without notice. All features, functions, and operations described herein may not be marketed by Adaptec, Inc. in all parts of the world. In some instances, photographs and figures are of equipment prototypes. Therefore, before using this document, consult your Adaptec, Inc. representative for information that is applicable and current. ADAPTEC, INC. DOES NOT ASSUME ANY RESPONSIBILITY OR LIABILITY FOR THE USE OF ANY PRODUCT OR PRODUCTS DESCRIBED HEREIN EXCEPT AS EXPRESSLY AGREED TO IN WRITING BY ADAPTEC, INC.

License Restrictions

The purchase or use of an Adaptec, Inc. product does not convey a license under any patent, copyright, trademark, or other intellectual property right of Adaptec, Inc. or third parties.

This product includes code licensed from RSA Security, Inc. Some portions licensed from IBM are available at http://oss.software.ibm.com/icu4j/

Trademark Acknowledgements

Adaptec, Inc. the Adaptec, Inc. Logo, and SANbloc are trademarks or registered trademarks of Adaptec, Inc. All other brand and product names may be trademarks of their respective companies.

Copyright Notice

Chapter 1. Introduction

The Spheras Storage Director™ (SSD) is a browser-based, remote configuration and management tool supporting external RAID controllers. The Spheras Storage Director console uses Java™ technology via an applet running in a supported World Wide Web browser environment.

1.1 About this Manual

This installation and user guide covers the steps involved to install and use the Spheras Storage Director Server and Client to configure, monitor, manage and maintain disk array controllers and the physical and logical drives that are connected to these controllers and the associated enclosures.

1.2 Conventions

Throughout this installation and user guide, the following conventions are used to describe possible effects on performance within the Spheras Storage Director console:

Note:	Supplementary information that can have an effect on system performance
Caution:	Notification that a proscribed action has the potential to adversely affect equipment operation, system performance, or data integrity
Warning:	Notification that a proscribed action will definitely result in equipment damage, data loss, or personal injury

1.3 Spheras Storage Director Components and Functions

Spheras Storage Director has two components. Each component handles specific tasks that are based on a selected function:

Spheras Storage Director Server

Spheras Storage Director Client

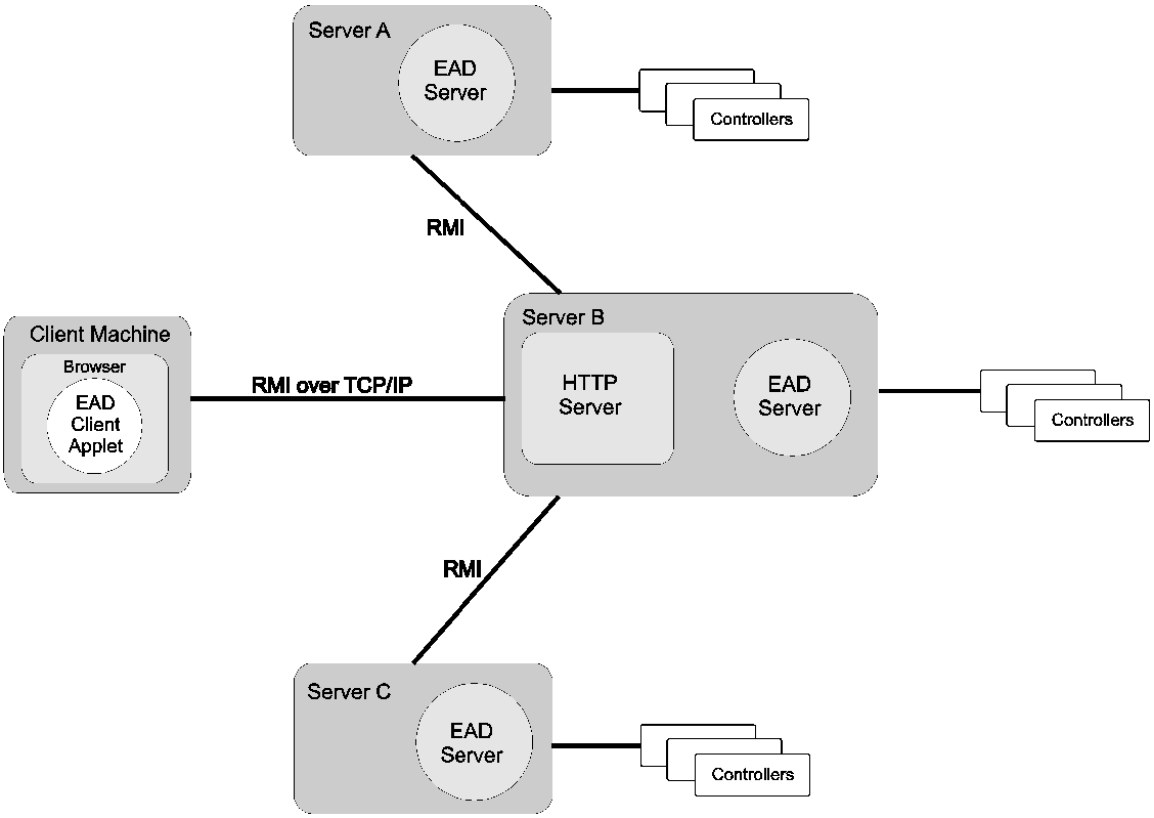

Figure 1 Spheras Storage Director Client and Server Relationship

This diagram shows how the client talks with the server. Spheras Storage Director uses Java RMI (Remote Method Invocation) as Internet connectivity paths that are used when the client talks with the server that launched it (Server B). It also shows how Server B will talk to the other servers in response to the user interacting with the other servers via the client machine.

1.1.1 Configuration Functions

Configuring a disk array is easily performed using the Configure Unused Storage with Assistance wizard. This wizard simplifies the process of setting up or reconfiguring a disk array. Follow the instructions as presented on the dialogs, answer a few questions, and the Configure Unused Storage with Assistance wizard does the rest. Use the Configure Unused Storage Manually Wizard for more control over how arrays are defined.

1.1.2 Monitoring Functions

The Spheras Storage Director Server component collects and disseminates information on resource utilization and array status. The Spheras Storage Director Client organizes this information through a graphical display. Errors and events are recorded in a log file and in the event content table. If a problem is serious enough to warrant immediate attention, operators can be notified by the Spheras Storage Director by using applications such as pagers and email.

1.1.3 Maintenance Functions

The Spheras Storage Director Client manages or performs maintenance on individual disk arrays and drives (with appropriate authentication), again through use of the graphical user

interface. This includes removing physical devices from operation in a functioning disk array (also known as "stopping" or off-lining a drive), rebuilding drives, selecting hot spares, and initiating a consistency check (or parity) on arrays that are configured to support redundancy. The Spheras Storage Director Server executes the management instructions specified by the Spheras Storage Director Client.

1.4 Spheras Storage Director Requirements

The Spheras Storage Director architecture is a client/server relationship. You can use a remote client to manage several servers. Hardware requirements and software requirements for the installation and operation of the Spheras Storage Director are described in the following sections.

Refer to your workstation documentation, server documentation, and to the operating system documentation for more information on hardware and operating system requirements.

1.1.4 Server Requirements

Spheras Storage Director Server software component (included on CD-ROM)

A PC with a 500MHz Pentium® III (or higher) class processor, at least 256 MB of system memory, a CD-ROM drive, and a minimum of 200MB of available space on a fixed disk.

One of the following operating systems:

Windows NT®4.0 with service pack 6a or later

Windows® 2000 SR3 or later

Windows XP SR1 or later

Red Hat Linux 7.2 with kernel 2.4.7-10 or Red Had Linux 8.0 with kernel 2.4.18-4 on an x86 platform

Solaris 8 or 9 with service packs on a SPARC platform

Either of the following HTTP servers: Apache version 2.0.x or Microsoft® IIS version 5.5 or higher.

1.1.5 Client Requirements

PC with a 500MHz Pentium III (or higher) class processor, at least 512 MB of system memory, and a minimum of 256 MB of available space on a fixed disk.

Browser requirements:

Windows: Netscape 6.0 or Internet Explorer 5.5 or higher with Java script and Java enabled.

Linux and Solaris: Only Netscape version 4.78 or 4.8 with Java script and Java enabled. To download 4.x versions of Netscape, see: http://wp.netscape.com/download/archive/index.html

100Base-T or better network data transfer rate.

A minimum display resolution of 800 x 600 pixels is recommended.

A display color setting of at least 65535 colors.

One of the operating systems listed in the server section above.

1.1.6 Using Host Bus Adapters and Drivers

A Host Bus Adapter (HBA) connects a host system to either a SCSI bus or a fibre loop. This device usually performs the lower layers of the SCSI or Fibre Channel protocol and normally operates in the initiator role. One or more HBAs may be integrated into a host/server computer, cluster of host computers, or SAN.

For current information on the latest HBAs and compatible drivers, see the following web site:

http://www.adaptec.com

1.1.7 Setting Up TCP/IP

The Transmission Control Protocol/Internet Protocol (TCP/IP) must be functioning properly to achieve optimal Spheras Storage Director performance. Set up your network hardware configuration and ensure that your TCP/IP connections are working. You can perform a test by pinging the IP addresses of the computers in your hardware configuration.

1.1.8 Using Controllers and Subsystems

The Spheras Storage Director supports 2 types of controller hardware configurations.

Simplex - A disk array system having only one controller or a system capable of having dual-active controllers but only one controller is present or active. A simplex system does not provide fault tolerance in the event of a controller failure. The single controller handles all I/O requests.

Dual-active - A controller configuration that interconnects two identical controllers that share a common set of disk drives. In addition to increasing overall performance, this method allows a surviving controller to take over resources of a failed controller. This failover process is transparent to the host.

The following is a list of supported subsystems:

SANbloc 2500

See the appropriate installation guide for your subsystem for detailed configuration information.

Spheras Storage Director may be used in a direct-connect loop or a fabric environment. A switch is a device used to connect a subsystem to an HBA. A switch is a Fibre Channel device that can be used to connect multiple ports on a Fibre Channel arbitrated loop or Fibre Channel fabric topology. (Fabric topology uses multiple switches to route frames between nodes in a Fibre Channel network.)

A fibre channel loop is a connection between multiple nodes that are interconnected and form a closed loop. Information is directed around the loop until the final destination is reached. The information is routed through and repeated by intermediate ports on the loop. These node ports are known as node loop ports (NL_Ports). All ports on the loop share the loop. Access to the loop is obtained through an arbitration process.

Chapter 2. Installation and Startup

This section assumes that the network administrator for your site will be performing these installation procedures.

Before installing the Spheras Storage Director Server and Client, read 1.4 Spheras Storage Director Requirements section. This checklist provides you with important hardware, software, and performance-related suggestions.

If you want to update your Spheras Storage Director software, go to 1.11 Updating Your Spheras Storage Director Server and Client.

1.5 Pre-Installation requirements

Spheras Storage Director Server Requirement

You must install one of the supported HTTP servers before installing the Spheras Storage Director Server. (For a list of the supported HTTP servers, see 1.1.4 Server Requirements on page 15.) An HTTP server is used for a remote system to download and execute the Spheras Storage Director Client applet. The HTTP server allows a Spheras Storage Director Server to be remotely accessed by various client computers.

Note: On Solaris systems, the Apache HTTP server is included on the installation CDs and is commonly installed when the Solaris operating system is installed.

It is outside the scope of this document to describe the installation of an HTTP server. See your HTTP server documentation for installation instructions. See Figure 1 on page 1 for more information about the relationship between the server and client.

2.1.1 Spheras Storage Director Client Requirement

Before installing the Spheras Storage Director Client, ensure that a web browser is installed on the client system. For a list of supported browsers, see 1.1.5 Client Requirements on page 15. Launch the browser to ensure that it is working. If not, consult your browser installation manual for troubleshooting information.

2.1.2 General Recommendations and Requirements

Important! This section contains requirements and recommendations to ensure proper operation and optimum performance of the Spheras Storage Director software.

You must install Spheras Storage Director Software on a system that is running a supported operating system. For detail on supported systems, see 1.4 Spheras Storage Director Requirements.

The Spheras Storage Director Client can use any network connection. However, for optimum performance, it is recommended that you use a network connection that provides a 100Base-T or better data transfer rate.

If you have multiple HTTP servers installed on the same system, they must be configured to communicate on different ports. Otherwise, the Spheras Storage Director client may not start.

For optimum performance, the server component and the client component should be installed on physically separate systems. While it is possible to install the server and client components on the same system; it is not recommended and may cause performance-related problems.

The Spheras Storage Director software is not designed to support the connection of multiple servers to a single storage subsystem. While it is possible to connect multiple servers to a single storage subsystem, this configuration can cause communication problems that affect the Spheras Storage Director software. In contrast, a single server can be connected to multiple storage subsystems.

The Spheras Storage Director software is a based upon a classic client/server model. Therefore, a path to the Server must be open for the Client to properly initialize. Failure to establish this path causes the Spheras Storage Director Client to hang waiting for communication with the Server. See your system administrator for details.

1.6 Installing the Spheras Storage Director Server and Client

After verifying the pre-installation requirements, you are ready to install the Spheras Storage Director Server component, the Client component, or both components. The same InstallShield Wizard is used to install both components. Install files for each of the supported operating systems are provided on the CD ROM that is shipped with your system.

2.1.3 Starting the Install Process

Note: You must install one of the supported HTTP servers before installing the Spheras Storage Director Server. (For a list of the supported HTTP servers, 1.4 Spheras Storage Director Requirements on page 15.

The program that is used to start the install process depends on the operating system that is used on your system. To proceed, select one of the following sections:

- Windows
- Linux
- Solaris-SparcWindows

Insert the Spheras Storage Director CD-ROM into your CD-ROM drive.

Click Start->Run and type the path of the setup file or click Browse to locate the following setup file:

x:\ssd\win\setup.exe

Where x represents the drive letter of your CD-ROM drive.

Click OK. The InstallShield setup window opens followed by the Spheras Storage Director Installer welcome window. Proceed to 1.7 Server and Client Installation Procedures on page 20.

Linux

You need to be in the Desktop Environment (KDE or Gnome) to install Spheras Storage Director Server with Linux. To install the Spheras Storage Director Server with Linux, follow these steps:

1. Open your desktop environment from the command line by typing:

startx

Note: For demonstration purposes, the K Desktop Environment (KDE) was used in this procedure.

2. Insert the CD-ROM into your CD-ROM drive.

3. Login as root.

4. Mount the CD:

mount /dev/cdrom /mnt/cdrom

(Many Linux environments will automount the CD drive.)

5. Locate the CDROM and find the setup.bin file in the following path:

/ssd/linux/setup.bin

6. Like Windows, there are several ways to run a program. Here are two methods:

Browse using a file manager to:

/mnt/cdrom/ssd/linux

Double click on setup.bin

In a console window, type:

cd /mnt/cdrom/ssd/linux

type:

./setup.bin

7. The InstallShield Wizard starts and the Spheras Storage Director Installer Welcome screen opens. Proceed to 1.7 Server and Client Installation Procedures on page 20.

Solaris SPARC

You need to be in the X Windows Environment to install Spheras Storage Director with Solaris SPARC. To install the Spheras Storage Director Server with Solaris SPARC, follow these steps:

1. Open your Common Desktop Environment.

2. Insert the CD-ROM into your CD-ROM drive.

3. Login as root.

4. Mount the CD. (Many Solaris environments have auto mounting so you may not have to complete this step.)

mount -F hsfs /dev/sr0 /mnt

5. In a console window, type:

cd cdrom/ssd/solaris

6. Type:

./setup.bin

7. The InstallShield Wizard starts and the Spheras Storage Director Installer Welcome screen opens. Proceed to the 1.7 Server and Client Installation Procedures on page 20.

1.7 Server and Client Installation Procedures

Notes:

During the install process, the InstallShield Wizard checks the following on your system:

- That a supported operating system is used
- That there is sufficient disk space for the installation
- That the user has administrator privileges on the system
- That services for any previous installations of Spheras Storage Director have been stopped

Error messages are posted if the system fails any of these checks. The dialogs for these messages are not shown in this section.

In this section, the screens that are shown are from a Windows system. The screens from all other operating systems are similar with the exception of the paths and path delimiters.

2.1.4 Server Installation Steps

This section describes the installation of the server component of the Spheras Storage Director. For information on installing the client component, see 2.1.5 Client Installation Steps on page 28.

1. After starting the installation using the setup program specified 2.1.3 Starting the Install Process on page 18, the InstallShield Wizard starts and the welcome screen is displayed.

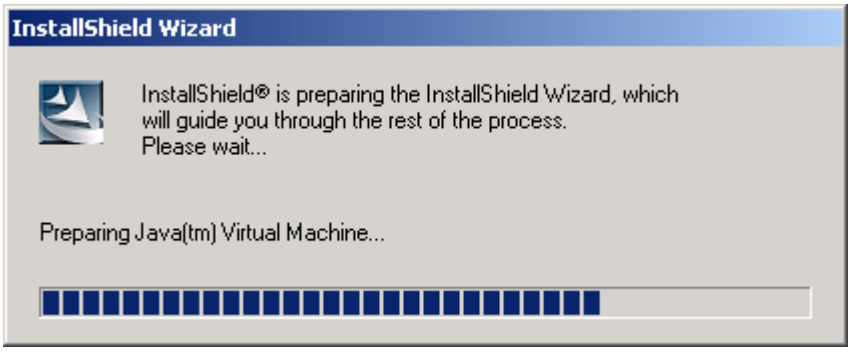

The InstallShield Wizard will first prepare the Java Virtual Machine.
It will then display the welcome/copyright screen.

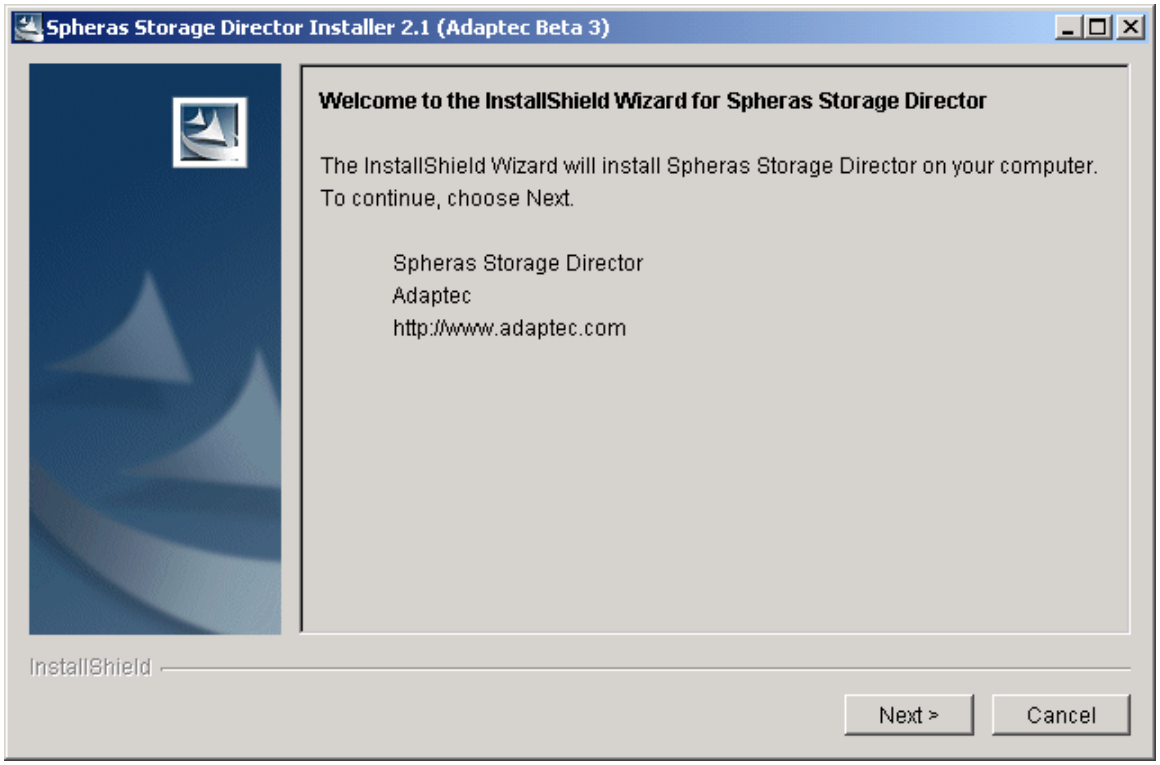

Figure 2 Installation Welcome

2. Click Next. The License Agreement is displayed.

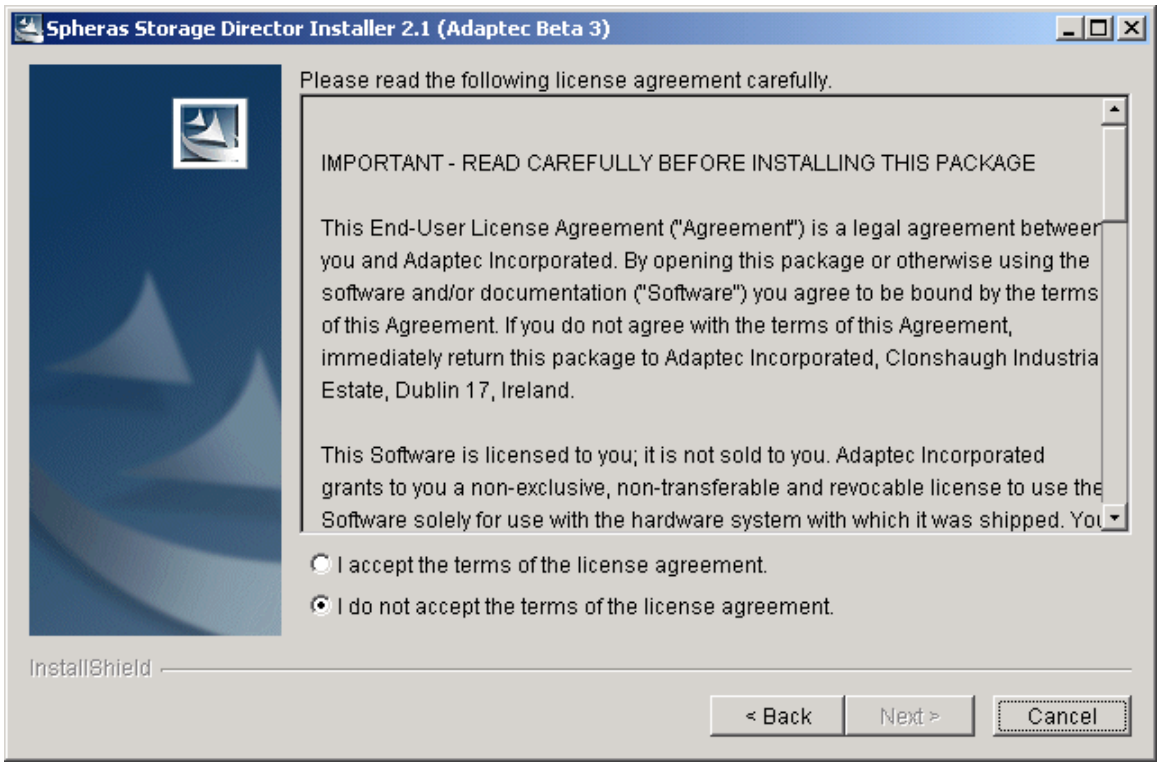

Figure 3 Spheras Storage Director license agreement

3. Read the license agreement. Select the "I accept the terms of the license agreement" radio button to accept the terms of the agreement and continue. If you select the "I do not

accept the terms of the license agreement," you will not be allowed to continue the installation. A dialog box will open asking if you wish to cancel the setup. Click OK to cancel the installation.

4. Click Next. In the Directory Name field, type your desired path; keep the default path or click Browse to navigate to the desired directory.

Note: Use the path name and delimiter that is appropriate for your operating system. For example, in Solaris the Directory Name would appear as /opt/Adaptec/SpherasStorageDirectorServer.

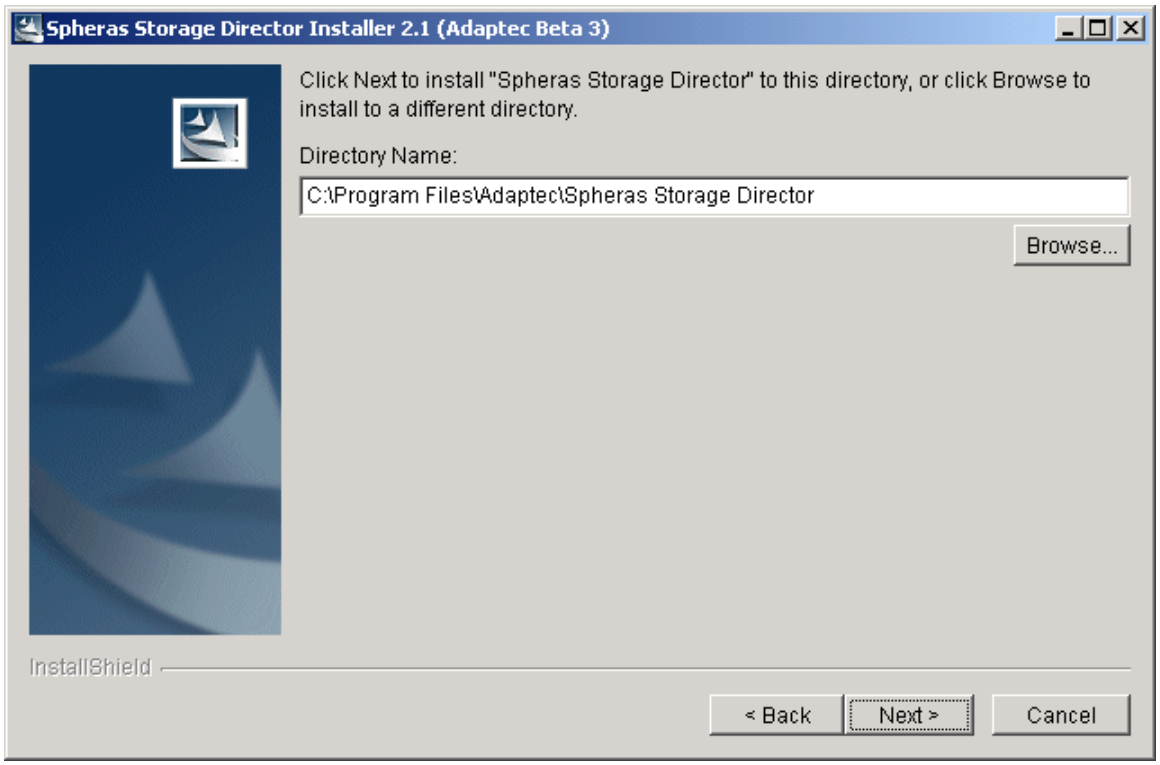

Figure 4 Choose the location for the server installation

5. Choose the type of installation

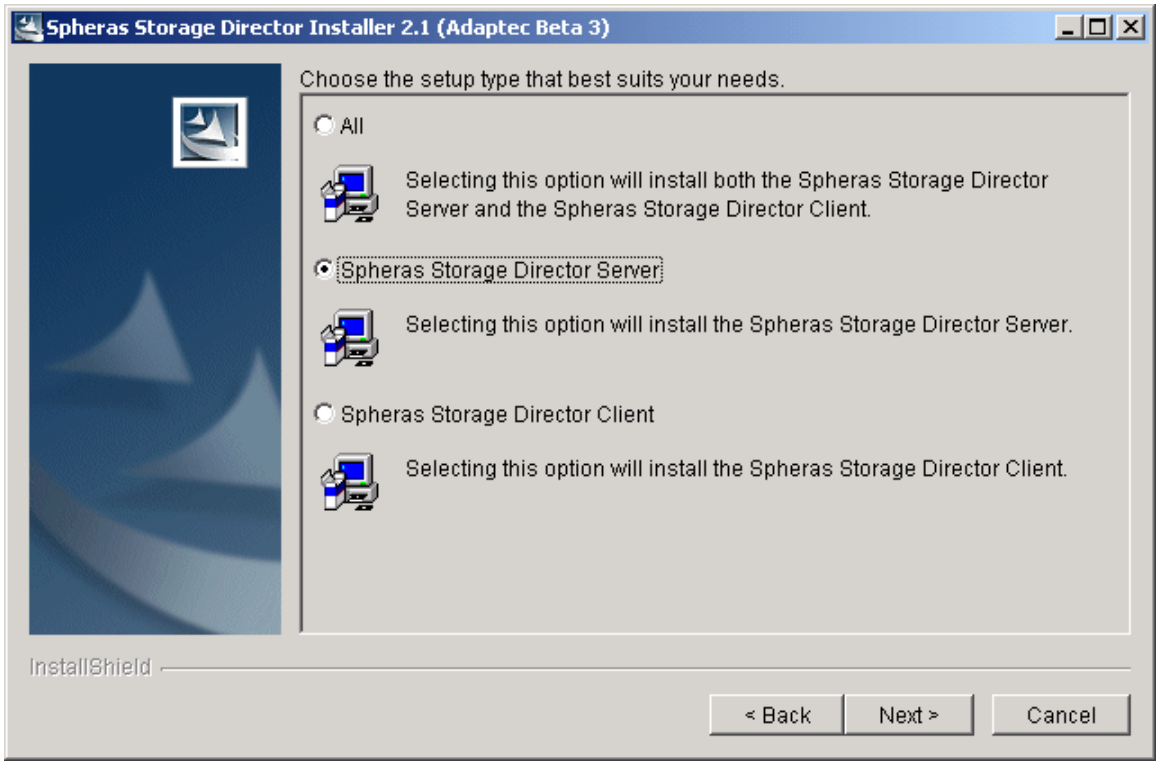

Figure 5 Choose Installation Type

6. Click Next. In the Java RMI (Remote Method Invocation) Port Number field, use the default port number or type a desired port number. The RMI port represents the connectivity path that the server uses to communicate with the client. All Spheras Storage Director Servers must use the same RMI port number to be controlled by a single Spheras Storage Director Client.

Note: The value that you specify for the RMI port cannot be less than 1024.

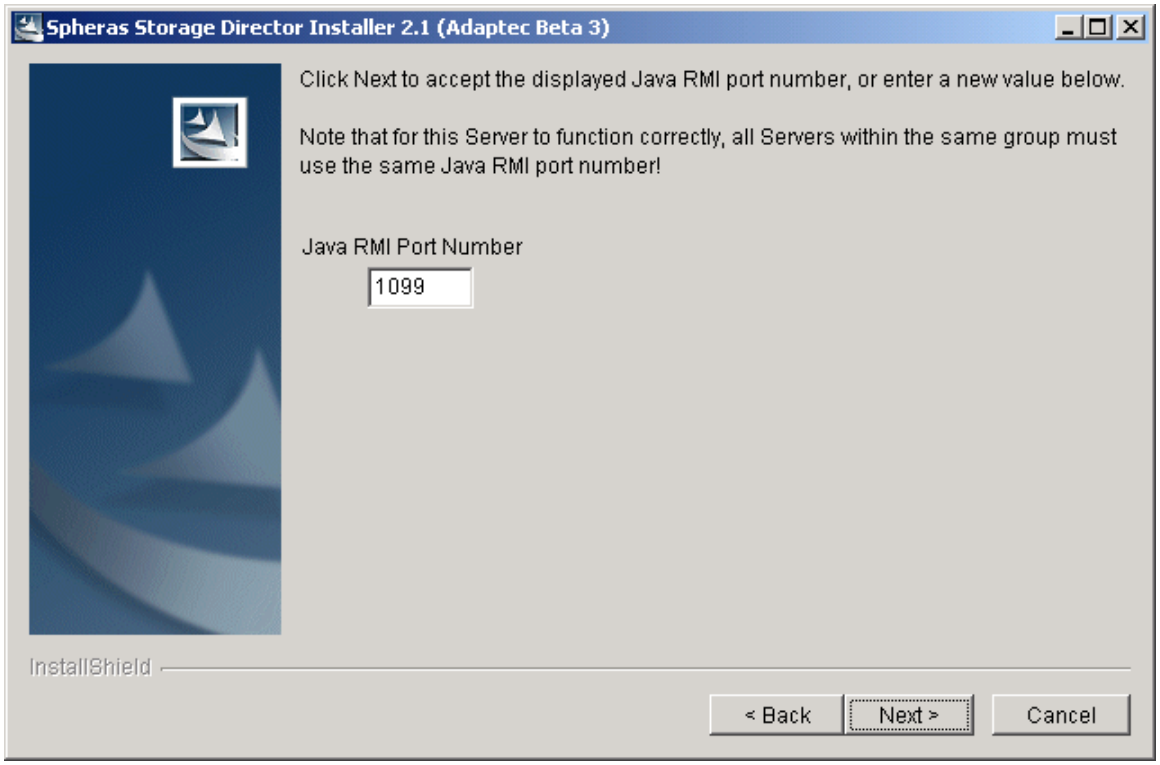

Figure 6 Set the Server RMI Port Number

7. Click Next. Select Yes to automatically configure your HTTP server(s). If both supported HTTP servers (Apache or Microsoft IIS) are properly installed and detected by the InstallShield Wizard, they both will be configured. However, each HTTP server must be configured to communicate on different ports. If you select No, you will need to manually configure your HTTP server after installing the Spheras Storage Director code.

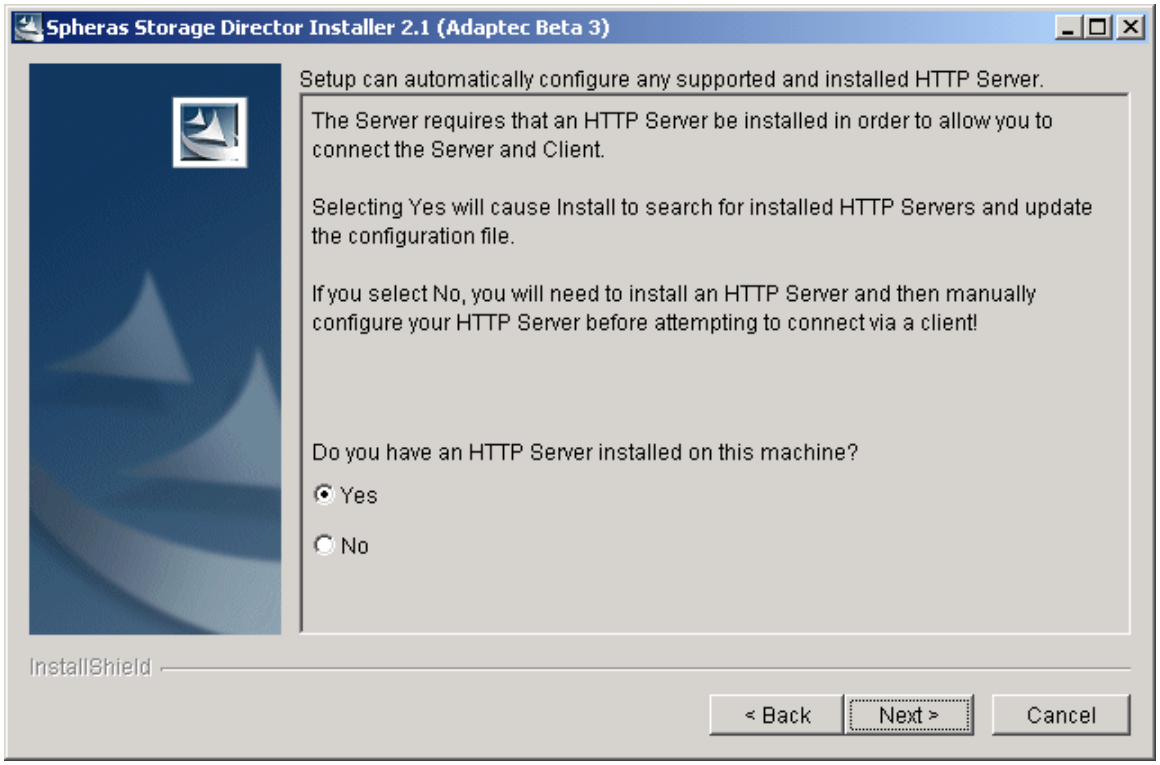

Figure 7 Selecting automatic configuration of an HTTP server

8. Click Next. A summary of what will be installed is displayed. Note that the total size of your installation may not match what is shown here.

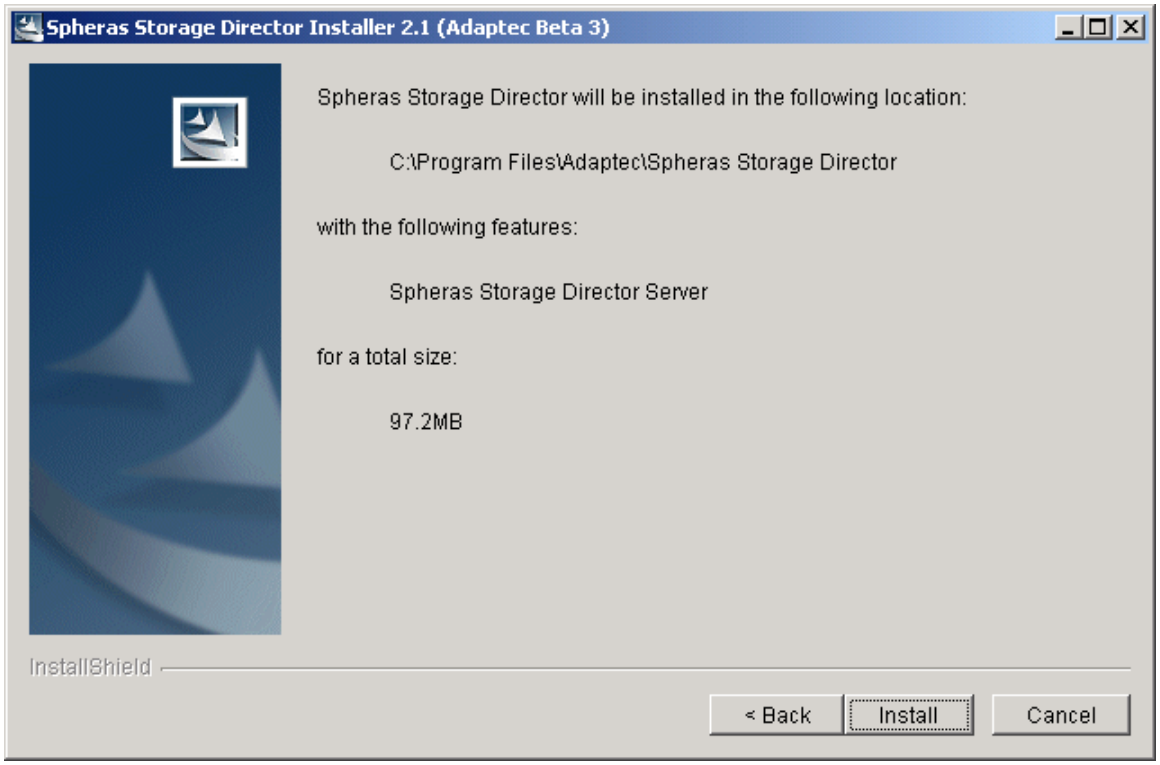

Figure 8 Installation summary

9. Click Install. The installation process begins. During installation, files are decompressed and copied to the destination folder that you previously chose. You will also see flashing blank windows during the installation. This is part of the automatic HTTP configuration.

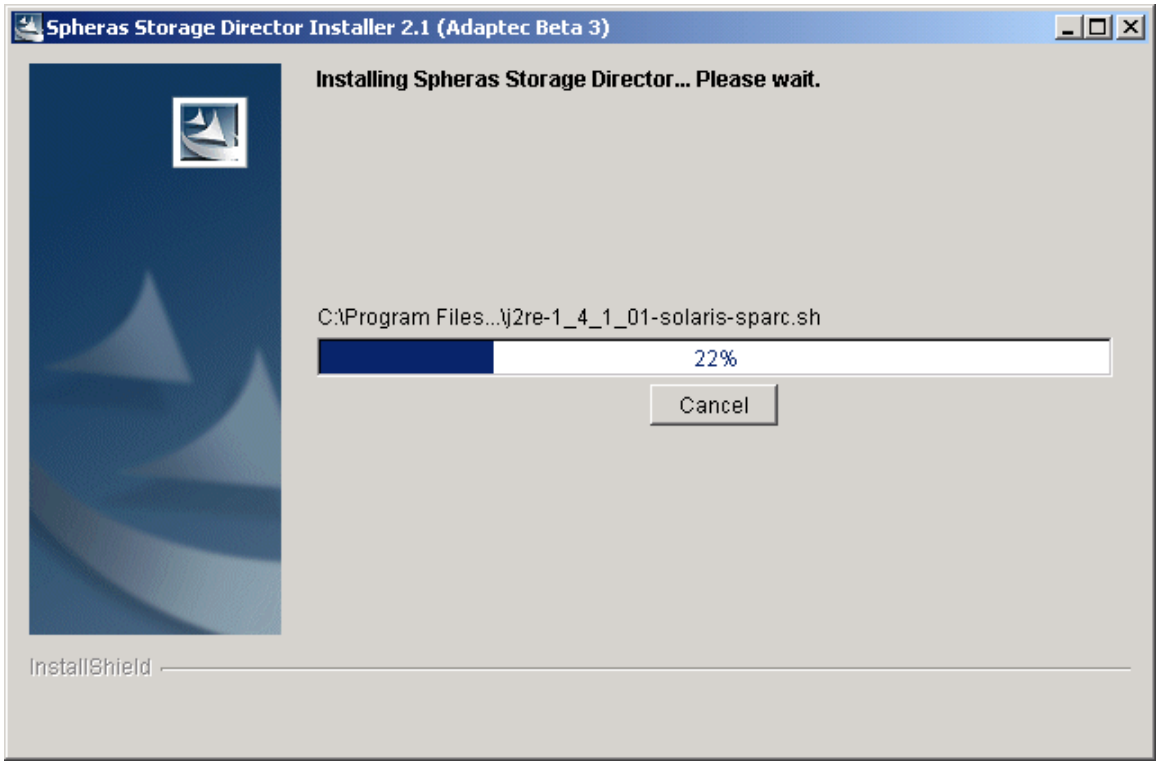

Figure 9 Installation progress

10. When the installation of the server code is complete, the following dialog is displayed. Select whether or not you want to view the readme file and click Next.

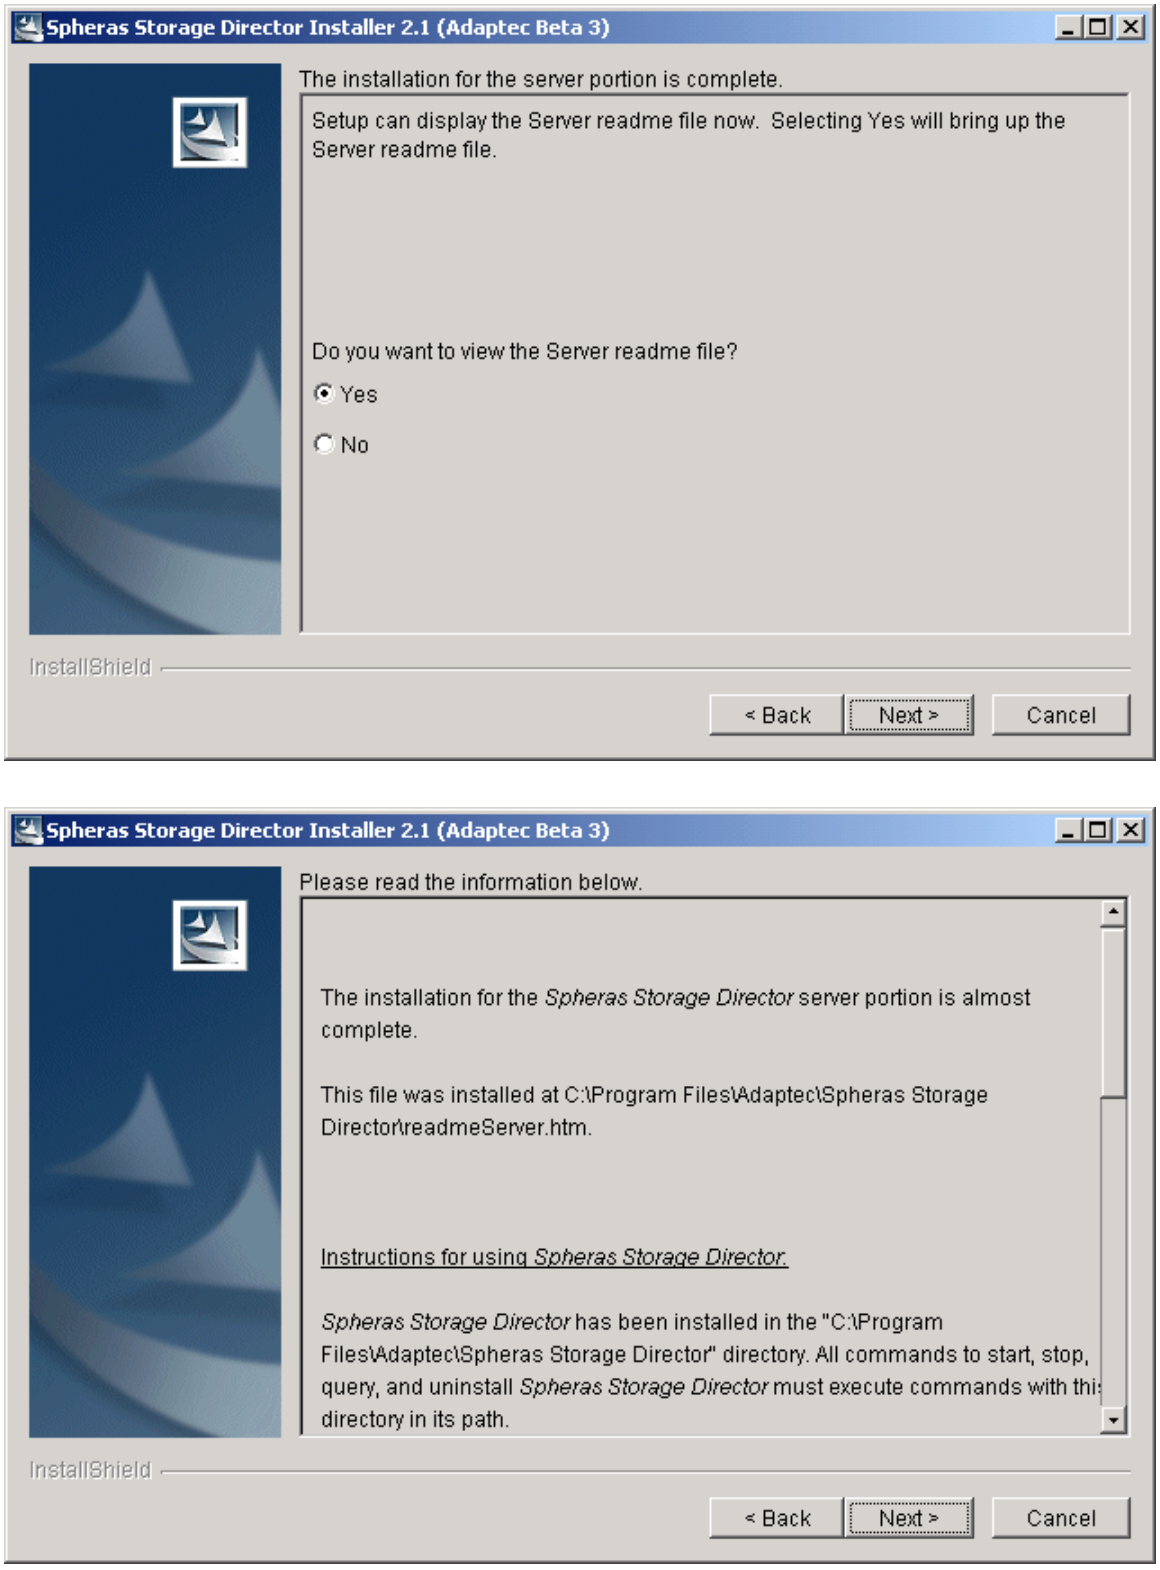

Figure 10 Server installation complete

11. When the installation is complete, the Finish dialog box opens. Click Finish to complete the installation.

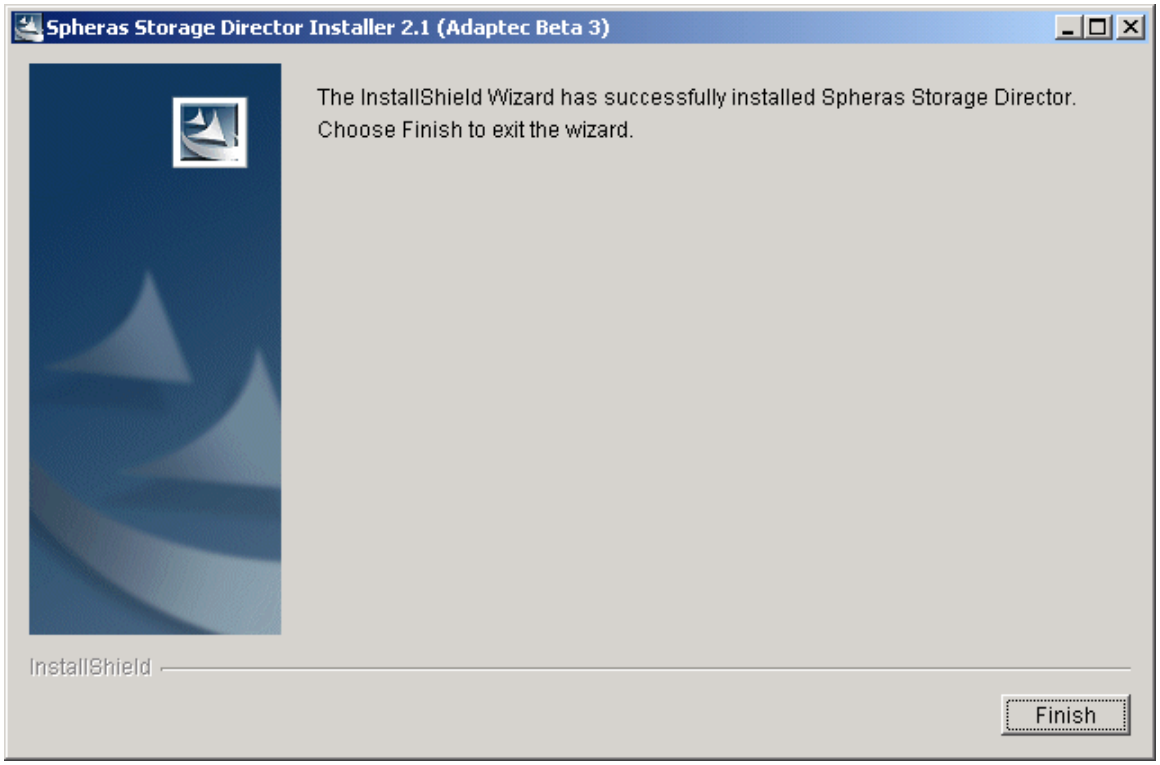

Figure 11 Installation Complete

12. Proceed to one of the following sections:

To install the client code, see 2.1.5 Client Installation Steps on page 28.

If you do not need to set up any additional clients, see:

- 1.8 Windows: Post-Installation Configuration and Verification on page 35
- 1.9 Linux: Post-Installation Configuration and Verification on page 36
- 1.10 Solaris Post-Installation Configuration and Verification on page 38

2.1.5 Client Installation Steps

This section describes the installation of the client component of the Spheras Storage Director. For information on installing the server component, see 2.1.4 Server Installation Steps on page 20.

1. After starting the installation using the setup program specified in 2.1.3 Starting the Install Process on page 18, the InstallShield Wizard starts and the welcome screen is displayed.

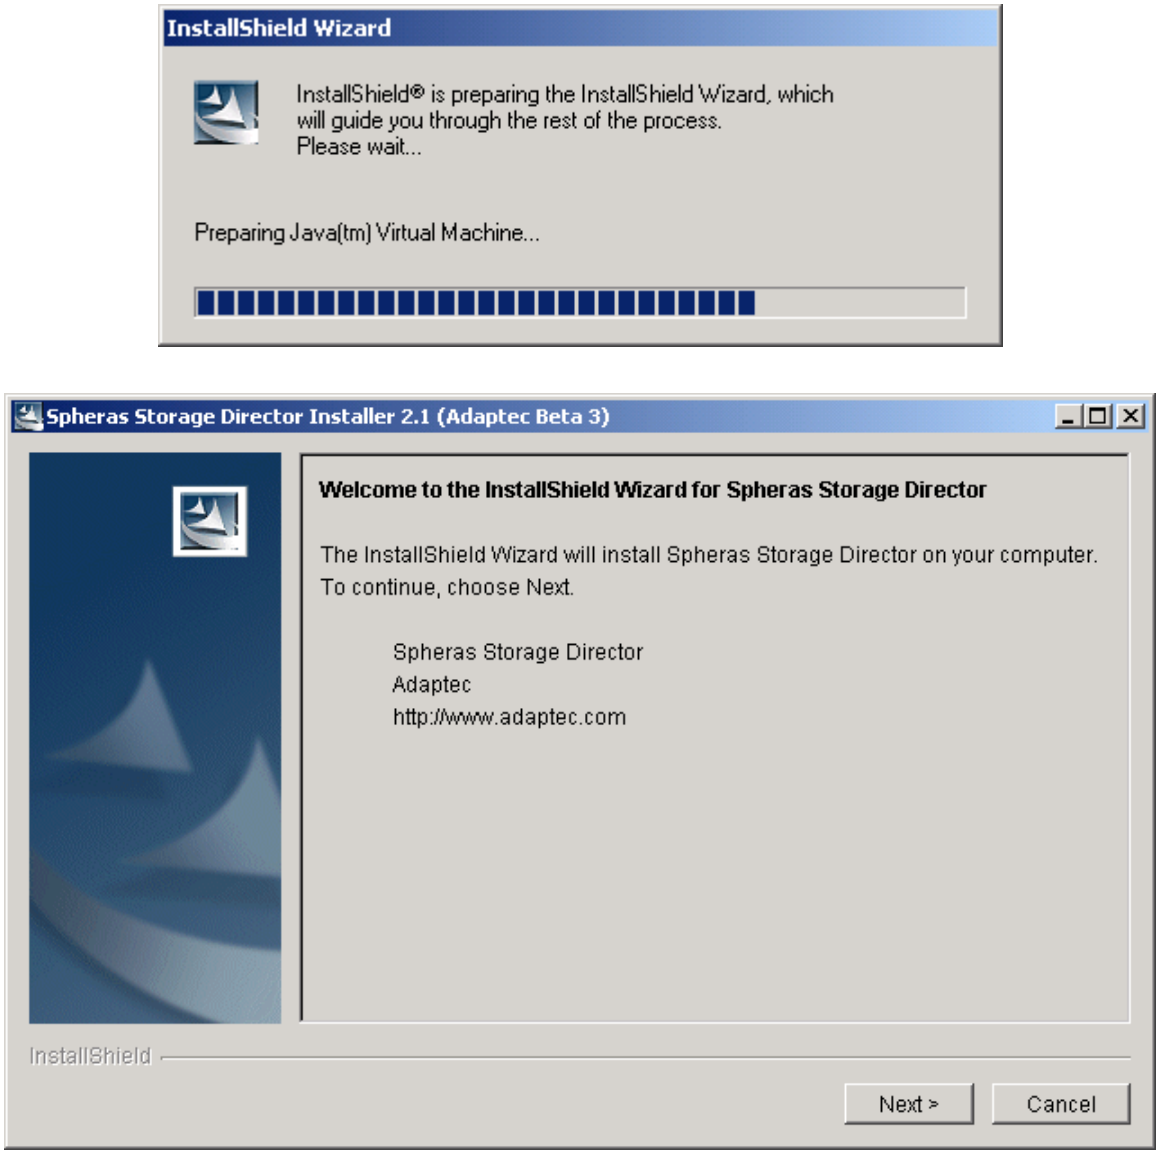

Figure 12 Installation Welcome

2. Click **Next**. The License Agreement is displayed.

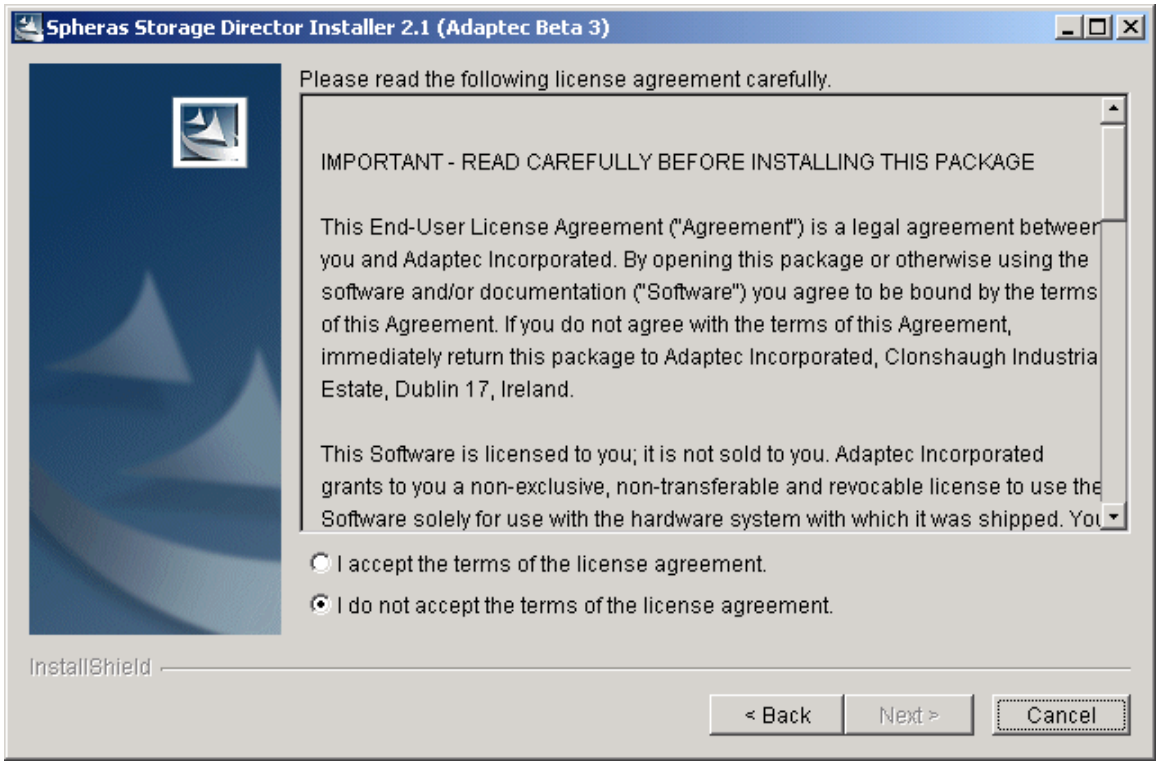

Figure 13 Spheras Storage Director license agreement

3. Read the license agreement. Select the "I accept the terms of the license agreement" radio button to accept the terms of the agreement and continue. If you select the "I do not accept the terms of the license agreement," you will not be allowed to continue the installation. A dialog box will open asking if you wish to cancel the setup. Click **OK** to cancel the installation.

4. Click **Next**. In the Directory Name field, type your desired path; keep the default path or click Browse to navigate to the desired directory.

Note: Use the path name and delimiter that is appropriate for your operating system. For example, in Solaris the Directory Name would appear as /opt/Adaptec/SpherasStorageDirectorServer.

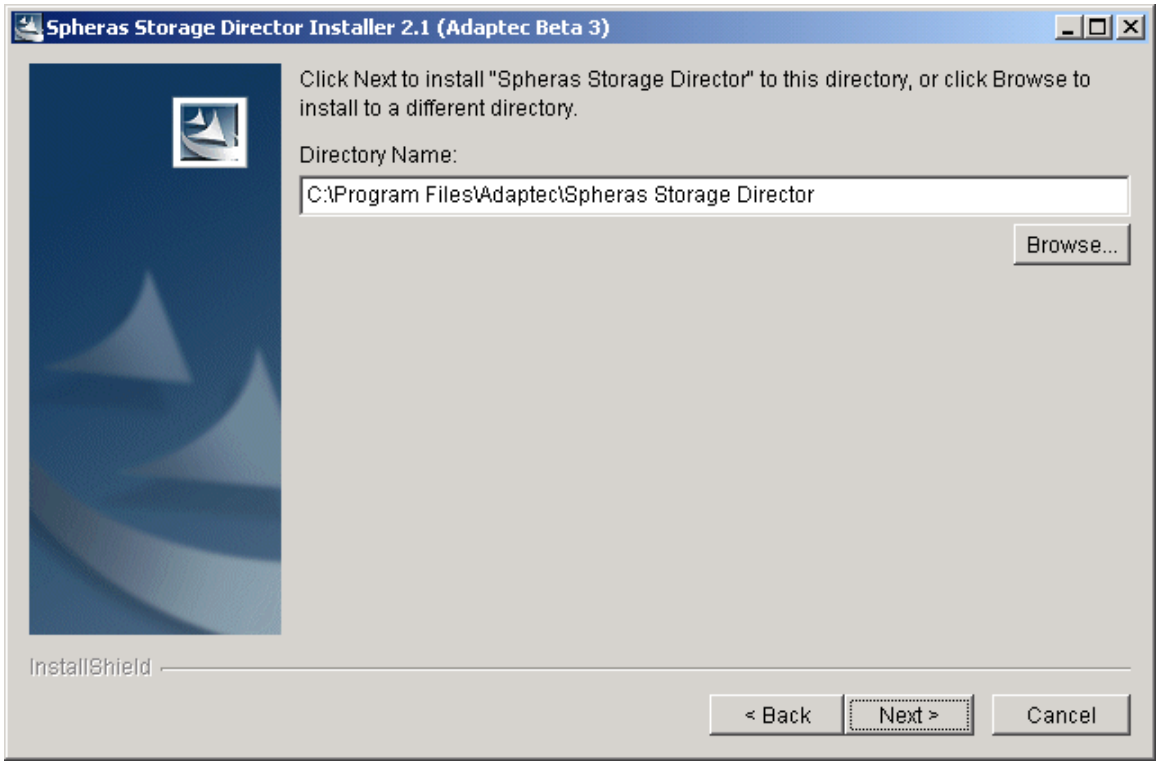

Figure 14 Choose the location for the client installation

5. Select the radio button for the component that you want to have installed.

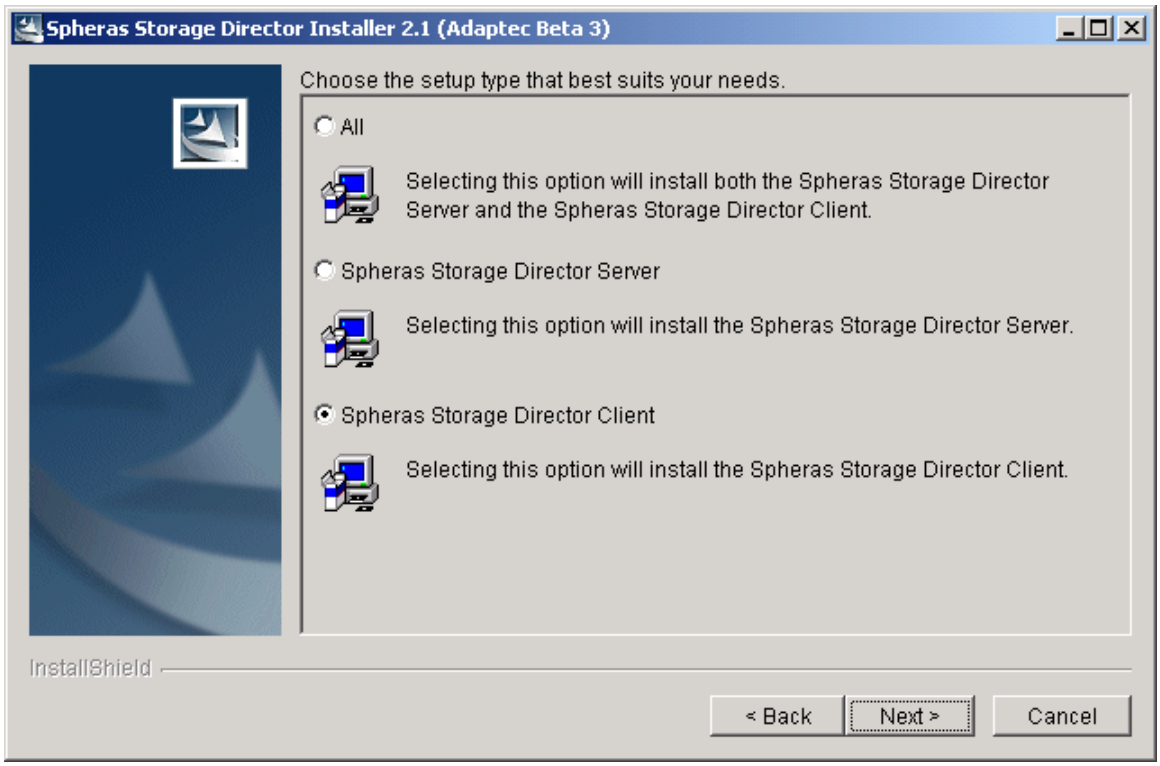

Figure 15 Choose the type of installation

6. Click **Next**. Enter the IP address or host name for one of the servers that you want to connect to. (You can add additional servers later. See: 1.19 Adding a New Server on page 48.

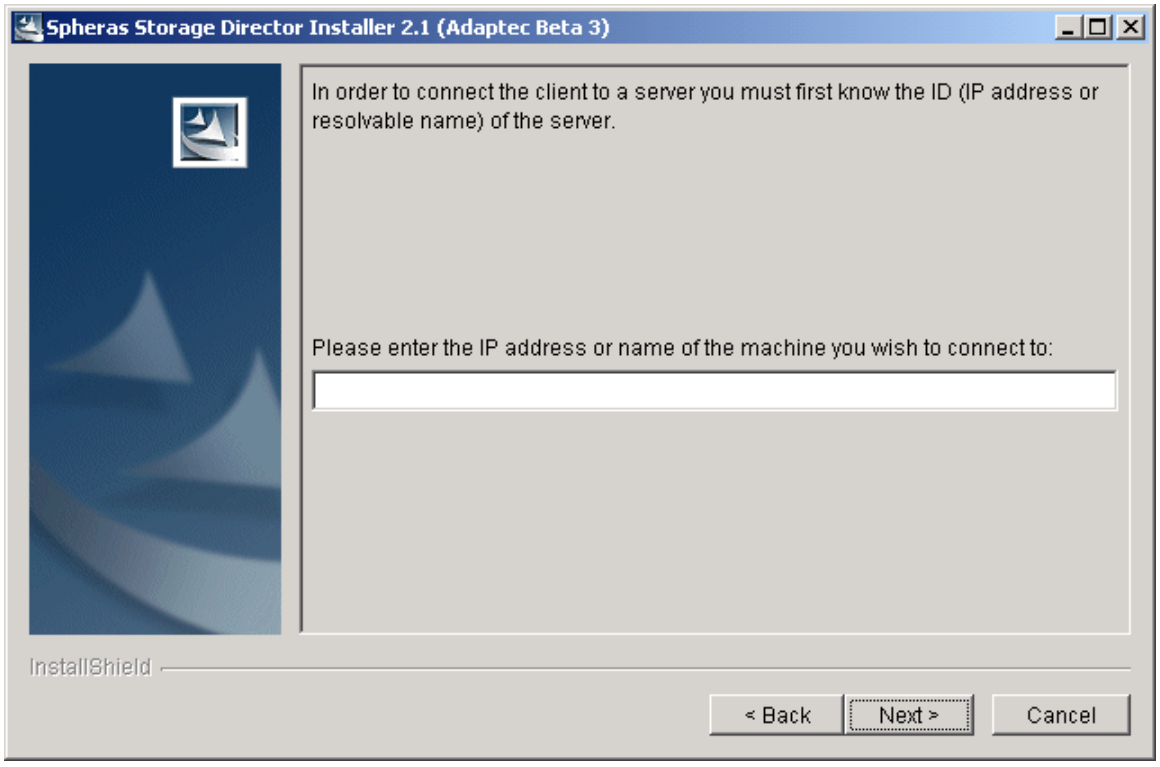

Figure 16 Specifying the server address

7. Click **Next**. A summary of what will be installed is displayed. Note that the total size of your installation may not match what is shown here.

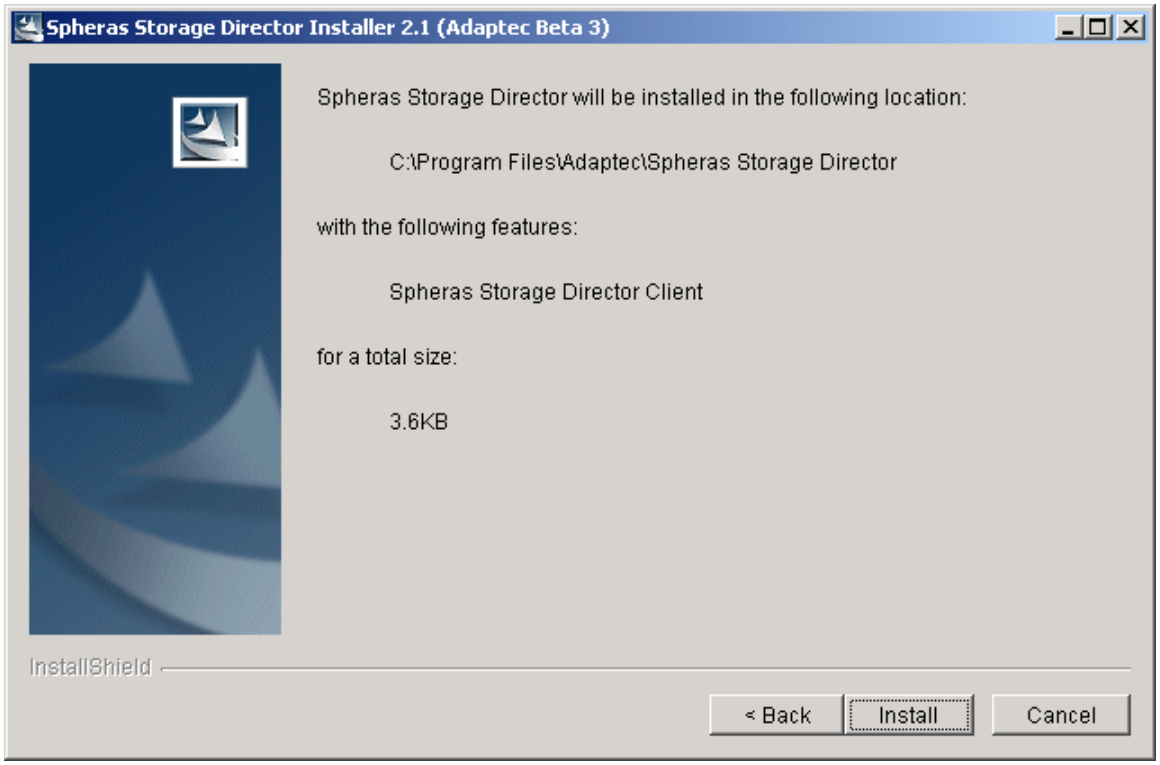

Figure 17 Installation summary

8. Click **Install**. The installation process begins. During installation, files are decompressed and copied to the destination folder that you previously chose.

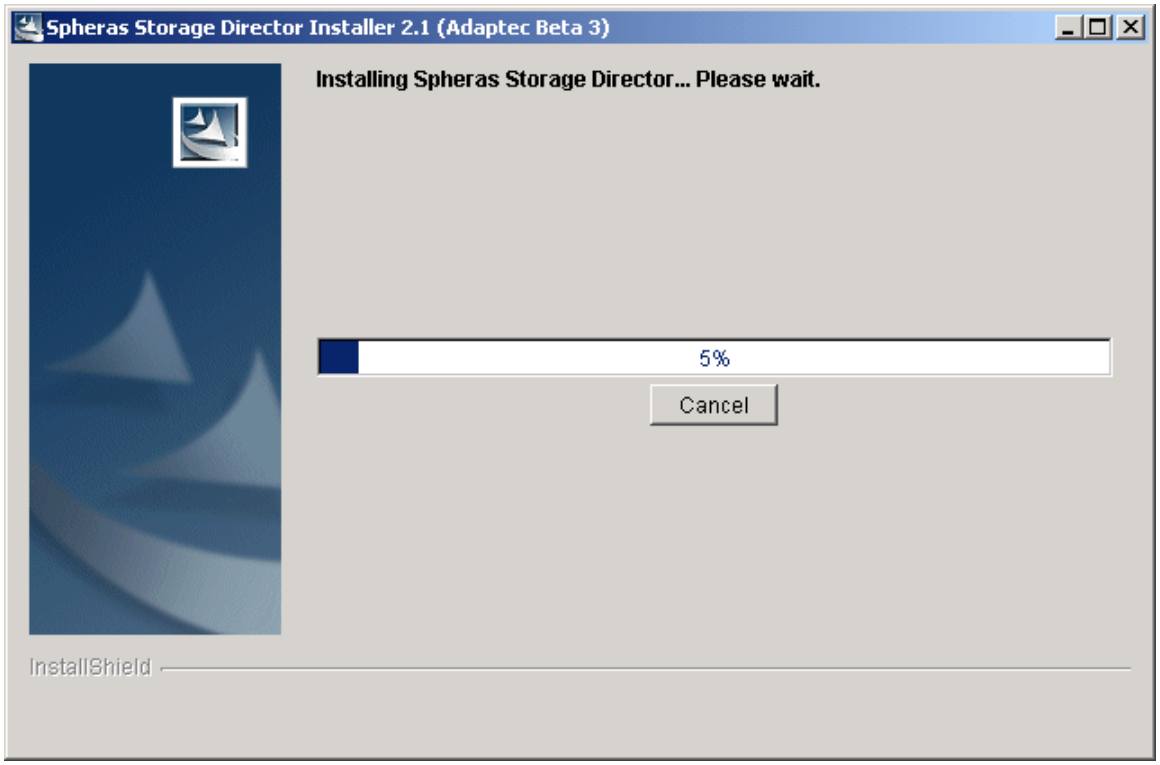

Figure 18 Installation progress

9. Client Readme

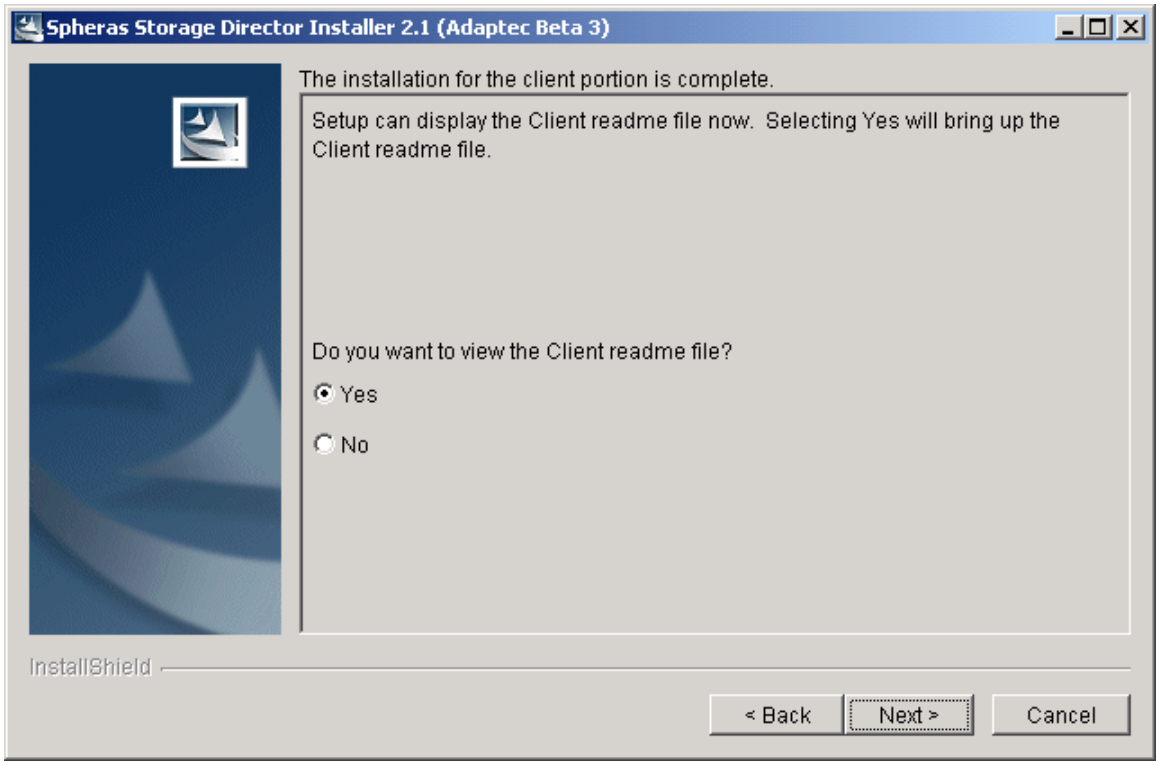

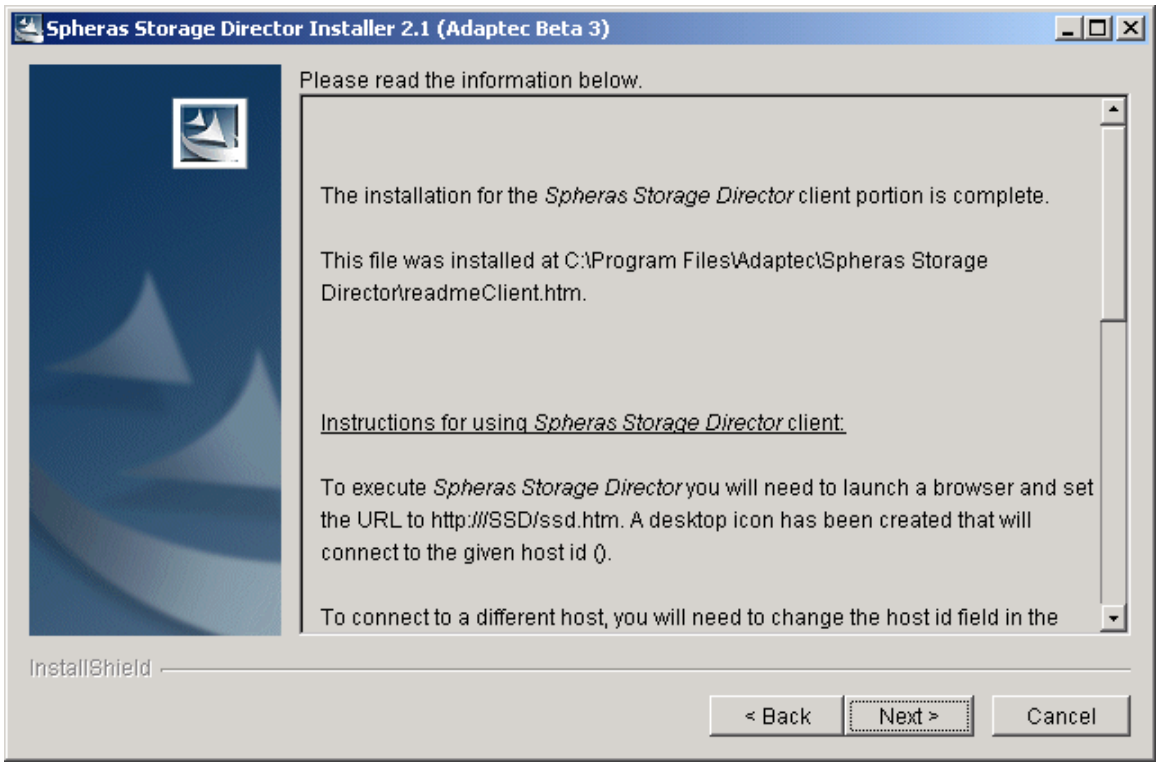

Figure 19 Client installation complete

10. When the installation is complete, the Finish dialog box opens. Click Finish to complete the installation.

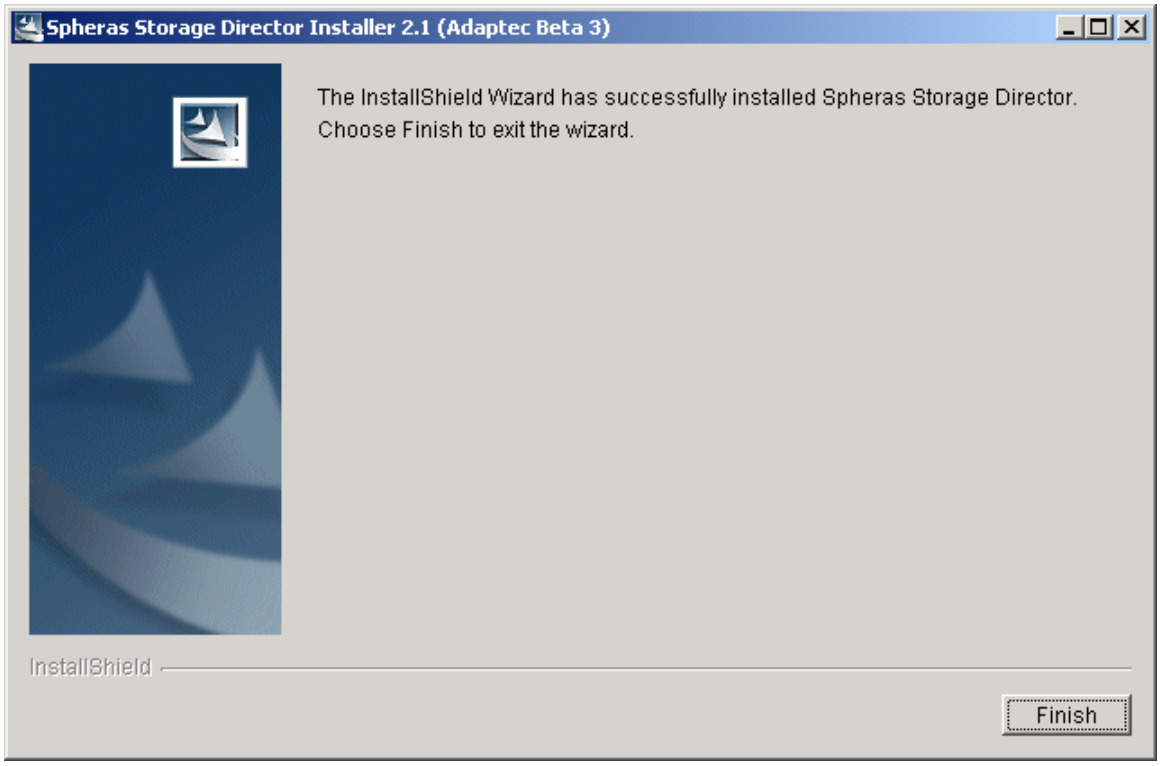

Figure 20 Installation complete

11. Proceed to one of the following sections:

    - 1.8 Windows: Post-Installation Configuration and Verification on page 35.
    - 1.9 Linux: Post-Installation Configuration and Verification on page 36.
    - 1.10 Solaris Post-Installation Configuration and Verification on page 38.

1.8 Windows: Post-Installation Configuration and Verification

Server Component

Verify or perform the following on systems that have the server component of the Spheras Storage Director installed. No additional configuration is required on systems that have the client component installed:

1. Edit the .\Program Files\Apache Group\Apache2\conf\httpd.conf file and ensure that the ServerName and any aliases are correct. For more information, see A.1.1 Windows - Manually Configuring an Apache HTTP Web Server on page 235.

2. Edit the .\WINNT\system32\drivers\etc\hosts file and ensure that it contains:

    - An entry for the server itself. The default setting is the loopback value of 127.0.0.1 for the name localhost.
    - The static IP address and name of the Spheras Storage Director Server. (This is necessary for Spheras Storage Director Clients to communicate with the server.)

3. Verify your HTTP server installation by stopping and restarting your HTTP server:
   Start-> Settings-> Control Panel-> Administrative Tools ->Services
   Scroll down the list and locate Apache2. Right-click on the service and select **Stop**. Right-click on the same service again and select **Start**.

4. If you did not install your HTTP server prior to installing the Spheras Storage Director server, your HTTP server was not configured automatically during the installation process. To configure your HTTP server, see Appendix A Manually Configuring HTTP Servers on page 235.

5. Set up a user ID on the system that is running the Spheras Storage Director server. The Spheras Storage Director server uses existing user IDs on the server running it to authenticate users and assign authorities to them. See 1.13 Security access levels on page 45.

6. Stop and restart the Spheras Storage Director Server service:
   Start-> Settings-> Control Panel-> Administrative Tools ->Services
   Scroll down the list and locate the Spheras Storage Director Server. Right-click on the service and select **Stop**. Right-click on the same service again and select **Start**.

Client Component

The Spheras Storage Director Client requires a fully-configured, JRE-enabled web browser. Each time that a Spheras Storage Director Client connects to a Spheras Storage Director Server the browser configuration is checked to ensure that an appropriate level of the JRE plugin is installed. If an appropriate level of the JRE plugin is not found, the "Java Runtime Environment Plugin Installation" starts automatically. For more information, see 1.12 Starting the Spheras Storage Director Client on page 42.

1.9 Linux: Post-Installation Configuration and Verification

Server Component

Verify or perform the following on systems that have the server component of the Spheras Storage Director installed:

1. Edit the /usr/local/apache2/conf/httpd.conf file and ensure that the ServerName and any aliases are correct. For more information, see A.2 Linux - Manually Configuring an Apache HTTP Web Server on page 241.

2. Edit the /etc/hosts file and ensure that it contains:

   - An entry for the server itself. The default setting is the loopback value of 127.0.0.1 for the name localhost.

   - The static IP address and name of the Spheras Storage Director Server. (This is necessary for Spheras Storage Director Clients to communicate with the server.)

3. Verify your HTTP server installation by stopping and restarting your HTTP server:

   Open a Terminal Window and type the following (these examples assume that you accepted the default path during installation):

cd /usr/local/apache2/bin

then type,

./apachectl stop

followed by:

./apachectl start

4. If you did not install your HTTP server prior to installing the Spheras Storage Director server, your HTTP server was not configured automatically during the installation process. To configure your HTTP server, see Appendix A, Manually Configuring HTTP Servers on page 235.

5. Set up a user ID on the system that is running the Spheras Storage Director server. The Spheras Storage Director server uses existing user IDs on the server running it to authenticate users and assign authorities to them. See 1.13 Security access levels on page 45.

6. Stop and restart the Spheras Storage Director Server service:

   Open a Terminal Window and type the following (these examples assume that you accepted the default path during installation):

cd /opt/Adaptec/SpherasStorageDirectorServer/

then type,

./serverctl stop

followed by:

./serverctl start

Client Component

The Spheras Storage Director Client requires a fully-configured, JRE-enabled web browser. Each time that an Spheras Storage Director Client connects to a Spheras Storage Director Server the browser configuration is checked to ensure that an appropriate level of the JRE plugin is installed. Following are suggestions for installing and configuring the JRE plugin for use with a Netscape browser:

1. Start Netscape and open the URL for an Spheras Storage Director Server. For example: http://server_machine_name/SSD/ssd.htm

2. Download the JRE installation files as instructed by the "Java Runtime Environment Plugin Installation" dialog. This dialog starts automatically if the client browser does not already have an appropriate JRE installed. Download the installation files to a known location. For example: /J2RE (a J2RE directory under the root directory).

3. Change the directory to the location used in the previous step:

cd /J2RE

4. Grant execute permission to the bin file:

chmod 555 file_name

5. Execute the bin file. You will be prompted to acknowledge the User License Agreement.

6. Issue an

rpm –ihv

command on the rpm file.

7. Verify that the plugin path name is /usr/java/j2rel1.4.1_01/plugin/i386/ns4 (assuming that Netscape 4 is being used) directory.

8. Edit the system profile (for example: /.profile or /etc/.profile) and include the plugin path and related export command, for example (assuming that Netscape 4 is being used):

NPX_PLUGIN_PATH=/usr/java/j2rel1.4.1_01/plugin/i386/ns4
export NPX_PLUGIN_PATH

9. Logout and log back into the system.

10. To verify that the plugin path has been set up correctly, issue:

set | grep -i NPX

11. Proceed to 1.12 Starting the Spheras Storage Director Client on page 42.

1.10 Solaris Post-Installation Configuration and Verification

This section contains additional items that must be configured or verified on Solaris systems.

2.1.6 Define Devices and Target/LUN Combinations

1. Modify the sd.conf file to define devices and target/lun combinations. In a console window enter:

cd /kernel/drv

Open the sd.conf file with an editor

Include entries for multiple lun support for the SCSI disk driver:

```
name="sd" class="scsi" class_prop="atapi" target=0 lun=0;
name="sd" class="scsi" class_prop="atapi" target=1 lun=0;
name="sd" class="scsi" class_prop="atapi" target=2 lun=0;
name="sd" class="scsi" class_prop="atapi" target=3 lun=0;
name="sd" class="scsi" target=0 lun=0;
name="sd" class="scsi" target=0 lun=1;
name="sd" class="scsi" target=0 lun=2;
name="sd" class="scsi" target=0 lun=3;
name="sd" class="scsi" target=0 lun=4;
name="sd" class="scsi" target=0 lun=5;
name="sd" class="scsi" target=0 lun=6;
name="sd" class="scsi" target=0 lun=7;
name="sd" class="scsi" target=1 lun=0;
 name="sd" class="scsi" target=1 lun=1;
name="sd" class="scsi" target=1 lun=2;
name="sd" class="scsi" target=1 lun=3;
name="sd" class="scsi" target=1 lun=4;
name="sd" class="scsi" target=1 lun=5;
name="sd" class="scsi" target=1 lun=6;
name="sd" class="scsi" target=1 lun=7;
```

2. The Spheras Storage Director uses generic SCSI device nodes for controller discovery. Edit the sgen.conf file to enable the generic SCSI device nodes. In a console window enter:

cd /kernel/drv

Open the sgen.conf file with an editor

Add the following lines:

device-type-config-list="direct";
inquiry-config-list="EUROLOGC", "FC250";*

Uncomment or add target/lun pairs where the controllers and luns exist on the SCSI bus. The contents of the sd.conf and sgen.conf files should match the desired configuration.

Where the controllers exist on the SCSI bus is dependent on the system's host bus adapter, SCSI configuration, and topology. Therefore, a specific example cannot be given. One possible example follows.

name="sgen" class="scsi" target=0 lun=0;
name="sgen" class="scsi" target=0 lun=1;
name="sgen" class="scsi" target=0 lun=2;
name="sgen" class="scsi" target=0 lun=3;
name="sgen" class="scsi" target=0 lun=4;
name="sgen" class="scsi" target=0 lun=5;
name="sgen" class="scsi" target=0 lun=6;
name="sgen" class="scsi" target=0 lun=7;

name="sgen" class="scsi" target=1 lun=0;
name="sgen" class="scsi" target=1 lun=1;
name="sgen" class="scsi" target=1 lun=2;
name="sgen" class="scsi" target=1 lun=3;
name="sgen" class="scsi" target=1 lun=4;
name="sgen" class="scsi" target=1 lun=5;
name="sgen" class="scsi" target=1 lun=6;
name="sgen" class="scsi" target=1 lun=7;
name="sgen" class="scsi" target=2 lun=0;
name="sgen" class="scsi" target=2 lun=1;
name="sgen" class="scsi" target=2 lun=2;
name="sgen" class="scsi" target=2 lun=3;
name="sgen" class="scsi" target=2 lun=4;
name="sgen" class="scsi" target=2 lun=5;
name="sgen" class="scsi" target=2 lun=6;
name="sgen" class="scsi" target=2 lun=7;
name="sgen" class="scsi" target=3 lun=0;
name="sgen" class="scsi" target=3 lun=1;
name="sgen" class="scsi" target=3 lun=2;
name="sgen" class="scsi" target=3 lun=3;
name="sgen" class="scsi" target=3 lun=4;
name="sgen" class="scsi" target=3 lun=5;
name="sgen" class="scsi" target=3 lun=6;
name="sgen" class="scsi" target=3 lun=7;

Reboot the system to reconfigure the environment to include the dev links for the devices that you entered in the sd.conf file:

reboot -- -r

Verify that sgen is loaded properly by running the following command on the server system:

/usr/sbin/devfsadm -v -i sgen

A non-zero return code indicates that sgen is not properly loaded.

Check /dev/scsi/direct to verify that sgen was properly loaded.

2.1.7 Server Component

Verify or perform the following after installing the server and client components:

1. Edit the /etc/appache/httpd.conf file and ensure that the ServerName and alias are correct. For more information, see A.3 Solaris - Manually Configuring an Apache HTTP Web Server on page, 243.

2. Edit the /etc/hosts file and ensure that it contains:

   An entry for the server itself. The default setting is the loopback value of 127.0.0.1 for the name localhost.

   The static IP address and name of the Spheras Storage Director Server. (This is necessary for Spheras Storage Director Clients to communicate with the server.)

3. Verify your HTTP server installation by stopping and restarting your HTTP server:

   Open a Terminal Window and type the following (these examples assume that you accepted the default path during installation):

cd /etc/apache2/bin

then type,

./apachectl stop

followed by:

./apachectl start

4. If you did not install your HTTP server prior to installing the Spheras Storage Director server, your HTTP server was not configured automatically during the installation process. To configure your HTTP server, see A.3 Solaris - Manually Configuring an Apache HTTP Web Server on page 243.

5. Set up a user ID on the system that is running the Spheras Storage Director server. The Spheras Storage Director server uses existing user IDs on the server running it to authenticate users and assign authorities to them. See 1.13 Security access levels on page 45.

6. Stop and restart the Spheras Storage Director Server service:

   Open a Terminal Window and type the following (these examples assume that you accepted the default path during installation):

cd /opt/Adaptec/SpherasStorageDirectorServer/

then type,

./serverctl stop

The status line displays if the server is running. If it is not started, type:

./serverctl start

2.1.8 Client Component

The Spheras Storage Director Client requires a fully-configured, JRE-enabled web browser. Following are suggestions for installing and configuring the JRE plugin for use with a Netscape browser:

1. Start Netscape and open the URL for a Spheras Storage Director Server. For example: http://server_machine_name/SSD/ssd.htm

2. Download the JRE installation files as instructed by the “Java Runtime Environment Plugin Installation” dialog. This dialog starts automatically if the client browser does not

already have an appropriate JRE installed. Download the installation files to a known location. For example: /J2RE (a J2RE directory under the root directory).

3. Change directory to the location that was used in the previous step:

cd /J2RE

The two files that were downloaded from the previous step should be: j2re-1_4_1_01-solaris-sparc.sh and j2re-1_4_1_01-solaris-sparcv9.sh

4. Grant execute permission to the executable files:

chmod 555 file_name

5. Execute j2re-1_4_1_01-solaris-sparc.sh first and then j2re-1_4_1_01-solaris-sparcv9.sh

Executing these files creates a subdirectory that contains the plugin file: /J2RE/j2re1.4.1_01

6. Edit the system profile (for example: /.profile or /etc/profile) and include the plugin path and related export command, for example (assuming that Netscape 4 is being used):

NPX_PLUGIN_PATH=/J2RE/j2re1.4.1_01/plugin/sparc/ns4
export NPX_PLUGIN_PATH

7. Logout and log back into the system.

8. To verify that the plugin path has been set up correctly, issue:

set | grep -i NPX

9. Proceed to 1.12 Starting the Spheras Storage Director Client on page 42.

1.11 Updating Your Spheras Storage Director Server and Client

This section assumes that you have been using the Spheras Storage Director Server and Client components and that you want to update your software with the latest version. It is recommended that you upgrade all of your Spheras Storage Director Servers with the latest version; failure to do this will affect the ability of the Spheras Storage Director Server to communicate with other Spheras Storage Director Servers.

Only the Spheras Storage Director Server requires updating. The Client is a Java Applet that communicates with the Server and does not require updating.

2.1.9 Updating Your Spheras Storage Director Server with Windows

To update your Spheras Storage Director Server software, follow these steps:

1. Go to Start->Settings->Control Panel->Administrative Tools-> Services.
2. In the Services dialog box, right-click the Spheras Storage Director Server and select Stop.
3. Close the dialog box.
4. Go to Start->Settings->Control Panel->Add/Remove Programs.
5. Select the Spheras Storage Director Server.
6. Click Change/Remove to uninstall the Spheras Storage Director Server.
7. To ensure compatibility with the new Spheras Storage Director Server, it is recommended that you remove any previous installations of JRE. The Spheras Storage Director Installation process can detect that a JRE is not installed and install a compatible version when required.
8. You are now ready to reinstall the Spheras Storage Director Server software. See 2.1.4 Server Installation Steps on page 20.
9. After you have installed the Spheras Storage Director Server launch the Client, see 2.1.5 Client Installation Steps on page 28.

2.1.10 Updating Your Spheras Storage Director Server with Linux and Solaris

To update your Spheras Storage Director Server software, follow these steps:

1. Go to the directory where the Spheras Storage Director Server is installed

/opt/Adaptec/Spheras Storage Director Server

2. On the command line, type:

./serverctl stop

3. On the command line, type:

./uninstall

4. To ensure compatibility with the new Spheras Storage Director Server, it is recommended that you remove any previous installations of JRE. The Spheras Storage Director installation process detects that JRE is not installed and installs a compatible version.
5. You are now ready to reinstall the Spheras Storage Director Server software. See 2.1.4 Server Installation Steps on page 20.
6. After you have installed the Spheras Storage Director Server launch the Client, see 2.1.5 Client Installation Steps on page 28.

1.12 Starting the Spheras Storage Director Client

The Spheras Storage Director Client is started using a browser (such as Netscape or Explorer) on any of the supported client operating systems to open the URL of a web server that has the Spheras Storage Director Server component installed.

Notes:

- If you are unable to connect to a Spheras Storage Director Server immediately after installing the Spheras Storage Director software, wait two minutes and then ensure that the Spheras Storage Director Server service is running. If problems persist, stop and restart the HTTP server.

- Your session of the Spheras Storage Director Client may end suddenly if you launch another URL. This happens because the program launching the URL may have a default setting configured to use the active browser to open it.

To start the Spheras Storage Director Client:

Open your supported browser (see, 1.4 Spheras Storage Director Requirements on page 15) and type the path of the Spheras Storage Director server. For example:

http://<your server name or IP address>/SSD/ssd.htm

If this is the first time the Spheras Storage Director client is being started, two dialogs are displayed:

If your browser does not meet the Java runtime environment (JRE) plugin requirements, you are presented with instructions for installing the JRE plugin. Follow the instructions that are presented in the dialogs.

Important: The correct level of the JRE plugin must be installed to use the Spheras Storage Director client. If you are prompted to install the JRE plugin, it means that the system has detected that you do not currently have the correct level installed. In this case, you must allow the JRE plugin installation to proceed. Select or specify the following during the installation.

1. Accept the Sun JRE license agreement.

2. Select Typical for the type of installation.

3. Accept the security warning regarding installation of the "Sun Microsystems, Inc. Java 2 Runtime Environment."

A Security Warning dialog asks if you want to install and run signed applets. The dialog box has four choices to select from:

- **Grant this session:** You will see the warning again next time you launch the client.

- **Deny:** You cannot run Spheras Storage Director

- **Grant always:** You will not see the dialog again.

- **More Info:** Certificate information.

It is recommended that you choose "Grant always."

The Spheras Storage Director console loads in its own application window. A splash screen is displayed immediately in your browser while the Spheras Storage Director console applet is loading. When the Spheras Storage Director Client is loaded, it will authenticate with the Spheras Storage Director Server. This authentication is based on a userid/password defined to the server through the operating system specific user management services. Once a userid has

been authenticated, the appropriate authorization for the userid will be determined from its group memberships on the server. See 1.13 Security access levels on page 45.

To learn more about the navigation features and options that the Spheras Storage Director Client provides, see Chapter 4, Navigating the Spheras Storage Director Client on page 55.

Chapter 3. Setting up the Spheras Storage Director Client

This section describes the security access levels, passwords, and method to sign onto the Spheras Storage Director Client. The process to add Servers and create Server Groups is also described.

1.13 Security access levels

The Spheras Storage Director Client authenticates with the Spheras Storage Director server. This authentication is based on user IDs and passwords that have been defined on the servers that host the Spheras Storage Director Server. The user IDs and passwords are defined on each server through the operating system's user management services. Said another way: the Spheras Storage Director uses server-level authentication, not domain-level authentication. Therefore, each server must have its own accounts.

Note: The user IDs and passwords must be defined using only standard English-language characters.

The Spheras Storage Director Client imposes no limit on the password length, however a password must be used.

Spheras Storage Director supports three access levels. On Windows systems, the groups are User, Power Users, and Administrators. On Unix systems, the groups are root, adm, and users.

Note: You can change the default values via the Server Properties dialog box.

User / users:	Allows viewing of the controller configuration, event logs, and long operation status (for example, the status of rebuilds, initializations, consistency checks, and capacity expansion). In essence, Users have Read-only privileges.
PowerUsers / adm:	Allows viewing of all the information that Spheras Storage Director provides. In addition, these users with this access level can define hot spares; start rebuilds, and run consistency checks.
Administrators / root:	Allows all the capabilities of User and Power User plus all of the Spheras Storage Director features that manipulate controllers and RAID subsystems.

1.14 Passwords

A password must be established for each user on the server host. Passwords can be any valid English-language characters that the server supports. Passwords are encrypted when transmitted across the network. The Spheras Storage Director Client imposes no limit on the password length.

Users with Administrator access are prompted to enter their password prior to applying changes that result in loss of data or loss of data availability.

1.15 Logging on to Spheras Storage Director

The Spheras Storage Director Client login dialog opens automatically upon Spheras Storage Director Client startup. When a user attempts to access remote server resources, the client automatically attempts a logon to that server by using the last successful user name and password. If the automatic logon attempt fails, you are prompted to enter a name and password to log onto the remote server. If you try to logon to a server on which you are already logged in, a warning appears. If you choose to continue the logon, it is treated as a logoff, followed by a log on.

The Spheras Storage Director client imposes no limit to the number of failed login attempts.

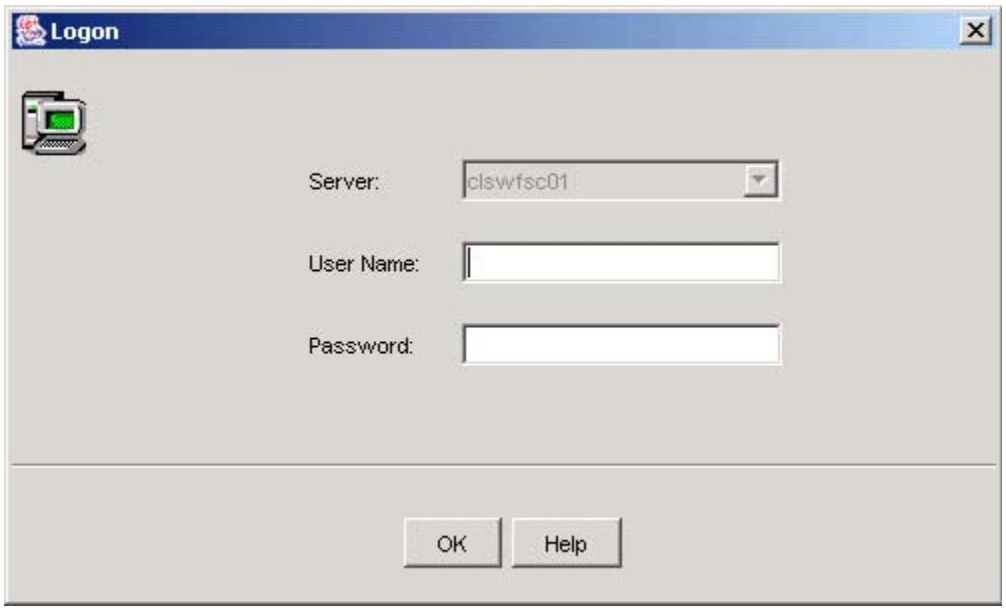

Figure 21 Logon Dialog

Use the following steps to logon to Spheras Storage Director:

1. Type your User name. Select a user name with access privileges on the server that correspond to the desired access level in Spheras Storage Director. User names are encrypted when transmitted across the network. There is no limit on the User name length. See 1.13 Security access levels on page 45 for more information.

2. Enter your password in the login dialog. Passwords are encrypted when transmitted across the network. See 1.14 Passwords on page 45 for more information.

3. Click **OK** or press the Enter key to logon.

1.16 Setting Up Server Groups

This section describes how to setup connections to server groups. A maximum of 100 server groups can be created. By creating server groups, you can organize your storage structure. For example, you could create server groups for business units, physical locations, or operating systems.

1. Click **Group->New** on the menu bar and the New Server Group dialog box opens.

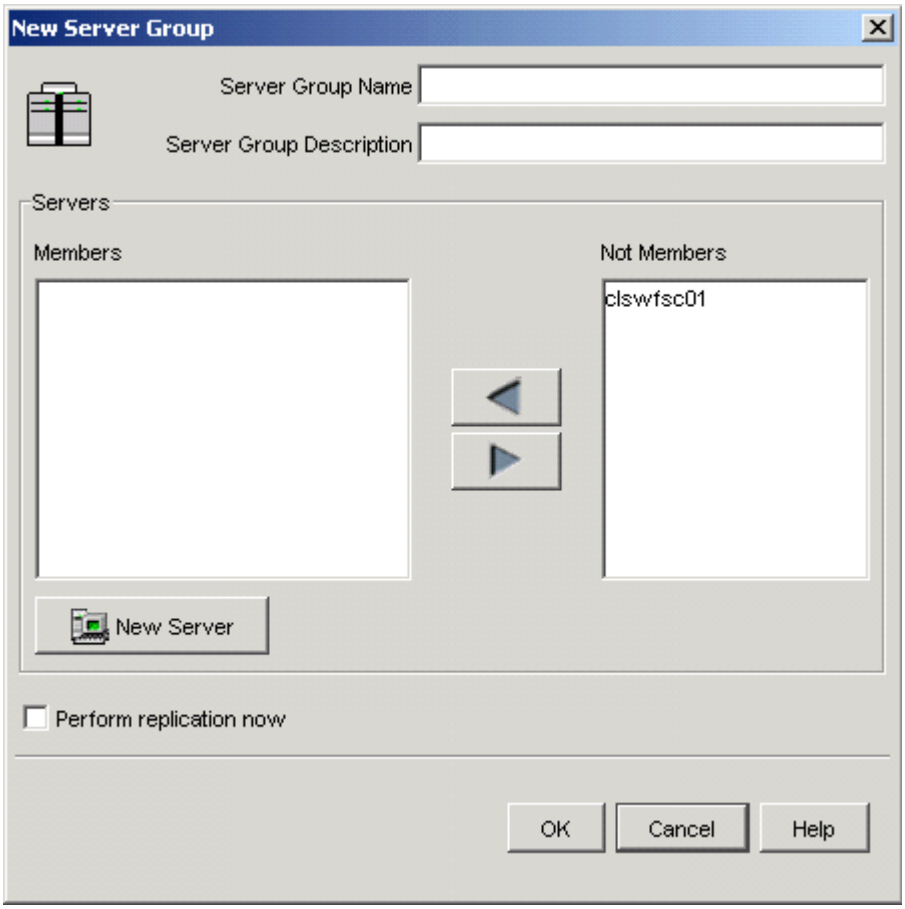

Figure 22 New Server Group Dialog Box

2. In the Server Group Name field, type the desired name of the server group you are adding. All characters are acceptable and there is no limit to the number of characters you can enter. Naming a server group is a way of identifying a site or collection of related server computers. For example, names like Payroll or Domestic Sales can be used.

3. The Server Group Description field is optional, type the desired description. All characters are acceptable and there is no limit to the number of characters you can enter.

4. Click **New Server** to add a new server if necessary. See 1.19 Adding a New Server on page 48 for details on how to add a new server.

5. Select **Perform Replication Now**. The server groups and server are copied to each of the defined Spheras Storage Director servers. It is recommended that you always enable this option.

6. Click **OK** or press the Enter key. The new server group appears in the navigation area.

1.17 Changing Server Group Properties

To view or change your Server Group Properties:

1. Select a server group from the navigation area or content area.

2. Click one of the following:

From the menu bar, click Group->Properties

From the tool bar, click Properties

Right-click the server group and select Properties

1.18 Deleting Server Groups

Deleting a Server Group allows you to delete any previously defined Server Group. However, you cannot delete the system-defined **All Servers** group, as this is part of the root directory of the navigation area. To delete a Server Group, follow these steps:

1. Select the Group that you want to delete in the navigation area

2. Click **Group->Delete** on the Menu bar and the Delete Group dialog box opens

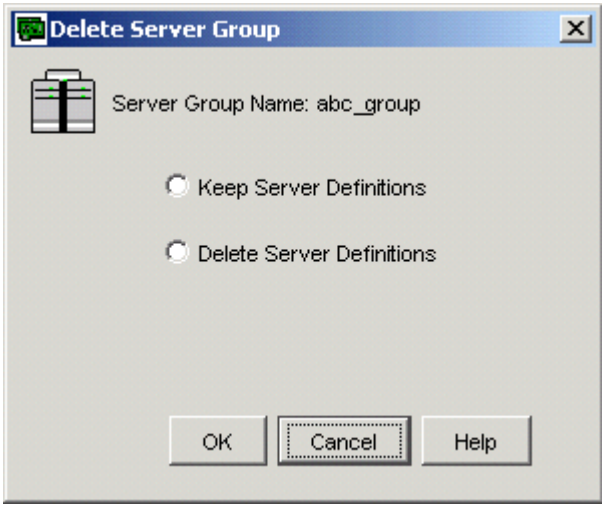

Figure 23 Delete a Server Group

3. Select one of the following delete options:

Keep Server Definitions - Keeps the server definitions as part of the All Servers group

Delete Server Definitions - If the servers in the selected groups are not part of other server groups, the server definitions are deleted.

4. Click **OK** or press the Enter key. The deleted group is removed from the navigation area.

1.19 Adding a New Server

This section describes how to define a new server. A maximum of 100 servers can be created per server group.

To define a new server:

1. Click **Server->New** on the menu bar, the New Server dialog box opens.

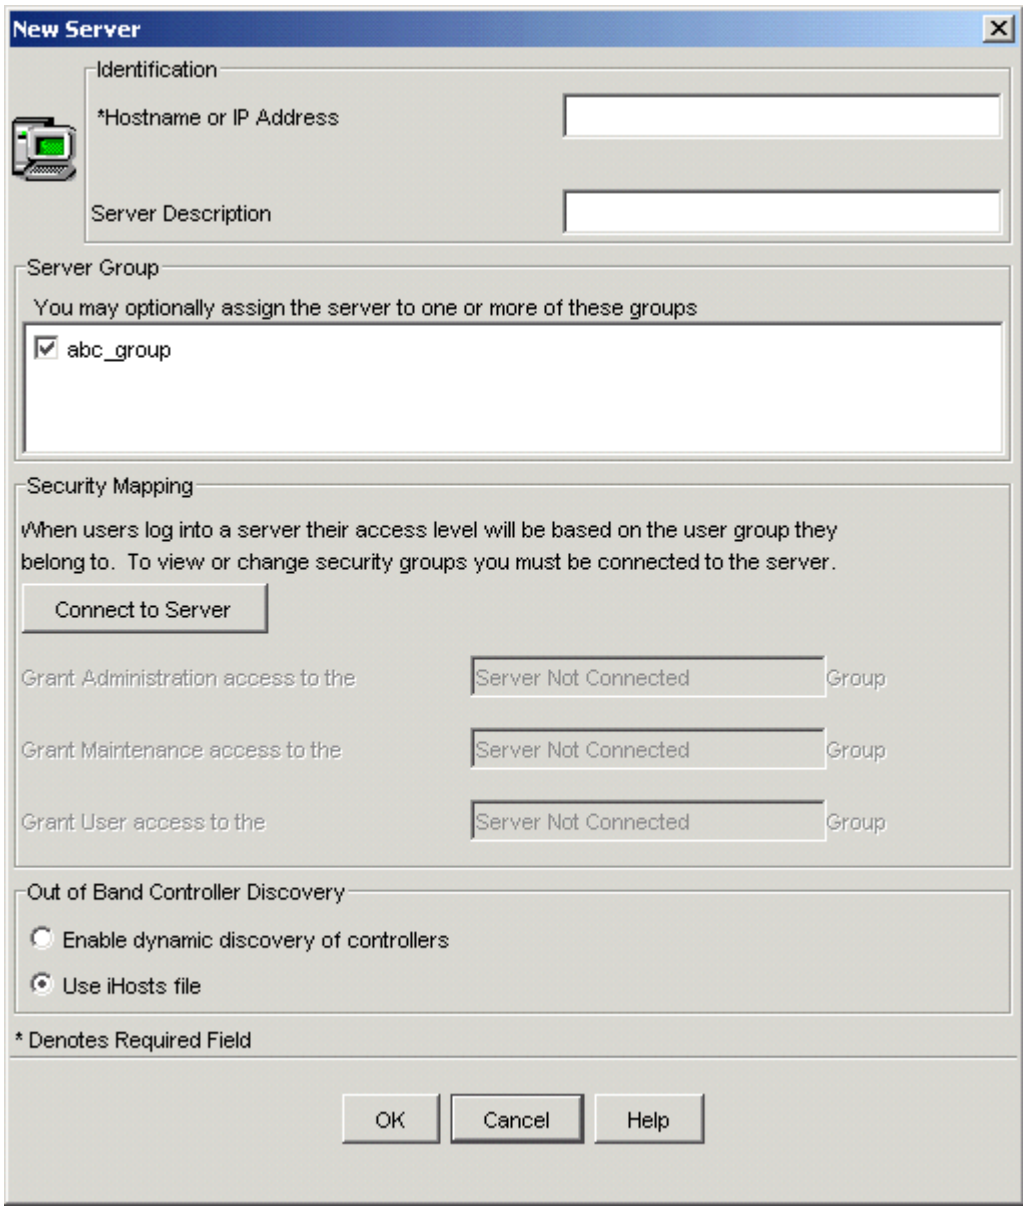

Figure 24 New Server Dialog Box

2. Type the IP Address or Host Name of the server you want to add. All characters are acceptable and there is no limit to the number of characters, you can enter.

3. The New Server Description field is optional. All characters are acceptable and there is no limit, to the number of characters you can enter.

4. Click **Connect to Server**. Once you are connected, the default security groups that you can enter on the new server are displayed in the Grant Administration Access fields. Change the Access fields as desired. See 1.13 Security access levels on page 45 for more information.

5. Select the type of Out of Band Controller Discovery that you want Spheras Storage Director to use. If you know the IP addresses of the controllers you want to communicate with, you can disable dynamic discovery of controllers and enable the iHosts file. See 1.24 Editing the iHosts File on page 52 for information on how to add controller IP

addresses. If you do not know the IP addresses of the controllers or the controller does not have an IP address, see 8 Select IP settings on page 88 for information on how to set up IP addresses.

- Dynamic discovery of controllers - allows the server to automatically discover all controllers that share its subnet.

- iHosts file (default) - only the IP addresses of the controllers that are listed in the iHosts file will be discovered. See 1.24 Editing the iHosts File on page 52.

Note: If a controller is connected to a host with both Fibre Channel and Ethernet (TCP/IP) connections, the Ethernet connection takes precedence and the Fibre Channel connection is not used.

6. Click **OK** or press the Enter key. The new server appears in the navigation area.

Notes:

1. When you log onto a server, your access is based on the user group in which you belong. To view or change security options, you must be connected to that server.

2. If you are using the dynamic discovery feature over a large network that contains numerous controller-based storage subsystems, the Spheras Storage Director interface may take a long time before the screen allows input from the user. Some of the reasons for the long delay include busy controllers and the available bandwidth on the network.

1.20 Changing Properties of Existing Servers

To change the properties of an existing server:

1. Select a server from the navigation area or the content area.

2. Click one of the following items:

   From the menu bar, click Server->Properties

   From the tool bar, click Properties

   Right-click on the server and select Properties

3. Assign a group for each access level.

   Grant Administration/root access to the Administrators group.

   Grant Maintenance/adm access to the Operators group.

   Grant User/users access to the Users group.

4. Select the type of Out of Band Controller Discovery that you want Spheras Storage Director to use. See step 5 on page 49.

Notes:

When you log onto a server, your access is based on the user group in which you belong. To view or change security options, you must be connected to that server.

If you are using the dynamic discovery feature over a large network that contains numerous controller-based storage subsystems, the Spheras Storage Director interface may take a long

time before the screen allows input from the user. Some of the reasons for the long delay include busy controllers and the available bandwidth on the network.

If a controller is connected to a host with both Fibre Channel and Ethernet (TCP/IP) connections, the Ethernet connection takes precedence and the Fibre Channel connection is not used.

1.21 Deleting a Server

Deleting a server, results in loss of the ability to manage any controllers that are visible to that server. Deleting a server also removes it from any server groups of which the server is part. Data that resides on the server is not affected; only the ability to manage controllers that are visible from the server is affected.

To delete a server, follow these steps:

1. Select a server from the navigation area or the content area.

2. From the menu bar, click Server->Delete or Right-click the server->Delete. A warning dialog box opens.

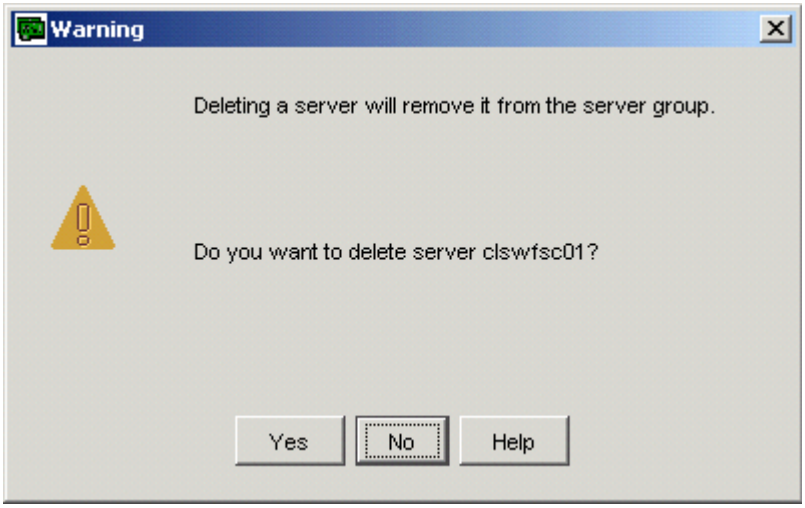

Figure 25 Deleting a Server Warning

3. Click **Yes** to delete the selected server or click **No** to close the dialog box without deleting a server.

1.22 Logging off of Spheras Storage Director

You should leave the Spheras Storage Director running as long as there are servers that you want to monitor or configure. If you exit Spheras Storage Director, you will be unable to receive events, or be informed, of errors or status, unless you restart the Spheras Storage Director.

To exit, click File->Close in the browser you are using.

1.23 Replicating a Server

Replication is similar to making a copy. This process sends a list of all the servers and the server groups that are seen on the client workstation to each of those servers. The servers receiving this list use it to update the current servers and server group lists. Therefore, if you run Spheras Storage Director from any of the servers, you see the same list of servers and server groups.

Replication provides an easy way to keep a consistent set of servers and server groups that are defined across all Spheras Storage Director Servers without having to configure each server independently.

To replicate an entire set of servers and server groups, follow these steps:

1. From the menu bar click Group->Replicate or right-click Group->Replicate in the navigation area. The Replicate Information dialog box opens confirming the replication.

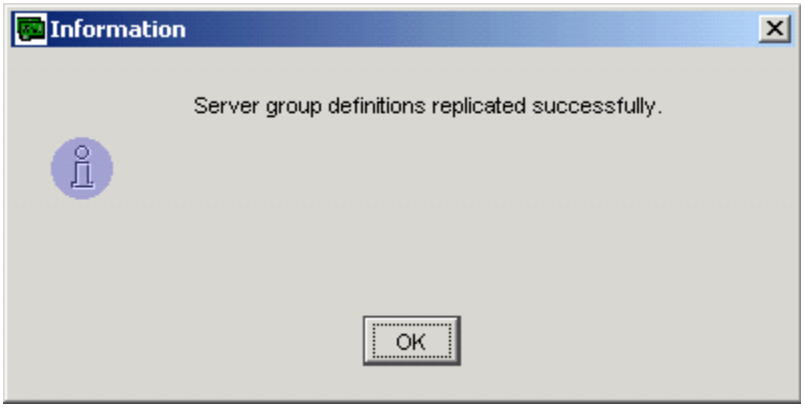

Figure 26 Replicate Information Dialog Box

2. Click OK.

Replication can also be done when creating a new server group. See 1.19 Adding a New Server on page 48.

1.24 Editing the iHosts File

The iHosts file contains the IP addresses for controllers that are not a part of the subnet of your Spheras Storage Director Server. This is referred to as out of band management. Editing an iHost file allows you to add IP Addresses to the file. By adding these IP Addresses, your Spheras Storage Director Server will be able to discover other controllers.

Note: If a controller is connected to a host with both Fibre Channel and Ethernet (TCP/IP) connections, the Ethernet connection takes precedence and the Fibre Channel connection is not used.

To edit an iHosts File follow these steps.

1. Click Server->Edit iHosts File. The Edit iHosts File opens.

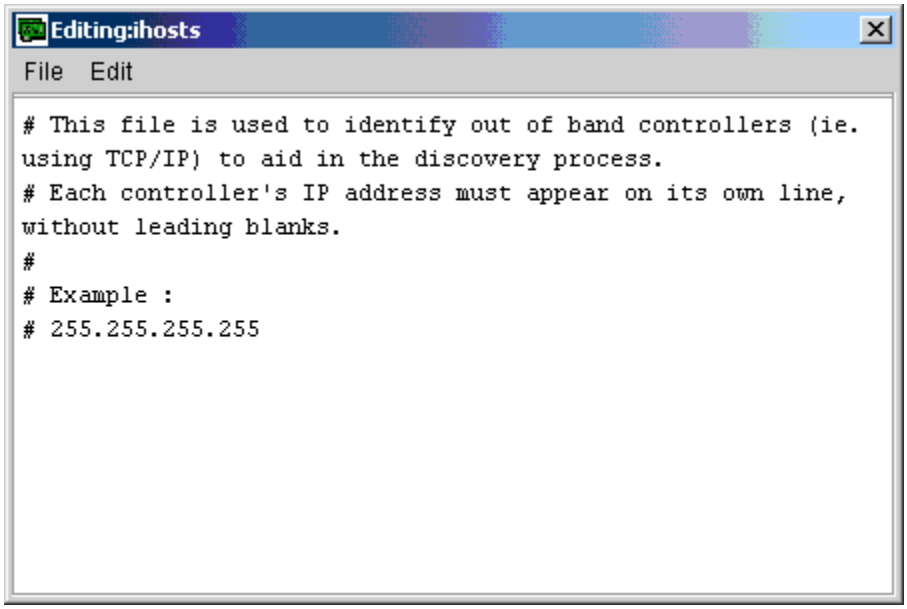

Figure 27 Edit an iHost File

2. Using the format of the Example given, type in your desired IP Addresses. You can also copy, cut, and paste your IP Address information. It is important to use the format shown in the example. Failure to do so may affect the ability of the server to discover controllers.

3. Remember to uncomment your IP addresses by removing the "#" symbol.

4. Select File->Save.

5. Select File->Exit to close the iHosts File. If you have not saved any changes, you are prompted to save your changes.

Chapter 4. Navigating the Spheras Storage Director Client

This section describes the navigation features and options that the Spheras Storage Director Client provides.

1.25 Components of the Spheras Storage Director Opening Screen

When starting up the Spheras Storage Director Client (with servers defined), the opening screen displays a menu bar, tool bar, path indicator, navigation area, content area, and status bar.

The following information describes the major components of the Spheras Storage Director Client windows.

1 **Menu Bar:** Each of the menus have several selections each. However, not all menu items are available all of the time. The menus that are displayed are dependent on your security access level and what you select in the navigation area or the content area. The contents of the menus and the functions that are performed by several of the selections are described later in this section.

2 **Toolbar:** The toolbar contains icons that start some of the common Spheras Storage Director functions. The identity and explanation of each icon is described later in this section.

3 **Path indicator:** Shows the path of the item selected from the navigation area. For example, if you choose an array from the navigation window, the path shows the following (where X is the device number): AllServers/MyServer/Controllers/ControllerX/Arrays/ArrayX

4 **Navigation area:** Provides a visual tree of the controller system. The following items are shown in the navigation area:

- Client Console
- Server Groups
- All Servers (predefined server group)
- Other server groups defined by the user
- Controllers
- Events
- Task Monitor
- Logical Drives
- Arrays
- Physical Drives
- BBUs

- Enclosures

Items that are not in an optimal state display a status indicator icon. The status indicators are described later in this section.

5 **Content area:** Displays the contents of the item selected in the navigation area (for example: server, controller, or array). The items that display vary according to the item that you select. For example, if you select a server, the content area displays all of the controllers that are on the server, the status of each controller, a brief description, and information about the event log on that server. If you select an item in the Physical View, the content area displays a graphic image of the enclosure that contains the controller and a status box that shows the contents and status of the selected controller. If you select an item in the Text View, a summary of the tasks, logical and physical drives, arrays, BBUs, and enclosures associated with the controller is displayed in a table format. The content area is described further in various sections of this manual.

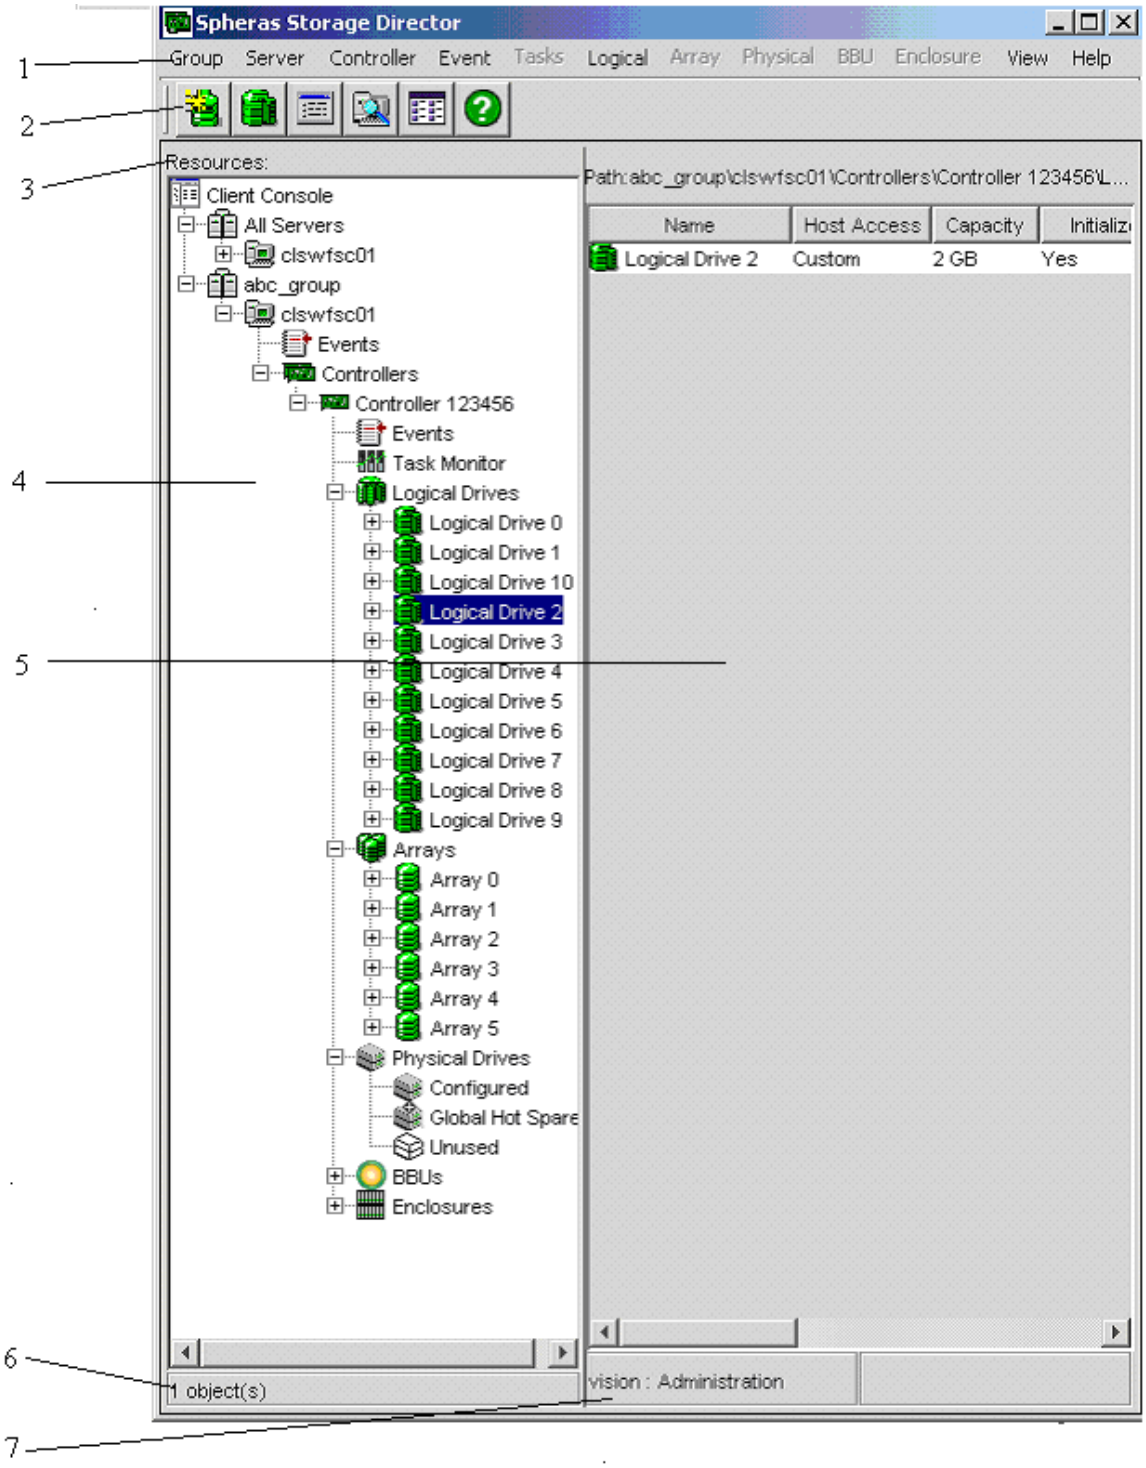

Figure 28 Spheras Storage Director Opening Screen

6 **Object status:** Shows the number of objects or selected objects in the content area.

7 **Logon status:** Shows the user name and the current access level.

1.26 Status Icons

The following icons display the status of the controller system. These icons will display on the node tree in the navigation area and in the content area. Some icons display only when using the Physical View.

Optimal State Icon	Warning State Icon	Error State Icon
		
All resources are optimal	At least one resource is in a Warning State	An error has occurred that did not cause data loss, usually a software error.
Serious State Icon	**Critical State Icon**	**Physical Drive Rebuild State**
		
At least one resource is in a Serious state.	At least one resource is in a Critical state.	The physical drive is in a rebuild state (Physical View Only).
Physical Drive Failed State	**Physical Drive Unconfigured State**	**Hot Spare Drive**
		
The physical drive is in a failed state (Physical View Only).	The physical drive is in an unconfigured state (Physical View Only).	The physical drive is a hot spare (Physical View Only).

Physical Drive Normal State	Physical Drive Not Selected State	Physical Drive Suspected Failed
		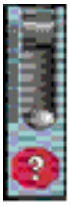
The physical drive is in a normal running state (Physical View Only).	The physical drive is in a not selected state (Physical View Only). This means that a physical drive is in the slot, but not part of the resource	The physical drive may have failed due to transient errors.

	or set of resources being viewed.
Physical Drive Not Supported  The physical drive is in a not supported state (Physical View only). This icon is also used with enclosure status.	

1.27 Menu Bar and Menus

Items within the pull-down menus are enabled and disabled based on the level of access that you have, the currently selected item in the navigation area and content area, and the state of the system. For example:

If you are logged on with User privileges and not logged onto any server, the only menu items available are Logon, View, and Help.

If you are logged on with Maintenance privileges, you see all the items from the User privileges plus the Replicate item in the Server Group menu.

If you are logged on with Administrator privileges, you see all the items from the Maintenance privileges plus the New Server and New Server Group menu.

Figure 29 Menu bar

4.1.1 Group Menu

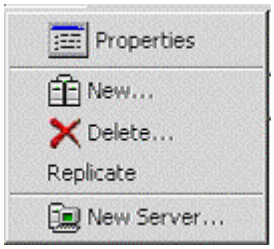

Figure 30 Group menu

The Group Menu contains the following options:

Properties: Displays the Server Group name, description, and server members and non-members.

New: Opens the New Server Group Window and allows you to create new server groups.

Delete: Opens the Delete Server Group dialog box and allows you to delete a server group. Also, gives you the option to delete all servers that were members of this server group.

Replicate: Replicates the entire set of server groups and servers to each known Spheras Storage Director Server.

New Server: Adds a server to the navigation node tree.

4.1.2 Server Menu

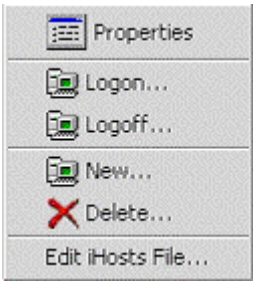

Figure 31 Server menu

The Server Menu contains the following options:

Properties: Displays the Server identification, description, other server groups, security mapping, and controllers for out-of-band discovery.

Logon: Opens the logon window.

Logoff: Opens the logoff window. (You need to logoff of a server if you want to logon again with different security permissions.)

New: Opens the New Server window and allows you to create new servers.

Delete: Deletes the selected server from the navigation node tree.

Edit iHosts File: Opens a text editor so you can edit the iHosts file. For more information on the iHosts file, see “Editing the iHosts File” on page 34.

4.1.3 Controller Menu

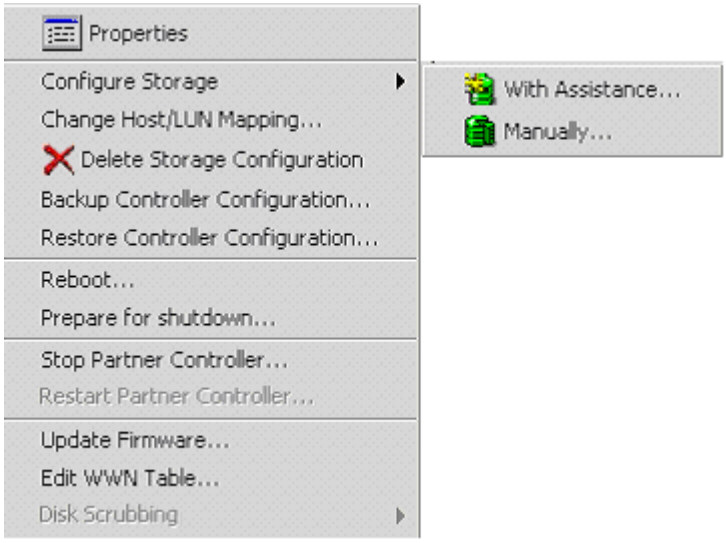

Figure 32 Controller menu

The Controller Menu contains the following options:

Properties: Allows you to change or modify various parameters for the selected controller.

Configure Unused Storage: Configures unused storage with assistance or manually.

Change Host/LUN mapping: Opens the Host/LUN mapping dialog box for you to modify the mapping.

Delete Storage Configuration: Deletes the configuration of the selected controller.

Backup Controller Configuration: Saves the configuration.

Restore Controller Configuration: Downloads a controller configuration from a file to the controller.

Reboot: Reboots the controller and brings the controllers out of conservative cache.

Shutdown: Sets the controller into conservative cache mode and prepares for a reboot. See "Shutting Down a Controller" on page 115 for more information.

Stop Partner Controller: If a running partner controller exists, this feature forces the partner controller to go Offline and puts the system in a failover state. This item is disabled if any of the following conditions exist:

- If a partner controller is not running
- If a long operation task is running.

Restart Partner Controller: If a running partner controller has stopped, this feature brings the partner controller back online and resumes

Duplex mode. This item is disabled if any of the following conditions exist:

- If a running partner controller exists
- If a long operation task is running.

Update Firmware: Updates the selected controller's firmware.

Edit WWN Table: Contains entries for every host that has ever accessed the controller. Allows you to remove old host references from the table.

Disk Scrubbing: Causes all RAID configured drives and hot spares (no un-configured drives) to be checked periodically, to identify potential problems before data is written to or accessed on those drives. Disk scrubbing is not enabled by default. See "Disk Scrubbing" on page 107 for more information.

4.1.4 Event Menu

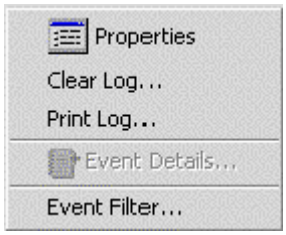

Figure 33 Event menu

The Event Menu contains the following options:

Properties: Allows you to view or change event log properties and event handling settings.

Clear Log: Clears the event log of the selected server.

Print Log: Opens a browser window to allow printing of the event log.

Event Details: Displays information about the generated events for the selected server.

Event Filter: Controls the level of severity of the messages that are displayed.

4.1.5 Tasks Menu

Figure 34 Task menu

The Task Menu has the following option:

Cancel: Cancels the selected task, such as a rebuild or a consistency check. Controller-initiated tasks, such as a Background Initialization, cannot be cancelled.

4.1.6 Logical Menu

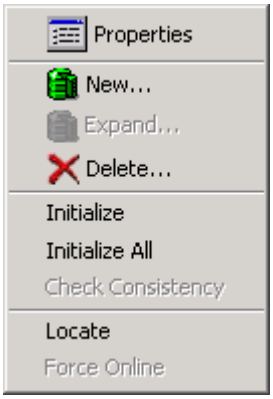

Figure 35 Logical menu

The Logical Menu has the following options:

Properties: Allows you to view or change logical drive properties.

New: Opens the Manual Configuration Wizard and allows you to create a new logical drive.

Expand: Opens the Manual Configuration Wizard and allows you to expand an existing logical drive.

Delete: Deletes the selected logical drive.

Initialize: Initializes the selected logical drives. See step 9 on page 72 for details about foreground and background initialization.

Initialize All: Initializes all logical drives on the controller.

Check Consistency: Performs a consistency (parity) check on the selected logical drive(s).

Locate: Blinks the LEDs of the physical drives contained by the selected logical drive.

Force Online: Forces the selected logical drive online. This option is only enabled if the drive is offline. This occurs when one or more physical drives, are offline.

4.1.7 Array Menu

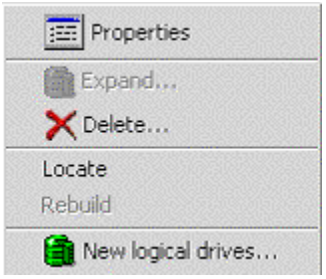

Figure 36 Array menu

The Array Menu has the following options:

Properties: Allows you to view array properties.

Expand: Opens the Manual Configuration Wizard and allows expansion of the selected array.

Delete: Deletes the selected array.

Locate: Blinks the LED(s) of the physical drive(s) contained by the selected array.

Rebuild: Rebuilds a critical logical drive when there is a hot spare available.

New Logical Drive: Opens the Manual Configuration Wizard and allows you to configure a new logical drive.

4.1.8 Physical Menu

Figure 37 Physical menu

The Physical Menu contains the following options:

Properties: Allows you to view physical drive properties.

Make Hot Spare: Makes the selected physical drive a hot spare.

Remove Hot Spare: Returns the physical drive to the unused status.

Locate: Blinks the LED(s) of the selected physical drive(s).

Spin Down: Prepares the physical drive for removal.

Force Online/Offline: Allows you to force a physical drive online or offline.

4.1.9 BBU Menu

Figure 38 BBU menu

The BBU Menu contains the following options:

Properties: Allows you to view or change BBU properties and launch a battery quick test if supported by the firmware.

Recondition Battery: Allows you to recondition the battery by automatically discharging and then recharging the battery.

4.1.10 Enclosure Menu

Figure 39 Enclosure menu

The Enclosure menu has the following option:

Properties: Allows you to view or change enclosure properties.

Locate: Blinks the LED(s) of the physical drive(s) contained by the selected enclosure.

4.1.11 View Menu

Figure 40 View menu

The View Menu contains the following options:

Text: Enabled when there is a text based view for the node selected in the navigation area.

Physical: Enabled when there is a physical based view for the node selected in the navigation area. Physical views are available for physical drives, unused drives, failed drives, and hot spares. This view is also available for an individual controller, logical drive or array when selected in the navigation area.

Tool Bar: Displays/hides the Tool Bar.

Status Bar: Displays/hides the Status Bar.

Refresh: Refreshes the current navigation or content areas information.

Repair Action: Enabled when a resource is in a critical or serious state. Displays problem explanation/recovery information from the online Help in a separate window.

4.1.12 Help Menu

Figure 41 Help menu

The Help Menu contains the following options:

Contents: Displays the contents of the online Help system in a separate window.

Index: Displays the index of the online Help system in a separate window.

Page Help: Displays information specifically related to the information currently being displayed in the navigation and content areas.

About: Opens a dialog box that displays the Spheras Storage Director version and copyright information.

1.28 Toolbar and Toolbar Icons

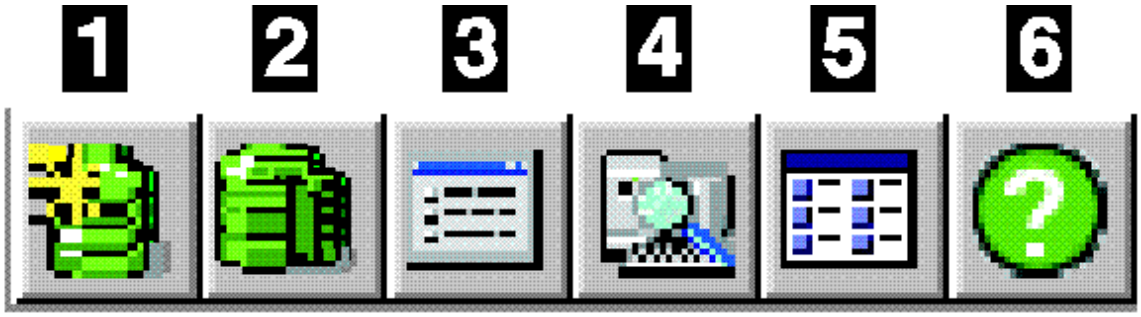

Figure 42 Toolbar

The Toolbar contains icons for the following functions:

1 Opens the Configure Unused Storage with Assistance Wizard.

2 Opens the Configure Unused Storage Manually Wizard.

3 Opens the property page of the selected object.

4 Refreshes the navigation and content area views.

5 Toggles between the text and physical views.

6 Displays Help for what is currently displayed in the content area.

1.29 Exiting Spheras Storage Director

We recommend leaving the Spheras Storage Director Client running as long as there are servers you wish to monitor or configure. If you do exit, you will be unable to receive events from the Spheras Storage Director Server and you will not be informed of errors or status unless you restart the Spheras Storage Director Client.

To exit the Spheras Storage Director, click the File->Close menu option of your browser or close the Spheras Storage Director Client screen.

Chapter 5. Configuration Activities

The following configuration activities are included in this section:

- Setting Event Log Properties
- Setting and Modifying Controller Properties on page 75
- Creating Storage Configurations on page 90
- Spanning in Spheras Storage Director on page 103
- Expanding a Logical Drive on page 103
- Deleting a Logical Drive on page 105
- Expanding an Array on page 105
- Deleting an Array on page 108
- Restoring a Controller Configuration from a Previously Saved File on page 109
- Backing Up a Controller Configuration on page 111

1.30 Setting Event Log Properties

The Event Log Properties dialog box provides you with options for setting desired event logging capabilities. This includes setting application launching to notify you of events that need immediate attention. For example:

- You may want to be notified immediately if a critical error occurs, so you would configure the system to launch an application to page you.
- If the log file is growing too quickly, you may want to remove all optimal events from it.

Note: The Client must be running in order to receive alarms from notification applications.

Use the following steps to configure Event Settings:

Open the Event Log Properties dialog box by selecting **Event-> Properties** on the menu bar.

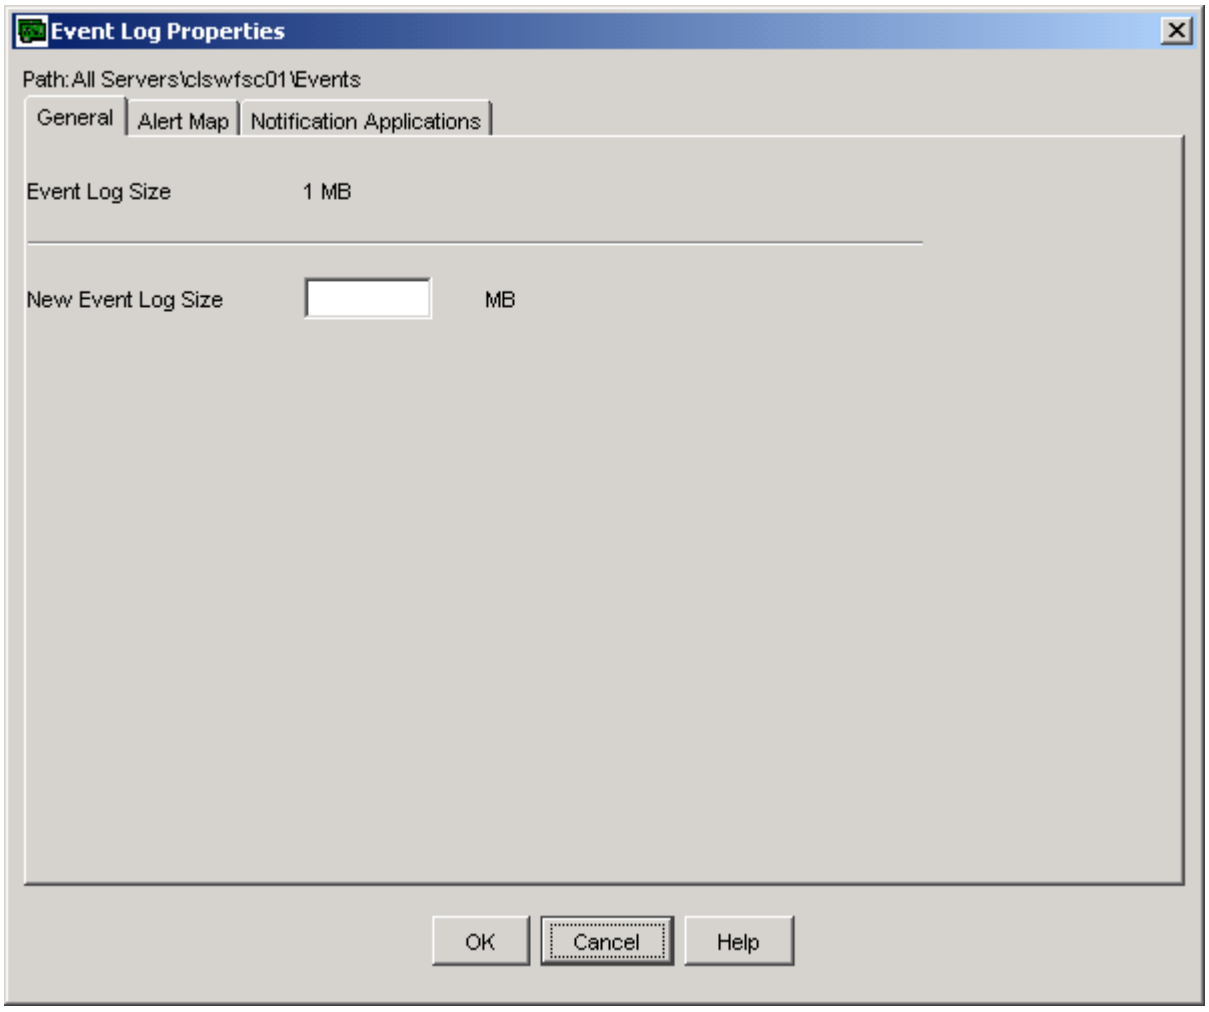

Figure 43 Event Log Properties Dialog Box: General Tab

In the Event Properties Dialog box, under the General tab, set the new event log size.

Select **Alert Map**.

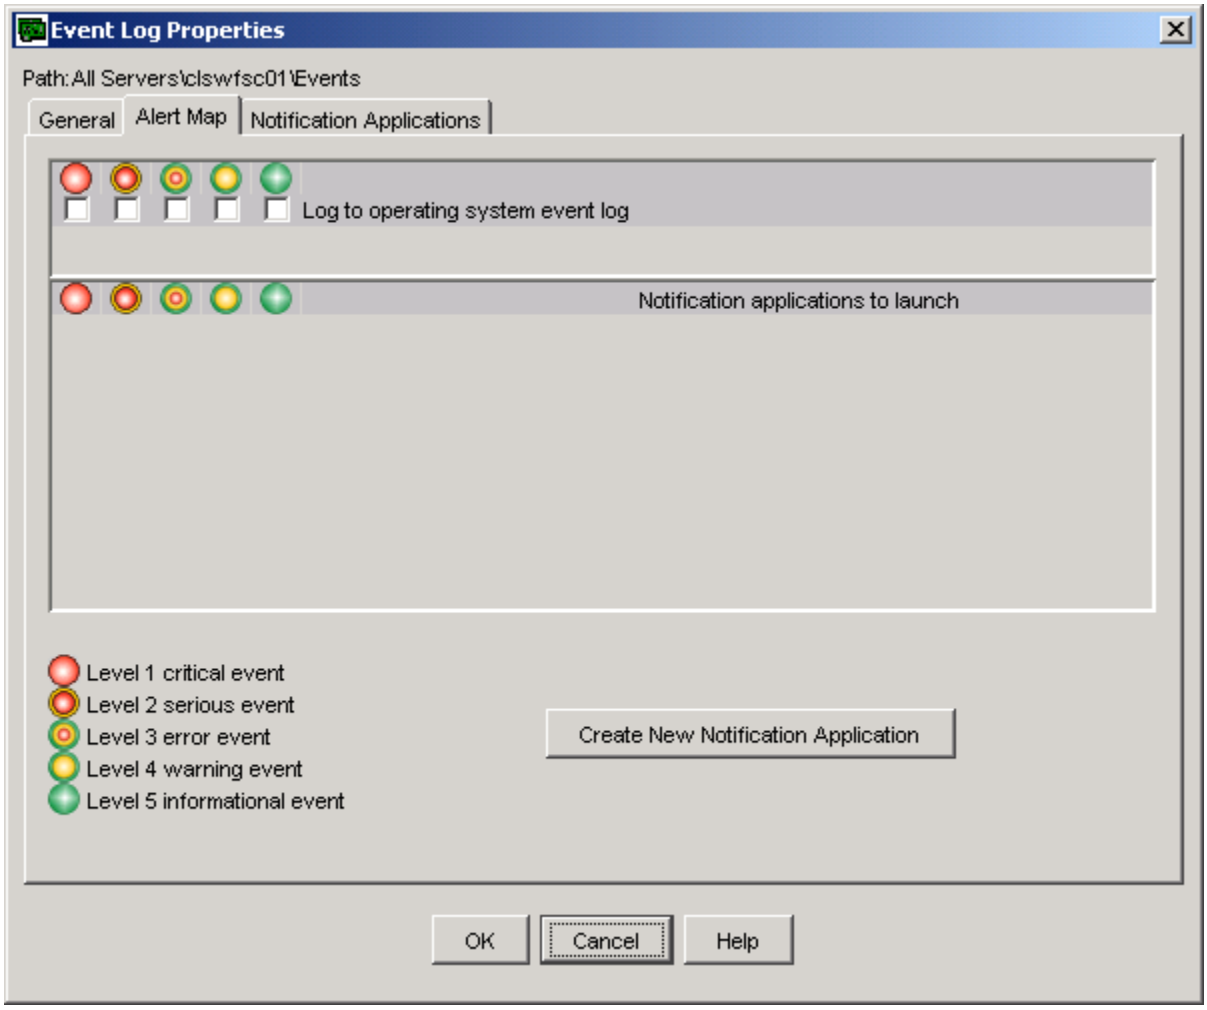

Figure 44 Event Log Properties Dialog Box: Alert Map Tab

Choose which event severity levels should be sent to the operating system event log by enabling or disabling the corresponding check boxes.

Enable or disable check boxes to launch the appropriate notification application corresponding to the event severity level.

Click **Create New Notification Application** to add an application to launch. The Add Notification Application dialog box opens.

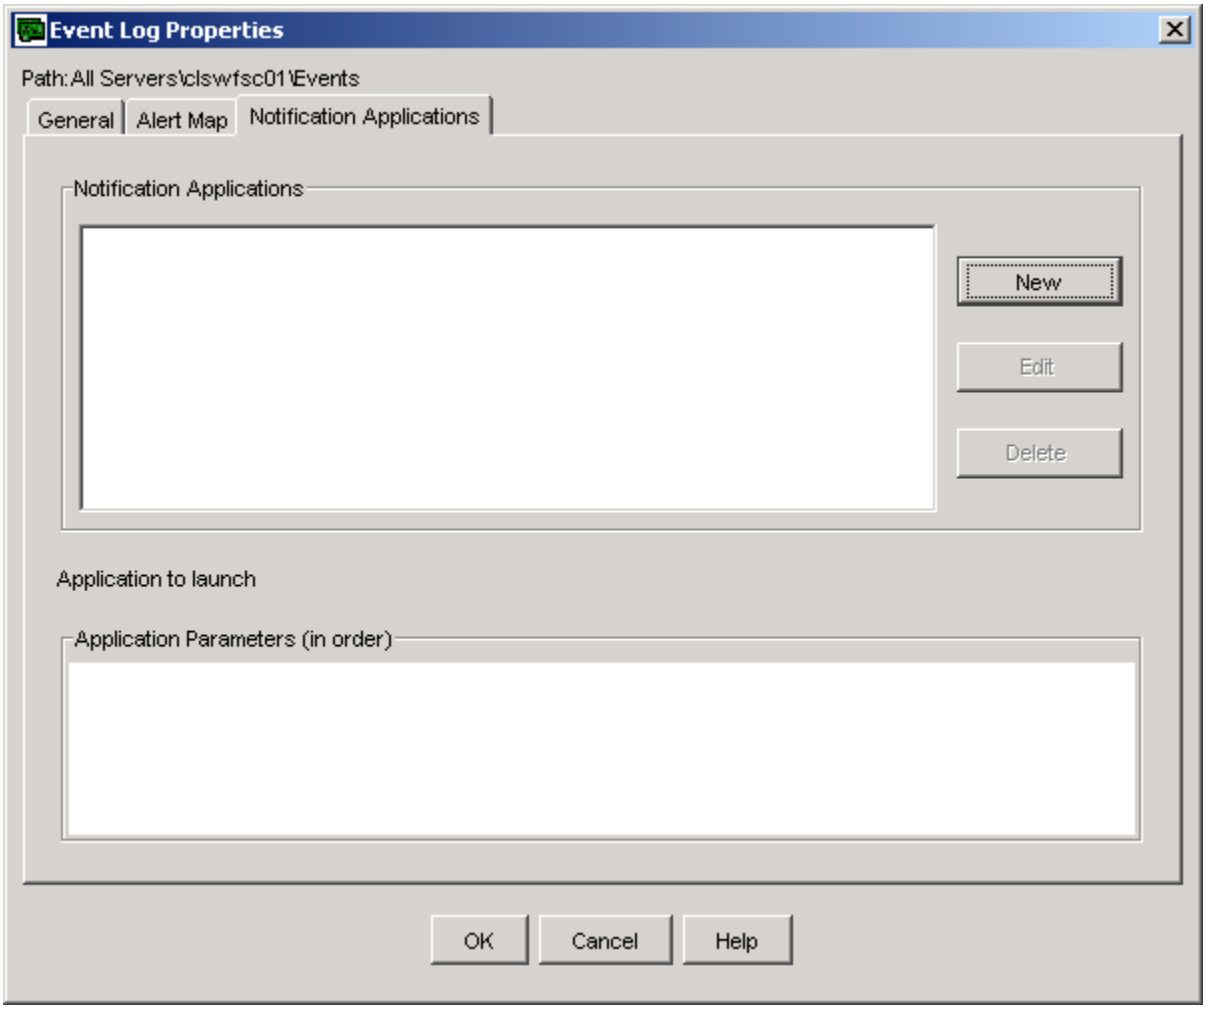

Figure 45 Event Properties Dialog Box: Notification Application Tab

Click **New** to add a new notification application.

Click **Edit** to edit an existing notification application.

Click **Delete** to delete a notification application.

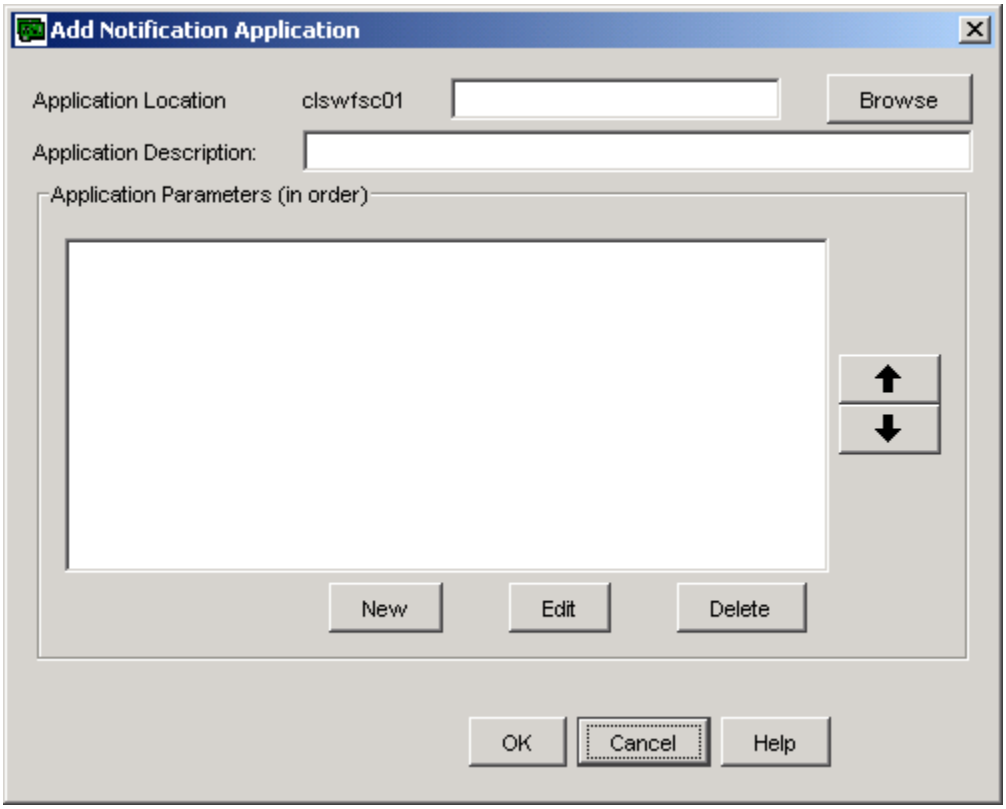

Figure 46 Add Notification Application Dialog Box

Type the name and path of the desired application to launch. If you don't remember the name or path of the application, click **Browse**. Use an existing application or create an application (such as a shell script) that sends e-mails, pages, or faxes.

Type a description of the application (for example: Bob's Pager).

Click **Delete** to delete an existing application parameter.

Click **Edit** to edit an existing application parameter.

Click **New** to add an Application Parameter.

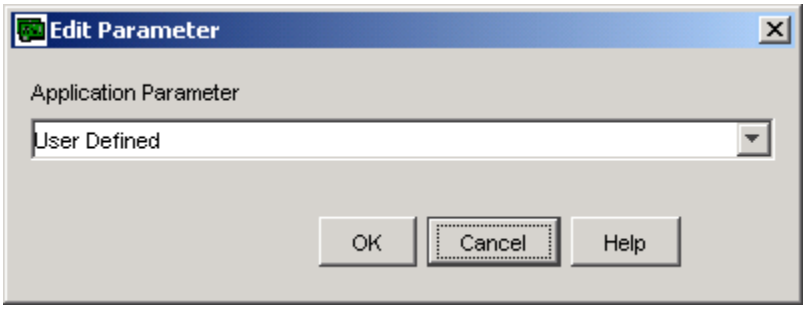

Figure 47 Add Parameter Dialog Box

Select one of the following system defined parameters to be passed to the application:

User Defined: You can type your own parameter string (for example: pager or email) by typing over the User Defined parameters.

Event Message ID: The number that is displayed in the ID column of the events table (for example, 389).

Event Message Severity: The severity level of the event.

Event Address: The contents of the address column of the events table (for example: ctl:0 chn:0 tgt:1).

Element ID: A text string that contains unique identifiers for each of the resources associated with the error. This is the controller World Wide Name. The enclosure World Wide Name, or the physical drive serial number, or both may also be displayed. See the serial number section of the Event Details dialog for the contents of this string.

Click **OK** or press the Enter key to save your settings.

Click **OK** or press the Enter key of the Add Notification Application dialog box to save your settings.

1.31 Launching Application Alerts with Spheras Storage Director Server

The Spheras Storage Director Client handles the launching of application alerts when controller events that match user specified severities occur. On the Spheras Storage Director Server side, you can modify sections of the cfg.ini file to launch application alerts from the server. The modification of the cfg.ini must be done with a text editor. To edit the cfg.ini file, follow these steps:

Locate the cfg.ini file. The file can is located in your Spheras Storage Director installation directory.

Using a text editor, open the cfg.ini file. An example file follows:

```
;EventAlertApplication#description
;This section holds the contents of the application used to notify
;the user of specified event severity.
;The SeverityCombination section is a bitmap, filled out as follows:
;  INFORMATION_EVENT = 1
;  WARNING_EVENT = 2
;  ERROR_EVENT = 4
;  SERIOUS_EVENT = 8
;  CRITICAL_EVENT = 16
;The values may be OR'd together to support multiple severity levels
;(eg. To launch an application whenever a
;SERIOUS_EVENT (8) or a CRITICAL_EVENT(16) occurs use a SeverityCombination
;value of 24).
;The user defined parameters must be separated by spaces
;
;The example below describes an Event Alert Application called "Page On Call
;Manager" which
;launches the application c:\pagemgr.bat whenever a critical or serious event
;is generated, and passes it 4 user
;generated parameters :
; AT&F&C1&D2L1
; ATDT
; 555-1234
; #
;and 4 system generated parameters
```

; EVENT_ID (the ID number representing that event)
; EVENT_SEVERITY (the severity level of the event, as defined above)
; EVENT_ADDRESS (address information for devices involved in the event
;(eg. "ctl: 0, chn: 0, tgt: 0")
; EVENT_ELEMENT_ID (world wide names and/or serial numbers of devices involved
;in the event (eg. "Controller world wide name:2000 0080e5 11ea07")

Locate the "Page On Call Manager" section in the file. Format this section of the cfg.ini file by using the following information as an example. The example below describes an Event Alert Application, which launches the application c:\pagemgr.bat and passes it the user-defined parameters:

[EventAlertApplication#Page On Call Manager]
UserDefinedParameters=AT&F&C1&D2L1 ATDT 555-1234 # EVENT_ID
EVENT_SEVERITY
EVENT_ADDRESS EVENT_ELEMENT_ID
SeverityCombination=24
ApplicationName=C:\pagemgr.bat

When you are finished editing the appropriate sections of the cfg.ini file, click **File->Save** to save your changes.

Click **File->Exit** to close the text editor.

1.32 Setting and Modifying Controller Properties

The Controller Properties dialog box allows you to view or change controller properties. Each Disk Array Controller is shipped from the factory with default settings, which have been found to work well in a majority of applications and environments. However, no two environments are the same, and you may want to modify certain controller properties before preparing your array. Follow the steps below to view or change your controller properties.

1. Open the Controller Properties; click **Controller->Properties** on the menu bar.

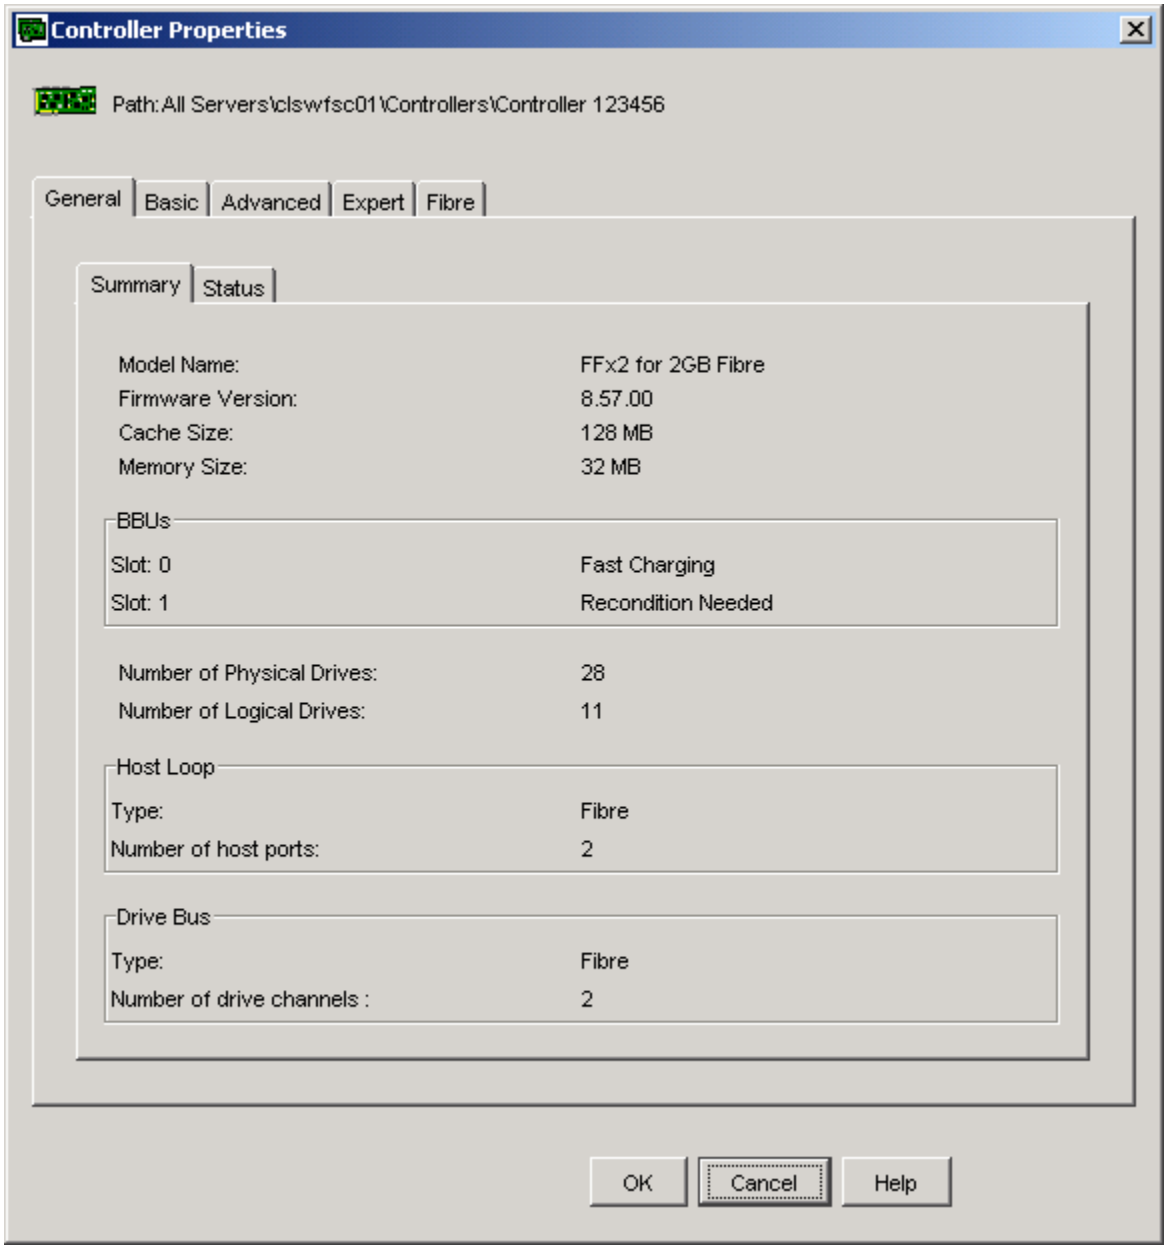

Figure 48 Controller Properties Dialog Box: General: Summary Tab

The Summary tab displays a summary of installed controller features and the status of the controller.

2. Select **Status**.

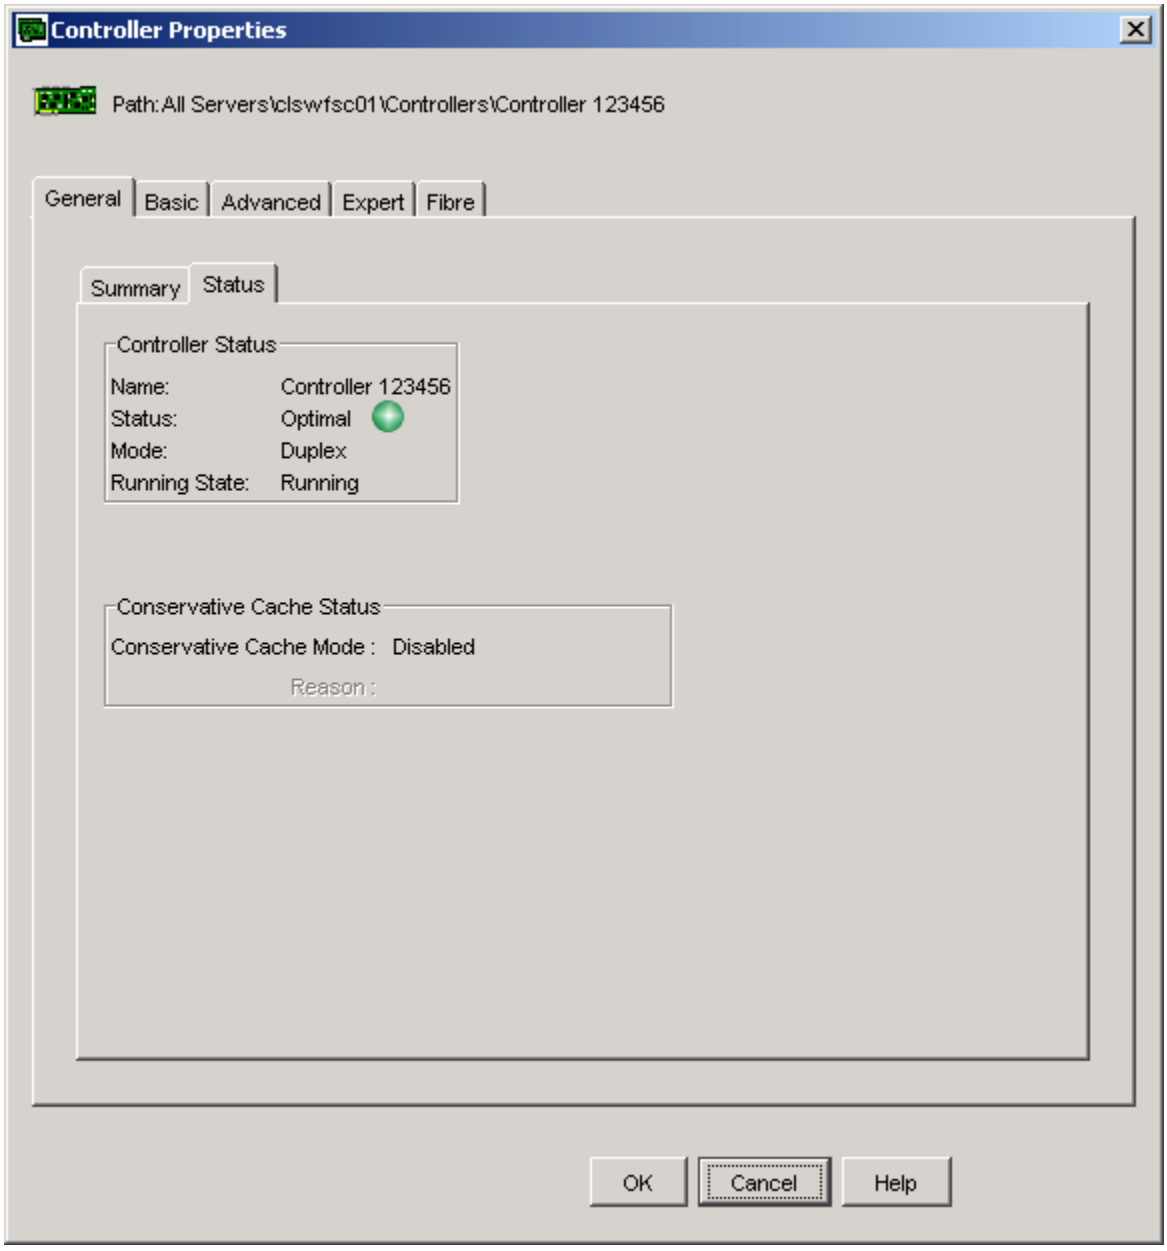

Figure 49 Controller Properties Dialog Box: General: Status Tab

The Status Tab displays the Name, Status, and Mode of the controller. It also displays whether Conservative cache is active or not and the reason. If the controller is not in conservative cache mode, the Reason field is greyed out.

Conservative cache mode is an operating mode in which logical devices, or system drives, configured with the write-back write caching policy are treated as though they were configured for write-through operation and the cache is flushed.

3. Select **Basic**.

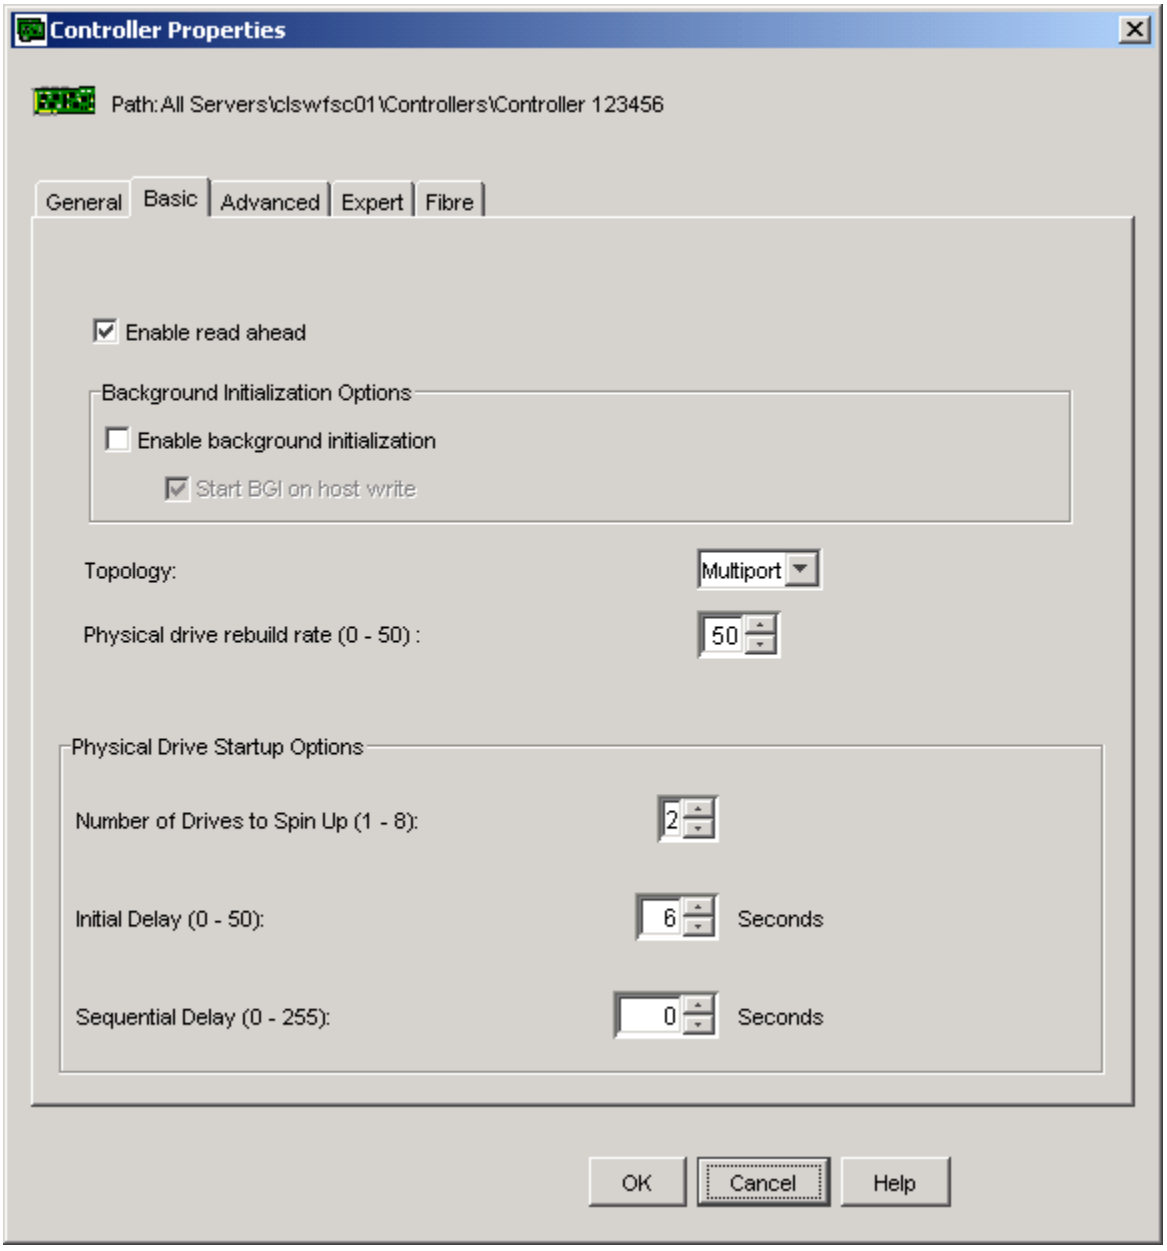

Figure 50 Controller Properties Dialog Box: Basic

In the Basic tab, you have the following options:

Enable read ahead: When enabled, reads data to the next stripe boundary (based on stripe size), and retains in cache any data beyond the size of the actual read request.

The Read Ahead function improves data retrieval performance by allowing the controller to read into cache a full stripe of data at a time. This greatly improves the percentage of cache hits. For small transfers, the read ahead algorithm helps with performance.

For example, if the stripe size is set to 8 KB and the host requests 1 KB of data, when this function is enabled the controller reads the full 8 KB. When the host requests the next 1 KB block, that data is already in the controller's cache. This function should remain enabled

during normal controller operation. A controller reset is required before this parameter takes effect.

Enable background initialization: When enabled, this option allows a single logical drive or multiple logical drives to be initialized as a background operation while reads and writes are being performed. Background initialization is invoked automatically upon receipt of the first read/write command from the host. When the controller starts up or after a configuration changes, background initialization automatically initializes any un-initialized system drives one at a time.

The secondary option, "**Start BGI on host write**" determines when the background initialization begins. If this box is checked, initialization of new logical drives starts when the controller receives a write command from the host.

The controller firmware allows the background initialization to be paused during a controller reset. The background initialization progress is saved to NVRAM. After the controllers have restarted, the background initialization resumes close to where the pause occurred.

Background initialization is paused by any of the following operations: initialize system drives (foreground initialization); consistency check; automatic rebuild; array expansion; and logical drive expansion. If one of these operations is started while background initialization is running, background initialization is paused until the interrupting operation is complete. At that point, background initialization continues initializing the system drive that it was initializing from the point of interruption. If the interrupting operation resulted in the system drive being initialized, however, background initialization skips that system drive and searches for the next un-initialized system drive and begins initializing.

Performance is degraded during background initialization because every write operation requires access to all drives in the RAID group. Once all system drives have been initialized, performance is restored.

Topology: Select one of the following topologies from the drop-down list.

Multi-Port. All ports are active. This topology does not provide transparent failover/failback and requires an alternate path driver in the host.

Multi-TID. All ports are active. This topology is supported on SANbloc hardware. This topology provides transparent failover/failback and should not be used in conjunction with an alternate path driver.

Notes:

- A controller reset is required to make a change to this option take affect.
- Only those topologies that are supported on the controller are listed in the drop-down menu.
- Inactive Port. Controller 0 port 0 and controller 1 port 1 are active. During failover, the inactive port on the partner takes over for the active port on the failed controller.

Note:

- Master/Slave. Both ports on controller 0 are active. During failover, controller 1 assumes control.

Side affects of changing topologies:

- When changing from inactive port to master/slave, logical drives associated with controller 1 port 1 are reassigned to controller 0 port 1.

- When changing from master/slave to inactive port, port mappings associated with controller 0 port 1 are reassigned to controller 1 port 1.

- When changing from master/slave or inactive port to Multi-TID or MultiPort, the inactive ports are enabled and all port mapping remains unchanged.

- When changing from Multi-TID or MultiPort to master/slave or inactive, all mappings on the inactive ports are lost.

Physical drive rebuild rate: A rate of 50 assigns the maximum allowable resources to a drive rebuild, allowing the rebuild to proceed at its fastest. Lowering the number devotes more resources to I/Os and consequently slows the Rebuild process. This parameter takes effect immediately without resetting the controllers.

Change any of the following Physical Drive Startup Options:

- **Number of Drives to Spin Up (1 - 8):** Number of drives to start at the same time.

- **Initial Delay:** Number of seconds to wait before spinning up the first disk set.

- **Sequential delay:** Number of seconds between consecutive device spin-up cycles.

4. Select **Advanced**.

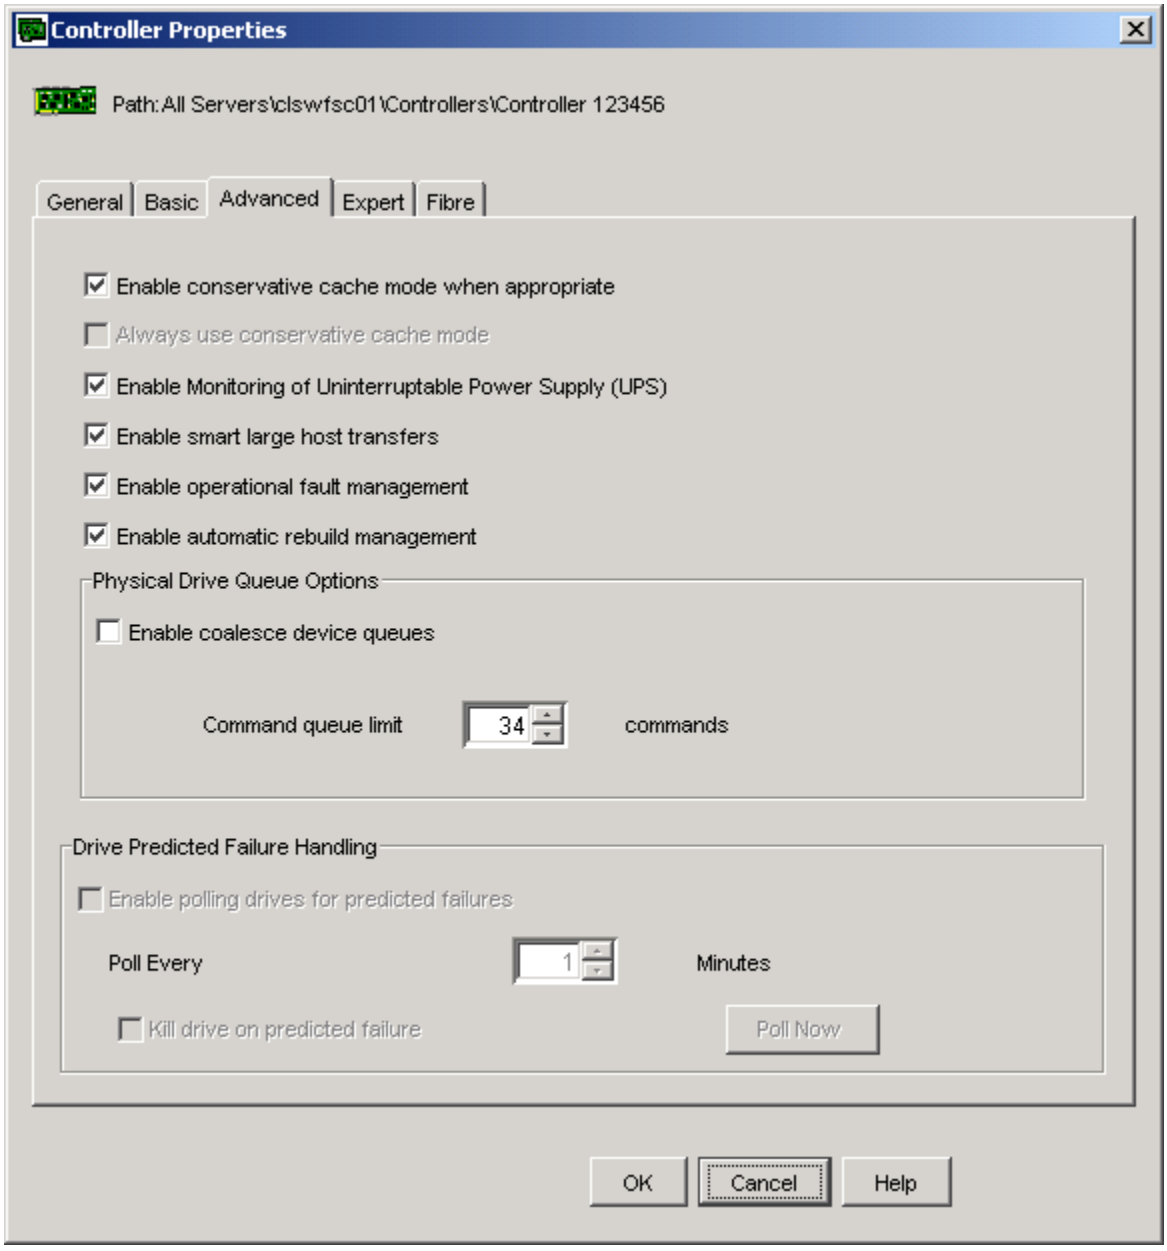

Figure 51 Controller Properties Dialog Box: Advanced

In the Advanced tab, you have the following options:

Conservative cache mode when appropriate: Turns on conservative cache if a controller failover condition exists, power supply fails, and other failed redundant parts, thus providing an extra degree of data safety.

Conservative cache mode is an operating mode in which logical devices, or system drives, configured with the write-back write caching policy are treated as though they were configured for write-through operation and the cache is flushed. When the condition causing conservative cache is resolved, the system drives are converted back to their original settings.

Conditions that cause conservative cache to turn on are:

- The Conservative Cache Mode Parameter is enabled for a dual-active controller configuration, and a controller failure has occurred.

- A MORE™ (Mylex® Online RAID Expansion) operation is initiated. During a MORE operation, the system automatically enters into conservative cache mode. After MORE completes, the system drives are restored to their original settings.

- A Power supply has failed (not the removal of a power supply).

- When an enclosure supporting redundant power (multiple power supply support) is operating on a single power supply.

- An out-of-limit temperature condition exists.

- The controller receives an indication of either an AC failure or a low battery from a UPS.

- The external hardware input signal has been asserted to force the controller into conservative cache mode.

Always use conservative cache mode: Forces the controller into conservative cache mode regardless of its current operating status. Conservative cache mode is used until you disable it by unselecting this option.

Monitoring of Uninterruptible Power Supply (UPS): Enables monitoring of uninterruptible power supply if the system enclosure supports it.

Smart large host transfers: When enabled, allows selection of Coalesce (fewer disconnects on large transfers) on host data transfers. When disabled, allows as available more disconnects on host data transfers. The Smart Large Host Transfers function takes effect for transfers larger than the stripe size.

Operational fault management:

Note: A controller reset is required to make a change to this option take affect.

This function should remain enabled during normal controller operation. Do not disable this function unless specifically instructed to do so as part of a trouble shooting diagnostic activity. Operational fault management allows the controller to take autonomous actions when a failure occurs. This option monitors and reports drive failures, background activity completion status, enclosure events, and so forth.

Automatic rebuild management: Works with SAF-TE and SES disk array enclosures to detect the removal of a failed drive and automatically perform a rebuild after installation of a replacement drive. Alternately, this function rebuilds to a drive that is assigned as a hot spare. The controller performs a rebuild operation automatically when a disk drive fails and the following conditions are true:

An online spare disk drive of identical or larger size is found attached to the same controller.

The failed disk drive is removed and replaced with a disk drive of identical or larger size.

All system drives that are dependent on the failed disk drive are configured as a redundant array; RAID 1, RAID 3, RAID 5, or RAID 0+1.

The Automatic Rebuild Management controller parameter is enabled.

The Operational Fault Management controller parameter is enabled.

Coalesce device queues: Joins data from adjacent I/Os into a single I/O to improve performance.

Command queue limit: Maximum allowed queue depth for tagged commands to each attached drive. Settings are from 1 (no tags) to 230.

Polling drives for predicted failures: When enabled (default), Predicted Drive Failures are written to the event log and drive error handling settings are enabled.

Poll every: Settings are from 1 to 254 minutes. The default value is 30 minutes.

Kill drive on predicted failure: When enabled, allows a physical drive to detect a Predicted Failure Analysis event and, if a spare exists, the spare drive is made offline and a rebuild begins immediately.

Poll now: When selected, the controller polls immediately and then reverts to the regular polling schedule. A message dialog box opens stating, "*Polling for predicted drive failures has started. Any events generated as a result of this polling can be viewed in the event log.*"

5. Select **Expert**.

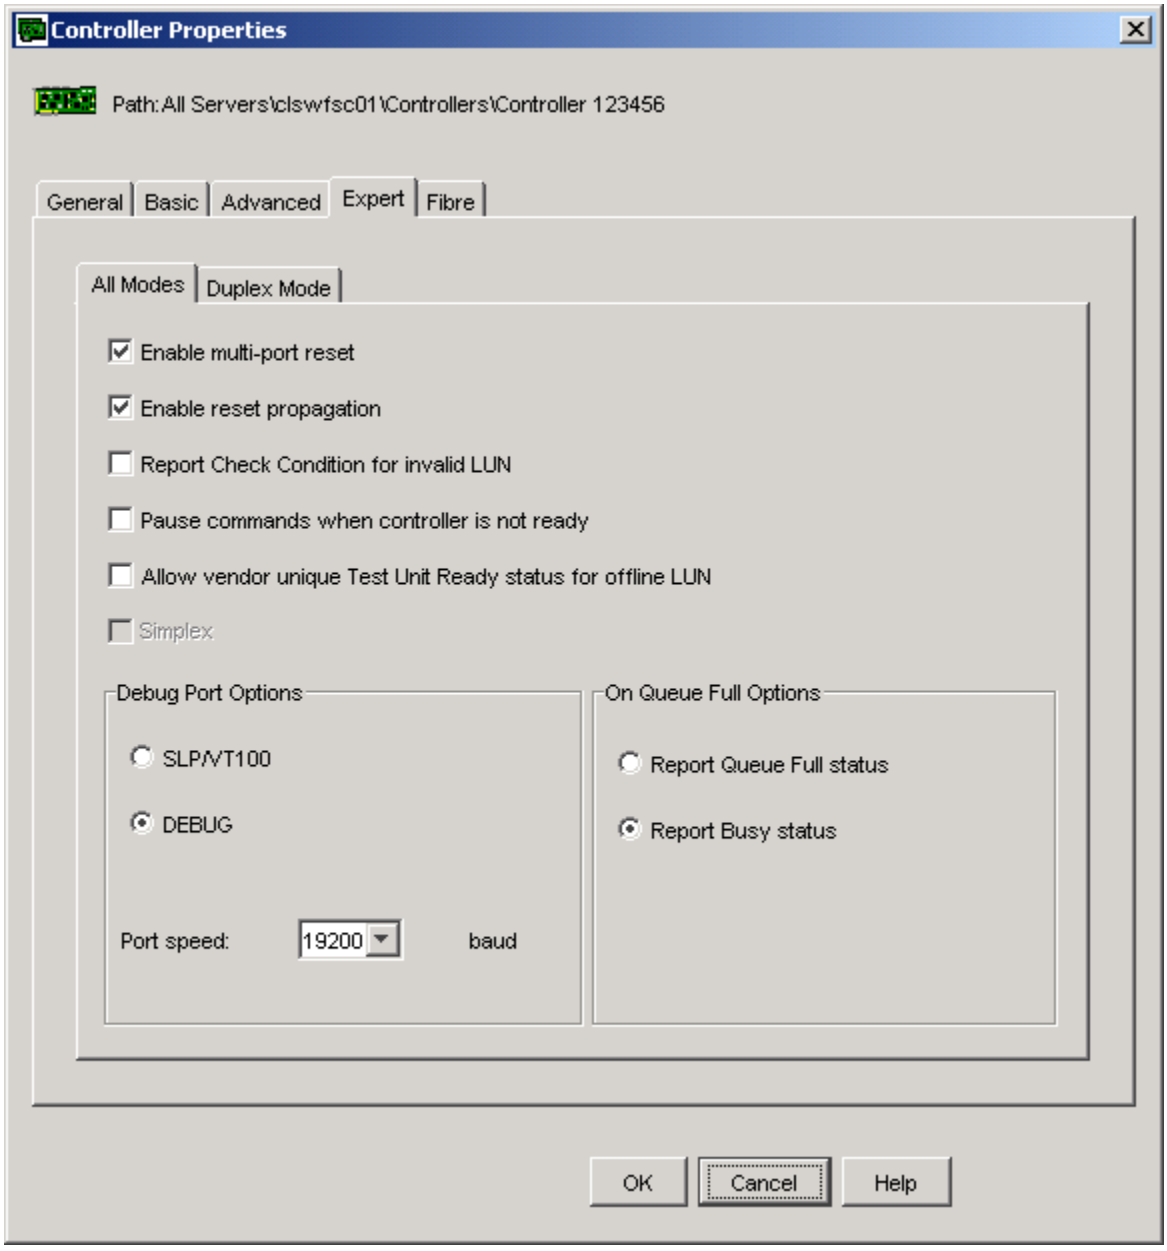

Figure 52 Controller Properties Dialog Box: Expert: All Modes Tab

In the All Modes Tab, you have the following options:

Multi-port reset: When enabled, allows execution of an internal reset on a port only if a logical drive has been reserved through that port.

Reset propagation: When enabled, a reset detected by one port may be propagated out of the other ports including ports on the partner controller.

Report check condition for invalid LUN:

Note: A controller reset is required to make a change to this option take affect.

When enabled, reports the check condition affecting the handling of the Inquiry command when the referenced LUN is invalid.

Pause commands when controller is not ready:

Note: A controller reset is required to make a change to this option take affect.

When enabled, if the controller is not fully started, certain commands encounter a pause.

Allow vendor unique Test Unit Ready status for offline LUN:

Note: A controller reset is required to make a change to this option take affect.

When enabled, causes all pending commands to be quietly dropped during failback. This is intended to help hosts that are confused with Queue Full.

Simplex:

Note: A controller reset is required to make a change to this option take affect.

A simplex controller handles all I/O requests. When selected, allows a controller pair in a failover state (one of the partner controllers has failed) to function as if it were a simplex controller. No warning events are generated regarding the failover state. This option is grayed out if the controller is simplex, or if the controller pair is in a dual active state.

Management Address: Use this field to specify the URL that is used by a Tivoli Storage Manager to launch a Spheras Storage Director client. This value is stored in the controller and is accessed by a Tivoli Storage Manager to get the URL. An example is: http:\\servername\ssd\ssd.htm

Note: The management address field is displayed only if it is supported by the controller. If the controller does not support this feature, the field is hidden and not just greyed out.

Debug Port Options:

Note: A controller reset is required to make a change to this option take affect.

Select SLP/VT100 emulation or Debug mode.

Debug Port Speed:

Note: A controller reset is required to make a change to this option take affect.

Sets the baud rate for the serial port from 2400 to 19200 baud. Speeds are available as supported by the selected controller.

On Queue Full Options:

Note: A controller reset is required to make a change to this option take affect.

When enabled, if a command is received and the controller detects a full queue, a *busy* status is returned rather than *queue full*. This helps hosts that do not accurately deal with *queue full* status.

6. Select Duplex Mode.

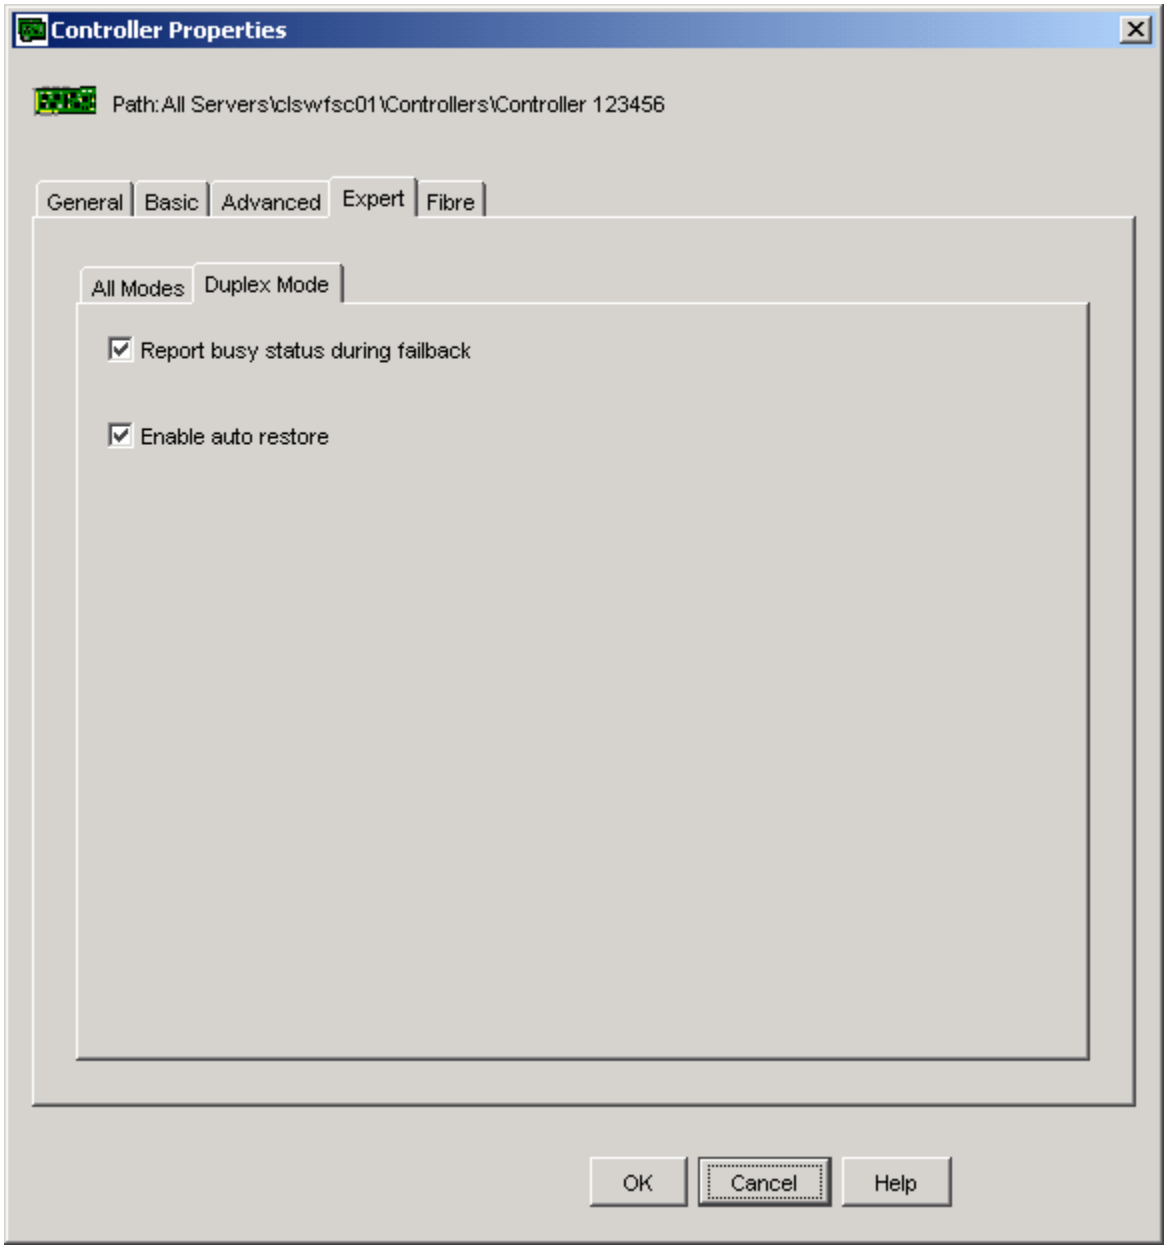

Figure 53 Controller Properties Dialog Box: Expert: Duplex Mode Dialog Tab

In the Duplex Mode Tab, you have the following options:

Report busy status during failback:

Note: A controller reset is required to make a change to this option take affect.

When enabled, causes all pending commands to be quietly dropped during failback. This is intended to help hosts that are confused by Queue Full.

Enable auto restore:

Note: A controller reset is required to make a change to this option take affect.

The default is disabled. When enabled in a dual-active controller system, Automatic Restore allows automatic recovery of a partner controller when a replacement is inserted.

7. Select **Fibre**.

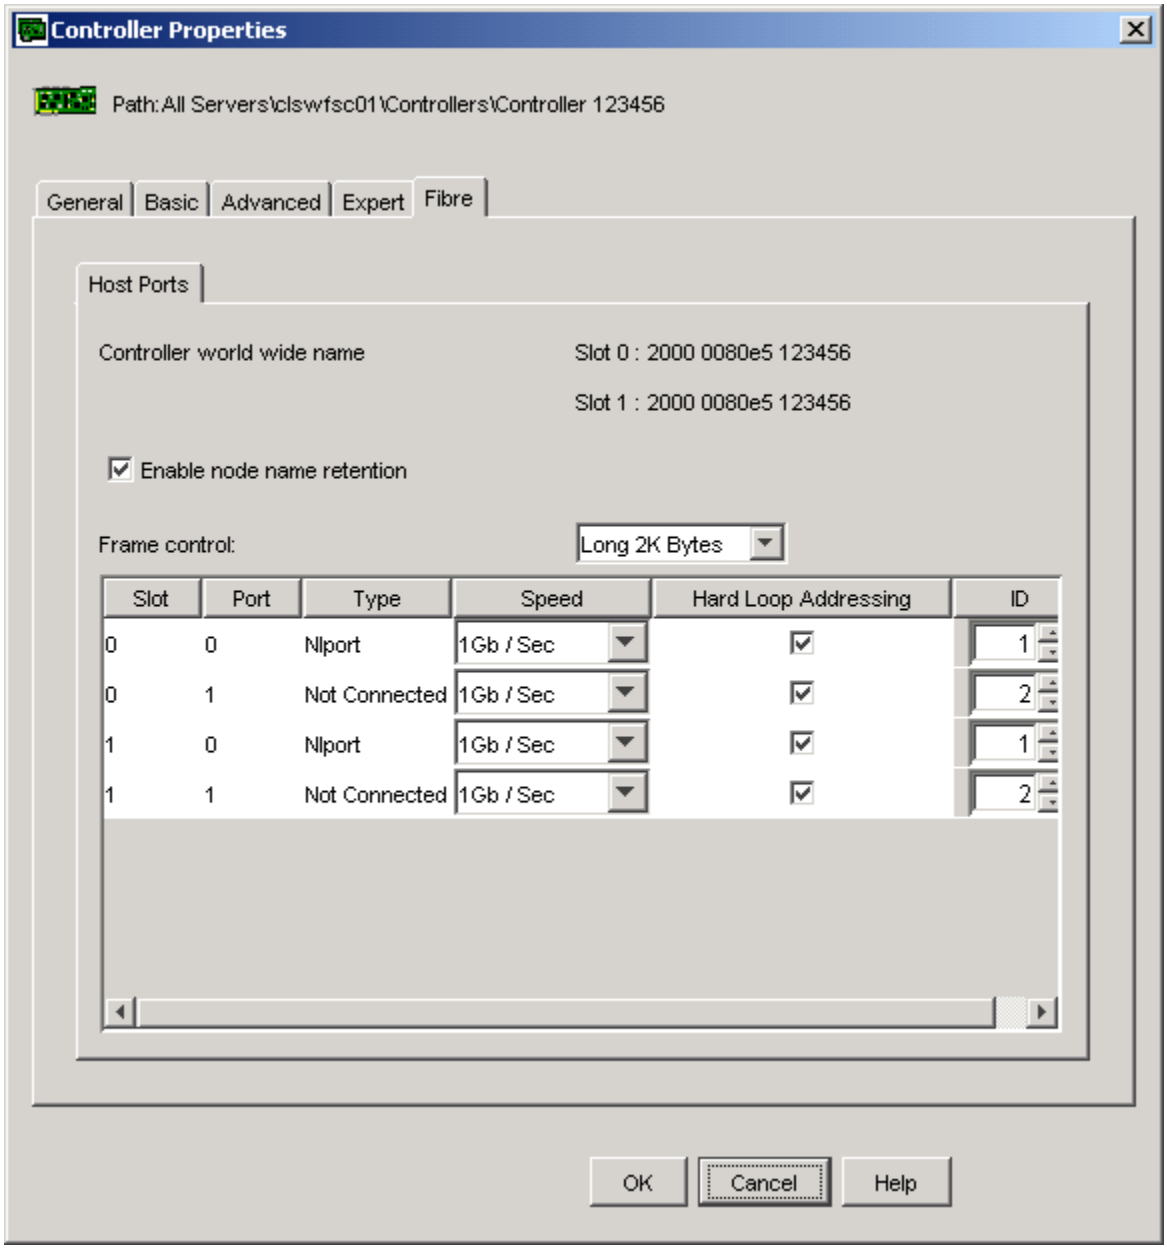

Figure 54 Controller Properties Dialog Box: Fibre: Host Ports Tab

In the Host Ports Tab, you have the following options:

Controller worldwide name: Displays the host identity of a Host Bus Adapter of a server on the fibre channel loop by slot number.

Node name retention:

Note: A controller reset is required to make a change to this option take affect.

When enabled, each controller shares its node name with its partner controller, and those names are used through all phases of failover and failback.

Frame control: Allows adjustment of the Fibre Channel chip's frame size.

Note: A controller reset is required to make a change to this option take affect.

Host port speed: Sets a 1 or 2 Gigabits per second for Auto Negotiate speed for the selected host ports.

Hard loop addressing:

Note: A controller reset is required to make a change to this option take affect.

Enables hard loop IDs on a controller.

ID field: is displayed only if Hard Loop Addressing is enabled. It allows you to select a value 0-125.

8. Select IP settings

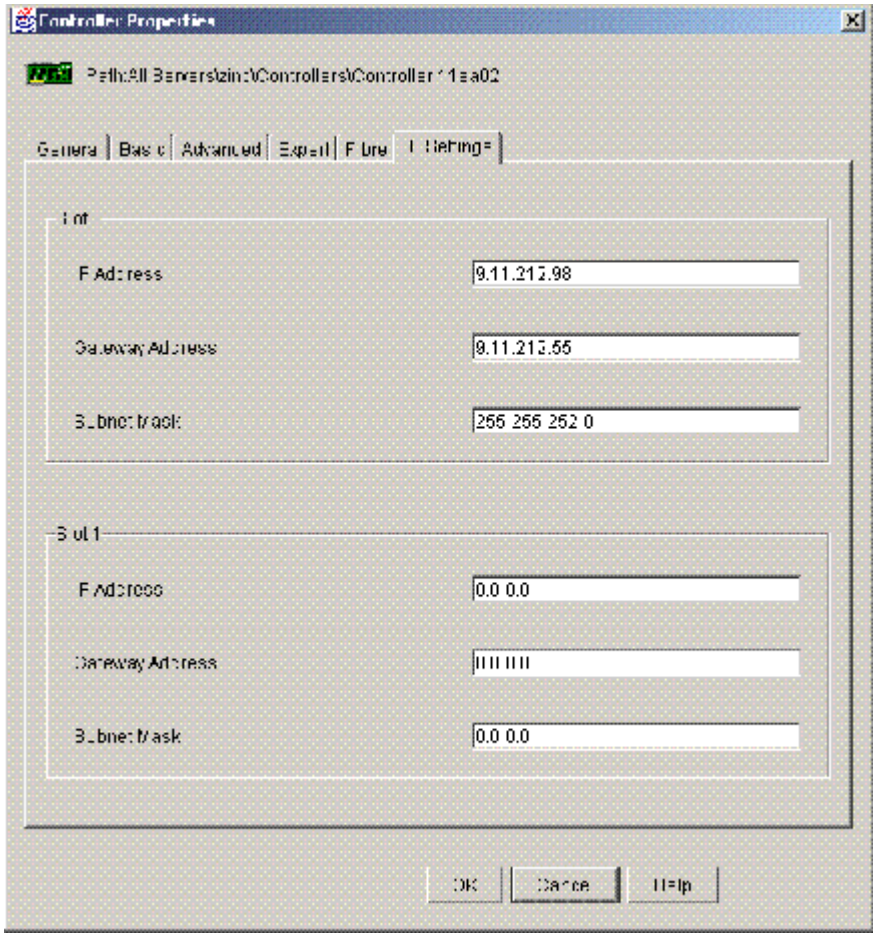

Figure 55 Controller Properties Dialog Box: IP Settings

Note: The IP Setting Tab is not shown if IP is not supported by the controller firmware.

All fields on this dialog support the Internet Protocol version 4. If an IP-configurable controller is not available in a slot, the slot selections are not available. Modifying these slot fields may affect the ability of the Spheras Storage Director Server to communicate with the controller.

If an IP configurable controller is discovered and an IP address has not yet been set for that controller, the status of the controller is *partially discovered*. The IP address may or may not be seen in the slot field. When a controller is partially discovered, no other controller properties dialogs are available. You must enter valid information in at least one of the slot fields and click **OK** before the other dialogs are enabled. If communication fails, you must again try and enter valid information in one of the slot fields.

For dual-active controllers, when one controller has been discovered, and the other only partially discovered, the state of the discovered controller is displayed in the Spheras Storage Director. You set the IP address of the partially discovered controller here. If you do not set an address and a failover occurs, the status of the dual-active pair reverts to partially discovered.

If you received an error message stating that the IP address was incorrectly entered, do the following:

Check that you have the correct IP address for the server you want to connect to, and re-enter the IP address.

Check your physical IP connection to make sure that it is intact and not loose.

Open a DOS Window and ping the IP address to see if it is functioning properly.

Click **OK** or press the Enter key to accept all of the changes in the General, Basic, Advanced, Expert, Fibre, and IP Controller Properties or **Cancel** to discard them.

If a parameter was changed that requires a controller reset, a warning dialog box opens. A reboot is required before the changes to the controller properties take effect.

Note: The Warning dialog and the reboot only occur when controller properties are changed.

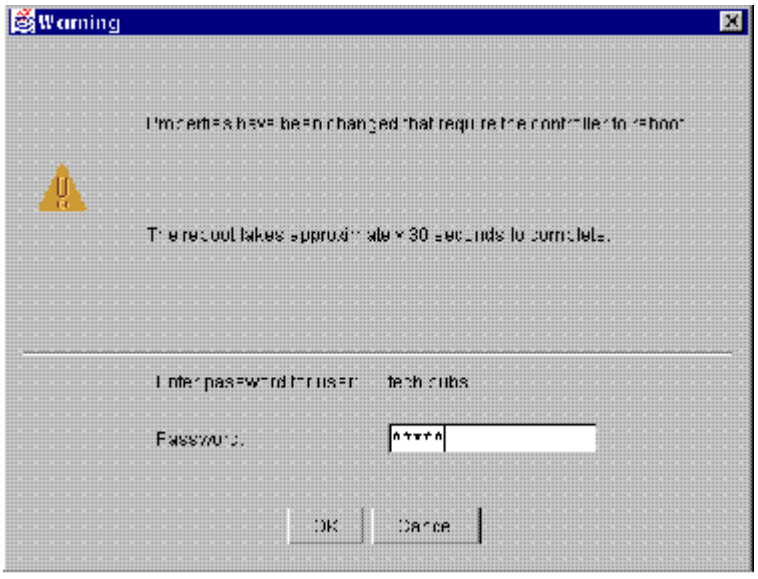

Figure 56 Controller Properties Warning Dialog Box

Type your password.

Click **OK**.

1.33 Creating Storage Configurations

The Assisted and Manual Configure Unused Storage Wizards allow you to configure RAID systems. The Configure Unused Storage with Assistance Wizard provides an option, which sets up a new array according to predefined parameters, and asks you questions to gather the key information necessary to build the array. The Manual Storage Configuration wizard provides you full control of all elements of the configuration process including SAN mapping.

It is recommended that you read both sections before proceeding to determine which configuration method is the most suitable for your storage application. Also, it is recommended that you refer to Appendix B, Detailed Information About Configurations on page 245, before beginning your configuration. The information provides the drawbacks and benefits of all RAID levels as well as Capacity, Availability, and Performance (CAP) strategies and other useful configuration information.

5.1.1 Multiple Configuration Tool Locking

Spheras Storage Director provides a Multiple Configuration Tool Lock feature that enables write access to be locked on a controller system. The configuration locking occurs when a configuration is being performed by another configuration tool. The following message is displayed *"Another configuration session is in progress. You must terminate it to continue. Terminate the other configuration session?"*

If this message is displayed, click **Yes** to cancel the other configuration session. Doing this allows your configuration to complete.

To cancel the configuration you just started, click **No**. Doing this allows the other configuration to complete.

If Spheras Storage Director is unable to terminate the other configuration session, the following message is displayed, *"Another Configuration session is in progress. Please try again after it has completed."*

5.1.2 Assisted Storage Configuration

This feature is only enabled if there are unused physical drives available. When using the 'Assisted Configuration Wizard', there may be situations where not all of the available physical drives are used when creating a configuration. For example, if you have 19 physical drives available, two arrays of nine drives are created leaving one physical drive unused. The unused physical drive may be used as a hot spare or left un-configured for later use.

Open the Configure Unused Storage With Assistance Wizard by clicking **Controller-> Configure Storage-> With Assistance...** on the menu bar or the icon as shown in Figure 57 Configure Storage With Assistance.

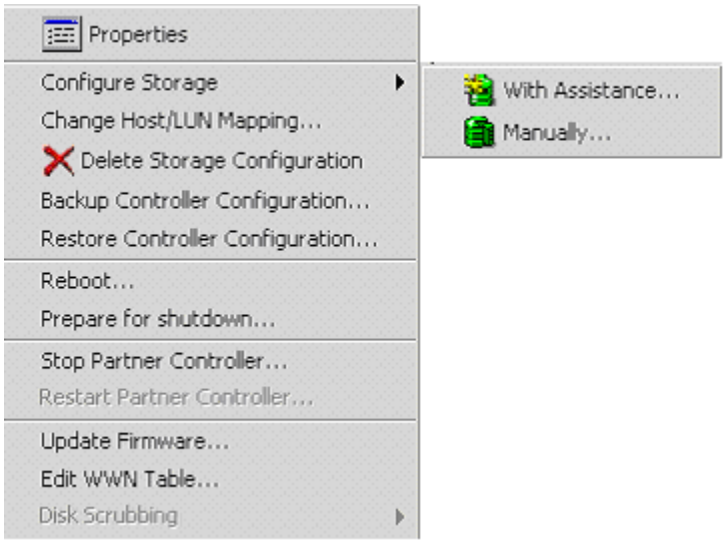

Figure 57 Configure Storage With Assistance

The Fault Tolerance dialog box opens.

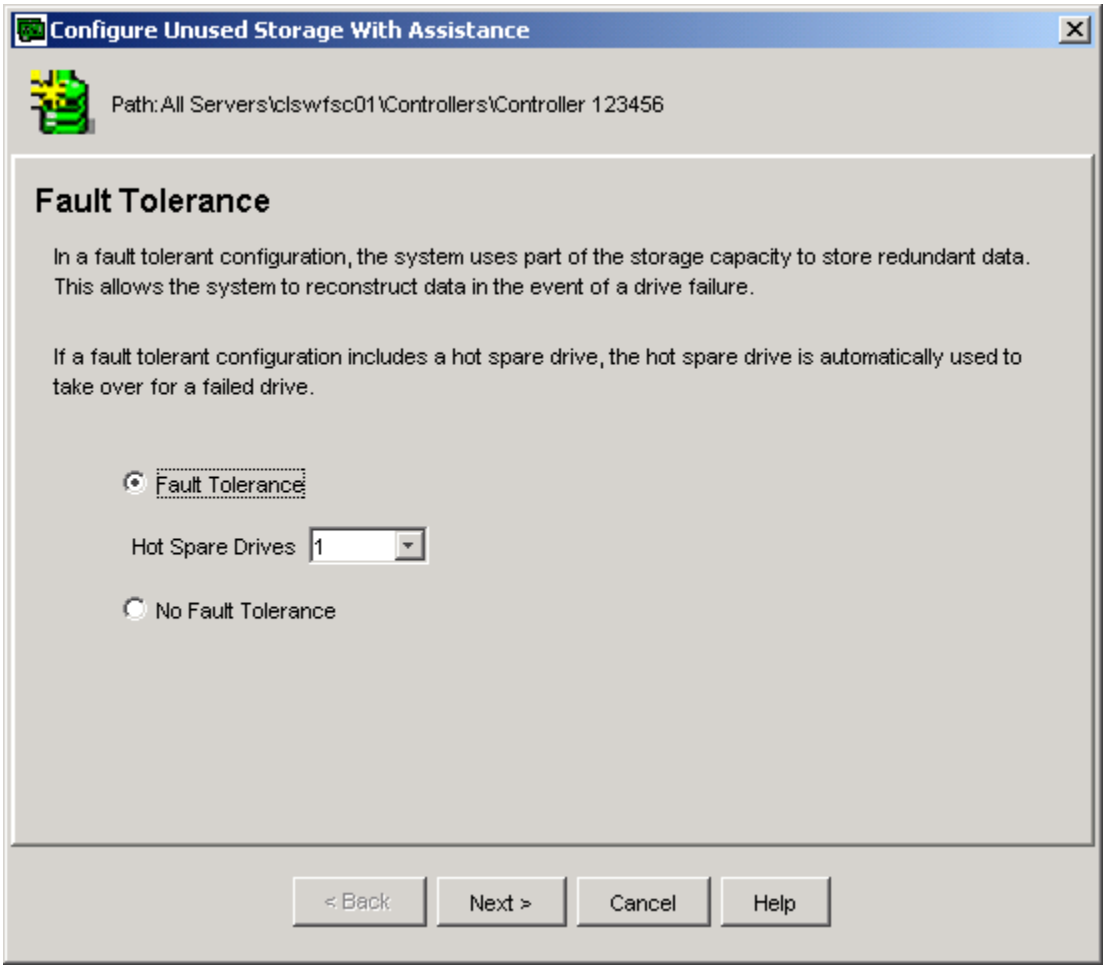

Figure 58 Configure Unused Storage with Assistance: Fault Tolerance Dialog Box

Configure the following options for fault tolerance:

Select **Fault Tolerance**. This option is disabled if there is only one available unused drive in the system.

Select the number of hot spares from the drop-down list. This option is only available with Fault Tolerance. The hot spare option is disabled if there are only two available unused drives in the system.

Click **Next** or press the Enter key.

If you selected Fault Tolerance from the Fault Tolerance dialog box, the following RAID Level dialog box opens.

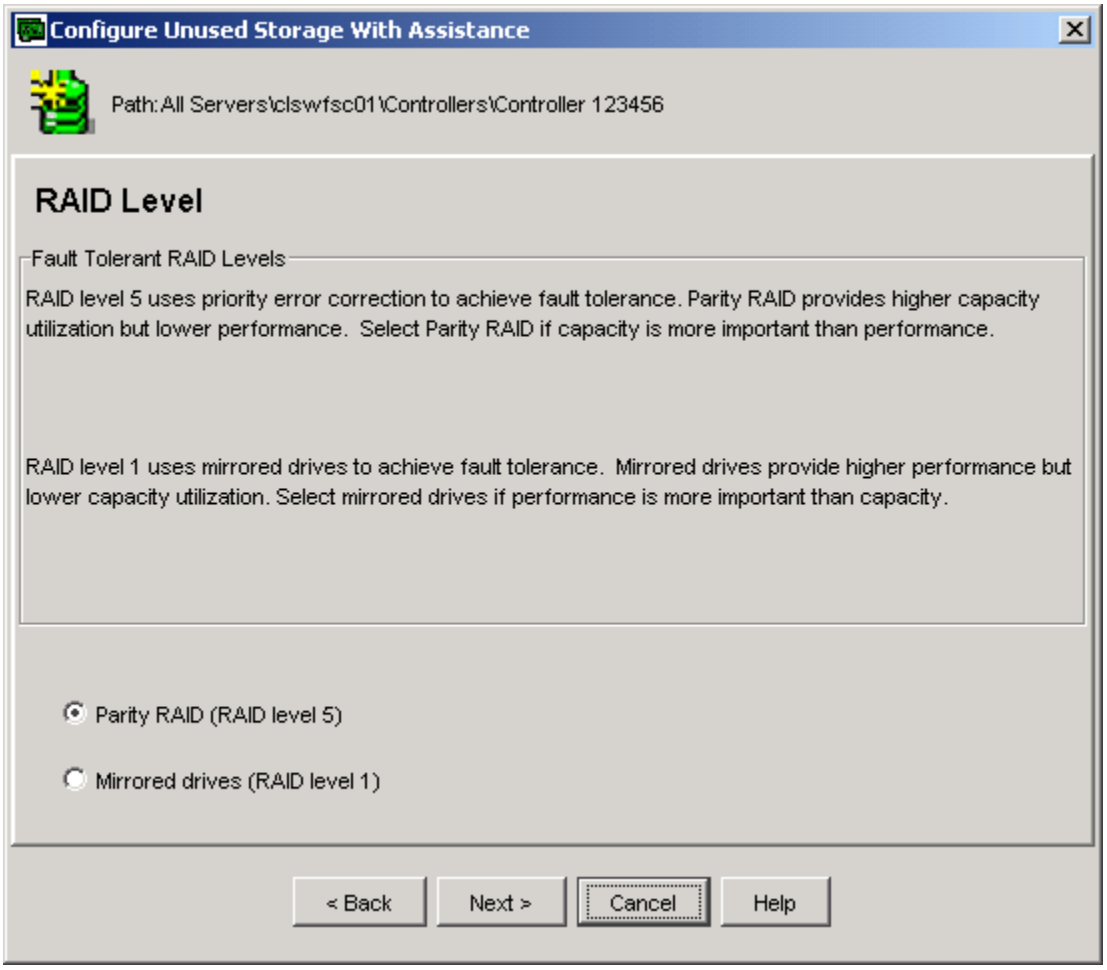

Figure 59 Configure Unused Storage with Assistance: RAID Level Dialog Box

Choose one of the following RAID Level options:

Parity RAID (RAID level 5). This option is the default if there are three or more unused drives available in the system.

Mirrored drives (RAID level 1).

If you selected the No Fault Tolerance check box, the following RAID level dialog box opens.

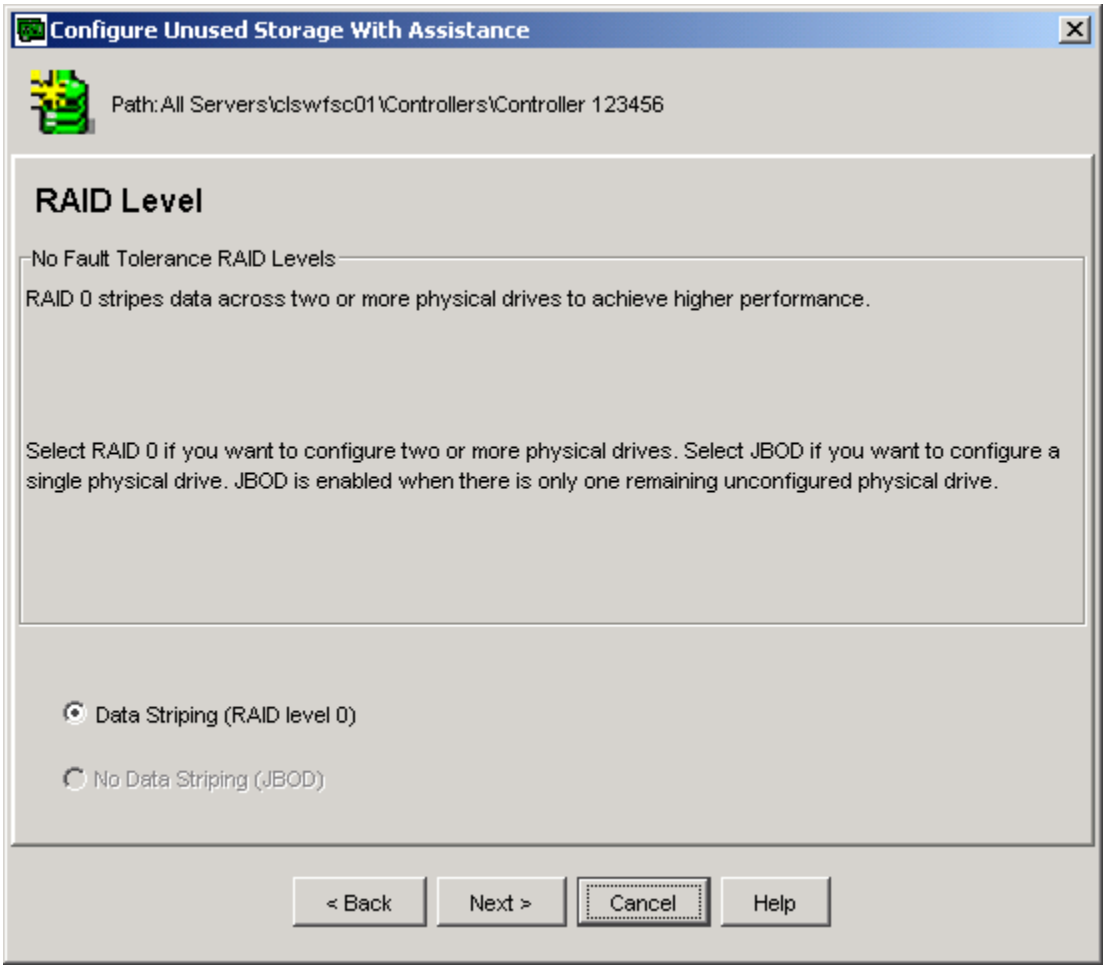

Figure 60 Configure Unused Storage with Assistance: No Fault Tolerance RAID Level Dialog Box

Choose one of the following RAID Level options:

Data Striping (RAID level 0). This option is selected by default if there is more than one unused drive available in the system.

No Data Striping (JBOD).

Click **Next** or press the Enter key. The Logical Drives dialog box opens.

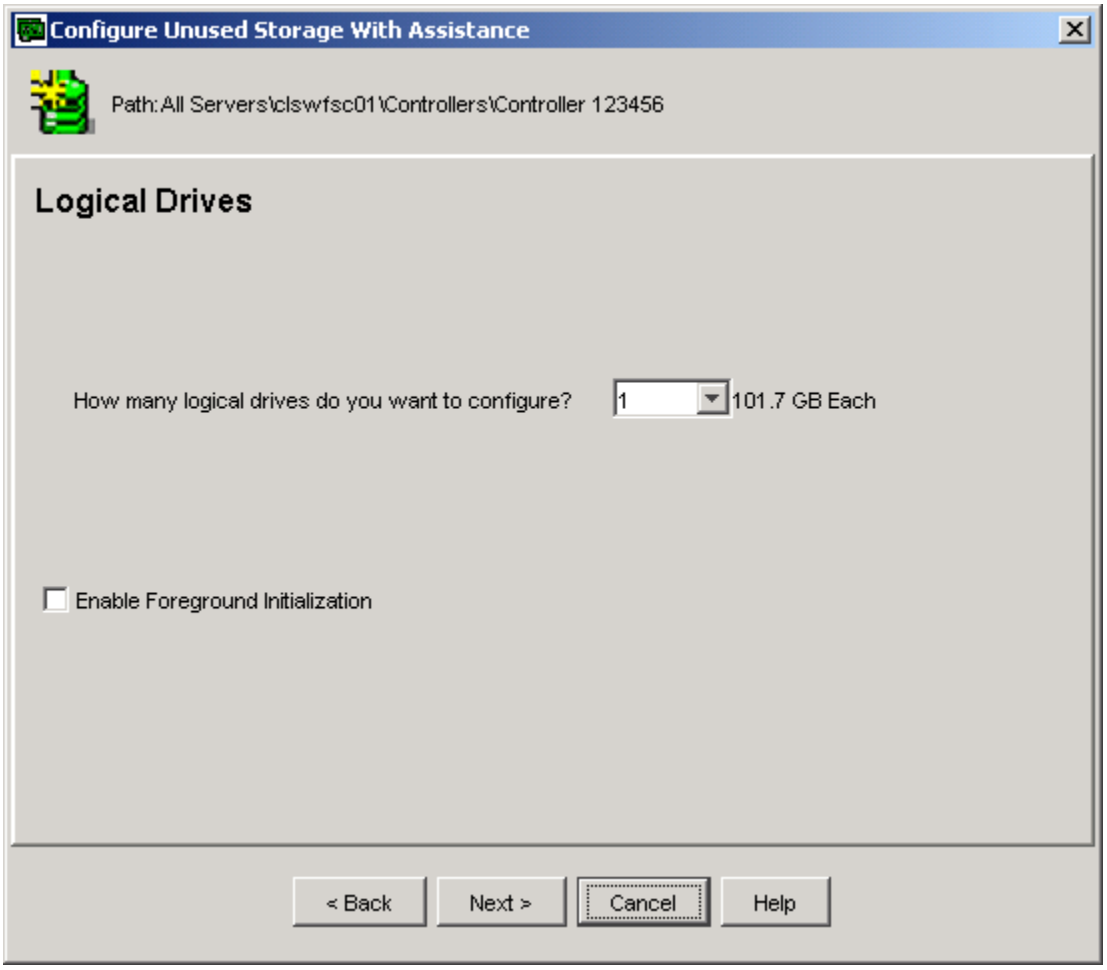

Figure 61 Configure Unused Storage with Assistance: Logical Drives Dialog Box.

Note: The dialog shown in this figure is displayed if the controller property Enable Background Initialization is enabled. If the controller property Enable Background Initialization is disabled, an additional check box is displayed which allows you to Enable Foreground Initialization.

From the drop-down list, select the number of logical drives you wish to create (1-32). The logical drive settings are set automatically as follows:

Write cache is enabled if a BBU or UPS is installed

Stripe size and cache line size are defaulted

All LUNs are assigned sequentially

All hosts are enabled.

Select **Enable Foreground Initialization** to start the initialization when you exit this Wizard. This check box is only displayed if the "Background Initialization" controller property is turned off.

With the size of system drives growing dramatically, a traditional initialization (foreground initialization) takes several hours to complete. The system drive is not accessible during a foreground initialization. Foreground initialization limits the number of simultaneous initializations to eight. Canceling a Foreground initialization cancels all foreground

initializations (see 1.48 Monitoring Long Operation Tasks on page, 134 to cancel a Foreground initialization).

Click **Next** or press the Enter key. The Summary dialog box opens.

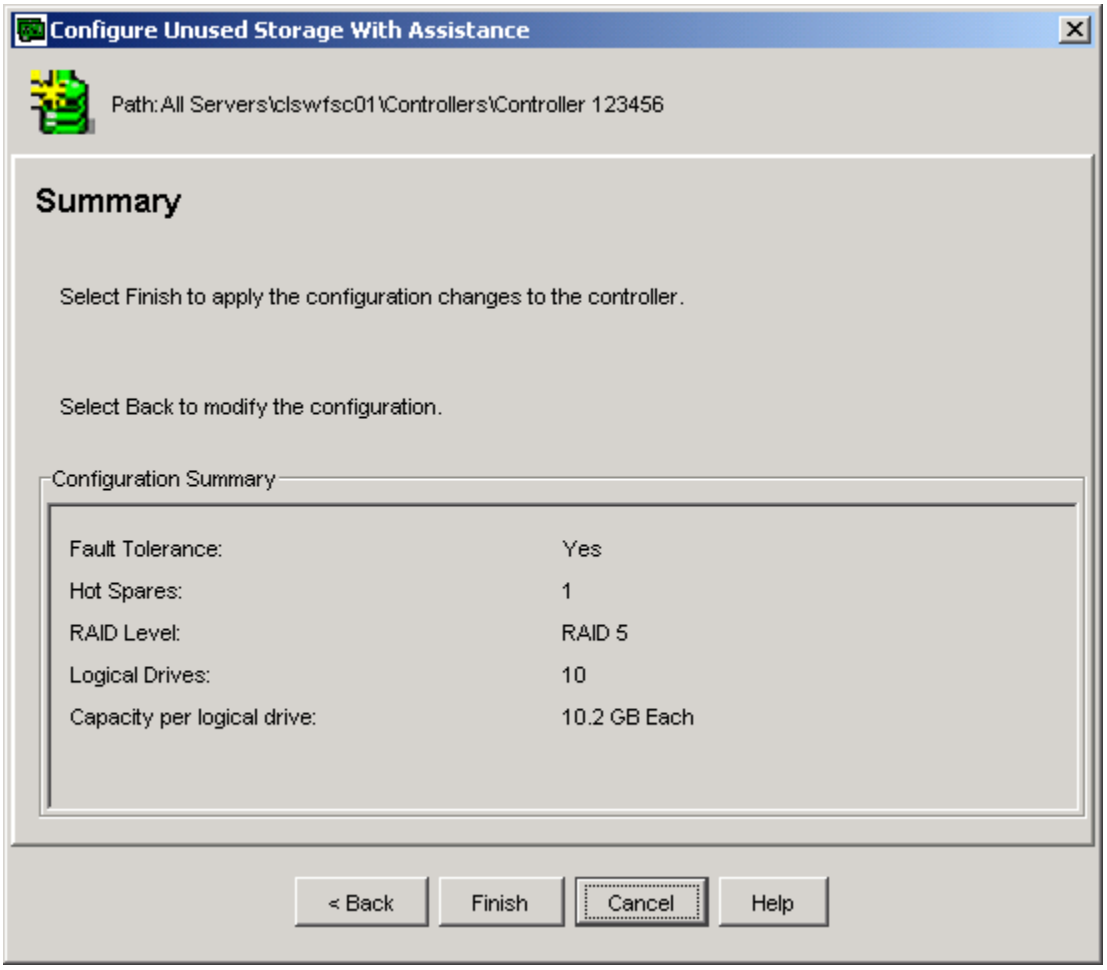

Figure 62 Configure Unused Storage with Assistance: Summary Dialog Box

Examine the configuration summary for details about the configuration you have set up. If you want to start over, click **Cancel**. If you want to make changes, click **Back** to return to the appropriate dialog box.

If you are satisfied with the configuration that you have set up, click **Finish** or press the Enter key to apply the new configuration to the controller.

5.1.3 Manual Storage Configuration

The Manual Storage Configuration wizard allows you to set up a new configuration manually deleting the previous configuration and data (if any). This feature also allows you to reconfigure an array if desired.

Open the Configure Unused Storage Manually Wizard by selecting **Controller-> Configure Storage-> Manually** on the menu bar or the icon as shown in Figure 63 Configure Storage Manually.

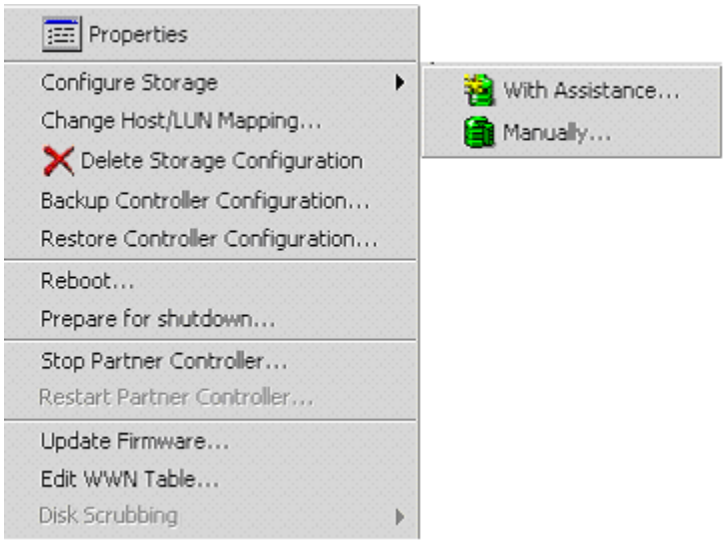

Figure 63 Configure Storage Manually

The Manual Storage Configuration dialog box opens.

Note: Only those previously-configured arrays that have unused capacity are shown in this dialog.

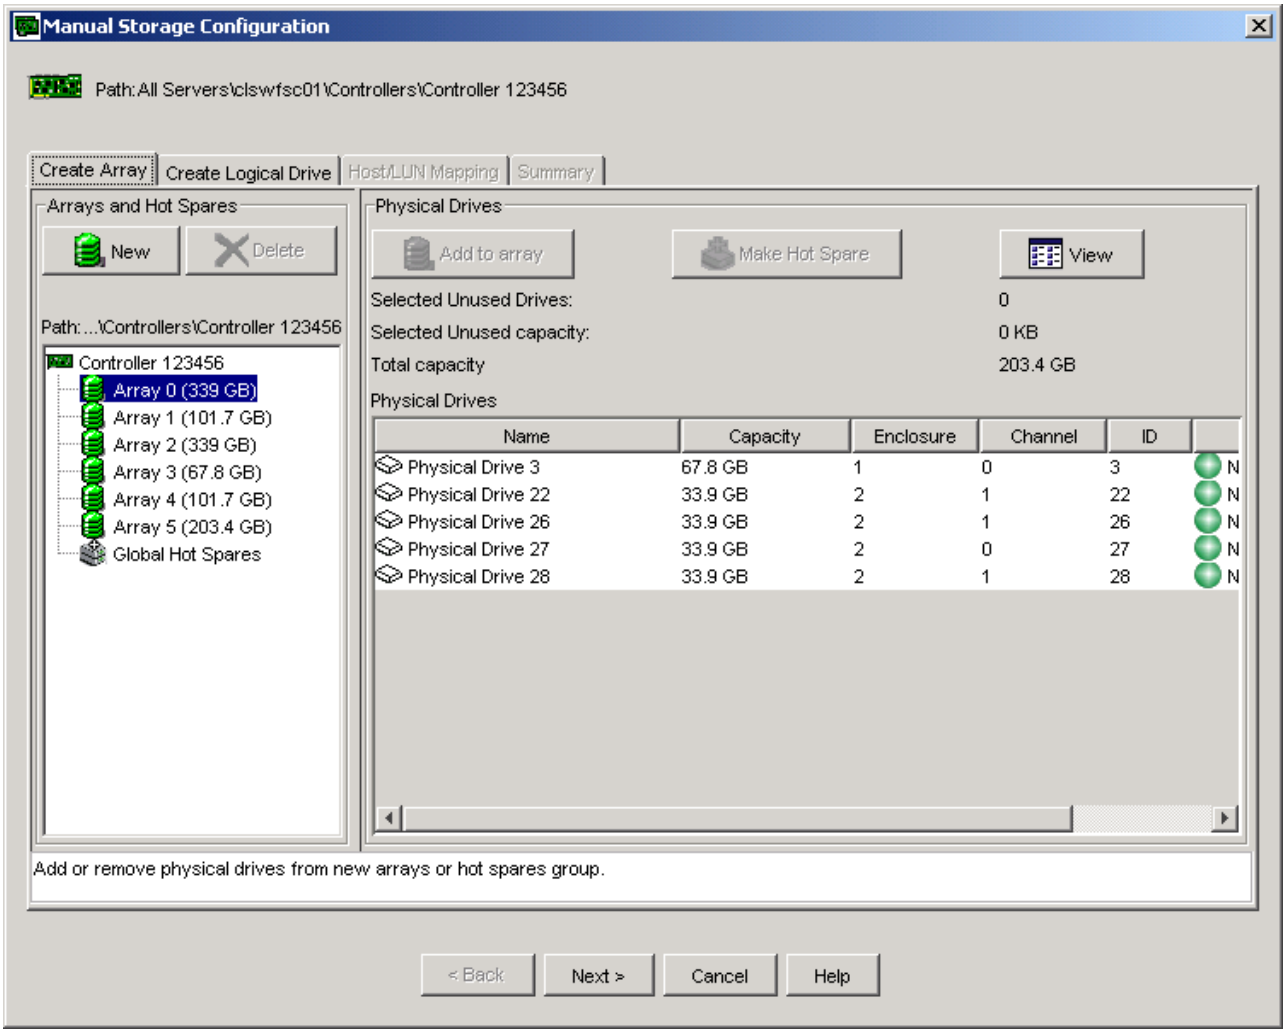

Figure 64 Manual Storage Configuration

To create a disk array, perform the following:

Spheras Storage Director automatically creates an array for you. The new array (New Array 1) is displayed in the tree node.

From the Physical Drives pane, select a physical drives to add to the array. Click **Add to Array** (or drag and drop). The physical drives are displayed in the tree node under the selected array.

Repeat steps 1-3 until you have the desired amount of disks in your array. If necessary, you can also create an additional array and add disks to the new array.

If you want to manually configure a hot spare, select an available physical drive and click **Make Hot Spare** (or drag and drop to the Hot Spare node). The hot spare drive is displayed in the tree node under the Hot Spares node.

Note: Before creating a hot spare, the Spheras Storage Director checks the size of the un-configured drive to ensure that it is large enough relative to the other configured drives on the controller. If the un-configured drive is not at least as large as the smallest configured drive that is allocated to an array on that controller, a warning message is displayed. In this situation, existing hot spares are not considered to be configured.

Click **View** to toggle between the physical view and the text view of the physical drives.

If you want to start over, click **Cancel** and reopen the dialog box.

When you have finished creating the disk arrays, click **Next** or the Create Logical Drive tab to continue your configuration.

The Create Logical Drive Tab displays.

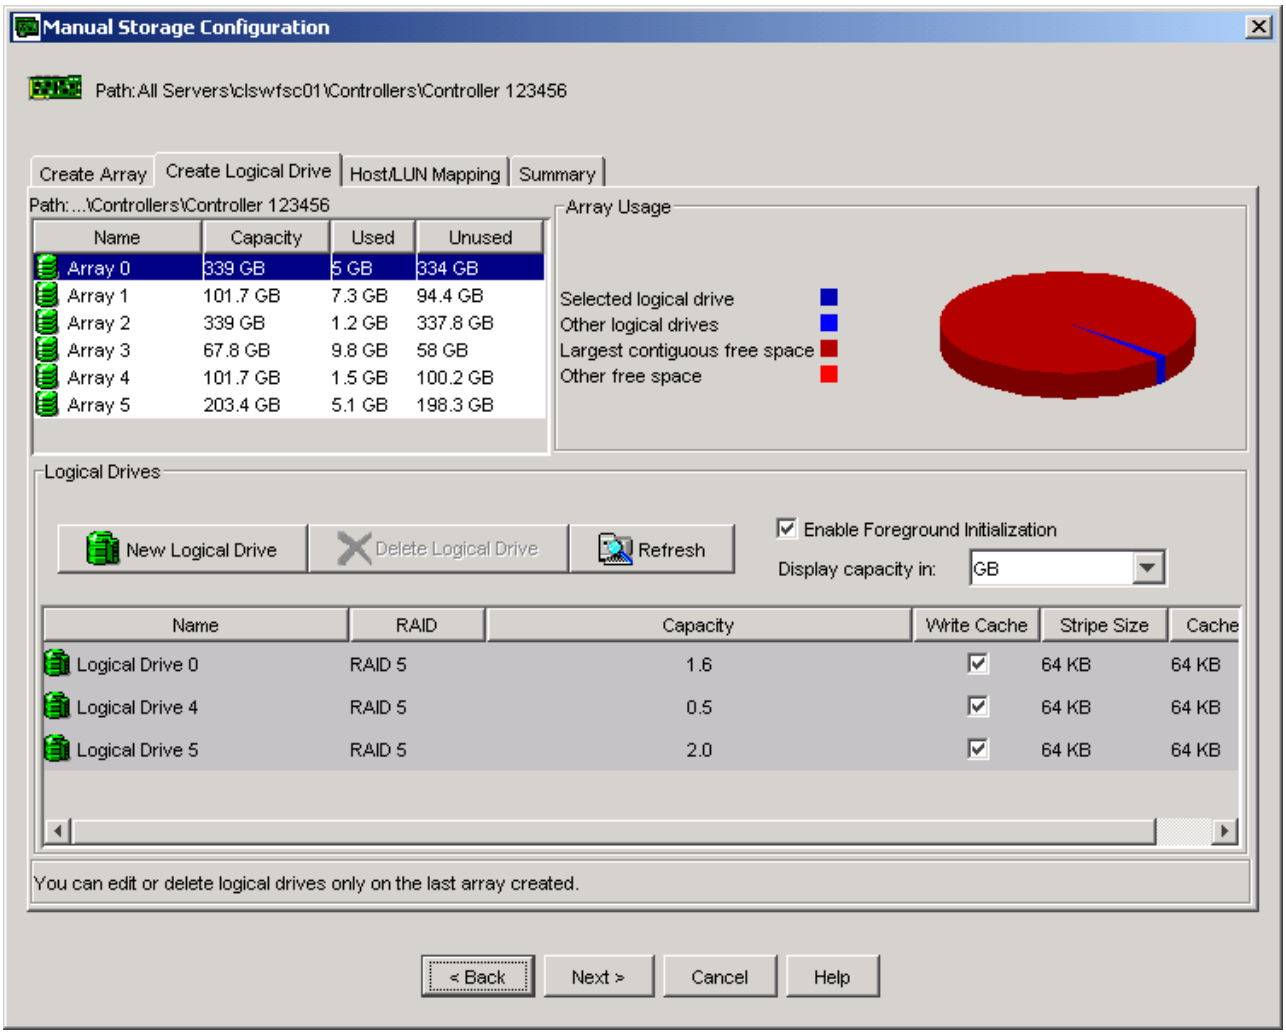

Figure 65 Manual Storage Configuration: Create Logical Drive Tab

The Create Logical Drives Tab is where you configure your disk arrays into logical drives. You must create at least one logical disk on each array. A maximum of 32 logical drives can be created per controller. Perform the following:

From the Display Capacity drop-down list (next to the Refresh button), select how you would like the capacity displayed.

Select the array that you want to configure.

Click **New Logical Drive** to create a logical drive.

From the RAID drop-down list, select a RAID level. Supported RAID levels are JBOD, RAID 0, RAID 1, RAID 3, RAID 5, RAID 0+1, RAID10, RAID 30, RAID 50. Only RAID levels compatible with the current array are available in the list.

Type the amount of available capacity for this logical drive. If your total configuration has only one logical drive that uses all available capacity, the default sizes should not be changed. If you intend to create additional logical drives, enter a smaller number to reserve the desired capacity. The Array Usage pie chart indicates how the capacity is being used.

Select **Write Cache** if you want this logical drive to use Write Back Caching. Write Back Caching improves performance, but puts data at risk. You should have a battery backup unit or uninterruptible power supply if selecting this option.

Select a stripe size. If this is a new configuration, this is an editable drop-down list. A change to one logical drive's stripe size is reflected by all other logical drives. Smaller stripe sizes provide better performance for random I/O; larger stripe sizes provide better performance for sequential transfers.

The Cache size is displayed for information only. For external controllers, the cache size and stripe size are the same. The cache size represents the size of the data in controller memory that is read or written at one time.

Select **Enable Foreground Initialization** to start the initialization when you exit this Wizard. This check box is displayed only if the "Background Initialization" controller property is turned off.

With the size of system drives growing dramatically, a traditional initialization (foreground initialization) takes several hours to complete. The system drive is not accessible during a foreground initialization. Foreground initialization limits the number of simultaneous initializations to eight. Canceling a Foreground initialization cancels all foreground initializations (see Monitoring Long Operation Tasks on page 134 to cancel a Foreground initialization).

When you are done configuring all of your logical drives, click **Refresh** to see the configuration changes reflected.

Click **Next** or the Host/LUN Mapping tab to continue your configuration. Or Click **Delete** to delete the logical drive and start over.

The Host/LUN Mapping Tab displays.

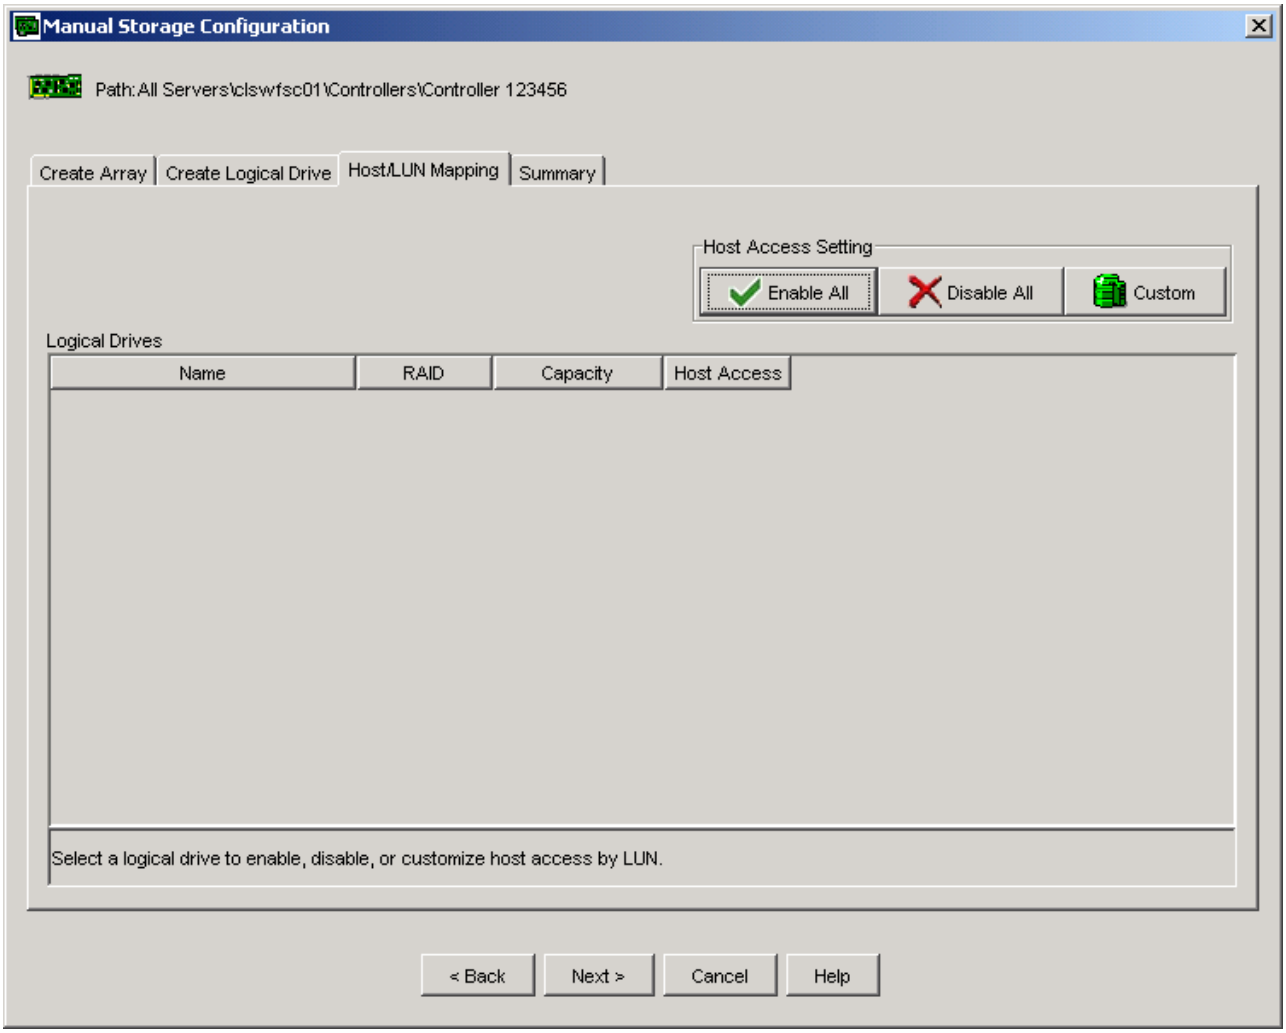

Figure 66 Manual Storage Configuration: Host/LUN Mapping Tab

The Host/LUN Mapping Tab is where you sequentially map LUN numbers to logical drives. The controller uses a World Wide Name (WWN) to uniquely identify host computers that have logged into the controller. A list of valid hosts and their corresponding WWNs is provided to Spheras Storage Director in order to configure the host to LUN mapping. Once logical drives are configured, the controller maintains a table of WWNs for each logical drive. This table defines the hosts that are granted access to each logical drive and the controller port and the LUN number.

Note: Logical drives may share LUNs as long as the same host does not access both logical drives.

You can set the Host Access by selecting one of the following buttons:

Enable All when a logical drive is selected gives all hosts access to the logical drive on the specified LUN.

Disable All when a logical drive is selected denies all host access to the logical drive.

Custom when a logical drive is selected, allows you to customize the host access settings by opening the Custom LUN Mapping Dialog box. Host access is determined independently for each controller port Enabled.

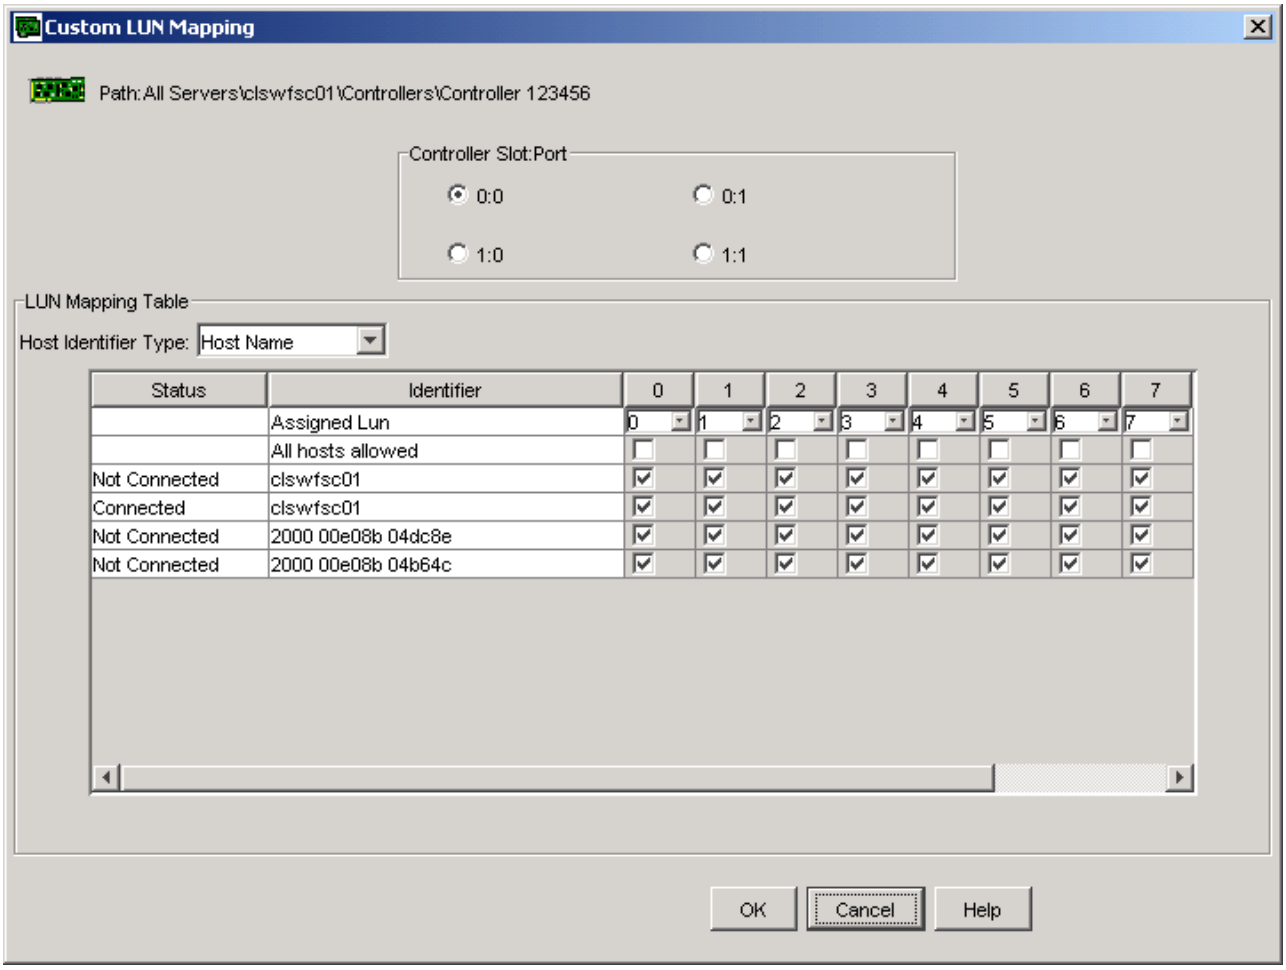

Figure 67 Host/LUN Mapping Tab: Custom LUN Mapping Dialog Box

The custom LUN Mapping dialog box allows you to manage host access to a logical drive/LUN number. Use the assigned LUN drop-down box to select the LUN number for a particular logical drive column. Then use the check boxes at the intersection of each logical drive column Figure 67 Host/LUN Mapping Tab: Custom LUN Mapping Dialog Box and host row to allow or deny host access. Select the All hosts allowed checkbox to give all hosts access to that logical drive/LUN. The individual host check boxes are disabled when the All hosts allowed checkbox is selected.

Select the following limits:

Controller Slot Port: Select the controller slot and port number to configure. Valid values depend on the mode of the controller and the topology setting.

Host Identifier Type: Select the way you would like the information displayed; select host name, or WWN from the drop down list.

Host Status: Displays the capability of a host to communicate with the controller.

Connected: Host is connected to the controller (logged on to the loop or on the same bus).

Not Connected: Host is not connected to the controller (Not logged on to the loop or on the same bus).

Identifier: Displays the LUN mapping identifiers that you can customize.

LUN Selection: Select available LUNs in the range of 0-255. LUNs already assigned to other logical drives mapped to the selected host are not displayed.

When you are done configuring all of your logical drives, click **Next** or the Summary tab to continue your configuration.

The **Summary** Tab displays.

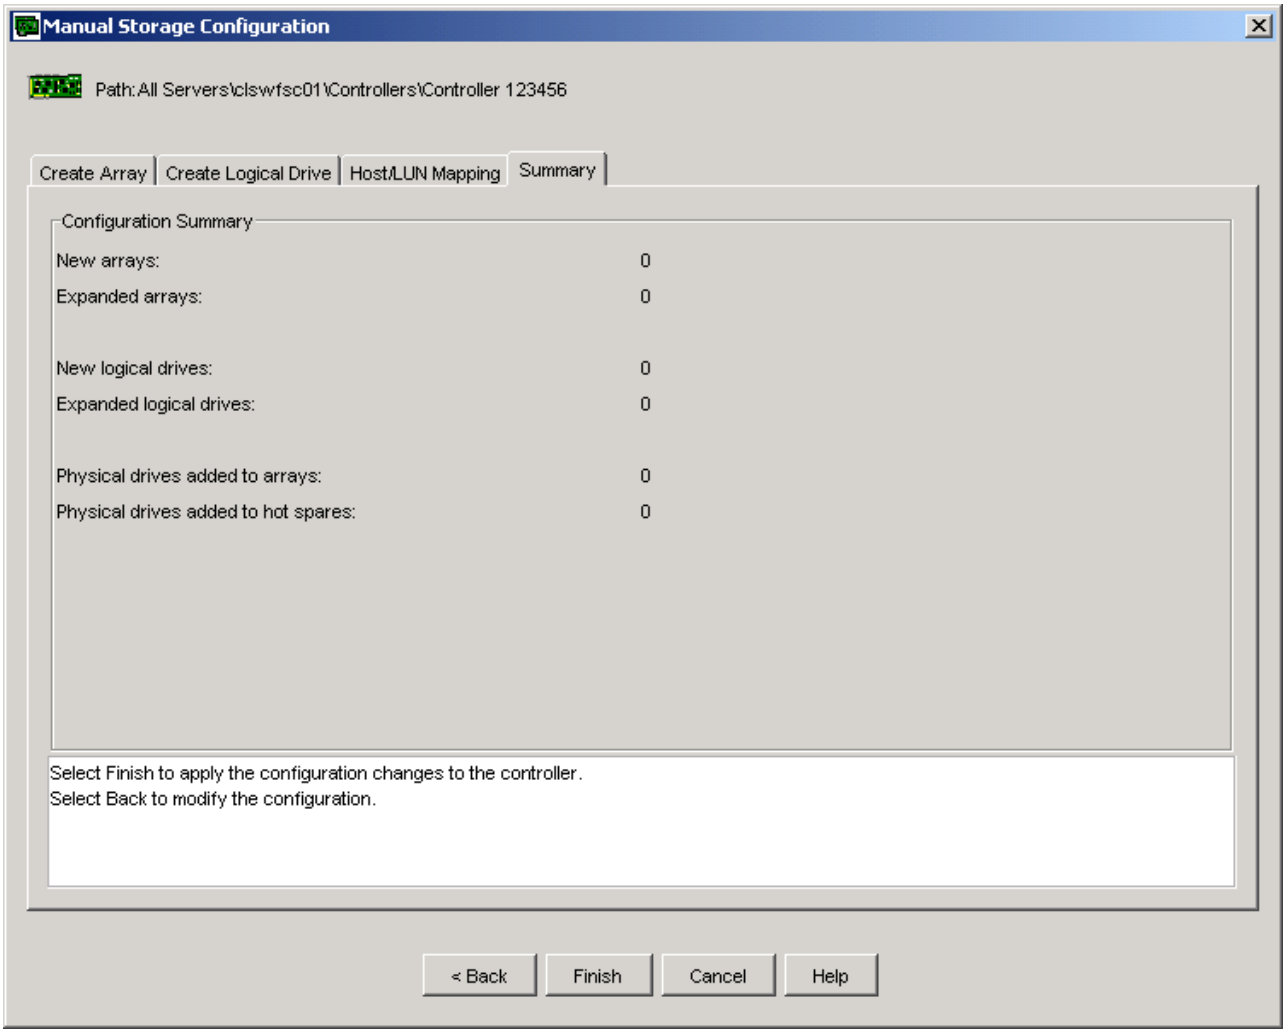

Figure 68 Manual Storage Configuration: Summary Tab

The Summary tab displays details about the configuration that will be set up according to your configuration choices:

If you want to change some of your choices, click **Back**

If you want to start over, click **Cancel**

If you accept the configuration as presented in the Summary tab, click **Finish** or press the Enter key to apply the new configuration to the controller.

1.34 Spanning in Spheras Storage Director

Spanning provides the ability to configure multiple disk arrays or parts of multiple disk arrays as one logical drive. Spanning is performed automatically in Spheras Storage Director whenever an array contains more than 16 physical drives.

1.35 Expanding a Logical Drive

Online capacity expansion is the ability to increase the size of a logical system drive while normal reads and writes occur. The expansion can be obtained when you change the RAID level, or increase the amount of space on the physical drives allocated to the logical drive, or both. When performing this operation, you cannot change the write cache, initialize drive, or stripe size of an existing logical drive. The user must stop the partner controller before starting a logical drive expansion process. When this occurs, a consistency check begins if the logical drive is in write-through cache mode. You must wait until the consistency check completes before expanding your logical drive. When the logical drive expansion is complete, you can restart the partner controller.

The following restrictions are enforced by the controller in the processing of this function:

- The maximum number of disk drives in the source RAID set is 16.
- The new disk drives must be added to the system while power is supplied.
- The disk drives being added must not already be part of an array.
- The system drive to be expanded must be online and in Optimal state.
- The capacity of each of the added disk drives must be greater than or equal to that of the smallest disk drive in the RAID set.

Using the enlarge system drive option, the system drive specified is enlarged as well as having its data striped across a larger number of physical disk drives. For this option, the system drive specified must be the only system drive on the pack of disk drives being migrated.

Add capacity input parameters and execution parameters are kept in NVRAM and COD. The process is restarted automatically by the firmware after a power failure.

In the event that this process must be terminated (for example, two disk drives fail and cannot be recovered), the controller automatically aborts the migration process.

Expanding a Logical Drive, Background initialization, Initialize system drive, Automatic rebuild, and Consistency check are mutually exclusive operations. Only one process can run at a time.

No configuration update commands issued from the host are allowed during this process.

Write-back is disabled during the process but resumes at the end of the operation.

Removing and replacing a controller, during migration can result in data loss. A hardware error, which renders the controller unusable or the NVRAM unreadable, also results in data loss during migration.

CAUTION: It is highly recommended that you backup all of your data and your current configuration before making modifications. See 1.40 Backing Up a Controller Configuration on page 111.

Follow these steps to perform a Logical Drive Expansion:

To expand a logical drive, the partner controller must be stopped. Click **Controller -> Stop Partner Controller.**

Note: It may take a few minutes to stop the partner controller.

In the navigation area, select the logical drive that you need to expand.

Click **Logical->Expand** from the menu bar.

The Manual Storage Configuration: Expand Logical Drive window opens.

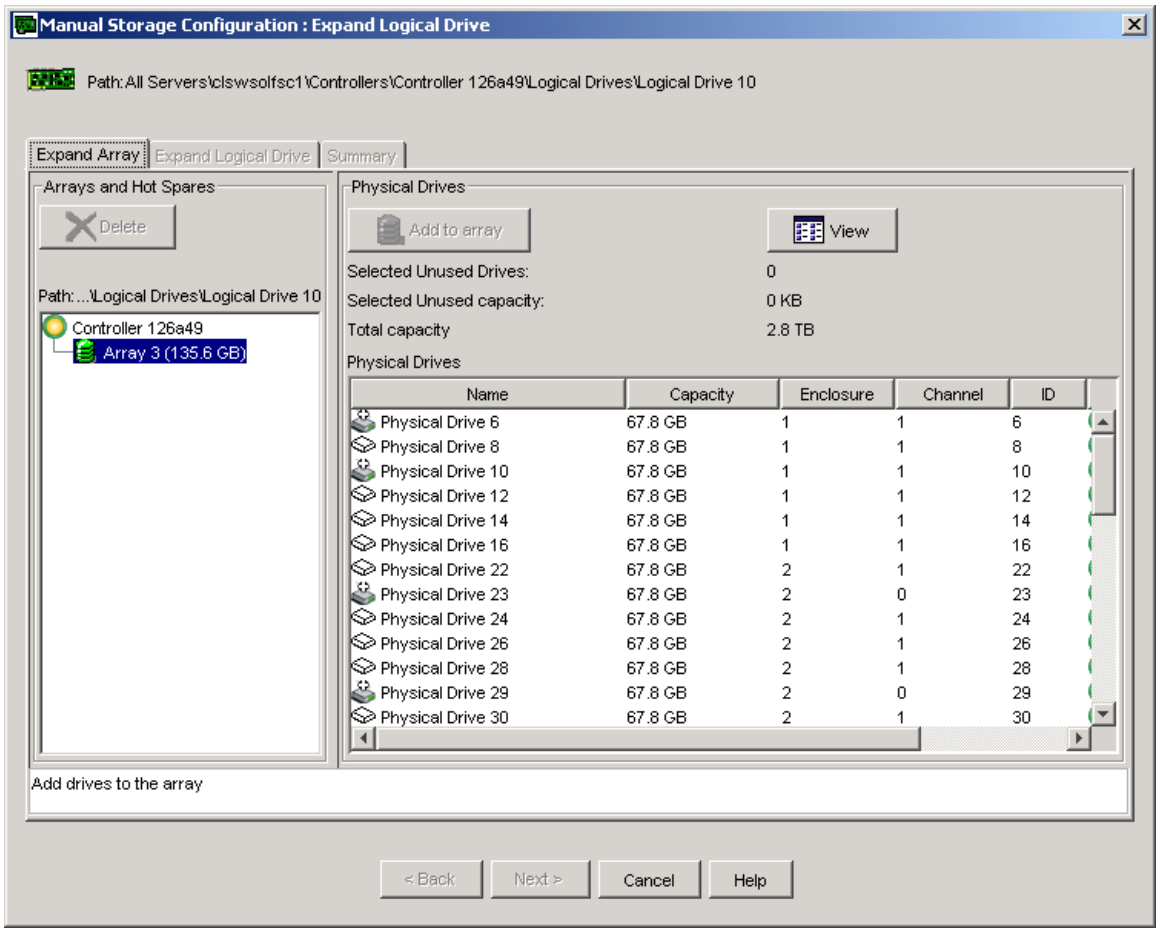

Figure 69 Manual Storage Configuration: Expand Logical Drive

Highlight the Physical Drive that you want to use in the expansion. Click the **Add to Array** button followed by the **Next** button.

Under the **Expand Logical Drive** tab, click **Next** followed by **Finish** to expand the capacity of your existing logical drive. If you want to add an additional logical drive to the existing configuration, click on the **New Logical Drive** button, select the RAID type of the new logical drive, followed by **Next** and **Finish**.

Notes:

During the Expansion process, a new array is created containing the newly added physical drives. This array is only present during the expansion process, and is removed upon completion of the expansion process.

When expanding a logical drive, you cannot create a new array or create another logical drive.

1.36 Deleting a Logical Drive

This feature allows you to delete one or more logical drives for rearranging your storage space.

It is recommended that you back up all of your user data on your arrays.

In the navigation area, right-click on the logical drive that you want to delete or select **Logical->Delete** from the menu.

Select Delete.

The Delete a Logical Drive Warning dialog box opens.

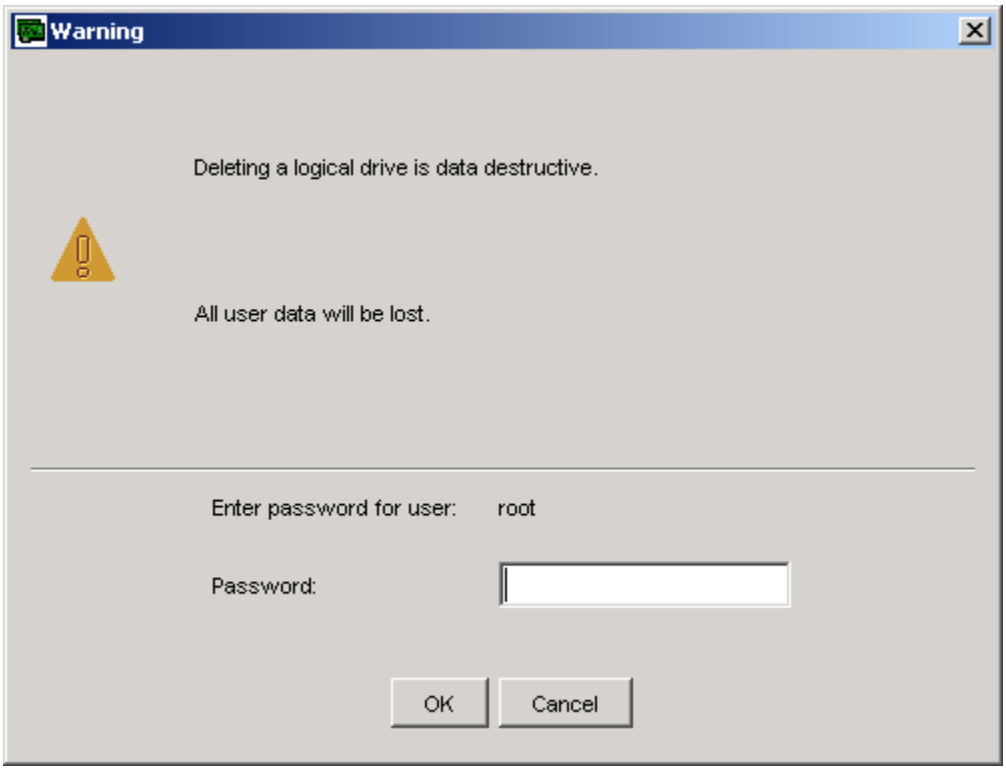

Figure 70 Deleting a Logical Drive

Type your Administrative password in the password field.

Click **OK**. The logical drive is deleted.

1.37 Expanding an Array

You can expand an existing array by adding physical drives to the array. The existing array can be expanded to a maximum of 16 drives. As part of the expand array process, the add

system drive option adds a new system drive using the increased capacity rather than increasing the size of the existing system drives. The system drive's data is striped across a larger number of physical disk drives than before the expansion operation. When the operation is complete, the new system drive partition needs to be created and formatted according to the operating system in use. During the RAID set expansion process, which includes re-striping data from the old (smaller) RAID set to the new (expanded) RAID set, the controller continues to service host I/O requests.

The user must stop the partner controller before expanding an array. When this occurs, a consistency check begins if the logical drive is in write-through cache mode. You must wait until the consistency check completes before expanding your logical drive. When the array expansion is complete, you can bring the partner controller back online.

The following restrictions are enforced by the controller in the processing of this function:

The maximum number of disk drives in the source RAID set is 16.

The new disk drives must be added to the system while power is supplied.

The disk drives being added must not already be part of an array.

The system drive to be expanded must be online and in Optimal state.

The capacity of each of the added disk drives must be greater than or equal to that of the smallest disk drive in the RAID set.

Using the enlarge system drive option, the system drive specified is enlarged as well as having its data striped across a larger number of physical disk drives. For this option, the system drive specified must be the only system drive on the pack of disk drives being migrated.

Add capacity input parameters and execution parameters are kept in NVRAM and COD. The process is restarted automatically by the firmware after a power failure.

In the event that this process must be terminated (for example, two disk drives fail and cannot be recovered), the controller automatically aborts the migration process.

Expanding a Logical Drive, Background initialization, Initialize system drive, Automatic rebuild, and Consistency check are mutually exclusive operations. Only one process can run at a time.

No configuration update commands issued from the host are allowed during this process.

Write-back is disabled during the process but resumes at the end of the operation.

Removing and replacing a controller, during migration can result in data loss. A hardware error, which renders the controller unusable or the NVRAM unreadable, also results in data loss during migration.

CAUTION: It is highly recommended that you backup all of your data and your current configuration before making modifications. See 1.40 Backing Up a Controller Configuration on page 111.

To perform an Array Expansion, follow these steps:

To expand an array, the partner controller must be stopped. Click **Controller -> Stop Partner Controller.**

Note: It may take a few minutes to stop the partner controller.

In the navigation area, select the array that you need to expand.

Click **Array->Expand** from the menu bar.

The Manual Storage Configuration: Expand Logical Drive window opens.

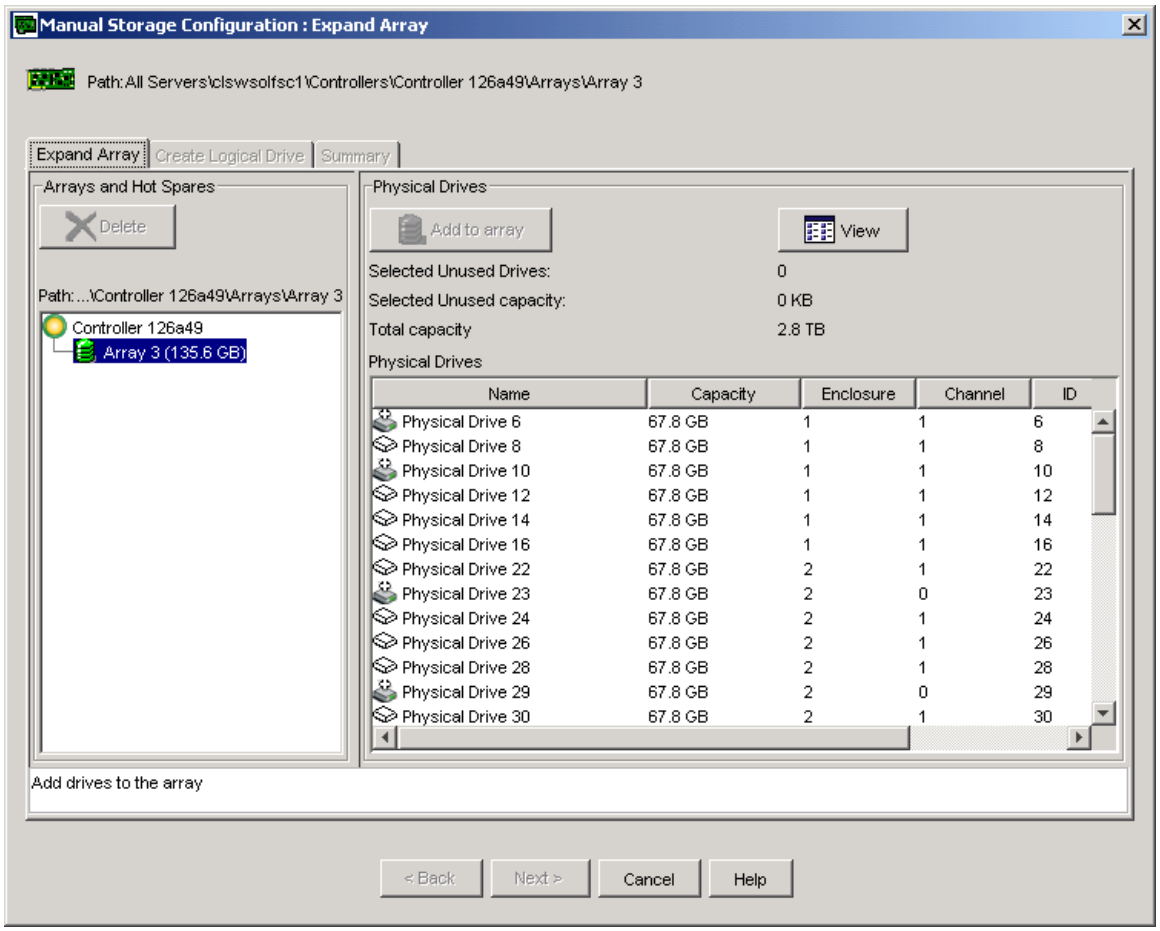

Figure 71 Manual Storage Configuration: Expand Logical Drive

Select which physical drives you want to add to the Array. Click **Next** to move to the **Create Logical Drive** tab.

If a logical drives exists on the array, you have two options:

Leave the existing logical drives unchanged.

Add one new logical drive.

Go to the Create Logical Drive tab if you wish to create a new logical drive. See Figure 65 Manual Storage Configuration: Create Logical Drive Tab for detailed configuration information.

Click **Next** to go to the **Summary** tab and click on the **Finish** button to complete the expansion of the array.

Note: When expanding an array, you cannot create a new array, but you can create a new single logical drive during the same session. You cannot expand the capacity of any existing logical drives during this session.

1.38 Deleting an Array

This feature allows you to delete one or more Arrays for the purpose of rearranging your storage space.

It is recommended that you back up all of your user data on your arrays.

You must first delete all of the logical drives in the array, see 1.36 Deleting a Logical Drive on page 105.

In the navigation area, right-click on the Array that you want to delete or select **Array->Delete** from the menu.

Select Delete.

The Delete an Array Warning dialog box opens.

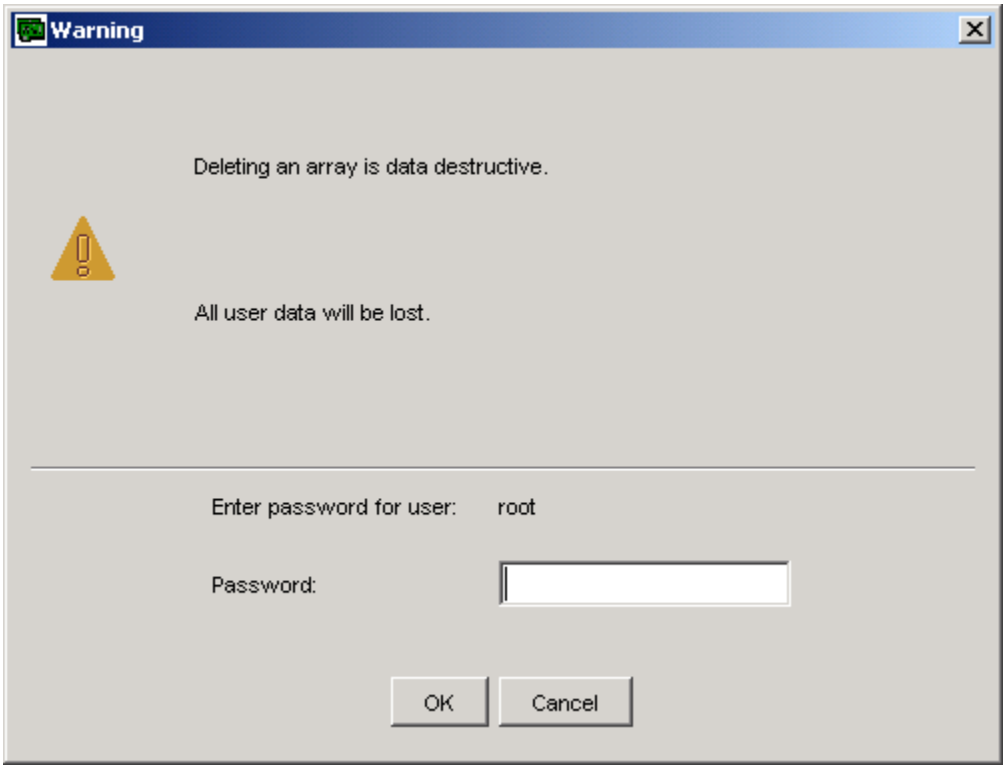

Figure 72 Delete an Array

Type your Administrative password in the password field.

Click **OK**. The array and its contents are deleted.

1.39 Restoring a Controller Configuration from a Previously Saved File

Restoring a configuration downloads a previously-saved controller configuration from a file on the server to the controller. It also provides you with both data recovery and identical system stamping that insures the logical drives and hot spares are defined in the same manner as the original system. The configuration is restored only if the physical disks are of equal or greater size and residing in the same slots and on the same channels. Do the following to restore a controller configuration:

Click Controller->Restore Controller Configuration as shown in Figure 73 Restore Controller Configuration.

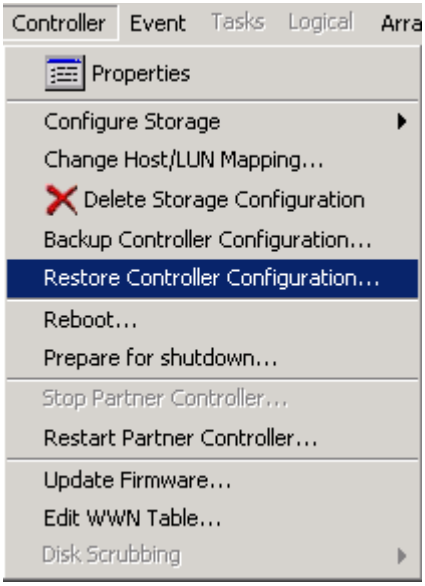

Figure 73 Restore Controller Configuration

The Restore Controller Configuration dialog box opens.

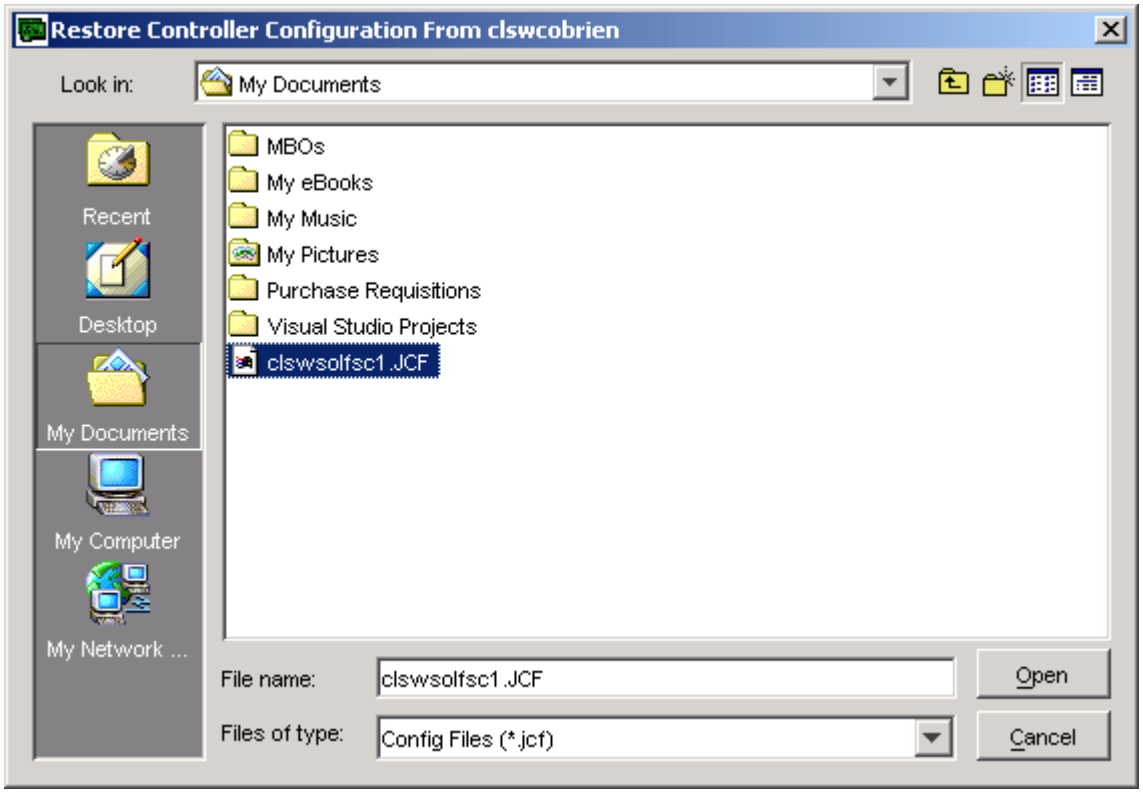

Figure 74 Restore Controller Configuration Dialog Box

Select the configuration file that you want to restore.

CAUTION: Restoring a controller configuration to an already-configured controller can cause loss of data.

Click **Open**.

The Summary tab of the Manual Configuration Wizard opens. The tab shows the configuration that you are about to restore.

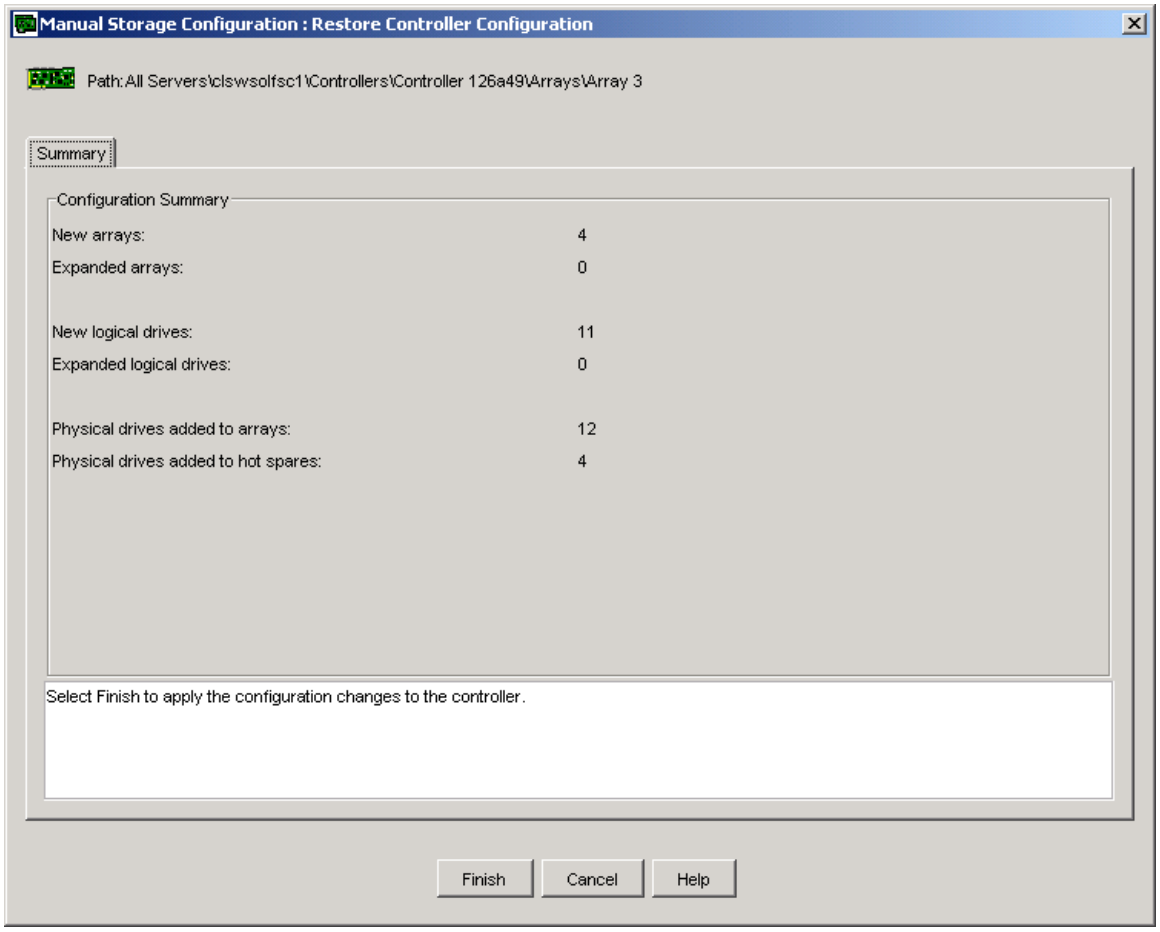

Figure 75 Restore Configuration Summary Tab

Click **Finish** to restore your configuration.

Click **Cancel** to keep the existing configuration.

1.40 Backing Up a Controller Configuration

Backing Up a configuration saves a controller configuration to the server. Do the following to back up a controller configuration:

Click Controller->Backup Controller Configuration as shown in Figure 76 Backup Controller Configuration.

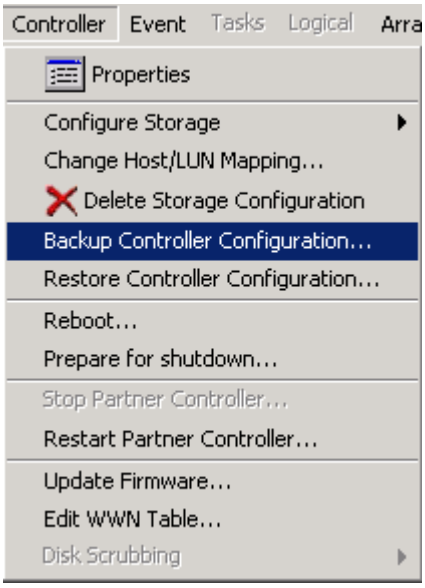

Figure 76 Backup Controller Configuration

In the Backup Configuration dialog box Figure 77 Backup Controller Configuration Dialog Box type a name for the configuration file you want to save to the controller.

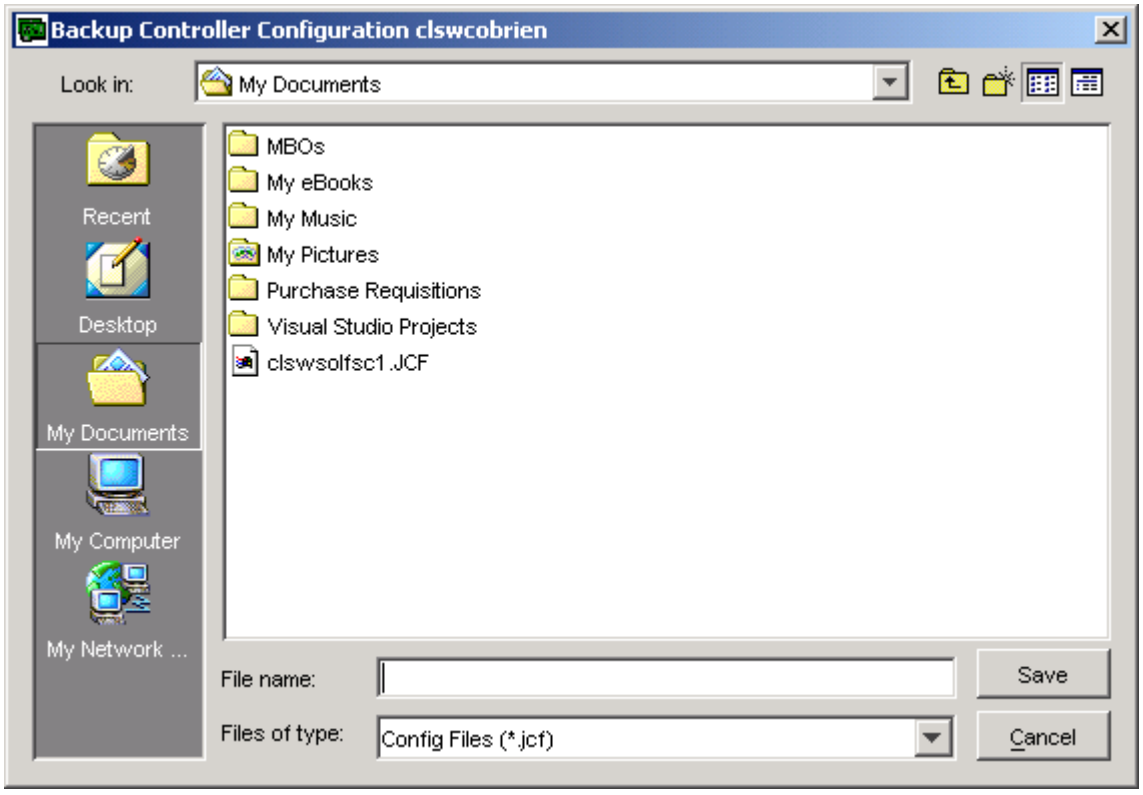

Figure 77 Backup Controller Configuration Dialog Box

Click **Save** to save the configuration file.

Chapter 6. Monitoring Activities

This section describes the following monitoring activities that Spheras Storage Director performs:

- 1.41 Monitoring All Servers
- 1.42 Monitoring Events (errors)
- 1.43 Monitoring a Controller on page 117 (monitoring controller activity, reviewing controller configuration)
- 1.45 Physical Disk and Logical Drive Monitoring on page 121 (reviewing physical device and logical drive information, and locating arrays)
- 1.46 Monitoring an Enclosure on page 128 (monitoring and managing enclosure information)
- 1.47 Monitoring the Battery Backup Unit on page 132
- 1.48 Monitoring Long Operation Tasks on page 134 (monitoring the status of ongoing processes, such as initialization, rebuild, consistency check, and expanding the capacity of an array)

1.41 Monitoring All Servers

To see a summary of the status of all servers, click on **All Servers** in the navigation area. The following information is displayed in the content area:

- The name of each server in this server group.
- The IP address of each server in this server group.
- The operating system on each server in this server group.
- The status of each server in this server group (optimal, serious, warning, or critical).
- The description of each server in this server group.

Note: The status that is displayed when **All Servers** is selected is the most severe status of the server itself and all of its controllers. Furthermore, the controllers status is the most severe status of the controller itself and all of its components. Therefore, examples of the status that is displayed include: Server Not Communicating, or Duplex Failed, or Physical Drive Failed.

1.42 Monitoring Events

The Spheras Storage Director server monitors the activity of all devices and controllers attached to the server. When a controller firmware generated event occurs, the Spheras Storage Director Server retrieves those events and forwards them to Spheras Storage Director client.

6.1.1 Viewing Events

Events are displayed in the content area of Spheras Storage Director. You can select the events node in the navigation area to see all events for all controllers on a server. If you want to view events on a specific controller, select the events node in the navigation area beneath the desired controller.

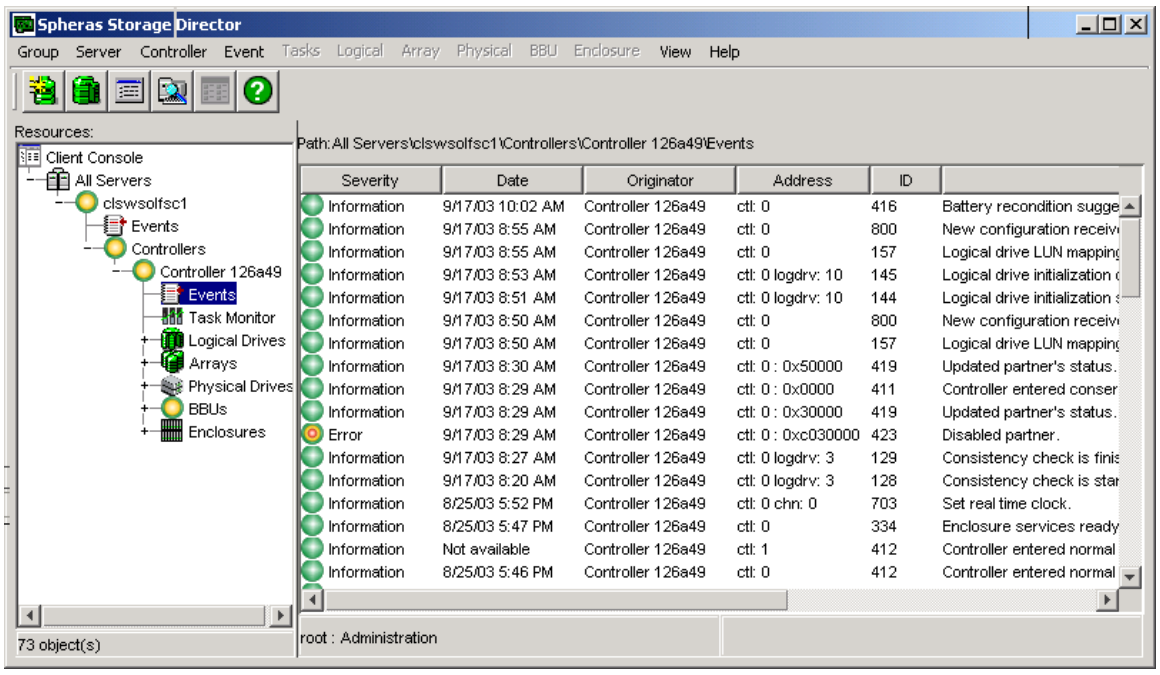

Figure 78 Selecting Events to View

Events display in the content area of Spheras Storage Director and can be any of the following. See 1.26 Status Icons on page 58 for a description of the icons representing the event status:

- Informational or Optimal (level 5)
  - An informational message showing a good status and no action is required.
  - Resources in an Optimal state are considered protected.
- Warning (level 4)
  - A warning or condition that does not require immediate user action, but requires close monitoring.
  - Resources in a Warning state are considered vulnerable.
- Error (level 3)
  - An error has occurred that did not cause data loss, usually a software error. You should retry the command.
- Serious (level 2)
  - A serious error device failure has occurred, usually a non-optimal logical drive or physical device failure. Action is needed to prevent data loss or loss of access to data. Hardware replacement may be required.
  - Resources in a Serious state are considered vulnerable.

- Critical (level 1)
  - A critical error or major device failure has occurred. Action is needed, but data or access to data may have already been lost. Call for service or replace hardware.
  - Resources in a Critical state are considered down.

The following information is displayed in the event table of the content area.

- The severity of the event.
- The date and time the event occurred.

Note: On Linux operating systems, the event log time-stamps may become off by one hour due to daylight saving time. For details on how the Linux system clock works, see

www.tldp.org/HOWTO/mini/Clock-2.html

- The controller where the event originated
- The controller address
- The event ID number
- The event description

6.1.2 Viewing Event Details

The Event Details dialog box displays event information for a single event. To access an Event Details dialog box, do the following:

1. Select an event in the content area that you want to see details about.
2. Click **Event->Event Details** from the menu bar.

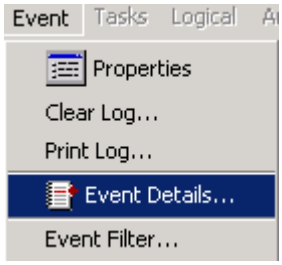

Figure 79 Select Event Details

The Event Details dialog box opens as shown in Figure 80 Event Details Dialog Box.

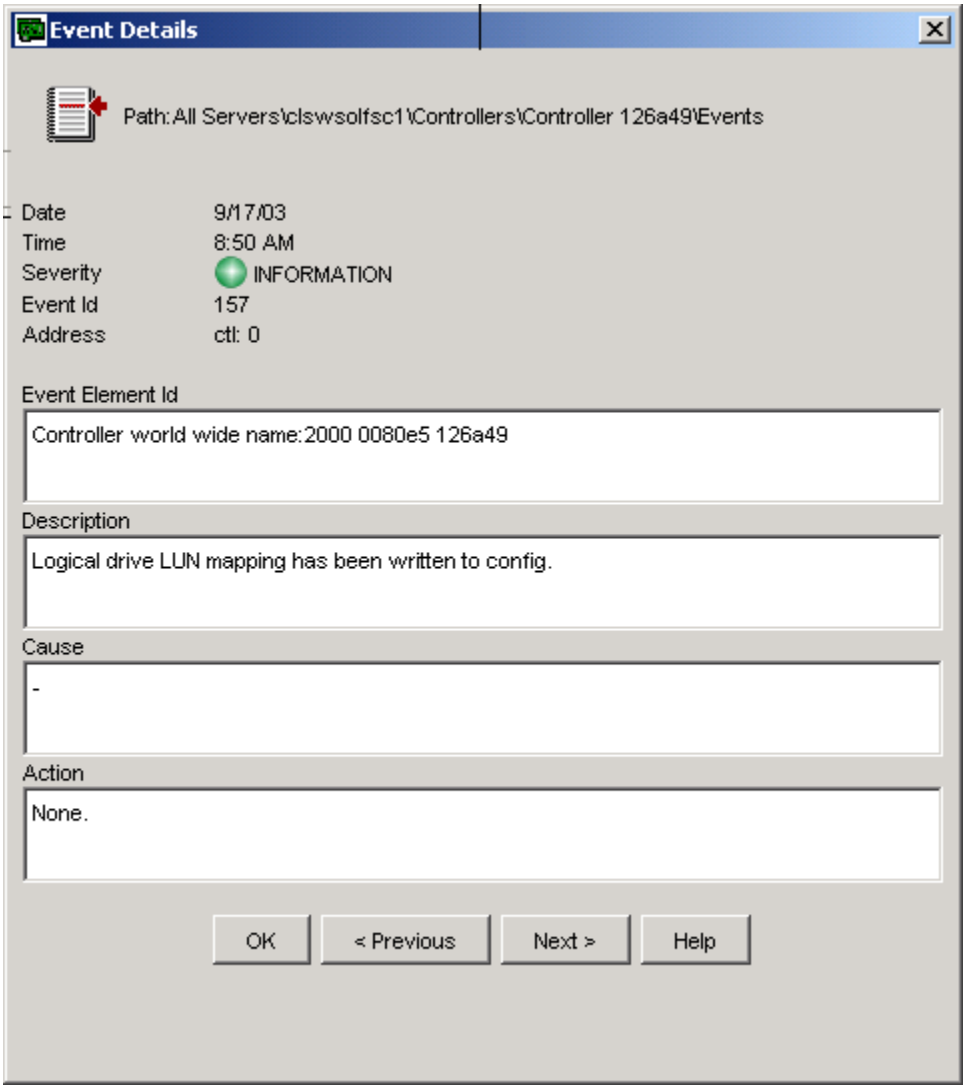

Figure 80 Event Details Dialog Box

The Event Details dialog box displays information about the selected event. The information includes if applicable (not shown in Figure 80 Event Details Dialog Box), the serial number of the physical drive and the World Wide Name of an enclosure. The **Previous** button changes the content of the display to the event that is earlier (downward) in the log. The **Next** button changes the display to the event that is later (upward) in the log. This information is helpful when troubleshooting an error. Click **Help** for recovery information regarding the event that you chose.

6.1.3 Filtering Events

The Event Filter dialog box allows you to choose which event severity levels to display in the content area.

1. Click **Event -> Event Filter** from the menu bar.

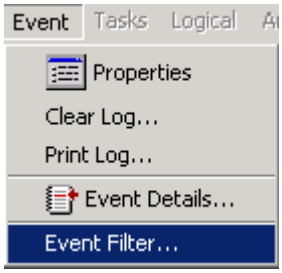

Figure 81 Select Event Filter

The Event Filer dialog box opens (Figure 82 Event Filter Dialog Box).

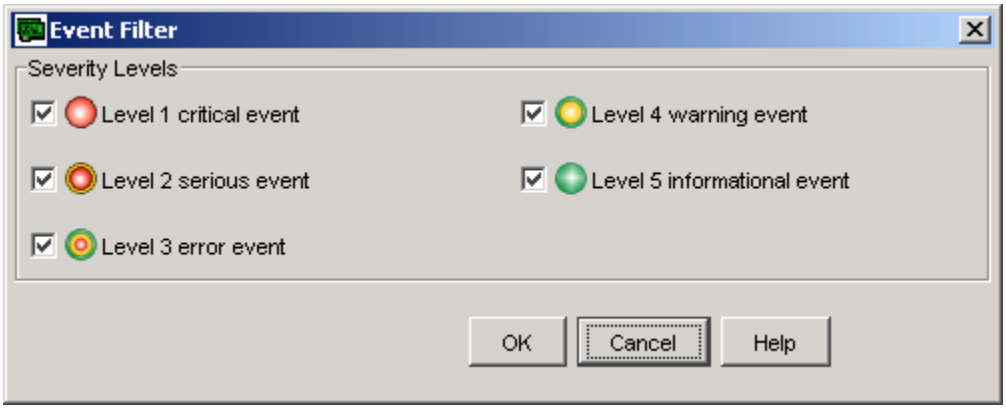

Figure 82 Event Filter Dialog Box

2. Check or uncheck the event types to display and click **OK**.

The events displayed in the content area are refreshed. The Event Filter settings are only saved for the current session. If you log out and log back onto the Client, your filter settings are lost.

6.1.4 Clearing the Event Log

To erase the entries in the event log, select **Event -> Clear Log** from the menu bar.

After you clear events for a particular server, you will not be able to retrieve them. If you have any doubt if will need the event log at a future date, use the **Print Log** menu option to print the log.

It is recommended that you clear the event log:

- When you have replaced a component of the Enclosure
- When the event log becomes too large to manage.

1.43 Monitoring a Controller

To monitor activities on a controller, select a controller from the navigation area.

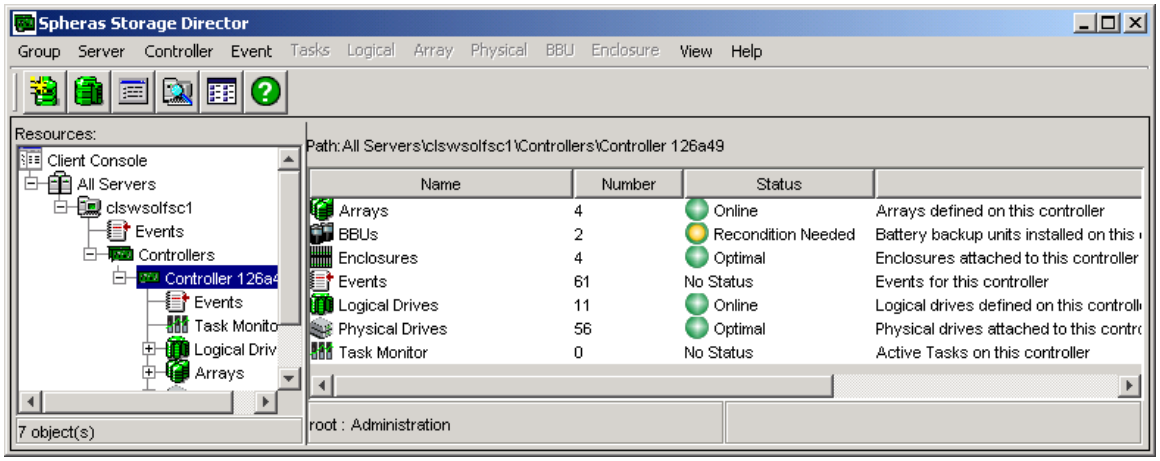

Figure 83 Monitoring a Controller (text view)

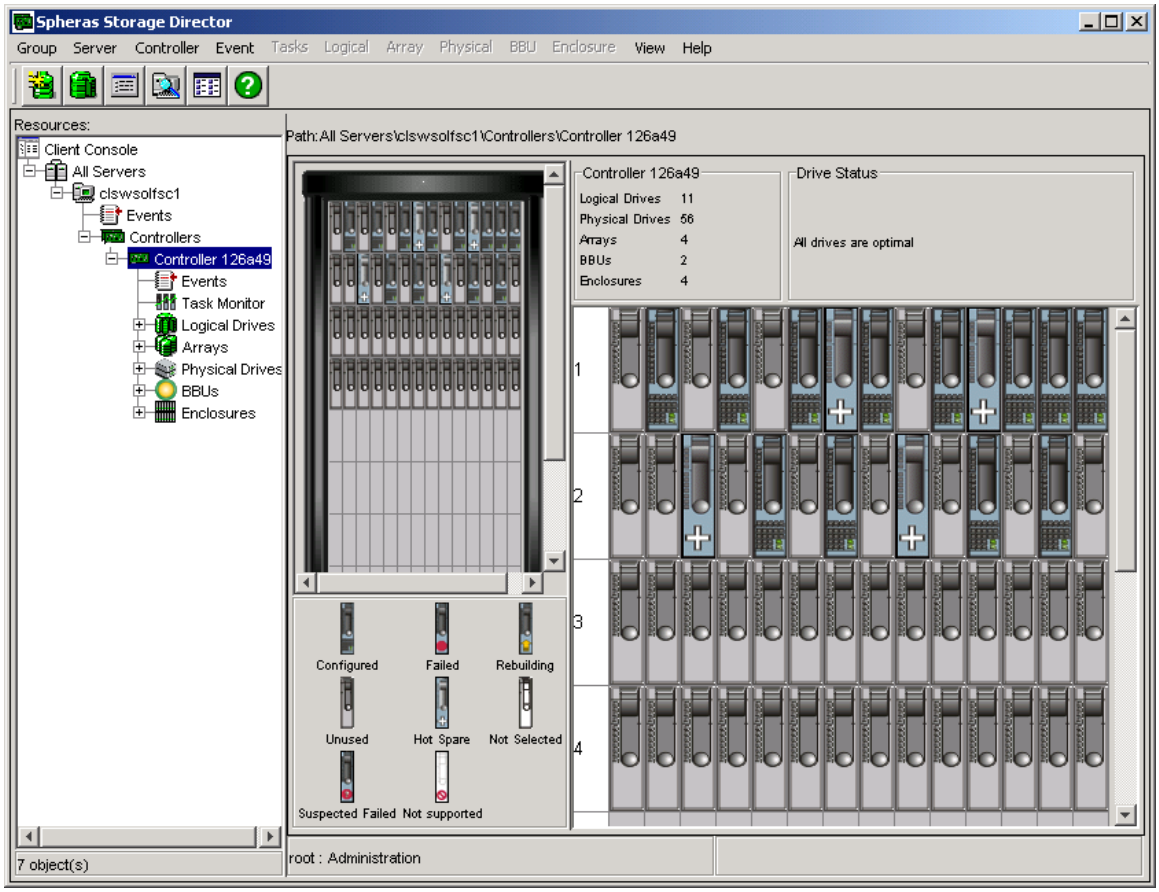

Figure 84 Monitoring a controller (physical view)

The content area displays all of the components of the controller including events related to the controller. The content area displays the information either in the Text View or the Physical View.

To get information about all controllers that are connected to a particular server, click on **Controllers** in the navigation area. This view displays the following information in the content area:

- The name of each controller.
- The model of each controller.
- The mode of each controller: simplex, duplex, or failover.
- The mode of the cache: normal or conservative cache.
- The status of each controller (determined by the most severe status of the controller itself or its components).
- The bus type of each controller: SCSI or Fibre.

If you are displaying the information using the Physical View, you can view controller information in the Controller Status pane. See "Setting and Modifying Controller Properties" on page 52 for information regarding the controller summary and status.

1.44 Monitoring an Array

This section describes how Spheras Storage Director monitors the capacity of an array.

To see a summary of the status of all arrays associated with a controller, click on **Arrays** in the navigation area. The following information is displayed in the content area:

- The name of each array.
- The capacity of each array.
- The unused capacity of each array.
- The status of each array. The status that is reported is determined by the most severe status of the array itself or all its logical drives. The logical drive status is the most severe status of the logical drive itself (for example, rebuilding) and that of its component physical drives (for example, physical drive failed).
- The number of logical drives in each array.
- The number of physical drives in each array.

To monitor activities on an individual array, select an array from the navigation area. There are two ways to view array information.

- Select an array and view the information in the content area
- Click **Array-> Properties** from the menu bar.

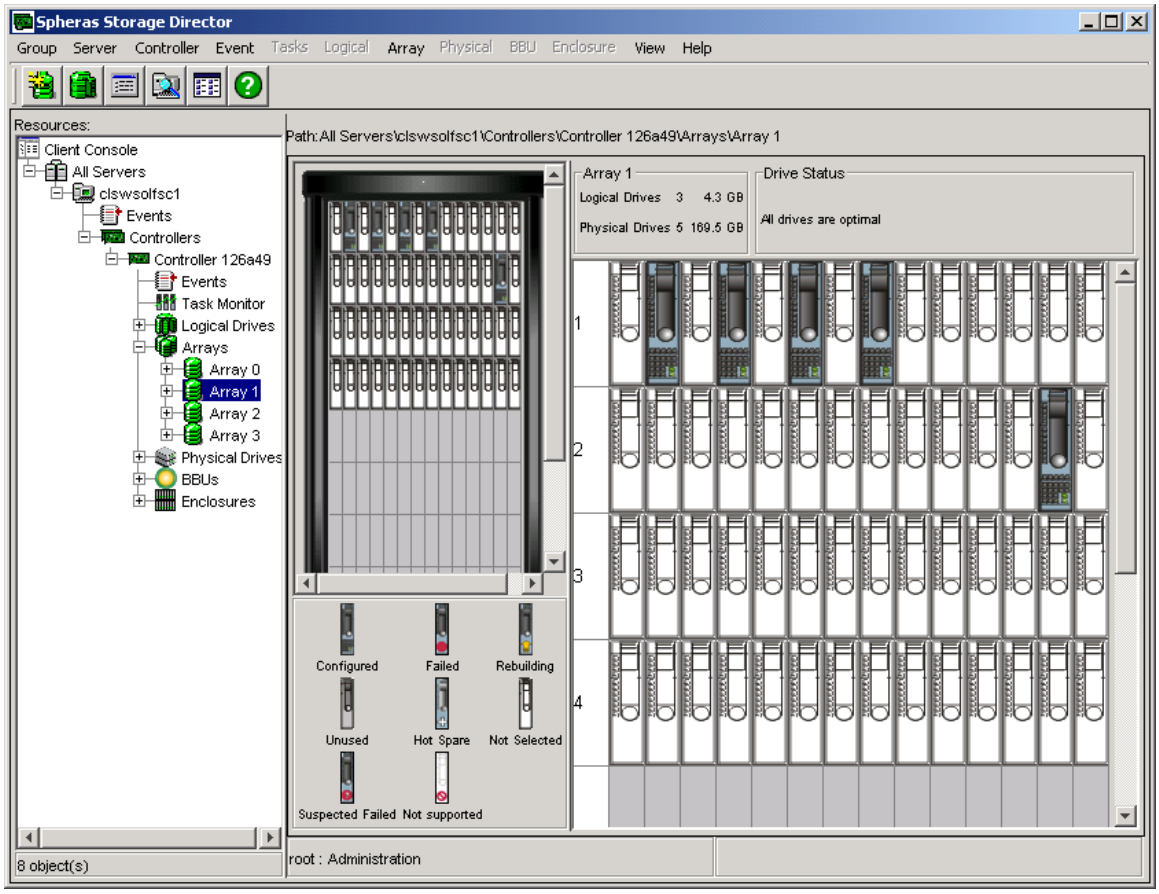

Figure 85 Displaying Array Information (physical view)

The content area displays array information either in the Text View or the Physical View. Both views contain the following information:

- The array number.
- The number and total capacity of the logical drives defined on the array.
- The number and total capacity of the physical drives on the array.
- The status of components in the array.

Note: In the Physical View shown in Figure 85 Displaying Array Information (physical view) there are drives shown with a status of Not Selected. This simply means that these drives are not part of the array that is being viewed.

To view Array Properties, click **Array->Properties** from the menu bar.

The Array Properties dialog box opens and displays the following information:

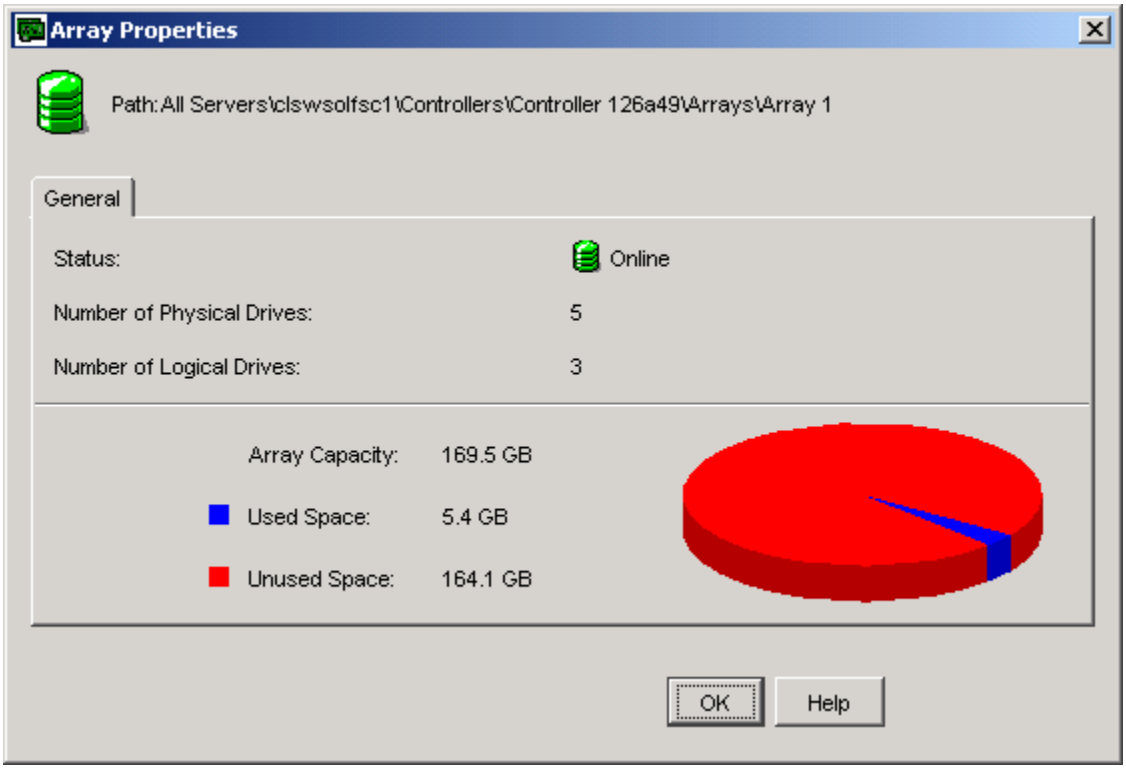

Figure 86 Select Array Properties

- The most severe status of any logical drive in the array.
- The number of physical drives in the array.
- The number of logical drives in the array.
- The total capacity of the array.
- The total amount of used and unused space held by this array.

1.45 Physical Disk and Logical Drive Monitoring

This section describes how Spheras Storage Director monitors physical disks and logical drives.

6.1.5 Displaying Physical Drive Information

There are two ways to view physical drive information:

- Select Physical Drives, Configured, Unused or Global Hot Spares from the navigation area and view the information in the content area
- Click on a physical drive in the content area and then click **Physical-> Properties** from the menu bar.

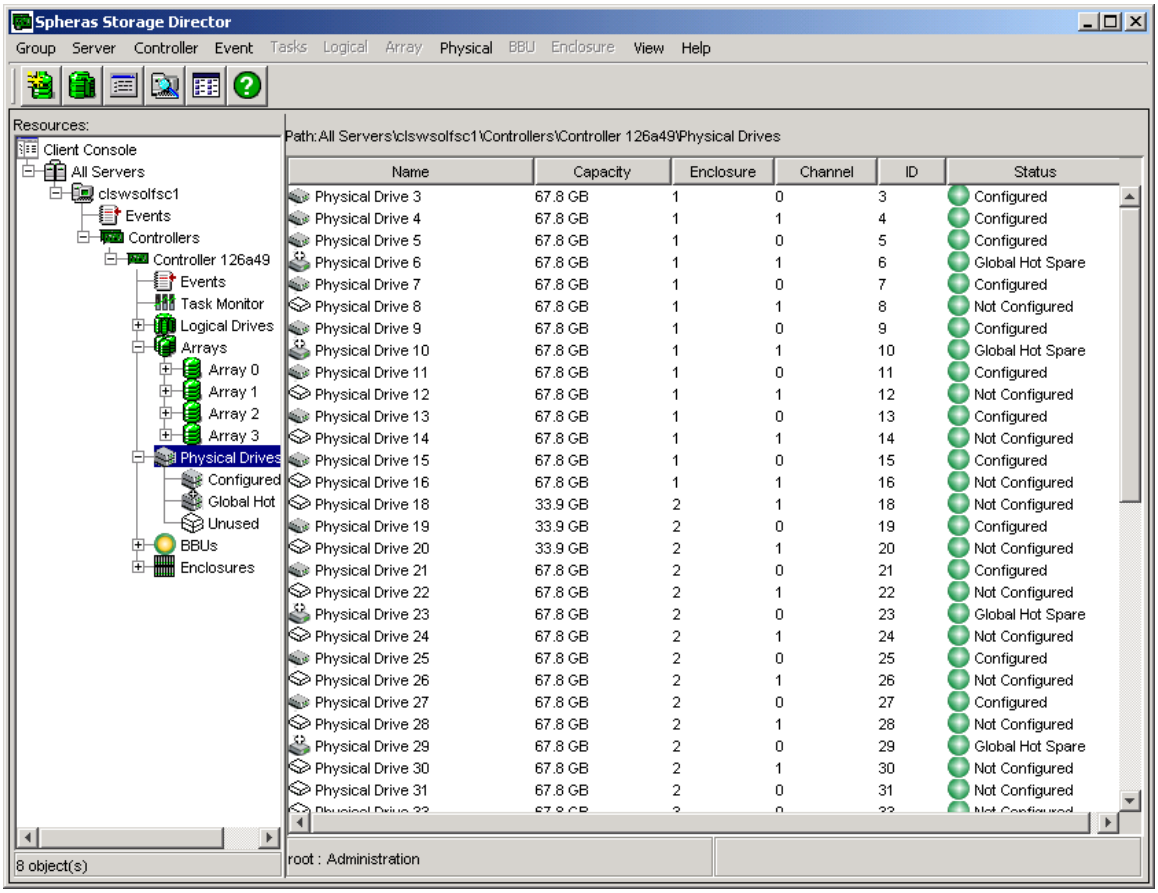

Figure 87 Displaying Physical Drive Information (text view)

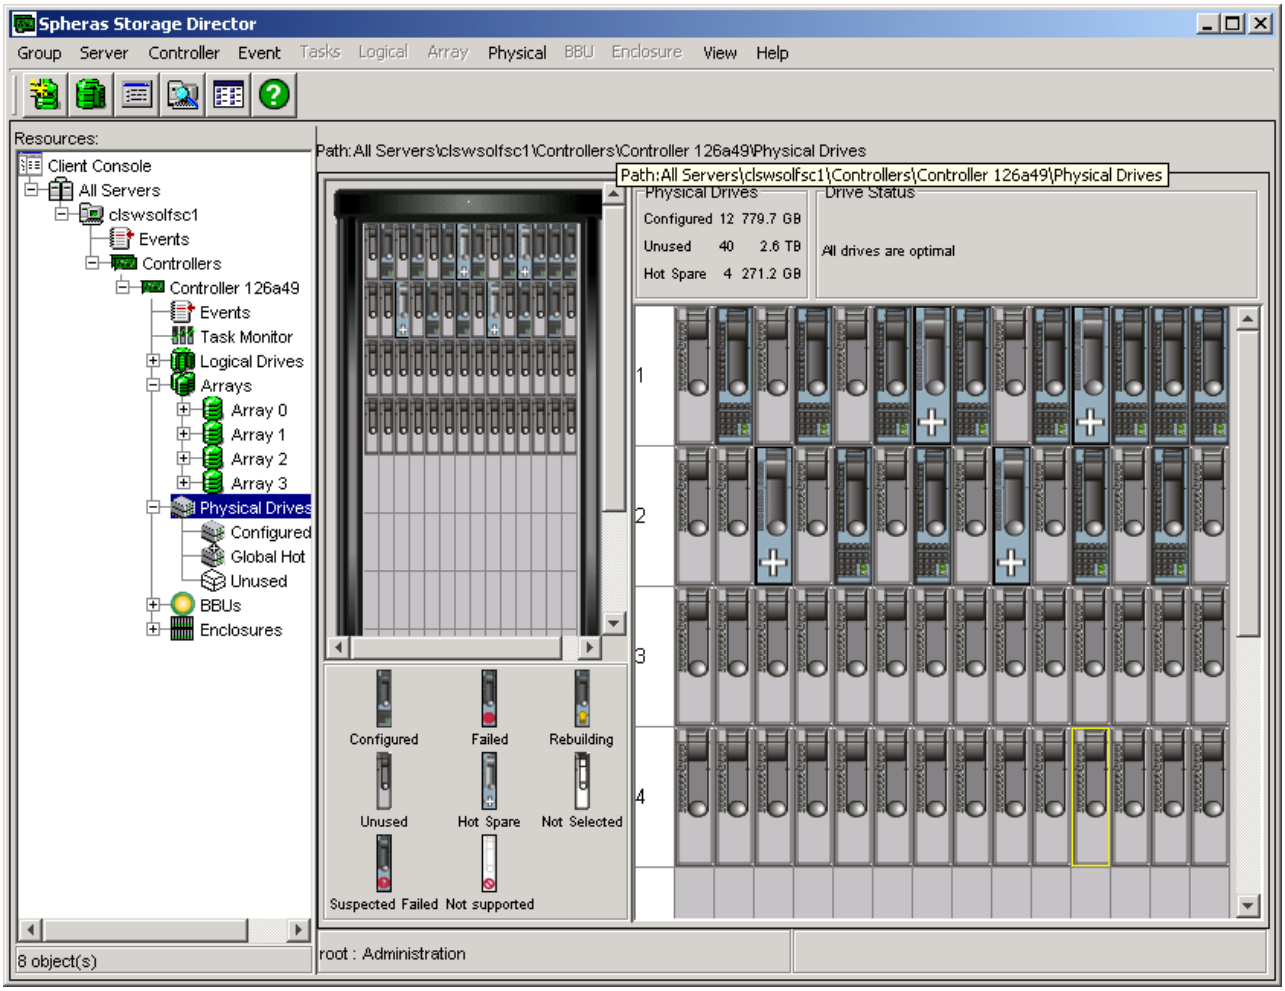

Figure 88 Displaying physical drive information (physical view)

The content area displays the physical drive information in either the Text View or the Physical View.

If you are displaying the information using the Text View, the content area displays information about the physical drive including name, capacity, enclosure, channel, ID, and Status (see Figure 87 Displaying Physical Drive Information (text view)).

If you are displaying the information using the Physical View, the content area displays information about the physical drive including how many drives are configured, unused, designated as hot spares, and their total capacity. Hover Help is also available for each physical drive pictured. Hover help is activated by moving your mouse over a particular physical drive. A text pop-up displays the drive name, size and channel.

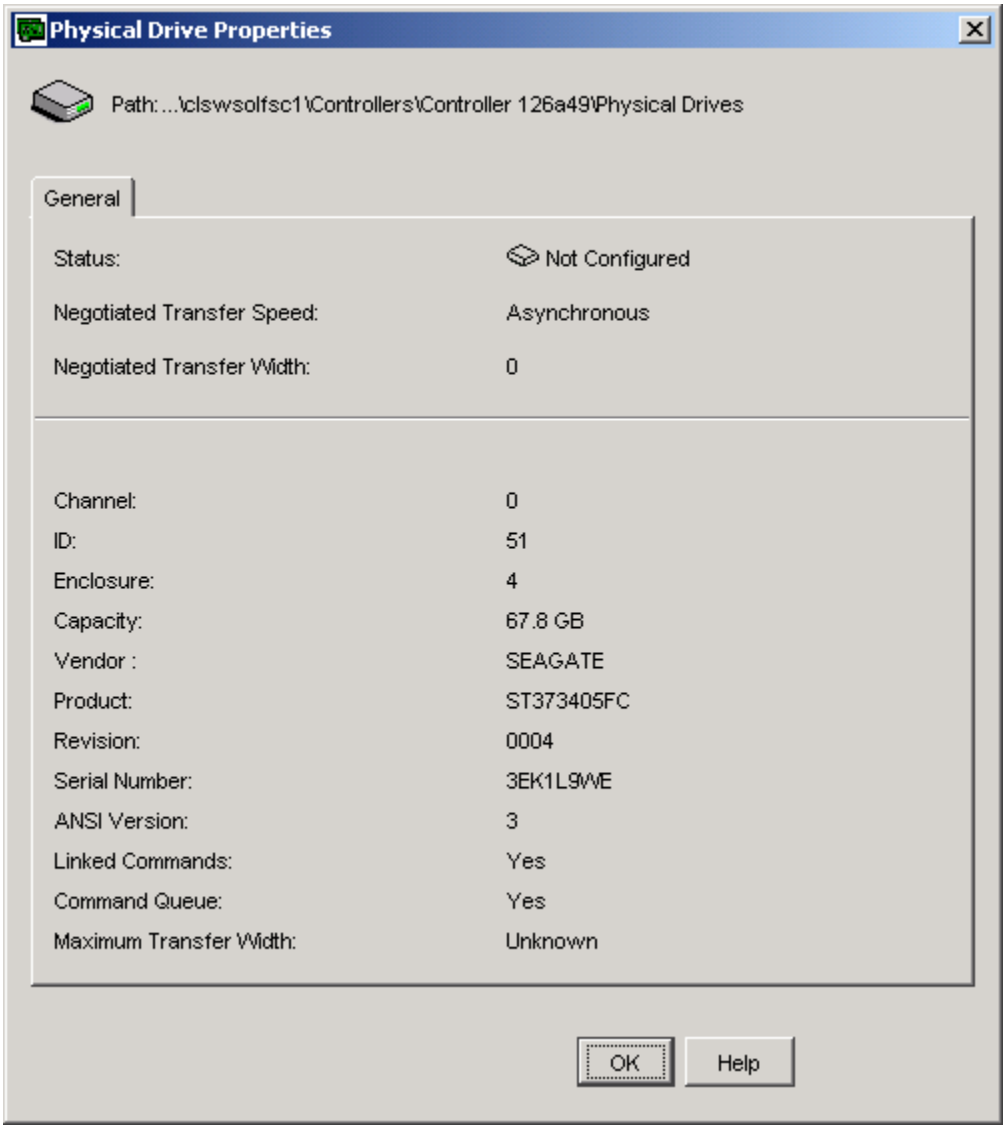

Figure 89 Physical Drive Properties Dialog Box

The Physical Drive Properties dialog box displays the following information:

- The current status of the disk (online, rebuilding, failed, or suspected failed)
- The drive channel negotiated transfer speed (MB/second)
- The negotiated drive channel bus width (in bits)
- The drive channel
- Loop ID number
- The associated enclosure number where the disk resides
- The physical capacity of the disk device
- The vendor or source of the drive

- The drive's product identification
- The revision level of the firmware on the drive
- The physical disk serial number
- The ANSI version supported
- Whether the following parameters are set to Yes or No: Linked Commands, and Command Queuing
- The maximum data transfer width (in bits)

6.1.6 Displaying Logical Drive Information

To monitor activities on a logical drive, select a logical drive from the navigation area.

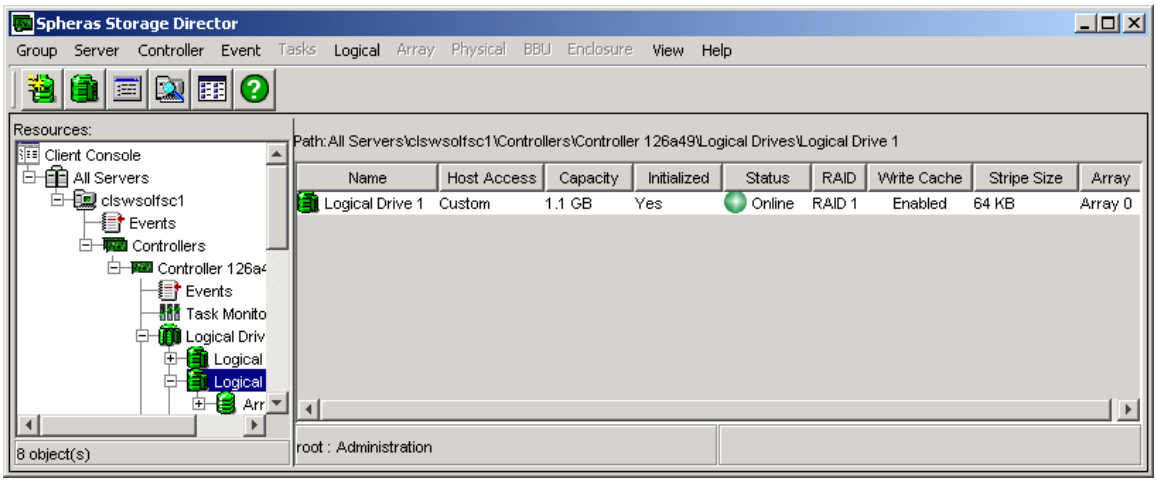

Figure 90 Displaying Logical Drive Information (text view)

The content area displays the logical drive information.

In the text view, the following information is displayed about the logical drive.

- The logical drive number
- The Host Access - displays Enabled, None, or Custom
- The capacity of this logical drive
- The operational status of this logical drive
- The RAID level at which the logical drive is configured
- The status of the write back cache (enabled or disabled)
- The stripe size in use by the logical drive
- The array number where the logical drive resides.

In the physical view, the following information is displayed:

- The LUN number
- The status of the host access (enabled or disabled)
- The capacity of this logical drive
- The status of the logical drive
- The RAID level of the logical drive
- The array number where the logical drive resides
- The drive status of physical drives within the logical drive..

There are two ways to view information about a specific logical drive:

- Select a logical drive in the navigation area and view the information in the content area
- Click **Logical->Properties** from the menu bar.

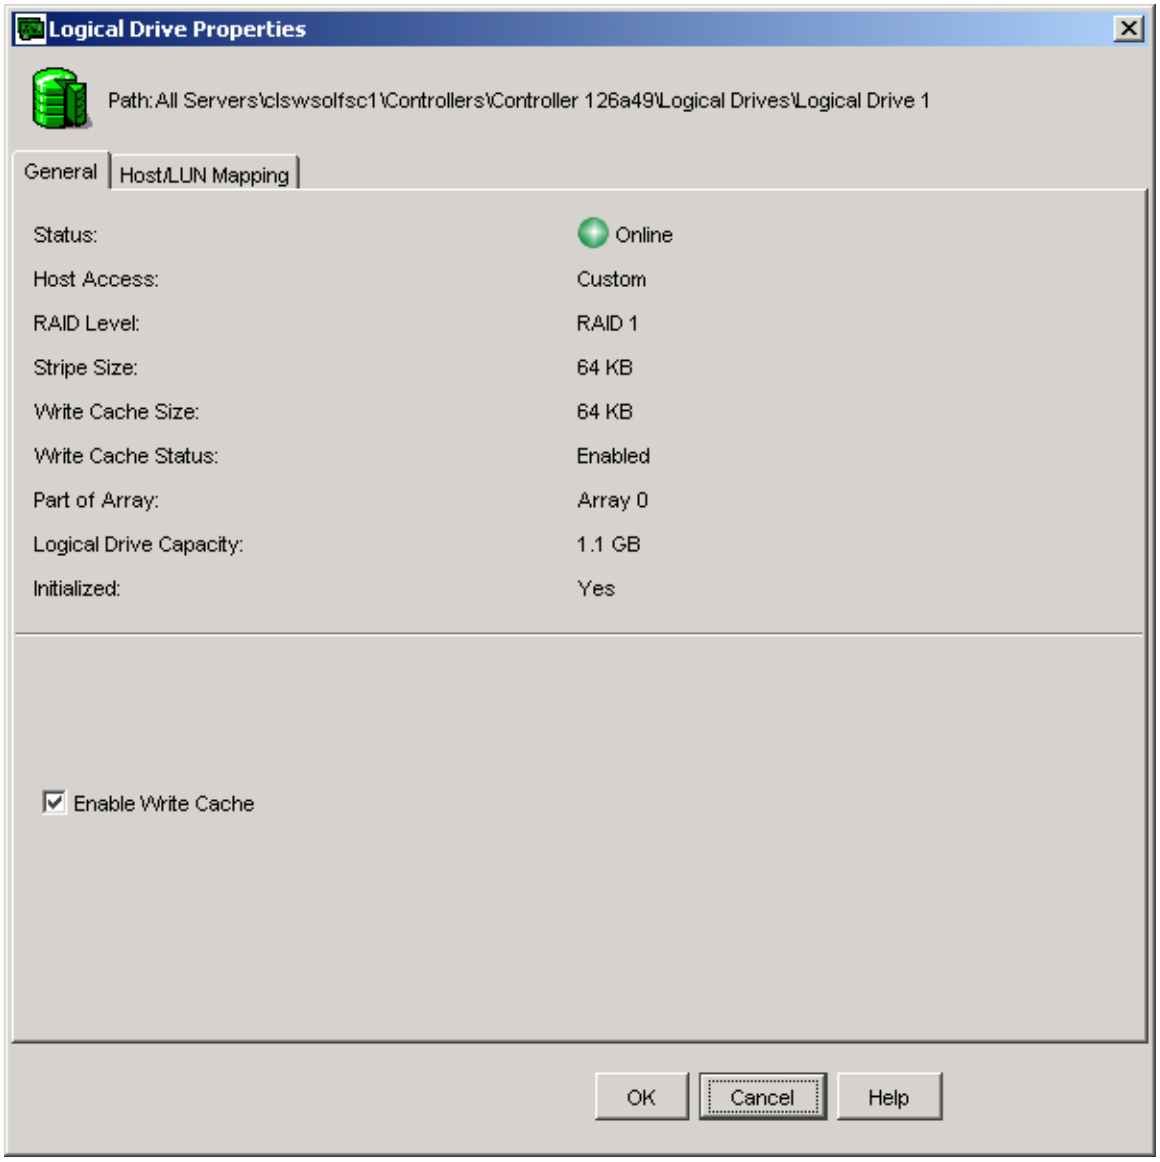

Figure 91 Logical Drive Properties

The Logical Drive Properties dialog box opens and displays information about the logical drive, see the previous information for a description of the contents. You can also enable or disable the write cache using the checkbox on this dialog.

The Host/LUN Mapping tab allows you to modify Host-to-LUN mapping values. See "5.1.3 Manual Storage Configuration" on page 95 for details about Host-to-LUN mapping.

To get general information about all logical drives, click on the Logical Drives icon in the navigation area. The following information is displayed:

- The name of each logical drive.
- The host access of each logical drive.
- The capacity or each logical drive.
- The status of each logical drive.

- The RAID level of each logical drive.
- Whether write back cache is enabled or disabled for each logical drive.
- The stripe size of each logical drive.
- The array number where each logical drive resides.

Click on an individual logical drive to view its information.

1.46 Monitoring an Enclosure

To monitor activities on an enclosure, select an enclosure from the navigation area. There are three ways to view enclosure information.

- Select the **Enclosures** icon in the navigation area and view the information in the content area
- Select an individual enclosure in the navigation or content area and view the information in the content area
- Click **Enclosure->Properties** from the menu bar.

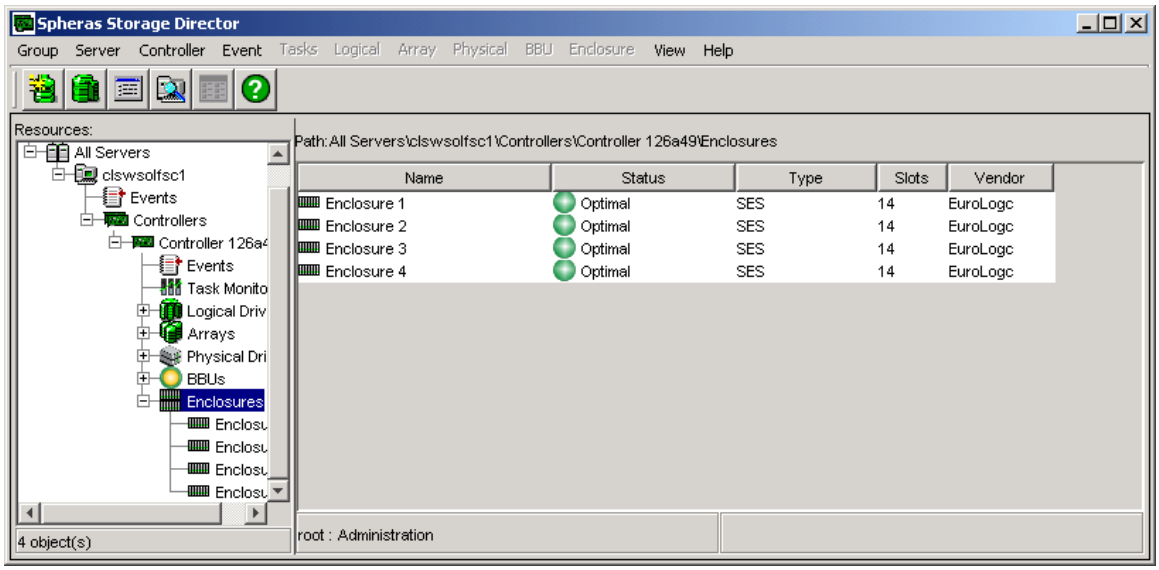

Figure 92 Displaying Enclosure Information

If you are displaying the information using the **Enclosures** icon, the content area displays general information about the enclosure including name (enclosure ID number), status, type of enclosure, number of slots, and enclosure vendor. The enclosure status that is reported is the most severe status of the enclosure itself or of the enclosure elements.

To view information about an individual enclosure, select an individual enclosure in the navigation area.

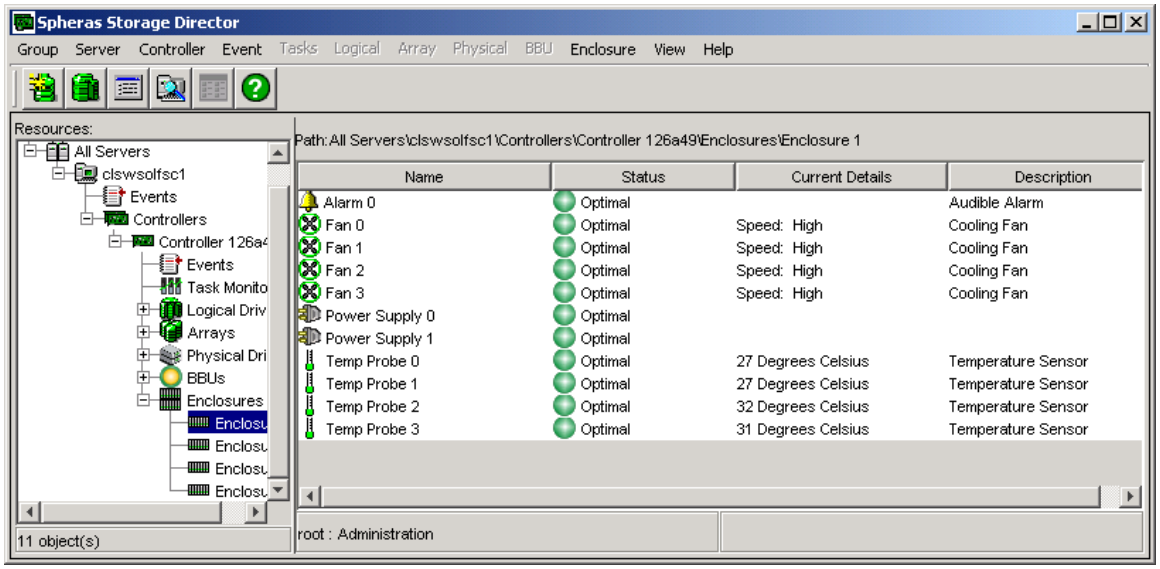

Figure 93 Individual Enclosure Information

When a specific enclosure is selected, the following is displayed:

- Alarms:
  - Name and icon
  - Status: Element Not Present, Optimal, Warning or Failed
  - Current details: blank or Alarm On
  - Description
- Fans:
  - Name and icon
  - Status: Element Not Present, Optimal, Critical, or Failed
  - Current details: Fan stopped, Speed: Low, or Speed: High
  - Description
- Power Supplies:
  - Name and icon
  - Status: Element Not Present, Optimal, or Failed
  - Current details: always blank
  - Description: always blank
- Temperature Sensors:
  - Name and icon
  - Status: Element Not Present, Optimal, Warning, or Failed
  - Current details: current temperature, Over Temperature Warning, Over Temperature Failed, Under Temperature Warning, or Under Temperature Failed
  - Description.
- UPS:

- Name and icon
- Status: Element Not Present, Optimal, Warning, or Failed
- Current details: always blank
- Description

To view Enclosure Properties and change the temperature thresholds, click **Enclosure->Properties**.

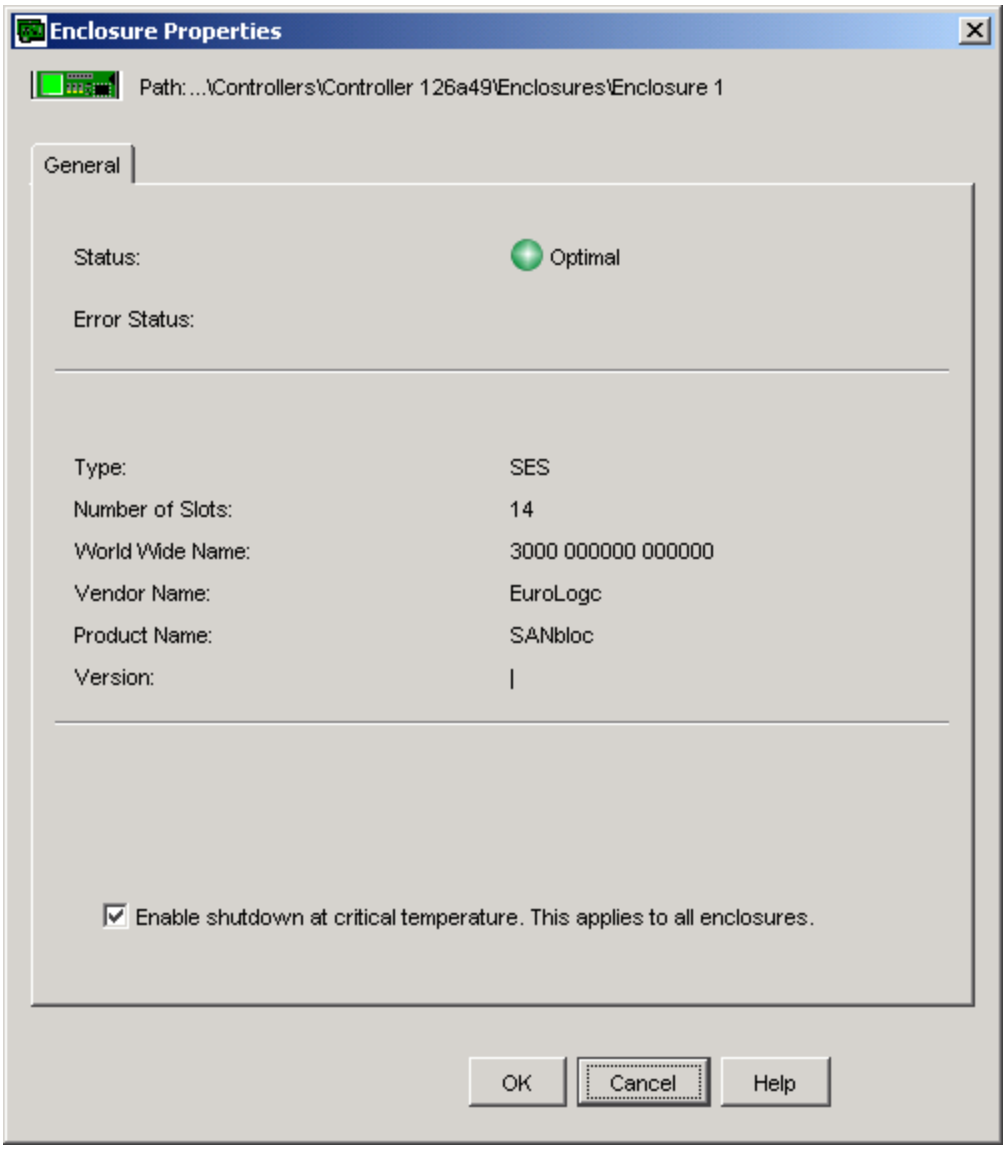

Figure 94 Viewing Enclosure Properties

– The Enclosure Properties dialog box opens and displays Figure 94 Viewing Enclosure Properties.

Notes:

1. When you make changes to the Enclosure Properties it takes a few minutes before the changes are displayed.

2. If the SSD server is unable to retrieve the information from the controller, the default values are displayed.

3. Some enclosures do not support setting temperature thresholds. For those enclosures, the values are not displayed in the dialog.

- Shutdown at critical temperature: Enable as needed. When a warning threshold is reached, the controller sends an event. When the critical threshold is reached, the controller spins down the drives and powers off the enclosure.

1.47 Monitoring the Battery Backup Unit

The optional Battery Backup Unit (BBU) maintains memory content in the presence of an ac power failure. The principal purpose of the BBU is to provide stable memory power during ac power glitches and short power outages, however, the BBU is capable of sustaining memory content for an extended period. The length of memory power backup is dependent upon the cache memory size and the particular battery pack used by the BBU.

To monitor the charge status on a Battery Backup Unit (BBU), select a BBU from the navigation area. To view BBU information, click **BBU->Properties** from the menu bar.

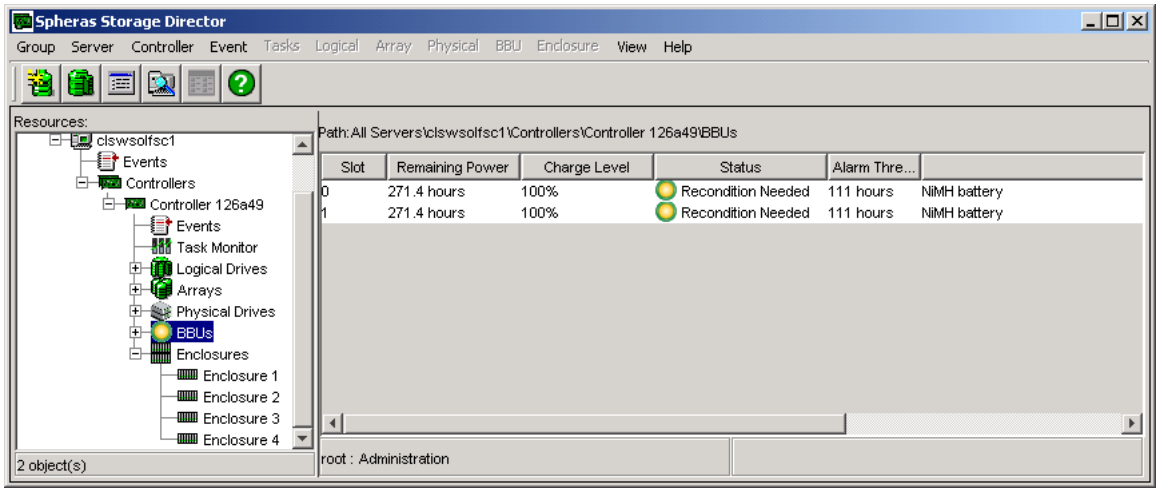

Figure 95 BBU Information

Note: The example shown above is displayed if you select **BBUs** from the tree. If you select a specific BBU, the slot number is not shown because it is part of the BBU name.

To view BBU Properties, click **BBU->Properties**.

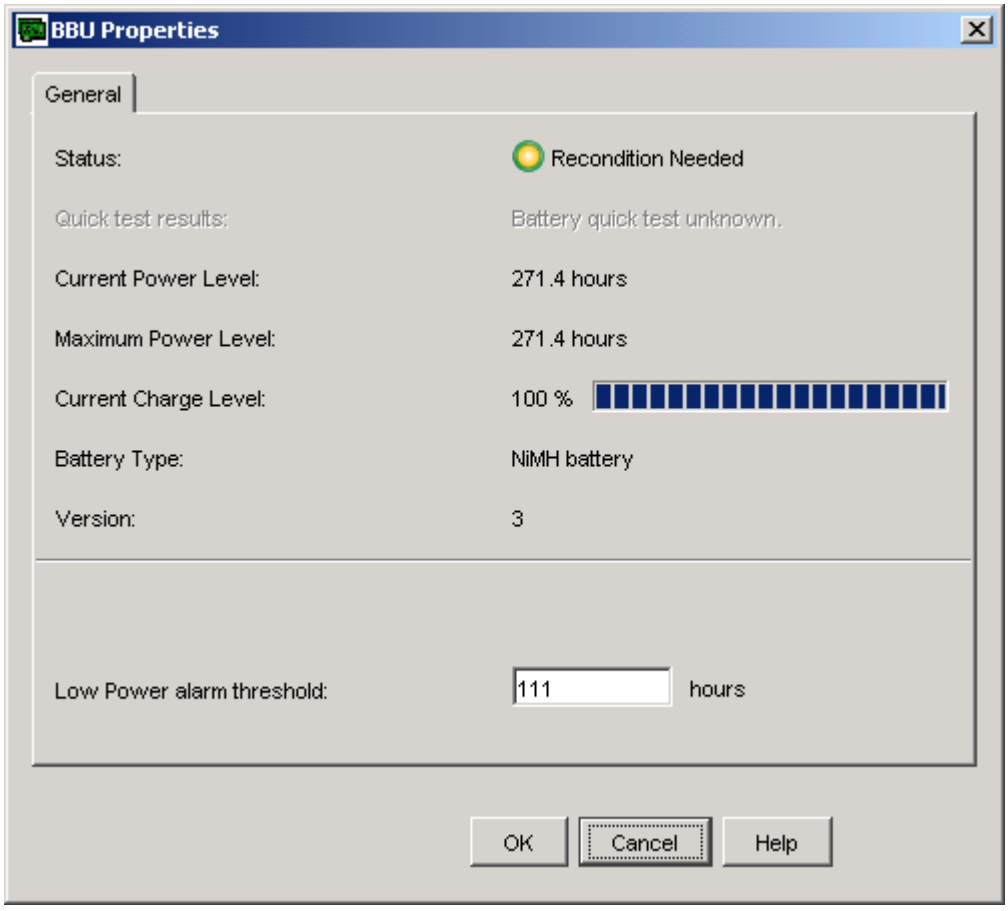

Figure 96 BBU Properties Dialog Box

The BBU Properties dialog box opens and the following information is displayed.

- The status of the BBU is:
  - Optimal
  - Low Power
  - Fast Charging
  - Discharging
  - Recondition Needed
  - Recondition (Fast Charging)
  - Recondition (Discharging)
  - Recondition Active
- The current power level of the battery expressed as hours or minutes of charge. This value changes as the battery is discharged or charged.
- The maximum power or highest level the battery has been charged up to expressed as hours or minutes of charge.
- The current charge level, where 100% means the battery has been charged up to its maximum level.

- The type of battery, which is dependent on controller hardware.
- Version number for the BBU type.
- The ability to edit the low power alarm threshold.
- The ability to schedule or run a quick test.

1.48 Monitoring Long Operation Tasks

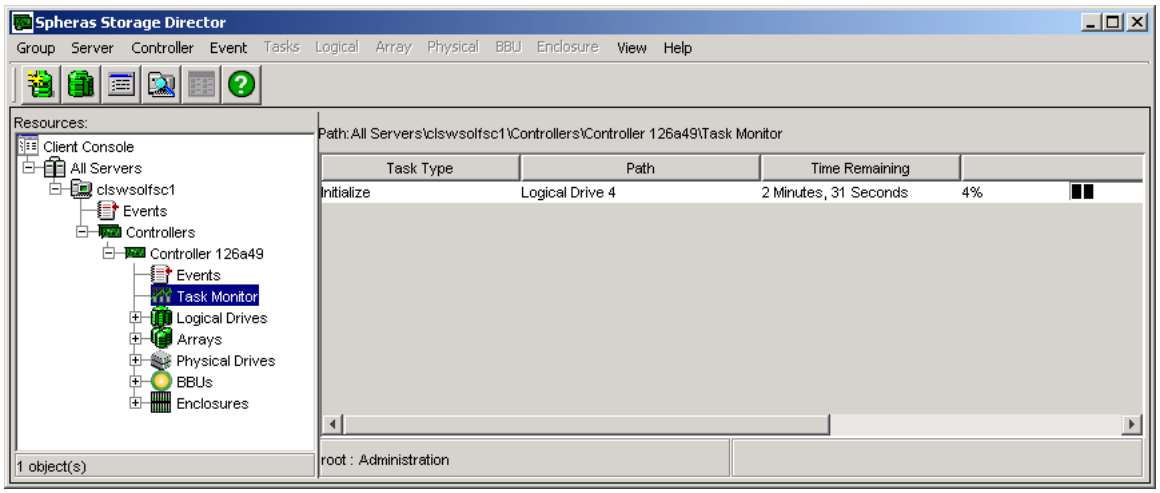

To monitor long operation tasks, select the Task Monitor from the navigation area.

Figure 97 Task Monitoring

The in progress tasks are displayed in the content area. The following task types are displayed.

- Logical drive initialization
- Logical drive consistency check
- Disk array capacity expansion
- Array rebuild
- BBU Fast Charging and BBU Discharging
- Disk Scrubbing
- Controller shutdown

The task monitor view displays the following information for each task in progress:

- The task type as listed above.
- The path where the task is taking place.
- Time Remaining to complete the task.

- A completion status for the task. For example, if the task is BBU Fast Charging, the completion status would display a percentage complete value.

To cancel a long operation task, click the task you wish to cancel and click **Task->Cancel**.

CAUTION: Canceling either, a logical drive rebuild or a logical drive consistency check, should only be done, if absolutely necessary. Canceling either of these long operations leaves the logical drive unstable, which may result in data loss.

Disk array capacity expansion cannot be canceled. This option is grayed out during the expansion process.

Event ID	Type of Event	Description	What to Do About It
388	Critical	Controller is dead. System is disconnecting from this controller.	Replace the controller and apply a configuration.
389	Warning	Controller has been reset.	Check the controller's configuration and restore as necessary. (Resetting a controller may cause loss of data or loss of access to data.)
390	Informational	Controller is found.	Nothing.
391	Critical	Controller is gone. System is disconnecting from this controller.	Replace the controller and apply a configuration.
395	Critical	Controller is gone. System is disconnecting from this controller.	Replace the controller and apply a configuration.
396	Informational	Controller powered on.	Nothing.
397	Informational	Controller is online.	Nothing.
398	Critical	Controller is gone. System is disconnecting from this controller.	Replace the controller and apply a configuration.
399	Serious	Controller's partner is gone, controller is in failover mode now.	Replace the controller, apply a configuration, and return the controller to duplex mode.

Event		
1	Name	MLXEV_PHYSDEV_ONLINE
	Description	A physical disk has been placed online.
	Cause	Rebuild completed. Physical disk was configured. Manual online was done.
	Action	None
	Address	ctl: 0 chn: 13 tgt: 14
	Severity	4 (Information)
	Source	ExPro InPro
2	Name	MLXEV_PHYSDEV_HOTSPARE
	Description	A physical disk was added as a hot spare.
	Cause	Device was configured. Manual hot spare was done. Automatic hot spare was done.
	Action	None
	Address	ctl: 0 chn: 13 tgt: 14
	Severity	4 (Information)
	Source	InPro
3	Name	MLXEV_PHYSDEV_HARD_ERROR
	Description	Physical disk error found.
	Cause	A bad sector was found on the physical disk. Mechanical failure on the physical disk. Host SCSI device detected illegal instruction. Target device generated unknown phase sequence.
	Action	If problem occurs frequently: – Replace the physical disk. – Contact your service representative.
	Address	ctl: 0 chn: 13 tgt: 14
	Severity	2 (Error)
	Source	ExPro
4	Name	MLXEV_PHYSDEV_PFA

	Description	Physical disk PFA condition found. This disk may fail soon.
	Cause	Physical disk predicted some future failure. External RAID logical device may have become critical.
	Action	Refer to the enclosure manufacturer’s service manual.
	Address	ctl: 0 chn: 13 tgt: 14
	Severity	2 (Error)
	Source	ExPro
5	Name	MLXEV_PHYSDEV_AUTO_REBUILD_START
	Description	An automatic rebuild has started.
	Cause	A physical disk failed and spare was available. A physical disk failed and no spare was available. A spare was added.
	Action	None
	Address	ctl: 0 chn: 13 tgt: 14
	Severity	4 (Information)
	Source	ExPro InPro
6	Name	MLXEV_PHYSDEV_MANUAL_REBUILD_START
	Description	A rebuild has started.
	Cause	Client started the rebuild on user's request. User replaced the failed device and 'raidbld' started the rebuild.
	Action	None
	Address	ctl: 0 chn: 13 tgt: 14
	Severity	4 (Information)
	Source	ExPro InPro
7	Name	MLXEV_PHYSDEV_REBUILD_DONE
	Description	Rebuild is over.
	Cause	Rebuild completed successfully.
	Action	None
	Address	ctl: 0 chn: 13 tgt: 14
	Severity	4 (Information)

	Source	ExPro InPro
8	Name	MLXEV_PHYSDEV_REBUILD_CANCELED
	Description	Rebuild is cancelled.
	Cause	User cancelled the rebuild. Higher priority rebuild started.
	Action	Restart the rebuild if required.
	Address	ctl: 0 chn: 13 tgt: 14
	Severity	4 (Information)
	Source	ExPro InPro
9	Name	MLXEV_PHYSDEV_REBUILD_ERROR
	Description	Rebuild stopped with error.
	Cause	Due to some unknown error on the controller, rebuild failed.
	Action	Try rebuild again.
	Address	ctl: 0 chn: 13 tgt: 14
	Severity	2 (Error)
	Source	ExPro InPro
10	Name	MLXEV_PHYSDEV_REBUILD_NEWDEV_FAILED
	Description	Rebuild stopped with error. New physical disk failed.
	Cause	New physical disk failed. New physical disk may not be compatible with MDAC hardware/firmware.
	Action	Replace the physical disk.
	Address	ctl: 0 chn: 13 tgt: 14
	Severity	2 (Error)
	Source	ExPro InPro
11	Name	MLXEV_PHYSDEV_REBUILD_SYSDEV_FAILED
	Description	Rebuild stopped because logical drive failed.
	Cause	At least one more physical disks failed in the array. Bad data table overflow.
	Action	It may not be possible to recover from this error. Contact your service representative.

	Address	ctl: 0 chn: 13 tgt: 14
	Severity	2 (Error)
	Source	ExPro InPro
12	Name	MLXEV_PHYSDEV_DEAD
	Description	A physical disk has failed.
	Cause	A physical disk failed. A user action caused the physical disk to fail.
	Action	Replace the physical disk.
	Address	ctl: 0 chn: 13 tgt: 14
	Severity	1 (Serious)
	Source	ExPro InPro
13	Name	MLXEV_PHYSDEV_FOUND
	Description	A new physical disk has been found.
	Cause	A physical disk has been powered on. A new physical disk has been added. Controller was powered on. Controller was added. System has rebooted.
	Action	None
	Address	ctl: 0 chn: 13 tgt: 14
	Severity	4 (Information)
	Source	ExPro InPro
14	Name	MLXEV_PHYSDEV_GONE
	Description	A physical disk has been removed.
	Cause	User removed an unconfigured physical disk. An unconfigured physical disk failed. A controller was removed. A controller powered off.
	Action	Replace the device if needed.
	Address	ctl: 0 chn: 13 tgt: 14
	Severity	4 (Information)
	Source	ExPro InPro
15	Name	MLXEV_PHYSDEV_UNCONFIGURED

	Description	A previously configured disk is now available.
	Cause	User set the physical device to unconfigured.
	Action	None
	Address	ctl: 0 chn: 13 tgt: 14
	Severity	4 (Information)
	Source	InPro
16	Name	MLXEV_PHYSDEV_EXPANDCAPACITY_START
	Description	Expand Capacity started.
	Cause	User started the RAID expansion operation. A suspended RAID expansion operation was started.
	Action	None
	Address	ctl: 0 chn: 13 tgt: 14
	Severity	4 (Information)
	Source	InPro
17	Name	MLXEV_PHYSDEV_EXPANDCAPACITY_DONE
	Description	Expand capacity completed.
	Cause	RAID expansion finished.
	Action	None
	Address	ctl: 0 chn: 13 tgt: 14
	Severity	4 (Information)
	Source	InPro
18	Name	MLXEV_PHYSDEV_EXPANDCAPACITY_ERROR
	Description	Expand capacity stopped with error.
	Cause	Multiple physical devices failed.
	Action	It may not be possible to recover from this error. Contact your service representative.
	Address	ctl: 0 chn: 13 tgt: 14
	Severity	2 (Error)
	Source	InPro
19	Name	MLXEV_PHYSDEV_COMMAND_TIMEOUT

	Description	SCSI command timeout on physical device.
	Cause	Physical device has been removed. Physical device failed. Command time out value is not correct.
	Action	Refer to the enclosure manufacturer's service manual.
	Address	ctl: 0 chn: 13 tgt: 14
	Severity	2 (Error)
	Source	ExPro InP_old
20	Name	MLXEV_PHYSDEV_COMMAND_ABORT
	Description	SCSI command abort on physical disk.
	Cause	User may have requested to abort the command. Firmware may have aborted the command to recover from error. The physical disk may have aborted the command.
	Action	None
	Address	ctl: 0 chn: 13 tgt: 14
	Severity	0 (Critical)
	Source	InP_old
21	Name	MLXEV_PHYSDEV_COMMAND_RETRIED
	Description	SCSI command retried on physical disk.
	Cause	The command may have timed out. Bus reset may have occurred. Device reset may have occurred.
	Action	None
	Address	ctl: 0 chn: 13 tgt: 14
	Severity	3 (Warning)
	Source	InP_old
22	Name	MLXEV_PHYSDEV_PARITY_ERROR
	Description	Parity error found.

	Cause	A physical device did not generate proper parity. The controller failed; did not check parity properly. Cable failed. Improper cable length. Another physical device interfered. Some outside environment affected the data on the cable (such as a radio frequency signal). Terminator is not connected. Improper termination.
	Action	Refer to the enclosure manufacturer's service manual.
	Address	ctl: 0 chn: 13 tgt: 14
	Severity	2 (Error)
	Source	ExPro InP_old
23	Name	MLXEV_PHYSDEV_SOFT_ERROR
	Description	Soft error found.
	Cause	An error was detected by a physical device and data was recovered.
	Action	Run consistency check. If problem occurs frequently, replace the physical device.
	Address	ctl: 0 chn: 13 tgt: 14
	Severity	3 (Warning)
	Source	InP_old
24	Name	MLXEV_PHYSDEV_MISC_ERROR
	Description	Miscellaneous error found.
	Cause	A physical device reported some error that does not fit in any category. Read/Write command time out. Data over run. Physical device was busy when host attempted to send command.
	Action	If problem occurs frequently, replace the physical device.
	Address	ctl: 0 chn: 13 tgt: 14
	Severity	3 (Warning)
	Source	InP_old
25	Name	MLXEV_PHYSDEV_RESET
	Description	SCSI device reset.
	Cause	Firmware has reset to recover from error. User has reset.

	Action	None
	Address	ctl: 0 chn: 13 tgt: 14
	Severity	4 (Information)
	Source	ExP_old InPro
26	Name	MLXEV_PHYSDEV_ACTIVESPARE
	Description	Active spare found.
	Cause	Physical disk was configured. Manual active spare was done. Automatic active spare was done
	Action	None
	Address	ctl: 0 chn: 13 tgt: 14
	Severity	4 (Information)
	Source	InP_old
27	Name	MLXEV_PHYSDEV_WARMSPARE
	Description	Warm spare found.
	Cause	Physical disk was configured. Manual warm spare was done. Automatic warm spare was done.
	Action	None
	Address	ctl: 0 chn: 13 tgt: 14
	Severity	4 (Information)
	Source	InP_old
28	Name	MLXEV_PHYSDEV_REQSENSE
	Description	Request sense data available.
	Cause	A physical device reported an error. Firmware reported an operational error.
	Action	Read the request sense data to understand the root cause.
	Address	ctl: 0 chn: 13 tgt: 14
	Severity	2 (Error)
	Source	InPro
29	Name	MLXEV_PHYSDEV_INIT_STARTED

	Description	Initialization started.
	Cause	Host started the initialization.
	Action	Wait until the initialization is completed. If the system is shut down before initialization is completed, the physical device can be made useful only by reinitializing it.
	Address	ctl: 0 chn: 13 tgt: 14
	Severity	4 (Information)
	Source	InPro
30	Name	MLXEV_PHYSDEV_INIT_DONE
	Description	Initialization completed.
	Cause	Physical device initialization completed successfully.
	Action	None
	Address	ctl: 0 chn: 13 tgt: 14
	Severity	
	Source	InPro
31	Name	MLXEV_PHYSDEV_INIT_FAILED
	Description	Initialization failed.
	Cause	Physical device could have problems supporting the SCSI format command.
	Action	Try to initialize again. Contact your service representative.
	Address	ctl: 0 chn: 13 tgt: 14
	Severity	2 (Error)
	Source	InPro
32	Name	MLXEV_PHYSDEV_INIT_CANCELED
	Description	Initialization canceled.
	Cause	User cancelled the operation.
	Action	Physical disk must be initialized again or the physical disk cannot be used.
	Address	ctl: 0 chn: 13 tgt: 14
	Severity	1 (Serious)

	Source	InP_old
33	Name	MLXEV_PHYSDEV_WRITEREC_DEAD
	Description	A physical disk failed because write recovery failed.
	Cause	Write recovery process failed.
	Action	Replace physical disk and rebuild it.
	Address	ctl: 0 chn: 13 tgt: 14
	Severity	1 (Serious)
	Source	InPro
34	Name	MLXEV_PHYSDEV_RESET_DEAD
	Description	A physical disk failed because SCSI bus reset failed.
	Cause	SCSI bus reset failed.
	Action	Replace physical disk and rebuild it.
	Address	ctl: 0 chn: 13 tgt: 14
	Severity	1 (Serious)
	Source	InPro
35	Name	MLXEV_PHYSDEV_DBLCC_DEAD
	Description	A physical disk failed because double check condition occurred.
	Cause	Double check condition occurred.
	Action	Replace physical disk and rebuild it.
	Address	ctl: 0 chn: 13 tgt: 14
	Severity	1 (Serious)
	Source	InPro
36	Name	MLXEV_PHYSDEV_REMOVED_DEAD
	Description	A physical disk failed because disk is missing.
	Cause	Access to the physical disk failed.
	Action	Replace physical disk and rebuild it.
	Address	ctl: 0 chn: 13 tgt: 14
	Severity	1 (Serious)
	Source	InPro

37	Name	MLXEV_PHYSDEV_GROSSERR_DEAD
	Description	A physical disk failed because of gross error on SCSI processor.
	Cause	Gross error occurred to the on-board SCSI processor.
	Action	Replace physical disk and rebuild it.
	Address	ctl: 0 chn: 13 tgt: 14
	Severity	1 (Serious)
	Source	InPro
38	Name	MLXEV_PHYSDEV_BADTAG_DEAD
	Description	A physical disk failed because of invalid tag.
	Cause	The device responded with an invalid tag.
	Action	Replace physical disk and rebuild it.
	Address	ctl: 0 chn: 13 tgt: 14
	Severity	1 (Serious)
	Source	InPro
39	Name	MLXEV_PHYSDEV_SCSITMO_DEAD
	Description	A physical disk failed because a command to the disk timed out.
	Cause	SCSI command timed out on the device.
	Action	Replace physical disk and rebuild it.
	Address	ctl: 0 chn: 13 tgt: 14
	Severity	1 (Serious)
	Source	InPro
40	Name	MLXEV_PHYSDEV_SYSRESET_DEAD
	Description	A physical disk failed because of the system reset.
	Cause	System reset occurred.
	Action	Replace physical disk and rebuild it.
	Address	ctl: 0 chn: 13 tgt: 14
	Severity	1 (Serious)
	Source	InPro
41	Name	MLXEV_PHYSDEV_BSYPAR_DEAD

	Description	A physical disk failed because of busy status or parity error.
	Cause	The physical disk returned busy status. The SCSI transaction with the physical disk met with parity error.
	Action	Replace physical disk and rebuild it.
	Address	ctl: 0 chn: 13 tgt: 14
	Severity	1 (Serious)
	Source	InPro
42	Name	MLXEV_PHYSDEV_BYCMD_DEAD
	Description	A physical disk set to failed state by host.
	Cause	Command from host set the physical disk to failed state.
	Action	Replace physical disk and rebuild it.
	Address	ctl: 0 chn: 13 tgt: 14
	Severity	1 (Serious)
	Source	InPro
43	Name	MLXEV_PHYSDEV_SELTMO_DEAD
	Description	A physical disk failed because access to the device met with a selection time out.
	Cause	Physical disk disconnected or powered off. Physical disk failure.
	Action	Replace physical disk and rebuild it. Check power and cabling.
	Address	ctl: 0 chn: 13 tgt: 14
	Severity	1 (Serious)
	Source	InPro
44	Name	MLXEV_PHYSDEV_SEQERR_DEAD
	Description	A physical disk failed because of a sequence error in the SCSI bus phase handling.
	Cause	Physical disk failure.
	Action	Replace physical disk and rebuild it.
	Address	ctl: 0 chn: 13 tgt: 14
	Severity	1 (Serious)

	Source	InPro
45	Name	MLXEV_PHYSDEV_UNKNOWNSTS_DEAD
	Description	A physical disk failed because device returned an unknown status.
	Cause	Bad physical disk or incompatible device.
	Action	Replace physical disk or the device and rebuild it.
	Address	ctl: 0 chn: 13 tgt: 14
	Severity	1 (Serious)
	Source	InPro
46	Name	MLXEV_PHYSDEV_NOTRDY_DEAD
	Description	A physical disk failed because device is not ready.
	Cause	Physical disk not spinning, just turned bad. Power to the physical disk failed.
	Action	Replace physical disk and rebuild it. Check power and rebuild device.
	Address	ctl: 0 chn: 13 tgt: 14
	Severity	1 (Serious)
	Source	InPro
47	Name	MLXEV_PHYSDEV_MISSING_DEAD
	Description	A physical disk failed because the physical disk was not found on start up.
	Cause	Physical disk not connected. Physical disk not responding. Clear configuration suspend mode command was invoked.
	Action	Check setup. Check the startup option parameters on the system.
	Address	ctl: 0 chn: 13 tgt: 14
	Severity	1 (Serious)
	Source	InPro
48	Name	MLXEV_PHYSDEV_CODWRFAIL_DEAD
	Description	A physical disk failed because write operation of the 'Configuration On Disk' failed.

	Cause	Failed physical disk. Device write protected.
	Action	Replace physical disk and rebuild it. Check the startup option parameters on the system.
	Address	ctl: 0 chn: 13 tgt: 14
	Severity	1 (Serious)
	Source	InPro
49	Name	MLXEV_PHYSDEV_BDTWRFAIL_DEAD
	Description	A physical disk failed because write operation of 'Bad Data Table' failed.
	Cause	Failed physical disk. Device write protected.
	Action	Replace physical disk and rebuild it. Check the startup option parameters on the system.
	Address	ctl: 0 chn: 13 tgt: 14
	Severity	1 (Serious)
	Source	InPro
50	Name	MLXEV_PHYSDEV_OFFLINE
	Description	Physical disk status changed to offline.
	Cause	–
	Action	None
	Address	ctl: 0 chn: 13 tgt: 14
	Severity	2 (Error)
	Source	ExPro
51	Name	MLXEV_PHYSDEV_STANDBY
	Description	Physical disk status changed to Hot Spare.
	Cause	–
	Action	None
	Address	ctl: 0 chn: 13 tgt: 14
	Severity	4 (Information)
	Source	ExPro

52	Name	MLXEV_PHYSDEV_REBUILD
	Description	Physical disk status changed to rebuild.
	Cause	–
	Action	None
	Address	ctl: 0 chn: 13 tgt: 14
	Severity	2 (Error)
	Source	ExPro
53	Name	MLXEV_PHYSDEV_ID_MISMATCH
	Description	Physical device ID did not match.
	Cause	–
	Action	None
	Address	ctl: 0 chn: 13 tgt: 14
	Severity	3 (Warning)
	Source	ExP_old
54	Name	MLXEV_PHYSDEV_FAILED_START
	Description	Physical disk failed to start.
	Cause	–
	Action	Reseat the physical disk. Replace the physical disk.
	Address	ctl: 0 chn: 13 tgt: 14
	Severity	2 (Error)
	Source	ExPro
55	Name	MLXEV_PHYSDEV_OFFSET_SET
	Description	Physical disk negotiated different offset than configured.
	Cause	–
	Action	None
	Address	ctl: 0 chn: 13 tgt: 14
	Severity	3 (Warning)
	Source	ExP_old

56	Name	MLXEV_PHYSDEV_SET_BUS_WIDTH
	Description	Physical disk negotiated different bus width than configured.
	Cause	–
	Action	None
	Address	ctl: 0 chn: 13 tgt: 14
	Severity	3 (Warning)
	Source	ExP_old
57	Name	MLXEV_PHYSDEV_MISSING_ONSTARTUP
	Description	Physical disk missing on startup.
	Cause	Physical disk missing.
	Action	Replace the physical disk or power-on all enclosures.
	Address	ctl: 0 chn: 13 tgt: 14
	Severity	1 (Serious)
	Source	ExP_old InPro
58	Name	MLXEV_PHYSDEV_REBUILD_START_FAILED
	Description	Rebuild startup failed due to lower physical disk capacity.
	Cause	Physical disk capacity not sufficient for rebuild.
	Action	Replace with a physical disk having sufficient capacity.
	Address	ctl: 0 chn: 13 tgt: 14
	Severity	2 (Error)
	Source	InPro
59	Name	MLXEV_PHYSDEV_MOVING_TO_OTHER_CHN
	Description	Physical disk is switching from one channel to the other channel.
	Cause	Physical disk removed or channel failed
	Action	Check Fibre Channel loop. Replace physical disk if necessary.
	Address	ctl: 0 chn: 13 tgt: 14
	Severity	3 (Warning)
	Source	ExPro

60	Name	MLXEV_PHYSDEV_OFFLINE_DEVICE_MADE_ONLINE
	Description	Temporary-Dead physical disk is automatically made online.
	Cause	Temporary-Dead state caused because of transient errors
	Action	Analyze event log to find out why the physical disk was marked DEAD.
	Address	ctl: 0 chn: 13 tgt: 14
	Severity	2 (Error)
	Source	InPro
61	Name	MLXEV_PHYSDEV_STANDBY_REBUILD_START
	Description	A standby rebuild has started.
	Cause	A physical disk failed and a spare was available.
	Action	None
	Address	ctl: 0 chn: 13 tgt: 14
	Severity	4 (Information)
	Source	InPro
62	Name	MLXEV_PHYSDEV_HOT_SPARE_SMALLER
	Description	Hot spare replaced with a smaller capacity physical disk.
	Cause	The new hot spare may have a smaller physical capacity than the physical disk it replaced. The controller coercion setting may have reduced the configurable size of the new hot spare.
	Action	None
	Address	ctl: 0 chn: 13 tgt: 14
	Severity	4 (Information)
	Source	ExP_old
63	Name	MLXEV_PPILOT_LOGICAL_DISK_ONLINE
	Description	A storage access path has come online.
	Cause	PATHpilot reported an access path has become available. Storage previously unavailable has become available again.
	Action	None This indicates that storage access has been restored.

	Address	PATHpilot: Curr HBTL: 12, 13, 14, 15
	Severity	4 (Information)
	Source	PATHpilot
64	Name	MLXEV_PPILOT_LOGICAL_DISK_OFFLINE
	Description	A storage access path has been lost.
	Cause	PATHpilot is unable to access certain storage on any path. Storage previously available is now unavailable.
	Action	Investigate possible causes (such as cabling, controller failure, HBA failure).
	Address	PATHpilot: Last HBTL: 12, 13, 14, 15
	Severity	1 (Serious)
	Source	PATHpilot
65	Name	MLXEV_PPILOT_LOGICAL_DISK_PATH_FAILOVER
	Description	A storage access path has shifted to the alternate path.
	Cause	PATHpilot is unable to access storage on the original path, but is able to use the alternate path. A PATHpilot failover has occurred.
	Action	Investigate possible causes (such as cabling, controller failure, HBA failure).
	Address	PATHpilot: Previous HBTL: 12, 13, 14, 15 Current HBTL: 20, 21, 22, 23
	Severity	2 (Error)
	Source	PATHpilot
66	Name	MLXEV_PPILOT_LOGICAL_DISK_PATH_FAILBACK
	Description	A storage access path has shifted to the primary path.
	Cause	PATHpilot is again able to access certain storage on the original path. A PATHpilot failback has occurred.
	Action	None This indicates that the primary path has been restored.
	Address	PATHpilot: Previous HBTL: 12, 13, 14, 15 Current HBTL: 20, 21, 22, 23

	Severity	4 (Information)
	Source	PATHpilot
67	Name	MLXEV_PHYSDEV_NON_REDUNDANT_ACCESS
	Description	Physical disk found on only one disk channel.
	Cause	Physical disk is connected on only one channel.
	Action	Inspect disk channel cables and related hardware for proper operation.
	Address	ctl: 0 chn: 13 tgt: 14
	Severity	2 (Error)
	Source	ExPro
68	Name	MLXEV_PHYSDEV_TYPE_INVALID
	Description	Physical disk type is not approved by vendor.
	Cause	Physical disk does not match a type approved by the vendor. The system rejects all physical disks after the first that is not approved by the vendor.
	Action	Contact RAID controller vendor to acquire a supported physical disk.
	Address	ctl: 0 chn: 13 tgt: 14
	Severity	4 (Information)
	Source	ExPro
69	Name	MLXEV_PHYSDEV_SOFT_ID
	Description	Physical disk has acquired an inappropriate loop ID. Enclosure disk-slot operations are disabled while this condition persists.
	Cause	Enclosure selector switch conflict. Physical disk hardware failure. Enclosure disk slot hardware failure.
	Action	Ensure that each disk enclosure selector switch is set to a unique number per enclosure manufacturer specification. Inspect physical disk connector. Replace physical disk. Replace disk enclosure.
	Address	ctl: 0 chn: 13 tgt: 14
	Severity	2 (Error)
	Source	ExPro
70	Name	MLXEV_PHYSDEV_PORT_FAILED

	Description	Physical disk port has failed or cannot operate at the configured channel speed.
	Cause	Physical disk hardware failure. Physical disk is not compatible with system. Enclosure disk slot hardware failure.
	Action	Replace physical disk. Replace disk enclosure.
	Address	ctl: 0 chn: 13 tgt: 14
	Severity	2 (Error)
	Source	ExPro
71	Name	MLXEV_CTLDEV_MIRROR_RACE_ENTRY_FAILED
	Description	Mirror race recovery failed for logical drive.
	Cause	A read or write operation to a physical disk failed while restoring redundancy.
	Action	Run consistency check and restore consistency.
	Address	ctl: 0 logdrv: 15
	Severity	2 (Error)
	Source	ExPro
72	Name	MLXEV_CTLDEV_CTRL_PARAMS_CHECKSUM_FAILED
	Description	Controller parameters checksum verification failed; restored default.
	Cause	NVRAM battery low. NVRAM hardware failure. Improper shutdown of the controller during controller parameter update.
	Action	Restore correct controller parameter settings. If problem persists, replace controller.
	Address	ctl: 0
	Severity	2 (Error)
	Source	InPro
73	Name	MLXEV_FIRMWARE_UPGRADE_STARTED
	Description	Online controller firmware upgrade has started.
	Cause	User has initiated an online firmware upgrade.
	Action	None

	Address	ctl: 0
	Severity	4 (Information)
	Source	ExPro
74	Name	MLXEV_FIRMWARE_UPGRADE_COMPLETE
	Description	Online controller firmware upgrade has completed successfully.
	Cause	Online controller firmware upgrade has completed without error. The partner controller will now be auto flashed.
	Action	None
	Address	ctl: 0
	Severity	4 (Information)
	Source	ExPro
75	Name	MLXEV_FIRMWARE_UPGRADE_FAILED
	Description	Online controller firmware upgrade has failed.
	Cause	Online controller firmware upgrade has failed. The original firmware will be reloaded.
	Action	Use the offline method to load the new firmware.
	Address	ctl: 0
	Severity	2 (Error)
	Source	ExPro
76	Name	MLXEV_CFG_COD_DROPPED
	Description	A COD with unsupported features has been detected.
	Cause	Firmware does not support certain features in that COD. COD import is to wrong system.
	Action	COD data is sequestered. Obtain compatible firmware then reimport the COD or import the COD to a different system.
	Address	ctl: 0
	Severity	4 (INFORMATION)
	Source	InPro
77	Name	MLXEV_CTLDEV_BBU_NEW_BATTERY
	Description	New battery found.

	Cause	Battery present in the controller could be a new or a replaced unit.
	Action	None.
	Address	ctl: 0
	Severity	4 (INFORMATION)
	Source	InPro
78	Name	MLXEV_CTLDEV_BBU_CHARGER_ERROR
	Description	Battery backup unit charger error.
	Cause	Battery backup unit failure. Hardware problems in battery backup unit.
	Action	Replace the battery. Replace the battery backup unit hardware. If the problem persists contact your service representative.
	Address	ctl: 0
	Severity	1 (Serious)
	Source	InPro
79	Name	MLXEV_CTLDEV_BBU_BATTERY_ERROR
	Description	Battery will not hold a charge.
	Cause	Battery failure.
	Action	Replace the battery. If the problem persists contact your service representative.
	Address	ctl: 0
	Severity	1 (Serious)
	Source	InPro
80	Name	MLXEV_FW_ENTERED_UNEXPECTED_STATE
	Description	Firmware entered unexpected state at run-time.
	Cause	Memory corruption. Hardware returned unexpected value.
	Action	If the controller stops responding reboot and try again. If problem persists contact your service representative.
	Address	ctl: 0 Version: 13.14-15 Param: 0x23222120
	Severity	1 (Serious)
	Source	InPro

	Paramtype	enum
	Params	0x90000001 : name_CHARGER_MANAGER : "Charger Manager received unexpected state value" 0x90000002 : name_LOAD_POWER_GAUGE : "Load Power Gauge State machine received unexpected state value" 0x90000003 : name_READ_CHARGE_LEVEL : "Read Charge Level State machine received unexpected state value" 0x90000004 : name_READ_GAS_GAUGE : "Read Gas Gauge received unexpected state"
81	Name	MLXEV_SYSDEV_REBUILD_FAILED_ON_PARTNER
	Description	Rebuild stopped on controller failure.
	Cause	Rebuild moved to the surviving controller because the partner, which was running the rebuild, failed.
	Action	None.
	Address	ctl: 0 logdrv: 15
	Severity	4 (Information)
	Source	ExPro
82	Name	MLXEV_SYSDEV_CHECK_CONSISTENCY_FAILED_ON_ PARTNER
	Description	Check Consistency stopped on controller failure.
	Cause	Check Consistency failed because the partner, which was running the Check Consistency, failed.
	Action	Restart the Check Consistency.
	Address	ctl: 0 logdrv: 15
	Severity	2 (Error)
	Source	ExPro
83	Name	MLXEV_SYSDEV_FOREGROUND_INIT_FAILED_ON_ PARTNER
	Description	Foreground initialization stopped on controller failure.
	Cause	Foreground initialization failed because the partner, which was running the foreground initialization, failed.
	Action	Restart the foreground initialization.
	Address	ctl: 0 logdrv: 15
	Severity	2 (Error)

	Source	ExPro
84	Name	MLXEV_SYSDEV_BACKGROUND_INIT_FAILED_ON_ PARTNER
	Description	Background initialization stopped on controller failure.
	Cause	Background initialization restarted on the surviving controller because the partner, which was running the background initialization, failed.
	Action	None.
	Address	ctl: 0 logdrv: 15
	Severity	4 (Information)
	Source	ExPro
85	Name	MLXEV_CFG_COD_GROUP_NOT_INSTALLED_PDD_ USED
	Description	Configuration group not installed due to physical disk conflict.
	Cause	Physical disk already configured in another group. Unable to map a physical device in this group.
	Action	Move to another controller. Remove one or more existing configuration groups and reboot with this group.
	Address	None.
	Severity	4 (Information)
	Source	ExPro
86	Name	MLXEV_CFG_COD_GROUP_NOT_INSTALLED_LDD_ USED
	Description	Configuration group not installed due to logical drive conflict.
	Cause	Logical drive already configured in another group. Unable to install all the logical drives in the group.
	Action	Move to another controller. Remove one or more existing configuration groups and reboot with this group.
	Address	None.
	Severity	4 (Information)
	Source	ExPro
87	Name	MLXEV_CFG_COD_GROUP_NOT_INSTALLED_SAN_MAP_CON FLICT

	Description	Configuration group not installed due to SAN map conflict.
	Cause	SAN map conflicted with that of another installed group.
	Action	Move to another controller. Remove one or more existing configuration groups and reboot with this group.
	Address	None.
	Severity	4 (Information)
	Source	ExPro
88	Name	MLXEV_CFG_COD_GROUP_NOT_INSTALLED_FOREIGN
	Description	Foreign configuration group not installed.
	Cause	Foreign COD group found but not installed.
	Action	Move to another controller Remove native configuration group and reboot with foreign group.
	Address	None.
	Severity	4 (Information)
	Source	ExPro
89	Name	MLXEV_CFG_COD_GROUP_NOT_INSTALLED_ALL_BAD_FOREIGN
	Description	No configuration group installed.
	Cause	All configuration groups are foreign and clash. Controllers may have changed.
	Action	Move to another controller Re-insert original controllers or remove some groups and reboot with one group.
	Address	None.
	Severity	4 (Information)
	Source	ExPro
90	Name	MLXEV_PPILOT_SECONDARY_DISK_NOT_FOUND
	Description	A secondary logical drive was not found during initialization.
	Cause	PATHPilot reported a secondary logical drive was lost during power cycle.
	Action	Investigate possible causes such as cabling, controller failure, or HBA failure.

	Address	PATHpilot: Last HBTL: 12,13,14,15
	Severity	4 (Information)
	Source	PATHpilot
91	Name	MLXEV_PPILOT_PRIMARY_DISK_NOT_FOUND
	Description	A primary logical drive was not found during initialization.
	Cause	PATHPilot reported a primary logical drive was lost during power cycle.
	Action	Investigate possible causes such as cabling, controller failure, or HBA failure.
	Address	PATHpilot: Last HBTL: 12,13,14,15
	Severity	3 (Warning)
	Source	PATHpilot
92	Name	MLXEV_PPILOT_NO_PATH_FOR_LD_FOUND
	Description	No paths were found for a logical drive during initialization.
	Cause	PathPilot reported all paths were lost to a logical drive during power cycle.
	Action	Investigate possible causes such as cabling, controller failure, or HBA failure.
	Address	PATHpilot: Last HBTL: 12,13,14,15
	Severity	3 (Warning)
	Source	PATHpilot
93	Name	MLXEV_PPILOT_RECOVERY_ON_STANDBY_PATH
	Description	A standby path has been recovered.
	Cause	PathPilot reported recovery of a standby path.
	Action	None. This indicates that storage access has been restored.
	Address	PATHpilot: Curr HBTL: 12,13,14,15
	Severity	4 (Information)
	Source	PATHpilot
94	Name	MLXEV_PPILOT_PATH_FAILURE_ON_STANDBY_PATH
	Description	A standby path failed.
	Cause	PathPilot reported a failure on a standby path.

	Action	Investigate possible causes such as cabling, controller failure, or HBA failure.
	Address	PATHpilot: Curr HBTL: 12,13,14,15
	Severity	3 (Warning)
	Source	PATHpilot
95	Name	MLXEV_REPLACED_PHYSDEV_SMALLER
	Description	User replaced a configured physical disk with a disk of smaller capacity.
	Cause	The new physical disk may have a smaller physical capacity than the physical disk it replaced. The controller's coercion setting may have reduced the configurable size of the new physical disk.
	Action	None.
	Address	ctl: 0 chn: 13 tgt: 14
	Severity	4 (Information)
	Source	InPro
96	Name	MLXEV_FIBREDEV_LOOPID_SOFTADDR_OCCURRED
	Description	Device loop ID conflict (soft addressing) detected.
	Cause	Device loop ID conflict detected on a disk channel resulted in soft addressing; potential data corruption.
	Action	Change index selector to enable hard addressing per enclosure manufacturer’s specification.
	Address	ctl: 0 chn: 13
	Severity	1 (Serious)
	Source	ExPro
97	Name	MLXEV_PHYSDEV_TRANSPORT_ERROR
	Description	One or more transport errors occurred during data transfer to or from a physical device in the last hour.
	Cause	Fibre Channel CRC errors were detected by the controller or a physical device during the last hour. Physical devices were installed or removed causing momentary Fibre Channel loop disruption. Physical devices are not properly seated in the enclosure. Fibre Channel cabling or associated Fibre Channel hardware was installed or removed. Fibre Channel cabling or associated Fibre Channel hardware is faulty. Enclosure, controller, or physical device is faulty.

	Action	Unless the problem occurs frequently, action is not required. Refer to the enclosure manufacturer's service manual.
	Address	ctl: 0 chn: 13 tgt: 14 param: 20
	Severity	4 (Information)
	Source	ExPro
	Paramtype	Value
98	Name	MLXEV_HOST_FIBRE_DOWN
	Description	The host Fibre Channel is down.
	Cause	Host Fibre Channel cable is disconnected or has failed. Host Fibre Channel device (server, switch, hub, host adapter) has failed. Host Fibre Channel interface on the controller has failed.
	Action	Inspect host Fibre Channel cables and associated hardware.
	Address	ctl: 0 chn: 13
	Severity	4 (Information)
	Source	ExPro
99	Name	MLXEV_HOST_FIBRE_UP
	Description	The host Fibre Channel is up.
	Cause	The host Fibre Channel loop initialized successfully.
	Action	None
	Address	ctl: 0 chn: 13
	Severity	4 (Information)
	Source	ExPro
100	Name	MLXEV_PHYSDEV_REDUNDANT_ACCESS_RESTORED
	Description	Redundant access to the physical disk has been restored.
	Cause	The physical disk is now accessible on both disk channels.
	Action	If a physical disk frequently looses redundant accessibility check I/O modules and back-end connections.
	Address	ctl: 0 chn: 1 tgt: 2
	Severity	4 (Information)
	Source	ExPro
101	Name	MLXEV_SYSDEV_PFA_START

	Description	Started to loop through each physical disk to perform predicted failure analysis.
	Cause	Controller started to perform predicted failure analysis on the physical disks.
	Action	None
	Address	ctl: 0
	Severity	4 (Information)
	Source	ExPro
102	Name	MLXEV_SYSDEV_PFA_STOP
	Description	Finished predicted failure analysis on every physical disk.
	Cause	Controller finished predicted failure analysis on the physical disks.
	Action	None
	Address	ctl: 0
	Severity	4 (Information)
	Source	ExPro
103	Name	MLXEV_SYSDEV_BKGND_PATROL_START
	Description	Background patrol started or was resumed.
	Cause	Background patrol started.
	Action	None
	Address	ctl: 0
	Severity	4 (Information)
	Source	ExPro
104	Name	MLXEV_SYSDEV_BKGND_PATROL_STOP
	Description	Background patrol has found no more physical disks or was commanded to stop.
	Cause	Background patrol finished or was stopped.
	Action	None
	Address	ctl: 0
	Severity	4 (Information)
	Source	ExPro

128	Name	MLXEV_SYSDEV_CHECK_START
	Description	Consistency check has started.
	Cause	User started a consistency check. 'raidbld' started a consistency check.
	Action	None
	Address	ctl: 0 logdrv: 15
	Severity	4 (Information)
	Source	ExPro InPro
129	Name	MLXEV_SYSDEV_CHECK_DONE
	Description	Consistency check has finished.
	Cause	Consistency check completed successfully without detecting any errors.
	Action	None
	Address	ctl: 0 logdrv: 15
	Severity	4 (Information)
	Source	ExPro InPro
130	Name	MLXEV_SYSDEV_CHECK_CANCELED
	Description	Consistency check is cancelled.
	Cause	User cancelled the consistency check.
	Action	Restart consistency check, if required.
	Address	ctl: 0 logdrv: 15
	Severity	2 (Error)
	Source	ExPro InPro
131	Name	MLXEV_SYSDEV_CHECK_ERROR
	Description	Consistency check on logical drive error.
	Cause	Inconsistent data were found. Bad sectors were found. A physical disk reliability problem.
	Action	See bad block and request sense table for more information.
	Address	ctl: 0 logdrv: 15
	Severity	2 (Error)

	Source	ExPro InPro
132	Name	MLXEV_SYSDEV_CHECK_SYSDEV_FAILED
	Description	Consistency check on logical drive failed.
	Cause	A logical device became critical. A logical device failed.
	Action	See request sense data for more information.
	Address	ctl: 0 logdrv: 15
	Severity	2 (Error)
	Source	ExPro InPro
133	Name	MLXEV_SYSDEV_CHECK_PHYSDEV_FAILED
	Description	Consistency check failed due to physical disk failure.
	Cause	A physical disk failed.
	Action	See request sense data for more information.
	Address	ctl: 0 logdrv: 15
	Severity	2 (Error)
	Source	ExPro InPro
134	Name	MLXEV_SYSDEV_OFFLINE
	Description	Logical drive has been made offline.
	Cause	One/multiple physical disk(s) failed.
	Action	It may not be possible to recover from this error. Contact your service representative.
	Address	ctl: 0 logdrv: 15
	Severity	1 (Serious)
	Source	ExPro InPro
135	Name	MLXEV_SYSDEV_CRITICAL
	Description	Logical drive is critical.
	Cause	One physical device failed.
	Action	Replace the physical device. Start the rebuild, if required.
	Address	ctl: 0 logdrv: 15

	Severity	2 (Error)
	Source	ExPro InPro
136	Name	MLXEV_SYSDEV_CRITICAL
	Description	Logical drive has been placed online.
	Cause	Rebuild completed. User set the physical disk online. New configuration was added.
	Action	None
	Address	ctl: 0 logdrv: 15
	Severity	4 (Information)
	Source	ExPro InPro
137	Name	MLXEV_SYSDEV_AUTO_REBUILD_START
	Description	An automatic rebuild has started on logical drive.
	Cause	A physical disk failed and a spare device was available. A spare physical disk was found and replaced the failed disk.
	Action	None
	Address	ctl: 0 logdrv: 15
	Severity	4 (Information)
	Source	InPro
138	Name	MLXEV_SYSDEV_MANUAL_REBUILD_START
	Description	A manual rebuild has started on logical drive.
	Cause	Client started the rebuild on user's request. User replaced the failed device and 'raidbld' started the rebuild.
	Action	None
	Address	ctl: 0 logdrv: 15
	Severity	4 (Information)
	Source	InPro
139	Name	MLXEV_SYSDEV_REBUILD_DONE
	Description	Rebuild on logical drive is over.
	Cause	Rebuild completed successfully only for this logical drive.
	Action	None

	Address	ctl: 0 logdrv: 15
	Severity	4 (Information)
	Source	InPro
140	Name	MLXEV_SYSDEV_REBUILD_CANCELED
	Description	Rebuild on logical drive has been cancelled.
	Cause	User cancelled rebuild. Higher priority rebuild started.
	Action	Restart the rebuild if required.
	Address	ctl: 0 logdrv: 15
	Severity	2 (Error)
	Source	InPro
141	Name	MLXEV_SYSDEV_REBUILD_ERROR
	Description	Rebuild stopped with error.
	Cause	Due to an unknown error on the controller rebuild failed.
	Action	Try rebuild again.
	Address	ctl: 0 logdrv: 15
	Severity	2 (Error)
	Source	InPro
142	Name	MLXEV_SYSDEV_REBUILD_NEWDEV_FAILED
	Description	Rebuild stopped with error. New physical disk failed.
	Cause	New physical disk failed. New physical disk is not compatible with MDAC hardware/firmware.
	Action	Replace the new disk.
	Address	ctl: 0 logdrv: 15
	Severity	2 (Error)
	Source	InPro
143	Name	MLXEV_SYSDEV_REBUILD_SYSDEV_FAILED
	Description	Rebuild stopped because logical drive failed.
	Cause	At least one more physical device failed in the array.

	Action	It may not be possible to recover from this error. Contact your service representative.
	Address	ctl: 0 logdrv: 15
	Severity	2 (Error)
	Source	InPro
144	Name	MLXEV_SYSDEV_INIT_STARTED
	Description	Logical drive initialization started.
	Cause	User started the initialization. Any previous data are lost.
	Action	None
	Address	ctl: 0 logdrv: 15
	Severity	4 (Information)
	Source	ExPro InPro
145	Name	MLXEV_SYSDEV_INIT_DONE
	Description	Logical drive initialization finished.
	Cause	Initialize operation completed successfully.
	Action	None
	Address	ctl: 0 logdrv: 15
	Severity	4 (Information)
	Source	ExPro InPro
146	Name	MLXEV_SYSDEV_INIT_CANCELED
	Description	Logical drive initialization cancelled.
	Cause	User cancelled the initialization.
	Action	Restart initialization if required.
	Address	ctl: 0 logdrv: 15
	Severity	2 (Error)
	Source	ExPro InPro
147	Name	MLXEV_SYSDEV_INIT_FAILED
	Description	Logical drive initialization failed.

	Cause	One/multiple physical device(s) failed. Controller has been removed. Controller has been powered off.
	Action	Refer to the device failure event.
	Address	ctl: 0 logdrv: 15
	Severity	2 (Error)
	Source	ExPro InPro
148	Name	MLXEV_SYSDEV_FOUND
	Description	A logical drive has been found.
	Cause	A new configuration has been added. MORE completed. A new controller has been plugged in. Controller has been powered on. System has rebooted.
	Action	None
	Address	ctl: 0 logdrv: 15
	Severity	4 (Information)
	Source	InPro
149	Name	MLXEV_SYSDEV_GONE
	Description	A logical drive has been deleted.
	Cause	A new configuration has been added. A new logical device has been deleted. Controller has been removed. Controller has been powered off.
	Action	None
	Address	ctl: 0 logdrv: 15
	Severity	4 (Information)
	Source	InPro
150	Name	MLXEV_SYSDEV_EXPANDCAPACITY_START
	Description	Expand capacity started.
	Cause	User started the online RAID expansion operation.
	Action	None
	Address	ctl: 0 logdrv: 15

	Severity	4 (Information)
	Source	ExPro InPro
151	Name	MLXEV_SYSDEV_EXPANDCAPACITY_DONE
	Description	Expand capacity completed.
	Cause	Online RAID expansion completed.
	Action	None
	Address	ctl: 0 logdrv: 15
	Severity	4 (Information)
	Source	ExPro InPro
152	Name	MLXEV_SYSDEV_EXPANDCAPACITY_ERROR
	Description	Expand capacity stopped with error.
	Cause	Multiple physical disks failed.
	Action	It may not be possible to recover from this error. Contact your service representative.
	Address	ctl: 0 logdrv: 15
	Severity	2 (Error)
	Source	ExPro InPro
153	Name	MLXEV_SYSDEV_BADBLOCK
	Description	Bad blocks found.
	Cause	Bad sector was found on a physical device during consistency check/rebuild/RAID expansion operation.
	Action	Run a consistency check with the restore option. Restore data from a backup.
	Address	ctl: 0 logdrv: 15
	Severity	2 (Error)
	Source	ExPro InP_old
154	Name	MLXEV_SYSDEV_SIZECHANGED
	Description	Logical drive size changed.
	Cause	A new configuration has been added. RAID expansion has added extra capacity.
	Action	None

	Address	ctl: 0 logdrv: 15
	Severity	4 (Information)
	Source	InP_old
155	Name	MLXEV_SYSDEV_TYPECHANGED
	Description	Logical drive type changed.
	Cause	A new configuration has been added. RAID migration completed. RAID expansion completed on RAID 1.
	Action	None
	Address	ctl: 0 logdrv: 15
	Severity	4 (Information)
	Source	ExPro InP_old
156	Name	MLXEV_SYSDEV_BADDATABLOCK
	Description	Bad data blocks found. Possible data loss.
	Cause	Bad blocks were found on multiple physical disks in same zone.
	Action	Restore data from a back up.
	Address	ctl: 0 logdrv: 15
	Severity	2 (Error)
	Source	InPro
157	Name	MLXEV_SYSDEV_WR_LUN_MAP
	Description	System drive LUN mapping has been written to configuration.
	Cause	–
	Action	None
	Address	ctl: 0
	Severity	4 (Information)
	Source	ExPro
158	Name	MLXEV_SYSDEV_DATAREAD_FROM_BLOCK_IN_BDT
	Description	Attempt to read data from block that is marked in Bad Data table.
	Cause	Attempt to read from block that is already marked bad in Bad Data table. Potential data loss.

	Action	Restore data from a back up.
	Address	ctl: 0 logdrv: 15
	Severity	1 (Serious)
	Source	InP_old
159	Name	MLXEV_SYSDEV_DATA_FOR_BLOCK_LOST
	Description	Data for disk block has been lost due to logical drive problem.
	Cause	Data retained in RAID cache for a Write-Back Logical Drive cannot be stored to the physical medium because of logical drive problem. The logical drive problem might be caused by multiple physical medium errors, multiple physical devices offline, or other reasons.
	Action	Ensure that all the physical disks related to the logical drive, disk channel, enclosure, or cabling are functional and accessible. Repair or replace them if necessary.
	Address	ctl: 0 logdrv: 15
	Severity	2 (Error)
	Source	InP_old
160	Name	MLXEV_SYSDEV_OFFLINE_DEVICE_MADE_AVAILABLE_WITH_DATALOSS
	Description	Temporary-Offline RAID5/RAID3 array is available to the user again with the possibility of data loss in the array.
	Cause	Temporary-Offline state caused by transient errors on physical disks.
	Action	Verify data from backup.
	Address	ctl: 0 logdrv: 15
	Severity	2 (Error)
	Source	InPro
161	Name	MLXEV_SYSDEV_OFFLINE_DEVICE_MADE_AVAILABLE
	Description	Temporary-Offline RAID0+1/RAID1/RAID0/JBOD array is available to the user again.
	Cause	Temporary-Offline state caused by transient errors on physical disks.
	Action	None
	Address	ctl: 0 logdrv: 15
	Severity	2 (Error)

	Source	InPro
162	Name	MLXEV_SYSDEV_STANDBY_REBUILD_START
	Description	A standby rebuild has started on logical drive.
	Cause	A physical disk failed and a spare disk was available.
	Action	None
	Address	ctl: 0 logdrv: 15
	Severity	4 (Information)
	Source	InPro
176	Name	MLXEV_SYSDEV_BG_INIT_STARTED
	Description	Logical drive background initialization started.
	Cause	User may have started background initialization. Firmware may have automatically started background initialization.
	Action	None
	Address	ctl: 0 logdrv: 15
	Severity	4 (Information)
	Source	ExPro
177	Name	MLXEV_SYSDEV_BG_INIT_STOPPED
	Description	Logical drive background initialization stopped.
	Cause	User may have stopped background initialization. Firmware may have automatically stopped background initialization.
	Action	None
	Address	ctl: 0 logdrv: 15
	Severity	4 (Information)
	Source	ExPro
178	Name	MLXEV_SYSDEV_BG_INIT_PAUSED
	Description	Logical drive background initialization paused.
	Cause	Background initialization paused due to a higher priority operation.
	Action	None
	Address	ctl: 0 logdrv: 15

	Severity	4 (Information)
	Source	ExPro
179	Name	MLXEV_SYSDEV_BG_INIT_RESTARTED
	Description	Logical drive background initialization restarted.
	Cause	Background initialization started after being paused.
	Action	None
	Address	ctl: 0 logdrv: 15
	Severity	4 (Information)
	Source	ExPro
180	Name	MLXEV_SYSDEV_BG_INIT_FAILED
	Description	Logical drive background initialization failed.
	Cause	Background initialization failed.
	Action	None
	Address	ctl: 0 logdrv: 15
	Severity	2 (Error)
	Source	ExPro
181	Name	MLXEV_SYSDEV_BG_INIT_COMPLETED
	Description	Logical drive background initialization completed.
	Cause	Background initialization completed successfully.
	Action	None
	Address	ctl: 0 logdrv: 15
	Severity	4 (Information)
	Source	ExPro
182	Name	MLXEV_SYSDEV_DATA_LOSS_LOW_BBU_CHARGE
	Description	Low battery charge level. Logical drive may have lost data.
	Cause	Controller was powered off for duration longer than battery capacity. User connected a new controller. User connected a new BBU battery.
	Action	Run consistency check to verify logical drive consistency. If needed, restore data from backup.

	Address	ctl: 0 logdrv: 15
	Severity	2 (Error)
	Source	InPro
256	Name	MLXEV_FMTFAN_FAILED
	Description	Fan failure.
	Cause	Cable connection broken. Fan failure.
	Action	Replace fan.
	Address	ctl: 0 enclosure: 13 unit: 14
	Severity	1 (Serious)
	Source	ExP_old InP_old
257	Name	MLXEV_FMTFAN_OK
	Description	Fan has been restored.
	Cause	Cable is connected properly. Faulty fan has been replaced.
	Action	None
	Address	ctl: 0 enclosure: 13 unit: 14
	Severity	4 (Information)
	Source	ExP_old InP_old
258	Name	MLXEV_AEMI_FAN_FAILED
	Description	Fan failure.
	Cause	Cable connection is broken. Fan failure.
	Action	Replace fan.
	Address	ctl: 0 chn: 13 tgt: 14 Fan: 15
	Severity	1 (Serious)
	Source	InP_old
259	Name	MLXEV_FMTFAN_NOTPRESENT
	Description	Storage cabinet fan is not present.

	Cause	Enclosure management connection is broken. Management hardware failure. Fan is not present.
	Action	Refer to the enclosure manufacturer's service manual.
	Address	ctl: 0 enclosure: 13
	Severity	4 (Information)
	Source	ExP_old InP_old
272	Name	MLXEV_FMTPOWER_FAILED
	Description	Power supply failure.
	Cause	Broken cable connection. Power supply failure.
	Action	Reconnect cable or replace the power supply as required.
	Address	ctl: 0 enclosure: 13 unit: 14
	Severity	1 (Serious)
	Source	ExP_old InP_old
273	Name	MLXEV_FMTPOWER_OK
	Description	Power supply has been restored.
	Cause	Faulty power supply has been replaced.
	Action	None
	Address	ctl: 0 enclosure: 13 unit: 14
	Severity	4 (Information)
	Source	ExP_old InP_old
274	Name	MLXEV_AEMI_PWR_SUPPLY_FAILED
	Description	Power supply failure.
	Cause	Broken cable connection. Power supply failure.
	Action	Reconnect cable or replace the power supply as required.
	Address	ctl: 0 chn: 13 tgt: 14 Power Supply: 15
	Severity	1 (Serious)
	Source	InP_old
275	Name	MLXEV_FMTPOWER_NOTPRESENT

	Description	Storage cabinet power supply is not present.
	Cause	Management connection is broken. Management hardware failure. Power supply is not present.
	Action	Refer to the enclosure manufacturer's service manual.
	Address	ctl: 0 enclosure: 13
	Severity	4 (Information)
	Source	ExP_old InP_old
288	Name	MLXEV_FMTHEAT_BAD
	Description	Over temperature. Temperature is above 70 °C.
	Cause	Room temperature is too high. Fan failure. Sensor failure.
	Action	Turn off the system and let it to cool down. Adjust the room temperature.
	Address	ctl: 0 chn: 13 tgt: 14 Heat Sensor: 15
	Severity	1 (Serious)
	Source	InP_old
289	Name	MLXEV_FMTHEAT_CRITICAL
	Description	Temperature is above 50 °C.
	Cause	Room temperature is high. Fan failure.
	Action	Replace fan. Turn off the system. Adjust the room temperature.
	Address	ctl: 0 enclosure: 13 unit: 14
	Severity	3 (Warning)
	Source	ExP_old InP_old
290	Name	MLXEV_FMTHEAT_OK
	Description	Normal temperature has been restored.
	Cause	Faulty fan has been replaced. Room temperature was reduced.
	Action	None

	Address	ctl: 0 enclosure: 13 unit: 14
	Severity	4 (Information)
	Source	ExPro InP_old
291	Name	MLXEV_AEMI_OVER_TEMPERATURE
	Description	Over temperature.
	Cause	Room temperature is too high. Fan failure.
	Action	Turn off the system and let it to cool down. Adjust the room temperature.
	Address	ctl: 0 chn: 13 tgt: 14 Heat Sensor: 15
	Severity	1 (Serious)
	Source	InP_old
292	Name	MLXEV_FMTHEAT_NOTPRESENT
	Description	Storage cabinet temperature sensor is not present.
	Cause	Enclosure management connection is broken. Management hardware failure. Sensor is not present.
	Action	Refer to enclosure manufacturer's service manual.
	Address	ctl: 0 enclosure: 13
	Severity	4 (Information)
	Source	ExP_old InP_old
304	Name	MLXEV_FMTSTWK_FAILED
	Description	Storage Works enclosure reported failure state.
	Cause	Power supply failed. Fan failure. Cabinet is too hot.
	Action	Refer to enclosure manufacturer's service manual.
	Address	ctl: 0 chn: 13
	Severity	1 (Serious)
	Source	InPro
305	Name	MLXEV_FMTSTWK_CRITICAL
	Description	Storage Works enclosure reported critical state.

	Cause	Not available.
	Action	Not available.
	Address	ctl: 0 chn: 13
	Severity	2 (Error)
	Source	InPro
306	Name	MLXEV_FMTSTWK_OK
	Description	Storage Works enclosure reported normal state.
	Cause	Problem has been rectified.
	Action	None
	Address	ctl: 0 chn: 13
	Severity	4 (Information)
	Source	InPro
307	Name	MLXEV_FMT_UPS_DISABLED
	Description	Uninterruptible power supply disabled.
	Cause	–
	Action	Refer to enclosure manufacturer’s service manual.
	Address	ctl: 0
	Severity	4 (Information)
	Source	ExPro
308	Name	MLXEV_FMT_UPS_AC_FAILED
	Description	Uninterruptible power supply ac failed.
	Cause	–
	Action	Refer to enclosure manufacturer’s service manual.
	Address	ctl: 0
	Severity	4 (Information)
	Source	ExPro
309	Name	MLXEV_FMT_UPS_BAT_LOW
	Description	Uninterruptible power supply battery low.
	Cause	–

	Action	Refer to enclosure manufacturer's service manual.
	Address	ctl: 0
	Severity	2 (Error)
	Source	ExPro
310	Name	MLXEV_FMT_UPS_FAILED
	Description	Uninterruptible power supply failed.
	Cause	–
	Action	Refer to enclosure manufacturer's service manual.
	Address	ctl: 0
	Severity	1 (Serious)
	Source	ExPro
311	Name	MLXEV_FMT_UPS_OK
	Description	Uninterruptible power supply normal.
	Cause	–
	Action	None
	Address	ctl: 0
	Severity	4 (Information)
	Source	ExPro
320	Name	MLXEV_ENCLFAN_FAILED
	Description	Fan failure.
	Cause	Cable connection broken. Fan failure
	Action	Replace fan.
	Address	ctl: 0 enclosure: 13 unit: 14
	Severity	1 (Serious)
	Source	ExPro InPro
321	Name	MLXEV_ENCLFAN_OK
	Description	Fan has been restored.
	Cause	Faulty fan has been replaced. Cable is connected properly.

	Action	None
	Address	ctl: 0 enclosure: 13 unit: 14
	Severity	4 (Information)
	Source	ExPro InPro
322	Name	MLXEV_ENCLFAN_NOTPRESENT
	Description	Fan is not present.
	Cause	Enclosure management connection is broken. Management hardware failure. Fan is not present.
	Action	Refer to enclosure manufacturer's service manual.
	Address	ctl: 0 enclosure: 13 unit: 14
	Severity	2 (Error)
	Source	ExPro InPro
323	Name	MLXEV_ENCLPOWER_FAILED
	Description	Power supply failure.
	Cause	Broken cable connection. Power supply failure.
	Action	Refer to enclosure manufacturer's service manual.
	Address	ctl: 0 enclosure: 13 unit: 14
	Severity	1 (Serious)
	Source	ExPro InPro
324	Name	MLXEV_ENCLPOWER_OK
	Description	Power supply has been restored.
	Cause	Faulty power supply has been replaced.
	Action	None
	Address	ctl: 0 enclosure: 13 unit: 14
	Severity	4 (Information)
	Source	ExPro InPro
325	Name	MLXEV_ENCLPOWER_NOTPRESENT
	Description	Power supply is not present.

	Cause	Enclosure management connection is broken. Management hardware failure. Power supply is not present.
	Action	Refer to enclosure manufacturer's service manual.
	Address	ctl: 0 enclosure: 13 unit: 14
	Severity	4 (Information)
	Source	ExPro InPro
326	Name	MLXEV_ENCLHEAT_BAD
	Description	Temperature is over safe limit. Failure imminent.
	Cause	Room temperature is too high. Fan failure. Sensor failure.
	Action	Turn off the system and let it to cool down. Adjust the room temperature.
	Address	ctl: 0 enclosure: 13 unit: 14
	Severity	1 (Serious)
	Source	ExPro InPro
327	Name	MLXEV_ENCLHEAT_CRITICAL
	Description	Temperature is above working limit.
	Cause	Room temperature is too high. Fan failure.
	Action	Adjust the room temperature. Replace fan. Turn off the system.
	Address	ctl: 0 enclosure: 13 unit: 14
	Severity	2 (Error)
	Source	ExPro InPro
328	Name	MLXEV_ENCLHEAT_OK
	Description	Normal temperature has been restored.
	Cause	Faulty fan has been replaced. Room temperature was reduced.
	Action	None
	Address	ctl: 0 enclosure: 13 unit: 14

	Severity	4 (Information)
	Source	ExPro InPro
329	Name	MLXEV_ENCLHEAT_NOTPRESENT
	Description	Temperature sensor is not present.
	Cause	Enclosure management connection is broken. Management hardware failure. Sensor is not present.
	Action	Refer to enclosure manufacturer's service manual.
	Address	ctl: 0 enclosure: 13 unit: 14
	Severity	4 (Information)
	Source	ExPro InPro
330	Name	MLXEV_ENCLACCESS_CRITICAL
	Description	Enclosure access critical.
	Cause	Enclosure management connection is broken. Management hardware failure.
	Action	Refer to enclosure manufacturer's service manual.
	Address	ctl: 0 enclosure: 13
	Severity	2 (Error)
	Source	ExPro InPro
331	Name	MLXEV_ENCLACCESS_OK
	Description	Enclosure access has been restored.
	Cause	Enclosure has been fixed or replaced.
	Action	None
	Address	ctl: 0 enclosure: 13
	Severity	4 (Information)
	Source	ExPro InPro
332	Name	MLXEV_ENCLACCESS_OFFLINE
	Description	Enclosure access is offline.
	Cause	Enclosure management connection is broken. Management hardware failure.
	Action	Refer to enclosure manufacturer's service manual.

	Address	ctl: 0 enclosure: 13
	Severity	1 (Serious)
	Source	ExPro InPro
333	Name	MLXEV_ENCLSES_SOFTADDR_OCCURRED
	Description	Enclosure soft addressing detected.
	Cause	Enclosure has duplicate loop IDs (soft addressing). Potential data corruption.
	Action	Refer to enclosure manufacturer's service manual.
	Address	ctl: 0 enclosure: 13
	Severity	1 (Serious)
	Source	ExPro
334	Name	MLXEV_ENCLACCESS_READY
	Description	Enclosure services ready.
	Cause	–
	Action	None
	Address	ctl: 0
	Severity	4 (Information)
	Source	ExPro
335	Name	MLXEV_ENCLHEAT_UNKNOWN
	Description	Access to temperature sensor has been lost.
	Cause	Module containing temperature sensor has been removed.
	Action	None
	Address	ctl: 0 enclosure: 13 unit: 14
	Severity	4 (Information)
	Source	ExPro
336	Name	MLXEV_ENCLPOWER_UNKNOWN
	Description	Access to power supply status information has been lost.
	Cause	Module establishing connectivity has been removed.
	Action	None

	Address	ctl: 0 enclosure: 13 unit: 14
	Severity	4 (Information)
	Source	ExPro
337	Name	MLXEV_ENCLFAN_UNKNOWN
	Description	Access to fan status information has been lost.
	Cause	Module establishing connectivity has been removed.
	Action	None
	Address	ctl: 0 enclosure: 13 unit: 14
	Severity	4 (Information)
	Source	ExPro
338	Name	MLXEV_ENCLOSURE_SHUTDOWN
	Description	Physical disks in enclosure are being spundown.
	Cause	Enclosure temperature is too high. Fan failure. Sensor failure.
	Action	Turn off the enclosure and repair the problem causing the critical over-temperature condition. After the problem is repaired, recover by power cycling the entire system including the controllers.
	Address	ctl: 0 enclosure: 13 unit: 14
	Severity	1 (Serious)
	Source	ExPro
339	Name	MLXEV_ENCLCOLD_WARN
	Description	Temperature is below working limit.
	Cause	Room temperature is too low. Sensor failure.
	Action	Inspect the enclosure environment for proper operation and repair as required. Increase the room temperature as required.
	Address	ctl: 0 enclosure: 13
	Severity	2 (Error)
	Source	ExPro
340	Name	MLXEV_ENCLCOLD_FAIL

	Description	Temperature is under safe limit. Failure is imminent.
	Cause	Room temperature is too low. Sensor failure.
	Action	Inspect the enclosure environment for proper operation and repair as required. Increase the room temperature as required.
	Address	ctl: 0 enclosure: 13
	Severity	1 (Serious)
	Source	ExPro
384	Name	MLXEV_SYSTEM_STARTED
	Description	Array management server software started successfully.
	Cause	The server system (or array management utility server) started.
	Action	If you did not expect a system reboot, investigate.
	Address	None
	Severity	4 (Information)
	Source	ExPro InP_old
385	Name	MLXEV_CTLDEV_WRITEBACK_ERROR
	Description	Write back error.
	Cause	Data cache write failed.
	Action	The data may have been lost. Restore the data from a backup.
	Address	ctl: 0
	Severity	2 (Error)
	Source	InP_old
386	Name	MLXEV_CTLDEV_STATE_TABLE_FULL
	Description	Internal log structures getting full; PLEASE SHUTDOWN AND RESET THE SYSTEM IN THE NEAR FUTURE.
	Cause	Too many configuration changes occurred since the last boot.
	Action	Reboot the system by power cycling when ever convenient.
	Address	ctl: 0
	Severity	3 (Warning)

	Source	Unspecified
388	Name	MLXEV_CTLDEV_DEAD
	Description	Controller is dead. System is disconnecting from this controller.
	Cause	Hardware failure.
	Action	Contact your service representative.
	Address	ctl: 0
	Severity	0 (Critical)
	Source	InP_old
389	Name	MLXEV_CTLDEV_RESET
	Description	Controller has been reset.
	Cause	Controller received a reset command.
	Action	If this was an unexpected event, refer to the enclosure manufacturer's service manual.
	Address	ctl: 0
	Severity	4 (Information)
	Source	ExPro InP_old
390	Name	MLXEV_CTLDEV_FOUND
	Description	Controller is found.
	Cause	New controller has been installed. Management software has restarted. System has rebooted.
	Action	None
	Address	ctl: 0
	Severity	4 (Information)
	Source	InP_old
391	Name	MLXEV_CTLDEV_GONE
	Description	Controller is gone. System is disconnecting from this controller.
	Cause	Controller has been powered off. Controller has been removed from the system.
	Action	None

	Address	ctl: 0
	Severity	0 (Critical)
	Source	InP_old
392	Name	MLXEV_CTLDEV_BBU_FOUND
	Description	Battery present.
	Cause	A battery module has been detected and is now operational.
	Action	None
	Address	ctl: 0
	Severity	4 (Information)
	Source	ExPro InPro
393	Name	MLXEV_CTLDEV_BBU_POWER_LOW
	Description	Battery power low.
	Cause	Battery power is low.
	Action	If this message occurs without power failure, replace the battery.
	Address	ctl: 0
	Severity	2 (Error)
	Source	ExPro InPro
394	Name	MLXEV_CTLDEV_BBU_POWER_OK
	Description	Battery power OK.
	Cause	Battery has enough power to enable the write data cache.
	Action	None
	Address	ctl: 0
	Severity	4 (Information)
	Source	ExPro InPro
395	Name	MLXEV_CTLDEV_POWER_OFF
	Description	Controller is gone. System is disconnecting from this controller.
	Cause	The connection to the controller has been lost.
	Action	None
	Address	ctl: 0

	Severity	0 (Critical)
	Source	Server
396	Name	MLXEV_CTLDEV_POWER_ON
	Description	Controller powered on.
	Cause	Controller was removed from the system. Controller has been powered off.
	Action	None
	Address	ctl: 0
	Severity	4 (Information)
	Source	Server
397	Name	MLXEV_CTLDEV_ONLINE
	Description	Controller is online.
	Cause	New controller has been installed.
	Action	None
	Address	ctl: 0
	Severity	4 (Information)
	Source	Server
398	Name	MLXEV_CTLDEV_OFFLINE
	Description	Controller is gone. System is disconnecting from this controller.
	Cause	Controller was set offline.
	Action	If you did not expect this, investigate.
	Address	ctl: 0
	Severity	0 (Critical)
	Source	Server
399	Name	MLXEV_CTLDEV_CRITICAL
	Description	Controller's partner is gone. Controller is in failover mode now.
	Cause	Controller was set offline.
	Action	If you did not expect this, investigate.
	Address	ctl: 0

	Severity	3 (Warning)
	Source	Server
400	Name	MLXEV_CTLDEV_BBU_RECOND_START
	Description	Battery reconditioning was started.
	Cause	User started a battery reconditioning.
	Action	None
	Address	ctl: 0
	Severity	4 (Information)
	Source	ExPro
401	Name	MLXEV_CTLDEV_BBU_RECOND_DONE
	Description	Battery reconditioning is finished.
	Cause	Battery reconditioning completed successfully.
	Action	None
	Address	ctl: 0
	Severity	4 (Information)
	Source	ExPro
402	Name	MLXEV_CTLDEV_BBU_RECOND_ABORT
	Description	Battery reconditioning is canceled.
	Cause	User cancelled the battery reconditioning.
	Action	Restart the battery reconditioning, if required.
	Address	ctl: 0
	Severity	4 (Information)
	Source	ExPro
403	Name	MLXEV_CTLDEV_INSTALLATION_ABORT
	Description	Installation aborted.
	Cause	Equipment improperly set up while powered off.
	Action	Re-establish proper setup.
	Address	ctl: 0
	Severity	1 (Serious)

	Source	InPro
404	Name	MLXEV_CTLDEV_FIRMWARE_MISMATCH
	Description	Controller firmware mismatch.
	Cause	Replacement controller with downlevel firmware installed.
	Action	Reload controller firmware.
	Address	ctl: 0 chn: 13 tgt: 14
	Severity	1 (Serious)
	Source	Unspecified
405	Name	MLXEV_CTLDEV_BBU_NORESPONSE
	Description	Battery removed.
	Cause	Battery physically removed.
	Action	Reinstall the battery.
	Address	ctl: 0
	Severity	2 (Error)
	Source	ExPro InP_old
406	Name	MLXEV_CTLDEV_WARM_BOOT_ERROR
	Description	WARM boot failed.
	Cause	Memory error detected during WARM boot scan. Possible data loss.
	Action	Restore data from a backup.
	Address	ctl: 0
	Severity	1 (Serious)
	Source	ExPro
407	Name	MLXEV_CTLDEV_BBU_CALIBRATE_START
	Description	Battery calibration cycle started.
	Cause	New battery detected.
	Action	None
	Address	ctl: 0
	Severity	4 (Information)

	Source	ExPro
408	Name	MLXEV_CTLDEV_BBU_CALIBRATE_DONE
	Description	Battery calibration cycle finished.
	Cause	Battery calibration completed successfully.
	Action	None
	Address	ctl: 0
	Severity	4 (Information)
	Source	ExPro
409	Name	MLXEV_CTLDEV_BBU_CALIBRATE_ABORT
	Description	Battery calibration cycle was canceled.
	Cause	User canceled the battery calibration cycle.
	Action	None
	Address	ctl: 0
	Severity	4 (Information)
	Source	ExPro
410	Name	MLXEV_CTLDEV_BBU_NO_BATTERY
	Description	Battery is not present.
	Cause	The battery electronics are present, but a battery was not detected.
	Action	Install or connect the battery.
	Address	ctl: 0
	Severity	2 (Error)
	Source	ExPro
411	Name	MLXEV_CTLDEV_CONSERV_CACHE_MODE
	Description	Controller entered conservative cache mode.

	Cause	User changed the setting. Controller preparing for shutdown. Battery low. RAID expansion in process. Over-temperature condition. No redundant power. Power supply failure. UPS ac failed. UPS battery low.
	Action	If this was an unexpected event, refer to the enclosure manufacturer's service manual.
	Address	ctl: 0 param: 0x2120
	Severity	3 (Warning)
	Source	ExPro
	Paramtype	bitfield
	Params	1 : name_FAULT_MANAGEMENT : "Fault Management" 2 : name_INPUT_SIG : "Hardware Input Signal" 4 : name_SHUT_DOWN : "Controller Shut Down" 8 : name_BBU_LOW : "BBU Battery Low" 16 : name_MORE : "More Active" 32 : name_OVER_TEMP : "Over Temperature" 64 : name_NON_REDUN_POWER : "Non-redundant Power" 128 : name_POWER_SUPPLY_FAIL : "Power Supply Failed" 256 : name_UPS_AC_FAIL : "UPS AC Failed" 512 : name_UPS_BAT_LOW : "UPS Battery Low"
412	Name	MLXEV_CTLDEV_NORMAL_CACHE_MODE
	Description	Controller entered normal cache mode.
	Cause	–
	Action	None
	Address	ctl: 0
	Severity	4 (Information)
	Source	ExPro
413	Name	MLXEV_CTLDEV_DEV_START_CMPLT
	Description	Controller device start complete.
	Cause	–

	Action	None
	Address	ctl: 0
	Severity	4 (Information)
	Source	ExPro
414	Name	MLXEV_CTLDEV_SOFT_ECC_CORRECTED
	Description	Soft ECC error corrected.
	Cause	Faulty memory module.
	Action	Replace memory module
	Address	ctl: 0
	Severity	2 (Error)
	Source	ExPro InPro
415	Name	MLXEV_CTLDEV_HARD_ECC_CORRECTED
	Description	Hard ECC error corrected.
	Cause	Faulty memory module.
	Action	Replace memory module.
	Address	ctl: 0 chn: 13 tgt: 14
	Severity	2 (Error)
	Source	ExP_old InPro
416	Name	MLXEV_CTLDEV_BBU_RECOND_NEEDED
	Description	Battery recondition suggested.
	Cause	–
	Action	None
	Address	ctl: 0
	Severity	4 (Information)
	Source	ExPro
417	Name	MLXEV_CTLDEV_REMOVED_PTNR
	Description	Controller's partner has been removed.
	Cause	Controller was removed or controller was ordered shutdown.
	Action	None

	Address	ctl: 0
	Severity	3 (Warning)
	Source	ExP_old
418	Name	MLXEV_CTLDEV_BBU_OUT_OF_SERVICE
	Description	BBU out of service.
	Cause	BBU will not be able to power the cache if ac power fails. Firmware will switch WriteBack logical drives to WriteThrough.
	Action	Replace BBU battery.
	Address	ctl: 0
	Severity	2 (Error)
	Source	InPro
419	Name	MLXEV_CTLDEV_UPDATE_PTNR_STATUS
	Description	Updated partner's status.
	Cause	–
	Action	None
	Address	ctl: 0 param: 0x2120
	Severity	4 (Information)
	Source	ExPro
	Paramtype	enum
	Params	0x0000 : name_UNKNOWN : “Unknown” 0x0100 : name_BOOTING : “Booting” 0x0200 : name_ACTIVE : “Active” 0x0300 : name_FAILED : “Failed” 0x0400 : name_REMOVED : “Removed” 0x0500 : name_INSERTED : “Inserted”
420	Name	MLXEV_CTLDEV_RELINQUISH_PTNR
	Description	Relinquished partner.
	Cause	Controller failback or array manager initiated
	Action	None
	Address	ctl: 0

	Severity	4 (Information)
	Source	ExPro
421	Name	MLXEV_CTLDEV_INSERTED_PTNR
	Description	Inserted partner.
	Cause	–
	Action	None
	Address	ctl: 0
	Severity	3 (Warning)
	Source	ExP_old
422	Name	MLXEV_CTLDEV_DUAL_ENABLED
	Description	Dual controllers enabled.
	Cause	–
	Action	None
	Address	ctl: 0
	Severity	4 (Information)
	Source	ExPro
423	Name	MLXEV_CTLDEV_KILL_PTNR
	Description	Disabled partner.
	Cause	–
	Action	Refer to enclosure manufacturer’s service manual.
	Address	ctl: 0 param: 0x2120
	Severity	2 (Error)
	Source	ExPro
	Paramtype	enum

	Params	0x0300 : name_PING_TIMEOUT : “ping timeout” 0x0301 : name_NEGOTIATE_GETCHUNK : “Couldn't allocate chunk of memory” 0x0302 : name_NEGOTIATE_ CABLE : “SCSI communication failed or cables are wrong” 0x0303 : name_ NEGOTIATE_HOSTID : “Host ID jumper mismatch” 0x0304 : name_NEGOTIATE_JUMPERS : “SLIP/DIFFL/FBR jumper mismatch” 0x0305 : name_ NEGOTIATE_DSKCHNS : “Number of disk channels present mismatch” 0x0306 : name_NEGOTIATE_HOSTCHNS : “Number of host channels present mismatch” 0x0307 : name_NEGOTIATE_FWVER : “FW version mismatch” 0x0308 : name_NEGOTIATE_FWTYPE : “FW header type mismatch” 0x0309 : name_NEGOTIATE_MEMSIZE : “Memory size mismatch” 0x030a : name_NEGOTIATE_MEMRD : “Memory read of partner failed” 0x030b : name_NEGOTIATE_CACHEMEMSIZE : “Cache memory size mismatch” 0x030c : name_COMMAND : “Received kill partner command” 0x030d : name_DURING_FAILBACK : “Failure handing over TID” 0x030e : name_ENTERING_NEXUS : “Negotiation finished, but nexus not entered in time” 0x030f : name_UNKNOWN : “Failed for unknown reason”
423 (Cont)	Params	0x0310 : name_WBSYNCH_CHAN0 : “Write-back sync to partner - chan 0” 0x0311 : name_WBSYNCH_CHAN1 : “Write-back sync to partner - chan 1” 0x0312 : name_WBSYNCH_CHAN2 : “Write-back sync to partner - chan 2” 0x0313 : name_WBSYNCH_CHAN3 : “Write-back sync to partner - chan 3” 0x0314 : name_WBSYNCH_CHAN4 : “Write-back sync to partner - chan 4” 0x0315 : name_WBSYNCH_CHAN5 : “Write-back sync to partner - chan 5” 0x0316 : name_NEGOTIATE_FWBUILD : “FW build mismatch” 0x0317 : name_NEGOTIATE_CABLE_CROSS : “Device cables are crossed” 0x0320 : name_FAILED_HOTPULL : “Partner removal detect while nexus active” 0x0321 : name_ABSENT_AT_BOOT : “Partner missing at negotiation time” 0x0322 : name_POWERFAIL_DURING_FAIL : “BBU - we power failed before fail-over finished” 0x0323 : name_POWERFAIL_DURING_RELINQ : “BBU - we power failed before relinquish finished” 0x0340 : name_LOCK_TIMEOUT : “Lock timeout” 0x0341 : name_LOCK_SCSI_FAIL : “Lock SCSI failed” 0x0342 : name_FOR_SIMPLEX_MIGRATION : “Kill partner so we can do migrate recovery” 0x0343 : name_FAILED_CTOC : “General ctoc message failure” 0x0344 : name_FLUSH_FAIL : “Flush Failed”

424	Name	MLXEV_CTLDEV_NEXUS
	Description	Dual controllers entered nexus.
	Cause	–
	Action	None
	Address	ctl: 0
	Severity	4 (Information)
	Source	ExPro
425	Name	MLXEV_CTLDEV_BAD_BOOTROM_IMAGE
	Description	Controller boot ROM image needs to be reloaded.
	Cause	Wrong firmware image file downloaded. MAC address changed.
	Action	Contact your service representative to reload the boot ROM image.
	Address	ctl: 0
	Severity	1 (Serious)
	Source	ExPro
426	Name	MLXEV_CTLDEV_BAD_MAC_ADDRESS
	Description	Controller is using default non unique world-wide name.
	Cause	MAC address lost or not set.
	Action	Contact your service representative to set the controller MAC address.
	Address	ctl: 0
	Severity	1 (Serious)
	Source	ExPro
427	Name	MLXEV_CTLDEV_MIRROR_RACE_RECOVERY_FAILED
	Description	Mirror race recovery failed.
	Cause	Some physical disks may have failed.
	Action	Run consistency check and restore consistency.
	Address	ctl: 0
	Severity	1 (Serious)
	Source	InPro

428	Name	MLXEV_CTLDEV_MIRROR_CRITICAL_DRIVE
	Description	Mirror race on critical logical drive.
	Cause	Logical device is critical.
	Action	Replace failed physical disk and rebuild.
	Address	ctl: 0
	Severity	1 (Serious)
	Source	InPro
429	Name	MLXEV_CTLDEV_IN_CLUSTER
	Description	Controller connected to cluster.
	Cause	One or more controllers detected during scan and entered nexus with those controllers.
	Action	None
	Address	None
	Severity	4 (Information)
	Source	InPro
430	Name	MLXEV_CTLDEV_NOT_IN_CLUSTER
	Description	Controller disconnected from cluster.
	Cause	All other controllers in cluster are inaccessible due to controller powered off, controller failure, or cable failure.
	Action	If this was an unexpected event, refer to service manual and correct the failure.
	Address	None
	Severity	4 (Information)
	Source	InPro
431	Name	MLXEV_SYSDEV_DATA_LOSS_IMPROPER_SHUTDOWN
	Description	Controller improperly shutdown. Data might have been lost.
	Cause	User improperly shutdown the controller, resulting in inconsistent logical drive and/or lost data.
	Action	Run consistency check to verify logical drive consistency. If needed, restore data from backup.
	Address	ctl: 0 logdrv: 15

	Severity	2 (Error)
	Source	InPro
432	Name	MLXEV_CTLDEV_AUTOMATIC_FLASH_STARTED
	Description	Dual-active automatic flash of replacement controller.
	Cause	Replacement controller is being flashed with survivor's firmware.
	Action	None
	Address	ctl: 0
	Severity	4 (Information)
	Source	ExPro
433	Name	MLXEV_CTLDEV_NEGOTIATION_FAILED_JUMPERS
	Description	Dual-active negotiation failed jumpers.
	Cause	Jumper settings do not match.
	Action	Replace, adjust, or add components so that the controllers have the same jumper settings.
	Address	ctl: 0
	Severity	2 (Error)
	Source	ExPro
434	Name	MLXEV_CTLDEV_NEGOTIATION_SAME_ID
	Description	Dual-active negotiation failed IDs.
	Cause	Controller IDs are not unique. Both controllers have the same ID.
	Action	Replace, adjust, or add components so that the controllers have the correct IDs.
	Address	ctl: 0
	Severity	2 (Error)
	Source	ExPro
435	Name	MLXEV_CTLDEV_NEGOTIATION_BOARD_TYPE
	Description	Dual-active negotiation failed board types.
	Cause	Board types differ. Both controllers must be of the same type.

	Action	Replace appropriate controller so that the board types are the same.
	Address	ctl: 0
	Severity	2 (Error)
	Source	ExPro
436	Name	MLXEV_CTLDEV_NEGOTIATION_DISK_CHANNELS
	Description	Dual-active negotiation failed disk channels.
	Cause	Number of disk channels differ. Both controllers must have the same number of disk channels.
	Action	Replace, adjust, or add components so that the controllers have the same number of disk channels.
	Address	ctl: 0
	Severity	2 (Error)
	Source	ExPro
437	Name	MLXEV_CTLDEV_NEGOTIATION_HOST_CHANNELS
	Description	Dual-active negotiation failed host ports.
	Cause	Host ports differ. Both controllers must have the same number of host ports.
	Action	Replace, adjust, or add components so that the controllers have the same number of host ports.
	Address	ctl: 0
	Severity	2 (Error)
	Source	ExPro
438	Name	MLXEV_CTLDEV_NEGOTIATION_MEMORY_SIZE
	Description	Dual-active negotiation failed memory size.
	Cause	Memory sizes differ. Both controllers must have the same control store memory size.
	Action	Replace, adjust, or add components so that the controllers have the same control store memory size.
	Address	ctl: 0
	Severity	2 (Error)
	Source	ExPro

439	Name	MLXEV_CTLDEV_NEGOTIATION_CACHE_SIZE
	Description	Dual-active negotiation failed cache memory size.
	Cause	Cache memory sizes differ. Both controllers must have the same cache memory size.
	Action	Replace, adjust, or add components so that the controllers have the same cache memory size.
	Address	ctl: 0
	Severity	2 (Error)
	Source	ExPro
440	Name	MLXEV_CTLDEV_MIRROR_RACE_TABLE_ERROR
	Description	Error in mirror race table.
	Cause	Logical drives appear in the mirror race table that do not exist in the configuration. Possible data integrity issue.
	Action	Check configuration for accuracy, run consistency check and restore consistency.
	Address	ctl: 0 logdrv: 15
	Severity	1 (Serious)
	Source	ExPro
441	Name	MLXEV_CTLDEV_STOP_REJECTED
	Description	A replacement controller attempted to stop the surviving controller.
	Cause	Incompatible replacement controller firmware or hardware.
	Action	Obtain a compatible replacement controller.
	Address	ctl: 0
	Severity	2 (Error)
	Source	ExPro
512	Name	MLXEV_SYSTEM_STARTED_NEW
	Description	System started.
	Cause	The server system (or array management utility server) started.
	Action	If you did not expect a system reboot, investigate.
	Address	–

	Severity	4 (Information)
	Source	InPro
513	Name	MLXEV_SYSTEM_SIZE_TABLE_FULL
	Description	Size table full.
	Cause	Too much physical device size information is defined.
	Action	Remove unused device information from this system.
	Address	None
	Severity	4 (Information)
	Source	InPro
514	Name	MLXEV_SYSTEM_USER_LOGGED_IN
	Description	User logged in.
	Cause	An array management utility user logged in on the server system.
	Action	Not available.
	Address	None
	Severity	4 (Information)
	Source	InPro
515	Name	MLXEV_SYSTEM_USER_LOGGED_OUT
	Description	User logged out.
	Cause	An array management utility user logged out of the server system.
	Action	Not available.
	Address	None
	Severity	4 (Information)
	Source	InPro
516	Name	MLXEV_SYSTEM_ALIVE
	Description	Server alive.
	Cause	Reconnected to server. Server rebooted.
	Action	None
	Address	None

	Severity	4 (Information)
	Source	InPro
517	Name	MLXEV_SYSTEM_DEAD
	Description	Lost connection to the server or server is down.
	Cause	Lost network connection to the server. Server shutdown.
	Action	None
	Address	None
	Severity	1 (Serious)
	Source	Client
518	Name	MLXEV_AUTOBOOT_CHANGED
	Description	Automatic reboot count has changed.
	Cause	Controller has rebooted. Automatic reboot has rearmed itself or was reconfigured.
	Action	None
	Address	ctl: 0 param: 0x2120
	Severity	4 (Information)
	Source	ExPro
	Paramtype	value
519	Name	MLXEV_BBU_BATT_TEST_START
	Description	Battery test has started.
	Cause	Battery test has started.
	Action	None
	Address	ctl: 0
	Severity	4 (Information)
	Source	ExPro
520	Name	MLXEV_BBU_BATT_TEST_COMPLETE
	Description	Battery test has completed.
	Cause	Battery test has completed.
	Action	See battery test status for completion mode.

	Address	ctl: 0
	Severity	4 (Information)
	Source	ExPro
521	Name	MLXEV_BBU_BATT_TEST_CANCELED
	Description	Battery test cancelled.
	Cause	Battery test cancelled.
	Action	None
	Address	ctl: 0
	Severity	4 (Information)
	Source	ExPro
522	Name	MLXEV_BBU_BATT_TEST_FAILED
	Description	Battery test failed, battery bad.
	Cause	Battery test failed, battery bad.
	Action	Replace battery.
	Address	ctl: 0
	Severity	2 (Error)
	Source	ExPro
640	Name	MLXEV_CHANNEL_FAILED
	Description	Channel failed.
	Cause	Cable disconnected.
	Action	Plug in cable.
	Address	ctl: 0 chn: 13
	Severity	3 (Warning)
	Source	Unspecified
641	Name	MLXEV_CHANNEL_OK
	Description	Channel online.
	Cause	Cable reconnected.
	Action	None.
	Address	ctl: 0 chn: 13

	Severity	3 (Warning)
	Source	Unspecified
642	Name	MLXEV_CHANNEL_SCSI_BUS_DEAD
	Description	Back end SCSI bus dead.
	Cause	Lost access to data on SCSI bus.
	Action	None
	Address	ctl: 0 chn: 13
	Severity	1 (Serious)
	Source	ExPro
643	Name	MLXEV_CHANNEL_SCSI_BUS_ALIVE
	Description	Back end SCSI bus alive.
	Cause	Regained access to data on SCSI bus.
	Action	None
	Address	ctl: 0 chn: 13
	Severity	4 (Information)
	Source	ExPro
644	Name	MLXEV_CHANNEL_FIBER_DEAD
	Description	Back end fibre dead.
	Cause	Lost access to data on Fibre Channel.
	Action	Refer to enclosure manufacturer’s service manual.
	Address	ctl: 0 chn: 13
	Severity	1 (Serious)
	Source	ExPro
645	Name	MLXEV_CHANNEL_FIBER_ALIVE
	Description	Back end fibre alive.
	Cause	Regained access to data on Fibre Channel.
	Action	None
	Address	ctl: 0 chn: 13
	Severity	4 (Information)

	Source	ExPro
700	Name	MLXEV_LOG_EMPTY
	Description	Event log empty.
	Cause	Tried to read past last entry.
	Action	None
	Address	ctl: 0
	Severity	4 (Information)
	Source	ExPro
701	Name	MLXEV_LOG_OUT_SYNC
	Description	Event log entries lost.
	Cause	Tried to read an entry that does not exist in the event log.
	Action	None
	Address	ctl: 0
	Severity	4 (Information)
	Source	ExPro
702	Name	MLXEV_LOG_REQUEST_SENSE
	Description	Request sense.
	Cause	A physical disk has generated an error.
	Action	Interpret the Key/ASC/ASCQ and take appropriate action.
	Address	ctl: 0 chn: 13 tgt: 14
	Severity	4 (Information)
	Source	ExPro
703	Name	MLXEV_LOG_SET_RTC
	Description	Set real time clock.
	Cause	Real time clock was set.
	Action	None
	Address	ctl: 0 chn: 13
	Severity	4 (Information)
	Source	ExPro

800	Name	MLXEV_CFG_NEW
	Description	New configuration received.
	Cause	A new configuration was downloaded to controller.
	Action	None
	Address	ctl: 0
	Severity	4 (Information)
	Source	ExPro
801	Name	MLXEV_CFG_CLEAR
	Description	Configuration cleared.
	Cause	Controller was told to clear the configuration.
	Action	None
	Address	ctl: 0
	Severity	4 (Information)
	Source	ExPro
802	Name	MLXEV_CFG_INVALID
	Description	Configuration invalid.
	Cause	The controller found an invalid configuration.
	Action	None
	Address	ctl: 0
	Severity	3 (Warning)
	Source	ExP_old
803	Name	MLXEV_CFG_COD_ACCESS_ERROR
	Description	Configuration on disk access error.
	Cause	The controller could not read the configuration off the disk.
	Action	None
	Address	ctl: 0
	Severity	3 (Warning)
	Source	ExP_old

804	Name	MLXEV_CFG_COD_CONVERTED
	Description	Configuration on disk converted.
	Cause	The controller converted a down level configuration on disk.
	Action	None
	Address	ctl: 0
	Severity	3 (Warning)
	Source	ExP_old
805	Name	MLXEV_CFG_COD_IMPORT_FAILED
	Description	Configuration on disk import failed.
	Cause	The controller could not import the configuration.
	Action	None
	Address	ctl: 0
	Severity	3 (Warning)
	Source	ExP_old
806	Name	MLXEV_DEBUG_DUMP_GENERATED
	Description	A debug dump exists on this system.
	Cause	The controller aborted and created debug dump information.
	Action	Contact field support for assistance in retrieving the data.
	Address	ctl: 0 abort code: 0x1314
	Severity	4 (Information)
	Source	ExPro
807	Name	MLXEV_DEBUG_DUMP_GENERATED_PARTNER
	Description	A debug dump exists on this system.
	Cause	The partner controller aborted and created debug dump information.
	Action	Contact field support for assistance in retrieving the data.
	Address	ctl: 0 abort code: 0x1314
	Severity	4 (Information)
	Source	ExPro

808	Name	MLXEV_CFG_COD_NOT_FOUND
	Description	Valid configuration on disk not found.
	Cause	No physical disk contains a valid configuration on disk area. Physical disks were never written with a configuration on disk. Physical disks were replaced while the controller system was offline.
	Action	If you expected the configuration on disk, investigate.
	Address	ctl: 0
	Severity	4 (Information)
	Source	ExPro
809	Name	MLXEV_CFG_COD_LUN_MAP_NOT_FOUND
	Description	Valid SANMap was not found in the configuration on disk.
	Cause	No physical disk contains a valid SANMap in the configuration on disk area. Physical disks from a system that did not use SANMapping may have been installed.
	Action	If you expected a valid SANMap, investigate.
	Address	ctl: 0
	Severity	4 (Information)
	Source	ExPro
896	Name	MLXEV_FATAL_HANG
	Description	Internal controller is hung.
	Cause	Internal controller is hung.
	Action	Power controller off and on.
	Address	ctl: 0 version: 13.14 – 15 param 0x23222120
	Severity	1 (Serious)
	Source	InPro
	Paramtype	enum

	Params	0x80000001 name_FATAL_HANG_80000001 : “Fatal Hang CfgMonitor_No NIOP” 0x80000002 name_FATAL_HANG_80000002 : “Fatal Hang CfgMonitor_Context size is Greater than Scratch Space” 0x80000011 : name_FATAL_HANG_80000011 : “Fatal Hang Cfg_ Cannot find unused Device addr in COD” 0x80000021 : name_FATAL_HANG_80000021 : “Fatal Hang COD_ Read Data no NIOP” 0x80000022 : name_FATAL_HANG_80000022 : “Fatal Hang COD_ Read Data Contest size is greater than Scratch space” 0x80000023 : name_FATAL_HANG_80000023 : “Fatal Hang COD_ Write Data no NIOP”
	Params	0x80000024 : name_FATAL_HANG_80000024 : “Fatal Hang COD_ Write Data Contest size is greater than Scratch space” 0x80000025 : name_FATAL_HANG_80000025 : “Fatal Hang COD_ Invalid Read/Write Data Area Code.” 0x80000031 : name_FATAL_HANG_80000031 : “Fatal Hang COD_ Upgrade no NIOP” 0x80000036 : name_FATAL_HANG_80000036 : “Fatal Hang COD_ Drv in WriteCFg List mismatch with ProcessCfg List” 0x80000047 : name_FATAL_HANG_80000047 : “Fatal Hang COD_ Drive RW StateFunction is NULL” 0x80000048 : name_FATAL_HANG_80000048 : “Fatal Hang COD_ WrCfg_ Drive not found” 0x80000051 : name_FATAL_HANG_80000051 : “Fatal Hang COD_ NIOP not Available in Remove BDT()” 0x80000052 : name_FATAL_HANG_80000052 : “Fatal Hang COD_ NIOP not avail in CreateBDTForGroup()” 0x80000053 : name_FATAL_HANG_80000053 : “Fatal Hang COD_ NIOP Not avail in RenameBDT()” 0x80000061 : name_FATAL_HANG_80000061 : “Fatal Hang COD_ Process No NIOP” 0x80000062 : name_FATAL_HANG_80000062 : “Fatal Hang COD_ Process context too large” 0x80000067 : name_FATAL_HANG_80000067 : “Fatal Hang COD_ Process Bad Group ID” 0x80000081 : name_FATAL_HANG_80000081 : “Fatal Hang FW_ Start MonitorNo NIOP” 0x80000082 : name_FATAL_HANG_80000082 : “Fatal Hang FW_ Start Monitor Context too large” 0x80000087 : name_FATAL_HANG_80000087 : “Fatal Hang FGI Status- Unknown RAID” 0x80000088 : name_FATAL_HANG_80000088 : “Fatal Hang BGI_ No NIOP” 0x80000091 : name_FATAL_HANG_80000091 : “Fatal Hang FW_ Mailbox Scheme Bad” 0x80000092 : name_FATAL_HANG_80000092 : “Fatal Hang FW_ Write Handler setup Failed” 0x80000093 : name_FATAL_HANG_80000093 : Fatal Hang FW_ IOP queue table is not aligned properly” 0x80000094 : name_FATAL_HANG_80000094 : “Fatal Hang FW_ R5 Check Consistency No Memory (no track avail)”

	Params	0x80000095 : name_FATAL_HANG_80000095 : “Fatal Hang FW_ R5 Check Consistency - Bad Memory Status” 0x80000096 : name_FATAL_HANG_80000096 : “Fatal Hang FW_BOOT_ No NIOP” 0x80000097 : name_FATAL_HANG_80000097 : “Fatal Hang FW_ Unimplemented opcode” 0x800000A1 : name_FATAL_HANG_800000A1 : “Fatal Hang EBHW_ SDRAM status Bad” 0x800000A2 : name_FATAL_HANG_800000A2 : “Fatal Hang EBHW_ SDRAM no Memory” 0x800000B1 : name_FATAL_HANG_800000B1 : “Fatal Hang EBHW_ No NIOP” 0x800000B2 : name_FATAL_HANG_800000B2 : “Fatal Hang EBHW_SPD Module Bank Density / SDRAM width Bad” 0x800000B3 : name_FATAL_HANG_800000B3 : “Fatal Hang EBHW_SPD_ Module Attribute Bad” 0x800000B4 : name_FATAL_HANG_800000B4 : “Fatal Hang EBHW_SPD_ Burst Length Bad” 0x800000B5 : name_FATAL_HANG_800000B5 : “Fatal Hang EBHW_SPD_ Bad Memory type” 0x800000B6 : name_FATAL_HANG_800000B6 : “Fatal Hang EBHW_SPD_ Row / Column Bad” 0x800000B7 : name_FATAL_HANG_800000B7 : “Fatal Hang EBHW_SPD_ Bad CAS Latency” 0x800000B8 : name_FATAL_HANG_800000B8 : “Fatal Hang EBHW_SPD_ Bad RAS2CAS Latency” 0x800000B9 : name_FATAL_HANG_800000B9 : “Fatal Hang EBHW_SPD_ Bad Row Precharge time” 0x800000BA : name_FATAL_HANG_800000BA : “Fatal Hang EBHW_SPD_ Bad Refresh Rate.” 0x800000C1 : name_FATAL_HANG_800000C1 : “Fatal Hang FBDMA_ FootBridge DAM parameter init failed (no mem)” 0x800000D1 : name_FATAL_HANG_800000D1 : “Fatal Hang MORE_ IQP Not Found” 0x800000D2 : name_FATAL_HANG_800000D2 : “Fatal Hang MORE_ cannot get new DP” 0x800000D3 : name_FATAL_HANG_800000D3 : “Fatal Hang MORE_ No mem” 0x800000D4 : name_FATAL_HANG_800000D4 : “Fatal Hang MORE_ Write Setup Failed” 0x800000D5 : name_FATAL_HANG_800000D5 : “Fatal Hang MORE_ Read Setup Failed” 0x800000E1 : name_FATAL_HANG_800000E1 : “Fatal Hang I960HWXOR_ Application Acceleration Unit Error” 0x800000E2 : name_FATAL_HANG_800000E2 : “Fatal Hang I960HWXOR _ Verify Parity Error”

	Params	0x800000E3 : name_FATAL_HANG_800000E3 : “Fatal Hang I960HWXOR _ Bad Address” 0x800000E8 : name_FATAL_HANG_800000E8 : “Fatal Hang FW_HW Board Model Not supported” 0x800000E9 : name_FATAL_HANG_800000E9 : “Fatal Hang FW_ Memory Size not supported” 0x800000F4 : name_FATAL_HANG_800000F4 : “Fatal Hang FW_ Bad Memory Allocation” 0x800000F5 : name_FATAL_HANG_800000F5 : “Fatal Hang FW_ Bad Memory Status” 0x800000F6 : name_FATAL_HANG_800000F6 : “Fatal Hang FW_ Bad Memory Unlock” 0x800000F7 : name_FATAL_HANG_800000F7 : “Fatal Hang FW_ Bad Memory Unuse” 0x800000F8 : name_FATAL_HANG_800000F8 : “Fatal Hang FW_ Illegal Attempt to clear Read Transit” 0x800000F9 : name_FATAL_HANG_800000F9 : “Fatal Hang FW_ Illegal Attempt to clear Write Transit” 0x800000FA : name_FATAL_HANG_800000FA : “Fatal Hang FW_ No SLP_Q entries” 0x800000FB : name_FATAL_HANG_800000FB : “Fatal Hang FW_ Cannot find group member when clearing GEN_ON” 0x800000FC : name_FATAL_HANG_800000FC : “Fatal Hang FW_ No continuous Block of Memory Avail.” 0x800000FD : name_FATAL_HANG_800000FD : “Fatal Hang FW_ No HW found(EastBay)” 0x800000FE : name_FATAL_HANG_800000FE : “Fatal Hang FW_ Reallocate Memory Failed” 0x800000FF : name_FATAL_HANG_800000FF : “Fatal Hang FW_ Not enough memory for C_L_D” 0x80000101 : name_FATAL_HANG_80000101 : “Fatal Hang FW_R5 Regenerate Dirty Block” 0x80000102 : name_FATAL_HANG_80000102 : “Fatal Hang FW_R5Regen_ No Master block found” 0x80000103 : name_FATAL_HANG_80000103 : “Fatal Hang FW_R5Regen_ Bad Sleep time” 0x80000104 : name_FATAL_HANG_80000104 : “Fatal Hang FW_R5Regen_ Bad Lock - other users on the line” 0x80000111 : name_FATAL_HANG_80000111 : “Fatal Hang SAFTE_ No NIOP” 0x80000112 : name_FATAL_HANG_80000112 : “Fatal Hang SAFTE_ No DP” 0x80000113 : name_FATAL_HANG_80000113 : “Fatal Hang SAFTE_ Too Many SAFTE Boxes” 0x80000121 : name_FATAL_HANG_80000121 : “Fatal Hang SMQS_ DAMD ID not free - Resume IO”

	Params	0x80000122 : name_FATAL_HANG_80000122 : “Fatal Hang SMQS_ Active Q length 0 - Resume IO” 0x80000131 : name_FATAL_HANG_80000131 : “Fatal Hang TBBU_ Bad LC link End” 0x80000132 : name_FATAL_HANG_80000132 : “Fatal Hang TBBU_ Bad Memory Size” 0x80000133 : name_FATAL_HANG_80000133 : “Fatal Hang TBBU_ No Memory to allocate” 0x80000134 : name_FATAL_HANG_80000134 : “Fatal Hang TBBU_ Unexpected Dirty Lines” 0x80000141 : name_FATAL_HANG_80000141 : “Fatal Hang BBU_ Bad Read Status” 0x80000142 : name_FATAL_HANG_80000142 : “Fatal Hang Invalid Phase Read” 0x80000151 : name_FATAL_HANG_80000151 : “Fatal Hang WbkMgr_ No Dirty lines” 0x80000152 : name_FATAL_HANG_80000152 : “Fatal Hang WbkMgr_ No IO Descriptor” 0x80000153 : name_FATAL_HANG_80000153 : “Fatal Hang WbkMgr_ No Dirty Tracks” 0x80000154 : name_FATAL_HANG_80000154 : “Fatal Hang WbkMgr_ Bad number of users” 0x80000155 : name_FATAL_HANG_80000155 : “Fatal Hang WbkMgr_ Bad Lock” 0x80000156 : name_FATAL_HANG_80000156 : “Fatal Hang WbkMgr_ Unexpected Dirty Lines” 0x80000201 : name_FATAL_HANG_80000201 : “Fatal Hang HWISP_ Total SCSI Active Exchange Count Underflow” 0x80000202 : name_FATAL_HANG_80000202 : “Fatal Hang HWISP_ Target Exchange Count Underflow” 0x80000203 : name_FATAL_HANG_80000203 : “Fatal Hang HWISP_ Invalid Data Direction” 0x80000204 : name_FATAL_HANG_80000204 : “Fatal Hang HWISP_ Invalid Command” 0x80000211 : name_FATAL_HANG_80000211 : “Fatal Hang HWISP_ NULL DP” 0x80000212 : name_FATAL_HANG_80000212 : “Fatal Hang HWISP_ NULL DP21” 0x80000213 : name_FATAL_HANG_80000213 : “Fatal Hang HWISP_ NULL DP22” 0x80000214 : name_FATAL_HANG_80000214 : “Fatal Hang HWISP_ NULL DP8020” 0x80000215 : name_FATAL_HANG_80000215 : “Fatal Hang HWISP_ NULL DP8021” 0x80000216 : name_FATAL_HANG_80000216 : “Fatal Hang HWISP_ NULL DP8022”

	Params	0x80000217 : name_FATAL_HANG_80000217 : “Fatal Hang HWISP_ NULL DP8031” 0x80000218 : name_FATAL_HANG_80000218 : “Fatal Hang HWISP_ NULL DP8032” 0x80000219 : name_FATAL_HANG_80000219 : “Fatal Hang HWISP_ NULL DP8033” 0x8000021A : name_FATAL_HANG_8000021A : “Fatal Hang HWISP_ NULL DP8034” 0x8000021B : name_FATAL_HANG_8000021B : “Fatal Hang HWISP_ NULL DP8035” 0x80000220 : name_FATAL_HANG_80000220 : “Fatal Hang HWISP_ Stale Entry 0x80000221 : name_FATAL_HANG_80000221 : “Fatal Hang HWISP_ Unexpected Return Value” 0x80000222 : name_FATAL_HANG_80000222 : “Fatal Hang HWISP_ Resource count Overflow” 0x80000223 : name_FATAL_HANG_80000223 : “Fatal Hang HWISP_ No Resource Avail.” 0x80000224 : name_FATAL_HANG_80000224 : “Fatal Hang HWISP_ Unknown Status” 0x80000225 : name_FATAL_HANG_80000225 : “Fatal Hang HWISP_ Notify Acknowledge” 0x80000226 : name_FATAL_HANG_80000226 : “Fatal Hang HWISP_ Unknown Interrupt Code” 0x80000227 : name_FATAL_HANG_80000227 : “Fatal Hang HWISP_ Un Acknowledged Event” 0x80000228 : name_FATAL_HANG_80000228 : “Fatal Hang HWISP_ Invalid RX_ID” 0x80000229 : name_FATAL_HANG_80000229 : “Fatal Hang HWISP_ Too many Scatter Gather Segments.” 0x80000231 : name_FATAL_HANG_80000231 : “Fatal Hang BkgPtrl_ No NIOP” 0x80000232 : name_FATAL_HANG_80000232 : “Fatal Hang BkgPtrl_ No Memory to Allocate” 0x80000238 : name_FATAL_HANG_80000238 : “Fatal Hang BkgPtrl_ Unexpected NULL” 0x80000239 : name_FATAL_HANG_80000239 : “Fatal Hang BkgPtrl_ Bad Patrol State”
	Params	0x8000023A : name_FATAL_HANG_8000023A : “Fatal Hang BkgPtrl_ Bad Physical State” 0x8000023B : name_FATAL_HANG_8000023B : “Fatal Hang BkgPtrl_ No DP” 0x8000023C : name_FATAL_HANG_8000023C : “Fatal Hang BkgPtrl_ No System Drive” 0x8000023D : name_FATAL_HANG_8000023D : “Fatal Hang BkgPtrl_ Bad Capacity” 0x8000023E : name_FATAL_HANG_8000023E : “Fatal Hang BkgPtrl_ Bad Logic State” 0x8000023F : name_FATAL_HANG_8000023F : “Fatal Hang BkgPtrl_ Device not found” 0x80000240 : name_FATAL_HANG_80000240 : “Fatal Hang BkgPtrl_ Bad Raid Level”
897	Name	MLXEV_FATAL_BRKP
	Description	Internal controller firmware breakpoint.

	Cause	Internal controller has encountered a firmware breakpoint.
	Action	Power controller off and on.
	Address	ctl: 0 version: 13.14 – 15 param: 0x23222120
	Severity	1 (Serious)
	Source	InPro
	Paramtype	value
912	Name	MLXEV_I960_HW_ERROR
	Description	Internal controller i960 processor error.
	Cause	Internal controller has encountered i960 processor specific error.
	Action	Power controller off and on.
	Address	ctl: 0 version: 13.14 – 15 param: 0x23222120
	Severity	1 (Serious)
	Source	InPro
	Paramtype	value
	Params	0x80008001 : name_HW_PATU_ERROR : “Error in Primary Address translation unit” 0x80008002 : name_HW_SATU_ERROR : “Error in Secondary Address Translation unit” 0x80008003 : name_HW_PBI_ERROR : “Error in Primary side of PCI to PCI Bridge Unit” 0x80008004 : name_HW_SBI_ERROR : “Error in Secondary side of PCI to PCI Bridge Unit” 0x80008005 : name_HW_ABI_ERROR : “Error in Bus Interface Unit” 0x80008006 : name_HW_DMA0_ERROR : “Error in DMA Unit , Ch0” 0x80008007 : name_HW_DMA1_ERROR : “Error in DMA Unit , Ch1” 0x80008008 : name_HW_DMA2_ERROR : “Error in DMA Unit, Ch2” 0x80008009 : name_HW_IntNMI_ERROR : “Error caused by Software interrupt” 0x8000800A : name_HW_ExtNMI_ERROR : “Error caused by External Source”
928	Name	MLXEV_SARM_HW_ERR
	Description	Internal controller Strong-ARM processor error.
	Cause	Internal controller has encountered Strong-ARM processor specific error.
	Action	Power controller off and on.

	Address	ctl: 0 version: 13.14 – 15 param: 0x23222120
	Severity	1 (Serious)
	Source	InPro
	Paramtype	value
944	Name	MLXEV_SCSI_HW_ERR
	Description	Internal controller back end hardware error.
	Cause	Timeout waiting for interrupt. Timeout waiting for command completion.
	Action	Power controller off and on.
	Address	ctl: 0 version: 13.14 – 15
	Severity	1 (Serious)
	Source	InPro
Total events = 276		

Chapter 7. Maintenance Activities

Maintenance processes include the following activities, that you may need to perform on arrays:

1.49 Running a Logical Drive Initialization

Logical drive initialization offers the ability to run a full foreground initialization of logical drives at the time of your choice, not just immediately following a new configuration. During the initialization process all data is cleared and zeros are written to the disks.

1. Select a logical drive from the navigation area or the content area.
2. Click **Logical->Initialize** from the menu bar.

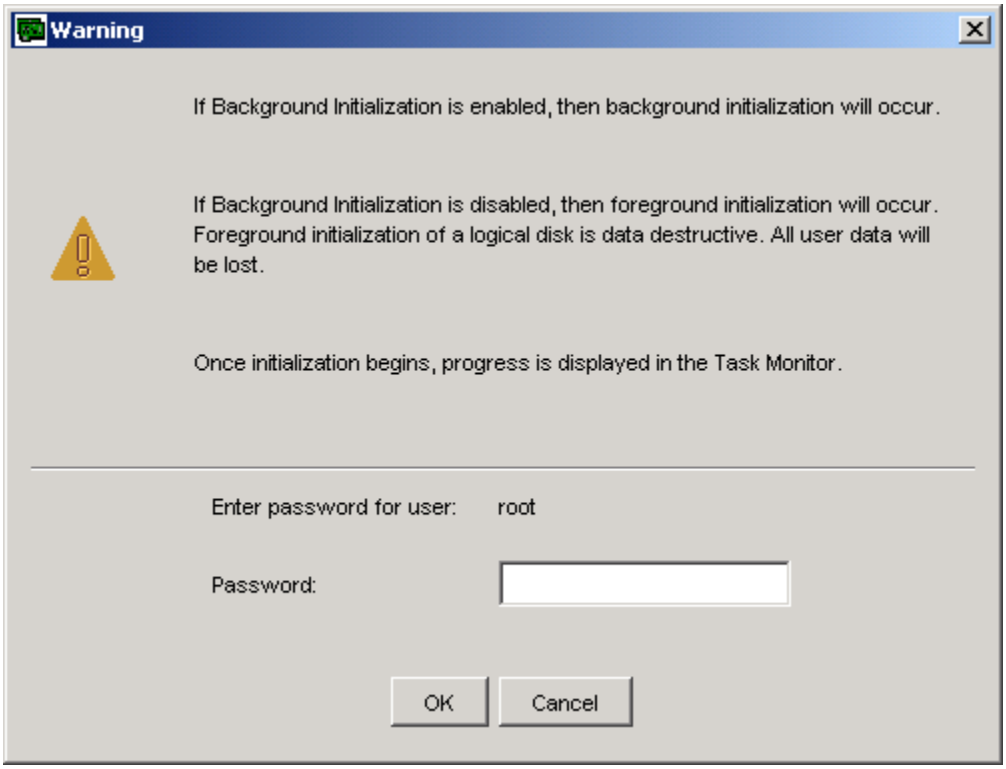

Figure 98 Logical Drive Initialization Warning

A Warning dialog box displays telling you that this process is data-destructive.

3. Type your password.

4. Click **OK** to confirm the initialization.

5. See "Monitoring Long Operation Tasks" on page 102 to monitor the progress of your initialization or cancel it.

1.50 Running a Logical Drive Consistency Check

Run a consistency check on each of your fault tolerant logical drives on a regular schedule. Deciding on how often to check your drives depends on the type of data. Run a consistency check more often if the cost of losing data is very high, and run it less often if the cost of losing data is low. Running a consistency check lessens the exposure to possible data loss situations. A consistency check scans the logical drive to determine whether consistency data has become corrupted and needs to be restored.

1. Select a logical drive from the navigation area or the content area.

2. Click **Logical->Check Consistency** from the menu bar.

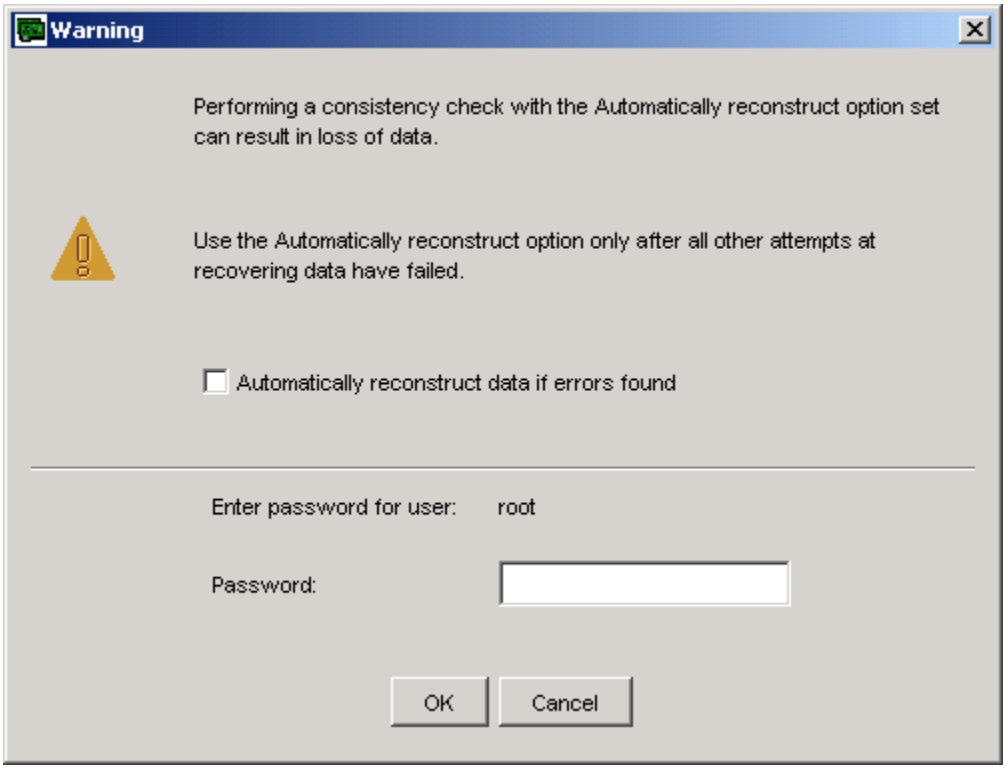

Figure 99 Logical Drive Consistency Check

A Warning dialog box displays telling you that this process could result in data loss.

3. Select the **Automatically reconstruct data if errors are found** checkbox if desired.

   CAUTION:
   Using this option can result in data loss.

4. Type your password.

5. Click **OK** to confirm the consistency check.

6. 1.48 Monitoring Long Operation Tasks on page 134, to monitor the progress of your Consistency Check or cancel it.

1.51 Running a Manual Rebuild

In the event of a physical disk drive failure, a Rebuild is the process of regenerating and writing data to a replacement disk drive. During a rebuild process, the system drive operates in an online critical and rebuilding, or degraded state. The physical device is operating in an online rebuild state.

- Arrays are rebuild candidates only if all of the logical drives use a redundant RAID level (1, 3, 5, 0+1).
- The Automatic rebuild management controller parameter is disabled.
- The Operational fault management controller parameter is enabled.
- The replacement drive must have a capacity that is at least as large as the consumed capacity of the failed drive.

Use the following steps to perform a rebuild:

1. Remove and replace the failed disk drive.

2. Allow sufficient time for the new disk drive to spin up. The drive should appear in an Un-configured state.

3. To rebuild an array, click **Array->Rebuild**.

   A message dialog box opens to confirm that the drive rebuild is in progress.

4. Click **OK** to complete the drive rebuild.

1.52 Disk Scrubbing

Note: Disk scrubbing is not available on all RAID controller firmware versions.

Disk Scrubbing periodically verifies all of the sectors, including COD (Configuration on Disk - retains the latest version of the saved configuration at a reserved location on every physical device), in your RAID configured and spare physical drives. Disk scrubbing works with all RAID levels. This feature detects any potential problems with the drives such as a mechanical problem or bad data access paths. This allows you to resolve potential errors before they can degrade your system. For setting Drive Error Handling parameters, see 1.32 Setting and Modifying Controller Properties on page 75, and locate the Advanced Tab Controller Properties.

To check the status of all physical drives connected to a controller (this includes spare drives), select a controller from the navigation area. To start or stop Disk Scrubbing, click **Controller->Disk Scrubbing** from the menu bar.

Select one of the following options:

- Run Continually - Starts Disk Scrubbing immediately and checks the status of all physical drives continually until you select Stop from the Disk Scrubbing menu.

- Run Once - Starts Disk Scrubbing immediately, checks the status of all physical drives once, and then stops. If this command is selected while a Run Continually operation is in progress, it is ignored.

- Stop - Stops Disk Scrubbing immediately. You can restart Disk Scrubbing using either Run Continually or Run Once.

Disk Scrubbing can also be viewed and stopped from the Task Monitor by clicking the Disk Scrubbing operation and clicking **Task->Cancel** from the Task menu.

1.53 Editing a WWN Table

The World Wide Name Table contains entries for every host, which has accessed the controller. Editing a World Wide Name Table (WWN) allows you to remove old host references from the table. To edit a WWN table follow these steps.

1. Click **Controller->Edit WWN Table**. The Edit WWN Table opens.

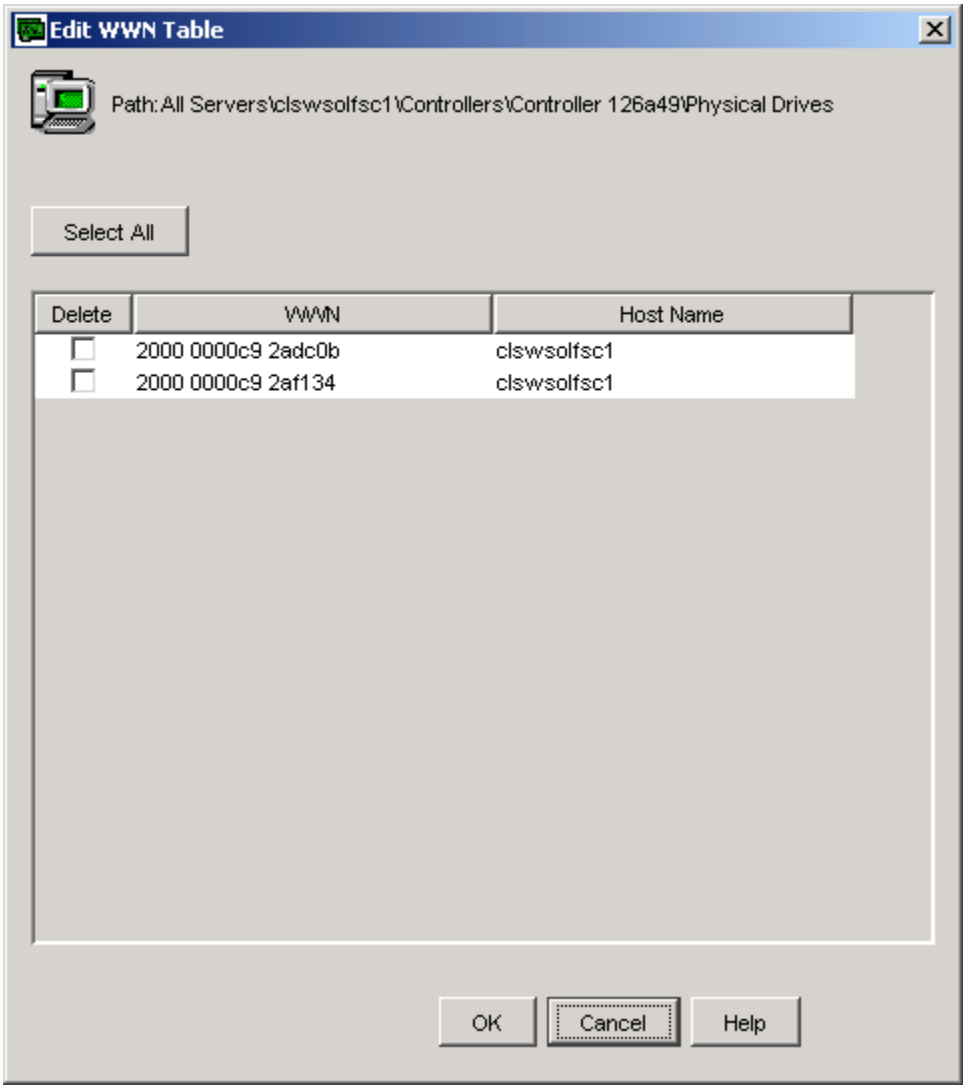

Figure 100 Edit a WWN Table

The WWN Table displays entries (WWN and Host Name) for every host that has ever accessed the controller.

2. Select the WWN that you want to delete.

3. Click **OK**. The disconnected hosts that you selected are removed from the table.

   Note: The controller does not allow the user to delete entries for hosts that are connected.

   The controller takes a few minutes to recognize this change. If you need to see the change immediately, reboot the controller by clicking **Controller->Reboot**.

1.54 Forcing a Physical Drive Online or Offline

Forcing a physical drive online or offline also affects the state of any logical drive currently utilizing the physical drive.

To force a physical drive offline or online do the following.

1. Select the physical drive that you want to force online/offline.

2. Click **Physical->Force Online** or **Physical->Force Offline** on the menu bar or right-click **Physical drive->Force Online/Offline**. An information dialog box opens telling you that the physical drive you have selected is now online/offline. The dialog box also displays the ID number of the physical drive.

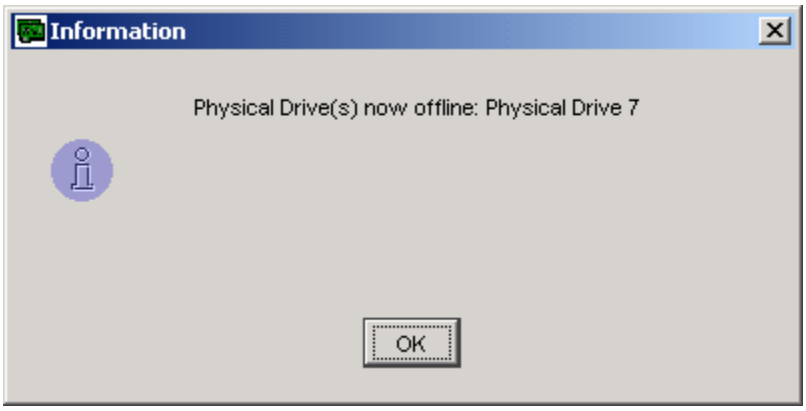

Figure 101 Forcing a Physical Drive Online/Offline Information Dialog box

3. Click **OK**.

1.55 Locating a Logical Drive, Physical Drive, or Enclosure

When you are performing maintenance tasks, you may need to locate a drive or enclosure in order to replace it.

To locate a logical or physical drive do the following.

1. Select a logical or physical drive to locate.

2. Either right-click the logical or physical drive and select locate or from the Logical or Physical Menu, select Locate.

3. A message dialog box opens (this example shows the Logical Drive message dialog box) telling you that the LEDs on the physical drives contained by the logical drive are now blinking. The blinking stops in 30 seconds. Look for the blinking LEDs on the physical drives.

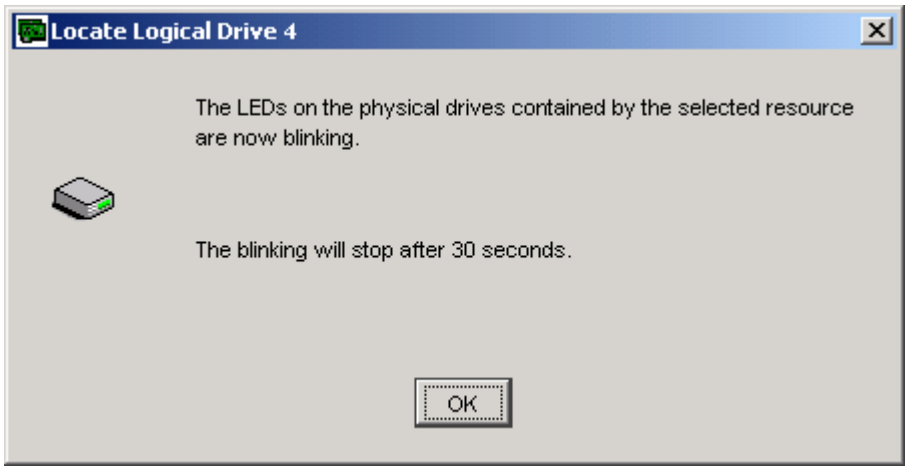

Figure 102 Locating a Logical Drive Message Box

4. Click **OK**.

To locate an enclosure, select the enclosure in the navigation area and select Locate from the Enclosure pull-down menu.

1.56 Updating a Controller's Firmware

The firmware update utility is used to update firmware by "flashing" the new code stored in a specified .ima file to the on-board flash memory. As maintenance releases of the firmware become available, this utility allows you to keep your controller current.

Note: Updating controller firmware requires the controller to reboot.

Use the following steps to perform a firmware update.

1. Click **Controller->Update Firmware** from the menu bar.

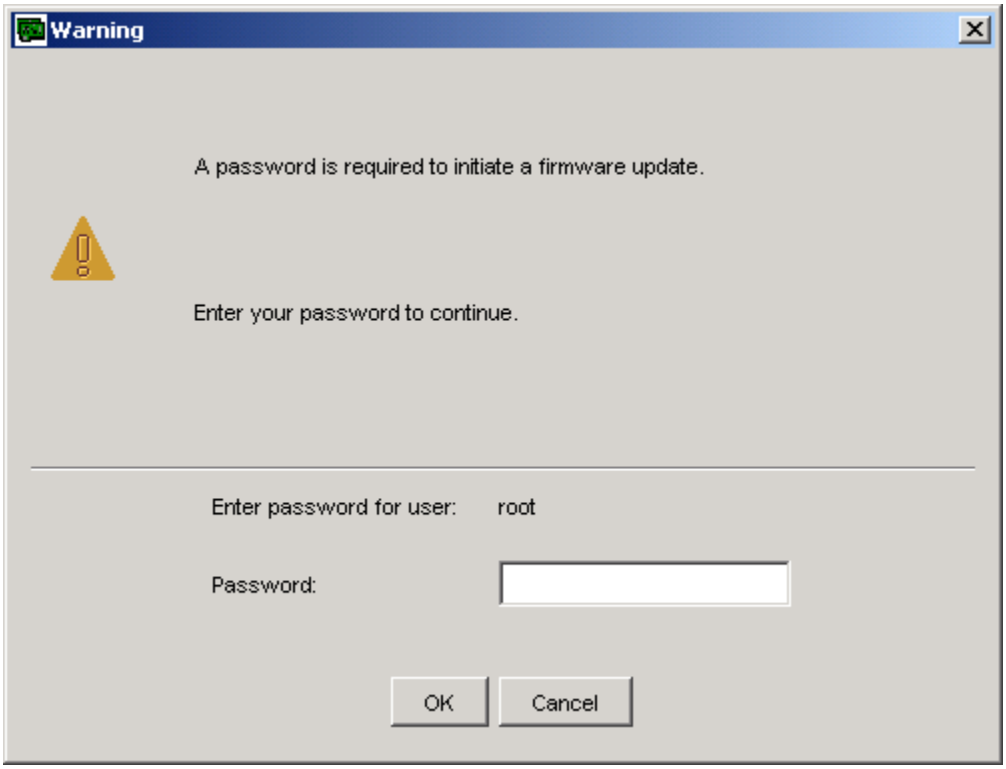

Figure 103 Update Firmware Password

A Warning dialog box displays telling you that a password is required to initiate a firmware update.

2. Type your password.

3. Click **OK**. The Update Firmware dialog box opens.

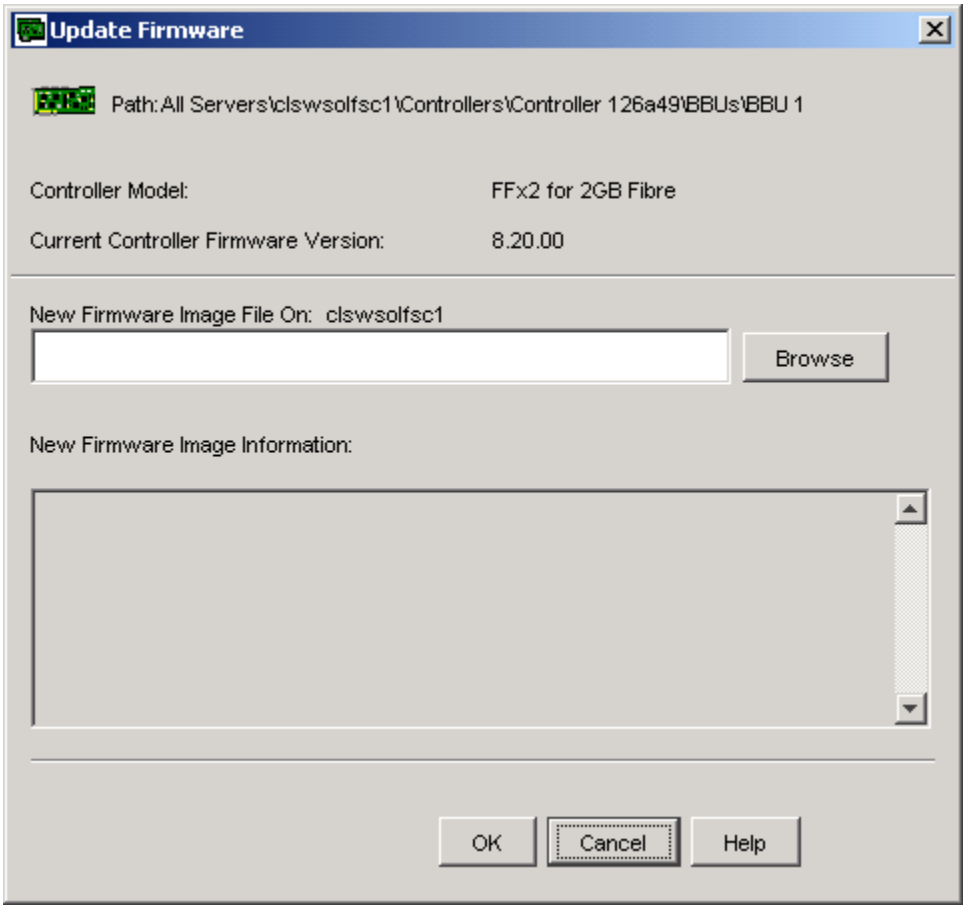

Figure 104 Firmware Update

4. Type the name of the appropriate image file (.ima) in the "New firmware image file on field," or click **Browse** to locate the file.

   Note: Spheras Storage Director does not impose any length limitation on the file path.

   If you select Browse, the Open Image file dialog box is displayed, see Figure 105 Open Image File Dialog Box.

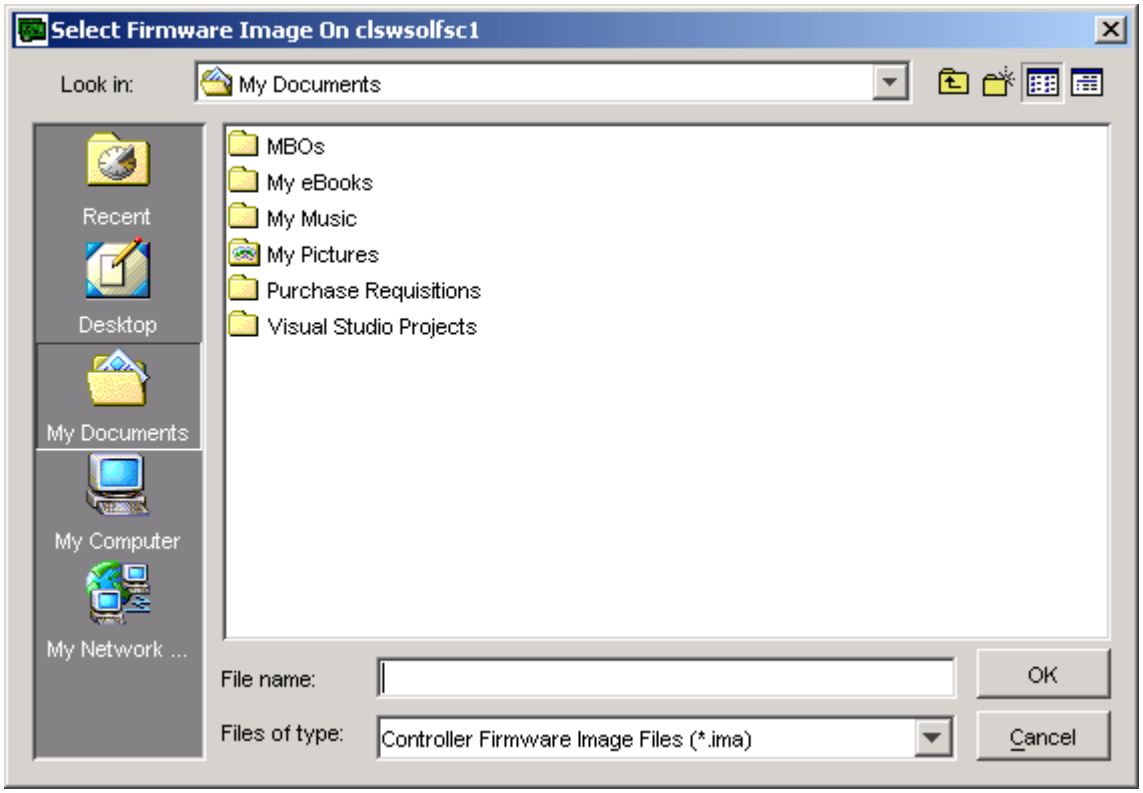

Figure 105 Open Image File Dialog Box

5. Navigate to and select the image file name.

6. Click **Open**. The .ima file name is displayed in the "New firmware image file on" field. The details of the firmware file are displayed in the "New firmware image information" field.

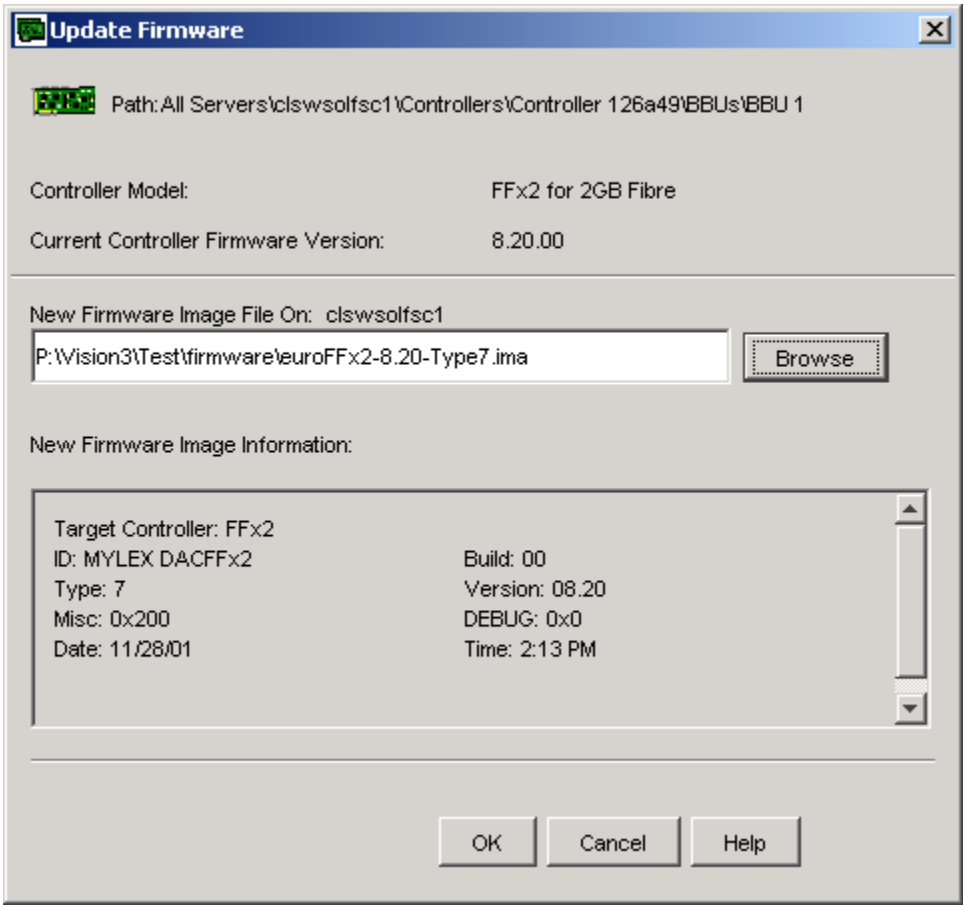

Figure 106 Upgrade Firmware Details

7. Click **OK**.

8. A warning message, *"This firmware update requires the controller to reboot. The reboot takes approximately 30 seconds to complete. Any I/O must be stopped."* is displayed.

9. Click **Yes** to complete the firmware update. Click **No** to cancel the firmware update.

1.57 Reconditioning a Battery Backup Unit

The status of the battery pack is continuously monitored by the *gas gauge*. The gas gauge does not read the actual charge on a battery pack, but measures and tracks the current into and out of the battery pack. The gas gauge is synchronized to the battery pack condition by performing a full discharge cycle followed by a charge cycle. A new battery pack must first be discharged and fully charged followed by another discharge and charge cycle. This process continues until the remaining capacity value is close to the full charge capacity value. A recondition cycle is performed the first time a battery pack is connected and repeated periodically. The frequency of recondition cycles is determined by the number of discharges. Any time a battery pack is disconnected and reconnected, a recondition cycle is performed. Additionally, when a controller or the module containing the battery pack is removed, a recondition cycle is performed.

Follow these steps to perform a Battery Recondition:

1. In the navigation area right-click the **BBU -> Recondition** or select **BBU -> Recondition Battery**.

2. An information dialog box opens confirming that the battery is now being reconditioned and to check the Task Monitor to view the status.

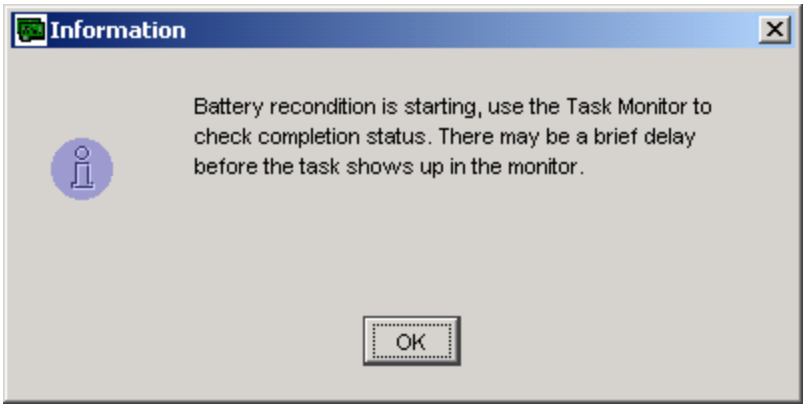

Figure 107 Reconditioning a Battery Information Dialog Box

3. Click **OK**.

4. Open the Task Monitor to view the status of the recondition. See 1.48 Monitoring Long Operation Tasks on page 134 for more information about the Task Monitor.

1.58 Shutting Down a Controller

The purpose of this activity is to cleanly prepare to shut down the controller system, leaving the BBU charged. In a duplex controller environment, both controllers are prepared for shutdown. Shutdown disables the controller cache and saves all data to physical drives prior to disconnecting the BBU.

1. Click **Controller->Shutdown** on the menu bar. The Controller Shutdown Warning dialog box opens.

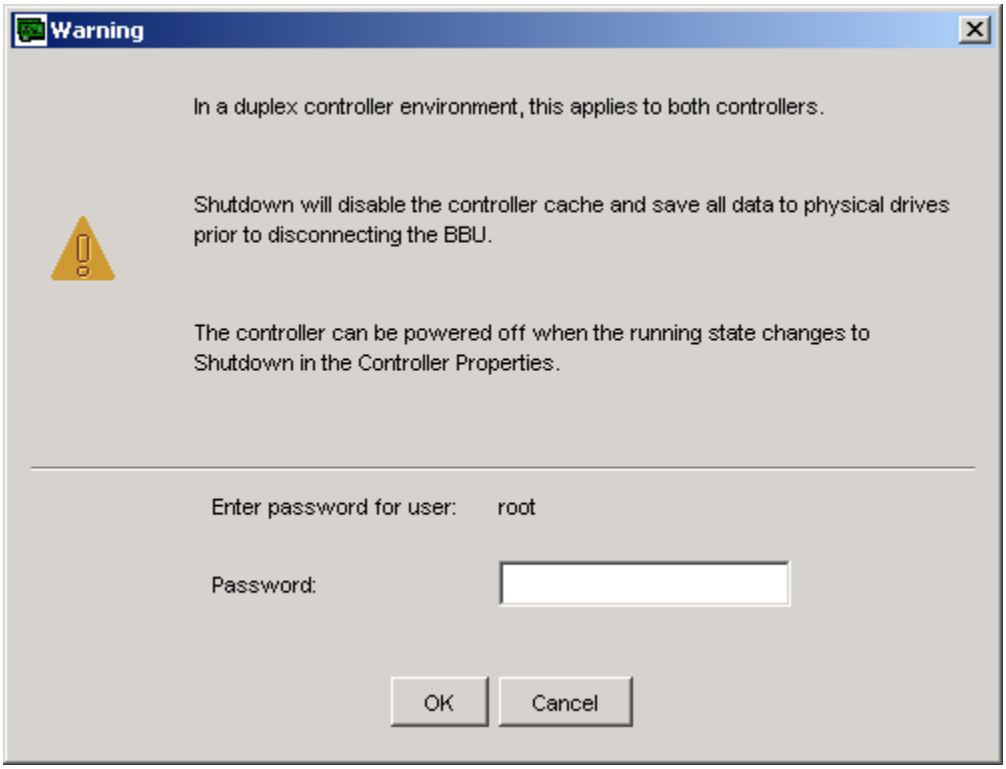

Figure 108 Controller Shutdown Warning Dialog Box

2. Type your Administrative password in the password field.

3. Click **OK**. The controller can be physically powered off when its status changes to Shutdown. See the Task Monitor to view the progress.

1.59 Repair Actions

7.1.1 *General* Repairs

For general repairs, first see 6.1.2 Viewing Event Details on page 115. The event detail tells you the event ID and gives you a description of problem that needs to be repaired. Using this information, take measures to fix the problem:

- Check, all cable connections, to be sure none is disconnected or broken.

- Check all generated events; there may be more than one problem.
- Replace parts as needed.

Remember that some events are the result of other events; one event can cause a cascade of others if left unfixed. For example, a problem with a hard drive causes a warning event, "Hard drive error found", which could lead to a consistency check error event, "Consistency check on logical drive error", which could lead to a serious event, "A hard disk has failed", which could lead to a critical event, "A logical drive has failed."

7.1.2 Repair an Array

Events that may cause an array to need repair:

- Expand capacity stopped
- Rebuild stopped
- New device failure
- Device failure

To repair an array, first see 6.1.2 Viewing Event Details on page 115. The event details tell you the event ID and give you a description of the event that caused the array to need repair. Using this information, take measures to fix the problem:

- Check, all cable connections, to be sure none is disconnected or broken.
- Check that all logical drives are present and functioning. If needed, see 7.1.3 Repair a Logical Drive on page 231.
- Replace any parts that need to be replaced, such as a failed fan.
- If expanding an array, make sure that the logical drives and physical drives can support the expansion. See 1.37 Expanding an Array on page 105 for more information.

7.1.3 Repair a Logical Drive

Events that may cause a logical drive to need repair:

- The logical drive is critical
- Logical drive failure
- Logical drive initialization failure
- Disk failure
- Consistency check error or failure
- A drive has been spun down

To repair a logical drive, first see 6.1.2 Viewing Event Details on page 115. The event details tell you the event ID and give you a description of the event that caused the logical drive to need repair. Using this information, take measures to fix the problem:

- Check, all cable connections, to be sure none is disconnected or broken.
- Check that all physical drives in the logical drive are present and functioning. Replace physical drives as needed.
- Replace any parts that need to be replaced, such as a failed disk.
- Check that the controller is installed and has power.
- If a consistency check had an error or failed, try again. If this problem persists, replace the disk drive. Check the errors other generated errors to identify the drive.
- If expanding a logical drive, make sure that the logical drive and physical drives can support the expansion. See 1.35 Expanding a Logical Drive on page 103 for more information.

7.1.4 Repair a Controller

To repair a controller, first see 6.1.2 Viewing Event Details on page 115. The event details tell you the event ID and give you a description of the event that caused the controller to need repair. Using this information, take measures to fix the problem:

- Typically, controller problems can be traced to a problem with the BBU, a logical drive, an array, a physical drive, or an enclosure. In the navigation area, expand the problem controller's tree and view its components in the Contents area. Select a component that is in a critical state, and view the repair action.
- Check all generated events; there may be more than one problem.
- Replace parts as needed.

Remember that some events are the result of other events; one event can cause a cascade of others if left unfixed. For example, a problem with a hard drive causes a warning event, "Hard drive error found", which could lead to a consistency check error event, "Consistency check on logical drive error", which could lead to a serious event, "A hard disk has failed", which could lead to a critical event, "A logical drive has failed."

7.1.5 Repair a BBU

Events that may cause a BBU to need repair:

- Power supply failure
- BBU power low
- BBU recondition needed
- BBU has failed

To repair a BBU, first see 6.1.2 Viewing Event Details on page 115. The event details tell you the event ID and give you a description of the event that caused the BBU to need repair. Using this information, take measures to fix the problem:

- If the BBU needs to be reconditioned or power is low, recondition the BBU. Remember that the battery must be fully reconditioned before being used. If a BBU recondition has

been canceled in the past, the BBU may not have enough power to sustain the write-back cache.

- Check, all cable connections, to be sure none is disconnected or broken.
- Replace the BBU as needed.

7.1.6 Repair an Enclosure

Events that may cause an enclosure to need repair:

- Fan failure
- Power supply failure
- Over temperature limit surpassed
- Enclosure access is critical or degraded
- BBU power low or BBU recondition needed or BBU has failed.

To repair an enclosure, first see 6.1.2 Viewing Event Details on page 115. The event details tell you the event ID and give you a description of the event that caused the enclosure to need repair. Using this information, take measures to fix the problem:

- If the BBU needs to be reconditioned or power is low, recondition the BBU.
- Remove access to the enclosure and power down.
- Check, all cable connections, to be sure none is disconnected or broken.
- Replace any parts that need to be replaced. Such as the fan, power supply, or BBU.

7.1.7 Repair a Server

Events that may cause a server to need repair:

- Lost connection to server
- Server is down
- One or more controllers visible to the server are having problems.

To repair a server:

- Check if the server's status indicates that the connection was lost. If this is the case, then the server may be down, the network connection to it is not working, or the Spheras Storage Director Server is not running on that server. Try reconnecting to the server.
- If the server's connection is good, then check if any of the server's controllers are in a critical state. In the navigation area, expand the server. Check if any controllers are in a critical state. If so, then one or more controllers are having problems. Expand the problem controller's tree to view its components. Click on the component that is in a critical state and view the repair action for it.

Appendix A Manually Configuring HTTP Servers

Note: Use the procedures in this section if you elected not to have your HTTP server automatically configured during installation of the Spheras Storage Director server or if you have trouble with your HTTP server. Refer to the following sections:

A.1. Windows - HTTP Web Server Configuration

In order for the Spheras Storage Director Client to interact with the Spheras Storage Director Server, an HTTP server must be installed on a minimum of one server. The Spheras Storage Director will work with either an Apache or Microsoft IIS server. If multiple HTTP servers are installed on the same system, they must be configured to communicate on different ports.

Use the following procedures according to which HTTP Web Server you are using.

A.1.1 Windows - Manually Configuring an Apache HTTP Web Server

To configure an Apache HTTP Web Server, follow these steps:

Using a text editor open the following file:

C:\Program Files\Apache Group\Apache\conf\httpd.conf

Locate the section that discusses defining aliases:

```
# Aliases: Add here as many aliases as you need (with no limit). The format is
# Alias fakename realname
#
# Note that if you include a trailing / on fakename then the server will
# require it to be present in the URL. So “/icons” isn’t aliased in this
# example, only “/icons/”. If the fakename is slash-terminated, then the
# realname must also be slash terminated, and if the fakename omits the
# trailing slash, the realname must also omit it.
#
Alias /icons/ "C:/Program Files/Apache Group/Apache2/icons/"
# Important: keep this line here

<Type your Alias here, see Step 3 for instructions>

<Directory “C:/Program Files/Apache Group/Apache/icons”>
  Options Indexes MultiViews
   AllowOverride None
  Order allow,deny
   Allow from all
   </Directory>
```

Note: Forward slashes (/) are used as the delimiter for all operating systems.

Alias /SSD/ "{your path}"

Windows example:

Alias /SSD/ "C:/Program Files/Adaptec/Spheras Storage Director Server/"

The alias allows your machine to access the Spheras Storage Director Server and launch the Spheras Storage Director Client. In the example above, replace the text {your path} with the full path where Spheras Storage Director Server has been installed on your system.

Scroll to find the Server Name variable

#
 # ServerName allows you to set a host name which is sent back to clients for
 # your server if it's different than the one the program would get (i.e., use
 # "www" instead of the host's real name).
#
 # Note: You cannot just invent host names and hope they work. The name you
 # define here must be a valid DNS name for your host. If you don't understand
 # this, ask your network administrator.
 # If your host doesn't have a registered DNS name, enter its IP address here.
 # You will have to access it by its address (e.g., http://123.45.67.89/)
 # anyway, and this will make redirections work in a sensible way.
#
 # 127.0.0.1 is the TCP/IP local loop-back address, often named localhost. Your
 # machine always knows itself by this address. If you use Apache strictly for
 # local testing and development, you may use 127.0.0.1 as the server name.
#
ServerName <Type Your IP Address or DNS Name Here>

Modify this variable by typing in your IP address or DNS name here.
If this variable is not changed, the Apache server will use it's own IP address (127.0.0.1) and run as a loop back server. If this happens, it will cause the SSD server not to function correctly or server failure.

Save the file.

Refresh the Apache HTTP Web Server. Go to the Control Panel and launch the Services dialog box. Locate "Apache" and click Stop, then click Start.

The Apache HTTP Web Server manual configuration is complete.

A.1.2 Windows - Manually Configuring a Microsoft IIS HTTP Web Server

To configure a Microsoft IIS HTTP Web Server, follow these steps:

Click Start->Programs->Administrative Tools->Internet Services Manager. The Internet Information Services dialog box opens.

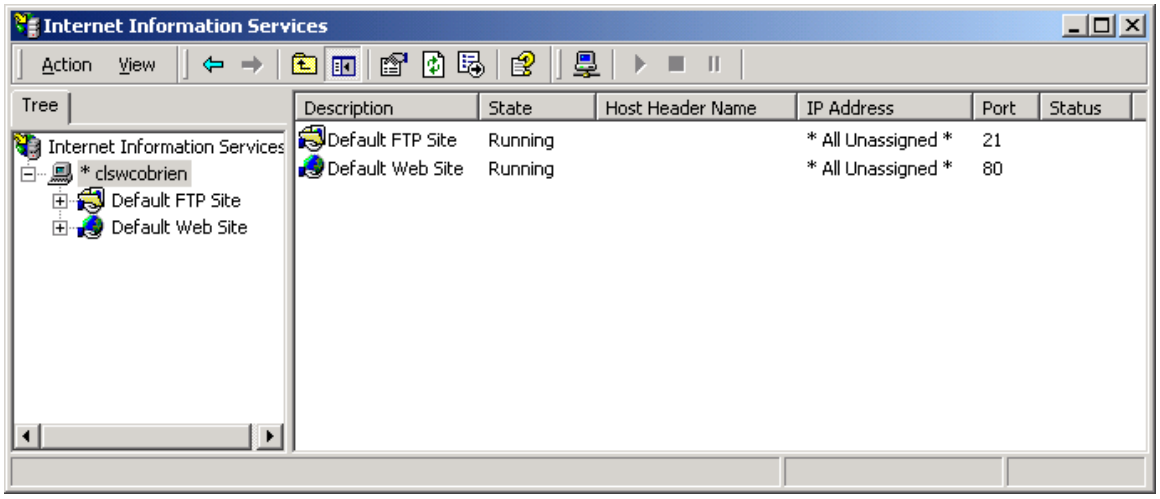

Figure 109 Internet Information Services

Internet Information Services Dialog Box

Locate your web sever in the left pane, in this example it's called Default Web Site, Figure 109 Internet Information Services.

Right-click Your Web Server->New->Virtual Directory.

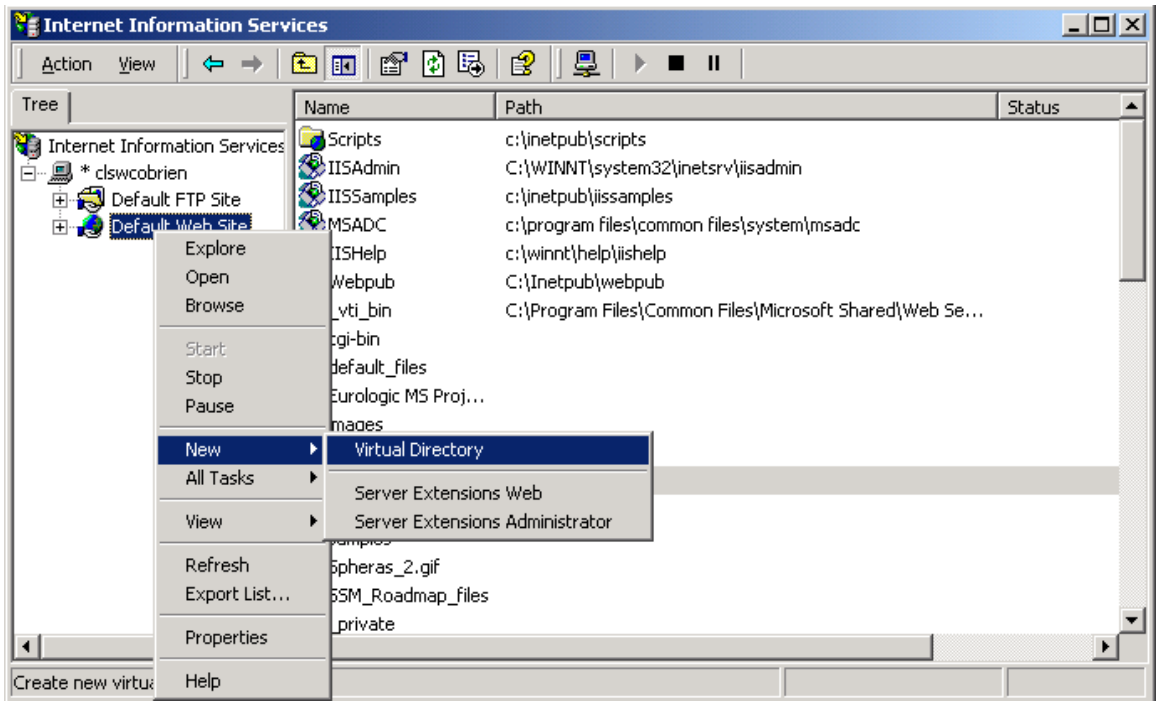

Figure 110 IIS – Selecting a new virtual directory

Select a New Virtual Directory

The Virtual Directory Creation Wizard Welcome Window opens.

Figure 111 Virtual directory creation welcome screen

Virtual Directory Creation Wizard Welcome Dialog Box

Click Next. The Virtual Directory Alias dialog box opens.

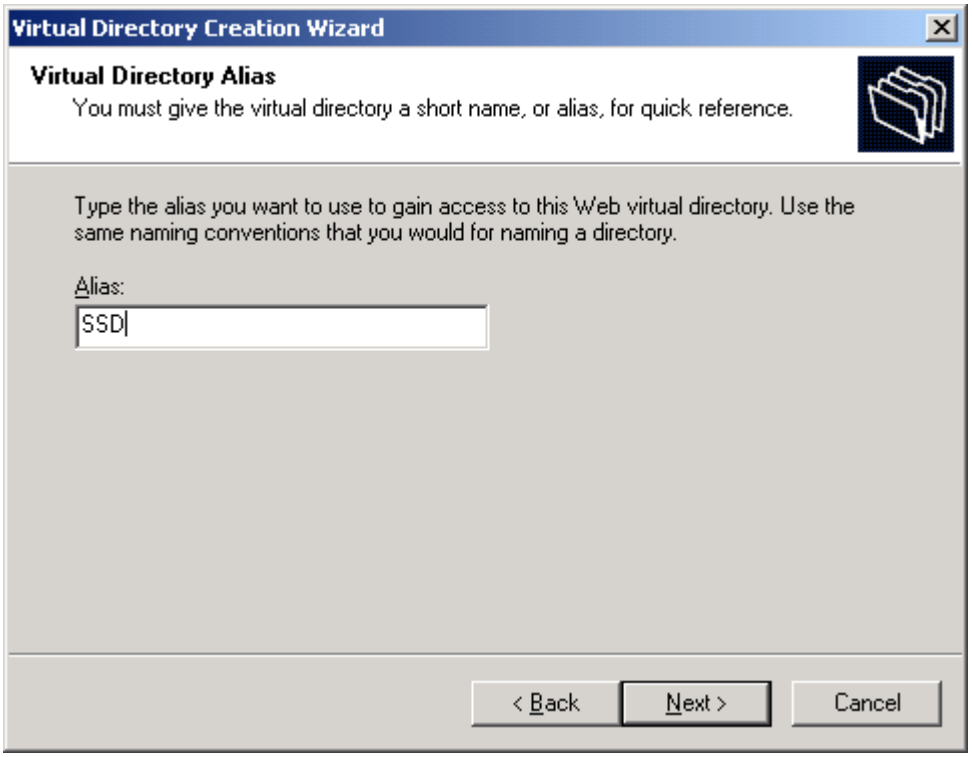

Figure 112 Spheras storage director alias

Virtual Directory Alias

Type an Alias for the Spheras Storage Director. The path will be the destination folder of your Spheras Storage Director Server, see Figure 113 Web site content directory.

Click Next. The Web Site Content Directory dialog box opens.

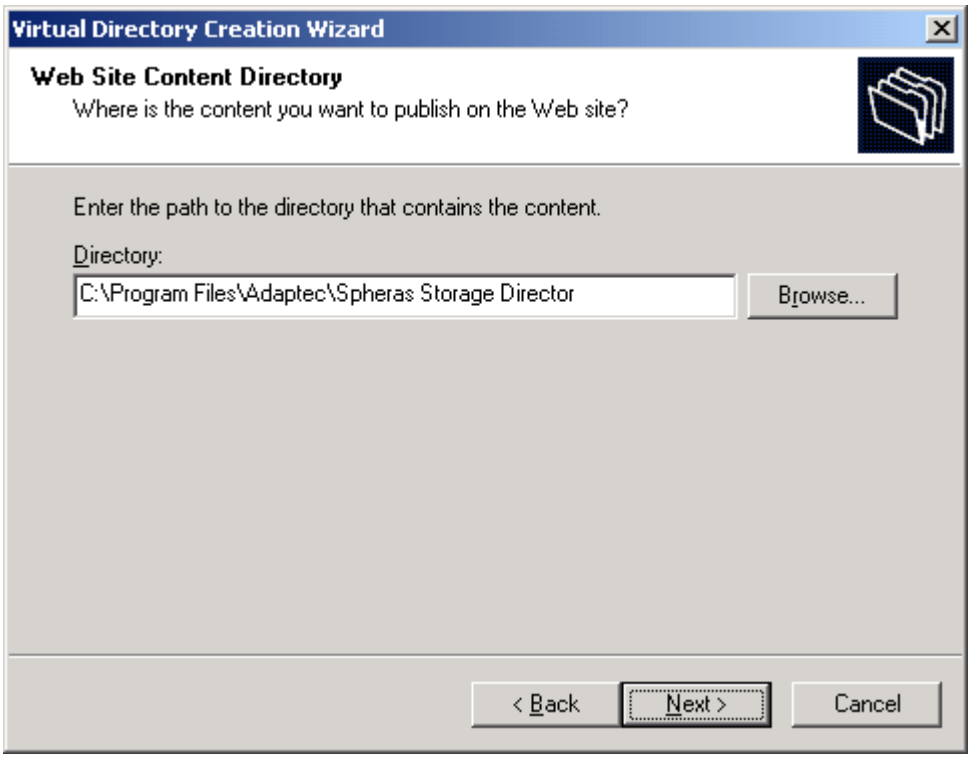

Figure 113 Web site content directory

Web Site Content Directory

In the Directory field, type the path; or click Browse to navigate to the directory. The path will be the destination folder of your Spheras Storage Director Server. See Figure 113 Web site content directory, for an example of a path.

Click Next. The Access Permissions dialog box opens.

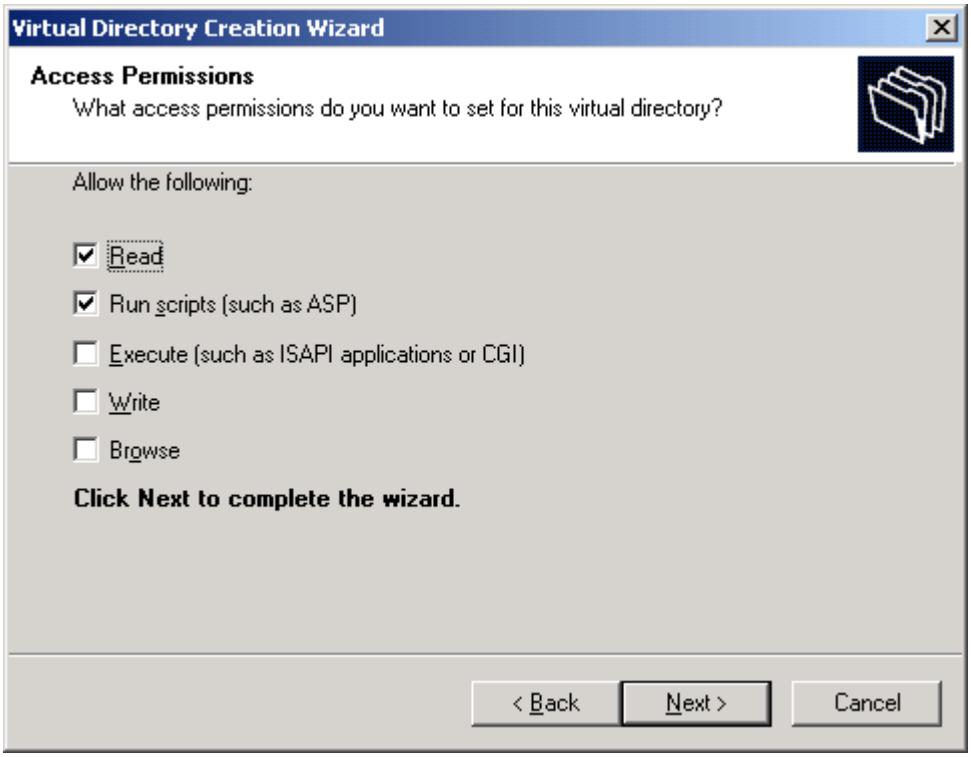

Figure 114 Virtual directory access permissions

Select Access Permissions

Select the Read option. This is the only option required by Spheras Storage Director. Select other options if desired.

Click Next.

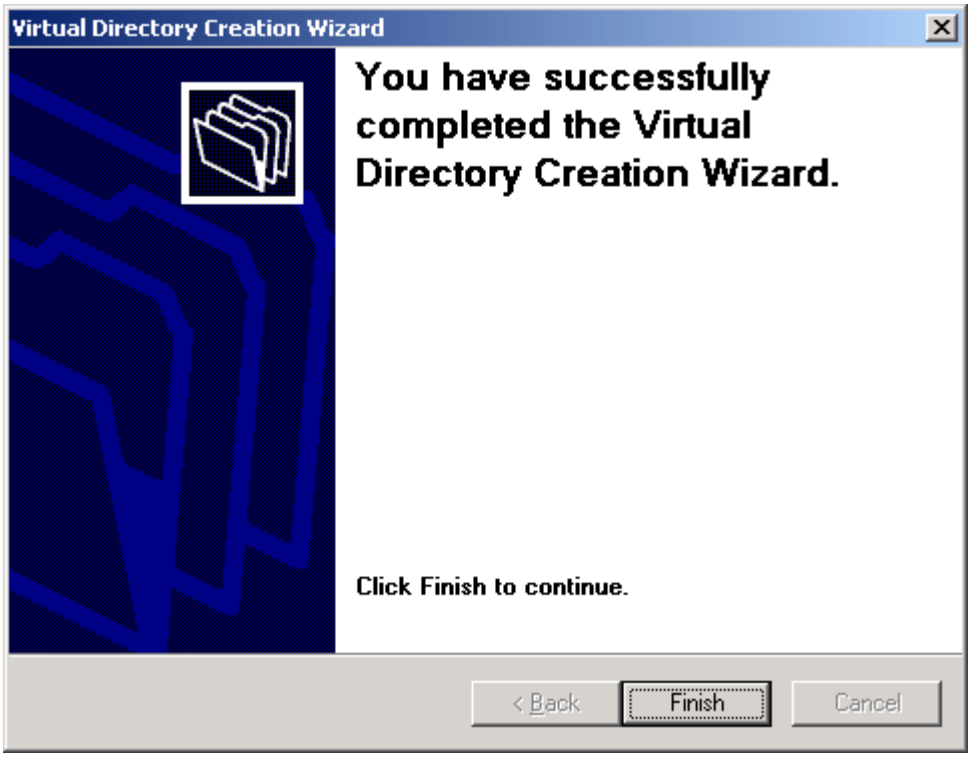

Figure 115 Virtual directory creation complete screen

Complete Your Virtual Directory Creation

Click Finish. The Microsoft IIS HTTP Web Server manual configuration is complete.

Refresh the Microsoft IIS HTTP Web Server. Go to the Control Panel and launch the Services dialog box. Locate "IIS Admin Service" and click Stop, then click Start. Alternately, you can use the Internet Services Manager to stop and restart your HTTP server.

A.2. Linux - Manually Configuring an Apache HTTP Web Server

In order for the Spheras Storage Director Client to interact with the Spheras Storage Director Server, an HTTP server must be installed on a minimum of one server. To configure an Apache HTTP Web Server, follow these steps:

Using a text editor open the following Apache file, the path may vary according to where your Apache software was installed.

/usr/local/apache/conf/http.conf

Locate the section that discusses defining aliases:

#
 # Aliases: Add here as many aliases as you need (with no limit). The format is
 # Alias fakename realname
#
 # Note that if you include a trailing / on fakename then the server will
 # require it to be present in the URL. So "/icons" isn't aliased in this
 # example, only "/icons/". If the fakename is slash-terminated, then the
 # realname must also be slash terminated, and if the fakename omits the
 # trailing slash, the realname must also omit it.
#

 # We include the /icons/ alias for FancyIndexed directory listings. If you
 # do not use FancyIndexing, you may comment this out.
#
Alias /icons/ "/usr/local/apache/icons/"
 # Important: keep this line here

<Type your Alias here, see Step 3 for instructions>

<Directory "/usr/local/apache/icons">
Options Indexes MultiViews
AllowOverride None
Order allow,deny
Allow from all
</Directory>

You must type the Alias exactly at the location it appears in the previous step.

Alias /SSD/ "{your path}"

Example:

Alias /SSD/ "/opt/Adaptec/SpherasStorageDirectorServer/"

This alias allows your machine to access the Spheras Storage Director Server and launch the Spheras Storage Director Client. In the example above, replace the text {your path} with the full path where the Spheras Storage Director Server is installed on your system.

Scroll to find the Server Name variable

#
 # ServerName gives the name and port that the server uses to identify itself.
 # This can often be determined automatically, but we recommend you specify
 # it explicitly to prevent problems during startup.
#
 # If this is not set to valid DNS name for your host, server-generated
 # redirections will not work. See also the UseCanonicalName directive.
#
 # If your host doesn't have a registered DNS name, enter its IP address here.
 # You will have to access it by its address anyway, and this will make
 # redirections work in a sensible way.
#
ServerName <Type Your IP Address or DNS Name Here>

Modify this variable by typing in your IP address or DNS name here.
If this variable is not changed, the Apache server will use it's own IP address (127.0.0.1) and run as a loop back server. If this happens, it will cause the SSD server not to function correctly or server failure.

Save the file.

Refresh the Apache HTTP Web Server the commands may vary according to which version of Linux you are running. Go to

/usr/local/apache/bin

and type

./apachectl stop

then type

./apachectl start

The Apache HTTP Web Server manual configuration is complete.

A.3. Solaris - Manually Configuring an Apache HTTP Web Server

In order for the Spheras Storage Director Client to interact with the Spheras Storage Director Server, an HTTP server must be installed on a minimum of one server. To configure an Apache HTTP Web Server, follow these steps:

Using a text editor open the following Apache file, the path may vary according to where your Apache software was installed.

usr\apache\conf\httpd.conf

Locate the section that discusses defining aliases:

```
#
# Aliases: Add here as many aliases as you need (with no limit). The format is
# Alias fakename realname
#
<IfModule mod_alias.c>
# Important: keep this line here
Alias /SSD/ "/opt/Adaptec/SpherasStorageDirectorServer/"    #
   # Note that if you include a trailing / on fakename then the server will
   # require it to be present in the URL. So "/icons" isn't aliased in this
   # example, only "/icons/"..
   #

   Alias /icons/ "/var/apache/icons/"


   <Directory "/var/apache/icons">
      Options Indexes MultiViews
      AllowOverride None
      Order allow,deny
      Allow from all
   </Directory>

Alias /SSD/ "{your path}"
```

This alias allows your machine to access the Spheras Storage Director Server and launch the Spheras Storage Director Client. In the example above, replace the text {your path} with the full path where the Spheras Storage Director Server is installed on your system.

Scroll to find the Server Name variable

```
#
# ServerName allows you to set a host name which is sent back to clients for
# your server if it's different than the one the program would get (i.e., use
# "www" instead of the host's real name).
#
# Note: You cannot just invent host names and hope they work. The name you
# define here must be a valid DNS name for your host. If you don't understand
# this, ask your network administrator.
# If your host doesn't have a registered DNS name, enter its IP address here.
# You will have to access it by its address (e.g., http://123.45.67.89/)
# anyway, and this will make redirections work in a sensible way.
#
ServerName <Type Your IP Address or DNS Name Here>
```

Modify this variable by typing in your IP address or DNS name here.
If this variable is not changed, the Apache server will use it's own IP address (127.0.0.1) and

run as a loop back server. If this happens, it will cause the SSD server not to function correctly or server failure.

Save the file.

Refresh the Apache HTTP Web Server. Go to

/usr/apache/bin

and type

./apachectl stop

then type

./apachectl start

Note: You may need to use chmod if the commands do not work.

Open your Netscape or Explorer browser and go to http://localhost. The following page should open if the Apache server is up and running.

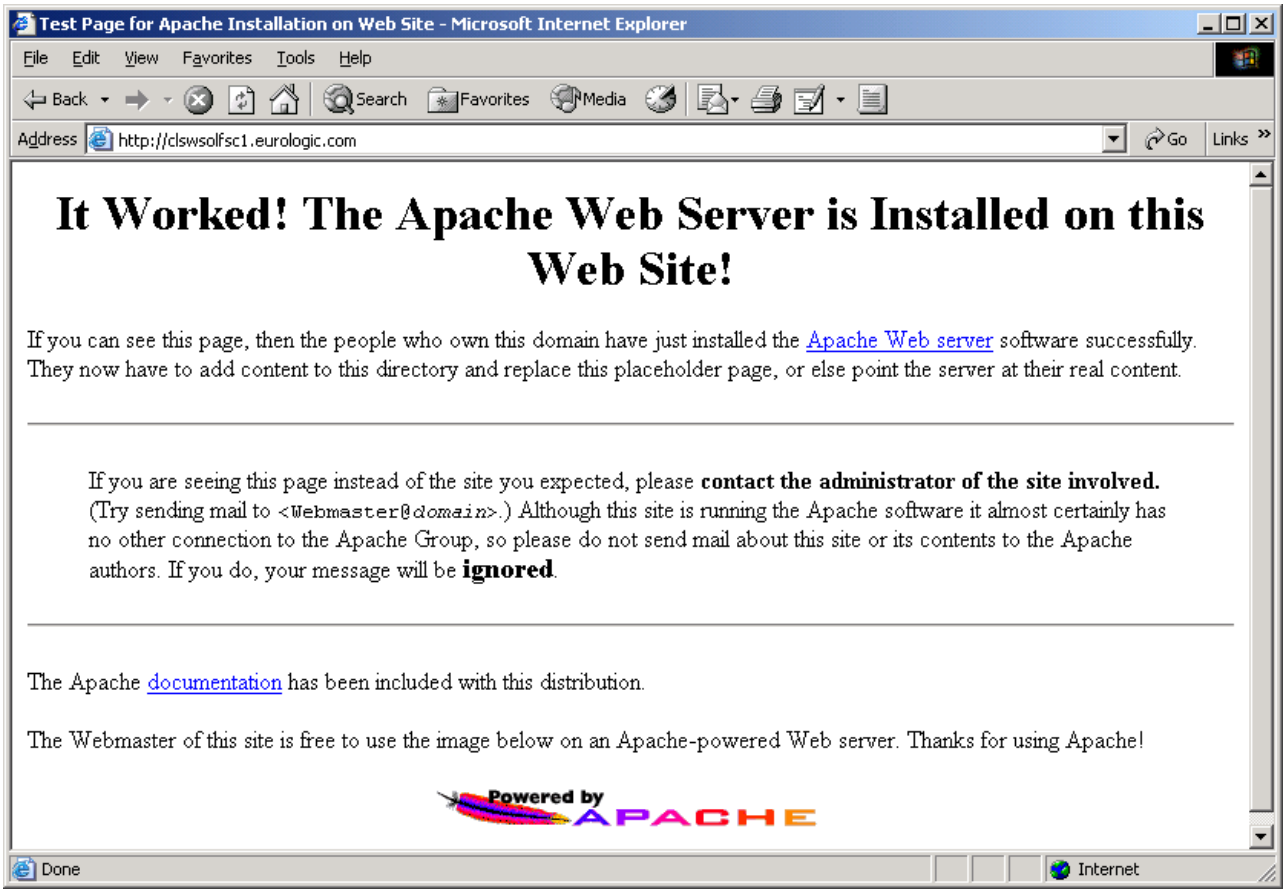

Figure 116 Apache Installation

The Apache HTTP Web Server manual configuration is complete.

Appendix B Detailed Information About Configurations

To determine the best configuration for your system, you need to answer some questions and make some decisions about the type of configuration you need; what you can afford; and what components you have.

The following sections provide basic RAID and configuration information.

B.1. What is the purpose of your array?

Think about this question before configuring the array. Is the array for increasing the capacity of your general-purpose file and print server? Will the array support a database application running 24 hours a day? Will the array contain large audio or video files to be played back on demand? Will the array serve as a repository for imaging systems?

Each of the above applications has a different access profile; the type and frequency of read and write activity performed against the array over time. Identifying the data access profile will help you determine a strategy favoring capacity, availability, or performance.

For example, so-called "video" servers typically write data (i.e., the video clip) infrequently, but play back the data very often. The ratio of reading to writing in this type of an array is far different from a general-purpose file server, which is handling small read and write operations all day long. In addition, the characteristics of the files themselves are very different, video and image files are typically very large when compared to letters, memos and spreadsheet files.

Correct installation of the disk array and the controller requires a proper understanding of RAID technology and the concepts described in this section.

RAID is an acronym for Redundant Array of Independent Disks. Disk array controllers implement several different versions of the Berkeley RAID technology.

When System Drives are defined or created, an appropriate RAID level must be selected. This decision is based on the following priorities:

- Disk capacity utilized (number of drives)
- Data redundancy (fault tolerance)
- Disk performance

The controller makes the RAID implementation and the disks' physical configuration transparent to the host operating system. This means the host operating system drivers and software utilities are unchanged, regardless of the RAID level selected.

B.2. RAID Levels

Selecting the proper RAID level for a specific data storage application requires consideration is given to the benefits of each technique.

RAID 0: Striping

RAID 1: Mirroring

RAID 3: Striping with parity on one drive

RAID 5: Striping with parity across all drives

RAID 0+1: Mirroring and Striping

JBOD: Single drive control

B.2.1 RAID 0: Striping

The controller stripes data across multiple drives

With Spanning = RAID 00

Benefits: Very high data throughput, especially for large files.

Drawbacks: Does not deliver any fault tolerance. If any drive in the array fails, all data is lost.

Uses: Intended for non-critical data requiring high performance.

Drives: Minimum, 2. Maximum, 16.

Fault Tolerance: No.

B.2.2 RAID 1: Mirroring

The controller duplicates data from one drive to another.

With Spanning = RAID 10

Benefits: Provides 100% data redundancy. Should one drive fail, the controller simply switches reads and writes to the other drive.

Drawbacks: Requires two drives for the storage space of one drive. While a controller is rebuilding a drive, users will experience reduced performance if they try to read or write data to the array.

Uses: When array availability is most important.

Drives: Minimum, 2. Maximum, 2.

Fault Tolerance: Yes

B.2.3 RAID 3: Striping with Parity

RAID Level 3 provides redundant information in the form of parity to a parallel access striped array, permitting regeneration and rebuilding in the event of a disk failure. One stripe of parity protects corresponding stripes of data on the remaining disks.

With Spanning = RAID 30

Benefits: Provides for high transfer rate and high availability, at an inherently lower cost than mirroring.

Drawbacks: Transaction performance is poor because all RAID Level 3 array member disks operate in lockstep.

Drives: Minimum, 3. Maximum, 16.

Fault Tolerance: Yes

B.2.4 RAID 5: Striping with Parity

The controller stripes blocks of data and parity information across all drives.

With Spanning = RAID 50

Benefits: Uses a fraction of the disk space required by RAID 1 to achieve data redundancy. RAID 5 provides good performance, for transaction processing applications, because each drive can read and write independently. Should a drive fail, the controller continues to allow reads and writes by calculating the missing information using the parity data. The controller recreates lost data on a replacement drive without interrupting access by users. If a hot spare is available, it can perform the rebuild automatically.

Drawbacks: Cannot match RAID 0 in write performance because of the processing required to compute and write error-correction data. While a controller is rebuilding a drive, users will experience reduced performance if they try to read or write data to the array.

Drives: Minimum, 3. Maximum, 16.

Fault Tolerance: Yes

B.2.5 RAID 0+1: Mirroring and Striping

The controller combines the functions of both data striping (RAID 0) and disk mirroring (RAID 1).

With Spanning = RAID 0+1+0

Benefits: Optimizes for both fault tolerance and performance. RAID 0+1 provides excellent performance for all data needs.

Drawbacks: Requires half the available disk space for data redundancy, the same as RAID level 1.

Drives: Minimum, 3. Maximum, 16.

Fault Tolerance: Yes

B.2.6 JBOD (Just a Bunch of Drives): Single drive control

The controller treats each drive as a stand-alone disk and provides a high-performance cache (sometimes referred to as a standalone drive).

Benefits: Cache reduces the amount of time the computer waits for a disk to get to the right place to read or write data.

Drawbacks: Does not provide data redundancy and does not use striping for performance enhancements.

Drives: Minimum, 1. Maximum, 1.

Fault Tolerance: No

B.3. RAID Rules

When connecting devices and configuring them to work with a RAID controller, follow this set of guidelines:

- To lessen their impact on the SCSI bus, connect SCSI-based tape, CD-ROM and other non-disk devices on a single channel, which preferably does not service any disk drives
- If using more than two drives, distribute them equally among all the channels on the controller. This results in better performance. Depending upon the model, disk array controllers have between one and five channels
- A maximum of 16 devices can comprise a Drive Group.
- A Drive Group can contain devices on any channel.
- Include all drives of the same capacity in the same Drive Group.
- If configuring for a standby or hot spare drive, make sure the standby drive size is greater than or equal to the capacity of the smallest drive in all the redundant Drive Groups.

- When replacing a failed drive, make sure the replacement drive size is greater than or equal to the size of the smallest drive in the effected Drive Group.

B.4. CAP Strategies Overview

Capacity, Availability, and Performance (CAP) should characterize your expectations of the disk array subsystem.

What is the best strategy for configuring the array? Do you want to access the maximum capacity of your disk drive investment? Or do you want to configure your array for a high degree of fault-tolerant operation? Perhaps performance is extremely important to you. Every installation has a different set of requirements, and most installations have a combination of these.

Unfortunately, these options are somewhat mutually exclusive and it isn't possible to configure an array with all of these characteristics optimized. For example, you can't have maximum capacity and maximum availability in a single array. Some of the disks must be used for redundancy, thus, capacity is reduced.

Similarly, configuring a single array for both maximum availability and maximum performance is not an option (unless you own a very large quantity of disk drives).

You will need to make compromises. Fortunately, your disk array controller is versatile enough to offer any of these preferences, either singly or in the most favorable combination possible.

Bottom line: You need to decide which capability is most important for your operational environment, and which is second most important. Once you've prioritized your expectations, it should be relatively easy to configure the array to meet your expectations. To help you decide, complete the CAP questionnaire.

B.4.1 CAP Questionnaire

Before configuring the array, complete the following CAP strategy questionnaire. After completing this questionnaire, take the information and use it as a guideline when choosing your configuration method.

I have ____ disks for the array.

Rank the following statements in order of importance: 1, 2, and 3

I want to emphasize capacity____ (see B.2 RAID Levels on page, 246)

I want to emphasize availability _____ (see B.2 RAID Levels on page, 246)

I want to emphasize performance _____ (see B.2 RAID Levels on page, 246)

I have ____ channels on the array controller (usually 1, 2, or 3).

I plan to ___ /not to ____ use a Hot Spare disk.

I do ____ /do not _____ have a battery backup for the controller cache.

I do ____ /do not _____ have a UPS for the array.

B.5. Configuring for Maximum Capacity

The following table shows the relationship between the different RAID levels, effective capacities offered for X number of drives of N capacity, and the computed capacities of five 2GB drives.

Maximum Capacity RAID LevelsSSD Installation and User Guide.doc

RAID Level	Effective Capacity	Example (capacity in GB)
0	X*N	5*2=10
1	(X*N)/2	5*2/2=5
3	(X-1)*N	(5-1)*2=8
5	(X-1)*N	(5-1)*2=8
0+1	(X*N)/2	(5*2)/2=5
JBOD	X*N	5*2=10

The greatest capacities (100%) are provided by RAID 0 and JBOD. Unfortunately, with these two solutions, there is no fault-tolerance. RAID 5 gives the next best capacity, followed by RAID 1 or RAID 0+1.

B.6. Configuring for Maximum Availability

When considering optimizing for availability, it is important to understand some of the terminology concerning the condition of array operation. These definitions are presented in the following table.

Maximum Availability Array Conditions

Array Condition	Meaning
Normal	The array is operating in a fault-tolerant mode and can sustain a drive failure without potential data loss.
Critical	The array is functioning and all data is available, but the array cannot sustain a drive failure without potential data loss.

Degraded	The array is functioning and all data is available, but the array cannot sustain a drive failure without potential data loss. Additionally, a reconstruction or rebuild operation is taking place, reducing the performance of the array. The rebuild operation takes the array from a critical condition to a normal condition.
Not Fault Tolerant	No fault-tolerant RAID levels are configured for any of the drives in the array.

An additional measure of fault-tolerance (or improved availability) is achieved using a Hot Spare (Standby) disk. This disk is powered-on but idle during normal array operation. If a failure occurs on a disk in a fault-tolerant set, the Hot Spare disk takes over for the failed drive. After completing an automatic rebuild cycle, the array continues to function in a fully fault-tolerant mode. This means the array can suffer a second drive failure and continue to function before any disks are replaced.

Impact of Controller Cache on Availability

Every disk array controller has a disk cache. This DRAM or EDRAM physical memory increases the performance of data retrieval and storage operations. The amount of disk cache varies with the controller model. The controller uses this memory to store disk writes.

The controller may report to the operating system that the write is complete as soon as the controller receives the data. This is referred to as Write Back (WB) cache. This improves performance, but exposes you to data loss if a system crash or power failure occurs before the data in the cache is written to disk.

To avoid potential loss of data, outfit your controller with a battery backup for the cache memory. The battery backup option enables the cache to be retained until the data can be written to disk. Your data is important. This simple and inexpensive insurance can prevent a power interruption from turning into a disaster.

Notes: If the array will not use write cache, battery backup is not necessary.

If write cache is used without the battery back-up option, an uninterruptible power supply (also known as a UPS) is required.

B.7. RAID Levels and Availability

To see the advantages of the RAID levels as they apply to availability, see B.2 RAID Levels on page, 246.

B.8. Configuring for Maximum Performance

To see the advantages of the RAID levels as they apply to performance, refer to the following table.

Maximum Performance RAID Level Advantages

RAID Level	Access Profile Characteristics

Maximum Performance RAID Level Advantages

RAID Level	Access Profile Characteristics
0	Excellent for all types of I/O activity.
1	Excellent for write-intensive applications.
3	Good for sequential or random reads and sequential writes.
5	Excellent for sequential or random reads and sequential writes.
0+1	Excellent for write-intensive applications.
JBOD	Mimics normal, individual disk performance characteristics.

Table 2 RAID level – access profile characteristics

B.9. Array Configuration Table: Overview

Consider these points when using the tables:

To determine the approximate capacity of a particular configuration, assume that each physical disk has a unit capacity of one. The value in the Effective Capacity column will indicate the effective available capacity for a Drive Group with a given RAID level.

A configuration, is fault-tolerant or has enhanced availability, if there is an asterisk after the value in the Effective Capacity column. The number of asterisks is proportional to the level of fault-tolerance, or availability, that a configuration will exhibit. The more asterisks, the better the fault tolerance.

To determine the performance of a given configuration, correlate the configuration's RAID level with the application. See CAP Strategies.

Drives that are not affiliated with a Drive Group are assumed to be hot spare, or standby drives, and provide an additional measure of fault-tolerance. The use of Standby drives with redundant RAID levels are indicated in the Possible RAID Levels column as SBY.

In most of the configurations, multiple RAID levels can exist for a given Drive Group. These multiple RAID levels can be implemented in the same Drive Group on different System Drives. A System Drive can be any percentage of a Drive Group. If different RAID levels are configured for a single Drive Group, then the effective capacity will be different than the value shown in the table. This is because the table does not consider the mixing of RAID levels within a Drive Group.

The maximum number of logical drives that can be created is 32.

B.10. Array Configuration Tables

The following tables (One Drive, Two Drives, Three Drives, Four Drives, Five Drives, Six Drives, Seven Drives, Eight Drives) give drives and drive groups, possible RAID levels for these drives or drive groups, and the effective capacity for these drives or drive groups. See also Example RAID configurations. For example:

For One Drive: If you have one drive, you would have one drive group (A:o), a JBOD RAID level (A:JBOD), and an effective capacity of 1 drive (A:1).

For Two Drives: If you have two drives, you could have two drive groups (A:o, B:o), both are JBODs (A:JBOD, B:JBOD), each with an effective capacity of 1 drive (A:1, B:1). Or, you could have one drive group (A:oo), a RAID 0 level (A:0), with an effective capacity of 2 drives (A:2). Or, you could have one drive group (A:oo), a RAID 1 level (A:1), with an effective capacity of 1 drive with fault tolerance (A:1*).

For Five Drives: If you have five drives, you could have five drive groups (A, B, C, D, E), all are JBODs, each with an effective capacity of 1 drive. Or, you could have one drive group (A), a RAID 5 level with a standby drive (spare), with an effective capacity of 3 drives with double the amount of fault tolerance. Or, you could have two drive groups (A, B), where drive group A is a RAID 0 with an effective capacity of 2 drives; drive group B is a RAID 0+1 with a standby drive (spare) and an effective capacity of 1 drive with double fault tolerance. Or, you could have one drive group (A): as a RAID 5, with a 4 drive effective capacity and fault tolerance; as a RAID 0+1, with a 2.5 drive effective capacity and fault tolerance; or as a RAID 0, with a 5 drive effective capacity.

Array Configuration: One Drive

Drives and Drive Groups	Possible RAID Levels	Effective Capacity
A:o	A:JBOD	A:1

Table 3 Array configuration – One drive

Array Configuration: Two Drives

Drives and Drive Groups	Possible RAID Levels	Effective Capacity
A:o	A:JBOD	A:1
B:o	B:JBOD	B:1
A:oo	A:0	A:2
	A:1	A:1*

Table 4 Array configuration – Two drives

Array Configuration: Three Drives

Drives and Drive Groups	Possible RAID Levels	Effective Capacity
A:o	A:JBOD	A:1

Array Configuration: Three Drives

Drives and Drive Groups	Possible RAID Levels	Effective Capacity
B:o	B:JBOD	B:1
C:o	C:JBOD	C:1
A:oo	A:0	A:2
O	A:1/SBY	A:1**
A:ooo	A:5	A:2
	A:0+1	A:1:5*
	A:0	A:3

Table 5 Array configuration – Three drives

Array Configuration: Four Drives

Drives and Drive Groups	Possible RAID Levels	Effective Capacity
A:o	A:JBOD	A:1
B:o	B:JBOD	B:1
C:o	C:JBOD	C:1
D:o	D:JBOD	D:1
A:oooo	A:5	A:3*
	A:0+1	A:2*
	A:0	A:4
A:ooo	A:5/SBY	A:2**
O	A:0+1/SBY	A:1:5**

Table 6 Array configuration – Four drives

Array Configuration: Five Drives

Drives and Drive Groups	Possible RAID Levels	Effective Capacity
A:o	A:JBOD	A:1

Array Configuration: Five Drives

Drives and Drive Groups	Possible RAID Levels	Effective Capacity
B:o	B:JBOD	B:1
C:o	C:JBOD	C:1
D:o	D:JBOD	D:1
E:o	E:JBOD	E:1
A:oooo	A:5/SBY	A:3**
O	B:0+1/SBY	A:2**
A:oo	A:0	A:2
B:oo	A:1/SBY	B:1**
O	B:0	B:2
	B:1/SBY	B:1**
A:ooooo	A:5	A:4*
	A:0+1	A:2.5*
	A:0	A:5
A:ooo	A:5	A:2*
B:oo	A:0+1	A:1.5
	B:0	B:2
	B:1	B:1*

Table 7 Array configuration – Five drives

Array Configuration: Six Drives

Drives and Drive Groups	Possible RAID Levels	Effective Capacity
A:o	A:JBOD	A:1
B:o	B:JBOD	B:1
C:o	C:JBOD	C:1
D:o	D:JBOD	D:1

Array Configuration: Six Drives

Drives and Drive Groups	Possible RAID Levels	Effective Capacity
E:o	E:JBOD	E:1
F:o	F:JBOD	F:1
A:oooooo	A:5	A:5*
	A:0+1	A:3*
A:ooooo	A:5/SBY	A:4**
O	A:0+1/SBY	A:2.5**
A:oooo	A:5	A:4*
B:oo	A:0+1	A:2
	B:0	B:2
	B:1	B:1*
A:ooo	A:5/SBY	A:2**
B:oo	A:0+1/SBY	A:1.5**
O	B:0	B:2
	B:1	B:1*
A:ooo	A:5	A:2*
B:ooo	B:5	B:2*
	A:& B: 5+0	A: & B: 4**
	A:0+1	A:1.5*
	B:0+1	B:1.5*
	A:& B: 0+1+0	A: & B: 3**

Table 8 Array configuration – Six drives

Array Configuration: Seven Drives

Drives and Drive Groups	Possible RAID Levels	Effective Capacity
A:o	A:JBOD	A:1

Array Configuration: Seven Drives

Drives and Drive Groups	Possible RAID Levels	Effective Capacity
B:o	B:JBOD	B:1
C:o	C:JBOD	C:1
D:o	D:JBOD	D:1
E:o	E:JBOD	E:1
F:o	F:JBOD	F:1
G:o	G:JBOD	G:1
A:ooooooo	A:5	A:6*
	A:0+1	A:3.5*
A:oooooo	A:5/SBY	A:5**
O	A:0+1/SBY	A:3**
A:oooo	A:5/SBY	A:3**
B:oo	A:0+1/SBY	A:2**
O	B:0	B:2
	B:1	B:1**
A:ooo	A:5/SBY	A:2**
B:ooo	B:5/SBY	A:2**
O	A: & B: 5+0/SBY	A: & B: 4***
	A: 0+1/SBY	A:1.5**
	B: 0+1/SBY	B: 1.5**
	A:& B: 0+1+0/SBY	A: & B: 3***

Table 9 Array configuration – Seven drives

Array Configuration: Eight Drives		
Drives and Drive Groups	**Possible RAID Levels**	**Effective Capacity**

Array Configuration: Eight Drives		
Drives and Drive Groups	**Possible RAID Levels**	**Effective Capacity**
A:o	A:JBOD	A:1
B:o	B:JBOD	B:1
C:o	C:JBOD	C:1
D:o	D:JBOD	D:1
E:o	E:JBOD	E:1
F:o	F:JBOD	F:1
G:o	G:JBOD	G:1
H:o	H:JBOD	H:1
I:o	I:JBOD	I:1
A:oooooooo	A:5	A:7*
	A:0+1	A:4*
	A:0	A:8
A:ooooooo	A:5/SBY	A:6**
	A:0+1/SBY	A:3.5**
O	A:0	A:7
A:ooooo	A:5/SBY	A:4**
B:oo	A:0+1/SBY	A:2.5**
O	B:0	B:2
	B:1/SBY	B:1**
A:ooo	A:5/SBY	A:2***
B:ooo	A:0+1/SBY	A:1.5***
o o	A:0	A:3
	A: & B: 5+0/SBY	A: & B: 4****
	A: 0+1+0/SBY	A: & B: 3****

Array Configuration: Eight Drives		
Drives and Drive Groups	**Possible RAID Levels**	**Effective Capacity**
	B:5/SBY	2***
	B: 0+1/SBY	B: 1.5***
A:oooo	A:5	A: 3*
B:oooo	B:5	B:3*
	A: & B: 5+0	A: & B: 6**
	A:0+1	A:2*
	B:0+1	B:2*
	A: & B: 0+1+0	A: & B: 4*
	A:0	B:4
A:oooo	A:5/SBY	A:3*
B:ooo	A:0+1/SBY	A:2*
O	B:5/SBY	B:2*
	B:0+1/SBY	B:1.5*
	A:0	A:4
	B:0	B:3

Table 10 Array configuration – Eight drives

Glossary

A

Active/Active

A synonym for Dual Active controllers. Under normal operating conditions, both controllers in a dual active controller configuration are actively processing I/O.

Active/Passive

A synonym for Hot Standby.

Active Termination

A type of terminator used in current SCSI channel setups, which utilizes an active Voltage regulator, thus closely matching cable impedance.

ANSI

American National Standards Institute, a standards organization that promotes and facilitates understanding among qualified members for the implementation of agreed upon code and signaling standards and conformity. ANSI represents the U.S. in the ISO (International Standards Organization).

Applet

Small application programs that are usually built into an operating system or a larger application program and designed to run from within another application. The built-in writing and drawing programs that come with Windows are sometimes called “applets.” Applications that run in a web browser are often Java-based applets.

Application Program

A program designed to perform a specific function directly for the user or, in some cases, for another application program. Examples of applications include word processors, database programs, Web browsers, development tools, drawing, paint and image editing programs, and communication programs.

Application Server

A centralized computer that holds and distributes application programs to users.

ASIC

Application-Specific Integrated Circuit, a chip created for a specific application.

Array

Multiple disk drives configured to behave as a single, independent disk drive. See also Disk Array.

Asynchronous Data Transfer

Data transfer not synchronized to a set timing interval. Asynchronous devices must wait for a signal from the receiving device after each byte of data.

Auto Swap

The exchange of a pre-installed replacement unit in a disk system for a defective one; the disk system performs the substitution without human intervention while continuing to carry out its normal functions (compare with Hot Swap).

Automatic Rebuild

The controllers provide automatic rebuild capabilities in the event of a physical disk drive failure. The controller performs a rebuild operation automatically when a disk drive fails and both of the following conditions are true: - A standby or hot spare disk drive of identical or larger size is found attached to the same controller; - • All system drives that are dependent on the failed disk drive are configured as a redundant array: RAID 1, RAID 3, RAID 5, or RAID 0+1. During the automatic rebuild process, system activity continues; however, system performance may degrade slightly.

Automatic Switchover

See Failover

B

Background Initialization

Where the initialization process of a disk array takes place in the background, allowing use of a disk array within seconds instead of several hours. Also known as Immediate RAID Availability.

BBU

Battery Backup Unit, provides a battery backup for data currently stored in the on-board cache memory during intermittent power loss to the controller. In the event of a power failure, the BBU can hold data in the cache for a certain amount of time. Once power is restored the data can be saved to a disk.

Benchmarks

A set of conditions or criteria against which a product or system is measured. Computer trade magazine laboratories frequently test and compare several new computers or computer devices against the same set of application programs, user interactions, and contextual situations. The total context against which all products are measured and compared is referred to as the benchmark. Programs can be specially designed to provide measurements for a particular operating system or application.

Berkeley RAID Levels

A family of disk array protection and mapping techniques described by Garth Gibson, Randy Katz, and David Patterson in papers written while they were performing research into I/O systems at the University of California at Berkeley. There are six Berkeley RAID levels, usually referred to as RAID Level 0 through RAID Level 5. See also RAID Levels.

BIOS

Basic Input/Output System, software that determines what a computer can do without accessing programs. The BIOS contains all the code required to control the keyboard, screen, drives, serial communications, and other functions. Usually the BIOS is built into a ROM chip installed on the motherboard so that the BIOS will always be available and not affected by disk failure. Sometimes the BIOS is recorded on a flash memory chip.

BIOS Configuration Utility

BIOS-based Configuration Utility, a utility program sequence used, upon powerup, for configuring various hardware elements in a system.

Booting (or Bootstrapping)

Loading operating system code and other basic software from a disk or other storage device to help a computer start.

Bridge RAID Controller

A device appearing as a single ID in a SAN, but which bridges to multiple devices, typically used to control external RAID subsystems (compare to Internal PCI-based RAID subsystems, see Internal RAID Controller). A bridge RAID controller is often referred to as an External RAID Controller.

Burst Data Rate

The speed at which a specific amount of data is sent or received in intermittent operations (compare to Sustained Data Transfer Rate).

Bus

A set of conductors that connect the functional units in a computer and are the channels through which data is transferred. There are several types of bus channels, including serial, parallel, PCI, ISA, EISA, and MCA. See also I/O Bus.

C

Cache

A temporary storage area for frequently accessed or recently accessed data. Cache is used to speed up data transfer to and from a disk. See also Caching.

Cache Flush

Refers to an operation where all unwritten blocks in a Write-Back Cache are written to the target disk. This operation is necessary before powering down the system.

Cache Line Size

Represents the size of the data "chunk" that will be read or written at one time, and is set in conjunction with stripe size. Under RAID EzAssist™ , the segment size (also known as Segment Size) should be based on the stripe size you selected. The default segment size for RAID controllers is 8K.

Caching

Allows data to be stored in a pre-designated area of a disk or RAM. Caching speeds up the operation of RAID systems, disk drives, computers and servers, or other peripheral devices. See also Cache.

CD-ROM

Compact Disk-Read Only Memory, a removable read-only storage device, similar to an audio compact laser disk, holding up to 640MB of data.

Channel

Any path used for the transfer of data and control of information between storage devices and a storage controller or I/O adapter. Also refers to one SCSI bus on a disk array controller. Each disk array controller provides at least one channel.

Cluster

A group of terminals or workstations attached to a common control unit or server, or a group of several servers, that share work and may be able to back each other up if one server fails.

Cold Swap

The physical exchange of a replacement unit in a storage system for a defective one. The exchange requires human intervention and power must be removed from the storage system in order to perform the exchange (compare to Hot Swap and Auto Swap).

Conservative Cache

An operating mode in which system drives configured with the write-back caching policy are treated as though they were configured for write-through operation and the cache is flushed.

Consistency Check

A process that verifies the integrity of redundant data. A consistency check on a RAID 1 or RAID 0+1 configuration (mirroring) checks if the data on drives and their mirrored pair are exactly the same. For RAID Level 3 or RAID Level 5, a consistency check calculates the parity from the data written on the disk and compares it to the written parity. A consistency check from utilities such as Global Array Manager®™ (GAM) or RAID EzAssist™ give the user the ability to have a discrepancy reported and corrected. See also Parity Check.

Controller

An adapter card, RAID controller, or other module that interprets and controls signals between a host and a peripheral device.

D

Data Transfer Rate

The amount of data per unit of time moved through a channel or I/O bus in the course of execution of an I/O load, usually expressed in MBps.

Device Driver

A software program that controls a particular type of device attached to a computer, such as a RAID subsystem, printer, display, CD-ROM, disk drive, etc.

Disk

A non-volatile, randomly addressable, re-writable data storage device, including rotating magnetic and optical disks as well as solid-state disks or other electronic storage elements.

Disk Array

A collection of disks from one or more commonly accessible disk systems. Disk arrays, also known as RAID, allow disk drives to be used together to improve fault tolerance,

performance, or both. Disk arrays are commonly used on servers and are becoming more popular on desktops and workstations. See also Array.

Disk Drive

A device for the electronic digital storage of information.

Disk Failure Detection

A RAID controller automatically detects SCSI disk failures. A monitoring process running on the controller checks, among other things, elapsed time on all commands issued to disks. A time-out causes the disk to be "reset" and the command to be retried. If the command times out again, the controller could take the disk "offline." DAC960 controllers also monitor SCSI bus parity errors and other potential problems. Any disk with too many errors will also be taken "offline." See also Offline.

Disk Media Error Management

Controllers transparently manage disk media errors. Disks are programmed to report errors, even ECC-recoverable errors. If ECC RAM is installed, the controller will correct ECC errors. When a disk reports a media error during a read, the controller reads the data from the mirror (RAID 1 or 0+1), or computes the data from the other blocks (RAID 3, RAID 5), and writes the data back to the disk that encountered the error. When a disk reports a media error during a write, the controller issues a "reassign" command to the disk, and writes the data out to a new location on the disk. Since the problem has been resolved, no error is reported to the system.

Disk System

A storage system capable of supporting only disks.

Drive Groups

A group of individual disk drives (preferably identical) that are logically tied to each other and are addressed as a single unit. In some cases this may be called a drive "pack" when referring to just the physical devices. All the physical devices in a drive group should have the same size; otherwise, each of the disks in the group will effectively have the capacity of the smallest member. The total size of the drive group will be the size of the smallest disk in the group multiplied by the number of disks in the group. For example, if you have 4 disks of 400MB each and 1 disk of 200MB in a pack, the effective capacity available for use is only 1000MB (5x200), not 1800MB.

Dual Active

A pair of components, such as storage controllers in a failure tolerant storage system, that share a task or set of tasks when both are functioning normally. When one component of the pair fails, the other takes the entire load. Dual active controllers (also called Active-Active controllers) are connected to the same set of devices and provide a combination of higher I/O performance and greater failure tolerance than a single controller.

E

Embedded Storage Controller

An intelligent storage controller that mounts in a host computer's housing and attaches directly to a host's memory bus with no intervening I/O adapter or I/O bus.

External RAID Controller

A RAID controller in its own enclosure, rather than incorporated into a PC or server. External RAID controllers are often referred to as a bridge RAID controller. SANArray FL, FF, FFx, and Pro FF2 controllers are external RAID controllers.

F

Failover

A mode of operation for failure tolerant systems in which a component has failed and a redundant component has assumed its functions.

Failover Port

A fibre channel port capable of assuming I/O requests for another, failed port on the loop. During normal operation, a failover port may be active or inactive. Failover ports assume the same loop ID and, optionally, the same node from the failed port.

Failure

A detectable physical change in hardware, requiring replacement of the component.

Fault Tolerance, Failure Tolerance

The ability of a system to continue to perform its function even when one of its components has failed. A fault tolerant system requires redundancy in disk drives, power supplies, adapters, controllers, and cabling. RAID controllers offer high levels of fault tolerance.

Fibre Channel

Technology for transmitting data between computer devices at a data rate of up to 2 Gbps (two billion bits per second), especially suited for connecting computer servers to shared storage devices and for interconnecting storage controllers and drives. Fibre Channel is expected to replace the Small System Computer Interface (SCSI) as the transmission interface between servers and clustered storage devices. It is also more flexible: devices can be as far as ten kilometers (about six miles) apart. The longer distance requires optical fiber as the physical medium; however, Fibre Channel also works using coaxial cable and ordinary telephone twisted pair wires.

Foreground Initialization

During the initialization process all data is cleared and zeros are written to the disks. With the size of system drives growing dramatically, a foreground initialization takes several hours to complete.

H

Hard Disk

A magnetically coated disk substrate that spins inside a disk drive and is used as the storage medium for digital data.

Hard Disk Drive (HDD)

The main data storage unit of a computer.

Host

Any computer system to which disks are attached and accessible for data storage and I/O.

Host I/O Bus

An I/O bus used to connect a host computer to storage systems or storage devices.

Hot Plug

The process of adding or removing a device from a bus while transactions involving other devices are occurring over the bus. See also PCI Hot Plug.

Hot Replacement of Disks

The design of all controllers allows for the replacement of failed hard disk drives without interruption of system service. In the event of a SCSI drive failure on a properly configured system (where the data redundancy features of the controller are used), the system generates a message to alert the system operator. When a replacement drive becomes available, the system operator can remove the failed disk drive, install a new disk drive, and instruct the controller to “rebuild” the data on the new drive, all without interrupting system operations. Once the rebuild is complete, the controller will be brought back into a fault tolerant state. See also Hot Swap.

Hot Spare

A physical disk drive not part of a system drive that the controller can use to automatically rebuild a critical system drive. The hot spare drive must have at least as much capacity as the largest disk drive in the array or the rebuild may not start.

Hot Standby

A redundant component in a fault tolerant storage system that has power applied and is ready to operate, but which does not perform its task as long as the primary component for which it is standing by is functioning properly. See also Hot Replacement of Disks and Hot Spare.

Hot Swap

The exchange of a replacement unit in a storage system for a defective unit. The exchange requires human intervention,

but the system can continue to perform its normal functions (compare with Auto Swap, Cold Swap, and Warm Swap).

I

Internal RAID Controller

A controller circuit board that resides inside a computer or server. An internal RAID controller resides on a bus, such as the PCI bus. Examples of internal RAID controllers include the AcceleRAID™ and eXtremeRAID® families.

I/O

Input/Output, the transmission of information between an external source and the computer.

I/O Bottleneck

Any resource in the I/O path whose performance limits the performance of a storage or I/O system.

I/O Bus

Any path used for the transfer of data and control information between I/O adapters and storage controllers or storage devices. See also Bus.

I/O Intensive

An application whose performance depends strongly on the performance of the I/O system.

I/O Per Second

Number of I/O transactions per second.

J

Java

A programming language expressly designed for use in the distributed environment of the Internet and enforces a completely object-oriented view of programming. Java can be used to create complete applications that may run on a single computer or be distributed among servers and clients in a network. It can also be used to build small application modules or applets for use as part of a Web page.

JBOD

Just A Bunch of Disks (Drives), a number of disk drives, usually in an enclosure. JBOD implies that the disks do not use RAID technology and function independently.

L

Latency

The time between the making of an I/O request and completion of the request's execution.

Short for rotational latency, the time between the completion of a seek and the instant of arrival of the first block of data to be transferred at the disk's read/write head.

Logical Drive

The logical devices that are presented to the operating system. System drives are presented as available disk drives, each with a capacity specified by the RAID controller.

LUN

Logical Unit Number, a SCSI representation of a system drive on a given channel and target ID. This may be a single device or an array of devices configured to behave as a single device.

LUN Mapping

A method whereby a LUN ID is assigned to a system drive, allowing a LUN to be made accessible through specific controllers and ports using system drive affinity.

M

Mirrored Cache

A cache memory that has duplicate data from another controller. In the event of failure of the original controller, the second controller can take the cached data and place it on the disk array.

Mirrored Hard Drive

Two hard drives the computer sees as one unit. Information is stored simultaneously on each drive. If one hard disk drive fails, the other contains all of the cached data and the system can continue operating.

Mirroring

Refers to the complete duplication of data on one disk drive to another disk drive, this duplication occurs simultaneously with each write operation: each disk will be the mirror image of the other (also known as RAID Level 1).

MORE

Mylex Online RAID Expansion, an advanced configuration mode that allows expansion of any unconfigured or hot spare drive into the expandable drive group while the controller is online with the host. For example, a system using a five-disk-drive RAID set can add another disk drive to create a six-disk-drive RAID set. The MORE operation can be performed on all RAID levels except JBOD. During the RAID set expansion process, which includes re-striping data from the old (smaller) RAID set to the new (expanded) RAID set, the controller continues to service host I/O requests.

NVRAM

Non-Volatile Random Access Memory, a memory unit equipped with a battery so that the data stays intact even after the main power had been switched off.

P

Parity

A method of providing complete data redundancy while requiring only a fraction of the storage capacity of mirroring. The data and parity blocks are divided between the disk drives in such a way that if any single disk drive is removed or fails, the data on it can be reconstructed using the data on the remaining disk drives. The parity data may exist on only one disk drive or be distributed between all disk drives in a RAID group.

Parity Check

a function used to verify the integrity of data on a system drive. It verifies that mirror or parity information matches the stored data on the redundant arrays. If the parity block information is inconsistent with the data blocks, the controller corrects the inconsistencies.

Partitioning

When the full usable storage capacity of a disk or array of disks appears to an operating environment in the form of several virtual disks whose entire capacity approximates that of the underlying disk or array.

PCI

Peripheral Component Interconnect, a standardized architecture that provides a high-speed data path between peripherals and the CPU. PCI is a high-performance, backplane interface, expansion slot architecture found on PCs, Macintosh, and UNIX® workstations. PCI cards are portable across hardware platforms with the help of various software drivers.

Physical Device

Any device connected to some kind of hardware. For example, SCSI disk, fibre disk, network disk, RAM disk, etc.

Physical Disk Drive

A single hard disk drive. Each physical disk drive is assigned a unique identification address.

Q

Queue

A line of things, commands, or data waiting to be handled, usually in sequential order starting at the beginning or top of the line or sequence.

R

RAID

Redundant Array of Independent Disks, a collection of two or more disks working together in an array. RAID controllers implement this technology to connect up to 15 SCSI devices per channel. The different forms of RAID implementation are known as “RAID levels.” See also Berkeley RAID

Levels, Disk Array, and RAID Levels. The system manager or integrator selects the appropriate RAID level for a system. This decision will be based on which of the following are to be emphasized: - Disk Capacity - Data Availability (redundancy or fault tolerance) -Disk Performance

RAID Controller

Low cost RAID controllers that use SCSI channels on the motherboard.

RAID Levels

Disk array controllers support four RAID Advisory Board approved (RAID 0, RAID 1, RAID 3, and RAID 5), two special (RAID 0+1, and JBOD), and three spanned (RAID 10, 30, and 50) RAID levels. All DAC960, AcceleRAID, and eXtremeRAID series controllers support these RAID levels. See also Berkeley RAID Levels. Level 0: Provides block "striping" across multiple drives, yielding higher performance than is possible with individual drives. This level does not provide any redundancy. Level 1: Drives are paired and mirrored. All data is 100 percent duplicated on a drive of equivalent size. Level 3: Data is "striped" across several physical drives. Maintains parity information which can be used for data recovery. Level 5: Data is "striped" across several physical drives. For data redundancy, drives are encoded with rotated XOR redundancy. Level 0+1: Combines RAID 0 striping and RAID 1 mirroring. This level provides redundancy through mirroring. JBOD: Sometimes referred to as "Just a Bunch of Drives." Each drive is operated independently like a normal disk controller, or drives may be spanned and seen as a single drive. This level does not provide data redundancy. Level 10: Combines RAID 0 striping and RAID 1 mirroring spanned across multiple drive groups (super drive group). This level provides redundancy through mirroring and better performance than Level 1 alone. Level 30: Data is "striped" across multiple drive groups (super drive group). Maintains parity information, which can be used for data recovery. Level 50: Data is "striped" across multiple drive groups (super drive group). For data redundancy, drives are encoded with rotated XOR redundancy.**Note:** The host operating system drivers and software utilities remain unchanged regardless of the level of RAID installed. The controller makes the physical configuration and RAID level implementation.

RAID Migration

A feature in RAID subsystems that allows for changing a RAID level to another level without powering down the system.

Recovery

The process of reconstructing data from a failed disk using data from other drives.

Redundancy

The inclusion of extra components of a given type in a system (beyond those the system requires to carry out its functions).

S

SAF-TE

SCSI Accessed Fault-Tolerant Enclosure, an "open" specification designed to provide a comprehensive standardized method to monitor and report status information on the condition of disk drives, power supplies, and cooling systems used in high availability LAN servers and storage subsystems. The specification is independent of hardware I/O cabling, operating systems, server platforms, and RAID implementation because the enclosure itself is treated as simply another device on the SCSI bus. Many other leading server, storage, and RAID controller manufacturers worldwide have endorsed the SAF-TE specification. Products compliant with the SAF-TE specification will reduce the cost of managing storage enclosures, making it easier for a LAN administrator to obtain base-level fault-tolerant alert notification and status information.

SCSI

Small Computer System Interface, a technological standard that defines connections between computers and peripheral devices.

SCSI Adapters

Storage controllers for managing SCSI devices.

SCSI Drive

A disk drive equipped with a SCSI interface (sometimes referred to as a SCSI Disk). Each disk drive will be assigned a SCSI address (or SCSI ID), which is a number from 0 to 7 (0 to 15 under wide or Ultra SCSI). The SCSI address uniquely identifies the drive on the SCSI bus or channel.

Sector

The unit in which data is physically stored and protected against errors on a fixed-block architecture disk.

Sequential I/O

A type of read and write operation where entire blocks of data are accessed one after another in sequence, as opposed to randomly.

Server

A computer program that provides and manages services to other computer programs on the same or other computers. The computer that a server program runs in is also frequently referred to as a server.

SES

SCSI Enclosure Services, a standard for SCSI access to services within an enclosure containing one or more SCSI devices. For disk drives, power supplies, cooling elements, and temperature sensors, the actions performed are the same as for SAF-TE monitoring. If a UPS is connected to any SES-monitored enclosures, and an AC failure or two minute warning is reported, conservative cache is enabled and all system drives are switched to write-through cache. Primarily used in fibre enclosures.

Session

The period of time between any two consecutive system shutdowns; system shutdown may be either a power off/on, or a hardware reset.

Standard Disk Drive

This term refers to a hard disk drive with SCSI, IDE, or other interface, attached to the host system through a standard disk controller.

Standby Replacement of Disks

See also Hot Spare. One of the most important features the RAID controller provides to achieve automatic, non-stop service with a high degree of fault-tolerance. The controller automatically carries out the rebuild operation when a SCSI disk drive fails and both of the following conditions are true: • A “standby” SCSI disk drive of identical size is found attached to the same controller; • All of the system drives that are dependent on the failed disk are redundant system drives, e.g., RAID 1, RAID 3, RAID 5, and RAID 0+1. Note: The standby rebuild will only happen on the SAME DAC960 controller, never across DAC960 controllers. During the automatic rebuild process, system activity continues as normal. System performance may degrade slightly during the rebuild process. To use the standby rebuild feature, you should always maintain a standby SCSI disk in your system. When a disk fails, the standby disk will automatically replace the failed drive and the data will be rebuilt. The system administrator can disconnect and remove the bad disk and replace it with a new disk. The administrator can then make this new disk a standby. The standby replacement table has a limit of 8 automatic replacements in any session (from power-on/reset to the next power-off/reset). When the limit of 8 is reached and a disk failure occurs, the standby replacement will occur but will not be recorded in the replacement table. To clear the “standby replacement” table, reboot the system from a DOS bootable floppy, run the configuration utility and select the option ’view/update configuration’ from the main menu. A red box labeled ’Drive Remap List’ will be displayed. Selecting the box will allow you to continue. You should save the configuration without making any changes, and exit the configuration utility. This will clear the replacement table. You may now proceed to boot your system and continue normal operations. In normal use, the replacement table limit of 8 should not cause any problems. Assuming that a disk fails about once a year (drives we support generally come with a 5-year warranty), the system would run continuously for a minimum of 8 years before the table would need to be cleared.

Storage Device

A collective term for disks, tape transports, and other mechanisms capable of non-volatile data storage.

Stripe Order

The order in which SCSI disk drives appear within a drive group. This order must be maintained, and is critical to the controller's ability to "rebuild" failed drives.

Stripe Size

The size, in kilobytes (1024 bytes) of a single I/O operation. A stripe of data (data residing in actual physical disk sectors, which are logically ordered first to last) is divided over all disks in the drive group.

Stripe Width

The number of striped SCSI drives within a drive group.

Striping

The storing of a sequential block of incoming data across multiple SCSI drives in a group. For example, if there are 3 SCSI drives in a group, the data will be separated into blocks and block 1 of the data will be stored on SCSI drive 1, block 2 on SCSI drive 2, block 3 on SCSI drive 3, block 4 on SCSI drive 1, block 5 on SCSI drive 2, and so on. This storage method increases the disk system throughput by ensuring a balanced load among all drives.

Sub-System Storage

A collection of disks providing data storage space to a system user.

Sustained Data Transfer Rate

A rate of data transfer defined for continuous operation at a maximum speed level.

Synchronous Data Transfer

Data transmission synchronized to a defined time interval, and is faster than asynchronous SCSI because there is no wait for acknowledgement of each byte from the receiving device (up to 20MHz).

T

Target ID

The SCSI ID of a device attached to a controller. Each SCSI channel can have up to 15 attached SCSI devices (target ID from 0 to 6 and 8 to 15).

TCP/IP

Transmission Control Protocol/Internet Protocol, the basic communication language or protocol of the Internet. It can also be used as a communications protocol in intranets and extranets. When set up with direct access to the Internet, a computer is provided with a copy of the TCP/IP program just as every other computer that you may send messages to or get information from also has a copy of TCP/IP.

Throughput

The number of I/O requests satisfied per unit of time (usually per second).

Transfer Rate

The rate at which data moves between the host computer and storage, input, or output devices, usually expressed as a number of characters per second.

W

Warm Swap

The exchange of a defective disk with a working disk while power is maintained to the system, but I/O operations have ceased.

Web

A number of computers, servers, or networks linked together in a matrix.

Write-Back Cache

A caching strategy whereby write operations result in a completion signal being sent to the host operating system as soon as the cache (not the disk drive) receives the data to be written. The target disk drive will receive the data at a more appropriate time in order to increase controller performance. An optional cache battery backup can be used to protect against data loss as a result of a power failure or system crash.

Write-Through Cache

A caching strategy whereby data is written to the SCSI drive before a completion status is returned to the host operating system. This caching strategy is considered more secure, since a power failure will be less likely to cause loss of data. However, a write through cache results in a slightly lower performance.

Index

F

G

H

I

J

L

M

N

O