

# Table of Contents

<b>1. Introduction.....</b>	<b>7</b>
1.1. Overview .....	7
1.2. Inter-Server Communication .....	7
1.2.1. Multicast.....	7
1.3. License Manager .....	7
1.3.1. Access License Limits.....	7
1.3.2. Feature License .....	8
<b>2. Quick Tour .....</b>	<b>11</b>
2.1. The Work Area .....	11
2.1.1. Tool Bar .....	12
2.1.2. Enclosure Section .....	13
2.1.3. Array and Logical Drive/Snapshot Section.....	18
2.1.4. Server Sidebar and Top Section .....	20
2.2. How to Use this Document .....	22
<b>3. Embedded StorView Setup .....</b>	<b>23</b>
3.1. Embedded Network Settings .....	23
3.1.1. Configuring Network Settings using the Wizard .....	23
3.1.2. Configuring Network Settings using the CLI.....	25
3.1.3. Configuring Network Settings using StorView .....	26
3.1.3.1. Using Dynamic IP (DHCP) .....	26
3.1.3.2. Using Static IP .....	26
3.1.4. Getting a New IP Address.....	27
<b>4. Getting Started .....</b>	<b>28</b>
4.1. Starting StorView.....	28
4.2. Upgrading the Access License .....	29
4.3. E-MAIL.....	30
4.3.1. Configuring E-Mail Notices .....	30
4.3.2. Deleting an E-Mail Addressee .....	31
4.4. SNMP .....	31
4.4.1. Configuring SNMP Traps.....	31
4.4.2. Deleting an SNMP Server .....	32
4.5. SYSLOG .....	32
4.5.1. Gathering Logs .....	32
4.5.2. Syslog Daemon .....	32
4.5.3. Operational View .....	32
4.5.4. Configuring Syslogd Server.....	33
4.5.5. Delete a Syslogd Server.....	33
4.6. Changing the Password.....	34
4.7. Monitoring Settings .....	34
4.7.1. Additional Monitoring Servers.....	35
4.7.2. Remove Monitored StorView Server IP .....	35
<b>5. Configuring a Storage Solution .....</b>	<b>36</b>
5.1. Creating Disk Arrays .....	36
5.1.1. RAID Levels.....	36
5.1.2. Terminology.....	37
5.1.3. Optimization and Drive Selection for RAID 5 and 6 .....	38
5.2. Create the Array .....	39
5.2.1. Configuring Array Writeback Cache.....	44

5.2.2.	Chunk Size Details .....	44
5.2.3.	Performance .....	45
5.3.	Initializing the Array.....	45
5.3.1.	Pause/Resume the Initialization .....	46
5.3.1.1.	Pause Initialization .....	46
5.3.1.2.	Resume Initialization .....	46
5.4.	Adding Hot Spare Drives .....	47
5.4.1.	Assigning a Global Spare.....	47
5.4.2.	Assigning a Dedicated Spare .....	48
5.4.3.	Removing a Spare .....	48
5.4.4.	Auto Spare.....	49
5.5.	Create the Logical Drive .....	50
5.6.	Saving the Configuration .....	51
5.6.1.	Saving the Configuration.....	52
<b>6.</b>	<b>Environments .....</b>	<b>54</b>
6.1.	Overview .....	54
6.2.	Controller Environmental.....	54
6.2.1.	Status .....	54
6.2.2.	Hardware/Firmware.....	55
6.2.3.	Configuration.....	55
6.2.4.	Operations .....	55
6.2.4.1.	Updating the Controller Firmware .....	55
6.2.4.2.	Update Expansion Firmware .....	56
6.2.4.3.	Determining the Version of Firmware on Expansion Modules .....	57
6.2.4.4.	Clearing the Controller Logs.....	58
6.2.4.5.	Exporting the Controller Logs.....	59
<b>7.</b>	<b>Controller Advanced Settings .....</b>	<b>60</b>
7.1.	Overview .....	60
7.2.	Advanced Settings.....	60
7.2.1.	Identity .....	61
7.2.2.	Fault Tolerance .....	62
7.2.3.	Host Ports .....	63
7.2.3.1.	RAID Controller Systems .....	63
7.3.	Advanced Power Management.....	64
7.3.1.	Drive Power Management (APM Level 1) .....	64
7.3.2.	APM Characteristics.....	65
7.4.	Advanced Performance Options.....	66
7.4.1.	Performances Options .....	66
7.4.1.1.	Synchronize Cache Writes to Disk .....	66
7.4.1.2.	Target Command Thread Balance .....	66
7.4.1.3.	Sequential Write Optimization .....	67
7.4.2.	Overload Management.....	67
<b>8.</b>	<b>Managing the Storage Solution .....</b>	<b>68</b>
8.1.	Overview .....	68
8.2.	Advanced Array Functions .....	69
8.2.1.	Deleting an Array .....	69
8.2.2.	Modifying Arrays .....	69
8.2.3.	Verify Parity .....	70
8.2.4.	Identifying Drive Members.....	71
8.2.5.	Rebuilding an Array .....	71
8.2.6.	Expanding an Array .....	72
8.2.7.	Trust an Array.....	73
8.3.	Advanced Power Management.....	74

8.3.1. Array Power Management (APM Level 2) .....	74
8.3.2. APM Characteristics .....	75
8.4. Restoring and Clearing the Configuration .....	76
8.4.1. Restoring the Configuration .....	76
8.4.2. Clearing the Configuration .....	77
8.5. Advanced Drive Options .....	78
8.5.1. Accessing the Drive Information Window .....	78
8.5.2. Locate Drive .....	79
8.5.3. Fail Drive .....	79
8.6. Advanced Logical Drive Functions .....	80
8.6.1. Viewing Unassigned Free Space .....	80
8.6.2. Expanding a Logical Drive .....	80
8.6.3. Deleting a Logical Drive .....	81
<b>9. SAN LUN Mapping .....</b>	<b>82</b>
9.1. Overview .....	82
9.2. Terminology .....	82
9.3. Accessing SAN LUN Mapping .....	83
9.4. Overview: SAN LUN Mapping Window .....	83
9.4.1. SELECT YOUR HBA and NAME YOUR HBA PORT Section .....	83
9.4.2. ADD MAPPINGS Section .....	84
9.4.3. MAPPINGS Section .....	84
9.5. Creating a SAN LUN Mapping .....	85
9.6. Deleting a SAN LUN Mapping .....	87
9.7. Modifying a SAN LUN Mapping .....	88
9.7.1. Re-Create the Mapping .....	88
<b>10. Using Snapshot .....</b>	<b>90</b>
10.1. Introduction .....	90
10.2. Enabling Snapshot .....	91
10.3. Performing a Snapshot .....	94
10.4. Mapping a Snapshot .....	96
10.5. Deleting a Snapshot .....	97
10.6. Performing a Snapback .....	99
10.7. Disabling Snapshot .....	101
<b>11. Failover, Performance &amp; Additional Functions .....</b>	<b>102</b>
11.1. Overview .....	102
11.2. How StorView Server Failover Works .....	102
11.3. StorView Performance Optimization .....	102
11.3.1. Execution Throttle .....	102
11.3.2. Scatter/Gather .....	102
11.4. Additional StorView Functions .....	103
11.4.1. About .....	103
11.4.2. Take Control Monitoring .....	103
11.4.3. Rescan .....	104
<b>12. Support &amp; Updates .....</b>	<b>105</b>
12.1. Tech Support .....	105
12.2. Updating Embedded StorView .....	106
<b>13. Event Logs .....</b>	<b>107</b>
13.1. Overview .....	107
13.2. Accessing and Navigating the StorView Event Log .....	108
13.3. Exporting the StorView Event Log .....	109
13.4. Clearing the StorView Event Log .....	111

13.5. Operating System Event Log .....	111
13.6. List of Events.....	112
13.6.1. Advanced Power Management Events .....	112
13.6.2. Alarm Events .....	113
13.6.3. Controller Events .....	114
13.6.4. Controller Port Events .....	125
13.6.5. Drive and Array Events .....	127
13.6.6. Enclosure Events.....	132
13.6.7. Persistent Reservation Events .....	136
13.6.8. Server Events.....	137
13.6.9. SnapShot Events .....	140
13.7. Failed Drives Codes .....	143
13.8. SAS Interface Error Codes.....	143
13.9. SAS Discovery Error Codes .....	144
<b>14. Statistics .....</b>	<b>145</b>
14.1. Overview .....	145
14.2. Access Statistics .....	146
14.3. Command Size - Alignment Statistics .....	147
14.4. Read-Ahead Statistics .....	148
14.5. Command Cluster Statistics .....	149
<b>15. Optimizing Write Performance .....</b>	<b>150</b>
15.1. Introduction.....	150
15.2. Sequential Access .....	151
15.2.1. Number of Outstanding Commands .....	151
15.3. Access Size.....	151
15.3.1. Access Alignment .....	151
15.4. Sub-Array.....	152
15.5. Multiple Drive Failures.....	152
15.5.1. Faster Rebuild .....	152
15.6. Summary.....	152
<b>16. Troubleshooting .....</b>	<b>153</b>
16.1. Problems You May Encounter .....	153
<b>17. More on SnapShot.....</b>	<b>156</b>
17.1. Overview .....	156
17.2. Terminology .....	156
17.3. Overwrite Data Area .....	157
17.4. Overwrite Data Area Location.....	157
17.5. Selecting the ODA Stripe Size .....	157
17.6. Overwrite Data Area Storage Capacity .....	158
17.7. ODA Assignment .....	158
17.8. Snapshots and Caching.....	158
17.8.1. Application Caches.....	158
17.8.2. Operating System Cache .....	158
17.8.3. RAID Controller Cache .....	158
17.9. Virtual LUNs .....	158
17.9.1. Virtual LUN Numbering.....	159
<b>18. Contacting Customer Support .....</b>	<b>160</b>
18.1. LaCie Technical Support Contacts .....	160
<b>19. Warranty Information .....</b>	<b>161</b>

## Notices

The information in this document is subject to change without notice.

While every effort has been made to ensure that all information in this document is accurate, the Authors accept no liability for any errors that may arise.

No part of this document may be transmitted or copied in any form, or by any means, for any purpose, without the written permission of the Authors.

Issue 4.0 January 30, 2009 Part No. 0094089-04A

## Acknowledgments

All names, brands, products or services are trademarks or registered trademarks of their respective companies.

## What is in this guide

This user guide gives you step-by-step instructions on how to setup and use the StorView Storage Management RAID Module software. This manual supports the RAID and Expansion enclosure storage systems listed in the Related Documentation section.

## Who should use this guide

This user guide assumes that you have a working knowledge of storage appliance products. If you do not have these skills, or are not confident with the instructions in this guide, do not proceed with the installation.

---

# 1. Introduction

## 1.1. Overview

StorView® Storage Management software is a full-featured graphical HTML-based software suite designed to configure, manage, and monitor storage subsystems. This user guide supports the RAID Module of StorView.

The RAID Module provides an extensive set of configuration and management tools. StorView's RAID Module is available in two versions: host-based and embedded. The host-based version is installed on the host computer system, while the embedded version is located in the RAID Controller firmware.

StorView's server component discovers storage solutions, manages and distributes message logs, and communicates with other server components installed on the same local network and external subnet networks. StorView has an HTML-based front end, accessed with a web browser, and provides the interface to the end user.

StorView incorporates web server software as part of the installation, Apache 2.0, which provides the interface between the server component and HTML interface. During installation the web server is automatically configured and requires no further management. The installation of the web server software is self contained and will not conflict with other web server software currently installed on your system. The license agreement is found in the preface section of the StorView Storage Management Software Installation Guide.

---

## 1.2. Inter-Server Communication

### 1.2.1. Multicast

StorView's server component uses multicasting technology to provide inter-server communication with other servers when the Global Access license is installed. During the server's initial start-up, it performs a multicast registration using the default multicast IP address of 225.0.0.225 on port 9191. Once registration is complete, the server is able to receive all packets sent to the multicast address.

All packets sent to the multicast address will remain in the local network, unless an explicit server IP address outside the subnet is added in the "Inter-Server Communication" Explicit StorView Server IPs Preference Settings. The inter-server communication abilities provide StorView with remote monitoring of other installations of StorView and their monitored storage systems.

StorView has the ability to communicate with any StorView installation on the local network. These other StorView server's are displayed on the Main screen and are listed under the "Other Servers"

section. They display the IP address, name, and an overall status of that server's monitored storage solution. They indicate the status of a monitored server storage solution, by the server icon changing to one of a few different states, refer to section [2.1.4. Server Sidebar and Top Section](#).

Each server sends a "check-in" packet in 10 second intervals. Once an initial check-in packet is received, all StorView servers will know the existence of the other servers. If a server fails to send three check-in packets, the other servers will mark that server as "missing." This is indicated by a white "Server" icon displayed on the Main screen under the "Other Servers" section.

When the server service that "owns" the monitored storage solution is down for any reason and three check-in packets are not received, the monitoring will automatically be transferred to another StorView server.

---

## 1.3. License Manager

StorView has two different licenses, an Access License and the Feature License.

---

### 1.3.1. Access License Limits

The StorView Access License has different limits for each of the two RAID module versions of StorView. The host-based version has Local Access and Global Access, and the embedded version has Remote Access and Global Access. Refer to section [4.2. Upgrading the Access License](#).

The Local Access is the default license of the host-based software. It provides the basic control, management and monitoring of the locally attached storage solution.

The Remote Access is the default license included with the embedded version. It provides the same functions as the Local Access license except it also provides support for server failover and remote login. Remote Access does not support features such as E-Mail notices, SYSLOG, and SNMP, however it does offer an upgrade path to the Global Access license.

The Global Access license is offered through a licensing program. Global Access provides the capabilities of the Local Access plus remote functions like login, configuration, monitoring, and alert notifications via E-Mail, SYSLOG and SNMP. With Global's remote login and management, the user

can focus or log-in to a different storage solution that is not locally attached to the host, this allows you to manage and monitor other remote solutions from just one location.

### 1.3.2. Feature License

The Feature License enables the full use of the SAN LUN Mapping feature and/or the Snapshot feature that are incorporated into the RAID Controller firmware. To access the License Information, click the Controller icon on the Main screen and then click the LICENSE INFORMATION button at the bottom of the Controller Information window.

The RAID Controller is provided with a 15 day grace period license for users that upgrade from a previous version and had a valid configuration. This time period allows the user to continue operating while they contact their provider to obtain a free unlimited license as part of the upgrade.

For new users, you may purchase a full unlimited license. During the grace period you will receive a daily reminder of the number of days remaining on the grace period in the SAN LUN Mapping window and through an event.

If, for some reason you allow the grace period to expire, your SAN LUN Mappings will be retained and protected until you enter a valid license key which will restore access to the mapped drives.

A limited Snapshot license is also included with the RAID Controller which allows for one Overwrite Data Area (ODA) to be created and four Snapshots per logical drive.

You can purchase a full Snapshot license which provides up to 256 Overwrite Data Areas (ODAs) to be created and up to 64 Snapshots for each of the logical drives until a maximum of 512 snapshots are created. You can setup and configure a variety of combinations depending upon your storage needs.

A licensing upgrade path to the maximum full Snapshot feature license is provided from the License Information screen by submitting a license key for a full license or by attempting to exceed the limited license, see [Fig. 06](#).

If you intend to use the Snapshot feature and choose to have access to the ODA logical drive from your operating system or VSS, you must SAN LUN Map the logical drive. Therefore, you will need a combined Snapshot license and a SAN LUN Mapping license.

To upgrade, you will need to provide your Configuration WWN to obtain a license key. The Configuration WWN is displayed in the License Information window.

continued on the next page >>

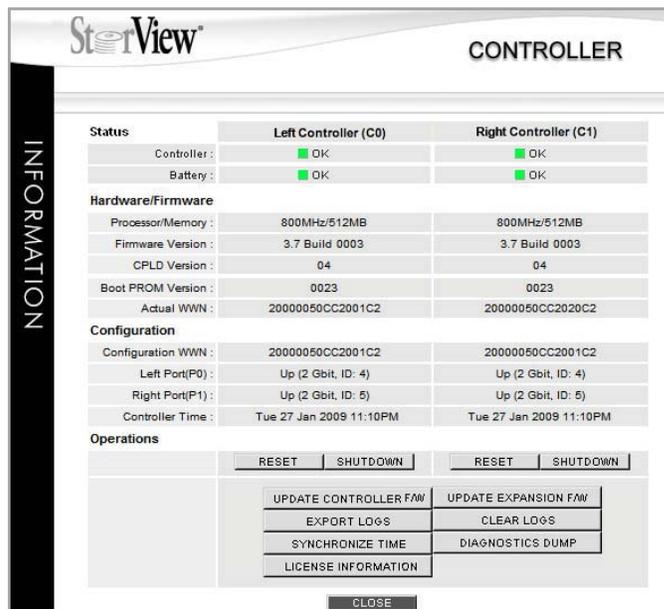


Fig. 01 - Controller Information Screen

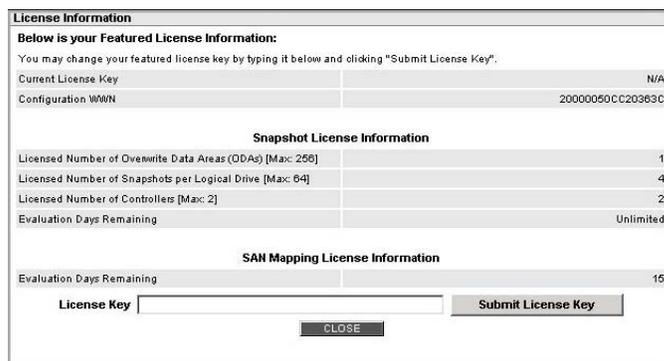


Fig. 02 - License Information Screen - 15 Day Grace Period License

The following illustration provides an explanation of the fields of the licensing Information window.

The first two items display the License Key, if a license key has been entered and accepted, and the Configuration WWN (used to create the license key). It is recommended to make a written record of the License Key and Configuration WWN.

Snapshot License Information section displays two possible license options, both will indicate "unlimited." One set of values represents the default limited license which provides for one ODA, four snapshots and the number of controllers that can be snapshotted (1 or 2). The other set of values represents a full unlimited license which provides for 256 ODAs with 64 snapshots per logical drive up to a maximum of a total of 512 snapshots and the number of controllers that can be snapshotted.

In the SAN LUN Mapping License Information section displays the number of days remaining on the 15 day grace period license. An event will occur each day redundantly warning about the remaining number of days. For those who purchased a full license, or those who upgraded, you will see "Unlimited."

To activate a feature license:

Enter the license key and click the Submit License Key button. The current license key will be displayed (Fig. 05).

[continued on the next page >>](#)

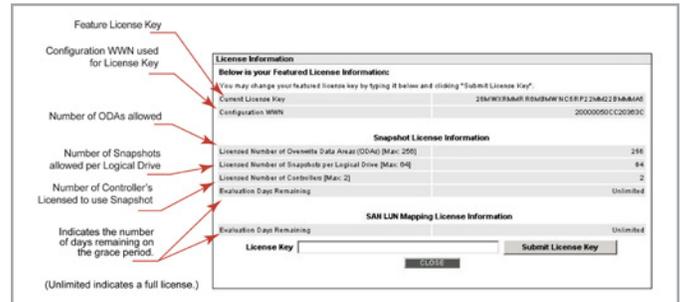


Fig. 03 - License Information Window Explained

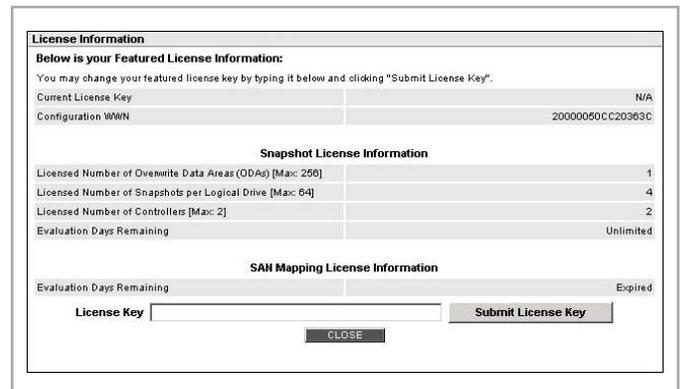


Fig. 04 - Basic Limited License - Expired

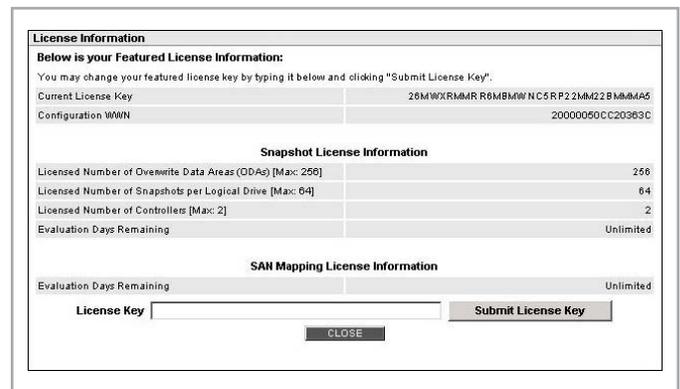


Fig. 05 - License Upgraded - Full License

If you attempt to exceed the limited Snapshot license, you will receive the message in [Fig. 06](#).

You will have the option to upgrade the Limited Snapshot license from that screen. Click the "UPGRADE LICENSE" button and enter a new license key to enable the maximum license. You may also upgrade the license at any time by clicking the Controller icon and then the LICENSE INFORMATION button (See [Fig. 06](#)).

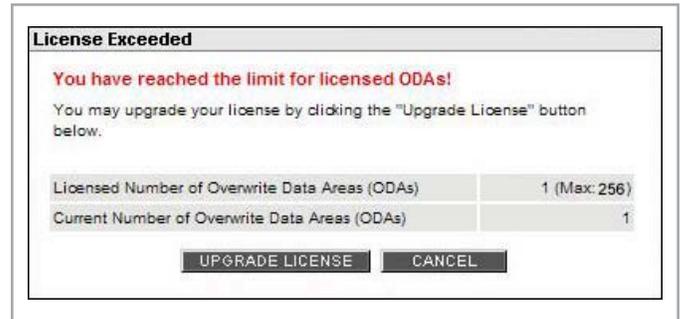


Fig. 06 - License Exceeded Message

## 2. Quick Tour

### 2.1. The Work Area

The StorView work area consists of the Main screen and its dialog windows for performing specific functions. On the Main screen you will find the Toolbar, sections for Arrays and Logical Drives, an Enclosure section and a Server section. Through the use of interactive components, animated icons, and mouse over information an intuitive, easy-to-use approach is provided to manage your storage solution.

Primary configuration functions include creating disk arrays, logical drives, LUN mapping and assignment of spare drives.

You also have access to advanced features that allow for array and logical drive expansion, optimizing controller parameters, rebuilding arrays, managing E-Mail notices of events, SYSLOG, and SNMP traps, reviewing event logs, and analyzing system statistics.

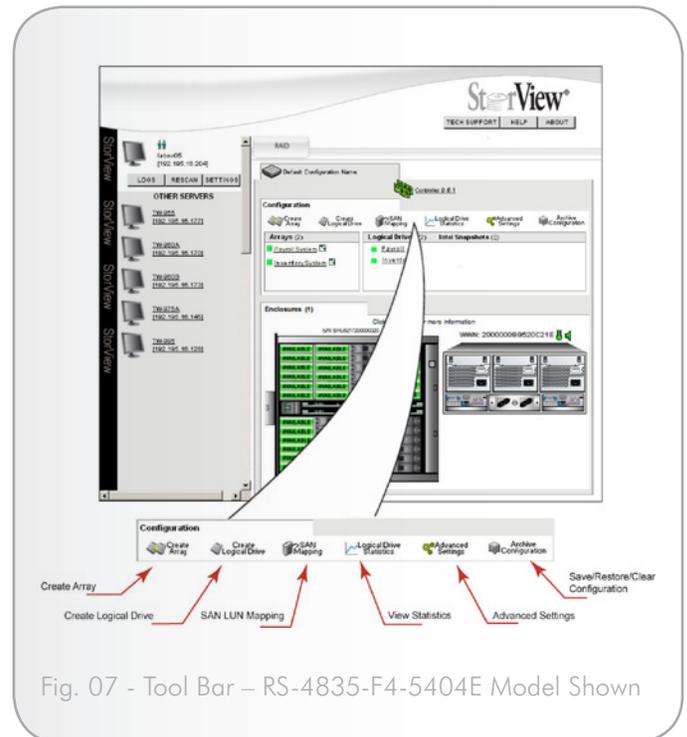


Fig. 07 - Tool Bar – RS-4835-F4-5404E Model Shown

### 2.1.1. Tool Bar

These icon buttons appear at the top of the work area and define the tool bar section of the work area. They provide a one click access to primary functions.

---

**NOTE:** Throughout the interface, holding the mouse pointer over an icon will display a pop-up window with information specific to the object.

---

Table 01 - Tool Bar Buttons

Button	Description
Create Array	This button will open the Create Array window allowing the user to create new disk arrays.
Create Logical Drive	This button will open the Create Logical Drive window allowing the user to configure new logical drives.
SAN Mapping	This button will open the SAN LUN Mapping panel which allows the user to further customize logical drive availability.
Logical Drive Statistics	This button opens the Statistics window.
Advanced Settings	This button opens a window from which you may change controller parameters.
Archive Configuration	This button will open a window from which you may choose to save, restore, or clear the configuration. Note when clearing a configuration this will delete all arrays and logical drives, as well as the data on those logical drives.

---

### 2.1.2. Enclosure Section

This section displays the front and rear views of the RAID enclosures and expansion enclosures.

continued on the next page >>

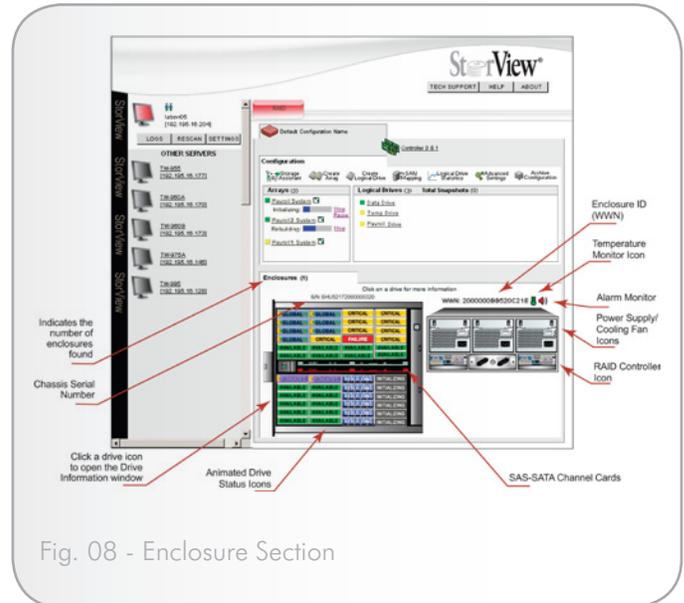


Fig. 08 - Enclosure Section

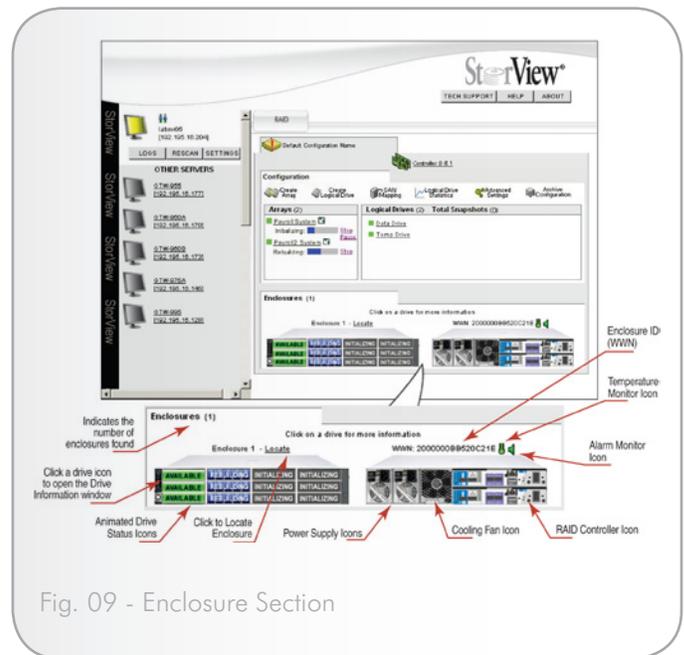


Fig. 09 - Enclosure Section

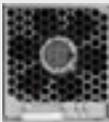
Table 02 - Enclosure Sections Descriptions

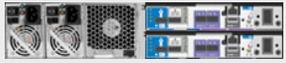
Enclosure Section	Icon	Description and Condition
Drive Status Icons: These icons appear on the front of the enclosure on the Main screen and provide an instant view of the status of disk drives as well as drive location and clicking the icon will open the Drive Information window.		
Array Member		Array Member - Disk drive is a member of an array.
Array Member - Spun down		Array Member (Spun-Down) - Disk drive is a member of an array but is spun down under APM.
Available		Available - Disk drive is online and available for use in an array or as a hot spare.
Available - Spun Down		Available (Spun-Down) - Disk drive is online and available for use in an array or as a hot spare but is spun down under APM.
Dedicated Spare		Dedicated Spare - Disk drive is marked as a dedicated spare to an array.
Dedicated Spare - Spun down		Dedicated Spare (Spun-Down) - Disk drive is marked as a dedicated "cold" spare to an array but is spun down under APM.
Empty		Empty - Disk drive slot is empty.
Failed		Failed - Disk drive has failed.
Global Hot Spare		Global Hot Spare - Disk drive is a global spare.
Global Hot Spare - Spun down		Global Hot Spare (Spun-Down) - Disk drive is a global "cold" spare but is spun down under APM.
Missing		Missing - Indicates that StorView is unable to determine the Initializing - Disk drive is a member of an array being initialized. status of the drive.
Initializing		Initializing - Disk drive is a member of an array being initialized.
Rebuilding		Rebuilding - Drive members of an array are in rebuild mode.
Locate		Locate - Clicking the "arrow" icon next to the "specific array" in the Arrays section will display an "arrow" icon on all the drive members of that array in the front enclosure view.
Critical		Critical - Drive(s) are members of a fault tolerant array and are in a non-fault tolerant state.*
Updating Firmware		Updating Firmware - This icon will appear when the subject drive's firmware is being updated.

continued on the next page >>

\*With RAID 6 disk arrays, the loss of a single drive does not place the array in a non-fault tolerant state. The drive status icon is the same for single or dual drive failures. To verify whether the array is in a Non Fault Tolerant vs. Fault Tolerant Warning state, identify the number of drive failures for the array. Two drive failures will put the RAID 6 array in a non fault tolerant state. Passing the mouse pointer over the status icon (yellow) will display the array's status, also in the Array Information screen the status is displayed.

Enclosure Section	Icon	Description and Condition
Failed Array Member		Failed Array Member - This icon will appear on all disk drives that are members of an array that has failed. For example if you remove a drive from a RAID 0 array or a drive in that array fails, the remaining drive members will have this icon displayed indicating that array has failed. If you accidentally remove the wrong drive in a critical redundant array (RAID 5) instead of the failed drive, that array will have failed and its member drives will have this icon displayed. Re-inserting the drive that was accidentally removed will put the drive members back to a critical state in which the array is being rebuilt.
Failed Array Member - Spun Down		Failed Array Member (Spun Down) - This disk drive is a member of an array. It has failed and was spun down under APM.
Expanding		Expanding - This icon is displayed on the drive members whose array is expanding.
Verifying		Verifying - This icon is displayed on the drive members whose array's parity data is being verified.
Unsupported		Unsupported - This icon is displayed when a drive does not match the other drives for the array. There are two images, one is for when the drive is a member of an array and the last one is when it is not a member of an array.
Unsupported		
Enclosure Component Icons: The following icons appear on the Main screen in the Enclosure section. They provide an instant view of the operating status of the enclosure components, and clicking the icons will open the Controller Information window.		
RAID Controller - Normal		Normal - RAID Controller is operating normally.
RAID Controller - Error		Error - A RAID Controller has failed in an Active-Active topology or the backup battery has failed.
RAID Controller - Empty		Empty - This icon represents the empty controller slot for future expansion. A blank plate is shown.
I/O Module - Normal		Normal - LaCie 12big rack serial I/O Module is operating normally.
I/O Module - Error		Error - The LaCie 12big rack serial I/O Module has failed.
I/O Module - Missing		Missing - The LaCie 12big rack serial I/O module is missing from the slot.
I/O Module - Empty		Empty - The LaCie 12big rack serial I/O module is removed and a blank plate is installed.
12big rack fiber RAID Controller - Normal		Normal - RAID Controller is operating normally.

Enclosure Section	Icon	Description and Condition
12big rack fiber RAID Controller - Error		Error - A RAID Controller has failed in an Active-Active topology or the backup battery has failed.
12big rack fiber RAID Controller - Empty		Empty - This icon represents the empty controller slot for future expansion. A blank plate is shown.
Disk I/O Module - Normal		Normal - Disk I/O Module is operating normally. (12big rack fiber only)
Disk I/O Module - Error		Error - Disk I/O Module has failed. (12big rack fiber only)
Disk I/O Module - Missing		Missing -Disk I/O is missing from the slot. (12big rack fiber only)
Disk I/O Module - Empty		Empty - Disk I/O is removed and a blank plate is installed. (12big rack fiber only)
Fan Icon - Normal		Normal - Both fans are operating normally.
Fan Icon - Fan 1 failed		Fan 1 Failed - One fan in the fan module has failed. The fan which failed will be indicated on the icon.
Fan Icon - Fan 2 failed		Fan 2 Failed - One fan in the fan module has failed. The fan which failed will be indicated on the icon.
Failure		Failure - Both fans in the fan module have failed or the cooling fan module has been removed.
Power Supply Icon - Normal		Normal gray icon indicates that the power supply is operating normally.
Power Supply Icon - Failure		A red flashing icon with "Failure" displayed indicates that the subject power supply has failed.
Power Supply Icon - Missing		A solid red icon indicates that the power supply is missing.

Enclosure Section	Icon	Description and Condition
Power Supply Icon - Unknown		Unknown - This icon indicates the enclosure power supply information from the SES processor or SAF-TE processes is missing or invalid.
Alarm Monitor Icon - Off		This icon indicates the alarm is silent.
Alarm Monitor Icon - On		This icon indicates the alarm is On (Continuous), On (Intermittent), or On (Remind).
Alarm Monitor Icon - Mute		This icon indicates the alarm is muted.
Alarm Monitor Icon - Disabled		This icon indicates the alarm is disabled. Note: Clicking the icon will open a window to manage the alarm.
Enclosure Temp. Icons: Enclosure temperature icon is displayed just above the rear enclosure icon and indicates the status of the enclosure temperature.		
Enclosure Temp. Icon - Normal		Normal - This icon indicates that the temperature is normal. It appears green.
Enclosure Temp. Icon - Warning		Warning - This yellow icon indicates that the enclosure temperature is approaching the established threshold.
Enclosure Temp. Icon - Failed		Failed - This red icon indicates that the enclosure temperature has reached or exceeded the enclosure temperature threshold. (If the fans are operating normally and the air flow temperature seems normal it may be an indication that the temperature sensor is faulty.)
Enclosure Temp. Icon - Missing		Missing - This icon indicates that the information from the SES regarding the sensors is invalid or missing.
Rear Enclosure Icon - 12big rack fibre		Normal - All components are operating normally.
		Communication Error- The SES process has lost communication with the enclosure, indicated by the icon becoming grey or dim. Or you have disabled "Enclosure Support" in the Controller Advanced Settings window.
Rear Enclosure Icon - 12big rack fibre expansion		Normal - All components are operating normally. (Used with 12big rack fibre only)
		Communication Error- The SES process has lost communication with the enclosure, indicated by the icon becoming grey or dim. Or you have disabled "Enclosure Support" in the Controller Advanced Settings window.

### 2.1.3. Array and Logical Drive/Snapshot Section

Table 03 - Array and Logical Drive/Snapshot icons

Array and Logical Drive Section	Description and Condition
<p>Array Status Icons: These icons are found in the Array section and depict the current state of the specific component.</p>	
<p> Green (Normal)</p>	<p>Status is ok.</p>
<p> Yellow (Warning)</p>	<p>Indicates a drive component in a RAID 1, 10, 5, 6 or 50 array has failed and the array is no longer fault tolerant, or the array is in a rebuild cycle.</p> <p>For RAID 6 arrays, it indicates 1 or 2 drives have failed. See Drive Status Icon "Critical" state.</p>
<p> Red (Error)</p>	<p>Indicates an array is invalid or offline due to an error:</p> <p>RAID 0 = One drive has failed.</p> <p>RAID 1/10 = Two drives have failed from the same pair.</p> <p>RAID 5 = Two drives have failed.</p> <p>RAID 6 = Three drives have failed.</p> <p>RAID 50 = Two drives have failed within the same sub-array.</p>
<p>Logical Drive Status Icons: These icons are found in the Logical Drive section and depict the current state of the specific component.</p>	
<p> Green (Normal)</p>	<p>Status is ok.</p>
<p> Yellow (Warning)</p>	<p>Yellow - Indicates the logical drive is part of an array that is degraded.</p> <p>For logical drives consisting of RAID 6 arrays, it indicates 1 or 2 drives have failed. See Drive Status Icon "Critical" state.</p>

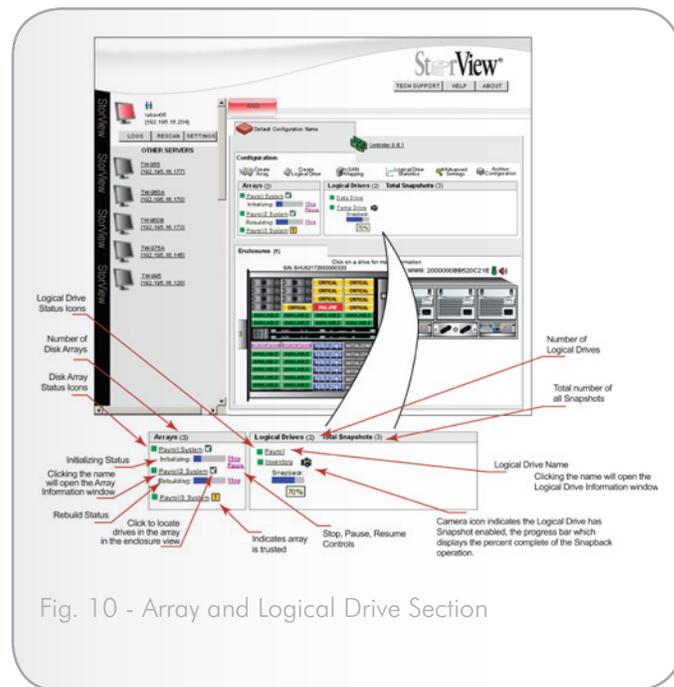


Fig. 10 - Array and Logical Drive Section

continued on the next page >>

Array and Logical Drive Section	Description and Condition
 Red (Error)	Red - Indicates the logical drive is part of an array which is invalid or offline,  RAID 0 = One drive has failed.  RAID 1/10 = Two drives have failed from the same pair.  RAID 5 = Two drives have failed.  RAID 6 = Three drives have failed.  RAID 50 = Two drives have failed in the same sub-array.
<p>ODA Status Icons: This icon is found in the Logical Drive section adjacent to the Snapshot enabled logical drive and depicts it's current state.</p> <p>During a Snapback operations, a status bar displaying the percent complete will appear under the Logical Drive icon.</p>	
 Normal	ODA is OK.
 Warning	ODA is reaching capacity. Free up ODA space by deleting unneeded snapshots. A warning will appear when a minimum amount of space remains on the ODA. The specific amount is dependant on the size of the ODA.
 Failed	The ODA has run out of space. All snapshots are invalid. In order to resume operations, delete all snapshots.
<p> Trusted Array Icon: This icon is found in the Array section, adjacent to a disk array. It indicates that the array has been trusted.</p>	

**NOTE:** At the top of the Arrays, Logical Drives and Snapshot section are values displayed in parenthesis next to each title. The value indicates the total number of Arrays for the storage solution, Logical Drives for the storage solution and total number of Snapshots taken for the storage solution.

### 2.1.4. Server Sidebar and Top Section

Table 04 - Server section icons

Server Section	Description and Condition
<p>StorView Server Icons: These icons will appear on the Server Sidebar section on the Main screen. They depict the current server and discovered servers. The icon will indicate the status of the server components by changing its color and state.</p>	
<p>Normal</p> 	<p>Normal Gray - Status is ok.</p>
<p>Warning</p> 	<p>Flashing Yellow - Indicates a server warning that a device connected is in degraded mode.</p>
<p>Error</p> 	<p>Flashing Red - Indicates a server error or device malfunction.</p>
<p>Remote StorView Servers Icon (Global License)</p>	
<p>Normal</p> 	<p>Normal Gray - Status is ok.</p>
<p>Warning</p> 	<p>Flashing Yellow - Indicates a server warning that a device connected is in degraded mode.</p>
<p>Error</p> 	<p>Flashing Red - Indicates a server error where a device has malfunctioned.</p>

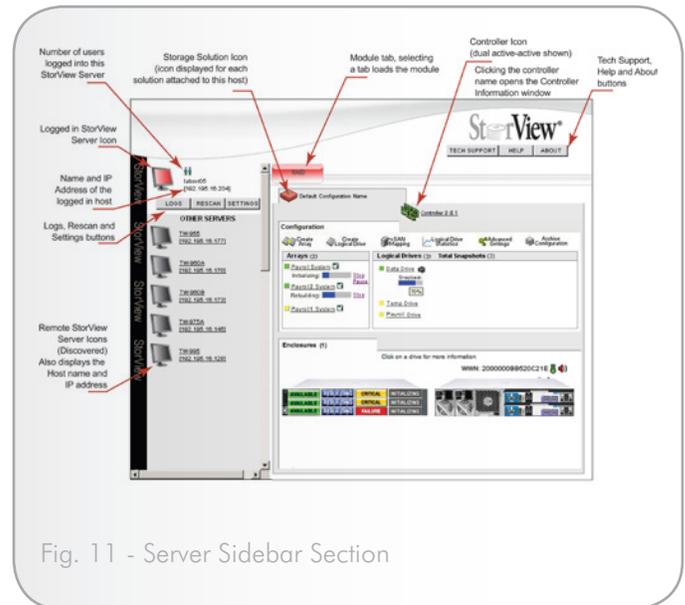


Fig. 11 - Server Sidebar Section

continued on the next page >>

Server Section	Icon	Description and Condition
Remote StorView Servers Icon		Flashing White - The server has not responded in at least 40 seconds and is considered missing. If you would like to remove the missing server from the list, click the Rescan button. This will refresh the screen with a current list of discovered servers.
User Icon (located adjacent to Server icon)		Represents each user logged into the StorView server you are monitoring. Placing the mouse pointer over the icon will display the IP address, host name and the user name.
Storage Solution Icon (displayed for each storage solution): You will also see the warning “!”, error, and unknown icons for unfocused storage solutions that are being monitored as well.		
Normal		Normal Gray - Status is ok.
Warning		Flash Yellow with red “!” - Indicates a component in the storage solution is in degraded mode.
Error		Flashing Red - Indicates a component in the storage solution has malfunctioned.
Unknown		Flashing Red with “?” - Indicates that the storage solution was there at startup but now cannot be located.
Storage Solution: Unmonitored		This icon indicates that another StorView server is monitoring this storage solution, or if you just performed a rescan then the StorView servers are still determining which StorView server will take control of the monitoring of the storage solution.
Controller Icons: These icons represents the RAID Controller installed in the enclosure. For duplex systems (Active-Active), a dual controller image is displayed.		
Normal		A green icon represents a normal operating system.
Normal		
Error		A flashing red icon appears if the controller’s backup battery unit has failed, or in Active-Active topologies when the partner controller has failed.
Module Tabs: RAID Module		The tab appears at the top of the Main window and when selected will focus the monitoring and management functions to specific systems types. This tab will flash yellow if a warning condition occurs and red if an error condition occurs.

## 2.2. How to Use this Document

The design of this user guide is to introduce StorView to its users, provide an explanation of the interface through this quick tour section and provide a step-by-step approach to configuring the network settings when using the embedded version. If you are using the host-based version you may skip Chapter 3.

The Getting Started chapter will walk you through starting StorView, upgrading the access license if necessary, and configuring E-Mail, SYSLOG, SNMP, and additional monitoring.

Chapter 5 will step you through the entire process of configuring the storage solution from defining disk arrays, assigning hot spares, and configuring the logical drives.

The remaining chapters deal with the more advanced features such as Snapshot, SAN LUN Mapping, controller environment monitoring and optimization, and modifying controller operational parameters. You will also find information on advanced power management and performance options in the latter chapters.

---

## 3. Embedded StorView Setup

### 3.1. Embedded Network Settings

Note If you are using only the host-based version of StorView, you may skip this chapter.

After completing your hardware set up, ensure you have a shielded Cat 5 or better Ethernet cable connected to the controller's networking port (StorView Embedded Module port (RJ-45)) and your network hub or host computer Ethernet port. If you are connecting to the network hub use a standard Ethernet cable, if you are connecting to your host port use a cross-over Ethernet cable. The first time you start Embedded StorView, you must configure the network settings.

Upon start-up, embedded StorView looks at the user Preferences Settings to determine if an IP address exists. If one is defined it will initialize the network interface using that IP address.

In the event an IP address is not defined, it attempts to get a DHCP IP address. You will need to contact your network administrator for the IP address assigned by the DHCP server. To identify the new IP address lease, one can look for "esv0" or "esv1" in your DHCP Manager software.

If an IP address cannot be determined, the software will use a default IP address of "10.1.1.5" for Controller 0 and "10.1.1.6" for Controller 1. If your network does not have a DHCP server, then set the host IP address to 10.1.1.20 and the subnet mask to 255.0.0.0 and run the wizard again. If an error is encountered, the embedded StorView Server will have the IP address of "10.1.1.7" assigned to it.

A tool is provided to configure new embedded StorView modules. From a Windows platform, it is accessed via the Embedded StorView Setup Wizard link on the disc navigation menu and on Linux it is accessed via a command line executable.

---

#### 3.1.1. Configuring Network Settings using the Wizard

When the Embedded StorView Setup program is run, it will broadcast UDP packets and any embedded StorView module will reply with UDP packets containing their information. A list of "Uninitialized Systems" is displayed. Uninitialized systems are those which have not had the default user name and password changed.

Even if a configuration is created with arrays and logical drives but the login name and password have not been changed, it will still be considered an uninitialized system. During the process of configuring an Embedded Module, you will be required to enter a "new"

password and confirm that password. The default password is the word "password."

The Embedded Module is identified by its MAC and IP address. It may be more helpful during setup to configure one Embedded Module at a time. You will find the MAC and IP address by accessing the VT-100 RAID Configuration Utility, selecting Diagnostics then choosing StorView Embedded Module Support. Select Enter StorView Embedded Module Menu Mode and choose View Network Settings.

[continued on the next page >>](#)

---

1. Insert the Software and Manuals Disc into your CD drive. The autorun program will automatically start the navigation menu.
2. Click the "Embedded StorView Setup Wizard" link to begin the configuration.
3. You are presented with a Welcome screen with instructions to proceed. Review the information.
4. Click the Next button and the wizard will begin searching for Embedded StorView Modules.

Those Embedded Modules with their default passwords intact are considered an Uninitialized System and will be displayed with their MAC address in the Uninitialized Systems window (Fig. 14).

5. From the Uninitialized Systems screen, select the MAC address of the Embedded Module you wish to configure.

If you wish to use a DHCP server to assign your IP address, click the check box "Use DHCP."

If you wish to manually configure your network setting, enter the correction information in the appropriate fields.

Enter a "new" password and enter the new password confirmation. Click the Configure button.

[continued on the next page >>](#)



Fig. 12 - Embedded StorView Setup Wizard - Welcome Screen

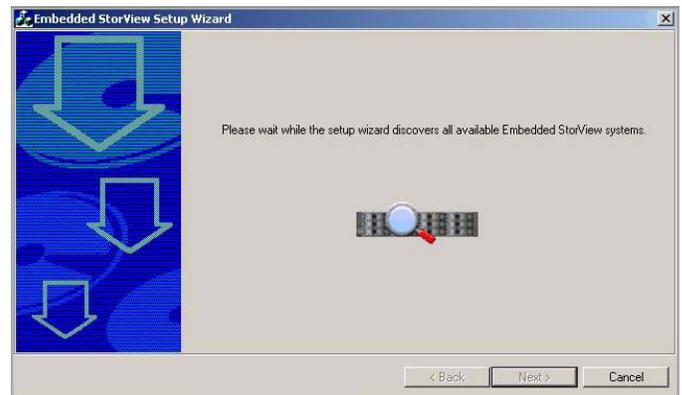


Fig. 13 - Embedded StorView Setup Wizard - Search Screen



Fig. 14 - Embedded StorView Setup Wizard - Uninitialized Systems Screen

- If you have additional Uninitialized Systems, select the next MAC address and choose the appropriate settings from the previous step.

Once you have completed configuring all the Uninitialized Systems and have clicked the Configure button, the wizard will display a popup message indicating all systems have been configured. It will then re-scan for any Uninitialized Systems and if none are found, you will be taken to the Initialized Systems screen (Fig. 17).

However, if you wish to only configure one Uninitialized System, the one you are configuring, and choose not to configure any other discovered Uninitialized System, click the Next button. The following prompt will appear. Click the YES button to continue or NO to cancel and return to the Uninitialized Systems screen.

**NOTE:** If someone plugs in a uninitialized module to the network (same subnet mask) during the time of the rescan or resets a system's password back to its defaults, you will be taken to the Uninitialized System screen again.

- Select the MAC address of the Initialized System you wish to start and click the Launch button.
- Your default web browser will open with a login screen. Enter the login name and password, then click OK. StorView will open in the browser at the Main screen.

**NOTE:** If you are finished starting up Embedded Modules, be sure to select the active Embedded StorView Setup Wizard window in the Windows Task Bar and click the Close button.

### 3.1.2. Configuring Network Settings using the CLI

From a Linux system using a command line interface (CLI), perform the following:

- Log in as "root" (password = `password`).
- Insert the Software & Manuals Disc into your CD drive.
- Change directories to the software location. Type:
 

```
cd [CDROM mount point path]/software/storview/embedded
```
- Execute the Setup Tool. From the command prompt, type:
 

```
./esvsetupcl-linux  
Embedded StorView Setup Tool. v1.0  
Entering Menu Mode.
```
- Follow the on screen prompts. At the conclusion of the setup



Fig. 15 - Request to Proceed Window



Fig. 16 - Request to Proceed Window



Fig. 17 - Embedded StorView Setup Wizard - Initialized Systems Screen

you will be instructed on how to proceed. StorView will start within your web browser.

**NOTE:** Host-based StorView is started by launching your web browser and entering the IP address of the local host followed by “:9292”. For more information regarding SSL (Secure Socket Layers), see section 4.1. *Starting StorView*.

### 3.1.3. Configuring Network Settings using StorView

1. Click the SETTINGS button on the Main screen and select the PREFERENCES tab.
2. The StorView Host Name field will have a default name, “esv0” for Controller 0 and “esv1” for Controller 1. If you wish to change this name, enter the desired name for this Embedded StorView server.

**NOTE:** If the StorView Host name displays “esvrr,” this indicates a problem with the Embedded StorView server module. Contact technical support.

#### 3.1.3.1. Using Dynamic IP (DHCP)

**NOTE:** StorView does not display the TCP/IP information when Dynamic IP (DHCP) is selected. You must use a third-party network administration program to obtain this information.

To use the DHCP Server network interface setting:

1. Click the option button next to Dynamic IP (DHCP).
2. Click the APPLY button to make the changes effective.
3. Click the CLOSE button.

#### 3.1.3.2. Using Static IP

To manually configure the network interface setting:

1. Click the option button next to Static IP.
2. Enter the desired IP address in the “IP Address” field and press the <Tab> key or click in the Subnet Mask field.
3. Enter the desired Subnet Mask in the “Subnet Mask” field and press the <Tab> key or click in the Default Gateway field.

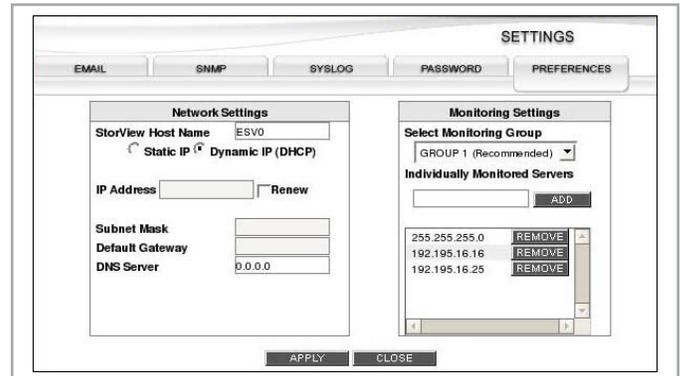


Fig. 18 - Settings Screen – Preference Tab (Dynamic IP Selected)

4. Enter the desired gateway or router address and press the <Tab> key or click in the DNS Server field.
  5. Enter the desired DNS Server IP address.
  6. Click the APPLY button to make the changes effective.
  7. Click the CLOSE button.
- 

### 3.1.4. Getting a New IP Address

If you are set up to receive your IP address using Dynamic IP (DHCP), you can force the embedded StorView server to obtain a new IP address from your DHCP server.

1. Click the SETTINGS button and select the PREFERENCES tab.
  2. Click the Renew check box.
  3. Click the APPLY button to make the changes effective.
  4. Click the CLOSE button.
-

## 4. Getting Started

### 4.1. Starting StorView

#### Host-based StorView

It is started by opening your web browser software and entering the IP address of the host attached to the storage system followed by the port number. The format is:

“HTTP://<explicit\_IP\_address>:9292,” “HTTP://127.0.0.1:9292”  
or “HTTP://localhost:9292.”

---

**NOTE:** For Microsoft Windows users, you may also launch StorView via the Start Menu. Click the Start button and select Programs, then choose StorView and select “StorView Manager Console.” Your web browser will open with a login screen.

---

To use SSL (Secure Mode), enter one of the following URL’s: “HTTPS://<your\_IP\_address>:9393,” or “HTTPS://<device\_fully\_qualified\_domain>:9393,” “HTTPS://localhost:9393,” or “HTTP://127.0.0.1:9393.”

You will get a warning notice that the certificate is not trusted. You can choose to allow the device only to be trusted or a root certificate that is provided to be trusted. To install a root certificate, enter the following URL: “HTTP://<machine\_fully\_qualified\_domain\_name>:9292/CA.cert” After the certificate is added to the trusted certificates, a secure browsing lock icon will appear on the web browser window indicating a secure connection.

---

**NOTE:** To trust the device, only you can modify the web browser Internet Options or Preferences settings for Network and alter the setting for web proxy.

---

---

**NOTE:** For Microsoft Windows Internet Explorer, be sure you have checked the box “Use SSL 2.0” under the Internet Options > Advanced > Security settings.

---

#### Embedded StorView

It is started by opening your web browser software and entering the IP address assigned to the controller’s embedded module followed by the port number. The format is “HTTP://10.11.48.120:9292.” For the first startup of the embedded version you will need to perform some network configuration of the module, refer to “Configuring Embedded Network Settings.”

To use SSL (Secure Mode), enter one of the following URL’s: “HTTPS://<your\_IP\_address>:9393” or “HTTPS://<device\_fully\_qualified\_domain\_name>:9393.”

You will most likely get a warning notice that the certificate is not trusted. You can choose to allow the device only to be trusted or a root certificate that is provided to be trusted. To install a root certificate, enter the following URL: “HTTP://<machine\_fully\_qualified\_domain\_name>:9292/CA.cert” After the certificate is added to the trusted certificates, a secure browsing lock icon will appear on the web browser window indicating a secure connection.

---

**NOTE:** For Microsoft Windows Internet Explorer, be sure you have checked the box for “Use SSL 2.0” under the Internet Options > Advanced > Security settings.

---

---

**NOTE:** If your network is using Microsoft Proxy Server, you may need to request IS to allow the port 9393. You may also consider setting the browser to force the intranet address to go direct bypassing the proxy. You may wish to consult the following Microsoft web page for additional information: <http://support.microsoft.com/kb/283284/en-us>.

---

Upon startup, you will be prompted for a user name and password. The default user name is “admin” and password is the word “password.”

It is started by opening your web browser and entering the explicit IP address assigned to the embedded StorView server followed by the port number (e.g., HTTP://10.11.48.120:9292).

---

## 4.2. Upgrading the Access License

Some capabilities of StorView are dependent on which access license is installed. If you have the Local Access license (host-based) or Remote Access (embedded) license installed, you are limited to local management and monitoring of the storage solution attached to the host system. The Global Access license enables the premium options of either version of StorView which provides full functionality and remote access, E-Mail notifications, SYSLOG, and SNMP.

If a remote login is attempted from another host system on the same network and you do not have a global license, you will see a message displayed regarding this. You will have the option to upgrade your license by entering your serial number and activation code. You can also upgrade your license from the

Other Servers section by clicking on the link in the notice displayed or the upgrade can be accessed by clicking the Settings button (E-MAIL tab, default, presents an upgrade option when using a local or remote access license).

1. Click on the "link" provided under "Other Servers." The Settings window will open with the E-MAIL tab selected. There you will enter the required information and click the ACTIVATE button.
2. Once you have completed the upgrade, the window will refresh with the E-MAIL, SYSLOG, and SNMP tabs now active. Click on the CLOSE button in the confirmation window.
3. Click the CLOSE button on the Settings window.
4. You can verify the change by clicking the About button and noting that it now displays "Global."

Also the notice displayed under the Other Servers section will now be removed and any remote discovered StorView Servers will be displayed.

Table 05 - License Types

License	Features	Version
Local Access	Configuration, GUI Monitoring, and Event Logs.	Host-based
Remote Access	All the features of Local, plus Remote Login and Monitoring Failover.	Embedded
Global Access	All the features of Local and Remote, plus E-Mail, SNMP, SYSLOG, and Other Servers list.	Host-based and Embedded

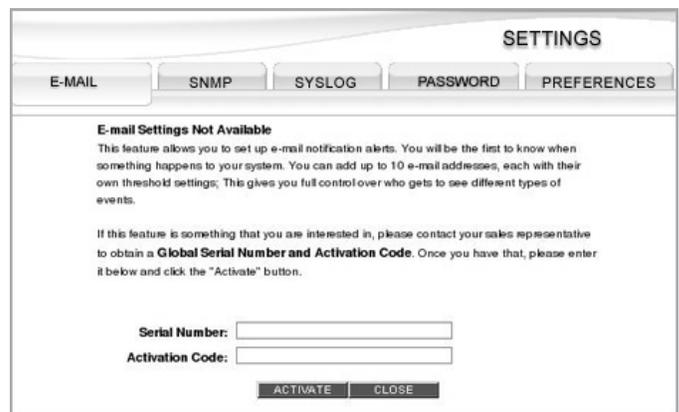


Fig. 19 - License Upgrade Screen

## 4.3. E-MAIL

### 4.3.1. Configuring E-Mail Notices

With a Global license installed, StorView provides you with the ability to establish up to ten E-Mail addresses where notices of events can be automatically sent.

To configure the E-Mail notifications perform the following:

1. From the Main screen click the SETTINGS button. The Settings window will open with the E-MAIL tab selected.

**NOTE:** All fields have a limit of 127 ASCII characters.

2. Enter the “name” or “IP address” of your E-Mail server. This will be the SMTP mail server name. E-Mail messages are sent to the E-Mail server using port 25.

**NOTE:** If your E-Mail server is not configured to receive on port 25, then E-Mail will not function properly.

3. The “From” field should be pre-configured using the software name followed by your fully qualified domain name. If you change this setting be sure it is compliant with your IS policy.
4. The “Reply-To” field should be completed using a desired e-mail address that recipients could respond to when they receive these automatically sent e-mails. This field by default will be blank.
5. If you would like a signature appended to the message, click the check box and type in the signature information in the scrollable window provided.
6. Enter the user E-Mail addresses. You may add up to ten (10) E-Mail addresses.

Type the full E-Mail address, and click one or more of the check boxes for the type of event to which the user is to be notified.

The types of events are: Information , Warning , and/or Error .

If you have more than five E-Mail recipients, you will need to click the button “6 - 10” to access the next five address blocks.

7. Once complete, click the APPLY button. You will receive a confirmation message that the changes were successfully completed. Click the CLOSE button.
8. Test the configurations by clicking the TEST button.  
You will receive a confirmation message that the test was successfully completed, and each addressee will receive a “Test Message” from the mail server. Click the CLOSE button.
9. Click the CLOSE button on the SETTINGS window.



Fig. 20 - Settings Screen – E-MAIL Tab

### 4.3.2. Deleting an E-Mail Addressee

1. From the Main screen click on the SETTINGS button. The SETTINGS window will open with the EMAIL tab selected.
2. Click the DELETE button next to the E-Mail Address name you wish to remove. The address is immediately removed.
3. Click the APPLY button to make the changes effective and a status pop-up notice will appear. Click the CLOSE button. Then click the CLOSE button on the SETTINGS window.



Fig. 21 - Settings Screen – Deleting an Addressee

## 4.4. SNMP

### 4.4.1. Configuring SNMP Traps

StorView can be configured to send SNMP traps to any network management system. These traps carry all the information that appears in the log entries for each level of severity.

**NOTE:** If your system requires a MIB file, it is located in the following directory "C:\Program Files\StorView>alerts-snmp.mib."

All SNMP traps sent from StorView will be received by the host SNMP Servers designated in the settings window for the specified port and community.

1. From the Main screen click on the SETTINGS button.
2. Click the SNMP tab.
3. Enter the SNMP Server name or IP address of the host you wish to receive SNMP traps.
4. Enter the IP port on which the SNMP Server expects to receive traps. The default is 162.
5. Enter the Community to which the trap belongs. The default is public.
 

SNMP Servers may belong to several different communities or receive packets for different communities.
6. Select the level of events you wish to be included in the traps. You can select from Information, Warning and Error types.
7. Click the APPLY button.
8. Test the configurations by clicking the TEST button. You will receive a confirmation message that the test was successfully completed, and each addressee will receive a "Test Message." Click the CLOSE button.
9. Click the CLOSE button on the SETTINGS window.

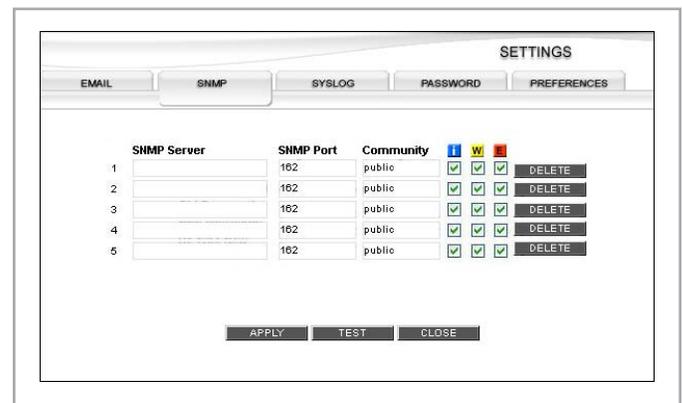


Fig. 22 - Settings Screen – SNMP Tab

### 4.4.2. Deleting an SNMP Server

1. From the Main screen click on the SETTINGS button.
2. Click the SNMP tab.
3. Click the DELETE button next to the SNMP Server you wish to remove. The server is immediately removed.
4. Click the APPLY button to make the changes effective and a status pop-up notice will appear. Click the CLOSE button. Then click the CLOSE button on the SETTINGS window.

## 4.5. SYSLOG

StorView's Syslogd agent is a module designed for StorView that enables sending RAID system event logs to all listening syslogd clients on all installed platforms.

Syslogd requires that the Global Access license be enabled before the feature is activated. Once activated on each installation of StorView, the locally monitored storage system events are then sent to the syslogd daemon running on the host system.

StorView syslogd is designed for greater flexibility in that it supports "many-to-many" mapping. You may have:

- ◆ Multiple syslogd servers that monitor a single RAID storage system for event logs.
- ◆ A single syslogd server that monitors multiple RAID storage systems for event logs.
- ◆ Multiple syslogd servers that monitor multiple RAID storage systems for event logs.

This "many-to-many" mapping provides an additional safety net in that it offers data redundancy. In this instance a failed syslogd server will not have any impact if another syslogd server is functioning normally.

If you are using a Linux operating system, the syslogd daemon is a component of your operating system. If you are using a Windows operating system you will need to install a separate tool (<http://syslogwin32.sourceforge.net/>) which runs as a Windows service.

### 4.5.1. Gathering Logs

A host based StorView running the syslogd agent with two in-band connections to RAID storage systems, each with a unique WWN ID. Assuming the syslog is configured to log/dispatch Error type events, it will gather all Error type events for both RAID storage systems and sent them directly to the intended host.

The StorView syslogd agent is global in scope, in that it will always

receive all error logs from the accessible storage systems. It is not possible to customize filtering of events for each managed storage system for that host.

With Embedded StorView, previously configured listening syslogd servers will persist even though the Embedded Module is flashed with new software as the file data is stored in NVRAM and is not affected by software upgrades.

However, with instances of installed host-based StorView, you must backup or make a copy of the syslog.ini file otherwise software updates will destroy the configuration information.

### 4.5.2. Syslog Daemon

The daemon listens for incoming event log messages on port 514. It logs them to a messages file. The syslogd messages consist of:

- ◆ Facility - indicates the type of syslog message. For example they can be kernel, user level, mail system and system daemon messages.
- ◆ Priority - indicates the sensitivity of the message, to include: Information, Warning, or Error.
- ◆ Timestamp - indicates the date and time of the event.
- ◆ Originating Device - specifies which device originated the event message.
- ◆ Message Content - provides a plain language content of the event message.

Syslogd daemon supports both Linux and Windows platforms.

### 4.5.3. Operational View

Whenever a new RAID storage system event is logged, based on the type of event, StorView syslogd will send the message out to all of its subscribers that match the type of event involved. If the event was an Error type, and all daemons are configured to receive that type, then they will all receive the event. If the event was an Information type, only those daemons configured to receive Information type will receive the event.

Event logs are sent in standard text format and have no encryption, in such security can be a concern. Additionally UDP packets are involved so there is no assurance that the event log arrived at the subscriber (syslogd server).

#### 4.5.4. Configuring Syslogd Server

1. If you have not yet enabled the Global Access license please do so at this time, see [4.2. Upgrading the Access License](#).

After activating the Access license, you will see the following screen when the SYSLOG tab is selected.

The fields are described in table [Table 06](#).

2. Enter the IP addresses of the host servers designated to be syslogd servers.
3. Enter the port ID of the recipient syslogd server. You must enter a correct value otherwise you will receive an error message.
4. Select the check boxes under the severity/priority icons.
5. Click the APPLY button to save your settings.
6. If you are finished completing the setup, click the CLOSE button on the SETTINGS window.
7. If you wish, test the settings by clicking the TEST button. You will receive a confirmation message that the test was successfully completed, and each syslogd server will receive a "Test Message." Click the CLOSE button.

#### 4.5.5. Delete a Syslogd Server

1. Access the Settings window and select the SYSLOG tab.
2. Click the DELETE button next to the "syslogd server" you wish to remove. The server name is immediately removed.
3. Click the APPLY button to make the changes effective and a status pop-up notice will appear. Click the CLOSE button. Then click the CLOSE button on the SETTINGS window.

Table 06- Syslogd Server Fields

Field	Description
syslogd Server	Allows the user to enter the IP address of the syslogd servers. An error message is displayed if an invalid server name is entered.
syslogd Port	Allows the user to enter a specified syslogd port. The default port is 514. An error message will appear if the enter is invalid or left blank. The range of port IDs are: 1 - 65535.
Information (I)	If checked the recipient would receive Informational type of event logs.
Warning (W)	If checked the recipient would receive Warning type of event logs.
Error (E)	If checked the recipient would receive Error type of event logs.
Delete Button	Clicking this button clears the previously entered IP address for the selected line and resets the check boxes to their default state.
Test Button	Send out a dummy test message to the specified recipient servers.



Fig. 23 - Settings Screen – Syslog Tab

## 4.6. Changing the Password

This option provides the ability to change the access password used at log in.

1. From the Main screen click on the SETTINGS button.
2. Click the PASSWORD tab at the top of the window.

**NOTE:** Passwords will not be displayed as you type them.

3. Type in the Old Password and press the <Tab> key or click in the next text box.
4. Type in the New Password and press the <Tab> key or click in the next text box.
5. Re-type the New Password and click the CHANGE button. You will receive a confirmation message that the changes were successful. Click the CLOSE button.
6. Click the CLOSE button on the SETTINGS window.

**NOTE:** If you lose or misplace your password, contact technical support for further instructions.

## 4.7. Monitoring Settings

The following options enable network administrators to make adjustments to the StorView server's multicast functionality. In the event there is a port conflict with the default multicast port, you have the ability to change this parameter.

**NOTE:** The Monitoring Settings are disabled with the Remote license, you must upgrade to a Global license to enable these features.

1. From the Main screen click on the SETTINGS button, then click the PREFERENCES tab.
2. Click the pull-down menu for "Select Monitoring Group" and choose Group 1, Group 2, or Group 3. Group 1 is port 9191, Group 2 is port 9192, and Group 3 is port 9193.
3. Click the APPLY button to make the changes effective, then click the CLOSE button. Click the CLOSE button on the SETTINGS window.



Fig. 24 - Settings Screen – Password Tab

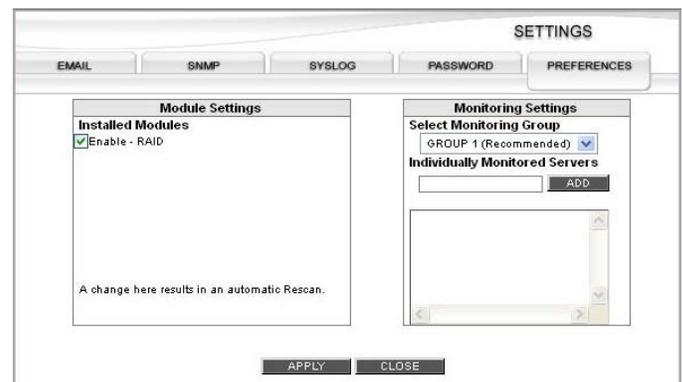


Fig. 25 - Monitoring Settings Screen – Preferences Tab (Host-based Version)

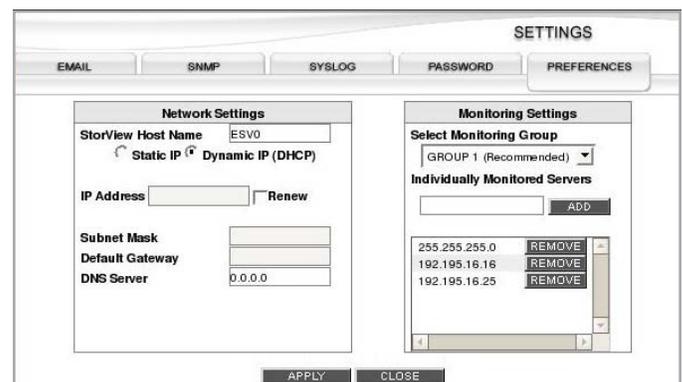


Fig. 26 - Monitoring Settings Screen – Preferences Tab (Embedded Version)

### 4.7.1. Additional Monitoring Servers

To specify additional StorView server(s) on a different subnet to be included in the receipt of StorView server packets, enter the IP addresses of those other StorView servers. You may add up to 10 additional monitored servers.

1. From the Main screen click on the SETTINGS button.
2. Click the PREFERENCES tab at the top of the window.
3. Enter the IP address in the “Individually Monitored Servers” field of another StorView server outside the subnet and click the ADD button.
4. Add additional explicit IP addresses of any other StorView server you wish to include to receive packets that is outside the subnet and click ADD button. Otherwise, skip to step 5.
5. Click the APPLY button.
6. Click the CLOSE button on the SETTINGS window.

### 4.7.2. Remove Monitored StorView Server IP

1. From the Main screen click on the SETTINGS button.
2. Click the PREFERENCES tab at the top of the window
3. Under the Individually Monitored Servers section, select the Server IP Address you wish to delete and click the REMOVE button.
4. Click the APPLY button. A confirmation window will appear, click the CLOSE button.
5. Click the CLOSE button on the SETTINGS window.



Fig. 27 - Settings Screen – Preferences Tab (Embedded Shown for Example Purposes)

## 5. Configuring a Storage Solution

### 5.1. Creating Disk Arrays

Configuring a storage solution requires some planning to ensure that you define the optimum settings for your storage solution. This includes choosing the correct RAID level and array options, assigning hot spares and creating the logical drives. More advanced features are available and discussed in later chapters.

This chapter will step you through this process, however you should have a basic understanding of RAID and storage concepts.

---

#### 5.1.1. RAID Levels

The following are the drive requirements for each supported RAID level.

Table 07- Drive Requirements by RAID Level

RAID Level	Minimum # of Drives	Maximum No. of Drives
0	1	16
1	2	2
5	3	16
6	4	16
50	6	16
10	4	16

---

### 5.1.2. Terminology

Table 08- RAID Terminology

Term	Description
Array	A group of disk drives that are combined together to create a single large storage area. Up to 64 arrays are supported, each containing up to 16 drives per array. There is no capacity limit for the arrays.
Back-off Percent	In order to allow drives from a different family or manufacturer to be used as a replacement for a drive in an array, it is recommended that a small percentage of the drive's capacity be reserved when creating the array. This is user selectable, from 0, 0.1, 1-10, and 90 percent. This is sometimes known as Reserved Capacity.
Cache Flush Array	This is the array that is used to automatically flush cache data in a situation where power has failed to some of the drives.
Chunk Size	This is the amount of data that is written on a single drive before the controller moves to the next drive in the stripe.
Initialization	RAID 5, 6, and 50 disk arrays must have consistent parity before they can be used to protect data. However, StorView will initialize all array RAID types regardless. If the user chooses the Trust option during array creation or stops the initialization, the array will be trusted. Note that any drive failure in a trusted array will result in data corruption.  It is possible to perform the initialization later. This recalculates the parity based on current data, ensuring data and parity are consistent.
Logical Drive Availability	To accommodate hosts with multiple ports and multiple host systems, it is possible to restrict a logical drive's availability to a particular HBA or controller port. Access can be enabled or disabled for each host port of each controller.
Mapped LUN Number	Each logical drive is presented to the host system with a unique LUN. In certain cases (such as after deleting another logical drive) it may be desirable to change the number that a logical drive is presented as. This can be done at any time, bearing in mind that any attached host systems may need to be rebooted or re-configured to maintain access to the logical drive.
RAID Level 0	RAID 0 is defined as disk striping where data is striped or spread across one or more drives in parallel. RAID 0 is ideal for environments in which performance (read and write) is more important than fault tolerance or you need the maximum amount of available drive capacity in one volume. Drive parallelism increases throughput because all disks in the stripe set work together on every I/O operation. For greatest efficiency, all drives in the stripe set must be the same capacity. Because all drives are used in every operation, RAID 0 allows for single-threaded I/O only (i.e., one I/O operation at a time). Environments with many small simultaneous transactions (e.g., order entry systems) will not get the best possible throughput.
RAID Level 1	RAID 1 is defined as disk mirroring where one drive is an exact copy of the other. RAID 1 is useful for building a fault-tolerant system or data volume, providing excellent availability without sacrificing performance. However, you lose 50 percent of the assigned disk capacity. Read performance is somewhat higher than write performance.
RAID Level 5	RAID 5 is defined as disk striping with parity where the parity data is distributed across all drives in the volume. Normal data and parity data are written to drives in the stripe set in a round-robin algorithm. RAID 5 is multi-threaded for both reads and writes because both normal data and parity data are distributed round-robin. This is one reason why RAID 5 offers better overall performance in server applications. Random I/O benefits more from RAID 5 than does sequential I/O, and writes take a performance hit because of the parity calculations. RAID 5 is ideal for database applications.
RAID Level 6	RAID 6 is the same as RAID 5 except that it uses a second level of independently calculated and distributed parity information for additional fault tolerance. This extra fault tolerance provides data security in the event two drives fail before a drive can be replaced. RAID 6 provides greater fault tolerance than RAID 5. There is a loss in write performance with RAID 6 when compared to RAID 5 due to the requirement for storing parity twice for each write operation. A RAID 6 configuration also requires N+2 drives to accommodate the additional parity data, which makes RAID 6 require more raw capacity than RAID 5 for an equivalent usable storage capacity.
RAID Level 10	RAID 10 is defined as mirrored stripe set.

---

Term	Description
RAID Level 50	This RAID level is a combination of RAID level 5 and RAID level 0. Individual smaller RAID 5 arrays are striped, to give a single RAID 50 array. This can increase the performance by allowing the controller to more efficiently cluster commands together. Fault tolerance is also increased, as one drive can fail in each individual array.
Stripe	The process of separating data for storage on more than one disk. For example, bit striping stores bits 0 and 4 of all bytes on disk 1, bits 1 and 5 on disk 2, etc.
Stripe Size	This is the number of data drives multiplied by the chunk size.
Sub-array	In RAID 50 applications, this is the name given when individual RAID 5 arrays that are striped together. Each sub-array has one parity drive.
Unassigned Free Space	The controller keeps a map of all the space that is not assigned to any logical drive. This space is available for creation or expansion. Each unassigned region is individually listed.

---

### 5.1.3. Optimization and Drive Selection for RAID 5 and 6

Typical RAID 5 or 6 implementations require a number of steps to write the data to the drives. In order to optimize your system performance based on the type of writes you expect in your operation, we have provided detailed information on optimizing the performance using full strip write operations in an appendix. If you intend to setup a RAID 5 or 6 disk array and wish to consider optimum performance, you will need to consider the number of data drives, parity drives, and chunk size. For a review, refer to chapter [15. Optimizing Write Performance](#). Additional information is provided at the appropriate step during configuration, see section [5.2.2. Chunk Size Details](#).

---

## 5.2. Create the Array

Configuring an array involves a few steps from one screen, the Create Array window. Here a performance profile is selected, the user chooses the disk drives and enters a unique name for the array. The remaining array parameters are pre-configured by the performance wizard. Although the settings can also be changed manually from the pre-selected values using drop-down menus and check boxes, the profile selections should be valid for most configurations. Once completed the Create button is clicked, the array is created and it automatically begins to initialize.

The Performance Wizard is a component of the Create Array window. When you open the Create Array window and select a performance profile, the wizard will make recommended settings for the remaining array parameters. The performance profiles are General, Database Server, Post Production Video, Video On Demand and Mail Server. When a profile is selected from the drop-down menu, all the parameters except selecting the disk drives are pre-configured.

If you change a parameter from a recommended setting, a warning icon appears under the Recommendations column for that item. This indicates a deviation from the recommended setting. The recommended settings are not mandatory and the user can choose to ignore them when they configure their arrays. Note that a warning message will appear at the bottom of the window above the CREATE button if any setting has deviated from those the wizard recommended.

See [Table 10 - Drive Requirements by RAID Level](#) at right for a description of the icons appearing in the Recommendations column.

The General performance profile is the default profile.

These array parameters define the details of the array as it is created and are saved in the configuration file. The configuration file is stored on all disk drives that are members of the array (regardless of whether the drives are in one or multiple enclosures). No changes are made to the configuration until the current process is saved, so it is possible to quit at any time without affecting the current configuration.

After making changes to your configuration, be sure to make a new backup copy of the configuration file, see section [5.6. Saving the Configuration](#). The ability of making a backup copy of the configuration allows you to quickly recover from a damaged configuration that was not self healing, and restore everything to the point in time when the configuration was last saved. This preserves the definition of the arrays, logical drives, Snapshots, SAN LUN Mappings and controller parameter settings.

---

**CAUTION:** A damaged configuration could result in loss of data.

---

Table 10 - Drive Requirements by RAID Level

Array and Logical Drive Section	Description and Condition
OK Icon 	This icon indicates the setting for the parameter matches the recommended setting for the selected profile.
Warning Icon 	This icon indicates a setting that deviates from the recommended setting for the selected profile.

continued on the next page >>

1. On the Tool Bar click the Create Array icon. The Create Array window will open.

**NOTE:** Passing the mouse pointer over each item in the Create Array window displays information about the object or item in the Notes box. In the above examples, the mouse pointer is over a disk drive which is displaying the drive information in the Notes box.

2. Select the profile you wish to use from the Performance Profile drop-down box.

The choices are:

Table 09- Creating an Array Choices

Profile	RAID Level	Chunk Size	Read-Ahead Cache	Write-back Cache	Min. # of Drives
General	5	256K	Auto	16MB	3
Database Server	1/10	64K	Auto	16MB	2
Post-production video	5	64K	Auto	MAX MB	3
Video On-Demand	5	256K	Auto	16MB	3
Mail Server	6	256K	Auto	16MB	4

The suggested best-suited application for each profile is:

- ◆ General - For file or web server applications.
- ◆ Database Server - For applications having high transaction rate with significant proportions of write data.
- ◆ Post Production Video - For applications with large 64K plus I/O operations.
- ◆ Video On Demand - For applications streaming data from the arrays. Only one logical drive should be built for optimum performance.
- ◆ Mail Server - For applications requiring high data resilience with high capacities.

**NOTE:** Making a change to a setting from the profile selected settings are not reset or changed when the user changes to another profile. If you altered the settings of a profile, close the window and click the Create Array button on the tool bar again. See [Fig. 31](#) for an example of mismatched settings and the on screen warning message.

3. Select drives to include in your array. Click on each drive that

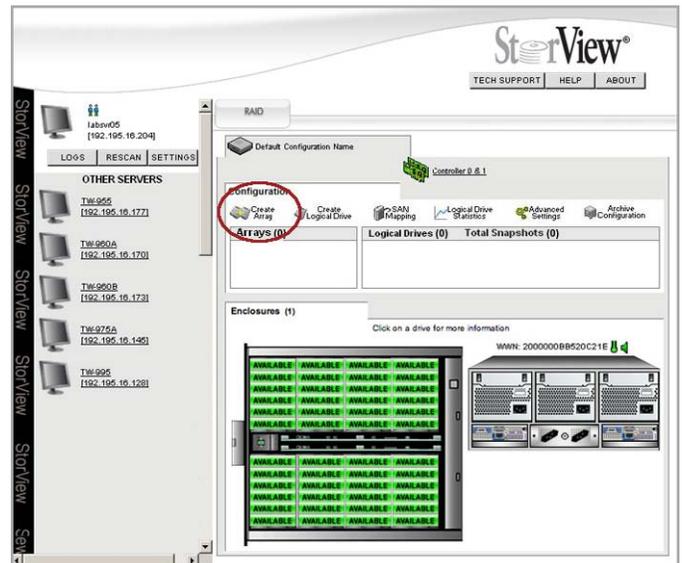


Fig. 28 - Main Screen – Create Array Icon

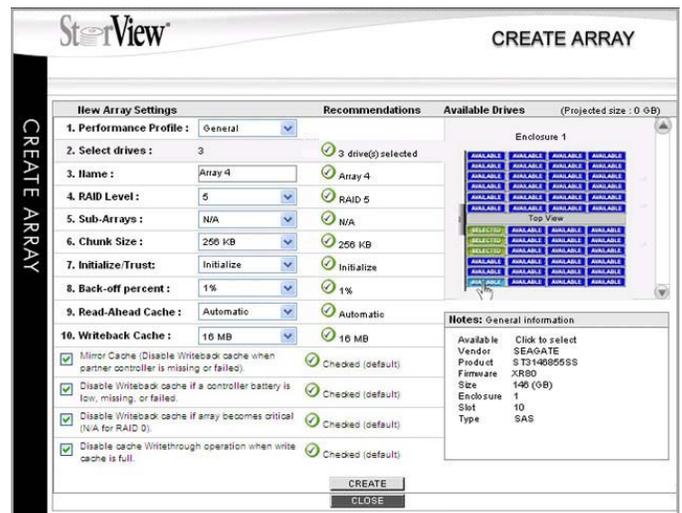


Fig. 29 - Create Array Window

has the “Available” icon displayed. The icon will change to “Selected.” As you select drives, the projected size of the array is displayed in the upper right corner of the window.

4. Enter a name for your array. You may use up to 32 characters (ASCII).

**NOTE:** The performance wizard at this point has pre-configured all the settings. If you are satisfied with the settings continue with step 11, otherwise continue with the steps below to manually configure your array parameters.

5. Select the RAID level for the array. Click the pull-down menu and choose from the available levels. These are based on the number of drives selected, refer to the “Drive Requirements” table at the beginning of this chapter.

(For RAID 50 arrays.) Create the sub-array(s). From the pull-down menu select the number of subarrays you wish to create for this array. When you click the “Create” button you will get a warning if you have selected more sub-arrays than allowed for the number of drives chosen. Reduce the number of sub-arrays.

**NOTE:** The Notes section under the Available Drives section will display information to help you with drive selection.

6. Choose the chunk size. Click the pull-down menu and select a chunk size (64K, 128K, or 256K). Chunk size is the amount of data that is written on a single drive before the controller moves to the next drive in the stripe.

For RAID level 0, 1, or 10 arrays, choose the correct size from the tables at right. For RAID 5/50/6 arrays, to achieve optimum write performance you should consider setting the chunk size based on the specified number of drives for a Full Stripe Write, for detailed information see section 5.2.2. [Chunk Size Details](#).

The idea behind optimum performance is you want to do as many full stripe writes as possible.

**NOTE:** The controller firmware will automatically set the chunk size if a smaller chunk size is selected rather than the size recommended for the number of drives and specified RAID level selected.

7. Select to Initialize the array. The default setting is to initialize, so you will just verify this setting. If it is not set, click the pull-down menu and choose Initialize. Initialization will begin automatically in the background once the array is created. You will have an option to stop or pause the initialization from the Main screen. If you Stop an initialization, the array will be trusted, see note below. As you create additional arrays, they too will begin initializing. The maximum number of arrays that can be initial-

Table 11- RAID 0

RAID 0				
# of drives	1	2	3	4+
Min. chunk size	256K	256K	128K	64K

Table 12- RAID 1 & 10

RAID 1	RAID 10			
# of drives	2	4	6	8+
Min. chunk size	256K	256K	128K	64K

ized in parallel is sixty-four (64).

---

**NOTE:** The Trust Array option should only be used in very special circumstances, see [8.2.7. Trust an Array](#).

---

8. Choose the “Back-off Percent” (reserved capacity) for the drives. The default is 1%. This determines how much drive capacity to reserve for future capacity expansions or to enable replacement drives of lesser capacity sizes to be used. The back-off percent option is not applicable to non-redundant array types.

A RAID 0 array is a non-redundant type of array and gains no benefit from establishing a reserve capacity. The values range from 0, 0.1, 1 through 9, and 90 percent.

- This determines how much drive capacity to reserve for future capacity expansions or replacement drives. This is not applicable to RAID 0 arrays.

---

**NOTE:** The back-off percent option is not applicable to non-redundant array types. A RAID 0 array is a nonredundant type of array and gains no benefit from establishing a reserve capacity.

---

9. Set the Read-Ahead Cache threshold. The choices are automatic, disabled, and four pre-determined sizes (256KB, 512KB, 1 MB, and 2 MB). Selecting Automatic, which is the recommended and default setting, allows the controller to determine the optimum size. Selecting Disabled will turn off the Read-Ahead Cache. Select from one of the predetermined sizes to optimize the read performance based on your data patterns.

---

**NOTE:** The Writeback Cache and Read-Ahead Cache thresholds work together to allow fine tuning of the I/O performance

---

10. Set the Writeback Cache options.

Click the pull-down menu to select from Disabled, or choose one of the pre-determined cache threshold sizes (1MB, 2MB, 4MB, 8MB, 16MB, 32MB, 64MB, 128MB, 256MB or MAX “MB”), see [5.2.1. Configuring Array Writeback Cache](#).

There are four additional options to be set: Writeback Cache: “Mirror Cache (Disable Writeback Cache when partner controller is missing or failed), [Checked - Default]” “Disable Writeback Cache if a controller battery is low, missing, or failed, [Checked - Default]” “Disable Writeback Cache if array becomes critical (not applicable to RAID 0), [Checked - Default]” and “Disable cache Writethrough operation when write cache is full. [Checked - Default]” Contact your supplier to obtain whitepapers on optimization and performance tuning, which provides more detail on these settings.

The Writeback Cache is used to optimize the write performance specific to your data patterns. In general, larger cache sizes will increase the write performance but may lower simultaneous read performance. The recommended size is 16 MB. The strategy of write operations results in a completion being sent to the host operating system as soon as the cache receives the data to be written. The disk drives will receive the data at a more appropriate time in order to increase controller performance and optimization of drive head seek. Refer to section [5.2.3. Performance](#).

---

**NOTE:** If you change any profile settings, you will get the following message at the bottom of the window.

---

[continued on the next page >>](#)

11. Click the CREATE button to complete this process.
12. You will see a confirmation message that the array was successfully created, click the CLOSE button.
13. Click the CLOSE button at the bottom of the Create Array window.

You can monitor the array initialization at the Main screen by observing a progress bar which appears under the array name displaying the percent complete. Also, in the enclosure front view, the affected disk drives being initialized display an animated icon indicating their initialization status.

You can stop or pause the Initialization process if you wish by clicking on the link located to the right of the progress bar.

Stopping the initialization will cause your array to be trusted. If you stop the initialization, you want to return at a later time to initialize the array, before you begin to use the array in a 'live' environment, see [5.3. Initializing the Array](#). A trusted array is indicated on the main screen with the following icon , adjacent to the specific array.

Pausing the initialization will halt the process until the resume option is selected. You can change the amount of processor time that is dedicated to the initialization for better initialization performance, see [7.2.2. Fault Tolerance](#).

**NOTE:** Some features are not available on older firmware based systems. In those cases, the unsupported feature will not appear in the user interface.

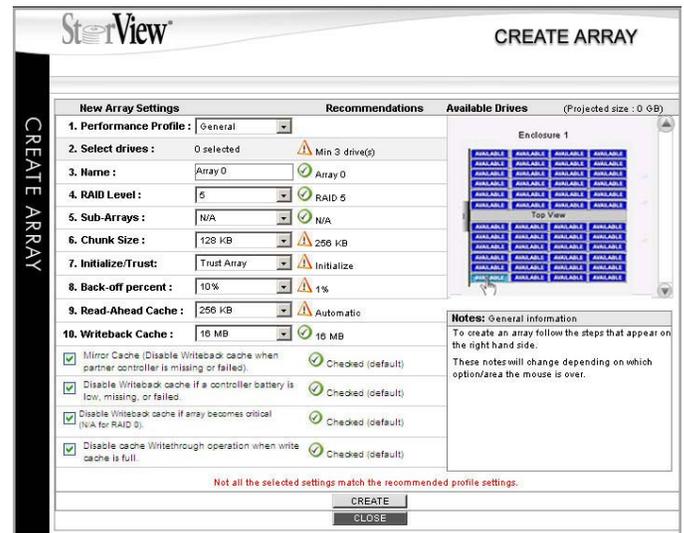


Fig. 30 - Monitoring the Initialization Process at the Main Screen

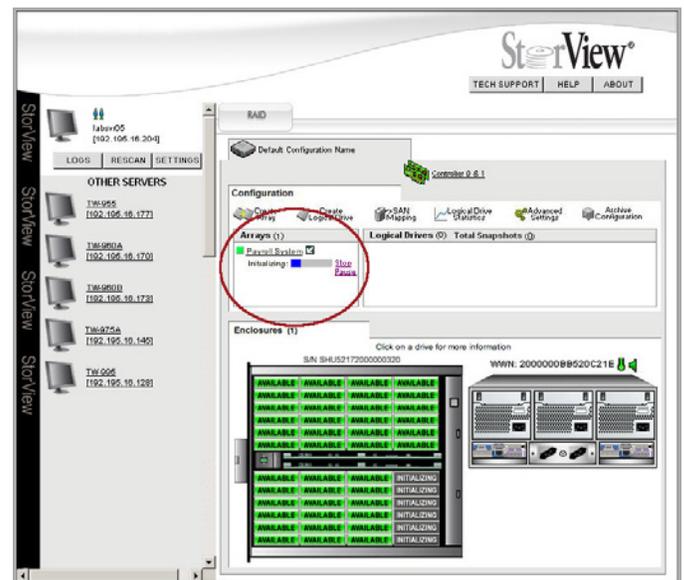


Fig. 31 - Reminder Message – Profiles Settings Mismatch Example

### 5.2.1. Configuring Array Writeback Cache

In a writeback cache operation, data is sent to the controller from the host and before actually sending the data to the drives, the controller immediately responds to the host confirming the data was received and written to the disk (even though the data may have not been written to the disk). The host may then send more data. This can significantly increase performance for host systems that only send a low number of commands at a time. The controller caches the data, and if more sequential data is sent from the host, it can cluster the writes together to increase performance further. If sufficient data is sent to fill a stripe in RAID 5, 6, and 50 configurations, the controller can perform a Full Stripe Write, which significantly reduces the write overhead associated with RAID 5, 6, and 50.

Disabling writeback cache ensures that the data is sent to the drives before status is returned to the host. With writeback cache enabled, if a short term power failure occurs, the battery back-up unit provides adequate power to ensure that cache is written to disk when the power is restored.

In duplex operations, the cache is mirrored to both controllers which provides further redundancy in the event of a single controller failure. Mirrored cache is designed for absolute data integrity. The cache in each controller contains both primary cached data for the disk groups it owns, and a copy of the primary data of the other controllers. Mirrored cache ensures that two copies of cache exist on both controllers before confirming to the operating system that the write operation has completed.

Normally, write-intensive operations will benefit from the higher performance when writeback cache is enabled on that array. Read-intensive operations, such as a streaming server, may not benefit from writeback cache.

A new performance option has been added which manages write-through operation when write cache is full. The default value is checked which disables the option providing typical workloads a performance improvement. When enabled, this option will force all write operations to go through the writeback cache and write command sorting provided thread balancing is disabled. The thread balancing option is found in the Performance Options window accessed through the Advanced Settings window.

---

### 5.2.2. Chunk Size Details

Typically, a Microsoft Windows access is 64 KB, therefore a stripe size of up to 1 MB would mean the controller has to cluster 16 commands to perform a full stripe write (actually 17 because of alignment).

If you were to use a larger stripe size, then you run the risk of not being able to cluster sufficiently for the application.

In cases where you are performing larger writes to the controller,

then you could go up to 4 MB for a stripe size, since you have more data to cluster.

It is recommended to keep the stripe size to 1-2 MB or less for general use, perhaps increasing it for specific applications such as large sequential accesses. This stripe size is actually the sub-stripe size in RAID 50 cases. A 4+1 array (4 data and 1 parity drives) with a 256 KB chunk has a 1 MB stripe size, as does an 8+1 array with 128 KB stripe size, and an 8+2 array with 256 KB chunk size.

Although 8+1 gives an even stripe size, this does not really matter for an operating system (OS) that writes in significantly smaller chunks. So, 8+1, 10+1, or 11+1 with a chunk size of 64 KB would be fine for an OS that does 64 KB writes. If the operating system can do much larger writes, then you may want to increase the chunk size. With writeback cache enabled, the controller can cache data and perform full stripe writes.

For a single enclosure example - best all around performance:

11+1 with 64 KB chunk size would yield a 704 KB stripe  
(11\*64=704)

10+2 with 128 KB chunk size would yield a 640 KB stripe  
(5\*128=640)

For most sequential accesses, the difference may not be too noticeable for different configurations. However, for large block random writes, it can help if the stripe size is similar to the I/O size. This allows the controller to perform a partial full stripe write, where it has most of the data for a full stripe, and can just read some from the drives to complete the stripe.

While Microsoft Windows does 64 KB accesses, these are not aligned. If the chunk size is 64 KB, most accesses crosses a chunk boundary and therefore involves two read/modify/write commands instead of one. However, since the controller can cluster, this problem is somewhat offset since the controller usually can cluster sufficiently to do full stripe writes. If it is completely random 64K access on Microsoft Windows, then a 64 KB chunk is not the best, rather 128 KB or 256 KB is better to minimize the number of commands that cross chunk boundaries.

Larger chunk sizes should be used if the operating system is writing large blocks, or with large sequential writes where the controller can cluster sufficiently. SATA disk drives have been found to perform significantly better as the I/O size is increased, so using as large a chunk size as feasible is better for these drives.

---

### 5.2.3. Performance

Fine tuning the arrays for optimum performance is a complex task that requires an understanding and analysis of your data structure, types of I/O, host system and on a per array basis when the settings for each disk array are created.

In addition to the write cache options, the RAID controller provides additional options to further tune your system for the best possible performance. Contact your supplier for a series of performance related whitepapers which will guide you through the Performance Options settings and how the different settings can affect your I/O performance dramatically. Refer to section [7.4. Advanced Performance Options](#).

---

## 5.3. Initializing the Array

Initializing an array clears all the data from the drives. This ensures the validity of the data and parity stored on the array member drives.

Two features of initialization are background and parallel. Once the array is created, initialization automatically begins in the background. While initialization is in progress, logical drives can be created and the disks are made immediately available to the operating system at which time data can be loaded.

As array's are created and the initialization begins, you can have up to sixty-four (64) arrays initialized in parallel at the same time. You may also choose to stop the initialization, or pause an initialization and then resume it at a later time. The controls for managing are displayed on the Main screen next to the "Array\_Name" after the initialization has started. If you Stop an initialization, the array will automatically be trusted, see note below.

The array can be initialized at a later time in which you could choose the option to Trust. This option should only be used in environments where the user fully understands the consequences of the function.

The trust option is provided to allow immediate access to an array for testing application purposes only.

---

**NOTE:** A trusted array does not calculate parity across all drives and therefore there is no known state on the drives. As data is received from the host parity is calculated as normal, but it occurs on a block basis. There is no way to guarantee that parity has been calculated across the entire stripe. The parity data will be inconsistent and so a drive failure within a trusted array will cause data loss. Before you use a trusted array in a "live" environment, you must initialize it. A trusted array is indicated on the main screen with the following icon, adjacent to the specific array.

---

1. Locate and click on the <Array\_Name> in the Array section on the Main screen you wish to initialize. This will open the Array Information window.
2. From the Array Information window, click the INITIALIZE button.
3. You will be prompted to enter your password to confirm you wish to initialize the array. Type your password and click GO. A confirmation message appears indicating the success of the operation, click the CLOSE button.
4. Click the CLOSE button on the Array screen. From the Main screen you can monitor the initialization.

Placing the mouse pointer over the progress bar will display the percent complete for the initialization progress in a pop-up window. The drive member icons of this array will change to an animated icon indicating the array is initializing.

You can stop the initialization process, if you wish, by clicking the Stop link located to the right of the progress bar.

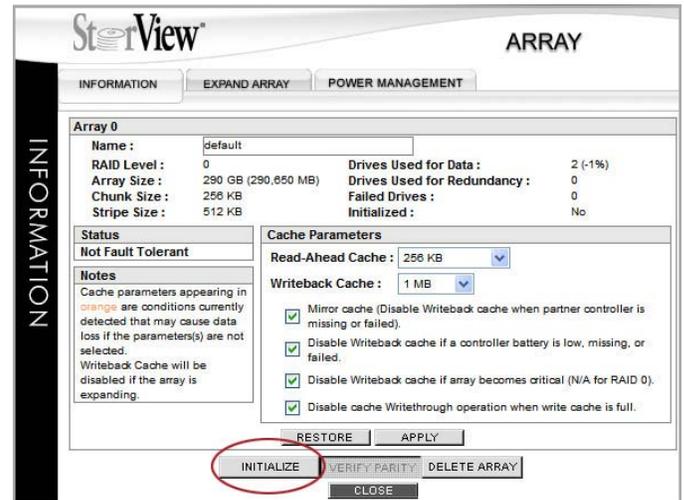


Fig. 32 - Array Information Window

### 5.3.1. Pause/Resume the Initialization

You can temporarily pause the initialization process, and resume the process at a later time.

#### 5.3.1.1. Pause Initialization

1. Click the Pause link located just to the right and below the progress bar. The "Pause" link will change to "Resume" and the progress bar will stop its forward motion.

#### 5.3.1.2. Resume Initialization

1. Click the Resume link located just to the right and below the progress bar. The initialization will continue from the point where it was paused.

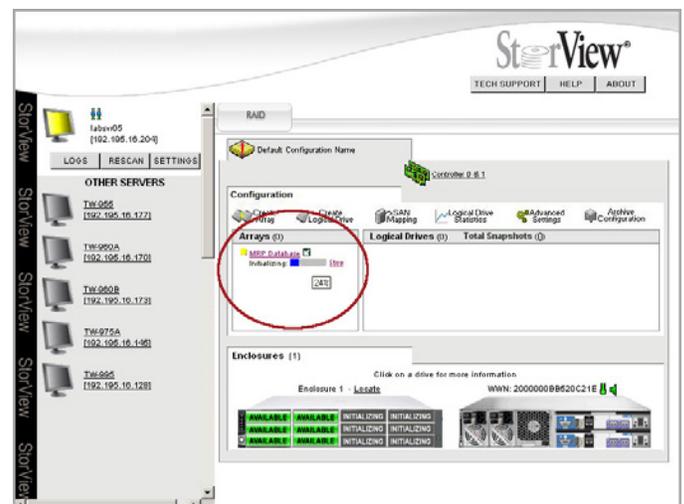


Fig. 33 - Monitoring the Initialization Progress

## 5.4. Adding Hot Spare Drives

Hot-swappable spare disk drives are supported. In the event of a drive failure, the controller will use either a defined global spare or a dedicated spare to replace a failed drive that is a member of a fault tolerant array. The process of configuring your redundant arrays includes assigning one or more drives as hot spares.

Global spares are not assigned to a specific array and they can be used by any array as the replacement drive, of course that is provided the spare drive is equal to or greater than the capacity of the array member drives.

A dedicated spare is assigned to a specific array and can only be used by that array. It should be the same size or greater as the other drive's that comprise the array.

Spare Drive Use Rules:

- Drives must be equal to or greater than the capacity of the array drive members.

**NOTE:** You may get a variety of messages, warnings, or failure notices when attempting to use a un-supported drive.

### 5.4.1. Assigning a Global Spare

1. From the Main screen, click a drive icon within the enclosure drive view indicating "Available" that you wish to make a global hot spare.

**NOTE:** There must be at least one drive online and available that meets the rules from above to be assigned as a hot spare, and a configuration must exist (at least one array defined).

2. From the Drive Information window, click the MAKE SPARE button.
3. A pop-up window will appear, select Global Spare from the drop down menu.
4. Click the CREATE button. You will see a confirmation window indicating the process was successful, click the CLOSE button.
5. Click the CLOSE button on the Drive Information window.

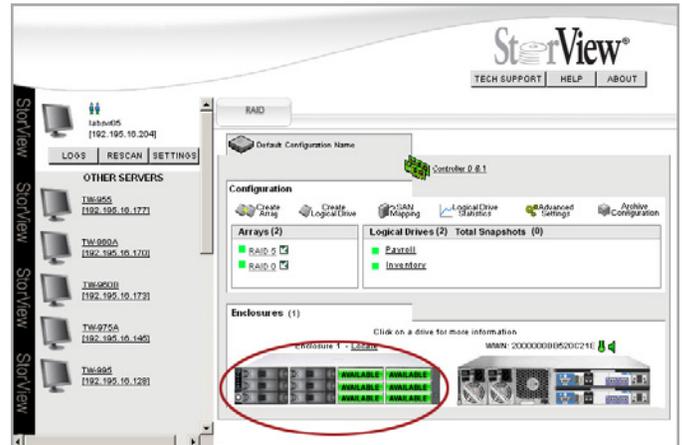


Fig. 34 - Main Screen

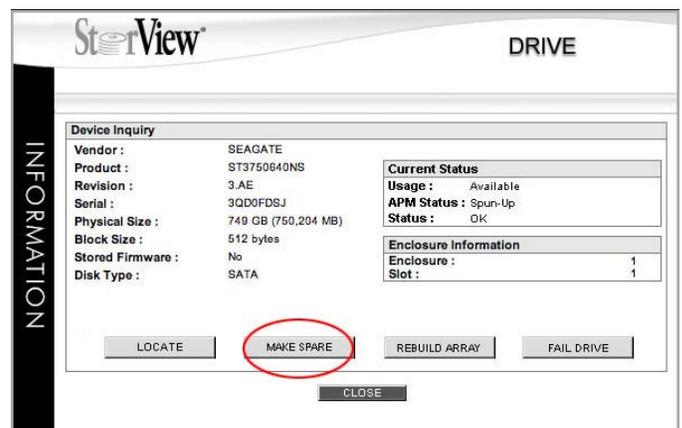


Fig. 35 - Drive Information Window

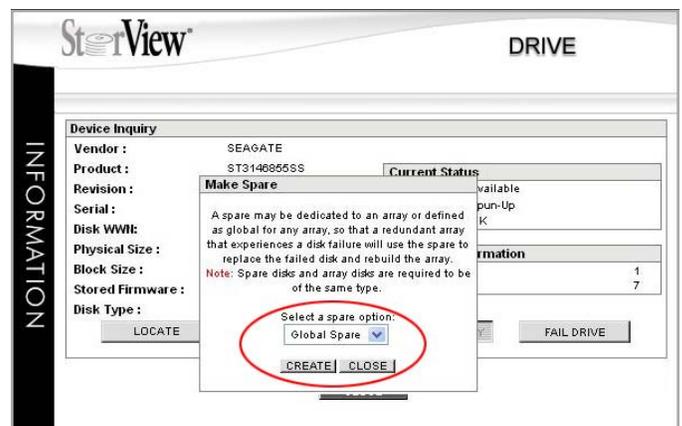


Fig. 36 - Selecting for a Global Spare

### 5.4.2. Assigning a Dedicated Spare

1. From the Main screen, click a drive icon in the enclosure front view indicating "Available" that you wish to make a dedicated hot spare.

**NOTE:** There must be at least one drive online and available that meets the rules for spare drives in section 5.4. Adding Hot Spare Drives to be assigned as a hot spare, and a configuration must exist (at least one array defined).

2. From the Drive Information window, click the MAKE SPARE button.
3. A pop-up window will appear, click the drop-down menu and select the array to which you wish to assign the dedicated spare.
4. Click the CREATE button. A confirmation window will appear indicating success, click the CLOSE button.
5. Click the CLOSE button on the Drive Information window.

**NOTE:** Only arrays that the spare drive is large enough to replace any member drive of that array or of the same drive type will be displayed in the pull down menu. For example, if you have two arrays one created using 500 GB drives (array 0) and one using 750 GB disk drives (array 1). If you have a 500 GB spare drive that you attempting to assign to an array, only Array 0 will be displayed because the drives in array 1 are of equal or lesser capacity drives then the spare. However, if you have a 750 GB spare drive both Array 0 and Array 1 will be displayed since the 750 GB spare is equal to or greater than any drive in either array.

### 5.4.3. Removing a Spare

This operation will remove the designation of the drive as a global or dedicated spare. The drive will then become online and available for other uses.

1. From the Main screen, click on a disk drive labeled "DEDICATED SPARE" or "GLOBAL SPARE" in the enclosure drive view that you wish to remove as a spare. The Drive Information window will open.
2. Click the REMOVE SPARE button. A confirmation window will appear indicating success, click the CLOSE button.
3. Click the CLOSE button on the Drive Information window.

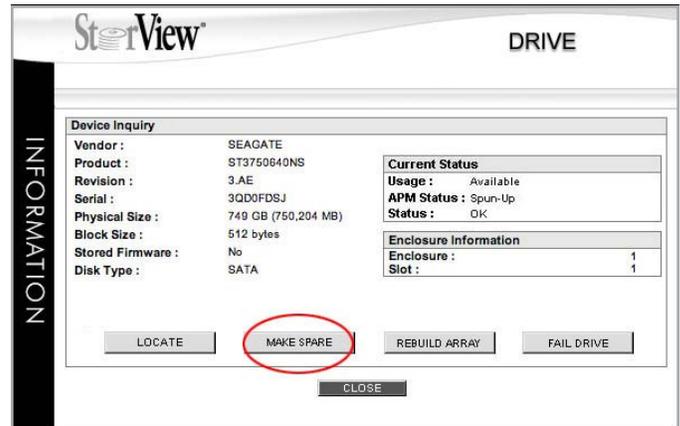


Fig. 37 - Main Screen

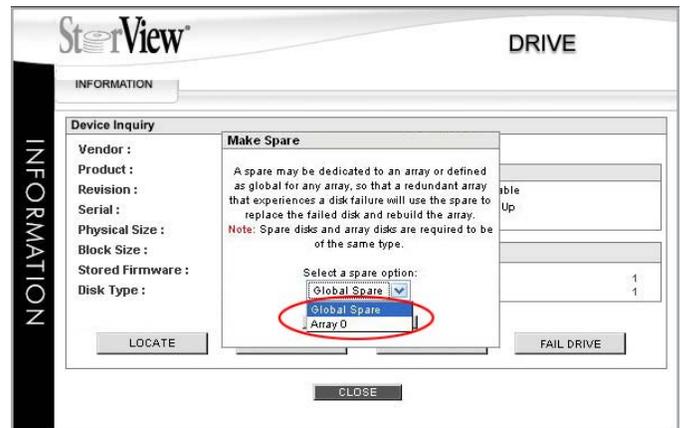


Fig. 38 - Drive Information Window

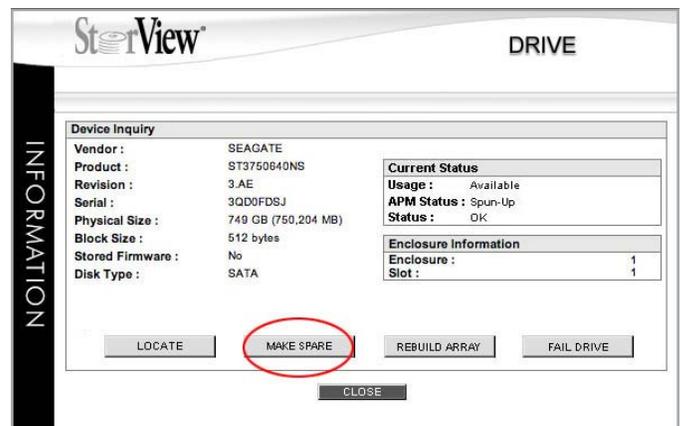


Fig. 39 - Drive Information window

### 5.4.4. Auto Spare

The Auto Spare option, when enabled, will automatically cause a replacement disk drive when inserted to be used as a dedicated hot spare for the failed drive and its array. When a new drive is inserted in place of the failed drive, a rebuild operation will begin automatically using the new drive. This option is useful when a global or dedicated hot spare drive is not assigned and you have a fault tolerant array that experiences a drive failure. This option allows the user to just insert a replacement drive and the rebuild will begin, instead of opening the Drive Information window for the replacement disk drive and assigning it as a hot spare. Refer to the Spare Drive Rules in section [5.4. Adding Hot Spare Drives](#).

1. To enable this feature, click the Advanced Settings icon on the Main screen Tool Bar. The Advanced Settings window will open.
2. Place the mouse pointer on the check box next to the Auto Spare parameter and click to place a check mark enabling the feature.
3. Click the APPLY button and click the CLOSE button on the confirmation window when it appears.
4. Then click the CLOSE button on the Advanced Settings window.

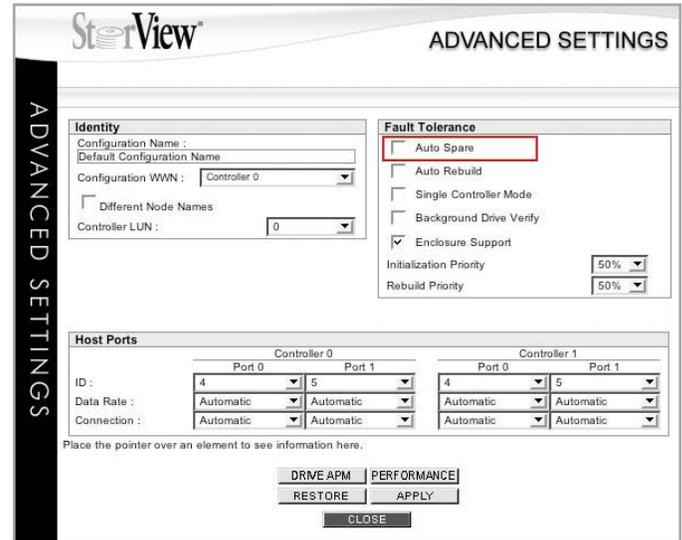


Fig. 40 - Advanced Settings Window

## 5.5. Create the Logical Drive

To complete the process of configuring your storage system, you will need to create one or more logical drives and assign them a LUN (logical unit number). This presents the “logical drive” or this LUN to the host operating system as a disk. You can also restrict access to the logical drive by assigning which controller port the logical drive is made available on. Only those host systems with access to the specific controller port(s) will have access to the logical drive.

A logical drive is defined or created from regions of an array, a whole array, or a combination of regions of different array(s) that can be made available as a single disk to one or more host systems. If you are creating a logical drive greater than 2 TB, please refer to your operating and file system documentation to verify it supports such sizes.

You may wish to avoid choosing combinations of a region from one array and from another array to create your logical drive. This will reduce data fragmentation.

1. From the Main screen click on the Create Logical Drive icon in the Tool Bar. The Create Logical Drive window will open.
2. Select the region or regions you wish to use for your logical drive from the list “Select Which Array(s) to use” window. You may hold the <Shift> or <Ctrl> key down to make multiple selections.
3. Enter a name for your logical drive, you may use up to 32 characters. The default names for logical drives follow the format “LDx.” Only 12 characters plus an ellipse is displayed when the name is longer. Holding the mouse pointer over the logical drive name on the Main screen will show the complete name in a popup.
4. Enter the size in GB’s for the logical drive capacity. As you select your regions, the maximum size is displayed to the right of the “Size:” field. You may use all or some of these regions for this logical drive. If you are creating a logical drive greater than 2,198 GB (2 TB), please refer to your operating and file system documentation to verify it supports such sizes.

---

**IMPORTANT INFO:** If you intend to use the snapshot feature and choose to have access to this logical drive from your operating system or VSS, you must SAN LUN Map the logical drive. You may do this before or after the snapshots have been taken.

---

5. Choose the desired Block Size. (Two sizes are available, 512 bytes and 4096 bytes.)
6. Select the LUN number for the logical drive from the “Mapped to” drop down menu.
7. Select the Controller Ports you wish to make the logical drive

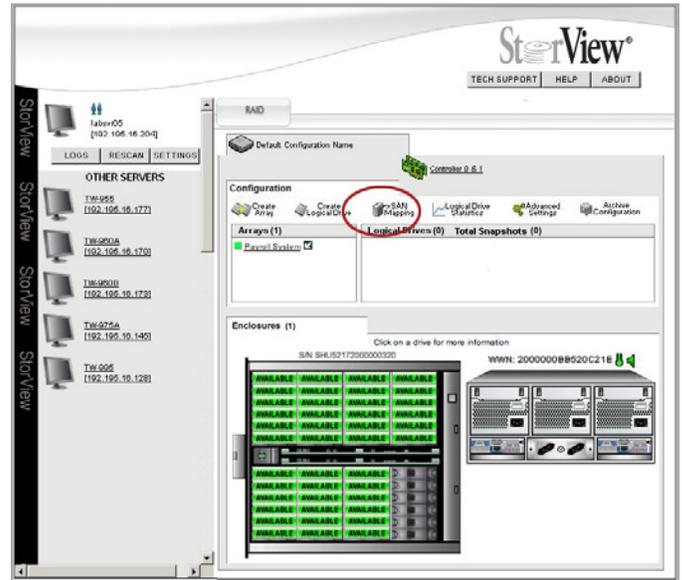


Fig. 41 - Main Screen

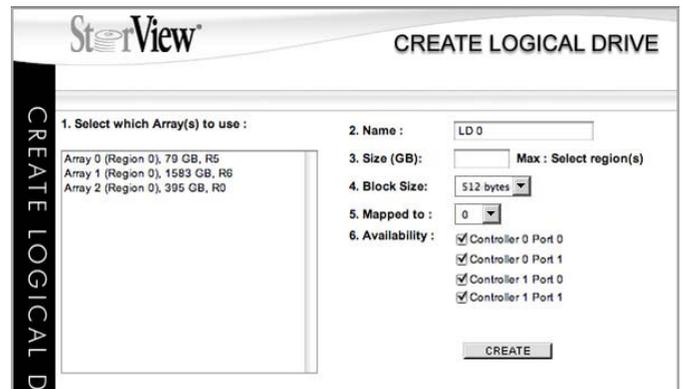


Fig. 42 - Create Logical Drive Screen

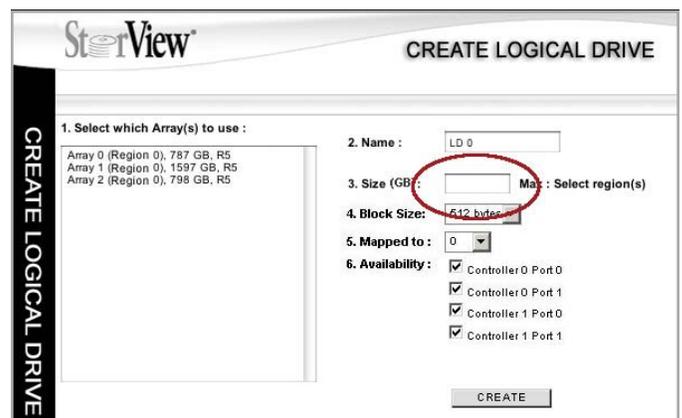


Fig. 43 - Defining the Logical Drive Capacity Screen

available through. Place a check mark next to the desired controller ports displayed.

The choices are: C0P0, C0P1, C1P0, and/or C1P1. Where C0 represents the left or lower controller and C1 represents the right or upper controller, and P0 indicates the port marked Host 0, FC Host 0, or "IN" 0, and P1 represents the port marked Host 1, FC Host 1, or "IN" 1 on each controller. There is a one-to-one relationship between the controller port selected and a data cable connected to that port. Ensure that the ports you select are properly cabled.

8. Click the CREATE button to finish creating the logical drive.

You will receive a screen prompt that the command was successful, click the CLOSE button. If the command was unsuccessful, review the settings for incorrect parameters and hardware for operational status.

9. You may continue to create more logical drives or exit by clicking the CLOSE button. In most storage system environments, creating the logical drives, assigning them their logical unit number (LUN) and setting the availability is sufficient to meet the requirements for setup. For more advanced and complex systems using storage area networks you may wish to perform the more advanced SAN LUN Mapping, see chapter 9. [SAN LUN Mapping](#). Otherwise access your operating system to make the new drives available for use.

---

## 5.6. Saving the Configuration

Saving the configuration information is a very useful feature of StorView. When you create or make changes to arrays, logical drives, hot spares, SAN LUN Mappings, feature licenses, or change the parameters of a specific controller setting, a file is written (known as the configuration) to all the disk drives that are members of the array. StorView has the ability to capture that file allowing you to save it to an external file. Should the situation occur where a configuration is cleared, you are instantly able to re-establish your storage system by restoring the configuration from the external file. It is recommended to periodically save the configuration.

---

**CAUTION:** If you cannot restore the configuration exactly as it was, you will not be able to restore access to the data and it will be lost.

---

### 5.6.1. Saving the Configuration

1. From the Main screen Tool Bar click the Archive Configuration icon.

The Configuration Archival Operations window appears.

2. Click the SAVE button. You may click the CLOSE button to cancel and return to the Main screen.
3. Click the DOWNLOAD button to continue with saving the configuration file, or click the CANCEL button to cancel.

continued on the next page >>

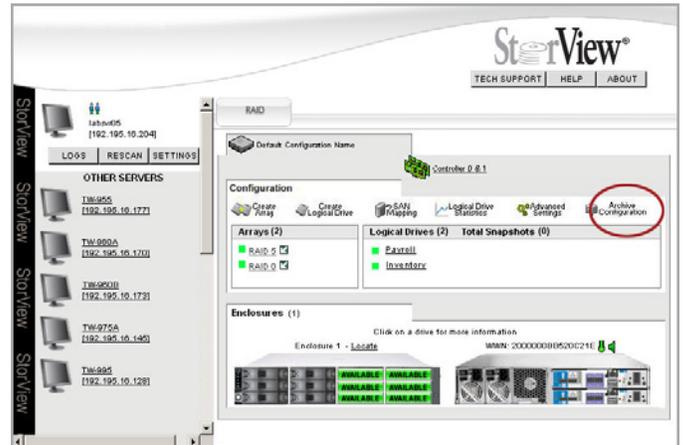


Fig. 44 - Main Screen



Fig. 45 - Configuration Archival Operations Window



Fig. 46 - Save Configuration Download Window

4. You are presented with the browser's standard "File Download" screen. Click the SAVE button to continue, or CANCEL to quit.
5. Next, you are presented with the "Save As" screen. If you wish to use the default file name, select the directory and click the Save button, otherwise enter the name you wish to use and specify the directory, and then click Save. Click the Cancel button to exit without making any changes.



Fig. 47 - Save Configuration File Screen

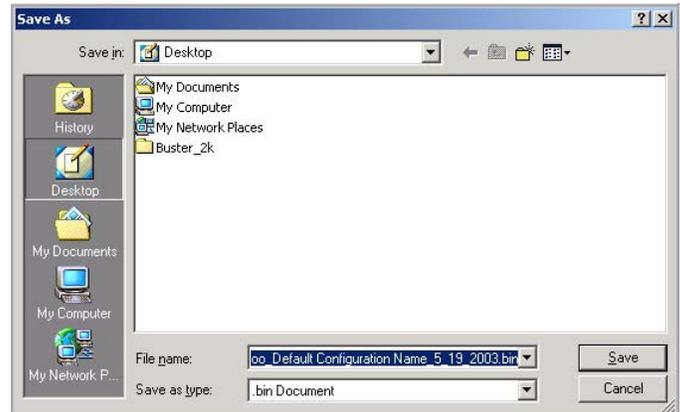


Fig. 48 - File Name Screen

## 6. Environments

### 6.1. Overview

The Controller Information window provides you with an instant view of the health status of the RAID Controller components and a mouse-over function for detailed information of the items listed in the Status section. User-controllable functions include updating RAID Controller firmware, updating the Expansion module firmware (Disk I/O module), synchronizing the date and time, resetting or shutting down the RAID Controller(s), and managing event log files.

### 6.2. Controller Environmentals

To view controller environmental conditions and manage controller environmental functions, click the Controller icon located just above the Tool Bar on the Main screen.

The Controller icon on the Main screen will flash red for a failure condition and yellow for an abnormal condition. If this occurs, click the icon and investigate the problem from the information provided. You may wish to investigate the Event log to verify one or more event(s) that changed the controller status, see section [13.6.6. Enclosure Events](#).

When the Controller Information window opens, the controller's status is displayed via status icons and a text message.

By passing the mouse pointer over each item in the Status group (controller and battery), a pop-up window will appear with specific detailed data.

#### 6.2.1. Status

This group of items are applicable to the functional status of the RAID Controller. They include general controller status and battery status. Placing the mouse pointer over the item will display a pop-up window with more detailed information.

Status icons appear adjacent to each item in the group along with a text message. Status icon and text message conditions are defined as green - normal, yellow - warning, and red - failed.

When the mouse pointer is rolled over the Battery status icon, the pop-up display provides general information about the battery, that is whether it is charged, charging in process, or has a fault. A warning (yellow) icon indicates the battery is low but is charging. An error (red) icon indicates the battery is low but has charging for over 24 hours and is most likely defective.

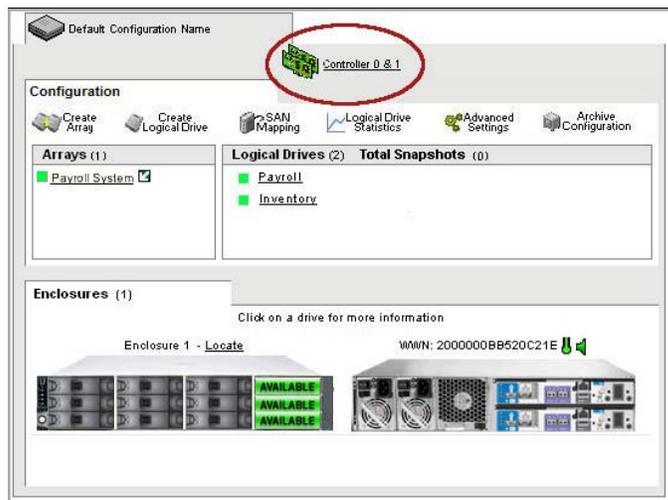


Fig. 49 - Main Screen – Arrays, Logical Drives, and Enclosure Sections

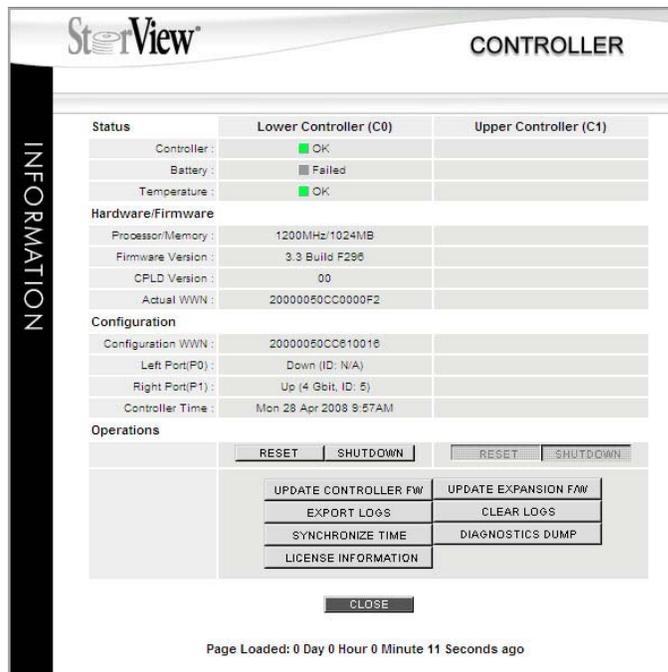


Fig. 50 - Controller Window with Pop-Up on Battery Status

### 6.2.2. Hardware/Firmware

This group of items are applicable to the controller’s physical memory and firmware. From this group you can determine the type of processor, onboard memory size, the firmware version, and the CPLD (Complex Programmable Logic Device) firmware version. The controller’s WWN (world wide name) is presented in this group.

### 6.2.3. Configuration

This group identifies the WWN (world wide name) assigned to the configuration, the speed of each port and the LUN assigned to the port. The configuration can be assigned the WWN of any controller preferably one of the controllers installed in the enclosure. For more information on the configuration WWN and port LUN assignment, see [7.2.1. Identity](#).

**NOTE:** Be careful when changing the Configuration WWN if you are using any of the Feature products licensed, since changing the WWN may invalidate the enabled feature such as Snapshot.

### 6.2.4. Operations

These items include a group of buttons that allow the user to reset and shutdown each controller individually, update the controller firmware (see [6.2.4.1. Updating the Controller Firmware](#)), export the controller logs (see [6.2.4.5. Exporting the Controller Logs](#)) and clear the log files (see [6.2.4.4. Clearing the Controller Logs](#)).

The Shutdown button under each controller will cause the controller to execute a graceful shutdown by flushing the cache first then executing the shutdown command.

**NOTE:** When choosing to shutdown the RAID Controller(s) for an extended period of time, be sure to verify that the Cache Active LED on the RAID Controller is not illuminated indicating that data is still being held in cache. If it is, wait until the controlling application has completed writing the data to the disk before continuing with the shutdown. This will ensure that the backup battery is not holding cache data and will prevent the battery from being drained.

The user can dump the controller’s diagnostic information into a file for use with technical support when troubleshooting a problem by clicking the DIAGNOSTICS DUMP button.

#### 6.2.4.1. Updating the Controller Firmware

This option provides the means to update the RAID Controller firmware.

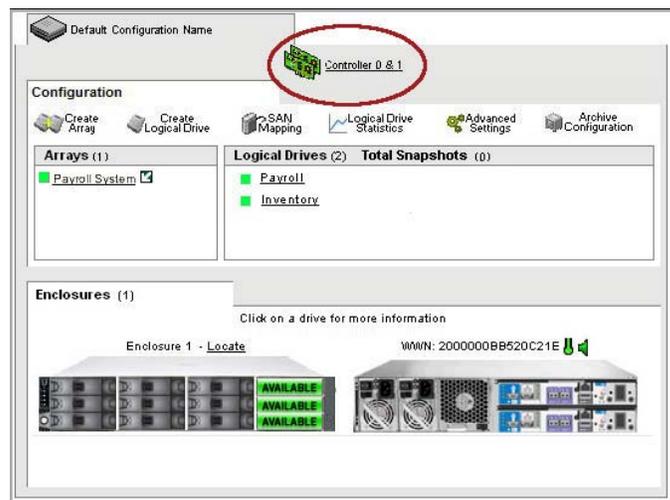


Fig. 51 - Main Screen – Arrays, Logical Drives, and Enclosure Sections

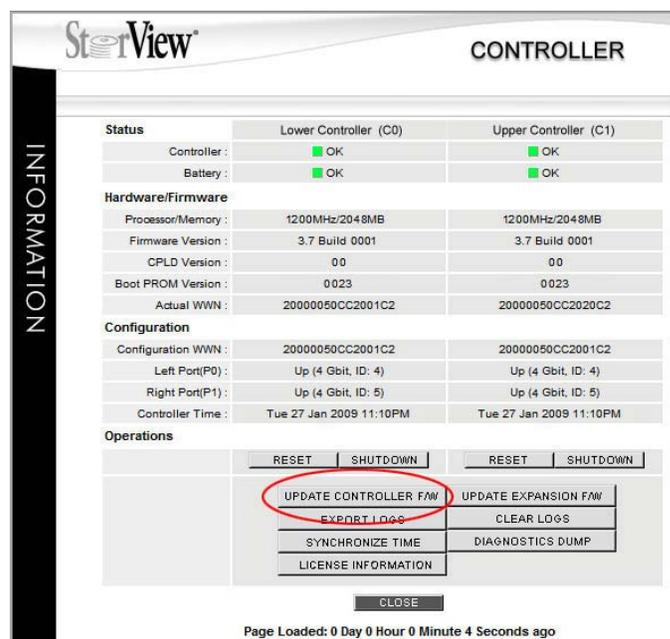


Fig. 52 - Controller Information Window

The controllers have the ability to automatically update their partners firmware in dual controller configurations, however the update process behaves differently under certain conditions. If one controller has a later version of firmware than its partner controller, during the startup process the later version firmware controller will automatically update the firmware on the lower version controller. If a partner controller fails in a dual controller configuration, when the failed controller is replaced regardless of its firmware version it will always be updated to match the surviving controllers firmware.

If you want to downgrade the firmware version you must shut down one controller and flash the operating controller. Then shut down the controller which was downgraded, bring the second controller up and flash its firmware to the lower version. Then start both controllers and resume operations.

1. Click the Controller icon located just above the Tool Bar.

The Controller Information window will open.

2. Click the "UPDATE CONTROLLER FW" button.
3. Locate the firmware file, click the Browse button. The file will be named similar to "EGCN-AxxxxAxx.xx.xxx.bin."

The file will be made available from the customer service representative or technical support.

4. After locating the file, click the UPLOAD button.

You will see an acknowledgement window appear indicating the status of the update, followed by the controller automatically restarting.

**NOTE:** For more information on License Information, see [1.3.2. Feature License](#).

### 6.2.4.2. Update Expansion Firmware

**IMPORTANT INFO:** The firmware on the Disk I/O modules must be the same version on ALL Disk I/O modules installed in ALL expansion enclosures in that storage solution. See section [6.2.4.3. Determining the Version of Firmware on Expansion Modules](#).

**NOTE:** In dual controller configurations, you must shutdown one RAID Controller and physically remove it from the enclosure before performing this procedure. Update the Disk I/O module on the "live" controller, then shut that controller down and remove it. Bring the other controller back online and repeat the procedure on the second Disk I/O module.

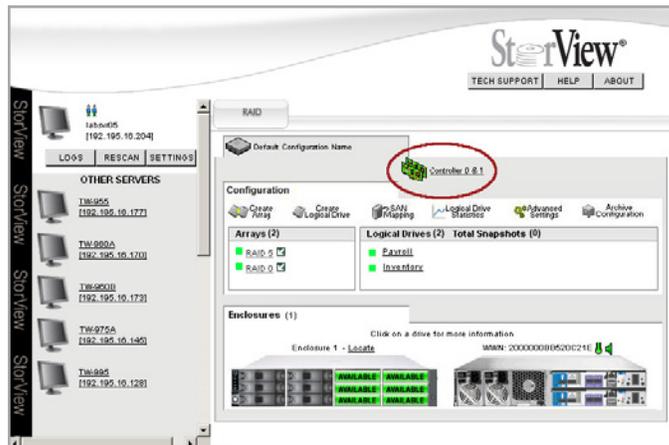


Fig. 53 - Controller Information Window

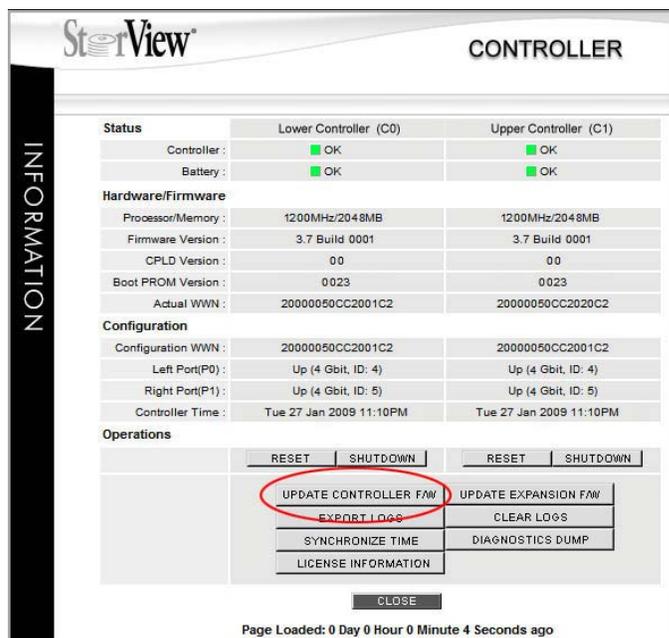


Fig. 54 - Controller Information Window

1. Click the Controller icon located just above the Tool Bar. The Controller Information window will open.
2. Click the UPDATE EXPANSION F/W button.
3. Locate the firmware file, click the Browse button. The file will be named similar to "cdfw\_1200\_SAS1.0\_01.07.bin."

The file will be made available from the customer service representative or technical support.

4. After locating the file, click the UPLOAD button.

You will see an acknowledgement window appear indicating the status of the update, followed by the expansion module automatically restarting. The RAID Controller will restart after the expansion module restarts.

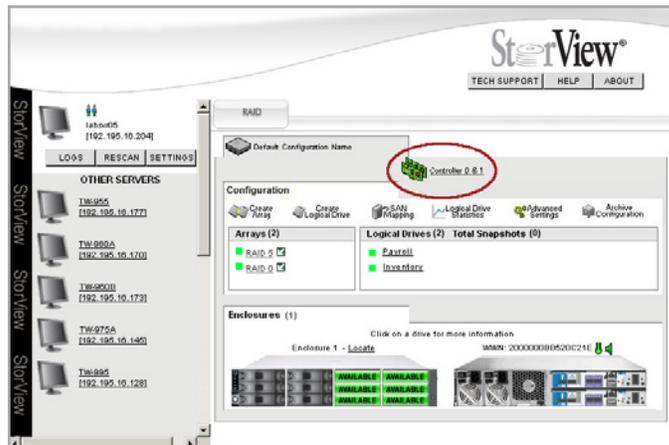


Fig. 55 - Main Screen

### 6.2.4.3. Determining the Version of Firmware on Expansion Modules

1. Click the Controller icon located just above the Tool Bar for the storage system you wish to examine. The Controller Information window will open.
2. Click the DIAGNOSTIC DUMP button.
3. Once complete, open the "DiagnosticsDump.log" file using a text editor.
4. Scroll down through the file and look for "List of Enclosures and Firmware Version" section and view the firmware version of each Disk I/O module.

**IMPORTANT INFO:** The firmware on the Disk I/O modules must be the same version on ALL Disk I/O modules installed in ALL expansion enclosures in that storage solution.

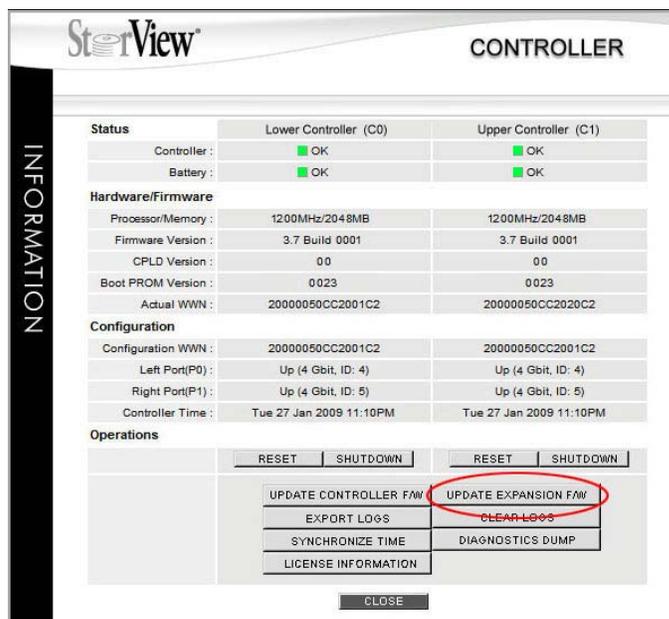


Fig. 56 - Controller Information Window

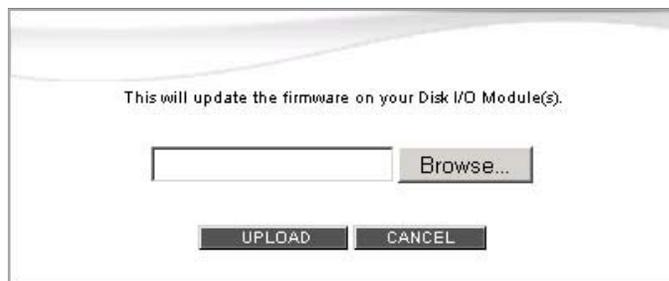


Fig. 57 - Main Screen

### 6.2.4.4. Clearing the Controller Logs

1. Click the Controller icon located just above the Tool Bar. The Controller Information window will open.
2. Click the CLEAR LOGS button.
3. A window opens prompting to confirm clearing the logs. Click the OK button.
4. Click the CLOSE button at the bottom of the Controller Information window.

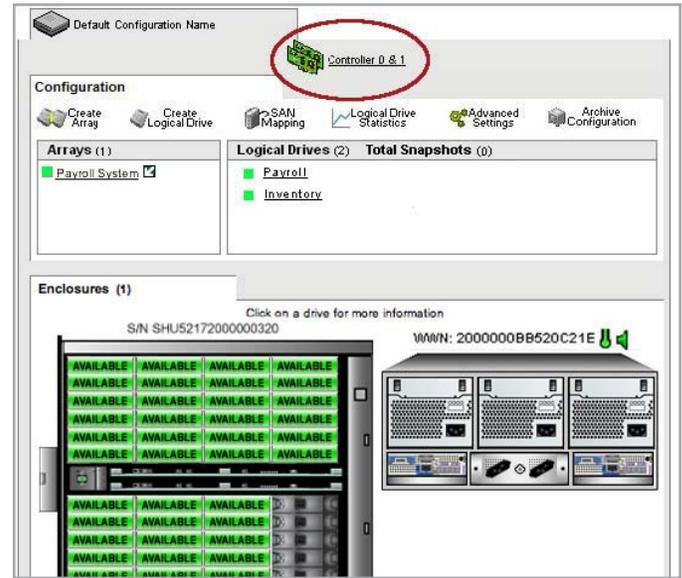


Fig. 58 - Main Screen – Arrays, Logical Drives, and Enclosure Sections

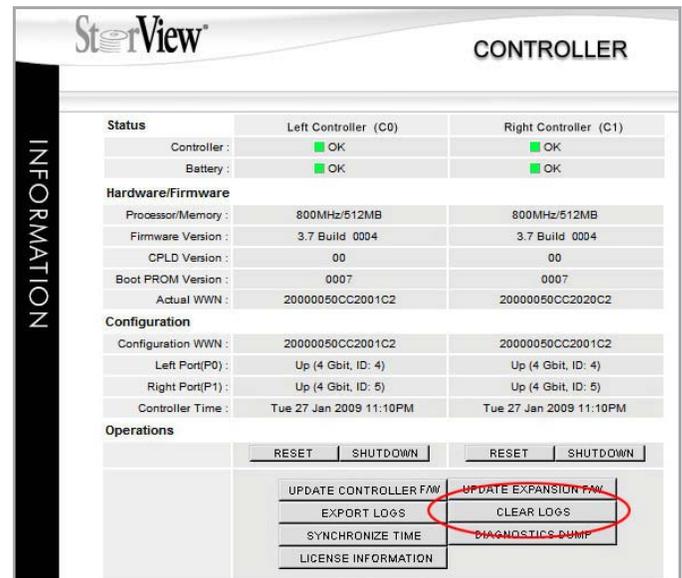


Fig. 59 - Controller Information Window



Fig. 60 - Controller Information Window

### 6.2.4.5. Exporting the Controller Logs

1. Click the Controller icon located just above the Tool Bar.  
The Controller Information window will open.
2. Click the EXPORT LOGS button.
3. Choose Save to Disk, enter a name and location for the file, then click the OK button.
4. Click the CLOSE button at the bottom of the Controller Information window.
5. Open the event log with a text editor program.

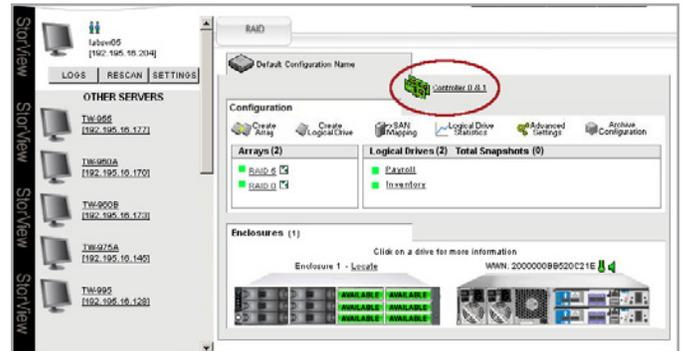


Fig. 61 - Main Screen – Arrays, Logical Drives, and Enclosure Sections

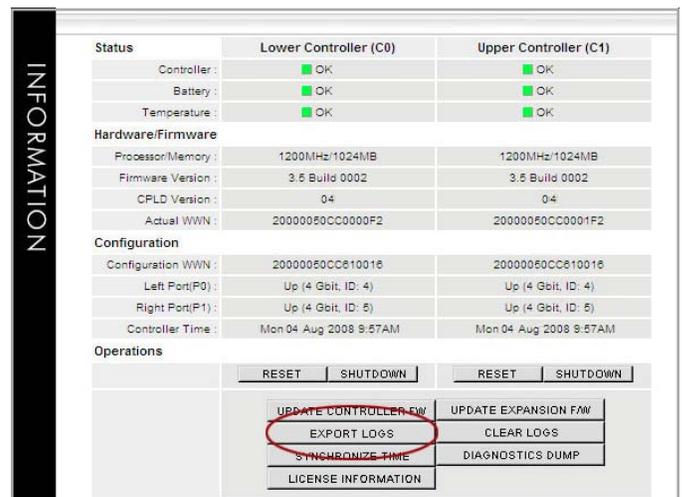


Fig. 62 - Controller Information Window

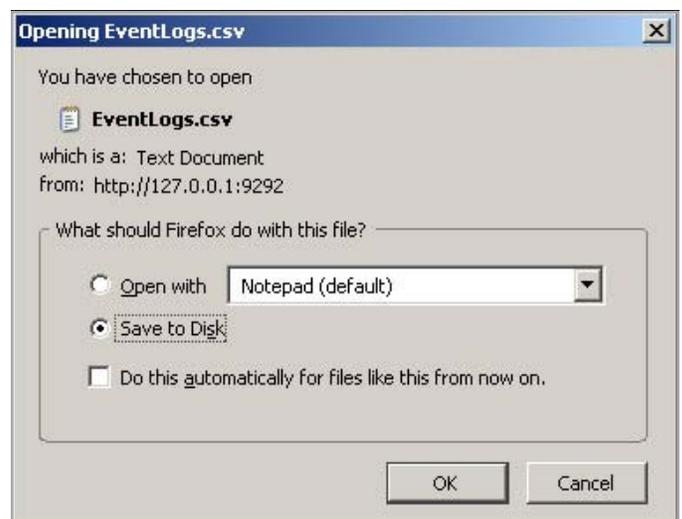


Fig. 63 - Confirmation window

## 7. Controller Advanced Settings

### 7.1. Overview

In order to customize your storage solution, you may want to make changes to the storage systems configuration parameters to optimize it for your application. This is accomplished through the Advanced Settings window. These functions include managing the identity, fault tolerance, host port settings, drive advanced power management and array performance tuning.

### 7.2. Advanced Settings

Click the Advanced Settings icon in the Tool Bar on the Main screen.

The Advanced Settings window will appear. From this window you are able to make changes to controller parameters, enable or disable Fault Tolerant features, and configure the controller's host ports.

The Advanced Settings window is divided into three sections; Identity, Fault Tolerance and Host Ports.

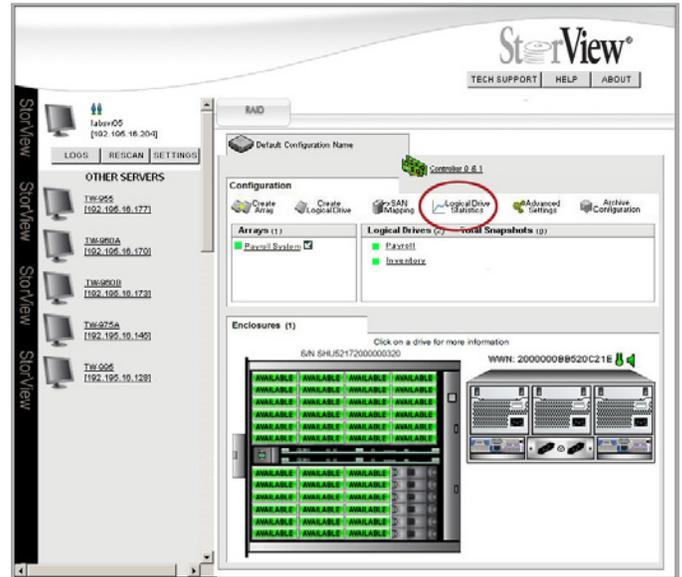


Fig. 64 - Main Screen – Arrays, Logical Drives, and Enclosure Sections

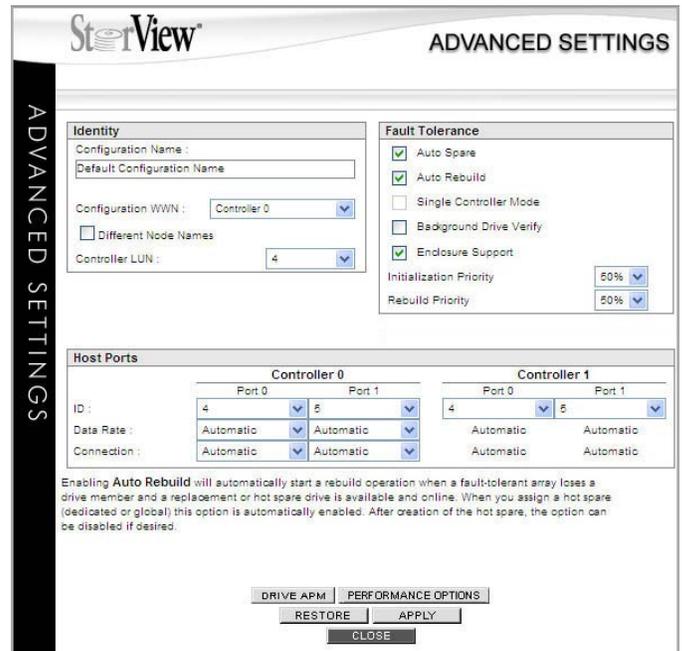


Fig. 65 - Advanced Settings Window

### 7.2.1. Identity

In the Identity section, you can make changes to the Configuration Name, assign the configuration the WWN of either controller, and set the LUN (logical unit number) for the controller.

The following describe the components of the Identity section:

- ◆ Configuration Name - This is the name you will assign to the configuration. The configuration contains all the information that defines the disk arrays, logical drives, SAN LUN Mapping, hot spare drives and controller specific settings.

If you wish to change the configuration name, enter the new name in the block provided. Click the APPLY button followed by the CLOSE button.

- ◆ Configuration WWN - This is the RAID Controller's WWN reported to the outside world to identify the StorView configuration.

If a new WWN is showing up in the list, it is possible for you to have a configuration WWN on your system where there is not a controller in the storage system with the same WWN. For example: if you have one or two controllers in your system with no configuration. Note there is no WWN in the Configuration WWN select box, only Controller 0/1. Then create an array. Then pull out Controller 0. (If you only had a single controller in the system then insert another physical controller. If you have dual controllers then just leave Controller 1 in there. Now go back and look at your Configuration WWN you will notice a different WWN for the configuration than the controller's WWN. The user can only set the Configuration WWN to Controller 0 or Controller 1's WWN.

However, if there is a stale WWN then it will show it in the list indicating to the user what it is currently set to. The Configuration WWN is what is reported to the outside world no matter what port you are plugged into the system. This way if you swap controllers (most likely because of a failure) your Configuration WWN will still report the same WWN as it did before so you will not have to change any mappings on your host or fibre channel switch.

If another controller was used to create the configuration, its WWN is displayed. You should assign the configuration WWN to the installed controller. In this case click the pull-down menu and select Controller 0 or Controller 1. Click APPLY and restart the StorView.

- ◆ Controller LUN - This option allows the user to set a specific LUN number or disable the Controller LUN. By default the Controller LUN is automatically assigned the next LUN number after the last logical drive.

In the event you have an operating system that is having a problem with the Controller LUN being displayed, click the pull-down selection and choose "Disabled."

- ◆ Different Node Name - Selecting this option allows the controller to report a different Configuration WWN for Port 0 and Port 1 (FC Host 0 and FC Host 1 connectors). Normally, when deselected, a host connected to either port will see the same Configuration WWN.

When enabled (selected) you will see a slightly different WWN for each port but the same Configuration name. This option is useful to users who are connecting the storage to a switch employing a fabric topology where the same WWN is not tolerated.

---

**IMPORTANT INFO:** If you are using the Snapshot feature, your feature license key is based on the Configuration WWN which is based on the RAID Controller's WWN used at the time of licensing. If you modify the Configuration WWN, the feature license key will no longer be valid and the snapshot function (feature) will be disabled.

---

### 7.2.2. Fault Tolerance

In the Fault Tolerance section, you can enable or disable controller features that improve its abilities to maintain a level of fault tolerance.

- ◆ Auto Spare - This option when selected allows the data to be rebuilt on the drive that is inserted into the slot from which the failed drive was removed. This is beneficial when a hot spare or global spare is not designated for a fault tolerant array and a drive fails in that array.
- ◆ Auto Rebuild - Selecting this option will automatically start a rebuild operation when a faulttolerant array loses a drive member and a replacement or hot spare drive is available and online. When you assign a hot spare (dedicated or global) this option is automatically enabled. After creation of the hot spare, the option can be disabled if desired.
- ◆ Single Controller Mode - When operating in the StandAlone mode (single controller configurations) selecting this option stops the controller from constantly checking for a partner controller. When operating a duplex Active-Active configuration, deselect this option.
- ◆ Background Drive Verification - This option is used to automatically verify the media on all drives in arrays in the background. If a media error is detected, the controller can automatically re-write the data, providing that the array is in fault tolerant mode.

**NOTE:** System performance may be impacted by Background Drive Verification, depending on the workload.

- ◆ Enclosure Support - Selecting this option will cause the enclosure components to be monitored by StorView. If you deselect this option StorView will not report the enclosure status, will not report enclosure events, and the image on the Main screen will be dimmed. This does not disable the audible alarm on the front panel.
- ◆ Initialization Priority - This option determines the amount of processor time allocated to the operation. The higher the value, the more time the processor will spend on the initialization operation, reducing the time to complete the operation. It is recommended to balance the two priority parameters in the event a rebuild and initialization were to occur simultaneously.
- ◆ Rebuild Priority - This option determines the amount of processor time allocated to the Rebuild operation. The higher the value, the more time the processor will spend on the rebuild operation, reducing the time to complete the operation. It is recommended to balance the two priority parameters in the event a rebuild and initialization were to occur simultaneously.
- ◆ Network Idle Timeout - This option is only displayed for Embedded StorView, and sets the maximum amount of time that the embedded software will run without any HTTP activity. If the time period is exceeded the software will assume the HTTP interface is not responding and will reboot the Embedded module in an attempt to recover. This will not impact RAID operations.

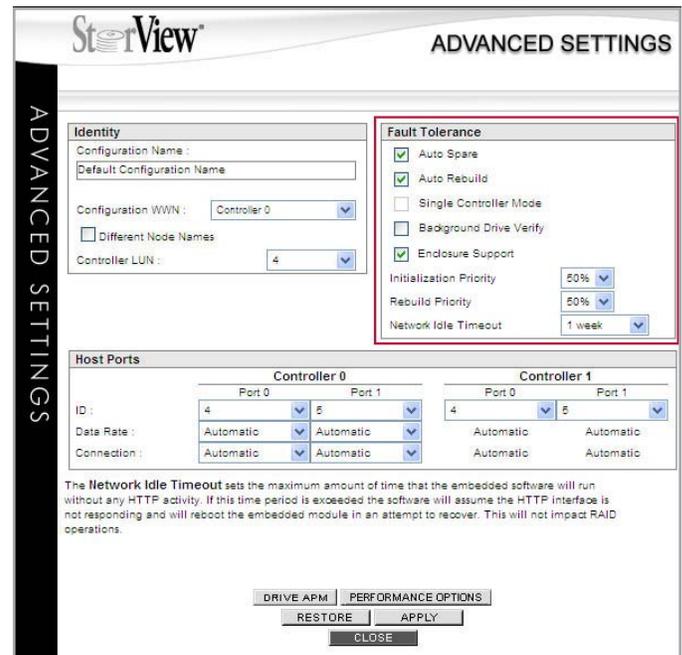


Fig. 66 - Advanced Settings Window

## 7.2.3. Host Ports

### 7.2.3.1. RAID Controller Systems

In the Host Ports section, you can change the ALPA ID assigned to each of the controller ports, set the type of connection, and set the data rate.

- ◆ Controller Port ID (P0) - This is the target ALPA for both controller(s) port 0. It can range from: Soft Address, or 0 - 125. The default is ID 4. This setting applies to Loop topology only.
- ◆ Controller Port ID (P1) - This is the target ALPA for both controller(s) port 1. It can range from: Soft Address, or 0 - 125. The default is ID 5. This setting applies to Loop topology only.

**NOTE:** The ports are identified in the software as P0 and P1, they refer to the label on the RAID Controller which indicates Host 0 and Host 1 connectors respectively. Controller 0 refers to the lower controller and Controller 1 is the upper controller.

- ◆ Controller Port Data Rate - Use the Automatic setting for most configurations. If you choose to use a specific setting (1 Gb, 2 Gb or 4 Gb) you override the automatic setting. Be sure the software setting here matches the hardware switch setting on the HBA ports.

**NOTE:** When using an Active-Active configuration (dual controllers) set the Controller Port Data Rate to one of the established settings. When the Automatic setting is used with Active-Active, it is possible for the speed to step down to 1Gb during a fail-back operation. Set the speed to either 2 Gb or 4 Gb to prevent this from happening.

- ◆ Connection - This option sets the type of connection that is being used from the host or switch. Use the Automatic setting for most environments where it will attempt to use Loop Only first then Point to Point. For custom settings, if you are connecting to a FL\_Port switch or NL\_Port HBA then select Loop Only, and if you are connecting to F\_Port switch or N\_Port HBA then select Point to Point.

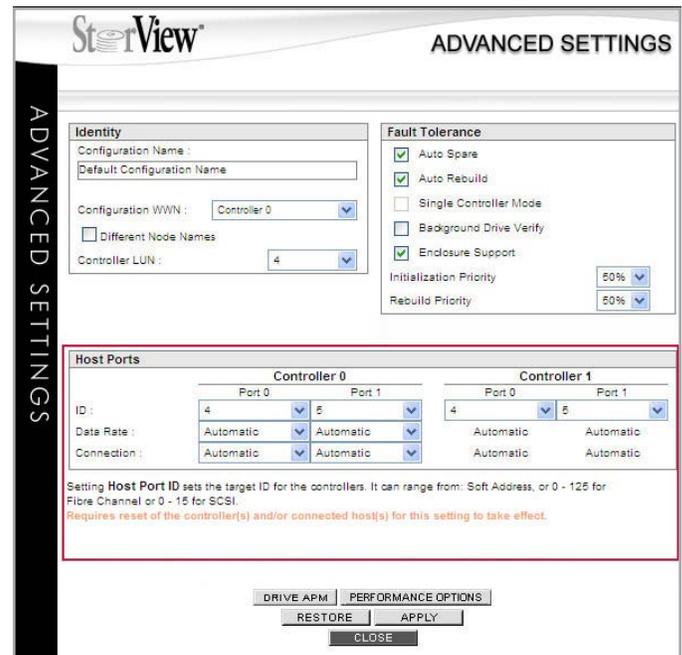


Fig. 67 - Advanced Settings Window

### 7.3. Advanced Power Management

Advanced Power Management brings to StorView and the storage systems two new levels of power management. These new power management features will help you extend the life of your disk drives, reduce power consumption, reduce noise levels, and excess heat, helping you save on energy costs. The first level is Drive Power Management and the second level is Array Power Management (refer to section 8.3.1. *Array Power Management (APM Level 2)* for details on Array APM.)

#### 7.3.1. Drive Power Management (APM Level 1)

Drive Power Management is accessed via the Advanced Settings button and the Drive APM button located at the lower position of the Advanced Settings window.

APM level 1 or Drive Power Management will manage these drive group types: hot spare drives, unused drives and failed drives. The user has the option to set policies for each drive group type managed by APM 1. For each drive group type, the policies are applied universally. Referring to Fig. 68 and Fig. 69 you can choose to enable or disable each drive group type and also enable periodic spin-up and test of the drives managed by APM. Below the drive group types section you will be able to customize your drives spinup and test policies.

**NOTE:** The power management feature will only spin down eight (8) unused or spare disk drives at a time. If there are more than eight unused or spare drives, APM will spin them down in groups of 8 or less. For example, if you have 10 unused drives, APM will spin down the first 8, then the remaining 2.

The time the page loaded, displayed across the bottom of the page, is designed to help you manage real time data for the opened window. It is a reminder to the user that you must manually refresh the page to ensure accurate data is presented. Be sure to reload the page when correct and current information is needed.

**IMPORTANT INFO:** You cannot enable APM features on drives that are also being controlled by the Snapshot feature.

#### To configure APM level 1:

1. From the Main screen, click the Advanced Settings button  on the tool bar .
2. From the Advanced Settings window click the DRIVE APM button located towards the bottom of the screen. This will open the Power Management window.

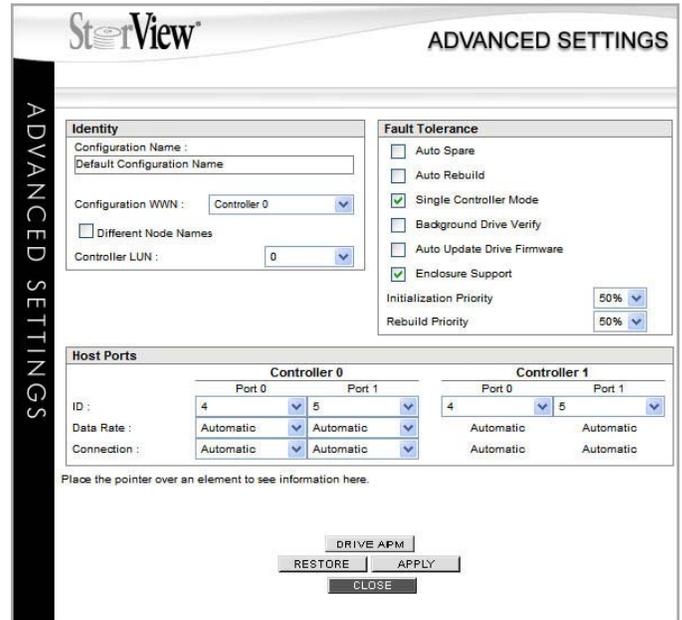


Fig. 68 - Advanced Settings – Drive APM Screen

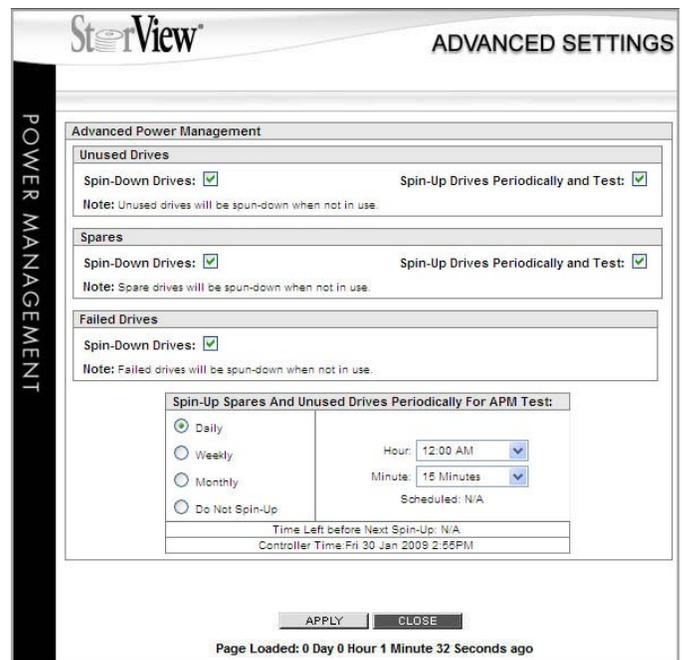


Fig. 69 - Advanced Settings – Drive APM Screen

- For each drive type (Unused Drives, Spares and Failed Drives) use the mouse pointer to select the check box adjacent to the type.

If you wish to have the drive type perform a periodic spin-up and self test, place a check mark in the box for each type, Unused and Spare drives only.

- Once you have enabled “Spin-Up Drives Periodically and Test,” at the lower half of the window choose the policy schedule you wish to use to manage your disks.

Use the mouse to select whether the drives should spin up and complete a test. You will choose whether the policy will be daily, weekly or monthly, and the time to perform the test.

- Click the APPLY button. Then click the CLOSE button on the confirmation window.
- Click the CLOSE button on the window.

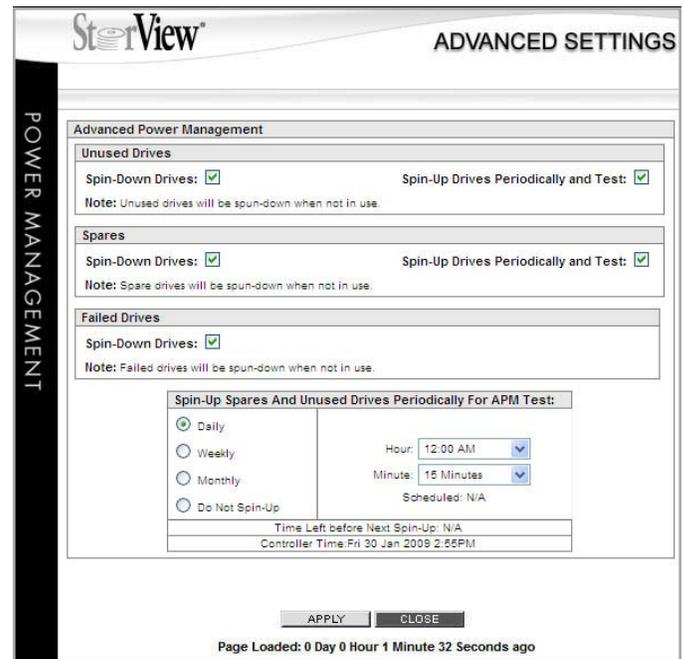


Fig. 70 - Power Management – Drive APM Screen

### 7.3.2. APM Characteristics

APM level 1 controlled drives will stay spun down after a configuration has been cleared. Since APM settings are written to meta-data and the user clears the configuration, then the APM level 1 settings will also be cleared. APM level 1 drives will not change their state as they are not updated.

The standard behavior is such that APM level 1 drives will stay spun-down if they are currently spun down even during configuration updates. They will be spun-up only if the user will require use of those drives, like creating an array.

If a drive is configured as a spare, it is known as a hot spare if it is spun-up and a cold spare if it is spundown.

Other characteristics are the drive icon changes to indicate a spun-down drive. If a drive is unused, a spare, or part of an array and has been spun down it will have a normal icon except with a blue tint overlay and a small “moon” icon to the left side of the icon. Refer to section 2.1.2. [Enclosure Section](#) for a list of all drive icons.

---

**IMPORTANT INFO:** When drives are spun down and depending upon the number of drives spun down, it could take longer than 1 and 1/2 minutes to spin the drives back up. If commands are issued to the drives before they are spun up, the commands will fail. You will need to reissue the commands once the drives are spun up and back online.

---

## 7.4. Advanced Performance Options

### 7.4.1. Performance Options

The following options will assist you in fine tuning the performance of your system. For specific details on the tuning process and applied metrics, refer to your supplier for subject related white papers.

#### 7.4.1.1. Synchronize Cache Writes to Disk

This option will determine whether there is normal processing of the SCSI Cache Synchronize command.

By placing a check mark in the box, the option is enabled. When enabled, a specific system command will cause the controller to flush the cache contents to the arrays.

When you disable this option, these system commands are disregarded and data is not immediately flushed to the arrays.

If you are operating a duplex system with backup battery units, disabling this option you may improve the system performance.

**CAUTION:** When disabled, a loss of power may result in data loss.

#### 7.4.1.2. Target Command Thread Balance

This option allows the controller to balance response times between multiple initiators.

When enabled (checked), the controller will attempt to reduce the maximum response time by servicing the oldest request first.

When disabled (unchecked), the controller will optimize throughput by sorting I/O requests to optimize disk access on any array with writeback cache enabled and write through disabled. This setting should provide optimized performance for typical workloads.

The default setting is Enabled (checked).

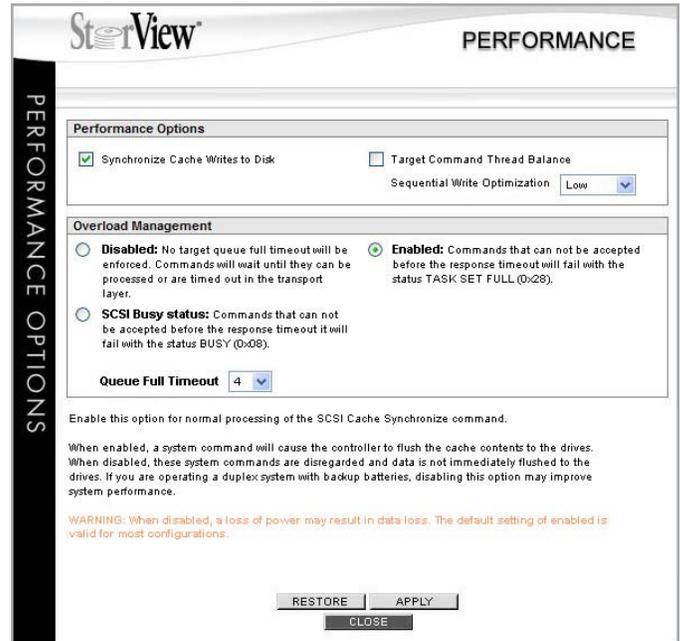


Fig. 71 - Performance Options – Synchronize Cache

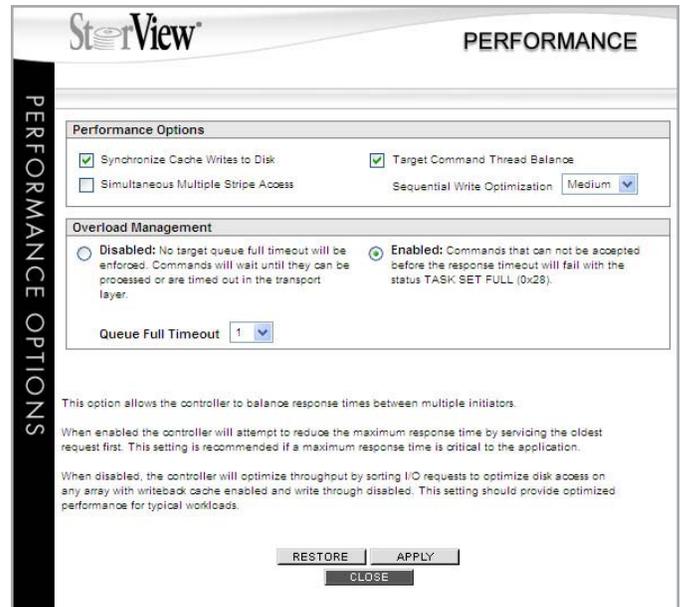


Fig. 72 - Performance Options – Target Command Thread Balance

### 7.4.1.3. Sequential Write Optimization

This option allows the user to set the amount of optimization the controller should perform for sequential write operations.

For best results, choose an optimization level that matches the amount of sequential write operations for your workload. If the workload is random or rare write operations, choose the DISABLED setting. Choose HIGH for mainly sequential writes, and MEDIUM or LOW for mixed work loads.

The default is LOW.

### 7.4.2. Overload Management

This option allows you to enable the response and set the response time out for SCSI Primary Commands.

Choose to enable this option when you want a response sent to the host within a specified amount of time because the target queue is full. The timeout periods are expressed in seconds.

The response will be either SCSI TASK SET FULL or SCSI BUSY. Click “Enable” for SCSI TASK SET FULL and “SCSI BUSY” for the SCSI BUSY response.

The Queue Full Timeout specifies the delay in seconds that commands will wait for space in the target queue before failing with the Target Queue Full response selected above. If a response of “Disabled” is selected, then the timeout value is not selectable.

Once enabled, you should select a timeout value shorter than the transport layer timeout of your host. Otherwise, the commands will be subject to the normal abort and recovery operations.

The default is enabled and a timeout value of 4 seconds.

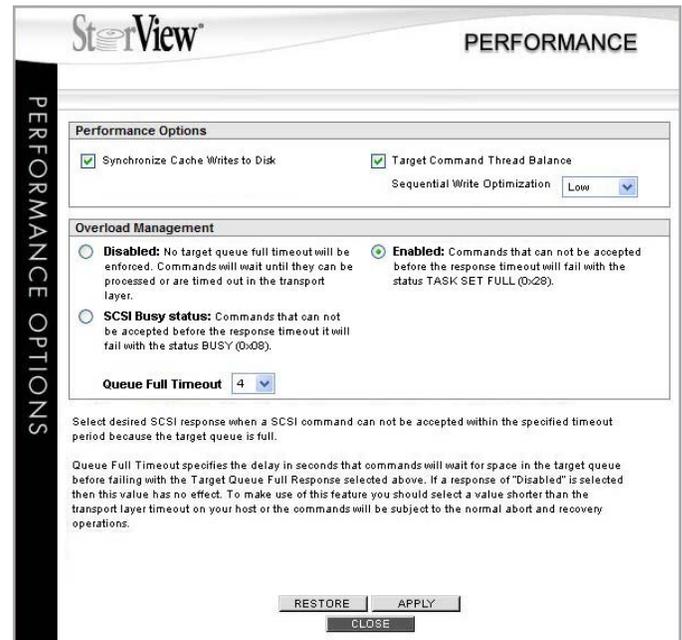


Fig. 73 - Performance Options – Sequential Write Optimization

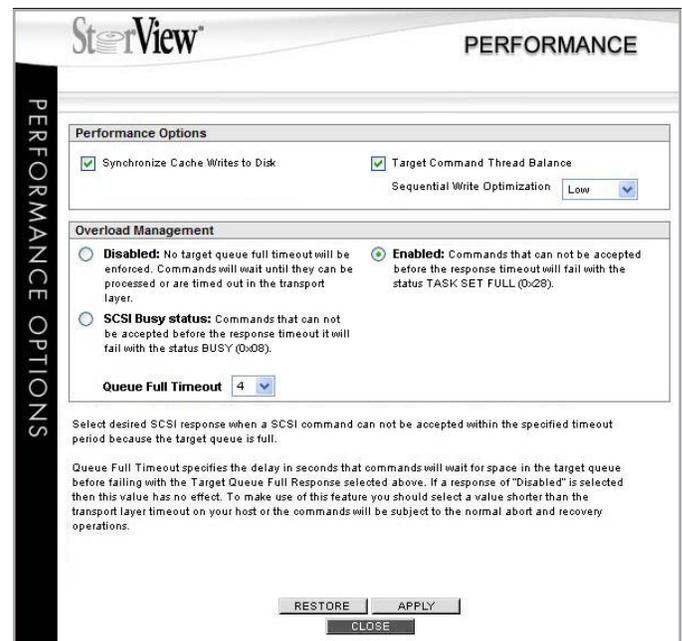


Fig. 74 - Overload Management

## 8. Managing the Storage Solution

### 8.1. Overview

In this chapter several areas are covered that pertain to managing your storage system. As a guide through this chapter the following are a list of the subjects included:

- ◆ Advanced Array Functions
    - Deleting an Array
    - Modifying an Array
    - Verifying an Array (Parity Check)
    - Identifying Drive Members
    - Rebuilding an Array
    - Expanding an Array
    - Trusting an Array
  - ◆ Advanced Power Management
    - Array Power Management (APM Level 2)
    - APM Characteristics
  - ◆ Restore and Clearing the Configuration
    - Restoring the Configuration
    - Clearing the Configuration
  - ◆ Advanced Drive Options
    - Accessing the Drive Information window
    - Locating a disk drive
    - Fail a disk drive
  - ◆ Advanced Logical Drives Functions
    - Viewing Unassigned Free Space
    - Expanding a Logical Drive
    - Deleting a Logical Drive
-

## 8.2. Advanced Array Functions

### 8.2.1. Deleting an Array

**CAUTION:** You must stop all host I/O operations prior to deleting an array.

**CAUTION:** Deleting an array will delete all data on the logical drives and those associated with that array. Be sure you have a backup of the data before proceeding.

**IMPORTANT INFO:** You cannot delete any array in the storage system while any of the following background operations are occurring: Expansion, Rebuild, Initialization, or Parity Check.

1. Stop all host I/O operations.
2. From the Main screen Configuration section, click the <Array\_Name> you wish to delete.
3. The Array Information screen will appear, click the DELETE ARRAY button.
4. A confirmation screen will appear, type your "password" and click the GO button. If the wrong password is entered, you will be prompted and the array will not be deleted.
5. Once the array has been successfully deleted, click the CLOSE button.

### 8.2.2. Modifying Arrays

Once the array has been created and is online you can make changes to the following:

- ◆ The name of the array.
- ◆ The Read-Ahead and Writeback cache parameters.

**NOTE:** You can restore the original settings by clicking the RESTORE button which will cancel any changes you have made as long as you have not clicked the APPLY button.

1. From the Main screen Configuration section, click the <Array\_Name> you wish to modify.
2. The Array Information window will open, type a new name for the array in the Name field and click the APPLY button.

**NOTE:** If the array was trusted or never initialized, you can initialize the array from this window by clicking the INITIALIZE button.

**NOTE:** If an array has been initialized the name of the button will change to "RE-INITIALIZE."

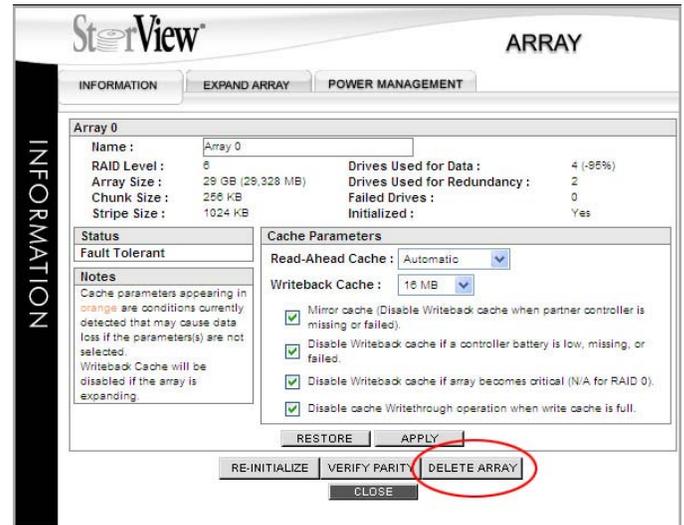


Fig. 75 - Array Information Window



Fig. 76 - Confirmation Window

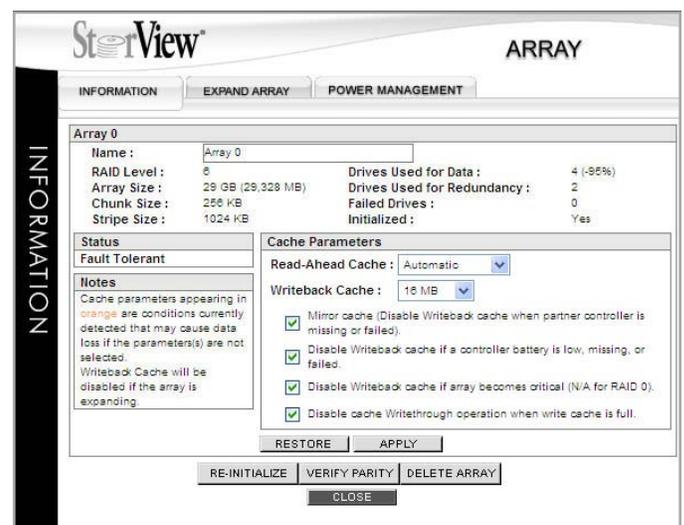


Fig. 77 - Array Information Window

### 8.2.3. Verify Parity

It is desirable to perform a parity check as a normal maintenance procedure to ensure data integrity. Also, if a RAID 5, 6, and 50 disk array experiences a situation where a controller is replaced after the controller is powered off with write operations in progress, it may be necessary to verify and correct the parity data on the array.

1. From the Main screen Configuration section, click the <Array\_ Name> you wish to verify parity data.
2. The Array Information window will open, click the VERIFY PARITY button.
3. Select a verify method from the drop-down list and click the VERIFY PARITY button.

The table below provides a description of each option.

Table 13 - Verify Parity Options

Option	Description
Check Parity	This option reads all the data and parity, calculates the XOR of the data, and compares it to the parity. If there is an error, it is displayed in the event log.
Rewrite Parity	This option reads all the data, calculates the XOR of the data, and writes this out as the new parity. This is the fastest to complete, since it does not have the overhead of a comparison.
Check and Rewrite Parity	This option reads all the data and parity, calculates the XOR of the data, and compares it to the parity. Then, if there is a discrepancy, it writes this out as the new parity and creates a log entry. This is the slowest to complete, since it has the overhead of a comparison as well as a rewrite.

During the verification, the drive members icons in the front enclosure view of that array will display an animated icon indicating a verification is in progress. Also, adjacent to the array name in the Main screen, a progress bar will indicate the percent complete. When you place the mouse pointer over the progress bar a pop-up will display the value of the percent complete.

You can stop the Verification process if you wish by clicking on the Stop link located to the right of the progress bar.

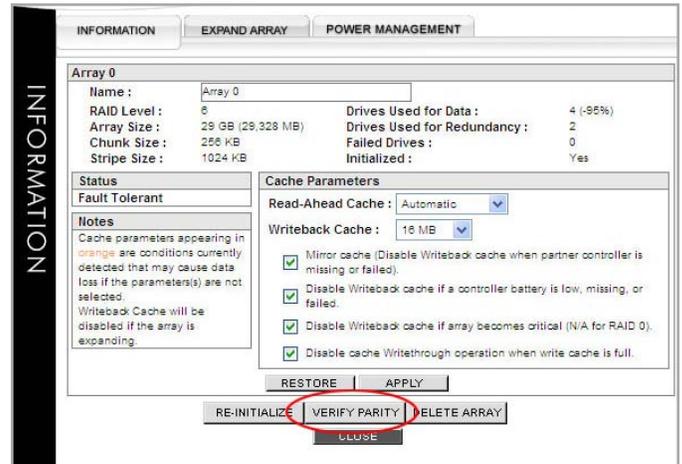


Fig. 78 - Array Information Window

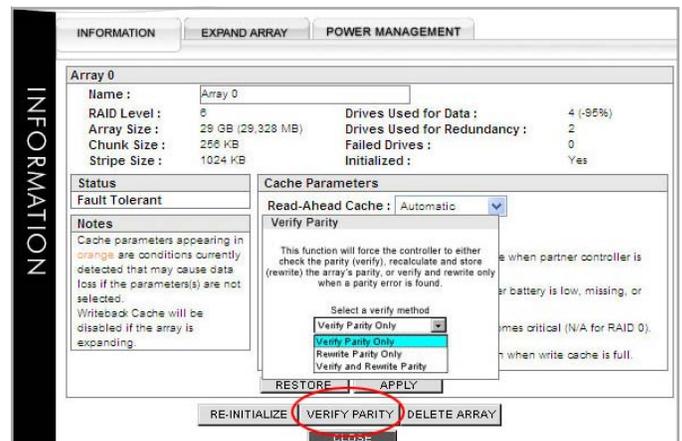


Fig. 79 - Verify Options Screen

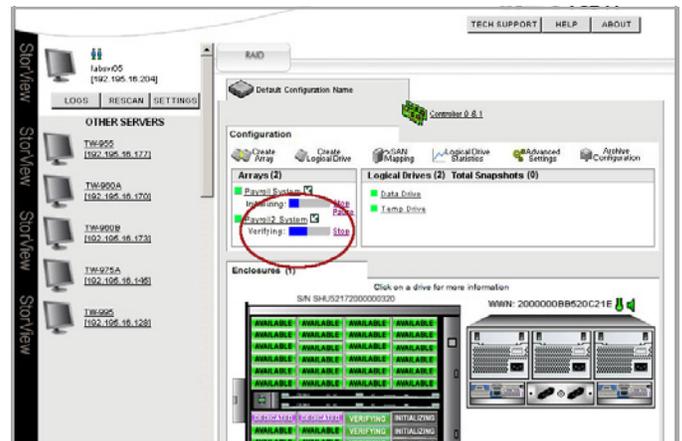


Fig. 80 - Monitoring Progress of Parity Verification

### 8.2.4. Identifying Drive Members

Should the need arise, you can quickly identify which drives in the enclosure are members of a specific array.

On the Main screen located on the right side of an Array name is an icon (Drive Identify icon), whose appearance is an arrow pointing to the lower left corner. This is used to turn on the identify function.

Clicking on the Drive Identity icon will cause all drive members of that array in the representation of the enclosure front view to have the “Drive Identity (arrow)” icon displayed on those drives. The icon also appears next to each logical drive created from the drives of the array.

You can also identify specific drives in an array by flashing its Drive Status LED, see 8.5.2. *Locate Drive*.

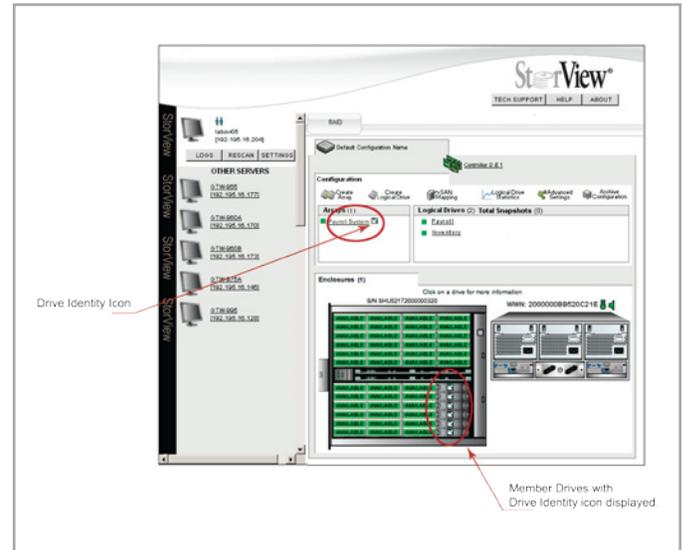


Fig. 81 - Identifying Member Drives Screen

### 8.2.5. Rebuilding an Array

This option is designed for situations where the user wants to manually start a rebuild operation.

**NOTE:** Only one rebuild can take place at a time. If a rebuild is currently underway, the second array scheduled for rebuild will be on-hold until the first array has completed the rebuild.

One scenario where this option would be useful is if you inadvertently pulled the wrong drive from a working array and that drive is now flagged as a failed drive, regardless of whether or not you re-insert the drive quickly. If you do not have a hot spare defined, the array will not automatically begin a rebuild operation. You must change the status of the flagged failed disk drive to a spare drive which will clear the condition and initiate a rebuild.

1. From the Main screen, identify the “failed” drive displayed in the enclosure front view and click that drive icon. The Drive Information window will open.
2. Click the REBUILD ARRAY button. A small window will appear.
3. Scroll down and choose the specific array that became critical from the removed/failed drive.
4. Click the REBUILD button.
5. A confirmation window will appear indicating the successful execution of the command, click the CLOSE button.
6. Click the CLOSE button on the Drive Information window.

**IMPORTANT INFO:** For a RAID 6 rebuild operations, if two (2) drives fail, and a rebuild is started after replacing both disk drives, only one drive rebuild is performed at a time. This means one drive will be rebuilt then the second drive rebuild will be started.

You can monitor the rebuild operation from the Main screen.

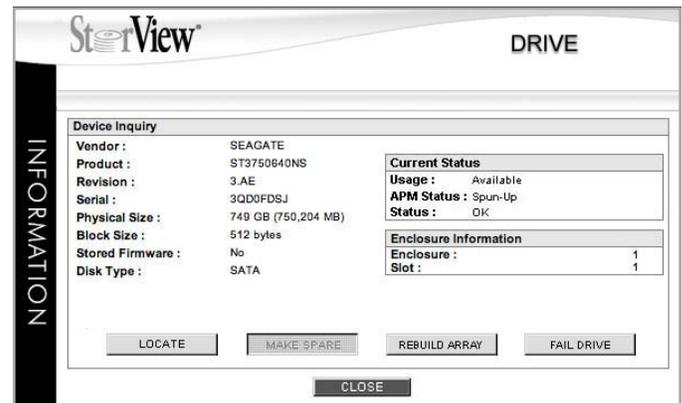


Fig. 82 - Drive Information Window

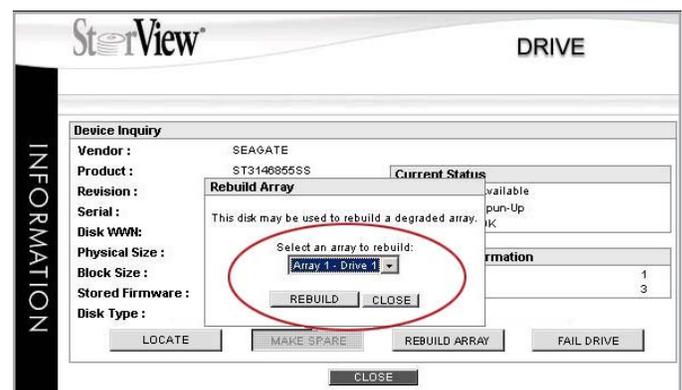


Fig. 83 - Drive Information Window

### 8.2.6. Expanding an Array

**Caution** You must stop all host I/O operations prior to expanding an array. After the expansion has reached 1%, then you can resume your host I/O operations.

The Expand Array feature is used to increase the capacity of an existing array. An array can be expanded to a maximum of 16 drives. Only one array can be expanded at a time.

---

**NOTE:** No configuration changes can be made to the arrays, logical drives, or SAN LUN Mapping while an expansion operation is in progress.

---

During the expansion process, data is re-striped across a new set of data drives, and new parity is calculated and written if necessary for fault tolerant arrays. If the array is a fault tolerant array, such as RAID level 1, 10, 5, 6, or 50, it will remain fault tolerant during the expansion.

Should a disk drive fail in a fault tolerant array during the expansion, the expand operation will continue as normal at which time it will flag the drive as failed and use the data and parity information to create the new data and parity stripe. After the expansion is complete, and if you had a hot spare designated, the automatic rebuild operation will commence bringing the non-fault tolerant expanded array back to a fault tolerant condition.

If a second drive failure occurs during expansion, that condition is not recoverable and you will have a total loss of data. You may wish to consider backing up the data prior to expanding an array. Although there is a level of protection during this operation without the backup, the best insurance is a valid backup.

---

**NOTE:** After the array expansion process has completed, if you are expanding for the purposes of new drive space you will need to create the appropriate logical drive(s) and define them in your operating system.

---

However, if the expansion is intended to increase the existing logical drive capacity you will need to perform a LUN Expansion. Afterwards a third-party volume/partition software product will be necessary to manipulate any existing partitions.

1. Stop all host I/O operations.
2. On the Main screen locate and click on the <Array\_Name> you wish to expand. This will open the Array Information window.
3. From the Array Information window, click the EXPAND ARRAY tab, see [Fig. 84](#).
4. Following the sequenced steps, click the Array Expansion Type

pull-down menu, choose the type of expansion applicable to your array.

5. Select the drives that will be used to expand the array.
6. Verify the changes you are about to make by examining the "Before Expansion" and "After Expansion" analysis.

[continued on the next page >>](#)

7. If your settings are correct, click the EXPAND button.
8. You will be prompted to confirm the Expand operation. Type your password and click the GO button.
9. You will receive a screen prompt that the command was successful, click the CLOSE button. If the command was unsuccessful, review the settings for incorrect parameters and hardware for operational status.

Once the expansion progress has reached 1% complete, you may resume normal host I/O operations.

### 8.2.7. Trust an Array

When you create an array, you have the option to trust the array. This option should only be used in environments where the user fully understands the consequences of the function. The trust array option is provided to allow immediate access to an array for testing application purposes only.

A trusted array does not calculate parity across all drives and therefore there is no known state on the drives. As data is received from the host parity is calculated as normal, but it occurs on a block basis.

There is no way to guarantee that parity has been calculated across the entire drive. The parity data will be inconsistent and so a drive failure within a trusted array will cause data loss.

1. On the Main screen in the Tool Bar, click the Create Array icon.
2. The Create Array window will appear, select your drives in the Available Drives section.
3. Enter a name for your array. You may use up to 32 characters (ASCII).
4. Select the RAID level for the array.
5. Enter the desired chunk size. Click the pull-down menu and choose from the available values.
6. At Item 7, use the pull-down menu and select "Trust Array."
7. Choose the "Back-off Percent" (reserved capacity) for the drives. The default is 1%.
8. Set the Read-Ahead Cache threshold.
9. Set the Writeback Cache options.

continued on the next page >>

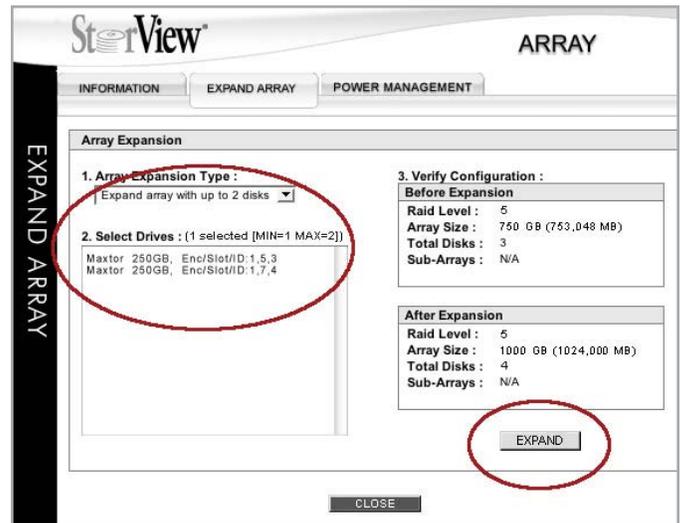


Fig. 84 - Expand Array Tab Selected

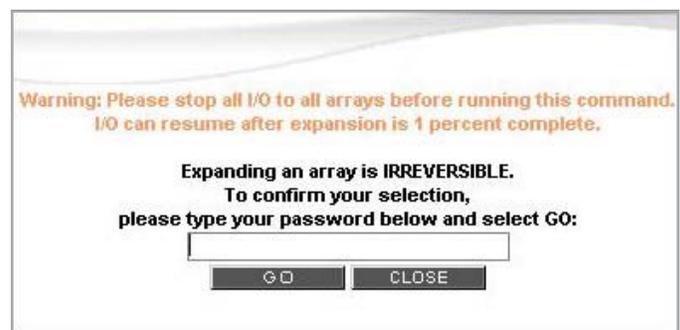


Fig. 85 - Expand Array Confirmation

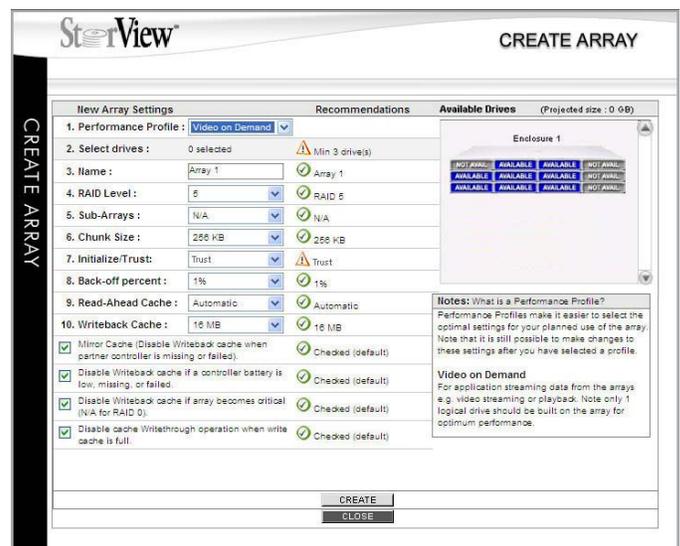


Fig. 86 - Create Array Screen

10. Click the CREATE button to create and trust the array.

**NOTE:** A trusted array is indicated on the Main screen with the following icon , which appears adjacent to the array name.

### 8.3. Advanced Power Management

Advanced Power Management brings to StorView and the storage systems two new levels of power management. These new power management features will help you extend the life of your disk drives, reduce power consumption, reduce noise levels, and excess heat helping you save on energy costs. The first level is Drive Power Management and the second is Array Power Management.

#### 8.3.1. Array Power Management (APM Level 2)

APM level 2 or Array Power Management, will individually manage the array's disk drives through a set of policies that establishes when the drives should be spun up and spun down. APM level 2 is implemented on a per array basis. Since APM level 2 deals with the arrays, each of the respective logical drives will be affected as well.

Important You cannot enable APM features on logical drives and their arrays that are also being controlled by the Snapshot feature. You must disable snapshot to utilize APM.

The information displayed across the bottom of the page is the time the page loaded, and is designed to help you manage real time data for an opened window. It is a reminder to the user that you must manually refresh the page to ensure accurate data is presented. Be sure to reload the page when correct and current information is needed.

#### To configure APM level 2:

1. From the Main screen, click the Array <name> you wish to set or manage APM policies.

[continued on the next page >>](#)

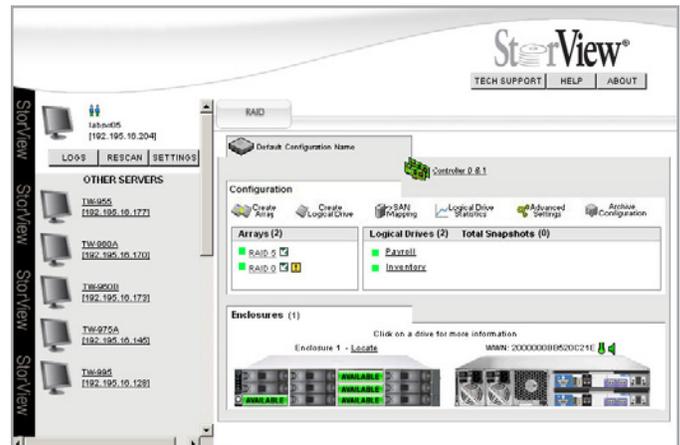


Fig. 87 - Main Screen - Trusted Array Example



Fig. 88 - APM-Snapshot Error Screen



Fig. 89 - Time Page Loaded Information

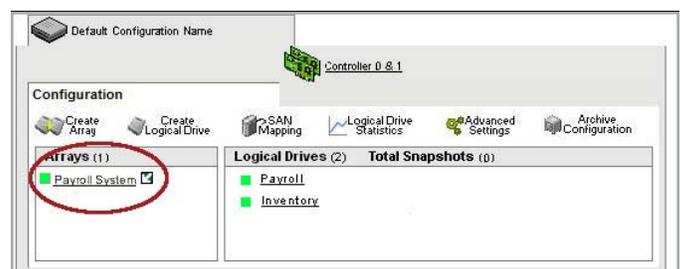


Fig. 90 - Time Page Loaded Information

2. From the Array Information window, click the POWER MANAGEMENT tab located at the top of the screen. This will display the Power Management settings for this array.
3. If not already enabled, click the check box next to "Advanced Power Management."
4. From the drop down menu's select the "Idle Time to Array Spin-Down" time period, as well as, the number of interval days and/or hours before the next APM test.
5. Click the APPLY button.  
To manually spin up or spin down an array, click the appropriate button located near the bottom of the page.
6. Click the APPLY button. Then click the CLOSE button on the confirmation window.
7. Click the CLOSE button on the window.

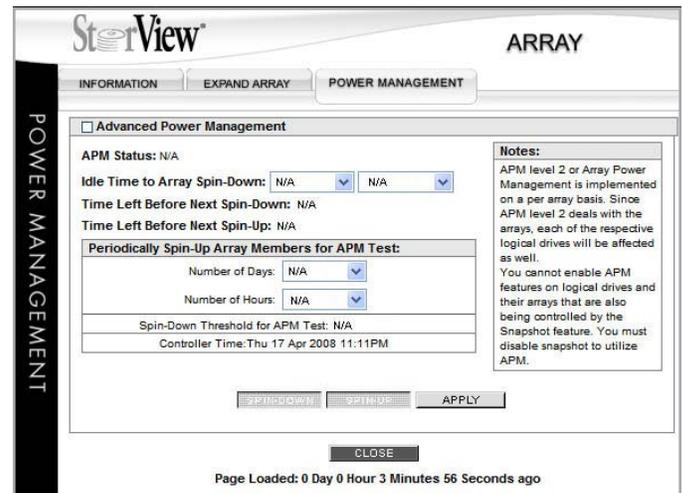


Fig. 91 - Array Power Management Screen - Unconfigured

### 8.3.2. APM Characteristics

APM level 1 controlled drives will stay spun down after a configuration has been cleared. Since APM settings are written to meta-data and the user clears the configuration, then the APM level 1 settings will also be cleared. APM level 1 drives will not change their state as they are not updated.

The standard behavior is such that APM level 1 drives will stay spun down if they are currently spun down even during configuration updates. They will be spun up only if the user will require use of those drives, like creating an array.

If a drive is configured as a spare, it is known as a hot spare if it is spun up and a cold spare if it is spun down.

If an APM level 2 policy is set on an array, that array will be spun down immediately if it has already been idle for that specified amount of time. APM uses controller statistics to determine the length of time the array has been idle and those statistics are tracked from the controller boot time. Even though APM level 2 is not enabled on a particular array, its status is still tracked.

Other characteristics are the drive icon changes to indicate a spun down drive. If a drive is unused, a spare, or part of an array and has been spun down it will have a normal icon except with a blue tint overlay and a small "moon" icon to the left side of the icon. Refer to section [2.1.2. Enclosure Section](#) for a list of all drive icons.

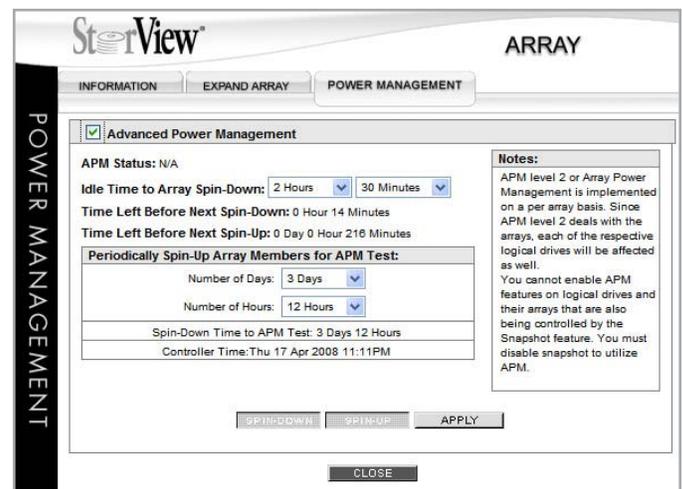


Fig. 92 - Array Power Management Screen - Configured

**IMPORTANT INFO:** When drives are spun down and depending upon the number of drives spun down, it could take longer than 1 and 1/2 minutes to spin the drives back up. If commands are issued to the drives before they are spun up, the commands will fail. You will need to reissue the commands once the drives are spun up and back online.

## 8.4. Restoring and Clearing the Configuration

**CAUTION:** If your configuration is cleared accidentally and you cannot restore the configuration exactly as it was, you will not be able to restore access to the data. This will result in data loss.

It is recommended to periodically save the configuration so that it is up to date, see 5.6. *Saving the Configuration*.

### 8.4.1. Restoring the Configuration

1. From the Main screen Tool Bar click the Archive Configuration icon. The Configuration Archival Operations screen appears.
2. Click the RESTORE button. You may click the CLOSE button to cancel and return to the Main screen.
3. The File upload window appears, click the Browse button.

You are presented with the browser's "Choose File" window. Select the appropriate file and click the Open button to continue, or Cancel to quit.

4. Click the UPLOAD button to continue to restore the configuration, or click the CANCEL button to quit.
5. After you have completed the configuration restoration, and if you had any RAID 5, 6, or 50 arrays defined, click the Array link on the Main screen for each RAID 5, 6, or 50 array. Perform a VERIFY PARITY operation before using those arrays. This will ensure that the data and parity data are correct.

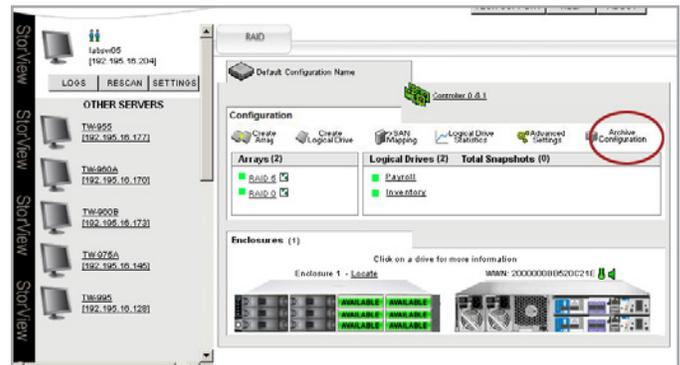


Fig. 93 - Main Screen



Fig. 94 - Configuration Archival Operations Window

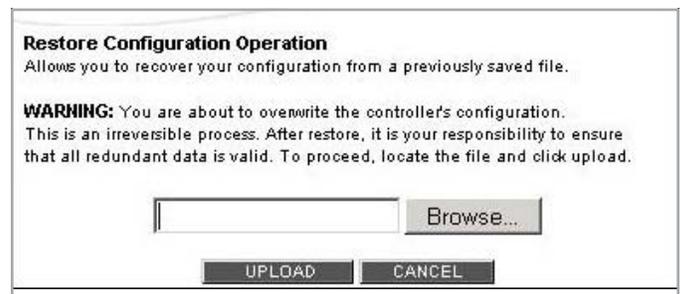


Fig. 95 - Restore Configuration Upload Window

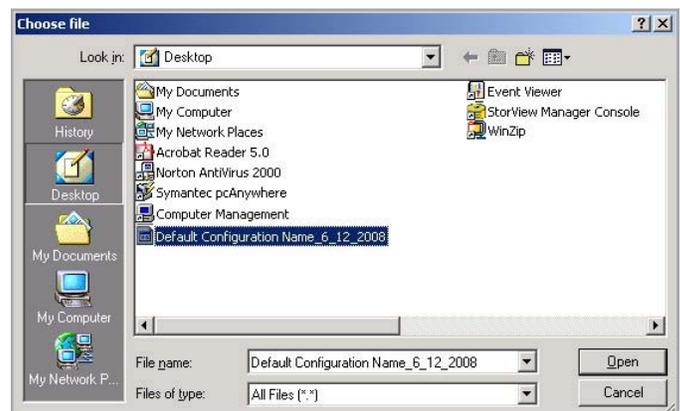


Fig. 96 - Restore Choose File Screen

### 8.4.2. Clearing the Configuration

Some conditions or situations may call for you to clear the entire configuration.

**CAUTION:** This process removes all arrays, logical drives, SAN LUN Mappings, Snapshots, Snapshot Feature License, etc. If there is any data on the drives, access to that data will be lost when the configuration is cleared.

1. From the Main screen Tool Bar click the Archive Configuration icon.

The Configuration Archival Operations window appears.

**IMPORTANT INFO:** Before you clear the configuration file, if you are using the Snapshot feature, ensure that you make a note of which WWN the configuration is based on, see [7.2.1. Identity](#). When you start your new configuration and if you intend to use the Snapshot feature, it will be necessary to set the Configuration WWN based to the same WWN as when the snapshot license key was issued. This is because the snapshot license key is generated so that it works only with an array using the configuration WWN as it was in the previous configuration.

2. Click the CLEAR button.

You may click the CLOSE button to cancel and return to the Main screen.

3. A pop-up window appears, type your password and click the GO button.

You will receive a confirmation of the operation. Click the CLOSE button.

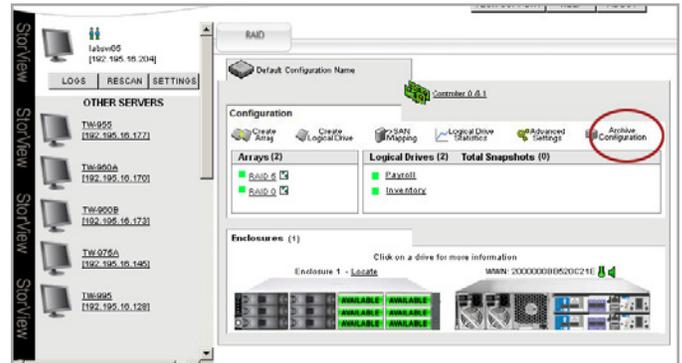


Fig. 97 - Main Screen



Fig. 98 - Configuration Archival Operations Window

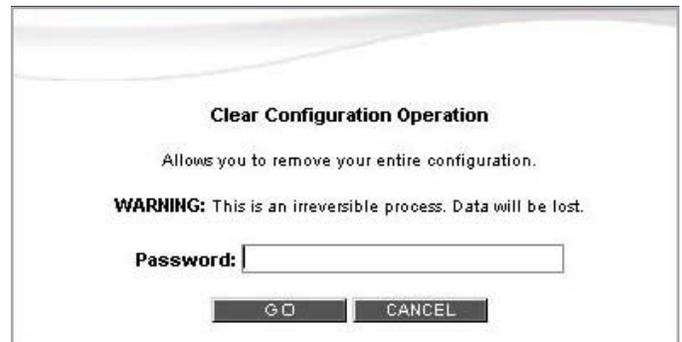


Fig. 99 - Restore Configuration Upload Window

## 8.5. Advanced Drive Options

The Drive Information window provides the user with the ability to view specific drive inquiry information and make changes to drive parameter settings. From the Drive Information window you will also find functional controls that allow you to locate a drive and execute a rebuild operation.

### 8.5.1. Accessing the Drive Information Window

1. From the Main screen, click on any disk drive icon displayed in the enclosure front view.

The Drive Information window will open.

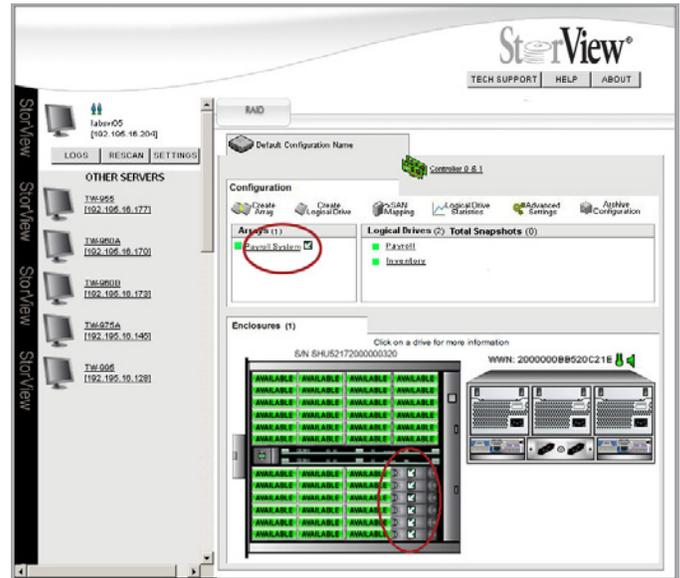


Fig. 100 - Main Screen

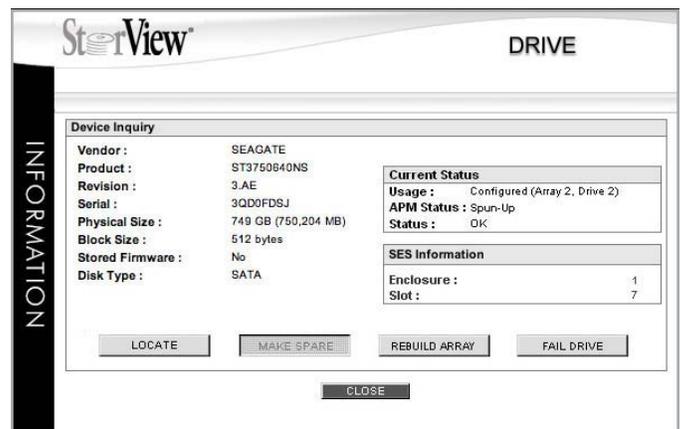


Fig. 101 - Drive Information Window

### 8.5.2. Locate Drive

This feature allows you to locate a drive installed in the enclosure using the Drive Activity LED.

1. To locate a disk drive, select a drive icon in the enclosure front view and click that drive icon.

The Drive Information window will open.

2. Click the LOCATE button.
3. A sub menu will open in the Drive Information window, from which you will select the time interval to blink the Drive's Activity LED. Select the time period you desire.
4. Identify the drive in the enclosure by its blinking Drive Activity LED. Refer to the hardware user's guide for details on Drive LEDs.

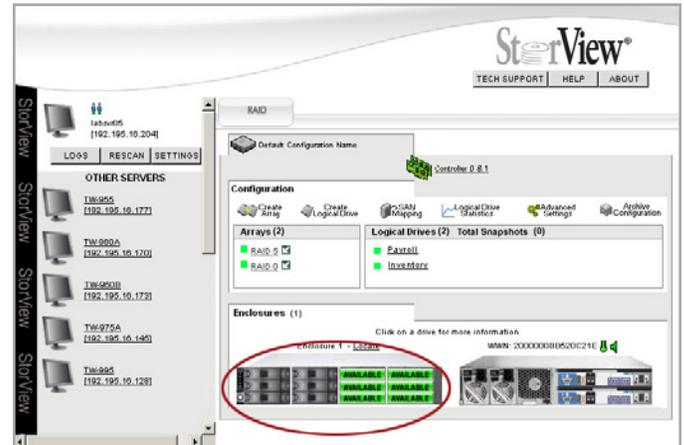


Fig. 102 - Main Screen

### 8.5.3. Fail Drive

This feature allows you to fail a drive that is a member of a redundant array. This provides the ability to safely fail a drive that is suspect of an impending failure or any other reason such as a drive inspection.

1. To fail a member disk drive, select a array member drive icon in the enclosure front view and click that drive icon.

The Drive Information window will open.

2. Click the FAIL DRIVE button.

The drive is then taken offline and marked as a failed drive. For RAID level 1/10 and 5, one drive can be failed, and for RAID level 6, two drives can be failed. Drives cannot be failed as members of a RAID 0 array.

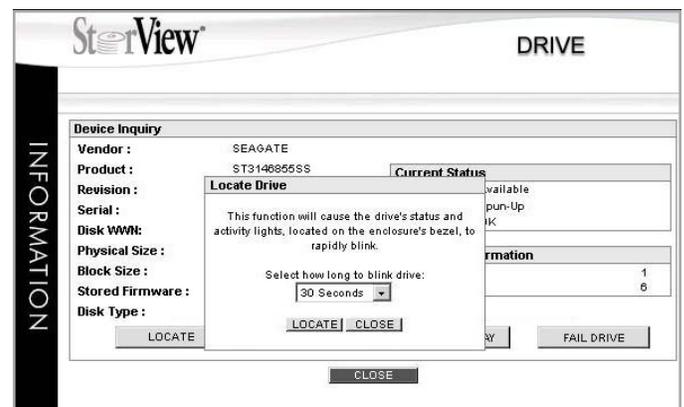


Fig. 103 - Drive Information Window - Locate Drive

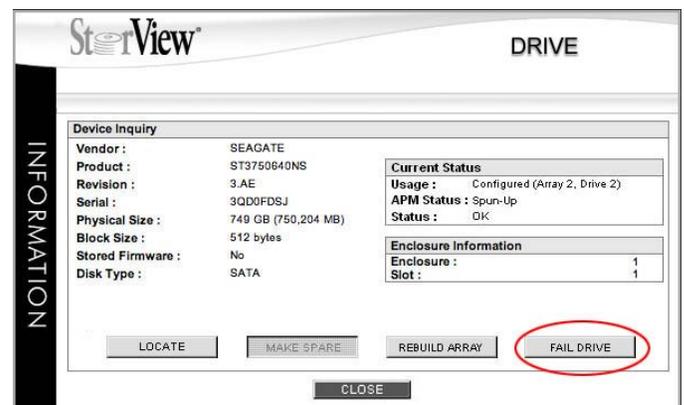


Fig. 104 - Drive Information Window - Fail Drive

## 8.6. Advanced Logical Drive Functions

### 8.6.1. Viewing Unassigned Free Space

Prior to creating or expanding a logical drive, you may wish to examine the unassigned free space. This will help you identify the available free space that can be used to create and expand logical drives.

The Create Logical Drive window is designed to display all available unused space or free regions.

1. From the Main screen in the Tool Bar click on the Create Logical Drive icon.

The available free space is displayed in the “Select which Array(s) to Use” scrollable window. You will see each array and how much space is available.

2. If you were just interested in the available free space, click the CLOSE button. Otherwise to continue with creating a logical drive, see [5.5. Create the Logical Drive](#).

### 8.6.2. Expanding a Logical Drive

**CAUTION:** You must stop all host I/O operations prior to expanding a logical drive.

Expanding a logical drive is a utility that allows you to take an existing logical drive and expand its capacity using free regions.

**NOTE:** After the expansion process has completed you will need to use a third-party volume/partition software product to manipulate any existing partitions.

1. Stop all host I/O operations.
2. From the Main screen in the Logical Drives section, click on a <logical\_drive\_name> that you wish to expand. The Logical Drive Information window will open.

[continued on the next page >>](#)

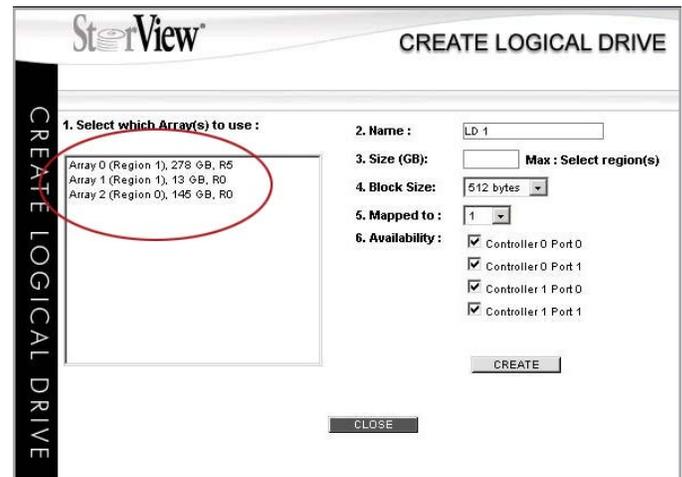


Fig. 105 - Create Logical Drive Window

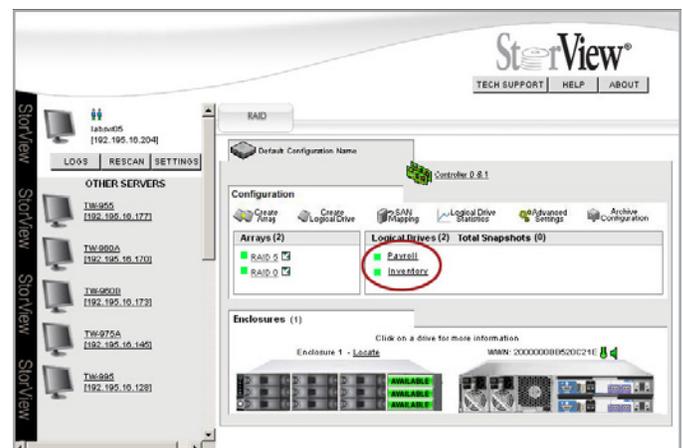


Fig. 106 - Main Screen

3. Locate the Expand section of the window (lower half), and follow the sequenced steps beginning at "Step 1" where you will choose a free space region to be used for the expansion.
4. In the "Add Capacity" box, enter the amount of the selected region to expand the logical drive. You may use the entire free region space or a portion of it, specify by entering the number of GBs to use.
5. Click the EXPAND button.
6. You will be prompted to enter your password to confirm the expansion. Type in your password and click the GO button.
7. You will receive a screen prompt that the command was successful, click the CLOSE button. If the command was unsuccessful, review the settings for incorrect parameters and hardware for operational status.

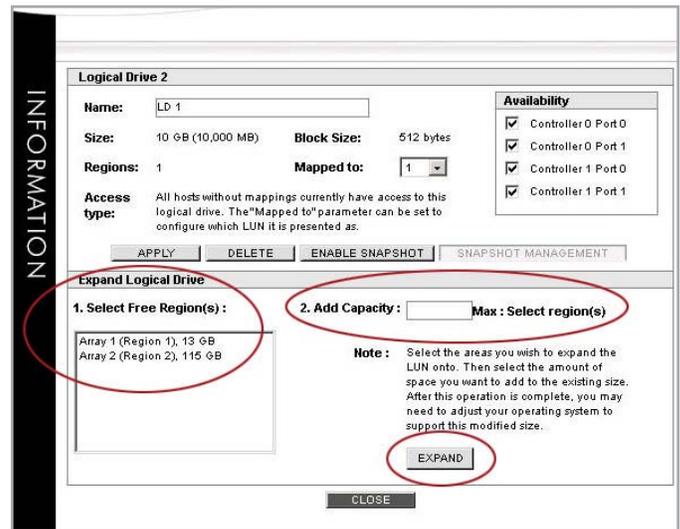


Fig. 107 - Logical Drive Information Screen (Not SAN LUN Mapped)

### 8.6.3. Deleting a Logical Drive

**CAUTION:** You must stop all host I/O operations prior to deleting a logical drive.

Deleting a logical drive is an option that allows the user to remove an existing logical drive that is no longer needed or desired. If the logical drive was previously used be sure to make a backup of any data on the logical drive. After deleting the logical drive, SAN LUN Mapping (if used) and the operating system will need to be modified due to the missing drive.

1. Stop all host I/O operations.
2. From the Main screen in the Logical Drives section, click on a <logical\_drive\_name> that you wish to delete.

The Logical Drive Information window will open.

3. In the Logical Drive section at the top, click the DELETE button.
4. You will be prompted to enter your password to confirm the deletion. Type in your password and click the GO button. Click the CANCEL button to exit without making any changes.
5. You will receive a screen prompt that the command was successful, click the CLOSE button. If the command was unsuccessful, review the settings for incorrect parameters and hardware for operational status.

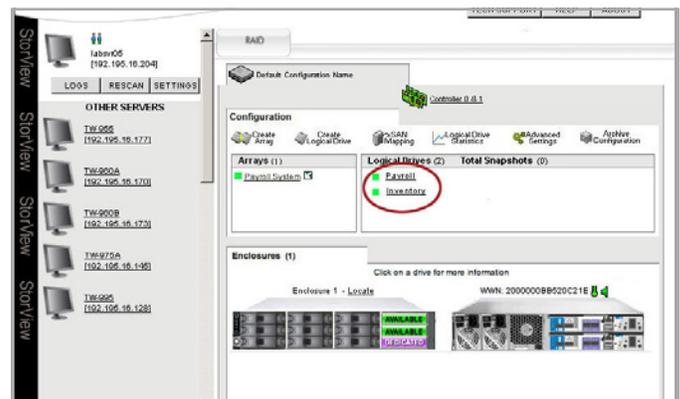


Fig. 108 - Main Screen

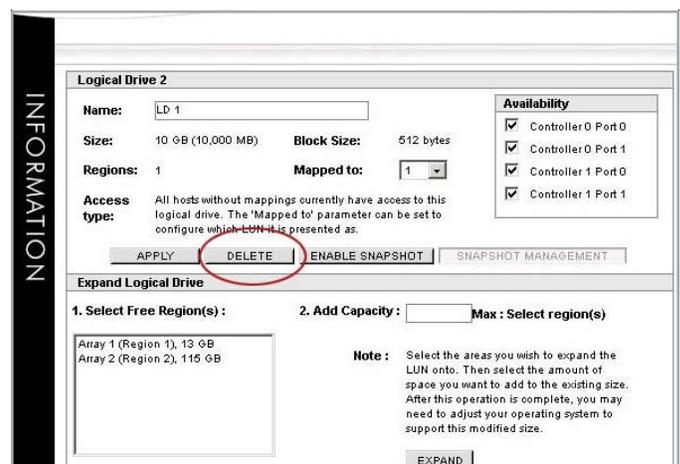


Fig. 109 - Main Screen

## 9. SAN LUN Mapping

### 9.1. Overview

SAN LUN Mapping is a feature licensed product. You must have a feature license validated for either the 15 day grace period, or an unlimited license. Refer to section [1.3.2. Feature License](#).

When attaching more than one host system to a storage system, it may be necessary to more precisely control which hosts have access to which logical drives. In addition to controlling availability on a controller port by port basis, it is also possible to further restrict access to a specific host system or single adapter in a host system by the use of SAN LUN Mapping. Up to 512 SAN LUN Mappings are supported.

### 9.2. Terminology

The following table describes the terminology relating to StorView's SAN LUN Mapping.

Table 14- SAN LUN Terminology

Term	Description
HBA Port Name (Port Name)	This is an eight byte hexadecimal number, uniquely identifying a single host HBA port. It incorporates the World Wide Name and two additional bytes which are used to specify the format and indicate the port number.
Mapping Name	A 32 character name that can be used to identify the host system.
Read/Write Access	A host may read and write to the logical drive.
Read Only Access	A host may only read from a logical drive.
*Used in another Mapping	This notation marks a logical drive that has been mapped to another Host HBA Port, but is available to be mapped to the selected Host HBA Port. You will receive a single warning about this during each session. And file sharing or path failover software is required to prevent data corruption.

### 9.3. Accessing SAN LUN Mapping

Clicking the “SAN Mapping” icon in the Tool Bar on the Main screen will open the SAN LUN Mapping window. Here you will find a list of the specific host HBA and their ports. You can view, name, create and remove mappings from this window. If no mappings are present, you may create a new mappings using the “ADD MAPPINGS” section. If a mapping exists, selecting a HBA will display the current mapping(s) and its parameters.

A list of current mappings for each HBA port is displayed with the specific port is selected. A graphical image depicts the basic HBA to Controller port connection and mapping details.

### 9.4. Overview: SAN LUN Mapping Window

Fig. 111 provides an explanation of the components of the SAN LUN Mapping window.

A graphical illustration of the physical connection from the Host HBA Port to the enclosure’s controller port is provided to help you visualize the topology being mapped.

The SAN LUN Mapping window is basically divided into sections. Each section allows you to perform a task in the process or creating a SAN LUN Mapping.

#### 9.4.1. SELECT YOUR HBA and NAME YOUR HBA PORT Section

In this section you will see the discovered HBAs and their ports in the “SELECT YOUR HBA” section. You must name the HBA port to be able to proceed, so in the box labeled “NAME YOUR HBA PORT:” enter the name you wish to use to identify the port. For example: Qlogic A Port 0, or Qlogic B Port 1. The graphic displays the WWN of the controller and a depiction of the cabled HBA to Controller Port.

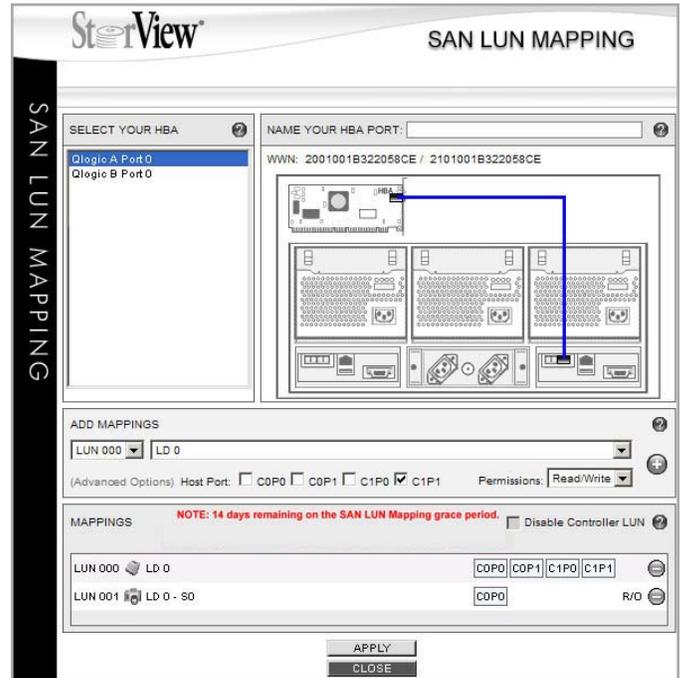


Fig. 110 - Logical Drive Information Screen (Not SAN LUN Mapped)

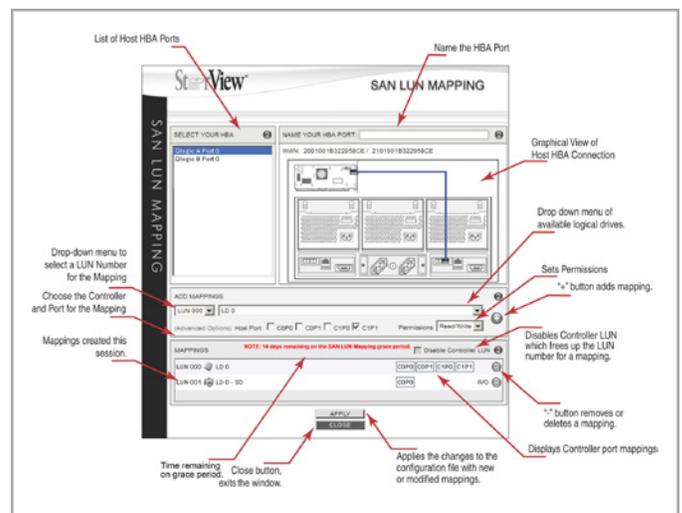


Fig. 111 - SAN LUN Mapping Screen Described

### 9.4.2. ADD MAPPINGS Section

In this section you select the LUN number to use for the mapping and choose a logical drive that has already been created.

Also you will select the controller and ports from which to make the mapping available on and set the permissions of the mapping. Click the button with the “plus” symbol will add this mapping to the MAPPINGS section.

You will have the following choices for your controller/port: C0P0, C0P1, C1P0 and C1P1. Where C0 represents the left or lower controller and C1 represents the right or upper controller, and P0 indicates the port marked Host 0, FC Host 0, or “IN” 0, and P1 represents the port marked Host 1, FC Host 1, “IN” 1 on each controller. Choose the ports on which you want the logical drive to be seen. Ensure you have cables connected to those controller ports.

---

**NOTE:** All the different port markings are provided since this document supports several controller model types.

---

### 9.4.3. MAPPINGS Section

In this section you will see the list of mappings you have created during this session. You can also disable the Controller LUN if a conflict occurs. The mapping displays the LUN number assigned, the logical drive name/number, which controller and ports it is available on and which the permissions are for the mapping.

A button with a “minus” sign is provided to remove a mapping.

---

**NOTE:** If you wish to make changes to an existing mapping you must first delete the mapping, then re-create the mapping with the new parameter settings.

---

Also in this section you will see the number of days remaining on the 15 day grace period license. With the unlimited full license you will not see a message. You will also receive a Controller related event every 24 hours reporting the number of days remaining on any of the temporary licenses.

---

## 9.5. Creating a SAN LUN Mapping

The following are the steps to create a SAN LUN Mapping. It is assumed a full license has been validated.

**IMPORTANT INFO:** If you intend to use the snapshot feature and choose to have access to this logical drive from your operating system or VSS, you must SAN LUN Map the logical drive. You may do this before or after the snapshots have been taken.

1. From the Main screen click the SAN Mapping icon in the Tool Bar. The SAN LUN MAPPING window will appear.
2. Select the Host HBA port.  
  
In the "SELECT YOUR HBA" section, choose an unnamed port and identify it using the displayed HBA Node WWN and Port WWN.
3. In the "NAME YOUR HBA PORT" section, enter a user defined name for the HBA Port. The name will appear in the "SELECT YOUR HBA" section, as you type.

You may use up to 32 ASCII characters, however only 26 characters are displayed in the "SELECT YOUR HBA" section. The minimum number of characters is five.

4. Repeat for any other HBA port.

[continued on the next page >>](#)

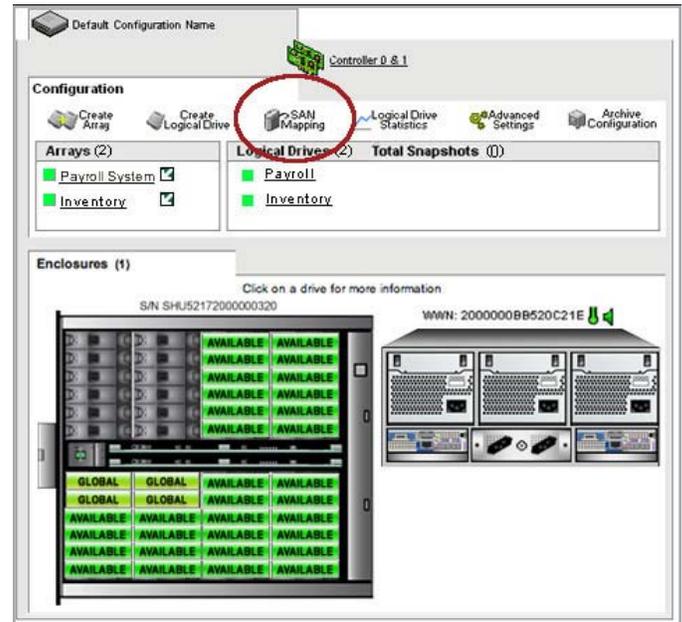


Fig. 112 - Main Screen – Enclosure/Array/Toolbar Section - SAN LUN Mapping Window

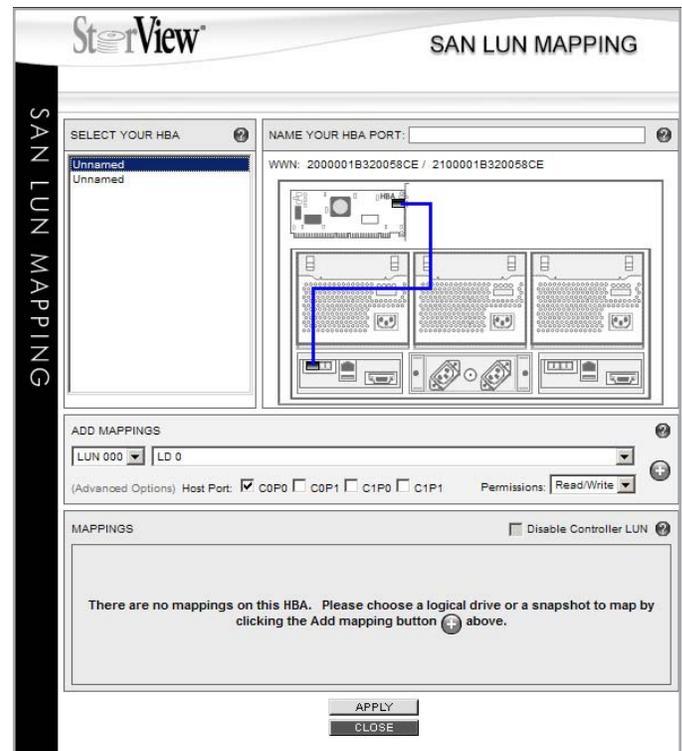


Fig. 113 - Main Screen – Enclosure/Array/Toolbar Section - SAN LUN Mapping Window

5. Add a mapping.

In the “ADD MAPPINGS” section, do the following:

- a. Choose the LUN (logical unit number) to present the mapped logical drive to the Host system. Click the drop-down menu and choose the desired number.
- b. Select the logical drive you wish to map. Click the pull down window and select from the list of logical drives displayed. Default logical drive names are LD0, LD1, LD2, LD3, etc.

**NOTE:** Logical drives marked with an asterisk “\*” indicate that another mapping for another Host HBA has been established for this logical drive. You can map it again to additional HBA’s but know that all Host HBAs mapped to this logical drive will see and have access to this logical drive. You will receive one warning about this process during a SAN LUN Mapping session. Also you must use file sharing or path failover software to ensure data integrity.

- c. Select an access permission for the mapping. Choose from the drop-down menu: Read/Write or Read Only.

**NOTE:** Microsoft Windows does not support Read Only permissions.

- d. Select the Host Port/Channel.

Click the check box(es) for the port that you wish the mapping to be available: C0P0, C0P1, C1P0, and/or C1P1. Where C0 represents the left or lower controller and C1 represents the right or upper controller, and P0 indicates the port marked FC Host 0, and P1 represents the port marked FC Host 1 on each controller. Ensure that the ports you select are properly cabled.

- e. Click the button with the “+” symbol on it. The mapping is added to the MAPPINGS section.

6. Review your mappings and settings, then click the APPLY button.

7. You will receive a confirmation, click the OK button to continue, or CANCEL to exit and return to the SAN LUN Mapping window.

8. You may continue to create more mappings by repeating the above procedures or end this session by clicking the CLOSE button.



Fig. 114 - Shared Mapping Warning

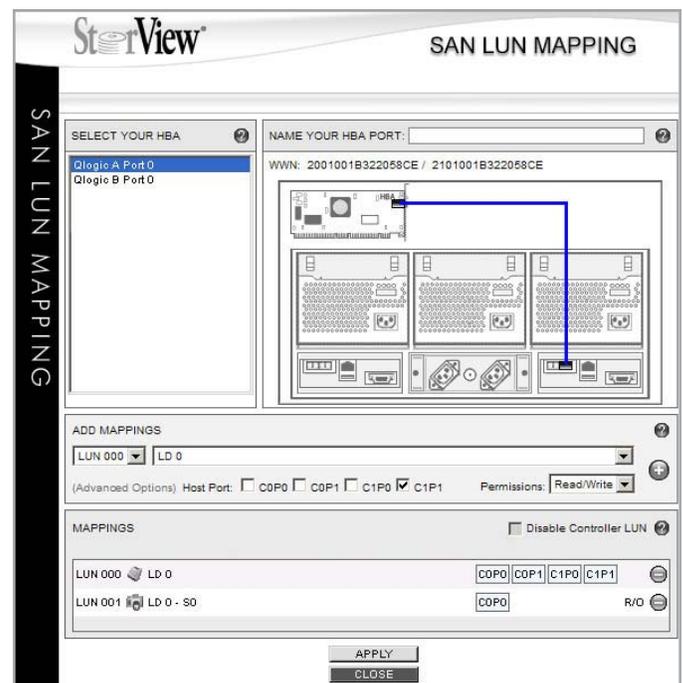


Fig. 115 - SAN LUN Mapping Window – Completed

## 9.6. Deleting a SAN LUN Mapping

**NOTE:** Prior to deleting SAN LUN Mappings, clear your web browser's cache (with Internet Explorer the function is known as deleting Temporary Internet Files), then proceed with removing the mappings.

1. From the Main screen click the SAN Mapping icon in the Tool Bar.
2. Select a Host HBA port in the "SELECT YOUR HBA" section that contains the mapping to be removed.
3. Select a Mapping to be removed from the "MAPPINGS" section.
4. Click the button with the "minus" symbol.
5. Click the APPLY button.
6. You will receive a confirmation, click the OK button to continue, or CANCEL to exit and return to the SAN LUN Mapping window.
7. You may continue to remove more mappings by repeating steps 3 and 4 above, or end this session by clicking the CLOSE button.

**NOTE:** After deleting the SAN LUN Mappings, the port mappings for that logical drive are disabled. Be sure to access each logical drive affected and configure the port mappings as desired. Refer to Create a Logical Drive for information on mapping the logical drive.

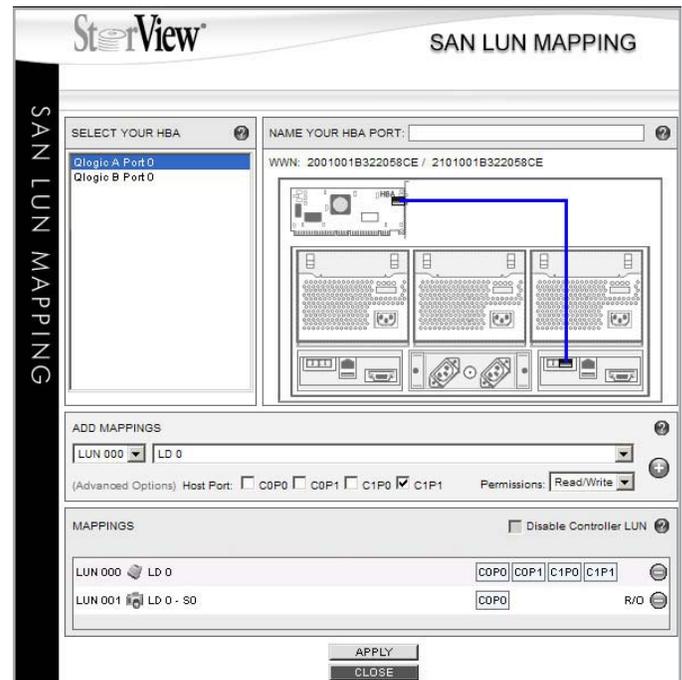


Fig. 116 - SAN LUN Mapping Screen

## 9.7. Modifying a SAN LUN Mapping

In order to make changes to an existing SAN LUN Mapping, you must first remove the existing mapping and then re-create the mapping with the changes.

**CAUTION:** Making changes to these mapping parameters may have an adverse affect on other mappings or to the operating system accessing the logical drive.

1. From the Main screen click the SAN Mapping icon in the Tool Bar.
2. Select a Host HBA port in the “SELECT YOUR HBA” section that contains the mapping to be removed.
3. Select a Mapping to be modified from the “MAPPINGS” section.

**NOTE:** Make a note of the settings for this mapping to use when you re-create the mapping.

4. Click the button with the “minus” symbol.
5. Click the APPLY button.
6. You will receive a confirmation, click the OK button to continue, or Cancel to exit and return to the SAN LUN Mapping window.

### 9.7.1. Re-Create the Mapping

1. From the Main screen click the SAN Mapping icon in the Tool Bar. The SAN LUN MAPPING window will appear.
2. Select the Host HBA port or Host HBA Initiator ID.

In the “SELECT YOUR HBA” section, choose an unnamed port and identify it using the displayed HBA Node WWN and Port WWN. This data is located above the graphical image of the storage system.

3. In the “NAME YOUR HBA PORT” section, enter a user defined name for the HBA Port. The name will appear in the “SELECT YOUR HBA” section, as you type.

You may use up to 32 ASCII characters, however only 26 characters are displayed in the “SELECT YOUR HBA” section. The minimum number of characters is five.

4. Repeat for any other HBA port.
5. Add a mapping.

In the “ADD MAPPINGS” section, do the following:

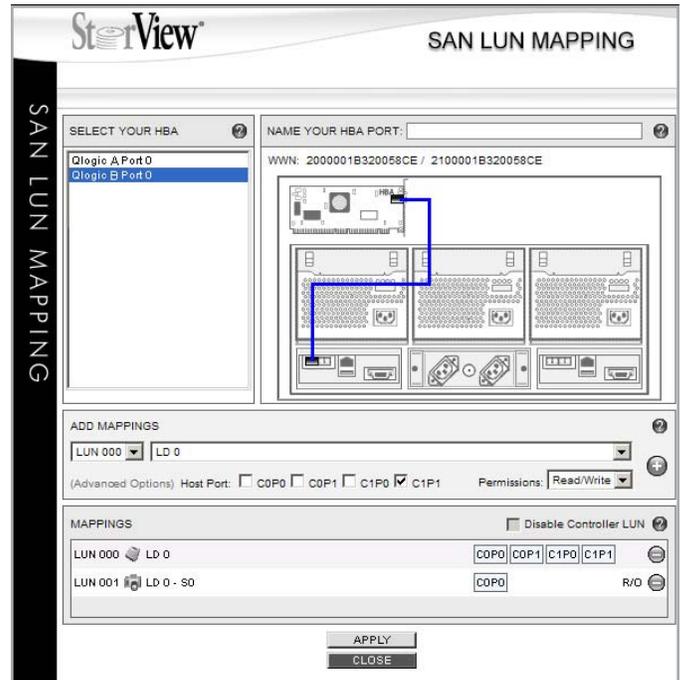


Fig. 117 - SAN LUN Mapping Screen

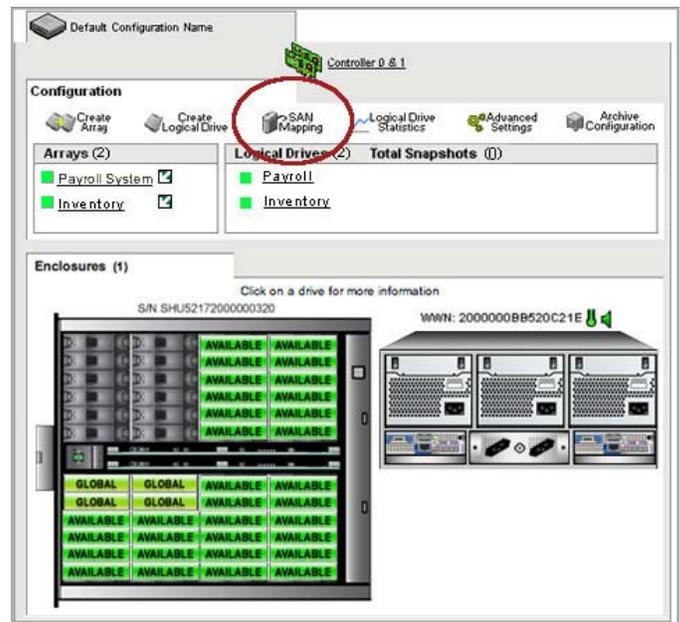


Fig. 118 - Main Screen – Enclosure/Array/Toolbar Section - SAN LUN Mapping Window

- a. Choose the LUN (logical unit number) to present the mapped logical drive to the Host system. Click the drop-down menu and choose the desired number.
- b. Select the logical drive you wish to map. Click the pull down window and select from the list of logical drives displayed. Default logical drive names are LD0, LD1, LD2, LD3, etc.

**NOTE:** Logical drives marked with an asterisk "\*" indicate that another mapping for another Host HBA has been established for this logical drive. You can map it again to additional HBA's but know that all Host HBAs mapped to this logical drive will see and have access to this logical drive. You will receive one warning about this process during a SAN LUN Mapping session. Also you must use file sharing or path failover software to ensure data integrity.

- c. Select an access permission for the mapping. Choose from the drop-down menu: Read/Write or Read Only.

**NOTE:** Microsoft Windows does not support Read Only permissions.

- d. Select the Host Port/Channel.

Click the check box(es) for the port that you wish the mapping to be available: C0P0, C0P1, C1P0, and/or C1P1. Where C0 represents the left or lower controller and C1 represents the right or upper controller, and P0 indicates the port marked FC Host 0 and P1 represents the port marked FC Host 1 on each controller. Ensure that the ports you select are properly cabled.

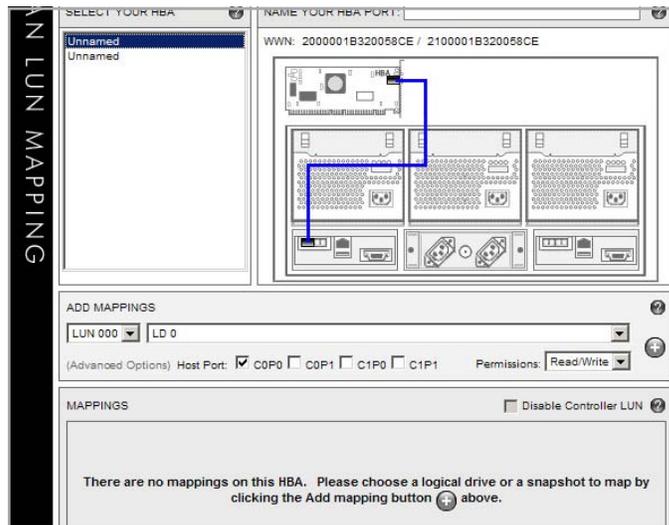


Fig. 119 - Main Screen – Enclosure/Array/Toolbar Section - SAN LUN Mapping Window



Fig. 120 - Shared Mapping Warning

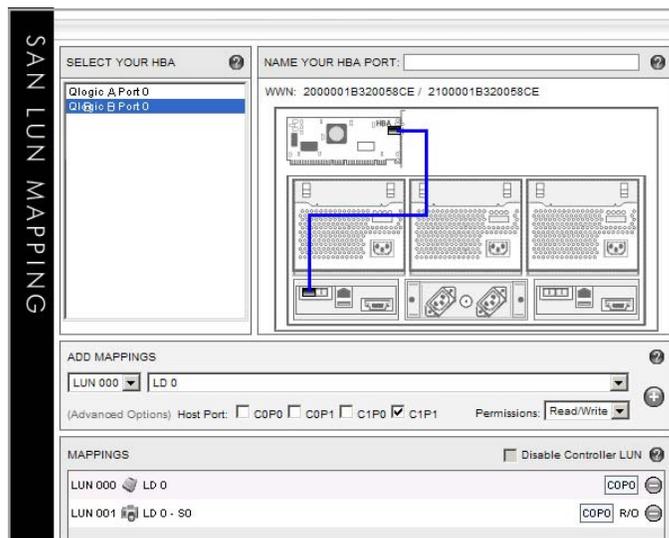


Fig. 121 - Shared Mapping Warning

## 10. Using Snapshot

### 10.1. Introduction

StorView's Snapshot is a licensed feature product. It has been designed to give the user considerable flexibility when setting up and performing snapshot and snapback operations. Snapshot fully supports Microsoft's Volume Shadow Copy Service (VSS).

A snapshot allows the user to create a point-in-time image of a logical drive that contains exactly the same data at the point the Snapshot was taken. Snapback allows the user to restore the logical drive to that point in time when the snapshot was taken.

For Windows-based systems, it is recommended to use the VSS provider to take and manage snapshots. Although you will use VSS to take the snapshots, StorView will still be used to create arrays, create logical drives, create ODA's, enable snapshot and perform snapback operations.

Unix-based systems will use the StorView snapshot functionality for their implementation of snapshot. This includes creating arrays, logical drives, ODA's, and enabling snapshot, in addition to snapshot management, taking the snapshots and performing snapback operations.

---

**NOTE:** APM level 2 cannot be enabled on an array that has Snapshot enabled.

---

For additional information, see [15. Optimizing Write Performance](#).

---

## 10.2. Enabling Snapshot

This process will enable the Snapshot feature by creating an Overwrite Data Area (ODA) and pairing it to a logical drive. You may create one ODA for the limited license and up to 256 ODA's with the maximum license.

**IMPORTANT INFO:** You cannot enable Snapshot support while an array expansion or array initialization is occurring.

Please wait for the operation to complete before attempting to enable Snapshot.

**IMPORTANT INFO:** If you intend to use the snapshot feature and choose to have access to this logical drive from your operating system or VSS, you must SAN LUN Map the logical drive. You may do this before or after the snapshots have been taken.

1. Ensure you have an array created to hold the ODA.

**NOTE:** It is recommended that you have an array created specifically for the Overwrite Data Area (ODA).

2. From the Main screen select a logical drive that will be snapshotted.

The Logical Drive Information window will appear.

**IMPORTANT INFO:** If this is the first time enabling Snapshot, you will have the ability to create one ODA and pair it to one logical drive. If you attempt to exceed the limits of the Free License, you will be prompted to enter your license key. To obtain a license key, contact your sales representative for feature licensing.

3. Click the ENABLE SNAPSHOT button.

**IMPORTANT INFO:** If you have not mapped the logical drive, you will receive the alert message in [Fig. 124](#).

When the Snapshot/ODA Management window opens, the following message will be displayed for logical drives that are not SAN LUN Mapped. If you are using a VSS Provider and the VSS Provider does not require the logical drive to be SAN LUN Mapped, ignore the message displayed.

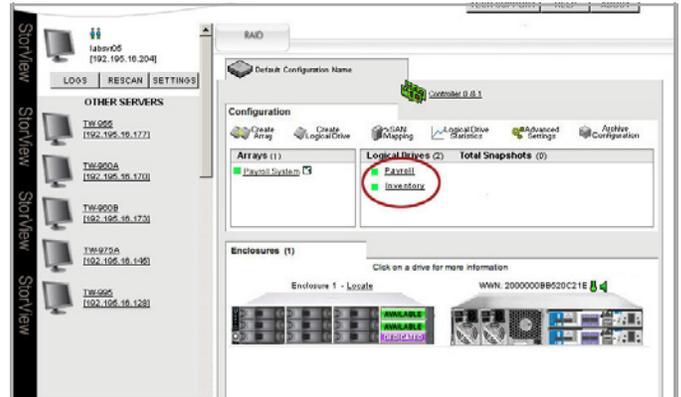


Fig. 122 - Main Screen

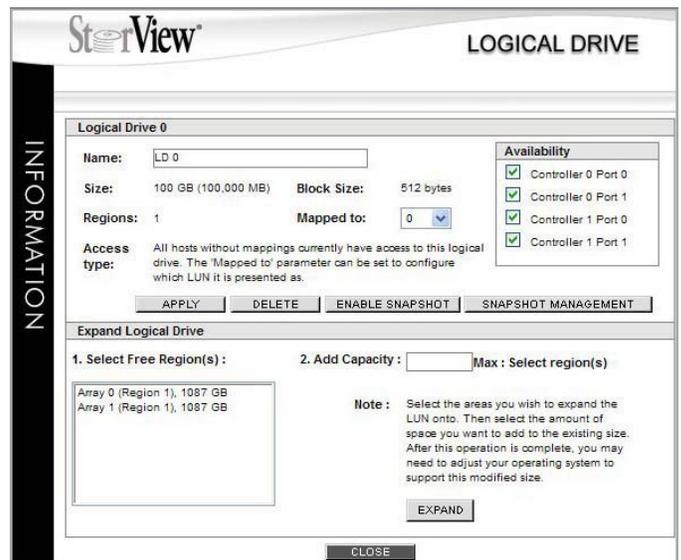


Fig. 123 - Logical Drive Information Window

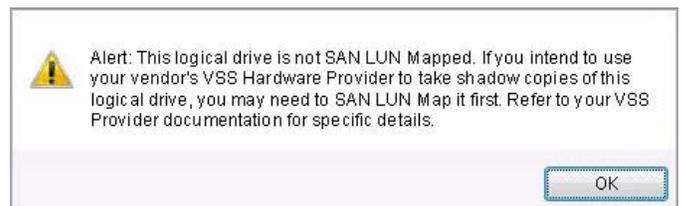


Fig. 124 - License Information Window

If you have exceeded the limited free license you will be prompted to upgrade by entering a license key. To upgrade continue with step 3, otherwise skip to step 5.

4. Enter your license key and click the Submit License Key button.  
You will receive confirmation that it was successful. If the procedure is unsuccessful ensure the key has been entered correctly and repeat. If it still remains unsuccessful contact customer support.

[continued on the next page >>](#)

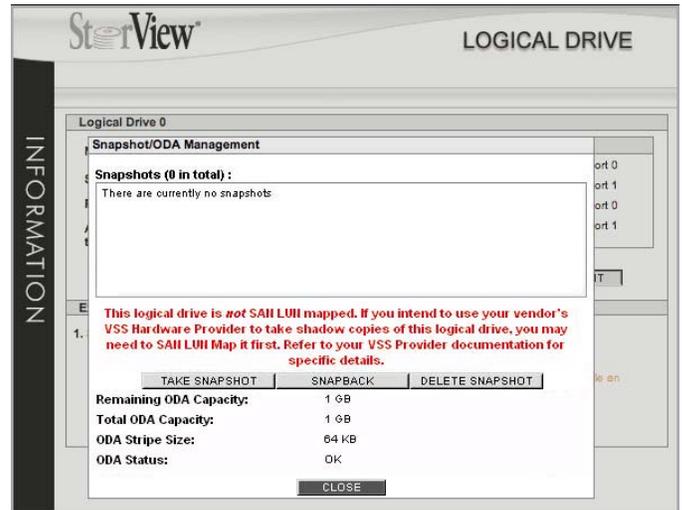


Fig. 125 - Logical Drive Information Window - SAN LUN Mapping Notice

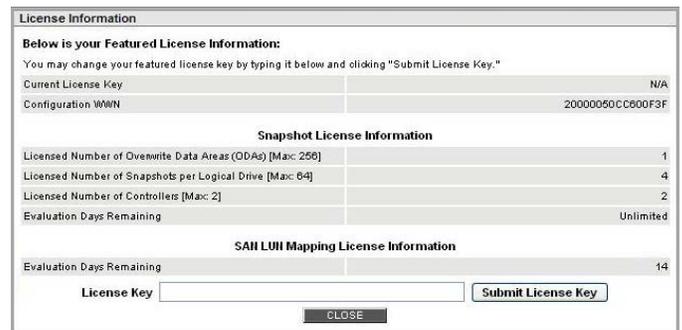


Fig. 126 - License Information Window

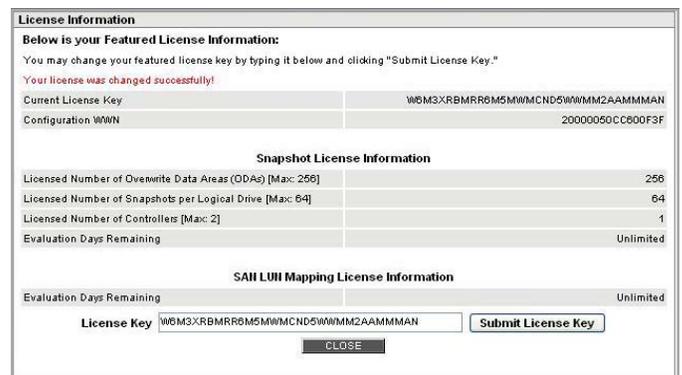


Fig. 127 - License Information Window

5. Select the array you wish to use for the Overwrite Data Area (ODA). This is where all the snapshot's will be stored for this logical drive.

**NOTE:** It is recommended that you have an array created specifically for the Overwrite Data Area (ODA). Choosing a different array than the array the logical drive uses will avoid any performance compromises. If necessary, return to the Main screen and create a new array for the ODA. (You will need to wait until the new array completes initializing and refresh the Logical Drive window for the new array to appear.)

By default the option "Show arrays used by this logical drive" is not selected. This is to deter the user from selecting an array that comprise the logical drive's arrays, as it is not the recommended choice. When the check box is selected in addition to the other disk arrays, the arrays for this logical drive are shown as well and can be selected.

6. Enter the size of the Overwrite Data Area in gigabytes (GBs).

It is highly recommended to consider future usage in order to determine the proper size to use based on the number of expected snapshots. Capacity planning is vitally important. If the Overwrite Data Area (ODA) size is exceeded, all snapshots associated with the logical drive become invalid.

7. From the drop-down menu, choose the desired stripe size (or write buffer size) for the Overwrite Data Area (ODA). The choices are: 64KB, 128KB, 256KB, 512KB, or 1MB.

It is recommended that you choose a stripe size based on your application's need. If you are making small changes, such as with a database, then you will want the smallest stripe size. If you are making large changes, such as with video application, then you will want a much larger stripe size.

**NOTE:** The stripe sizes available in the drop-down menu will change dynamically based on the size of the ODA entered. For example if you specify an ODA greater than 549 GBs, the selection for 64 Kb will not be displayed. If you enter an ODA size greater than 1 terabyte, then both 64 Kb and 128 Kb stripe sizes will not be displayed in the drop-down menu.

8. Click the ENABLE button. You will see a confirmation window.
9. Click the CLOSE button.

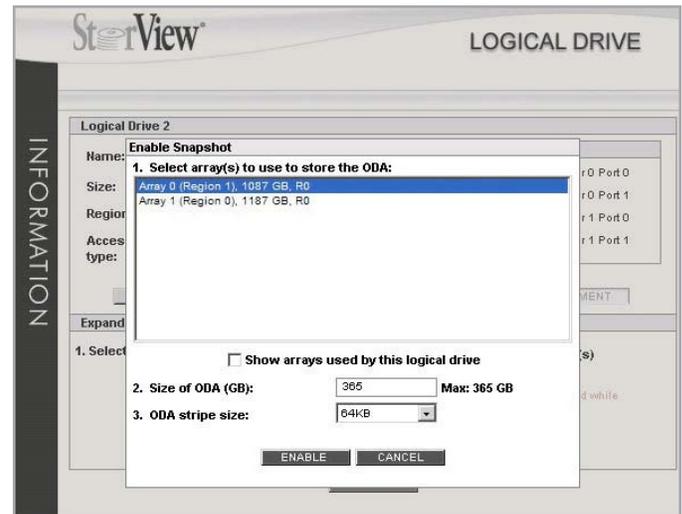


Fig. 128 - Logical Drive Information Window - Enable Snapshot

Table 15- Overwrite Data Area (ODA) Stripe Sizes

Stripe Size	Maximum Overwrite Data Area Size (ODA Size)
64 KB	549 GB
128 KB	1 TB
256 KB	2 TB
512 KB	4 TB
1 MB	8 TB

### 10.3. Performing a Snapshot

This process will take a point-in-time image of the selected logical drive. Snapshot's are automatically named, starting with "Snap 00" and will increment for each additional snapshot (Snap 01, Snap 02, etc.) for that logical drive, unless a Snapshot has been deleted then it will automatically name the snapshot with the next available name. The Time/Date stamp of the Snapshot is the governing authority as to which is the latest and correct file.

**IMPORTANT INFO:** If you are using Microsoft Windows Server and VSS (Volume Shadow Copy Service), it is recommended to use VSS as the preferred method for taking snapshots.

You can view your snapshots from the Snapshot/ODA Management window. However, it is important to note that the name will not always be the next sequential numbered name.

**IMPORTANT INFO:** You must ensure that applications writing to the logical drive are synchronized with the snapshot.

This can be performed in a variety of ways depending on the applications using the logical drive where the snapshot is being initiated. Buffers may need to be flushed for a consistent point-in-time image of the files on the logical drive.

1. From the Main screen select the logical drive to which you wish to take a snapshot.
2. From the Logical Drive window, click the SNAPSHOT MANAGEMENT button.
3. Click the TAKE SNAPSHOT button to begin the snapshot procedure.

You will receive a notice, prompting you to ensure the applications have been synchronized and the logical drive is dismounted to ensure a valid snapshot is taken.

You will see a confirmation window once it is complete. Click the CLOSE button.

continued on the next page >>

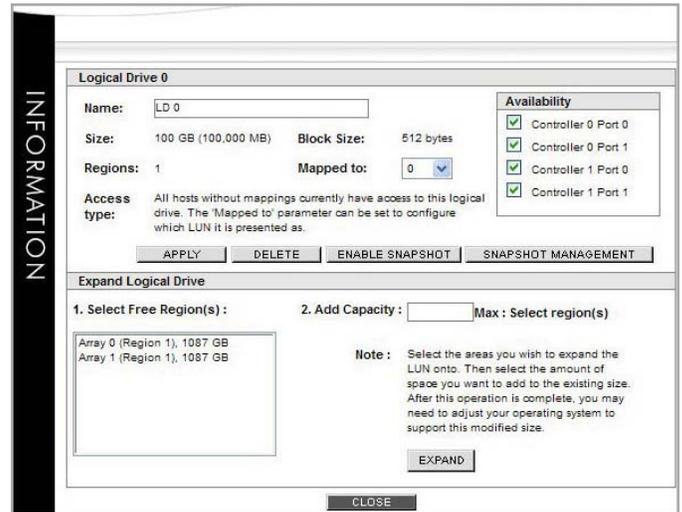


Fig. 129 - Logical Drive Information Screen - Before Taking a Snapshot

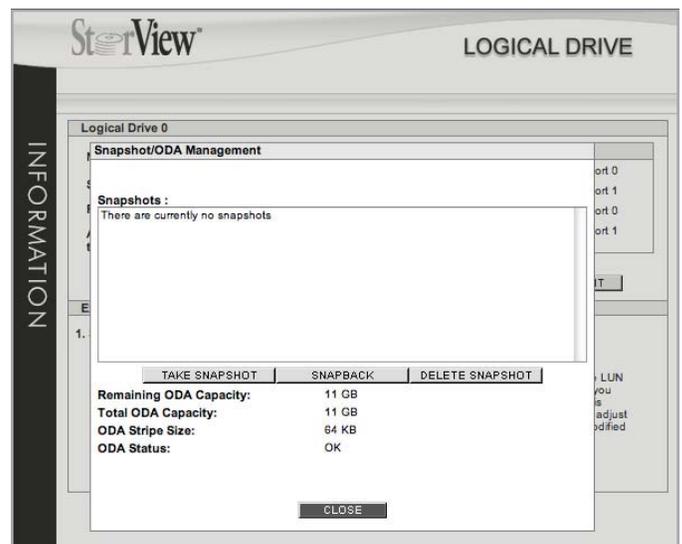


Fig. 130 - Taking a Snapshot

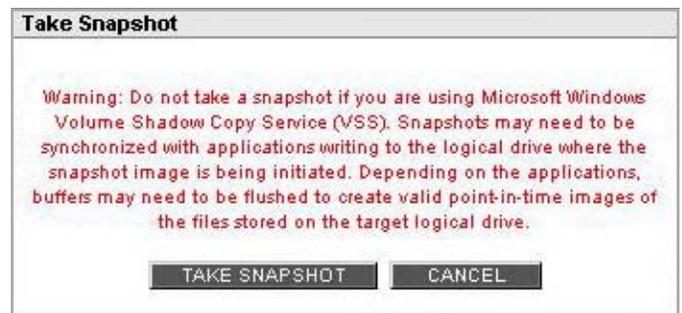


Fig. 131 - Snapshot Message

4. After the snapshot is complete, you will see the following in the Snapshot/ODA Management window.
5. Click the CLOSE button on the Snapshot/ODA Management window.
6. Click the CLOSE button on the Logical Drive Information window.

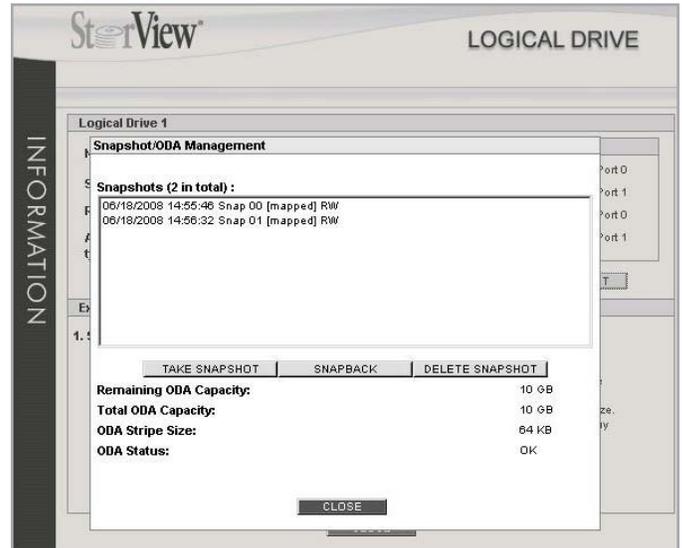


Fig. 132 - After Taking a Snapshot

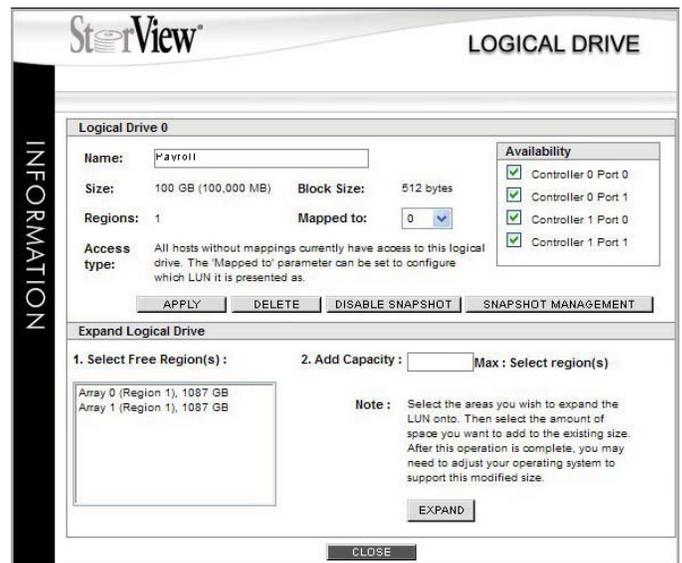


Fig. 133 - Logical Drive Information Screen - After Taking a Snapshot

## 10.4. Mapping a Snapshot

To make the snapshot available to the host, you will need to perform a SAN LUN Mapping of the snapshot.

**IMPORTANT INFO:** If you intend to use the snapshot feature and choose to have access to this logical drive from your operating system or VSS, you must SAN LUN Map the logical drive. You may do this before or after the snapshots have been taken.

1. From the Main screen click the SAN Mapping icon in the Tool Bar.
2. Select the Snapshot from the drop-down Logical Drive menu that you wish to map.

In the example in [Fig. 135](#), the “LD0-S0” entry represents Snapshot 0 of Logical Drive 0.

It is important to note that snapshots are not governed by their numbers (S0-S63) as to whether one is a later version than another. Also, if you delete a snapshot, the remaining snapshots do not automatically re-number. Instead the next snapshot will be assigned the lowest available number. So if you had 10 snapshots (S0-S9) and deleted snapshot S02, the next snapshot taken will be given the name S02. To determine a specific snapshot in time, refer to the snapshots time-date stamp. Choose the host port check boxes you wish the snapshot logical drive to be available, then click the APPLY button.

3. Click the CLOSE button.
4. Access your host operating system and application software to backup or retrieve data from the snapshot logical drive.

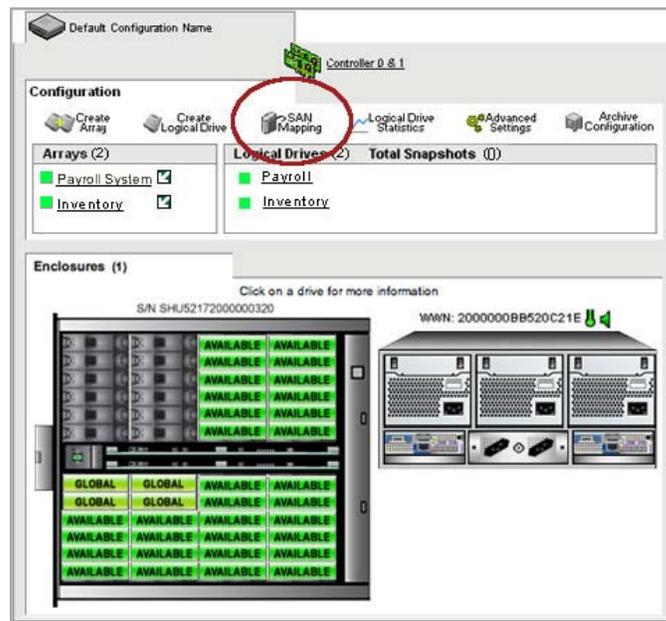


Fig. 134 - Main Screen – Enclosure/Array/Toolbar Section - SAN LUN Mapping Window

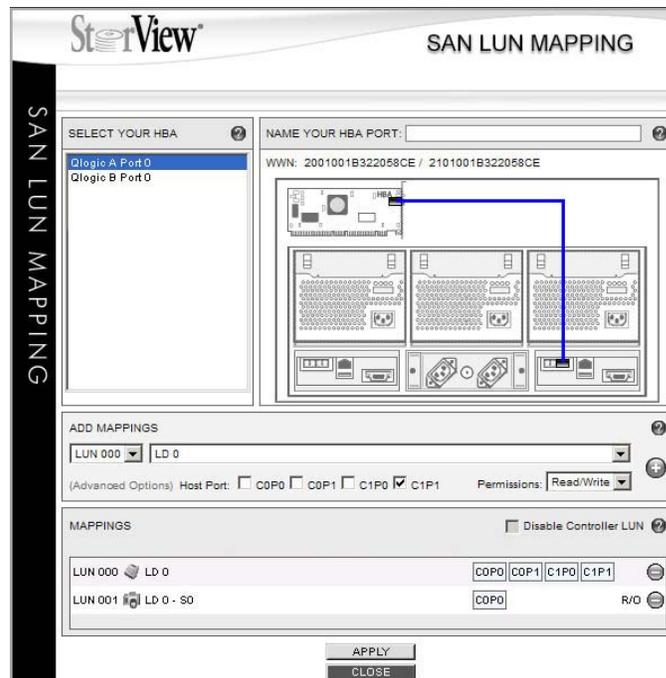


Fig. 135 - Selecting Snapshot to Map

## 10.5. Deleting a Snapshot

This option allows you to delete a snapshot taken on a specific logical drive. This procedure will not affect the other snapshots. After the snapshot is deleted, it makes room for a new snapshot.

1. From the Main screen select the logical drive with the Snapshot icon adjacent to it.

The Logical Drive Information window will appear.

2. Click the SNAPSHOT MANAGEMENT button.

[continued on the next page >>](#)

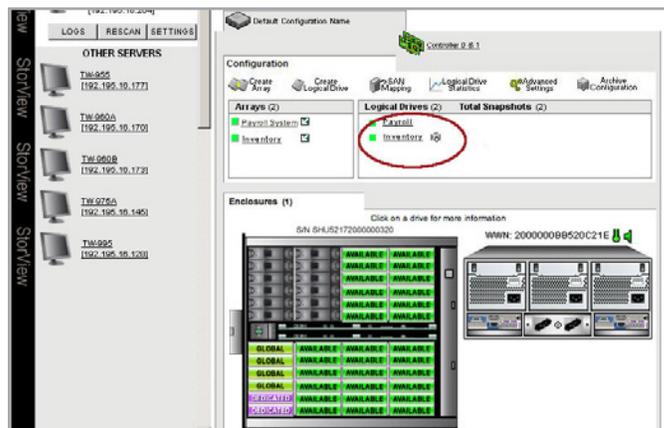


Fig. 136 - Main Screen

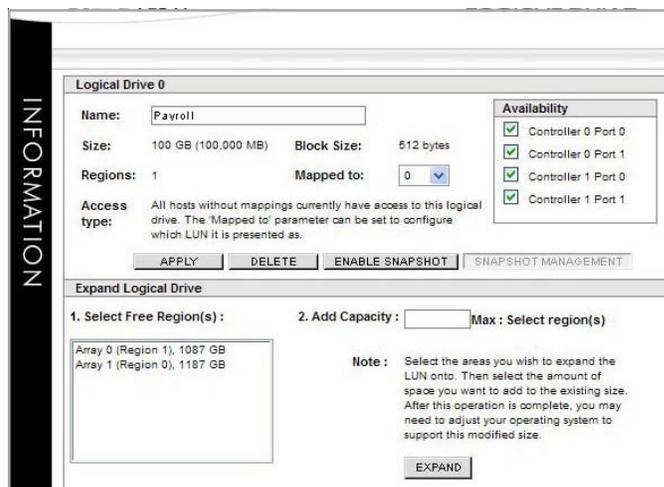


Fig. 137 - Logical Drive Information Window

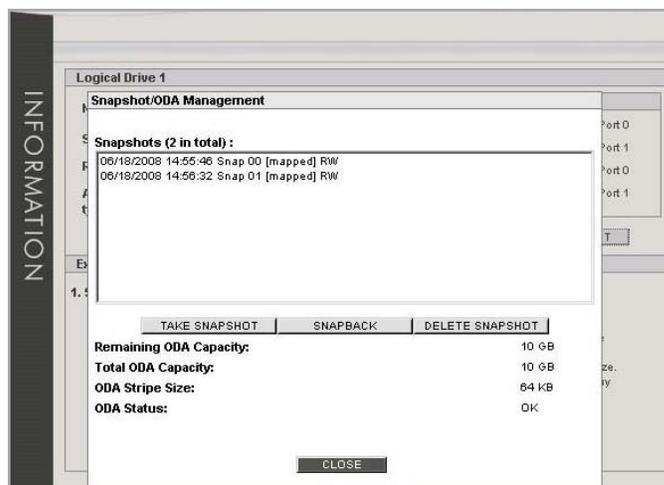


Fig. 138 - Snapshot Information Window

3. Select the snapshot you wish to delete from the “Snapshots” list window, then click the DELETE SNAPSHOT button.

You may choose one snapshot or all snapshots. To choose all snapshots select the first and while holding down the shift key select the last snapshot.

You will see a notice window identifying the snapshot(s) about to be deleted.

4. Click the DELETE SNAPSHOT button.
5. You are prompted to enter your password. Type your password and click the GO button.
6. You will see a message that the execution was successful, click the CLOSE button.

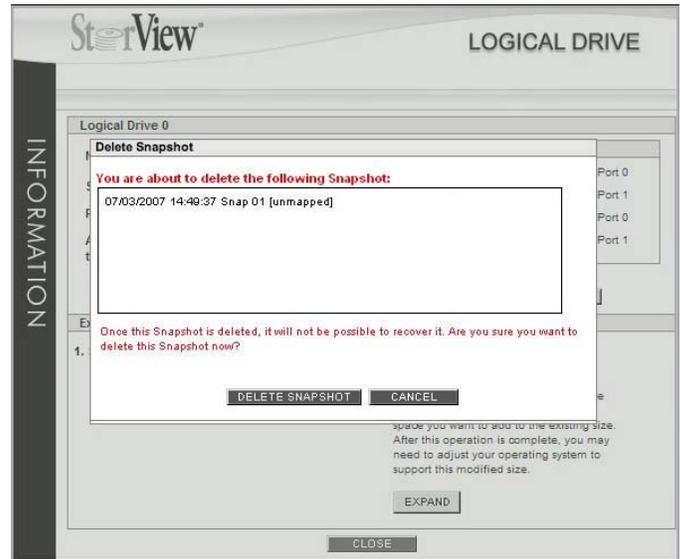


Fig. 139 - Snapshot Warning Window



Fig. 140 - Password Confirmation Window

## 10.6. Performing a Snapback

This option allows you to restore the logical drive to the state when the snapshot was taken.

**IMPORTANT INFO:** Before you perform a snapback operation, be sure to back up the data on the existing logical drive first. When the snapback command is issued, the existing data will be overwritten and replaced with an exact copy of the Snapshot image.

**IMPORTANT INFO:** Use great care when using the Snapback feature. Refer to the example below: In this example a 200GB logical drive is created and mapped. Two Volumes or Disks are created, 100GB each, from the logical drive completely separate of each other and labeled them D: and E:. The logical drive has snapshot enabled and therefore means both disks or volumes are covered by a single snapshot or shadow copy. The disks or volumes are used independently and snapshots are taking whenever necessary. At a point in the time you decide to recover volume or disk D: using StorView's Snapback feature. Remember that the entire logical drive will be restored which could affect and/or LOSE data on the E: volume or disk. It is important that you understand that the snapback feature will restore the data on both volumes D: and E: at the point in time the snapshot was taken. You must be aware of the state of data when creating multiple disks or volumes of the same logical drive, and using snapshot and snapback features.

1. If necessary, perform a backup of the logical drive before performing the Snapback operation.
2. From the Main screen select the logical drive with the Snapshot icon adjacent to it.

The Logical Drive Information window will appear.

3. Click the SNAPSHOT MANAGEMENT button.

[continued on the following page >>](#)



Fig. 141 - Main Screen – Arrays and Logical Drives: Selecting Logical Drive

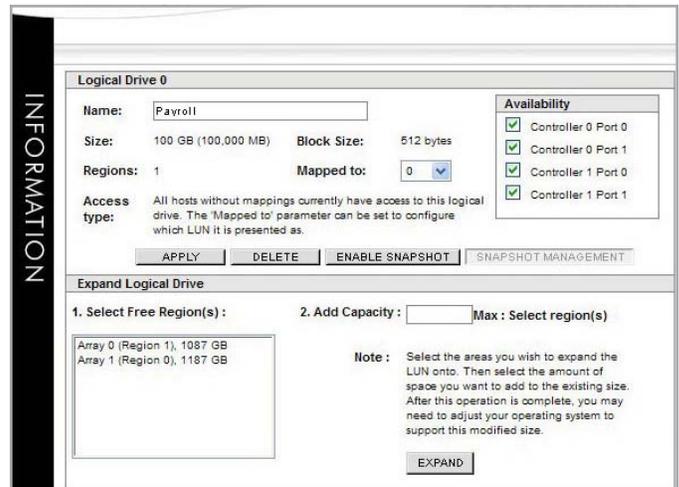


Fig. 142 - Logical Drive Information Window

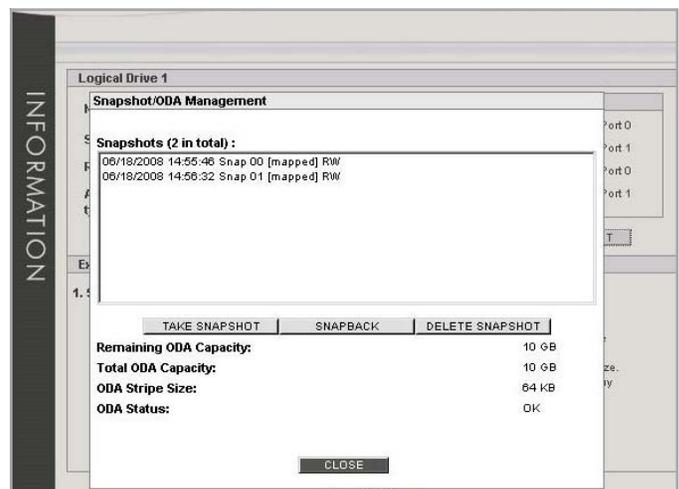


Fig. 143 - Snapshot Information Window

4. Select the snapshot you wish to snapback from the "Snapshots" list window, then click the SNAPBACK button.
5. Select which type of Snapback you wish to perform, Original or Modified and click the appropriate button.
6. Click the CLOSE button on the Snapshot/ODA Management window.
7. Click the CLOSE button on the Logical Drive Information window.
8. If the snapback is rather large and takes a little bit of time, you will see a progress display on the Main screen, located under the logical drive icon.

The logical drive is now restored to it's point-in-time of the snapshot.

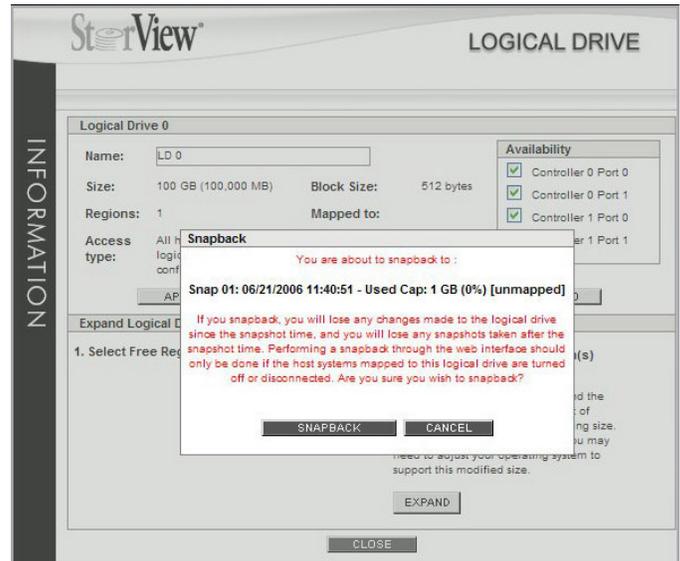


Fig. 144 - Snapshot Warning Window



Fig. 145 - Snapback Options Window



Fig. 146 - Snapback Progress Display

## 10.7. Disabling Snapshot

This operation will remove the Snapshot feature which results in deleting an ODA. The process will unpair the ODA from the logical drive. When performing this operation, all snapshots for this logical drive will be lost.

1. From the Main screen select the logical drive to which you wish to disable Snapshot support.

The Logical Drive Information window will appear.

2. Click the DISABLE SNAPSHOT button.
3. You will be prompted to enter your password. Enter the password and click the GO button.

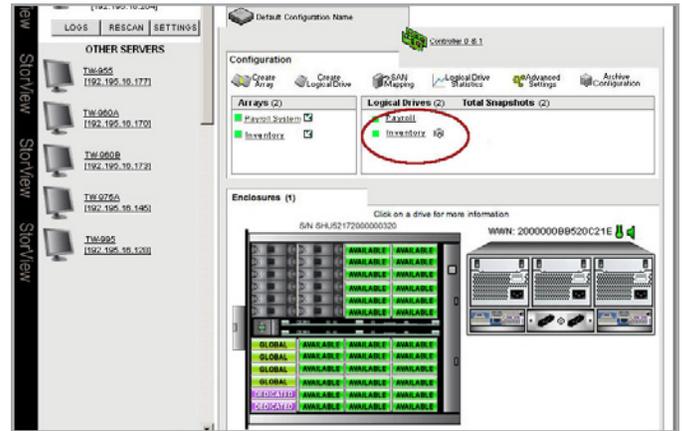


Fig. 147 - Main Screen – Arrays, Logical Drives, and Enclosures Sections

**CAUTION:** Disabling Snapshot support will result in the loss of all Snapshots for this logical drive.

4. Click the DISABLE button.
5. You will receive a confirmation window, click the CLOSE button.
6. Click the CLOSE button on the Snapshot/ODA Management window.
7. Click the CLOSE button on the Logical Drive Information window.

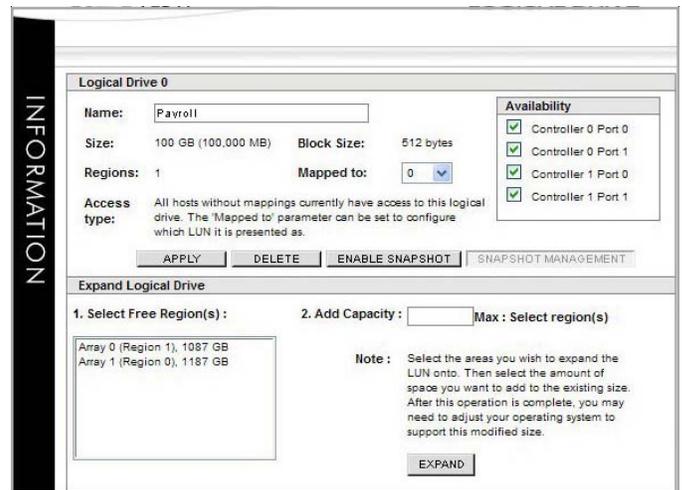


Fig. 148 - Logical Drive Information Window

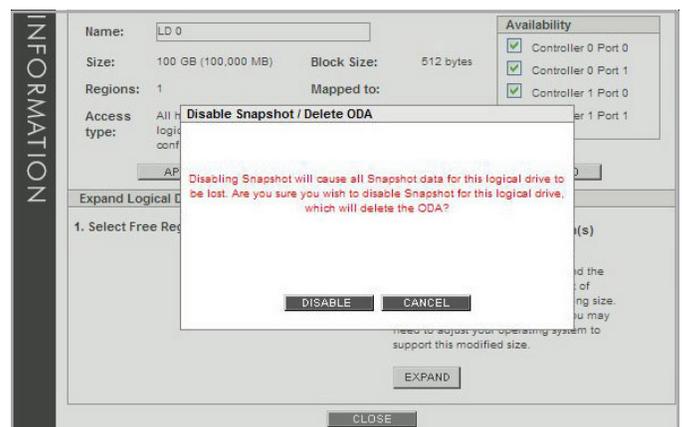


Fig. 149 - Disable Snapshot Window

# 11. Failover, Performance & Additional Functions

## 11.1. Overview

In this chapter you will find more information on the following:

- ◆ How StorView failover operates
- ◆ Performance optimization processes
- ◆ Additional Functions
  - About
  - Take Control Monitoring
  - Rescan
  - Alarm Mute and Enable/Disable

---

## 11.2. How StorView Server Failover Works

The failover feature of the StorView RAID Module must have one of the following setups.

If you are using the host-based version of StorView, two or more host servers must be directly connected to the RAID storage system enclosure and each host must have a copy of StorView installed with a Global Access license.

If you are using the embedded version of StorView, you must have dual controllers installed (duplex topology) each with embedded StorView installed.

At startup, each StorView server will create a device list of all the attached storage systems. It then sends the list out on the network as a device list packet. The other StorView servers on the network will then respond with their device list packets. Since the StorView servers are attached to the same storage solution, they will have the same or similar devices in their device list packet. The identical devices in each device list packet will be flagged. After analysis, the StorView server with the lowest serial number or address will take control of those devices. The other StorView server will indicate on its interface that another StorView server is monitoring the storage solution.

During normal operations, the StorView server(s) send “check-in” packets every 10 seconds. If three consecutive check-in packets for a specific StorView server are not received, its devices are flagged and the StorView server with the lowest serial number or address will take control of those devices.

---

## 11.3. StorView Performance Optimization

There are some parameters that can be adjusted on the host HBA and operating system to increase the performance of StorView. They are HBA Execution Throttle setting and the operating system’s Scatter/Gather registry setting.

---

### 11.3.1. Execution Throttle

To improve general I/O performance by allowing more commands on the fibre bus or the SCSI bus, we recommend changing your host bus adapter’s execution throttle parameter to 256. Refer to your host HBA documentation for more information.

---

### 11.3.2. Scatter/Gather

(Microsoft Windows Only) To increase general I/O performance by allowing larger data transfers we recommend editing the “MaximumSGList” parameter in the registry. The recommended hexadecimal value is “ff.” The path is:

`HKEY_LOCAL_MACHINE/System/CurrentControlSet/Services/<name of HBA driver>/Parameters/Device/`

Refer to your Microsoft Windows operating system documentation for more information on editing the registry and your host HBA adapter documentation.

---

## 11.4. Additional StorView Functions

### 11.4.1. About

Clicking this button displays the software version information and the type of license installed. When using embedded StorView the About window also provides the control button to update the software, see section [12.2. Updating Embedded StorView](#).

1. From the Main screen, click the ABOUT button, located in the upper right corner of the window under the StorView logo.

The About window is displayed. The license type for this installation is indicated below the version number in parenthesis.

2. Click the CLOSE button on the About window.

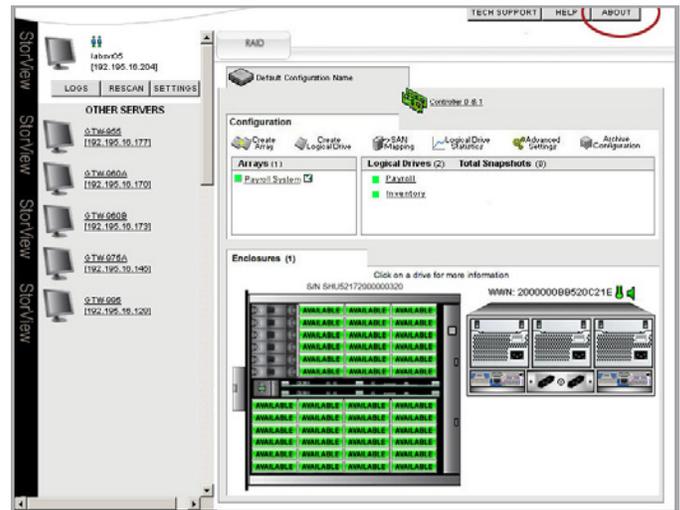


Fig. 150 - Main Screen

### 11.4.2. Take Control Monitoring

When multiple host servers are physically attached to the same storage system or the hosts are attached to a fibre switch in which the switch is physically attached to the storage solution, the installed StorView server with the lowest serial number or IP address will take control of those devices.

If you wish to take control of the storage solution from another StorView server, click the "Take Control" link from the Main screen message displayed on that StorView server.

After clicking the "Take Control" link on the Main screen, StorView will perform a scan and reload the configuration for this system. The StorView Server which previously had control will now display the message that the selected storage solution is being monitored by another StorView server.

You would also see this condition if this StorView server failed to send the required three consecutive check-in packets and it fails over to another StorView server attached to the same storage system. Once the problem is resolved on this host StorView Server, you can take back control again with that StorView server.



Fig. 151 - About Window (Host-based StorView)

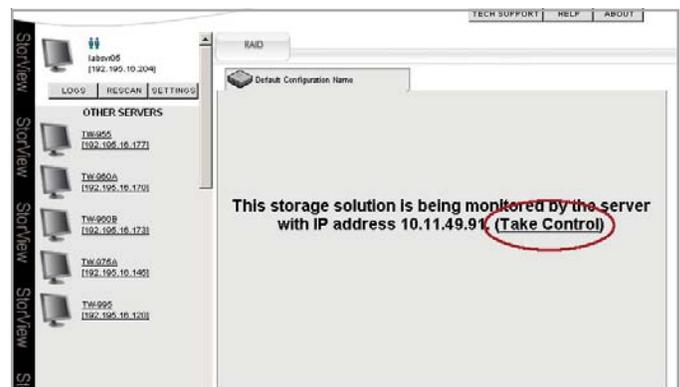


Fig. 152 - Main Screen Take Control - Monitoring Screen

### 11.4.3. Rescan

This Rescan function will force a search to re-discover storage solutions in order to refresh the display and update status information.

From the Main screen, click the RESCAN button located on the far left side of the screen. After a few moments the Main screen is re-displayed with updated information.

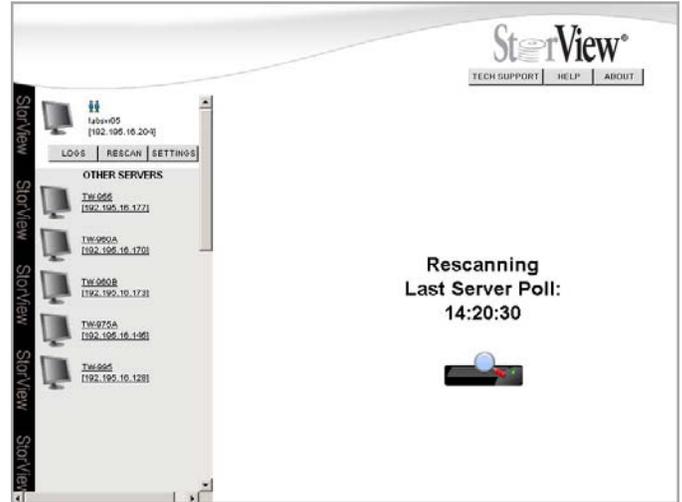


Fig. 153 - Main Screen in Rescan Mode

## 12. Support & Updates

### 12.1. Tech Support

This feature allows the user to provide technical support personnel with event and configuration information to assist with troubleshooting.

1. From the Main screen, click the TECH SUPPORT button, located in the upper right corner of the window under the StorView logo.
2. Enter the requested information for each field.

The "Problem" field is scrollable, allowing you to review the information you will be sending.

---

**NOTE:** The gathering of this information may take a few minutes.

---

3. Click the DOWNLOAD button.

You will receive a screen prompt to save the file on your local disk. Enter a name for the file and click Save. The software will create a file with your user data, a capture of the event log files, and a capture of all configuration information. Technical support representatives will be able to work with this information to assist with solving your problem.

4. Click the CLOSE button on the Technical Support window.
5. When requested by a technical support representative, E-mail the saved file.

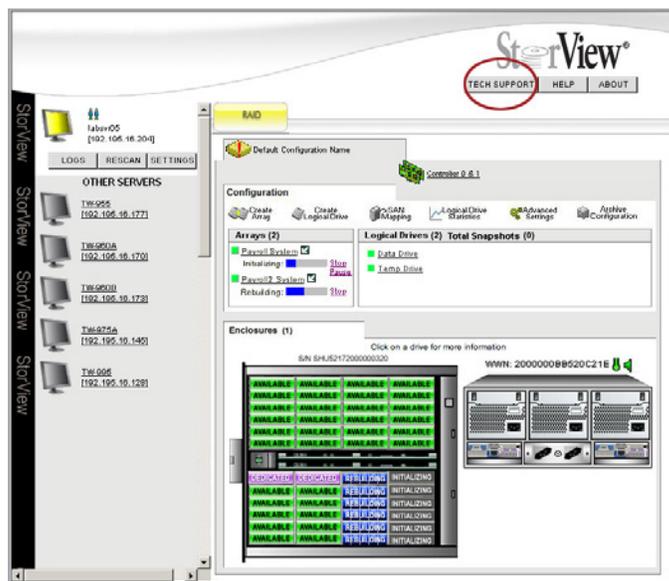


Fig. 154 - Main Screen

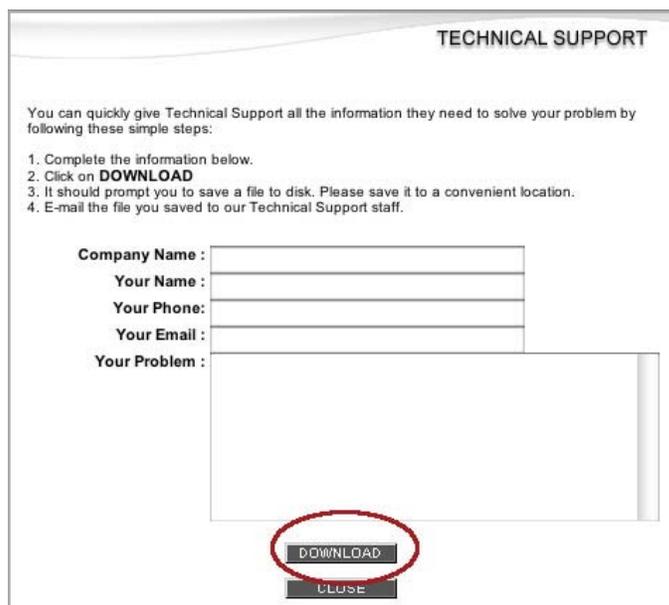


Fig. 155 - Tech Support Screen

## 12.2. Updating Embedded StorView

To update the Embedded StorView software:

1. Click the About button.

---

**CAUTION:** Ensure there is uninterrupted power during the update.

---

2. Click the UPDATE button.
3. Type the name of the firmware file or click the Browse button and locate the file. The file name will be similar to storview-3.09.xxx-zzz-nb-en.bin.
4. Enter your login password and click the UPLOAD button.

Once the update is complete, StorView Server will automatically restart. This process will not affect I/O activity.

---



Fig. 156 - About Screen (Embedded StorView)

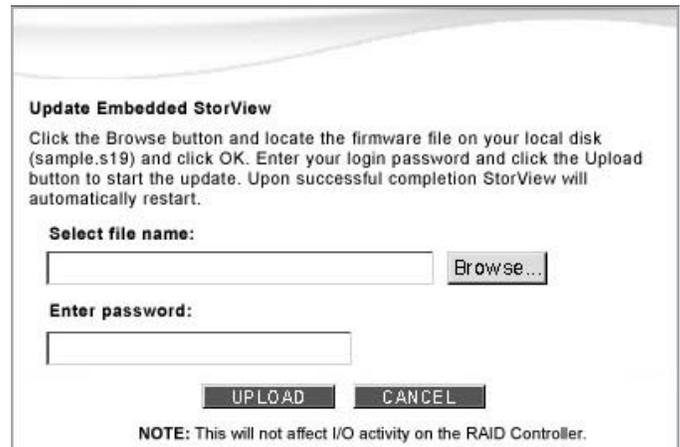


Fig. 157 - About Update Screen (Embedded StorView)

## 13. Event Logs

### 13.1. Overview

StorView has the ability to manage the events generated by the controller and enclosure components.

StorView also has its own unique set of events that are related to the StorView server component of the software. Events can be used to assist with troubleshooting, monitoring components, and determining the status of hardware components, logical drives and arrays. The following event types are logged:

- ◆ Advanced Power Management
- ◆ Alarms
- ◆ Controller, Controller Ports, and Expansion modules
- ◆ Drive and Array
- ◆ Enclosure components
- ◆ Fibre Loop and SAS/SATA Bus (Drive and Host)
- ◆ Network communication (applicable to the host-based version only)
- ◆ Persistent Reservations
- ◆ StorView server (applicable to the host-based version only)
- ◆ Snapshot

There are two event logs maintained: one set of log entries the controller maintains and one set StorView maintains. There are some differences and limitations between the controller set of event logs and StorView's set of event logs. The differences include the type of events logged and in some cases the ease of interpretation. StorView event logs provide a user friendly language to describe the event.

The controller's maximum event log size is 4096 entries with the oldest events being overwritten as the log reaches the size limit. Some repetitive events are appended to previous events, so entries are not used up unnecessarily. The controller logs are managed by clicking the Controller icon and accessing the Operation tab. From there you can export the controller logs to an external file or clear the log entries.

Enclosure event monitoring can be disabled which reduces the polling that StorView performs thereby increasing the performance that may be necessary in certain applications. This is accomplished by deselecting the Enclosure Support option in the Advanced Settings window. Disabling this option will stop enclosure component monitoring, which can be noted on the Main screen by the dimming of the enclosure rear view graphics and a notation above the graphic stating "Enclosure Information Not Available" see [6.2. Controller Environmentals](#).

StorView's event log will maintain the controller's compilation of events and the software's specific events. The controller's compilation of events include Controller Events (those unique to the RAID Controller), Drive Events and Host Events, and if the Enclosure Support option is enabled, enclosure component events.

---

**NOTE:** If the "Enclosure Support" option is disabled, E-Mail notifications established for those enclosure events will not occur.

---

The StorView server will also perform a synchronization of its event log to the controller log when the StorView server starts. Since the controller(s) can continue to operate when StorView Server is shut down, the StorView log would have missing events during this down period. The event synchronization feature of StorView will append the log with the controller events that occurred while the StorView server was shutdown.

The time stamp for each event in the StorView log is the exact time the event was received by StorView, and can be slightly off for the actual time it occurred in the controller log. After synchronization, events that occurred while the StorView server was down are marked with an additional string in the event description which displays the actual time stamp that event occurred. The string will be in the form of an asterisk followed by the time and date in parenthesis. At the bottom of the Event Log window you will find the footnote "\* Indicates event occurred while Server module was down." This indicates that those events with this extra time stamp in the description are the results of a synchronization and displays the exact time the event actually occurred.

StorView's event log has a maximum size limited only by the available disk space, therefore the log events in StorView will require regular maintenance to ensure the list is manageable and doesn't fill to disk capacity. You can export the log files to a comma delimited file prior to clearing them for later use. StorView also supports the SYSLOG feature. Refer to Section see [4.5. SYSLOG](#).

---

### 13.2. Accessing and Navigating the StorView Event Log

To access the Event Logs, click on the LOGS button located under the focused StorView server icon on the Main screen.

The embedded version of StorView does not support some components of Event Log window. The unsupported components include the Log # column, Log # window, Jump To button, Export button, and the Clear Log button. Clear log is now handled by the clear controller log function found in the Environmentals section.

Table 16- Event Type Icons

Icon	Description
	Information
	Warning
	Error

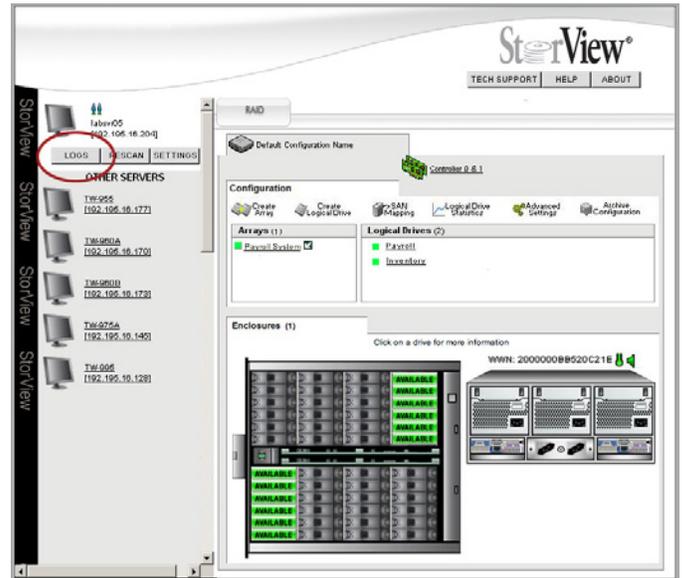


Fig. 158 - Main Screen

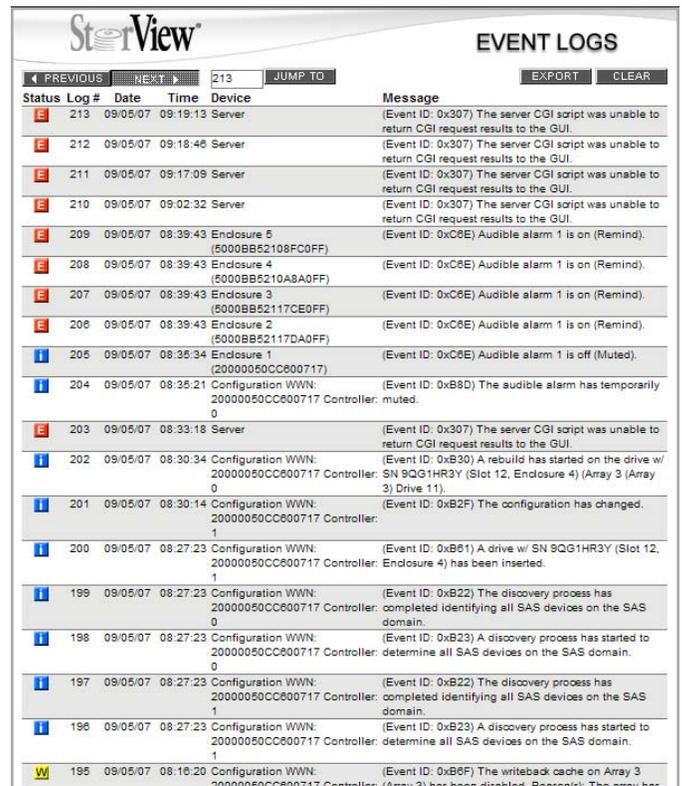


Fig. 159 - Event Log Screen

### 13.3. Exporting the StorView Event Log

**NOTE:** This option is applicable to the host-based version only.

The event logs can be exported to a comma delimited file for use in third-party software products when using the host-based version of StorView. Some web browsers provide more options for the format of the file. Refer to your browser software for specific details.

1. To export the log file, click the LOGS button on the Main screen for the storage system you are logged into.
2. Click the EXPORT button in the Event Logs window.

continued on the next page >>

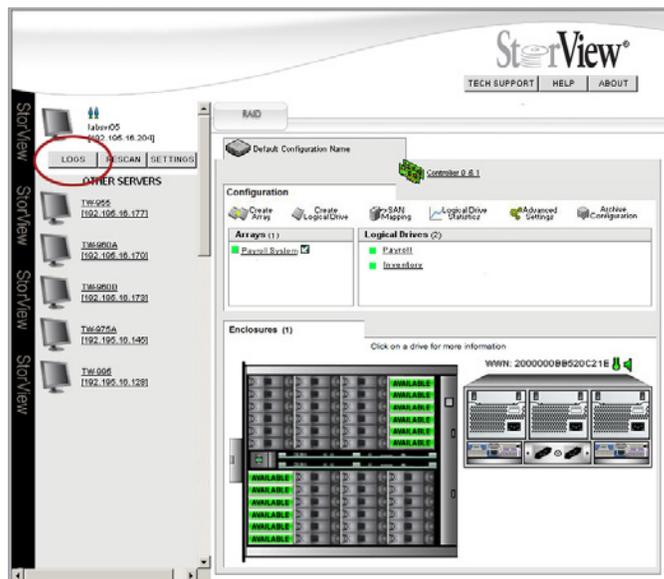


Fig. 160 - Main Screen

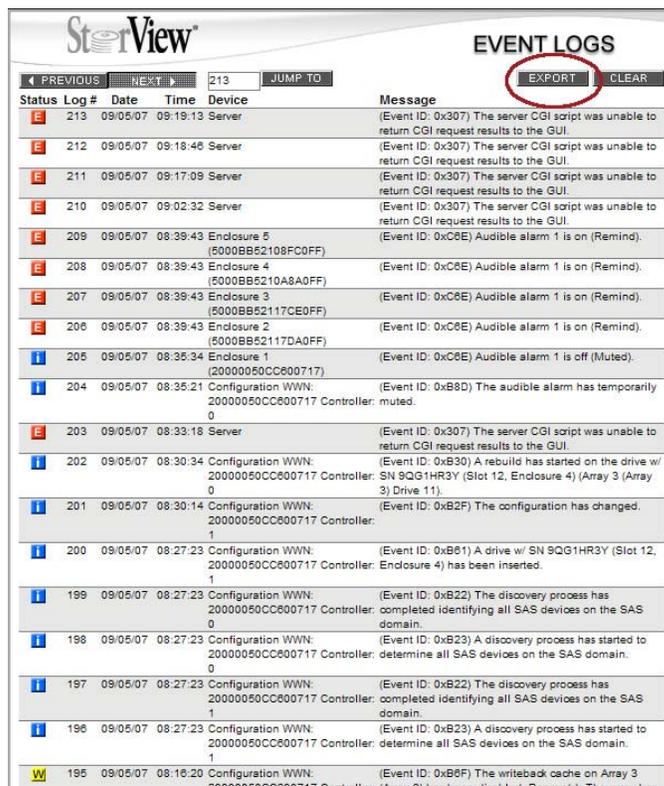


Fig. 161 - Event Log Screen

The examples on the next page show some web browser file export options, your browser may be slightly different.

- ◆ If you are using Internet Explorer as your web browser, you will see [Fig. 162](#). Choose to save the file or open it. The saved file format will be a comma delimited format.
- ◆ If you are using a Mozilla type web browser, you will see [Fig. 163](#). This product provides you with a few more options through the “Advanced” button.

Click the Advanced button and select the file format type, creator application, and other options, as desired.

3. Click the CLOSE button on the Event Logs window.



Fig. 162 - Example of Export Log Event Options - Internet Explorer

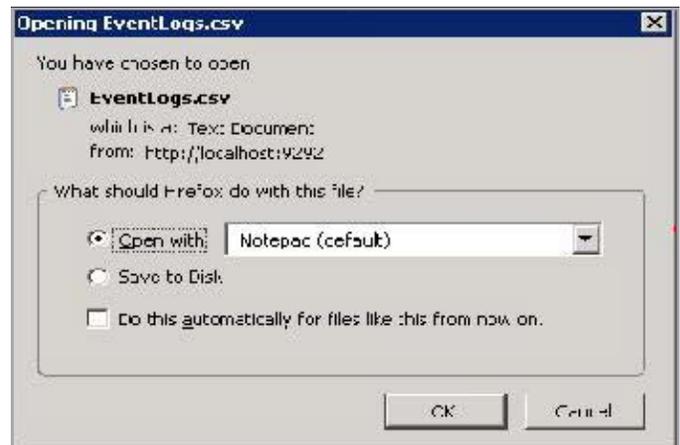


Fig. 163 - Example of Export Log Event Options - Internet Explorer

### 13.4. Clearing the StorView Event Log

You can clear the StorView event log retained on the host server from this window when using the hostbased version of StorView.

1. To clear the log file, click the LOGS button on the Main screen for the storage system you are logged into. The Event Logs window will open.
2. Click the CLEAR button on the upper right corner of the Event Logs window. The following window appears.
3. Click the OK button to clear the log files and continue.
4. Click the CLOSE button at the bottom of the Event Logs window.

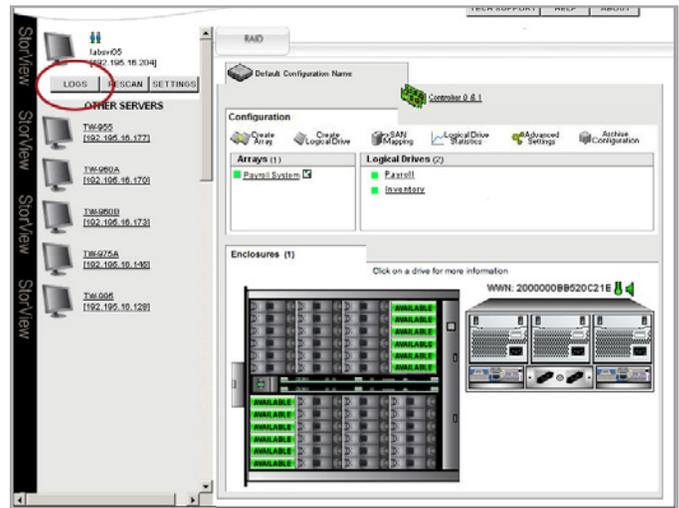


Fig. 164 - Main Screen

### 13.5. Operating System Event Log

StorView is capable of passing all the events to the host operating system event log. Accessing the operating system event logs will display events generated by StorView. Each event is identified by an Event ID. In the tables for the events you will see the Event type followed by its ID. The ID is given in the format of its hexadecimal value and its equivalent decimal value in parenthesis. The Event number is how the events are displayed in the operating system event log, and the decimal value is the format the OS event log will use to display the Event ID.

You can double-click the specific event in the operating system log and it will display a window with a plain english description of the event.

Also, you can use the tables to locate the event ID and determine the possible cause of that event and suggested actions to take if necessary.

StorView events are placed into the application event logs.

To shutoff OS event logging, edit the following file using a text editor:

```
<install directory>/db/server.ini
```

1. Change the field "UseOsEventLog" from "true" to "false."

UseOsEventLog = true enables event logs to be sent to the OS Event log, and,

UseOsEventLog = false disables event logs being sent to the OS Event log

2. At the Main screen click the RESCAN button. After the rescan is complete events will no longer be sent to the Windows operating system event log.

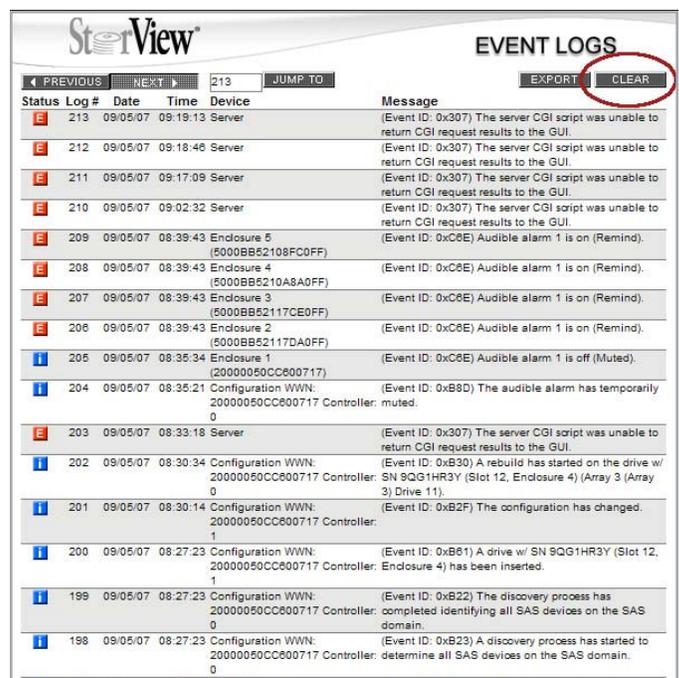


Fig. 165 - Event Logs Window (Clear Logs)



Fig. 166 - Confirmation Window

## 13.6. List of Events

Events in this chapter are categorized and listed in the order of their individual Event Type [ID]. As this is a comprehensive list of all events, not every event is applicable to every controller or enclosure model.

### 13.6.1. Advanced Power Management Events

These events are related to Advanced Power Management, APM level 1 and APM level 2. The ID is given in the format of its hexadecimal value which is seen in Unix operating systems and its equivalent decimal value in parenthesis which is seen in the Microsoft Windows operating system.

Table 17- Advanced Power Management Events

Advanced Power Management Event Messages	Type [ID]	Cause	Action
APM drive test failed (test result: w) on drive w/ xxxxxxxx (Slot y, Enclosure z). Sense Data aa/ bb/cc.	Error [0xB1D (2845)]	Unrecoverable problem on APM level 1 drive.	Check the drive.
Drive w/ SN <Serial number> (Slot <Slot number>, Enclosure <Enclosure number>) executes a spin up cycle.	Information [0xB93 (147)]	User or software commanded a drive spin up.	No action necessary.
Drive w/ SN <Serial number> (Slot <Slot number>, Enclosure <Enclosure number>) executes a spin down cycle.	Information [0xB93 (147)]	User or software commanded a drive spin down.	No action necessary.
These drives execute a spin down cycle (unused drives: <Unused drive count>, spares: <Spare count>, failed drives: <failed drive count>).	Information [0xB94 (148)]	User or software commanded spin down on the identified drive(s).	No action necessary.
These drives execute a spin up cycle (unused drives: <Unused drive count>, spares: <Spare count>, failed drives: <failed drive count>).	Information [0xB94 (148)]	User or software commanded spin up on the identified drive(s).	No action necessary.
<Array drive count> drives from array with ID <Array number> execute a spin down cycle.	Information [0xB95 (149)]	User or software commanded spin down on the drives in the identified array.	No action necessary.
<Array drive count> drives from array with ID <Array number> execute a spin up cycle.	Information [0xB95 (149)]	User or software commanded spin up on the drives in the identified array.	No action necessary.
APM drive test started on <Drive count> drives.	Information [0xB96 (150)]	User or software executed a drive test.	No action necessary.
APM drive test started on <Drive count> drives from array with ID <Array number>.	Information [0xB96 (150)]	User or software executed a drive test on the drives in the selected array.	No action necessary.

Advanced Power Management Event Messages	Type [ID]	Cause	Action
APM drive test completed on <Drive count> drives.	Information [0xB97 (151)]	User or software executed drive test is complete.	No action necessary.
APM drive test completed on <Drive count> drives from array with ID <Array number>.	Information [0xB97 (151)]	User or software executed drive test on the drives in the selected array is complete.	No action necessary.
The expansion cable was connected to Enclosure <Enclosure number>.	Information [0xB98 (152)]	The user connected an expansion cable.	No action necessary.
The expansion cable was removed from Enclosure <Enclosure number>.	Information [0xB98 (152)]	The user disconnected an expansion cable.	No action necessary.
The expansion cable was removed from Enclosure <Enclosure number> or failed for more than 10 times in an hour.	Warning [0xB98 (2968)]	The user disconnected an expansion cable greater than 10 times or the cable failed for more than 10 times in an hour.	Replace cable.

### 13.6.2. Alarm Events

These events are related to the audible alarm reported by the SES processor. The ID is given in the format of its hexadecimal value which is seen in Unix operating systems and its equivalent decimal value in parenthesis which is seen in the Microsoft Windows operating system.

Table 18- Alarm Events

Alarm Event Messages	Type [ID]	Cause	Action
The audible alarm is temporarily muted.	Information [0xB8D (2957)]	The user muted the alarm. The audible alarm is temporarily muted.	No action necessary.
Alarm Disable.	Information [0xB8E (2958)]	The user disabled the alarm. The audible alarm has been disabled and will silence all alarm events. The disabled alarm will be cleared when the Alarm is enabled or the controller is reset.	No action necessary.
Alarm Enable.	Information [0xB8F (2959)]	The user disabled the alarm. The audible alarm has been enabled and will allow all alarms to be heard.	No action necessary.
Alarm <x> is Off (Muted). Where <x> is the alarm number. Currently there is only one alarm in the enclosure system, however this is a provision for future expansion.	Information [0xC6E (3182)]	No status condition being reported. Alarm silenced.	No action necessary. User pressed the Alarm Mute button on the front panel.

Alarm Event Messages	Type [ID]	Cause	Action
Alarm <x> is Intermittent.	Warning [0xC6E (3182)]	A status condition caused the alarm to sound every two minutes until muted.	Press the Alarm Mute button on the front panel and isolate the cause of the alarm.
Alarm <x> is Remind.	Warning [0xC6E (3182)]	A status condition that caused the alarm to sound is continuing to remind the user.	Press the Alarm Mute button on the front panel and isolate the cause of the alarm.
Alarm <x> is On Continuous.	Error [0xC6E (3182)]	A status condition caused the alarm to sound.	Press the Alarm Mute button on the front panel and isolate the cause of the alarm.

### 13.6.3. Controller Events

The following table provides a brief description of the events which relates to all models of the RAID Controller and the Configuration. The ID is given in the format of its hexadecimal value which is seen in Unix operating systems and its equivalent decimal value in parenthesis which is seen in the Microsoft Windows operating system.

Table 19- Controller Events

Controller Events Messages	Type [ID]	Cause	Action
There was a fatal Watchdog Error.  Additional Info: (Advanced hex data for customer service or engineering use.).	Error [0xB01 (2817)]	Internal hardware or firmware failure.	Replace the controller.
There was a fatal ECC Error.  Additional Info: (Advanced hex data for customer service or engineering use.).	Error [0xB01 (2817)]	Fault SDRAM or damaged internal bus.	Replace the controller.
There was a fatal Host Fibre Channel Interface Error on Loop <xx>.  Additional Info: (Advanced hex data for customer service or engineering use.).	Error [0xB01 (2817)]	Internal hardware or firmware failure on the coprocessor.	Replace the controller.  Contact technical support.
There was a fatal Coprocessor Error.  Additional Info: (Advanced hex data for customer service or engineering use.).	Error [0xB01 (2817)]	Internal hardware or firmware failure on the coprocessor.	Replace the controller.

Controller Events Messages	Type [ID]	Cause	Action
<p>There was a fatal Data Abort: &lt;yyyyyyyy&gt;.</p> <p>Additional Info: (Advanced hex data for customer service or engineering use.).</p>	Error [0xB01 (2817)]	This fatal controller event log error indicates an unexpected interrupt, exception, or error status internal to the policy processor or a related coprocessor.	No action necessary.
<p>There was a fatal Prefetch Abort: &lt;yyyyyyyy&gt;.</p> <p>Additional Info: (Advanced hex data for customer service or engineering use.).</p>	Error [0xB01 (2817)]	Indicates an unexpected interrupt, exception, or error status internal to the policy processor or a related coprocessor.	No action necessary.
<p>There was a fatal Software Interrupt: &lt;yyyyyyyy&gt;.</p> <p>Additional Info: (Advanced hex data for customer service or engineering use.).</p>	Error [0xB01 (2817)]	Indicates an unexpected interrupt, exception, or error status internal to the policy processor or a related coprocessor.	No action necessary.
<p>There was a fatal Bad Instruction: &lt;yyyyyyyy&gt;.</p> <p>Additional Info: (Advanced hex data for customer service or engineering use.).</p>	Error [0xB01 (2817)]	Indicates an unexpected interrupt, exception, or error status internal to the policy processor or a related coprocessor.	No action necessary.
<p>There was a fatal IRQ INTERRUPT: &lt;yyyyyyyy&gt;.</p> <p>Additional Info: (Advanced hex data for customer service or engineering use.).</p>	Error [0xB01 (2817)]	Indicates an unexpected interrupt, exception, or error status internal to the policy processor or a related coprocessor.	No action necessary.
<p>There was a fatal ATU Error: &lt;yyyyyyyy&gt;.</p> <p>Additional Info: (Advanced hex data for customer service or engineering use.).</p>	Error [0xB01 (2817)]	Indicates an unexpected interrupt, exception, or error status internal to the policy processor or a related coprocessor.	No action necessary.
<p>There was a fatal PBI Error: &lt;yyyyyyyy&gt;.</p> <p>Additional Info: (Advanced hex data for customer service or engineering use.).</p>	Error [0xB01 (2817)]	Indicates an unexpected interrupt, exception, or error status internal to the policy processor or a related coprocessor.	No action necessary.
<p>There was a fatal AAU Error: &lt;yyyyyyyy&gt;.</p> <p>Additional Info: (Advanced hex data for customer service or engineering use.).</p>	Error [0xB01 (2817)]	Indicates an unexpected interrupt, exception, or error status internal to the policy processor or a related coprocessor.	No action necessary.

Controller Events Messages	Type [ID]	Cause	Action
There was a fatal Internal Error Log Continuation: <yyyyyyyy>.  Additional Info: (Advanced hex data for customer service or engineering use.).	Error [0xB01 (2817)]	Indicates an unexpected interrupt, exception, or error status internal to the policy processor or a related co-processor.	No action necessary
There was a fatal F/W Error: <xx><yyyyyyyy>.  Additional Info: (Advanced hex data for customer service or engineering use.).	Error [0xB01 (2817)]	Indicates an unexpected interrupt, exception, or error status internal to the policy processor or a related co-processor.	No action necessary.
There was a fatal Firmware Error 0xB<xx>.  Additional Info: (Advanced hex data for customer service or engineering use.).	Error [0xB01 (2817)]	Indicates an unexpected interrupt, exception, or error status internal to the policy processor or a related co-processor.	No action necessary.
A fatal PCIe Interface Error: Link Training Error.	Error [0xB01 (2817)]	Indicates the PCIe bus of the policy processor is not fully operational.	Replace the controller.
A fatal PCIe Interface Error: Bad Link Width.	Error [0xB01 (2817)]	Indicates the protocol processor was unable to negotiate for the proper link width resulting in degraded controller performance.	Replace the controller.
There was a fatal PCIe Interface Error: Correctable Error.  Additional Info: (Advanced hex data for customer service or engineering use.).	Error [0xB01 (2817)]	Cannot clear PCIe correctable errors.	Replace the controller.
There was a fatal PCIe Interface Error: Presence Mismatch.  Additional Info: (Advanced hex data for customer service or engineering use.).	Error [0xB01 (2817)]	Indicates a connection problem between the two controllers.	Replace the controller.
There was a fatal PCIe Interface Error: Fault Detected on Other (Partner) Controller.  Additional Info: (Advanced hex data for customer service or engineering use.).	Error [0xB01 (2817)]	Indicates a connection problem between the two controllers.	Replace the controller.

Controller Events Messages	Type [ID]	Cause	Action
There was a fatal PCIe Interface Error: No Initialization Information from Other (Partner) Controller.  Additional Info: (Advanced hex data for customer service or engineering use.).	Error [0xB01 (2817)]	Indicates a connection problem between the two controllers.	Replace the controller.
There was a fatal PCIe Interface Error: Information Mismatch on Other (Partner) Controller.  Additional Info: (Advanced hex data for customer service or engineering use.).	Error [0xB01 (2817)]	Indicates a compatibility problem between the two controllers.	Replace the controller.
The controller's internal temperature <aa>C has exceeded the maximum limit. The controller will shutdown to prevent damage.	Error [0xB03 (2819)]	Blocked fan.	Check enclosure for sufficient air flow.
		Failing fan.	Check for a failed Power Supply/ Cooling module or Cooling Fan module, if found replace module.
		Elevated ambient temperature.	Check the ambient temperature of the environment, decrease the local ambient temperature.
The Temperature Sensor x from Controller y reported a reading of z degrees Celsius and it has exceeded the maximum limit (n degrees Celsius). The controller will shutdown to prevent damage.	Error [0xB03 (2819)]	Blocked fan.	Check enclosure for sufficient air flow.
		Failing fan.	Check for a failed Power Supply/ Cooling module or Cooling Fan module, if found replace module.
		Elevated ambient temperature.	Check the ambient temperature of the environment, decrease the local ambient temperature.
The controller's internal temperature <aa>C is approaching the maximum limit. You should check the cooling system for problems.	Warning [0xB04 (2820)]	Blocked fan.	Check enclosure for sufficient air flow.
		Failing fan.	Check for a failed Power Supply/ Cooling module or Cooling Fan module, if found replace module.
		Elevated ambient temperature.	Check the ambient temperature of the environment, decrease the local ambient temperature.
The Temperature Sensor x from Controller y reported a reading of z degrees Celsius and is approaching the maximum limit (n degrees Celsius). You should check the cooling system for problems.	Warning [0xB04 (2820)]	Blocked fan.	Check enclosure for sufficient air flow.
		Failing fan.	Check for a failed Power Supply/ Cooling module or Cooling Fan module, if found replace module.
		Elevated ambient temperature.	Check the ambient temperature of the environment, decrease the local ambient temperature.

Controller Events Messages	Type [ID]	Cause	Action
The partner controller has failed or has been removed.	Error [0xB08 (2824)]	Failure or removal of one controller (partner) in an Active-Active configuration.	Re-install the controller. or Replace the controller.
This controller has not received a response from the other (partner) controller in the allotted time, and therefore it has been disabled.	Error [0xB09 (2825)]	Failure or removal of one controller (partner) in an Active-Active configuration.	Replace the controller.
The controller's <x> voltage reading measures <aa>V which exceeds the limit.	Error [0xB19 (2841)]	Voltage regulator hardware failure.  Enclosure 5V or 12V problem in the power supply.	Replace the controller.  Replace the defective Power Supply or Power Supply/Cooling module.
Internal transfer error.	Error [0xB1A (2842)]	Hardware problem.	Replace the controller.
Controller mismatch detected.  "This" or "The other" controller was shut down.	Error [0xB1C (2844)]	RAID Controllers with different SAS controller chips have been installed. This is not a supported configuration.	Replace one of the controllers so that both controllers are of the same type.
The discovery process has completed identifying all SAS devices on the SAS domain.	Information [0xB22 (2850)]	A SAS discovery was completed.	No action necessary.
The discovery process has started to determine all SAS devices on the SAS domain.	Information [0xB23 (2851)]	A SAS discovery was initiated.	No action necessary.
The other (partner) controller has been inserted.	Information [0xB29 (2857)]	Partner controller has been inserted.	No action necessary.
The other (partner) controller has passed its self-test and is now ready (failback).	Information [0xB2A (2858)]	Partner controller is ready to fail back.	No action necessary.
The other controller has shut itself down, either because of a failure or user request.	Warning [0xB2B (2859)]	A controller failure.  User request to shutdown.	Replace Controller.  No action necessary.
A stripe synchronization of a RAID set has started. This occurs when a controller fails, or after a controller is powered off with write commands in progress.	Information [0xB2C (2860)]	A controller fails or is powered off during a write operation.	No action necessary.
A stripe synchronization of a RAID set has completed.	Information [0xB2D (2861)]	A controller fails or is powered off during a RAID write operation.	No action necessary.
The configuration has changed.	Information [0xB2F (2863)]	A change in the configuration has occurred.	If you are using the Save Configuration feature, re-save your configuration information - it no longer matches, otherwise no action is necessary.
The controller is flushing the partner's mirrored cache to the drives. There are <xx> cache entries totalling <yy> 512-byte blocks.	Information [0xB35 (2869)]	Failure or removal of the partner controller.	No action necessary.

Controller Events Messages	Type [ID]	Cause	Action
The controller has completed flushing the partner's mirrored cache to the drives.	Information [0xB36 (2870)]	Completion of mirrored cache flushing.	No action necessary.
Line Power Mode is now active, the battery is not required.	Information [0xB42 (2882)]	Battery charging complete.  Special feature mode. Controller can be operated with battery removed and no LED status indication occurs.	No action necessary.
The backup battery unit attached to this controller is now functioning correctly.	Information [0xB42 (2882)]	Battery OK. Charged for <x> seconds.	No action necessary.
The controller has been powered off.	Information [0xB50 (2896)]	Removal of controller or power.	No action necessary.
The controller has been powered on.	Information [0xB51 (2897)]	The controller was powered on.	No action necessary.
The controller self-test was successfully completed.	Information [0xB52 (2898)]	Self-test completion on startup.	No action necessary.
The controller self-test has failed.	Error [0xB53 (2899)]	Self-test failure on startup.	Replace the controller.
The controller's NVRAM has been reset.	Information [0xB54 (2900)]	Occurs first time after production.	No action necessary.
The controller has an invalid World Wide Name.	Error [0xB55 (2901)]	Occurs first time after production.	Contact technical support.
The Event Log has been cleared.	Information [0xB56 (2902)]	The user has cleared the event log.	No action necessary.
The controller has been reset.	Information [0xB57 (2903)]	User initiated a reset to the controller.	No action necessary.
The controller has been shutdown gracefully.	Information [0xB58 (2904)]	User initiated a controller shutdown.  The controller temperature threshold was exceeded and the controller shut itself down.	No action necessary.  Check for a failed Cooling Fan or Power Supply/Cooling module, replace as needed. Check for blocked air flow, correct as needed.  Check for high ambient temperature, reduce the environments ambient temperature.
All identified enclosures have at least two communication paths to their SES devices.	Information [0xB5B (2907)]	SES initialization.	No action necessary.
Failover started.	Information [0xB5C (2908)]	Failure or removal of the partner controller.	No action necessary.
Failover completed.	Information [0xB5D (2909)]	Completion of failover process.	No action necessary.

Controller Events Messages	Type [ID]	Cause	Action
Failback started.	Information [0xB5E (2910)]	Partner controller started failback.	No action necessary.
Failback completed.	Information [0xB5F (2911)]	Completion of failback process.	No action necessary.
The controller firmware has been upgraded to version <xxx>.	Information [0xB60 (2912)]	User upgraded the controller firmware.	No action necessary.
The controller battery backup unit is charging.	Information [0xB62 (2914)]	Battery charging started.	No action necessary.
Flushing of the battery protected cache has started.  There are <xx> cache entries totalling <yy> 512-byte blocks.	Information [0xB63 (2915)]	Failure of power with writeback cache present.	No action necessary.
Flushing of the battery protected cache has completed.	Information [0xB64 (2916)]	Completion of cache flushing.	No action necessary.
The cache data being preserved by the controller's battery was lost. There were <xx> cache entries totalling <yy> 512-byte blocks.	Error [0xB65 (2917)]	Failure of power for an extended time with writeback cache present.	Check the file system.
The cached data was lost. There were <xx> cache entries, totalling <xx> MBs.	Error [0xB65 (2917)]	Failure of power for an extended time with writeback cache present.	Check the file system.
The cached data was lost. There were <xx> MBs lost in Array <yy>.	Error [0xB65 (2917)]	Failure of power for an extended time with writeback cache present.	Check the file system.
The controller has been shutdown either locally or remotely. The controller temperature was exceeded and the controller shut itself down.	Information [0xB65 (2917)]	Internal hardware or firmware failure.	Replace the controller.
An SDRAM ECC error - bit <xx> at address <xx,xx,xx,xx.xx> has been detected and corrected.	Warning [0xB72 (2930)]	SDRAM error.	If it repeats, replace the controller.
A configuration parameter has been changed: <array name> (Array <number>) has been trusted due to a cancellation of an initialization.	Information [0xB74 (2932)]	A user cancelled an initialization.	No action necessary.
Hardware Error  Additional Info: (Advanced hex data for customer service reps or engineer use.).	Error [0xB7A (2938)]	The controller will continue to function, however the SES temperature sensing may not function properly.	Replace the controller.
Running <version_number> CEMI firmware version: Release <num>, Build <num>.	Information [0xB89 (2953)]	CEMI firmware is current.	No action necessary.

Controller Events Messages	Type [ID]	Cause	Action
CEMI firmware upgrade from <original CEMI FW version>: Release <num>, Build <num> to new version: Release <num>, Build <num> failed because there is no response from the processor.	Error [0xB89 (2953)]	Firmware failed to upgrade due to nonresponding controller.	No action necessary.
"Automatic" or "Manual" CEMI firmware upgrade from <original CEMI FW version>: Release <num>, Build <num> to new version: Release <num>, Build <num> <completed successfully>, <is underway>, or <failed>. "Automatic" or "Manual," depending on whether the upgrade attempt was automatic or manual.	Error [0xB89 (2953)]	Firmware is attempting to upgrade.	No action necessary.
Occurs when there is a remote request to restart Embedded module.  A remote request for a Embedded module reset was successful.	Information [0xB8A (2954)]	Embedded module reset was successful.	No action necessary.
Occurs when there is a remote request to restart Embedded module.  A remote request for a Embedded module reset failed.	Error [0xB8A (2954)]	Embedded module reset failed.	Try again, if no success replace RAID Controller.
Occurs when there is a remote request to restart Embedded module.  A remote request for an Embedded module reset was blocked, CEMI is updating.	Error [0xB8A (2954)]	Embedded module reset was unsuccessful because the CEMI is being currently programmed.	Retry sending the request for a reset.
Embedded module reset generated an unknown log entry.	Error [0xB8A (2954)]	Embedded module reset was unsuccessful because of unknown reasons.	Try again, if no success replace RAID Controller.

Controller Events Messages	Type [ID]	Cause	Action
<p>When the firmware queries the PMC firmware version, these events are logged:</p> <p>PMC images are up to date. CRC: &lt;16 digit hex CRC code&gt;.</p> <p>PMC firmware image of type &lt;%s&gt; &lt;%s&gt;. Where the first &lt;%s&gt; is one of: "bootrom" or "initstring" or "application" and the second &lt;%s&gt; is one of: "is up to date, current image's CRC: &lt;16 digit hex CRC code&gt;," or "was updated successfully from old image CRC: &lt;16 digit hex CRC code&gt; to new image CRC: &lt;16 digit hex CRC code&gt;."</p>	Information [0xB8B (2955)]		No action necessary.
<p>When the firmware queries the PMC firmware version, these events are logged:</p> <p>PMC firmware image of type &lt;%s&gt; &lt;%s&gt;.</p> <p>Where the first &lt;%s&gt; is one of: "bootrom" or "initstring" or "application" and the second &lt;%s&gt; is "failed to update from old image CRC: &lt;16 digit hex CRC code&gt; to new image CRC: &lt;16 digit hex CRC code&gt;."</p>	Error [0xB8B (2955)]		No action necessary.
<p>Invalid SAS Disk I/O or EBOD I/O module found in Enclosure: &lt;xxx&gt;.</p>	Error [0xB90 (2960)]	An unsupported EBOD I/O module was detected in the expansion enclosure.	Replace the subject EBOD I/O module with a correct supported module.
<p>Drive power down/up maintenance is carried out on the drive (Slot &lt;x&gt;, Enclosure &lt;y&gt;).</p>	Information [0xB91 (2961)]	Drive maintenance.	No action necessary.
<p>Drive maintenance is forcing the drive (Slot &lt;x&gt;, Enclosure &lt;y&gt;) to fail.</p>	Error [0xB91 (2961)]	Drive failure unknown.	Replace disk drive.

Controller Events Messages	Type [ID]	Cause	Action
There is a mismatch of controller cache size causing the shutting down of this controller. This controller (version: <This Controller Version>, cache size: <This Controller Cache Size>) versus partner controller (version: <Partner Controller Version>, cache size: <Partner Controller Cache Size>).	Error [0xB92 (2962)]	A mismatch was detected between Controller 0 and 1 SDRAM Memory module sizes. Both SDRAM Memory modules in a dual controller (Active-Active) configuration must be the same size.	When this unsupported configuration occurs, Controller 1 is always shut-down and Controller 0 is always kept running, regardless of either controller's memory size. With Controller 1 shutdown, remove and install the same size SDRAM Memory module as what is installed in Controller 0. After the memory module is replaced, reinsert Controller 1.
There is a mismatch of controller cache size causing the shutting down of the partner controller. This controller (version: <This Controller Version>, cache size: <This Controller Cache Size>) versus partner controller (version: <Partner Controller Version>, cache size: <Partner Controller Cache Size>).versus partner controller (version: %2X, cache size: %s).After the memory module is replaced, reinsert Controller 1.	Error [0xB92 (2962)]	A mismatch was detected between Controller 0 and 1 SDRAM Memory module sizes. Both SDRAM Memory modules in a dual controller (Active-Active) configuration must be the same size.	When this unsupported configuration occurs, Controller 1 is always shut-down and Controller 0 is always kept running, regardless of either controller's memory size. With Controller 1 shutdown, remove and install the same size SDRAM Memory module as what is installed in Controller 0.
The header of the expander SPIMEM downloaded image is invalid.	Error [0xB99 (2969)]	Corrupt file.	Obtain a new copy of the file.
The expander SPIMEM was updated successfully.	Information [0xB99 (2969)]	Firmware was updated successfully.	No action necessary.
The ID of the expander SPIMEM download image does not match.	Error [0xB99 (2969)]	Corrupt file.	Obtain a new copy of the file.
The CRC of the expander SPIMEM download image does not match.	Error [0xB99 (2969)]	Corrupt file.	Obtain a new copy of the file.
Failed to verify the expander BOOTROM flash.	Error [0xB9A (2970)]	File failed to verify.	Replace the controller.
Failed to read from the expander EEPROM I2C.	Error [0xB9B (2971)]	Internal communication error.	Replace the controller.
The expander EEPROM was updated successfully.	Information [0xB9B (2971)]	Firmware was updated successfully.	No action necessary.
The controller watchdog interrupt warning was repeated <xx> times.	Warning [0xB9C (2972)]	Internal errors are being reported.	Replace the controller.
Unable to delete any array while an array expansion background operation is in progress.	Information [0xB9E (2974)]	Indicates the user attempted to delete any array while an array expansion background operation is in progress.	Wait for the background process has completed and try the array deletion again.

Controller Events Messages	Type [ID]	Cause	Action
Unable to delete any array while an array rebuild background operation is in progress.	Information (2974)	[0xB9E] Indicates the user attempted to delete any array while an array rebuild background operation is in progress.	Wait for the background process has completed and try the array deletion again.
Unable to delete any array while an array initialization background operation is in progress.	Information (2974)	[0xB9E] Indicates the user attempted to delete any array while an array initialization background operation is in progress.	Wait for the background process has completed and try the array deletion again.
Unable to delete any array while an array parity check background operation is in progress.	Information (2974)	[0xB9E] Indicates the user attempted to delete any array while an array parity check background operation is in progress.	Wait for the background process has completed and try the array deletion again.
The data rate on Host Port <x> has been limited to <n> Gbps instead of the <n> Gbps supported by the controller.	Warning (2975)	[0xB9F] The SFP inserted in the host port does not support the data rate set for that port.	User should be aware with the current SFP the port is speed limited, otherwise replace with a higher speed SFP more compatible with the data rate set for the host port.
This controller has temporarily paused its bootup process because its NVRAM timestamp <xxx> is different from the partner controller <xxx>. The partner controller's state: <value>.	Information (2976)	[0xBA0] The controller initialization is on hold due to differences between this controller and the other controller's NVRAM timestamp.	No action necessary.
This controller continues its bootup process.	Information (2977)	[0xBA1] The controller has exited the state of paused initialization and is continuing the boot process.	No action necessary.
This controller's NVRAM timestamp is not valid.	Warning (2978)	[0xBA2] The other controller has a valid timestamp which will be used or the configuration is not valid.	No action necessary.
The other controller's NVRAM timestamp is valid.	Information (2979)	[0xBA3] The other controller's NVRAM valid timestamp will be used and will be given control.	No action necessary.
SAN LUN Mapping trial period has expired.	Information (2984)	[0xBA8] The SAN LUN Mapping licensed trial period or 15 day grace period has ended.	Purchase a valid SAN LUN Mapping license to reinstate functionality. Contact technical support.
SAN LUN Mapping trial period expires in <xx> days.	Information (2984)	[0xBA8] This is a daily reminder that the SAN LUN Mapping is in its licensed trial period or 15 day grace period and will expire in the specified number of days.	Purchase a valid SAN LUN Mapping license to reinstate functionality. Contact technical support.

### 13.6.4. Controller Port Events

These events are related to the host side Controller Port or Loop. The ID is given in the format of its hexadecimal value which is seen in Unix operating systems and its equivalent decimal value in parenthesis which is seen in the Microsoft Windows operating system.

Table 20 - Controller Port Events

Controller Port Events	Type [ID]	Cause	Action
Host Loop 0 is not initializing correctly.	Error [0xB15 (2837)]	The loop is not initializing.	Check/replace the cable. Replace the Controller.
SAS Port 0 will not initialize.	Error [0xB15 (2837)]	The port is not initializing.	Check/replace the cable. Replace the Controller.
Host Loop 1 is not initializing correctly.	Error [0xB16 (2838)]	The loop is not initializing.	Check/replace the cable. Replace the Controller.
SAS Port 1 will not initialize.	Error [0xB16 (2838)]	The port is not initializing.	Check/replace the cable. Replace the Controller.
Host Loop 0/1 acquired Loop ID <xx> because we were not able to get Loop ID <xx> (as specified in the controller settings).	Error [0xB17 (2839) [0xB18 (2840)]	Address conflict with either host adapter or other device on the same loop.	Resolve address conflict.
A LIP has occurred on Drive Loop <xx>.	Information [0xB21 (2849)]	The loop is disrupted or the host is booting.	No action necessary.
The Drive Loop (xx) is up.	Information [0xB22 (2850)]	Loop is up.	No action necessary.
The Drive Loop (xx) is down.	Information [0xB23 (2851)]	Loop is down.	Check/replace the cable.
A LIP has occurred on Host loop <xx>. Reason: <type>, The LIP was repeated <yy> times.	Information [0xB24 (2852)]	A LIP was generated so that a loop port could acquire a physical address on an arbitrated loop.  A LIP was generated by port ID: <xx> so that the loop would be reinitialized.  A LIP was generated because of a loop failure.  A LIP was generated by port ID: <xx> because of a loop failure.	No action necessary.
Host Loop <xx> is now up.	Information [0xB25 (2853)]	Loop is becoming ready.	No action necessary.
SAS Host Port <xx> is now up.	Information [0xB25 (2853)]	SAS port is becoming ready.	No action necessary.
Host Loop <xx> is down.	Information [0xB26 (2854)]	Loop is going down.	Check/replace the cable.

Controller Port Events	Type [ID]	Cause	Action
SAS Host Port <xx> is down.	Information [0xB26 (2854)]	SAS port is going down.	Check/replace the cable.
A host has accessed a Logical Drive <yy> for the first time, or for the first time following a reset or LIP. It accessed it through Host Loop <xx> (ID <zz>) with the SCSI command <check condition, busy, or task set full>.	Information [0xB2E (2862)]	First access by a particular host after a LIP or reset.  A host has accessed a Logical Drive <yy> for the first time, or for the first time following a reset. ID <zz> accessed it through Host Channel <nn> with the SCSI command 0x<zz>.	No action necessary.
Host Loop <num> has reported an error status of 0x<xx> to a particular command.	Error [0xB37 (2871)]	This may indicate a loop reset or LIP during a command, or a loop failure. Repeat Count = <count>.	No action necessary.  However, this event is usually followed by another event that indicates the real problem.
SAS Host Port <num> has reported an error <port>. It was repeated <number of times the error happened> times.	Error [0xB37 (2871)]	This may indicate a reset or LIP during a command. Repeat Count = <count>.	No action necessary.  However, this event is usually followed by another event that indicates the real problem.
Drive Loop <num> has reported an error status of <error status> to a particular command. This may indicate a loop reset or LIP during a command, or a loop failure.  Repeat Count = <number of times the error happened>.	Error [0xB38 (2872)]	The drive loop encountered a loop reset or LIP during a command or a loop failure has occurred.	No action necessary.
Host Loop <num> has reported an invalid status of 0x<xx> to a particular command.problem.	Error [0xB39 (2873)]	This indicates a firmware error in the host fibre channel chip.	Contact technical support.
The controller has generated a LIP on Drive Loop <xx>, due to a loop error.	Error [0xB3C (2876)]	Controller initiated a LIP on Drive Loop 1 due to a loop error. Could be caused by a disk drive being pulled or any drive side loop LIP.	No action necessary.  However, this event is usually followed by another event that indicates the real
The controller has generated a LIP on Host Loop <xx>, due to a loop error.	Error [0xB3D (2877)]	Controller initiated a LIP.	No action necessary.

Controller Port Events	Type [ID]	Cause	Action
The host system w/ WWN: <xxxxxxxxxxxxxxxx> and Loop ID of <xx> has logged into the controller through Host Loop <xx>.  These events will only be listed for HBAs with SAN LUN Mappings.	Information [0xB3F (2879)]	Host systems logs into the controller.	No action necessary.
SAS Host System <xxxxxxxxxxxxxxxx> has logged into <y>. (ID: <z>).	Information [0xB7C (2940)]	Host system has logged into the controller port <y>.	No action necessary.
SAS Host System <xxxxxxxxxxxxxxxx> has logged out of <y>. (ID: <z>).	Information [0xB7D (2941)]	Host system has logged out of the control- ler port <y>.	No action necessary.

### 13.6.5. Drive and Array Events

These events are related to the drives, loops (where applicable) and disk arrays. The ID is given in the format of its hexadecimal value which is seen in Unix operating systems and its equivalent decimal value in parenthesis which is seen in the Microsoft Windows operating system.

**NOTE:** On SATA disk drives the event message displays the drive's Serial Number.

Table 21 - Drive and Array Events

Drive and Array Events	Type [ID]	Cause	Action
The drive w/SN <xxxxxx> (Slot <nn>, Enclosure <nn>) (<Array Name> Drive <member index>) has failed due to an unrecoverable error. Sense Data: <xx>/<xx>/<xx>.	Error [0xB0A (2826)]	Typically due to a nonrecoverable media error or hardware error.	Replace the disk drive.
The drive w/SN <xxxxxx> (Slot <number>) (Drive <number>) has been marked as failed because it was removed.	Error [0xB0B (2827)]	Drive has been removed or bypassed by the user, or has a serious hardware error.  Removal of cables connecting the enclo- sures.	Replace the disk drive.  Replace the cables.

Drive and Array Events	Type [ID]	Cause	Action
Rebuilding has failed due to an unrecoverable error on the new drive w/SN <xxxxxxx> (Slot <nn> (Drive <number>) in the array.	Error [0xB0C (2828)]	Typically due to a nonrecoverable media error, or hardware error.	Replace with a new drive and initiate a rebuild.
Rebuilding has failed due to an unrecoverable error on the new drive w/SN <xxxxxxx> (Slot <nn>, Enclosure <nn>) (<Array Name> Drive <num>).	Error [0xB0D (2829)]	Typically due to a nonrecoverable media error or hardware error.	Backup all data and restore to a new array.
The drive w/SN <xxxxxxx> (Slot <nn>, Enclosure <nn>) (Slot <number>) (Drive <number>) has failed due to a time-out.	Error [0xB0E (2830)]	Drive error.	Replace the disk drive.
Drive Loop 1 is not initializing correctly.	Error [0xB10 (2832)]	Loop error.	Check the cables.
Disabled Enclosure <name> Slot <nn> due to excessive errors.	Error [0xB13 (2835)]	This indicates that the controller has shut-down the slot due to multiple errors from the drive.	Remove the drive in the identified slot will reenable the PHY.
Drive Loop 1 has exceeded the allowable error count. The controller will not use this loop for data transfers. After two hours have elapsed, the loop will be re-enabled.	Error [0xB14 (2836)]	Too many errors on the loop.	Wait two hours and try again.
Array <name> is in a critical state.	Error [0xB1B (2843)]	Drive removal or failure.	Replace the disk drive and rebuild the array.
The drive w/SN <xxxxxxx> (Slot <number>) returned a bad status while completing a command. SCSI Info:  Operation <type>, Status <type>.	Error [0xB27 (2855)]	Unknown status returned by the disk drive.	Contact technical support and provide them with a copy of the event log.
The drive w/SN <xxxxxxx> (Slot <nn>, Enclosure <nn>) timed out for the SCSI Operation <type>.	Warning [0xB28 (2856)]	Drive hardware error or bus error.	Check cabling and ensure the disk drives are properly seated.
A rebuild has started on the drive w/SN <xxxxxxx> (Slot <nn> (Drive <number>).	Information [0xB30 (2864)]	A rebuild has started.	No action necessary.

Drive and Array Events	Type [ID]	Cause	Action
A rebuild has completed on (Array <Name> Drive <number>).	Information [0xB31 (2865)]	A rebuild has completed.	No action necessary.
A rebuild has re-started on the drive w/SN <xxxxxx> (Slot <nn> (Drive <number>)).	Information [0xB32 (2866)]	A rebuild has started.	No action necessary.
Array <name> has started initializing.	Information [0xB33 (2867)]	Initialization has started.	No action necessary.
Array <name> has completed initializing.	Information [0xB34 (2868)]	Initialization has completed.	No action necessary.
The controller has detected a data underrun from the drive w/SN <xxxxxx> (Slot <nn>, Enclosure <nn>) for the SCSI Op Code 0x<xx>. This is caused by the controller detecting a bad CRC in a frame, and usually indicates a link problem, either with cabling or an enclosure.	Error [0xB3B (2875)]	Bus error.	Check cabling and ensure that the disk drive is properly seated in its slot.
An unrecoverable drive error has occurred as a result of a command being issued. This may be due to a drive error in a non-fault tolerant array, such as RAID 0, or when the array is already in a degraded mode. The controller will pass the status from the drive back to the host system, to allow the host recovery mechanisms to be used. Details: Host Loop <x>, Host Loop ID <y>, Mapped LUN Requested <z>, Op Code <zz>, Sense Data <uu>.	Error [0xB40 (2880)]	Typically due to a nonrecoverable media error, hardware error, or loop (bus) error.	No action necessary.
A RAID parity check has started on <Array Name>. Type of parity check = <paritytype>.	Information [0xB43 (2883)]	Parity check started.	No action necessary.

Drive and Array Events	Type [ID]	Cause	Action
A RAID parity check has completed on <Array Name>. Type of parity check = <paritytype>. Error Count = <zz>.	Information [0xB44 (2884)]	Parity check completed.	No action necessary.
A RAID parity check has been aborted on <Array Name>.  Type of parity check = <paritytype>. Error Count = <zz>.	Information [0xB45 (2885)]	Parity check canceled by the user.	No action necessary.
A drive w/SN (Slot <nn>, Enclosure <nn>) has been inserted.	Information [0xB61 (2913)]	Drive was inserted.	No action necessary.
The controller has started updating a drive's firmware. Drive w/SN <xxxxxx> (Slot <nn> ID: <zz> Firmware Version: <yy.yy.yyyy>).	Information [0xB66 (2918)]	The controller has started updating a drive's firmware. Drive <w/sn <xxxxxx> or w/wwn <xxxxxxxxxxxxxxxx>, Slot <nn>, Enclosure <nn>, Firmware Version: <xxxx>.	No action necessary.
The controller has finished updating a drive's firmware. Drive SN: <xxxxxx> ID: <zz> (Slot <number>) Firmware Version: <yy.yy.yyyy>.	Information [0xB67 (2919)]	The controller has finished updating a drive's firmware. Drive w/SN <xxxxxx> or WWN <xxxxxxxxxxxxxxxx>, Slot <nn>, Enclosure <nn>, Firmware Version: <xxxx>.	No action necessary.
An array expansion has started on Array <name>.	Information [0xB68 (2920)]	Expansion has started.	No action necessary.
An array expansion has completed on Array <name>.	Information [0xB69 (2921)]	Expansion has completed.	No action necessary.
An array expansion has restarted on Array <name>.	Information [0xB6A (2922)]	Expansion has restarted.	No action necessary.
The writeback cache on Array <name> has been disabled.  Reason(s): (See reasons).	Warning [0xB6F (2927)]	Disabling of writeback cache for the indicated reasons: <ul style="list-style-type: none"> <li>◆ The partner controller has failed.</li> <li>◆ The battery is not charging or present.</li> <li>◆ The array has become critical.</li> <li>◆ A "prepare for shutdown" was received by the controller.</li> </ul>	<ul style="list-style-type: none"> <li>◆ Replace the failed controller.</li> <li>◆ Charge the backup battery or re-install the battery.</li> <li>◆ Resolve the array issue and rebuild the array.</li> <li>◆ No action necessary.</li> </ul>

Drive and Array Events	Type [ID]	Cause	Action
The writeback cache on Array <name> has been re-enabled.	Information [0xB70 (2928)]	Re-enabling of writeback cache.	No action necessary.
Because of a background verify failure, data blocks at LBA <yyy> from drive SN: <xxxxxx> (Slot <number>) (Consecutive Number of Allocations <zz>) have been reallocated.	Warning [0xB71 (2929)]	Data Blocks at this location have media errors. The data is being reallocated to different blocks.	No action necessary.
A rebuild was aborted on (Array <yy> Drive <ww>).	Information [0xB73 (2931)]	A rebuild was canceled by the user.	No action necessary.
SATA Drive Error: (Slot <number>) Information <description>.	Error [0xB75 (2933)]	Drive or SATA link error.	No action necessary.
A drive w/ SN <xxxxxx> (Slot <nn>) has been removed.	Error [0xB76 (2934)]	A drive w/ SN <xxxxxx> (Slot <nn>, Enclosure <nn>) has been removed.	No action necessary.
There was a bad block during a rebuild on Array <nn>, Drive <mm>, LBA <xx xx xx xx xx>, Block Count <xx>.	Error [0xB78 (2936)]	A bad block was detected during the rebuild operation. Data loss will occur with that data stripe.	Replace drive after rebuild. Restore lost data from a known good backup.
An unsupported drive w/ SN <xxxxxx> (Target ID <num>) has been inserted.	Error [0xB7B (2939)]	A SATA drive with an A/A MUX Transition card was installed in a system with an older model controller.	Replace the SATA drive with one that has a A/P MUX Transition card installed.
<Array Name> is in an unsupported configuration with a mixture of A/P and A/A SATA drives.	Error [0xB8C (2956)]	The array is marked invalid because it contains SATA drives with a mixture of Active-Active and Active-Passive MUX Transition cards.  A SATA drive with an Active-Active MUX Transition card is not supported in older 12 drive enclosures.	Replace the offending drive with one that contains a matching MUX Transition card as the remaining drive members.
<Array Name> is in an unsupported A/P SATA drive mutiplexer configuration.	Error [0xB8C (2956)]	The array is marked invalid because it contains a SATA drive with an unsupported Active-Passive MUX Transition card.	Replace the offending drive with one that contains a matching MUX Transition card.
<Array Name> is in an unsupported A/A SATA drive mutiplexer configuration.	Error [0xB8C (2956)]	The array is marked invalid because it contains a SATA drive with an unsupported Active-Active MUX Transition card.	Replace the offending drive with one that contains a matching MUX Transition card.

### 13.6.6. Enclosure Events

These events are related to the enclosure components reported by the SES processor. The ID is given in the format of its hexadecimal value which is seen in Unix operating systems and its equivalent decimal value in parenthesis which is seen in the Microsoft Windows operating system.

Table 23 - Enclosure Events

Enclosure Events	Type [ID]	Cause	Action
Enclosure<xx> (w/ WWN:<xxxxxxxxxxxxxxxx>) timed out on SCSI com- mand 0x02X.  <b>NOTE:</b> This event is only valid for the expansion en- closure or daisy-chained systems.	Warning [0xB79 (2937)]	This error is generated when a command to the Enclosure processor time-outs.  Faulty cable or drive malfunctions could be the cause of this error.	Verify if the system has gone to SES_LEVEL_1. If it has, verify the configuration.  You may occasionally see this error during drive inser- tion, failover/failback or drive removal. As long as the system remains at SES_LEVEL_3 the user does not need to inter- vene. If this event is periodically posted in the event log the user may have a hard drive or EBOD I/O module problem. The system should be inspected to isolate the problem to either drives or EBOD I/O module.
Power supply <zz> is OK.	Information [0xB6B (2923)]	Normal condition reported.	No action necessary.
Power supply <zz> is op- erating outside of its speci- fication.	Warning [0xB6B (2923)]	The specific power supply has failed. The specific power supply has the mains power removed.	Replace the power supply mod- ule.  Ensure that the specific power supply mains power is applied.
Power supply <zz> is in a critical state.	Warning [0xB6B (2923)]	The specific power supply has failed.	Replace the power supply.
		The specific power supply has the mains power removed.	Ensure that the specific power supply mains power is applied.
		The power supply was removed.	Insert the power supply.
Power supply <zz> is not installed.	Error [0xB6B (2923)]	The power supply was removed.	Re-insert the power supply, con- nect the power cord.
An Expansion enclosure was connected to Enclosure ax.  The enclosure count in- creased from yy to zz.	Information [0xB98 (2968)]	Expansion cable was connected.	No action necessary.
An Expansion enclosure was removed from Enclosure xx.  The enclosure count de- crease from yy to zz.	Warning [0xB98 (2968)]	Expansion cable was disconnected.	Re-connect the SAS patch cable.

Enclosure Events	Type [ID]	Cause	Action
The expansion cable was removed from Enclosure <xz> or failed for more than 10 times in an hour.	Warning [0xB99 (2969)]	The user disconnected an expansion cable greater than 10 times or the cable failed for more than 10 times in an hour.	Replace cable.
The expansion cable was removed from Enclosure <xx> or failed for more than 10 times in a hour.	Information [0xB99 (2969)]	The firmware image was not corrupt.	No action necessary.
The expander SPIMEM image is valid. Current version: xxxxx, new version: xxxxx.	Information [0xB99 (2969)]	Firmware was updated successfully.	No action necessary.
The expander SPIMEM was updated successfully. Current version: xxxxx, new version: xxxxx.	Error [0xB99 (2969)]	Corrupt file.	Obtain a new copy of the file.
The header of the expander SPIMEM download image is invalid. Current version: xxxxx, new version: xxxxx.	Error [0xB99 (2969)]	Corrupt file.	Obtain a new copy of the file.
The ID of the expander SPIMEM download image does not match. Current version: xxxxx, new version xxxxx.	Error [0xB99 (2969)]	Corrupt file.	Obtain a new copy of the file.
The CRC of the expander SPIMEM download image does not match. Current version: xxxxx, new version: xxxxx.	Error [0xB99 (2969)]	Failed to update.	Obtain a new copy of the file.
Failed to read from the expander SPIMEM. Current version: xxxxxx, new version xxxxx. (RS-1220-F4-6412E Only)	Error [0xB99 (2969)]	Failed to update.	Replace the RAID Controller.
Failed to erase the expander SPIMEM. Current version: xxxxx, new version: xxxxx. (RS-1220-F4-6412E Only)	Error [0xB99 (2969)]	Failed to update.	Replace the RAID Controller.
Failed to program the expander SPIMEM page. Current version: xxxxx, new version: xxxxx.  (RS-1220-F4-6412E Only)	Error [0xB99 (2969)]	Failed to update.	Replace the RAID Controller.

Enclosure Events	Type [ID]	Cause	Action
Failed to verify the expander SPIMEM image. Current version: xxxxx, new version: xxxxx.  (RS-1220-F4-6412E Only)	Error [0xB99 (2969)]	Failed to update.	Replace the RAID Controller.
Failed to verify the expander SPIMEM image trailer. Current version: xxxxx, new version: xxxxx.  (RS-1220-F4-6412E Only)	Error [0xB99 (2969)]	Failed to update.	Replace the RAID Controller.
The status returned by SAS-SATA Channel Card z from Enclosure xx: Unsupported. SAS Address: <SN>.  (RS-4835-F4-5402E Only)	Information [0xC4F (3151)]	-	-
The status returned by SAS-SATA Channel Card z from Enclosure xx: OK. SAS Address: <SN>. (RS-4835-F4-5402E Only)	Information [0xC4F (3151)]	Normal condition reported.	No action necessary.
The status returned by SAS-SATA Channel Card z from Enclosure xx: Critical. SAS Address: <SN>.  (RS-4835-F4-5402E Only)	Warning [0xC4F (3151)]	The specified SASSATA Channel Card has been removed.	Install the SAS-SATA Channel Card.
The status returned by SAS-SATA Channel Card z from Enclosure xx: Non-critical. SAS Address: <SN>.  (RS-4835-F4-5402E Only)	Information [0xC4F (3151)]	-	-
The status returned by SAS-SATA Channel Card z from Enclosure xx: Unrecoverable. SAS Address: <SN>.  (RS-4835-F4-5402E Only)	Error [0xC4F (3151)]	-	-
The status returned by SAS-SATA Channel Card z from Enclosure xx: Not Installed. SAS Address: <SN>.  (RS-4835-F4-5402E Only)	Error [0xC4F (3151)]	The specified SASSATA Channel Card was not installed at power up.	Install the SAS-SATA Channel Card.

Enclosure Events	Type [ID]	Cause	Action
The status returned by SAS-SATA Channel Card z from Enclosure xx: Unknown. SAS Address: <SN>. (RS-4835-F4-5402E Only)	Error [0xC4F (3151)]	The specified SASSATA Channel Card has detected an unknown error.	Replace the SAS-SATA Channel Card.
The status returned by SAS-SATA Channel Card z from Enclosure xx: Not Available. SAS Address: <SN>. (RS-4835-F4-5402E Only)	Error [0xC4F (3151)]	-	-
Power supply <zz> is in an unrecoverable state.	Error [0xC6B (3179)]	The power supply was removed.	Re-insert the power supply, connect the power cord.
Power supply <zz> is not installed.	Error [0xC6B (3179)]	The power supply was removed.	Re-insert the power supply, connect the power cord.
Fan <zz> is OK.	Information [0xC6C (3180)]	Normal condition reported.	No action necessary.
Fan <zz> does not support status detection.	Information [0xC6C (3180)]	Unknown status.	No action necessary.
Fan <zz> is in a critical state. Additional information may appear:  The fan is stopped.  The fan is running at its lowest speed.  The fan is running at its second lowest speed.  The fan is running at its third speed level.  The fan is running at its fourth speed level.  The fan is running at its fifth speed level.  The fan is running at its intermediate speed.  The fan is running at its highest speed.	Error [0xC6C (3180)]	A specific fan failure.	Replace the cooling fan module.
Temperature sensor <zz> is OK.	Information [0xC6D (3181)]	Temperature sensors are reporting normal temperatures in the enclosure.	No action required.

Enclosure Events	Type [ID]	Cause	Action
Temperature <zz> is operating outside of specifications.	Warning [0xC6D (3181)]	Temperature sensors are reporting enclosure temperatures have reached the threshold of 50°C (122°F).	<p>Ensure that both cooling fans are operating normally. (Replace if needed.)</p> <p>If the environment ambient temperature is high, reduce the ambient temperature.</p> <p>Ensure that the airflow is not blocked or restricted on the enclosure.</p>
Temperature sensor <zz> is in a critical state.	Error [0xC6D (3181)]	Temperature sensors are reporting enclosure temperatures have reached the threshold of 70°C (158°F).	<p>Automatic system shutdown will begin. In Active-Active controller configurations, one controller will shutdown its partner and shutdown the drives, then itself.</p> <p>Ensure that the cooling fans are operating normally. (Replace if needed.)</p> <p>If the environment ambient temperature is high, reduce the ambient temperature.</p> <p>Ensure that the airflow is not blocked or restricted on the enclosure.</p>

### 13.6.7. Persistent Reservation Events

These events are related to persistent reservation. The ID is given in the format of its hexadecimal value which is seen in Unix operating systems and its equivalent decimal value in parenthesis which is seen in the Microsoft Windows operating system.

Table 24 - Persistent Reservation Events

Persistent Reservation Events	Type [ID]	Cause	Action
Incorrect Reservation Key. (Controller: <num>, Port <num>, Logical Drive <num>).	Information [0xB87 (2951)]	A Persistent Reservation issue has occurred.	No action necessary.
Incorrect Reservation Type. (Controller: <num>, Port <num>, Logical Drive <num>).	Information [0xB87 (2951)]	A Persistent Reservation issue has occurred.	No action necessary.

Persistent Reservation Events	Type [ID]	Cause	Action
I_T Nexus is not Reserved. (Controller: <num>, Port <num>, Logical Drive <num>).	Information [0xB87 (2951)]	A Persistent Reservation issue has occurred.	No action necessary.
I_T Nexus is not Registered. (Controller: <num>, Port <num>, Logical Drive <num>).	Information [0xB87 (2951)]	A Persistent Reservation issue has occurred.	No action necessary.
Reservation Conflict Detected. (Controller: <num>, Port <num>, Logical Drive <num>).	Information [0xB87 (2951)]	A Persistent Reservation issue has occurred.	No action necessary.
Persistent Reservation Check Condition Detected. (Controller: <num>, Port <num>, Logical Drive <num>).	Information [0xB87 (2951)]	A Persistent Reservation issue has occurred.	No action necessary.
Invalid Service Action Reservation Key.  (Controller: <num>, Port<num>, Logical Drive <num>).	Information [0xB87 (2951)]	A Persistent Reservation issue has occurred.	No action necessary.

### 13.6.8. Server Events

**NOTE:** These events are applicable to the host-based version only.

The following table provides a brief description of the events which relate to the server software component. The ID is given in the format of its hexadecimal value which is seen in Unix operating systems and its equivalent decimal value in parenthesis which is seen in the Microsoft Windows operating system.

Table 25 - Server Events

Server Events	Type [ID]	Cause	Action
The server has been started.	Information [0x101 (257)]	The server started successfully.	No action necessary.
A user (admin) has logged into the Server from Host: <name>, IP address: xx.xx.xx.xx.	Information [0x101 (257)]	The specified user has logged in.	No action necessary.

Server Events	Type [ID]	Cause	Action
The server has been shut-down.	Information [0x102 (258)]	The host is shutting down or a user stopped the server service.	No action necessary.
A system rescan has been initiated.	Information [0x103 (259)]	The system rescan is starting note: a shut-down and start up event will follow.	No action necessary.
The serial number and key entered are incorrect. Remote features are temporarily disabled.	Error [0x104 (260)]	The serial number and key specified in the server settings file is not a correct match.	Open the software in a browser from the local console, click the Settings button. Re-enter the serial number and key. If you still have problems, contact technical support.
All event logs cleared.	Information [0x105 (261)]	A user cleared the event logs.	No action required.
A user <name> (IP:<IP_address>, Host: <host name>) has updated their password.	Information [0x106 (262)]	A specific user has updated their password.	No action required.
A user <name> (IP:<IP_address>, Host: <host name>) unsuccessfully attempted to update their password.	Information [0x107 (263)]	A specific user tried to change the password but verification failed or old password failed.	Ensure the user is authorized or needs assistance.
A user <name> has logged into the server from IP address: <IP_address>.	Information [0x108 (264)]	A specific user has logged into the server from the specified IP address location. The host name was not available.	No action required.
A user (<name>) has logged into the Server from Host: <host name>, IP address:<IP_address>.	Information [0x108 (264)]	A specific user has updated their password.	No action required.
A user <name> has been logged out of the server from IP address: <IP_address>.	Information [0x109 (266)]	A specific user has logged out.	No action required.
A user <name> (IP: <IP_address>,;) has been logged out of the server from Host: <host name>, IP address: <IP_address>.	Information [0x109 (266)]	A specific user has logged out.	No action required.
The multicast address is not configured. No communication will take place with other servers.	Warning [0x201 (513)]	The setting "MulticastPort" is not setup in the db/IP.db file.  The setting was removed or the file became corrupt.	Open the file <software folder>\db\IP.db in a text editor, and add the following line: MulticastPort=9191.  Save the file and start the software and press the RESCAN button or restart the service.

Server Events	Type [ID]	Cause	Action
Multicast socket creation failure: <reason>	Warning [0x202 (514)]	Could not setup the necessary communication paths to talk to other servers. The specific cause will be specified in the <reason> appended to the message.	<p>Change the multicast port used. Note that all servers must communicate with each other on the same multicast port. If the setting is changed on one, they must be changed on all servers.</p> <p>Open the file &lt;software folder&gt;\db\IP.db in a text editor, and change the following line: MulticastPort=&lt;port number&gt;.</p> <p>Save the file and start the software and press the RESCAN button or restart the service.</p>
Shared memory used for the server and the GUI to communicate could not be created nor located.	Error [0x301 (769)]	Tried to create/locate the shared memory used for the server and server CGI script to communicate.	Quit and all connections accessing this server, wait 1 minute and restart them.
A CGI request was initiated from the GUI but the associated shared memory could not be opened.	Error [0x302 (770)]	The server script creates new shared memory for the server to access parameter passed from the GUI. This shared memory could not be found.	<p>Retry request.</p> <p>Quit and all connections accessing this server, wait 1 minute and restart them.</p> <p>Restart the system.</p>
A CGI request was initiated from the GUI but the associated parameters could not be located.	Error [0x303 (771)]	The server script creates new parameter for the server to access parameter passed from the GUI. This parameters could not be found.	<p>Retry request.</p> <p>Quit and all connections accessing this server, wait 1 minute and restart them.</p> <p>Restart the system.</p>
The server performed a CGI request but an internal error prevented the server from returning the results.	Error [0x304 (772)]	The request was performed but the server encountered an error that prevented completion.	<p>Retry request.</p> <p>Quit and all connections accessing this server, wait 1 minute and restart them.</p> <p>Restart the system.</p>
The server performed a CGI request but the shared memory needed to return the results could not be created.	Error [0x305 (773)]	The server completed the request and attempted to create shared memory to send the results back to the server CGI. The attempt failed.	<p>Retry request.</p> <p>Quit and all connections accessing this server, wait 1 minute and restart them.</p> <p>Restart the system.</p>

Server Events	Type [ID]	Cause	Action
The server failed to complete a CGI request before its allowed time expired.	Error [0x306 (774)]	The server completed the request but by the time it was done the server CGI gave up.	<p>Retry request.</p> <p>Quit and all connections accessing this server, wait 1 minute and restart them.</p> <p>Restart the system.</p>
The server CGI script was unable to return CGI request results to the GUI.	Error [0x307 (775)]	The server passed the completed request to the server CGI but for some reason the request wasn't completed correctly by the server CGI script.	<p>Retry request.</p> <p>Quit and all connections accessing this server, wait 1 minute and restart them.</p> <p>Restart the system.</p>
Start up is complete.	Information [0xA01 (2561)]	The RAID controller module has been loaded by the server.	No action necessary.
Several failures encountered while trying to communicate with the RAID controller.	Error [0xA03 (2563)]	Several commands have been sent through a known good path but they have failed.	Ensure that the paths and hardware are operational.
Successful communication with controller after several failures.	Information [0xA04 (2564)]	A known failed path is now functional and passed retest.	No action necessary.

### 13.6.9. SnapShot Events

The following table provides a brief description of the events which relate to the Snapshot component. The ID is given in the format of its hexadecimal value which is seen in Unix operating systems and its equivalent decimal value in parenthesis which is seen in the Microsoft Windows operating system.

Table 26 - SnapShot Events

SnapShot Events	Type [ID]	Cause	Action
Logical Drive <ID #>'s Snapshot capacity is at <x>%.	Warning [0xB81 (2945)]	<p>ODA capacity has reached the indicated threshold.</p> <p>The capacity value displayed is the current unused/remaining space at the time of the event.</p> <p>The amount of space varies based on the ODA size. For ODA sizes of 1- 50 GB, the values will be 40, 30, 20, 10. For ODA sizes 50-500 GB the values will be 15, 10, 5, 2. And for ODA sizes greater than 500 GB, the values will be 4, 3, 2, 1.</p>	Delete some snapshots for this snapshotted logical drive.

SnapShot Events	Type [ID]	Cause	Action
Logical Drive <ID #>'s ODA is invalid.	Error [0xB81 (2945)]	ODA is no longer valid because of a RAID or snapshot error.	<ul style="list-style-type: none"> <li>◆ Restore from backup.</li> <li>◆ Establish a larger storage area for this snapshotted logical drive.</li> </ul>
Logical Drive <ID #>'s ODA has been overrun.	Error [0xB81 (2945)]	<p>ODA is no longer valid.</p> <p>Capacity was exceeded during last snapshot.</p>	<ul style="list-style-type: none"> <li>◆ Restore from backup.</li> <li>◆ Establish a larger storage area for this snapshotted logical drive.</li> </ul>
Logical Drive <ID #>'s had an ODA unassigned from it.	Information [0xB82 (2946)]	ODA was unassigned.	No action necessary.
Logical Drive <ID #>'s had an ODA assigned from it.	Information [0xB82 (2946)]	ODA was assigned.	No action necessary.
Snapshot <#> on Logical Drive <ID #> is being deleted.	Information [0xB84 (2948)]	Snapshot deletion has started.	No action necessary.
Snapshot <#> on Logical Drive <ID #> has been deleted.	Information [0xB84 (2948)]	Snapshot was deleted.	No action necessary.
Snapshot <#> on Logical Drive <ID #> has been created.	Information [0xB84 (2948)]	Snapshot was created.	No action necessary.
A Snapback (based on Snapshot <x>) on Logical Drive <ID #> has been started <x,y>.	Information [0xB84 (2948)]	A snapback operation was started by the operator.	No action necessary.
A Snapback (based on Snapshot <x>) on Logical Drive <ID #> has completed <x,y>.	Information [0xB84 (2948)]	A snapback operation has successfully completed.	No action necessary.
Snapshot <x> had an error on Logical Drive <ID #> (ODA #, LUN #, Error=<error code>).	Error [0xB85 (2949)]	There is a problem with the logical drive being snapshotted.	Inspect the logical drives and arrays.
License has been modified - the new level is <#>.	Information [0xB86 (2950)]	A license update/upgrade was performed.	No action necessary.

SnapShot Events	Type [ID]	Cause	Action
Snapshot metadata was left in the controller cache and the power was removed and the battery was drained before the cache was flushed. A flush occurs every two seconds. Metadata is no longer valid. Fatal loss of snapshots will be the result.	Error [0xB88 (2952)]	Snapshot metadata was in the controller cache when the power was removed before the cache flushed. Followed by the backup battery was drained with a loss of the metadata.	Loss of snapshot, perform a new snapshot and ensure the power remains constant.
There is a mismatch of snapshot versions.	Error [0xB88 (2952)]	Snapshot metadata has encountered a version mismatch. There is a incompatible version of StorView and/ or RAID Controller firmware installed.	Update the controller firmware to at least version 3.04.xxxx or later, and the StorView software to at least version 3.09.xxxx or later.
Metadata in the cache is invalid or out of date.	Warning [0xB88 (2952)]	A controller was inserted that contained old snapshot metadata in its cache.	No action necessary.
Snapshot metadata is corrupted.	Error [0xB88 (2952)]	A snapshot metadata is corrupted and cannot be used.	No action necessary.
Snapshot metadata recovery was started.	Information [0xB88 (2952)]	A snapshot metadata recovery from cache was automatically started by the controller.	No action necessary.
Snapshot metadata recovery from cache just completed.	Information [0xB88 (2952)]	A snapshot metadata recovery from cache was completed.	No action necessary.

## 13.7. Failed Drives Codes

The controller maintains a list of failed drives. Drives are listed in the following format:

Failed Drive:xx SN: <xxxxxxx>

Reason Code

The reason code may be one of the following:

Table 27 - Failed Drives Codes

Failed Drives Codes	Reason	Action
Drive Time-out	The drive has either timed out or been removed.	Re-insert the disk drive.  or Replace the disk drive.
Command: xx Sense Key: yy Ext Sense: zz	The drive has failed for the specified command, with the indicated SCSI sense key and extended sense key.	Replace the disk drive.

## 13.8. SAS Interface Error Codes

The controller maintains a list of internal SAS Interface error codes. For these error codes the RAID Controller will automatically recover either by a reset or re initialization.

Table 28 - Error Code Message

Error Code Message
There was an Inter-Controller Link Error.
There was an Inter-Controller Link Initialization Error.
There was a SAS Discovery Timeout.
There was a SAS Discovery Error.
There was a SAS Controller Reinitialization Error.
There was a SAS Interface Error xxxxxxxx *

\* This is a generic undecoded error directly from the SAS interface chip where the xxxxxxxx is an unknown hex value.

## 13.9. SAS Discovery Error Codes

The controller maintains a list of internal SAS Discovery error codes. There is no action necessary with these error codes.

Table 29 - SAS Discovery Error Codes

Error Code Message
An error has been detected on port 0/1 of the SAS domain or with a SAS device during the discovery process: Max SATA Targets.
An error has been detected on port 0/1 of the SAS domain or with a SAS device during the discovery process: Multiple Paths.
An error has been detected on port 0/1 of the SAS domain or with a SAS device during the discovery process: Table to Table.
An error has been detected on port 0/1 of the SAS domain or with a SAS device during the discovery process: Multiple Subtractive.
An error has been detected on port 0/1 of the SAS domain or with a SAS device during the discovery process: SMP CRC Error.
An error has been detected on port 0/1 of the SAS domain or with a SAS device during the discovery process: SMP Function Failed.
An error has been detected on port 0/1 of the SAS domain or with a SAS device during the discovery process: Route Table Index Missing.
An error has been detected on port 0/1 of the SAS domain or with a SAS device during the discovery process: Expander Out of Entries.
An error has been detected on port 0/1 of the SAS domain or with a SAS device during the discovery process: SMP Timeout.
An error has been detected on port 0/1 of the SAS domain or with a SAS device during the discovery process: Expander Error.
An error has been detected on port 0/1 of the SAS domain or with a SAS device during the discovery process: Ports With Same Address.
An error has been detected on port 0/1 of the SAS domain or with a SAS device during the discovery process: Unaddressable Device Found.
An error has been detected on port 0/1 of the SAS domain or with a SAS device during the discovery process: Loop Detected.

---

**NOTE:** A single event may contain multiple errors, in those situations the error is appended to the base message separated by a semicolon.

---

Example: An error has been detected on port 0 of the SAS domain or with a SAS device during the discovery process: Max SATA Targets; SMP Function Failed; Loop Detected.

---

## 14. Statistics

### 14.1. Overview

StorView and the RAID Controller will monitor all incoming commands and calculate various statistics.

The statistics monitored include:

- ◆ Command Count
- ◆ Command Alignment
- ◆ Command Size
- ◆ Read-Ahead Statistics
- ◆ Write Clustering Statistics
- ◆ RAID 5/50/6 Write Statistics

From the Main screen click the Logical Drive Statistics icon in the Tool Bar.

The controller maintains individual access statistics for all logical drives, controllers, and individual or all ports. These can be used to help balance the load from the host.

You may also export the statistical data to a comma delimited file for use in third-party software products.

---

## 14.2. Access Statistics

These statistics are for both reads and writes, and can be used to tune the operating system for optimum performance.

Every time statistics are viewed, the controller first outputs the time since the statistics were last reset.

However, the statistics can be cleared at any time. This is useful in determining the access pattern for a particular test or period of time.

Table 30 - Statistics Descriptions

Statistic	Description
Reads	This is the average number of MBs transferred in the last few seconds from the logical drives, controllers or ports. This value is expressed in MB/seconds.
Writes	This is the average number of MBs transferred in the last few seconds to the logical drives, controllers or ports. This value is expressed in MB/seconds.
No. of Operations	This is the total number of read and write accesses that have occurred since these statistics were reset, or the controller was last powered on.
Bytes Transferred	This is the total number of bytes read and written since these statistics were reset, or the controller was last powered on.

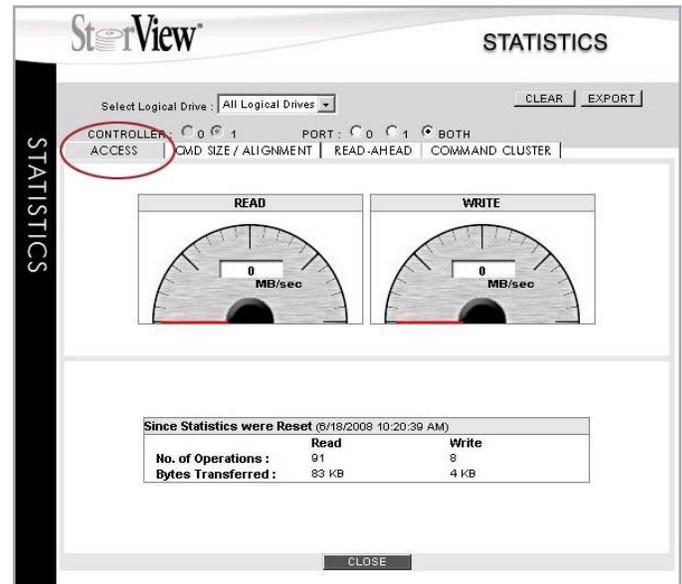


Fig. 167 - Statistics Screen – Access Tab

### 14.3. Command Size - Alignment Statistics

Command size statistics express the percentage of commands whose size is as specified. The Alignment statistics is the percentage of commands whose address aligned on the specified address boundary.

Table 31 - Alignment Statistics

Statistic	Description
Command Size	Expressed in the percentage of commands whose size is specified for reads and writes. The values are displayed with a horizontal bar for each value. The lack of a bar displayed for a specific value indicates it is 0% (or less than 1%). For example, consider a read or write command from a host system with Logical Block Address (LBA) 0x0000070, and access size 0x80, expressed in decimal 128. Using 512 byte blocks on the disk drives, it can be seen that this is a read of 64 Kbytes, which is the command size.
Alignment	This is the percentage of commands whose address is aligned on the specified chunk boundary. The alignment of a command from a host system is determined by the command's address. In an optimal system, a write of one chunk of data would reside exactly within a chunk on one disk. However, if this is not the case, this write will be split up into two separate writes to two different data drives. This of course will have a negative effect on performance. To overcome these problems, the user can, with more sophisticated operating systems, set the access size and alignment to an optimal value. These statistics can help the user to tune the operating system.
How to Use Command Size and Alignment	To calculate the alignment, we check the LBA for the largest number of blocks that will evenly divide into it, in powers of 2. So, we can see that in this case, the alignment is $0x10 = 16$ blocks. This equates to 8K. The alignment, in conjunction with the access size, gives an indication of how many drives are involved in an access. In the above example, consider a RAID 5/50/6 array with a chunk size of 64K. In this case, the above access will actually involve 2 data drives, since it needs to access 8K in the first drive ( $0x80 - 0x70 = 0x10$ blocks = 8K), and the remaining 56K in the next drive ( $0x70$ blocks = 56K). This is clearly inefficient, and could be improved by setting the alignment to 64K on the operating system. If that is not possible, using a larger chunk size can help, as this reduces the number of accesses that span chunks. Aligning an access on the same value as the access size will improve performance, as it will ensure that there are no multi-chunk accesses for commands that are smaller than a chunk size.

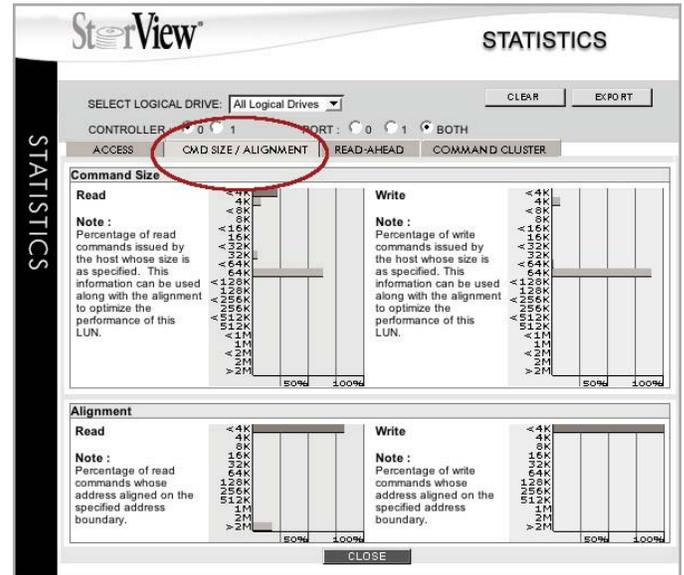


Fig. 168 - Statistics Screen – Command Size/Alignments Tab

## 14.4. Read-Ahead Statistics

If sequential read commands are sent to the controller, it assumes that the commands which follow may also be sequential. It can then go and perform a read of the data, before the host requests it. This improves performance, particularly for smaller reads. The size of the read-ahead is calculated based on the original command size, so the controller does not read too much data. The controller maintains statistics for all read-ahead commands performed.

Table 32 - Read-Ahead Statistics

Statistic	Description
Sequential Command Interval	<p>In determining whether to perform a read-ahead or not, the controller will search back in the command queue whenever it receives a new read command that is not satisfied by an existing read-ahead cache buffer. In a multi threaded operating system, commands from one thread may be interspersed with commands from another thread.</p> <p>This requires that the controller not just check the immediately previous command. The controller will search back for a number of commands, to see if the new command is exactly sequential to any one of these previous commands. If it is, then the controller determines that the data access pattern is sequential, and so performs a read-ahead. These statistics record the average number of commands the controller must search back for when it finds a sequential command match, the maximum number, and also the percentage for each one of these values. These give an indication of the multi threaded nature of the host.</p>
Read-Ahead Command Hit Rate	<p>This is the percentage of read command hits versus the total number of read commands that have been issued. This gives an indication of the sequential nature of the data access pattern from the host.</p>
Read-Ahead Command Efficiency	<p>This is the percentage of the number of read command hits versus the projected number of read-ahead command hits. This is a measure of the efficiency of the read-ahead algorithm. A low value means that much of the data that the controller reads in the read-ahead command is not actually requested by the host.</p>

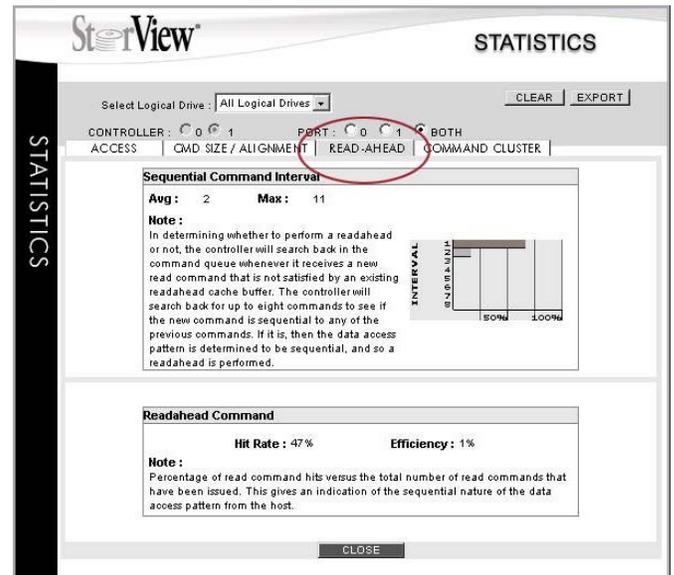


Fig. 169 - Statistics Screen – Read-Ahead Tab

## 14.5. Command Cluster Statistics

To increase performance, the controller can cluster sequential write commands together to create a larger write command. This results in less commands being sent to the disk drives. Additionally, if sufficient data is clustered by the controller, then it can perform a full stripe write for RAID5/50/6 arrays. This significantly improves performance. In cases where the host does not send a sufficient number of outstanding writes, writeback cache can be used to delay the write to disk, increasing the likelihood of clustering more data.

Table 33 - Command Cluster Statistics

Statistic	Description
Write Cluster Rate	This is the percentage of the number of write commands that are part of a cluster versus the total number of write commands that have been issued. This gives an indication of the sequential nature of the data access pattern from the host, and of the performance of the writeback cache.
RAID 5/50/6 Full Stripe Write Rate	This is the percentage of the amount of data that is written as a full stripe write versus the total amount of data written. This gives an indication of the sequential nature of the data access pattern from the host, and of the performance of the writeback cache, for RAID 5/50/6 arrays.
Command Cluster Count	When the controller clusters a write command, it may cluster a large number of them together. These statistics record the average and maximum number of commands the controller clusters, and also the percentage for each one of these values.
Command Cluster Interval	<p>In determining whether to cluster write commands or not, the controller will search back in the command queue whenever it receives a new write command. In a multi threaded operating system, commands from each thread may be interspersed with commands from another thread. This requires that the controller not just check the immediately previous command. The controller will search back for a number of commands, to try to determine if the new command is exactly sequential to any one of these previous commands. If it is, then the controller determines that it can cluster these commands.</p> <p>These statistics record the average and maximum number of commands the controller must search back for when it finds a sequential command match, and also the percentage for each one of these values.</p>

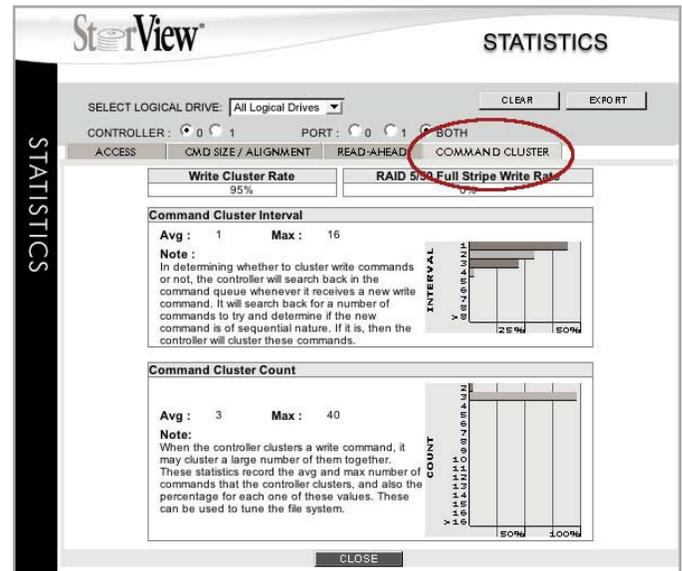


Fig. 170 - Statistics Screen – Command Cluster Tab

## 15. Optimizing Write Performance

### 15.1. Introduction

With a typical RAID 5 implementation, there are a number of steps that are performed when data is written to the media. Every write from the host system will typically generate two XOR operations and their associated data transfers, to two drives. If the accesses are sequential, the parity information will be updated a number of times in succession. However, if the host writes sufficient data to cover a complete stripe, the parity data does not need to be updated for each write, but it can be recalculated instead. This operation takes only one XOR operation per host write, compared to two for a standard RAID 5 write. The number of data transfers necessary are also reduced, increasing the available bandwidth. This type of write access is termed a “Full Stripe Write.”

---

**NOTE:** This appendix was written for RAID 5. However, with RAID 6 which uses two parity drives, the size of the stripe based on the number of data drives multiplied by the chunk size principle is the similar.

---

The illustration at right displays the distribution of data chunks (denoted by C<sub>x</sub>) and their associated parity (denoted by P(y-z)) in a RAID 5 array of five drives. An “array” is defined as a set of drives, on which data is distributed. An array will have one RAID level. A “chunk” is the amount of contiguous data stored on one drive before the controller switches over to the next drive. This parameter is adjustable from 64K to 256K, and should be carefully chosen to match the access sizes of the operating system. A Stripe is a set of disk chunks in an array with the same address. In the example at right, Stripe 0 consists of C0, C1, C2, and C3 and their associated parity P(0-3).

Maximum performance will be achieved when all drives are performing multiple commands in parallel. To take advantage of a Full Stripe Write, the host has to send enough data to the controller. This can be accomplished in two ways. First, if the host sends one command with sufficient data to fill a stripe, then the controller can perform a Full Stripe Write. Alternatively, if the host sends multiple sequential commands, smaller than a stripe size (typically matching the chunk size), the controller can internally combine these commands to get the same effect. In the above example, if a 256K chunk size is used, then the stripe size is 1MB (4 chunks \* 256K). So, for maximum performance, the host could either send 5 \* 1 MB write commands, or 20 \* 256K write commands.

The effectiveness of the controller’s ability to perform a Full Stripe Write depends on a number of parameters (see following sections).

---

P (20-23)	C20	C21	C22	C23	Stripe 5
C16	C17	C18	C19	P (16-19)	Stripe 4
C12	C13	C14	P (12-15)	C15	Stripe 3
C8	C9	P (8-11)	C10	C11	Stripe 2
C4	P (4-7)	C5	C6	C7	Stripe 1
P (0-3)	C0	C1	C2	C3	Stripe 0

Fig. 171 - Distribution of Data and Parity in a RAID 5 with Five Drives

## 15.2. Sequential Access

If the commands sent from the host are not sequential, the controller will not be able to cluster them together. So, unless each individual access is sufficient to fill a stripe, a Full Stripe Write will not occur.

### 15.2.1. Number of Outstanding Commands

For the controller to successfully cluster commands, there has to be a number of write commands sent simultaneously. Setting the host to send up to 64 commands should prove adequate. Alternatively, enabling writeback cache will have a similar effect, as the controller can then cluster sequential commands even if the host only sends a small number of commands at a time.

## 15.3. Access Size

With very small accesses, it is necessary to have a large number of commands to cluster together to fill up a full stripe. So, the larger the access size the better. It is best to use an access size that will fill a chunk. Of course, even if a stripe is not filled up, small sequential writes will still benefit from command clustering.

### 15.3.1. Access Alignment

The alignment of a command from a host system is determined by the command's address. In an optimal system, a write of one chunk of data would reside exactly within a chunk on one disk. However, if this is not the case, this write will be split up into two separate writes to two different data drives. This will have a negative effect on performance. To overcome these problems, the user can, with more sophisticated operating systems, set the access size and alignment to an optimal value.

As can be seen from [Fig. 172](#), to get the highest performance from this system, it is necessary to have a number of stripes being written in parallel. As the array expands, with more and more drives, the number of commands (and amount of sequential data) necessary to do this increases.

In [Fig. 172](#), we can see that seven chunks of sequential data are necessary to fill a stripe. To have multiple commands active for all disk drives, this requires more data than for the case with five drives. As can be seen, this number will increase as the number of drives increases. If a large number of drives are used, it may get difficult to achieve maximum performance, as it becomes more difficult to cluster a large number of commands to achieve a Full Stripe Write.

P (56-62)	C56	C57	C58	C59	C60	C61	C62	Stripe 8
C49	C50	C51	C52	C53	C54	C55	P (49-55)	Stripe 7
C42	C43	C44	C45	C46	C47	P (42-48)	C48	Stripe 6
C35	C36	C37	C38	C39	P (20-23)	C61	C62	Stripe 5
C28	C29	C30	C31	P (28-34)	C32	C33	C34	Stripe 4
C21	C22	C23	P (21-27)	C24	C25	C26	C27	Stripe 3
C14	C15	P (14-20)	C16	C17	C18	C19	C20	Stripe 2
C7	P (7-13)	C8	C9	C10	C11	C12	C13	Stripe 1
P (0-6)	C0	C1	C2	C3	C4	C5	C6	Stripe 0

Fig. 172 - Distribution of Data and Parity in a RAID 5 with Eight Drives

## 15.4. Sub-Array

The difficulty in realizing the maximum performance possible introduces the concept of a sub-array.

Suppose an array consisted of two RAID 5 sets, see [Fig. 173](#). If these are then striped, the resulting array would appear as shown at right. In this case, in order for a Full Stripe Write to be performed, it is still only necessary to cluster four write commands together, as opposed to the seven necessary as indicated at right. The array of drives appears as two separate sub-arrays, each with its own rotating parity.

It can be seen that the more sub-arrays used, the more likely it is for a Full Stripe Write to occur, and hence the higher the performance. It is recommended to use either four or five drives in a sub-array, for best performance. [Fig. 174](#) shows that even with 15 drives, it is still possible to perform Full Stripe Writes, by clustering together 4 chunks of data.

P (40-43)	C40	C41	C42	C43	P (44-47)	C44	C45	C46	C47	Stripe 5
C32	C33	C34	C35	P (32-35)	C36	C37	C38	C39	P (36-39)	Stripe 4
C24	C25	C26	P (24-27)	C27	C28	C29	C30	P (28-31)	C31	Stripe 3
C16	C17	P (16-19)	C18	C19	C20	C21	P (20-23)	C22	C23	Stripe 2
C8	P (8-11)	C9	C10	C11	C12	P (12-15)	C13	C14	C15	Stripe 1
P (0-3)	C0	C1	C2	C3	P (4-7)	C4	C5	C6	C7	Stripe 0

Fig. 173 - Distribution of Data and Parity in a RAID 5 with Ten Drives and Two Sub-Arrays

P	C60	C61	C62	C63	P	C64	C65	C66	C67	P	C68	C69	C70	C71	Stripe 5	
C48	C49	C50	C51	P	C52	C53	C54	C55	P	C56	C57	C58	C59	P	C47	Stripe 4
C36	C37	C38	P	C39	C40	C41	C42	P	C43	C44	C45	C46	P	C47	Stripe 3	
C24	C25	P	C26	C27	C28	C29	P	C30	C31	C32	C33	P	C34	C35	Stripe 2	
C12	P	C13	C14	C15	C16	P	C17	C18	C19	C20	P	C21	C22	C23	Stripe 1	
P	C0	C1	C2	C3	P	C4	C5	C6	C7	P	C8	C9	C10	C11	Stripe 0	

Fig. 174 - Distribution of Data and Parity in a RAID 5 with Fifteen Drives and Three Sub-Arrays

## 15.5. Multiple Drive Failures

In a configuration with multiple sub-arrays, it is possible for the array to sustain multiple drive failures, provided that there is only one failure in each sub-array.

### 15.5.1. Faster Rebuild

A rebuild operation must read data and calculate parity from all the remaining drives in the array. If multiple sub-arrays are used, this means that it is only necessary to read the data from the remaining drives in the sub-array, not all of the drives in the array. This increases both the rebuild speed and the speed of access to missing data, which also has to be recreated from the remaining drives.

## 15.6. Summary

In summary, for maximum performance, it is recommended to use four or five drives in a sub-array. If there are more than five drives in a sub-array, it is better to use a smaller chunk size, say 64K or 128K, as this will lead to more Full Stripe Writes.

## 16. Troubleshooting

### 16.1. Problems You May Encounter

This appendix provides typical solutions for problems you may encounter while using StorView to control and manage the storage systems. Also refer to the Event chapter, and review the cause and actions for each event listed.

Table 34 - Troubleshooting

Symptom	Reason	Solution
Continuous indications that the partner controller has failed or is missing.	<p>A partner controller in an Active-Active configuration has failed or was removed.</p> <p>Operating in a Stand-Alone configuration with Single Controller Mode not selected.</p>	<p>Until the partner controller is replaced, temporarily enable Single Controller Mode in the Controller Parameters tab.</p> <p>Be sure to disable this option when the partner controller is to be replaced.</p> <p>If you are operating in a Stand-Alone configuration, enable the Single Controller Mode setting in the Controller Parameters tab.</p>
Password Error	<p>Password not accepted at log in.</p> <p>Password was forgotten or lost.</p>	<p>Password is case sensitive, ensure that the password is entered correctly.</p> <p>Contact technical support for the procedures to recover from a lost or missing password.</p>
Lost communication with the RAID Controllers.	Service is hung.	<p>Restart the StorView service. Access the Control Panel and double-click on Services. Locate the StorView Service and click Stop. Once the service has stopped, click Start and retry the connection by clicking the RESCAN button on the StorView Main screen.</p> <p>On Linux system access the process viewer and stop the StorView Process. Restart the process and click the RESCAN button on the StorView Main screen.</p>
Hot spare not automatically starting when drive failure occurs in a redundant array in which a global or dedicated hot spare is defined.	<p>The Auto Rebuild option is not enabled in the Controller Parameters.</p> <p>Hot spare disk drive is too small to be used for the drive replacement.</p> <p>Waiting for a valid replacement drive to be inserted.</p>	<p>Open the Controller Information window (click the Controller icon), place a check mark in the box by clicking the check box on the Auto Rebuild parameter.</p> <p>Ensure that the disk drive defined as a hot spare is equal to or greater than the size of the drive members of the arrays.</p> <p>Auto Rebuild is not selected and no hot spare drive is assigned, but Auto Hot Spare is enabled. The array will begin rebuilding once a valid replacement drive is inserted in the drive slot of the failed drive.</p>

Symptom	Reason	Solution
Consistently occurring time out errors when the browser window is open.	Host HBA parameter settings are not configured for best performance optimization.	<p>Access your Host HBA settings and make the following changes: Execution Throttle improve general I/O performance by allowing more commands on the fibre bus. Do this by changing your host bus adapter's execution throttle parameter to 256.</p> <p>Scatter/Gather: (Microsoft Windows) Increase the general I/O performance by allowing larger data transfers. Do this by editing the "MaximumSGList" parameter in the register. The recommended hexadecimal value is "ff." The path is: <a href="#">HKEY_LOCAL_MACHINE/System/CurrentControlSet/Services/&lt;name of HBA driver&gt;/Parameters/Device/.</a></p>
Shared Memory Error is displayed.	The CGI script manager may have not released a segment of shared memory.	This may occur when heavy I/O is happening at the same time you are accessing StorView. If this occurs you will need to stop and then restart the StorView Server service.
After switching drives and/or controllers from one storage solution enclosure to another, one of the solutions reports that the storage solution is being monitored by another host.	Multiple Configuration WWNs being used.	If you have been interchanging configured drives or controllers between storage solutions you may have a situation where multiple solutions are now sharing the same Configuration WWN. This can be corrected by changing the Configuration WWN value found in the Controller Parameters on either of the storage solutions. After making this change, all participating host systems will require a reboot.
Inadvertently pulled the incorrect drive from the enclosure and the array is dead.	Possible incorrect drive identification and removal.	<p>If by mistake you remove a working drive member instead of the failed drive, this can cause the array to fail. In most cases you can simply re-insert that drive that was incorrectly removed and the array will return to the same state it was in prior to removing the drive.</p> <p>For RAID 5/50/6 arrays, a drive failure will put the array in a critical state, if a hot spare was available the array should go into a rebuild mode. If you inadvertently remove one of the known good drives that is in the process of rebuilding, the rebuild operation will stop. Once you re-insert the incorrectly removed drive the array will return to the critical state and the rebuild process will start again. If you did not have a hot spare assigned, the array will be in a critical state. If you inadvertently remove a known good drive instead of the failed drive the array will change to a failed array state. Re-inserting that inadvertently removed drive will put the array back into a critical state. Replacing the failed drive will cause the array to begin a rebuild operation provided that you assign it as a hot spare or, if the Auto Hot Spare option was enabled the rebuild will begin automatically as the new replacement drive is installed.</p>

continued on the next page >>

Symptom	Reason	Solution
Inadvertently pulled the incorrect drive from the enclosure and the array is dead. (continued)	Possible incorrect drive identification and removal. (continued)	<p>For RAID 0 arrays, if you inadvertently remove a known good drive, the array will become dead. Once you re-insert the incorrectly removed drive the array will return to its working state.</p> <p>For RAID 1/10 arrays, if you inadvertently remove a known good drive, the array will become failed. Once you re-insert the incorrectly removed drive the array will return to its previous state. If the array was critical, you can then replaced the known failed drive with a working drive and assign it as a hot spare and the array will begin rebuilding.</p> <hr/> <p><b>NOTE:</b> For all arrays removing a drive as described above will cause all current processing I/O from the controller to stop. Any I/O in progress may have been lost or cause a corrupt file. Be sure to verify all data stored during this type of incidence to ensure data reliability.</p> <hr/>
The controller's IDs and/or Configuration WWN was changed and now there is a communication failure.	When you changed the controller IDs, a new nexus is established which requires the operating system and software to establish new communication paths.	If you using Microsoft Windows you can use the StorView "RESCAN" feature to relocate the storage solution(s).
StorView displays a message: "No storage solution found."	The host operating system is not able to see the storage solution. Ensure that the Fibre devices appear in your HBA's BIOS.	<p>Ensure that you have the latest driver installed for your HBA. Probe the SCSI enclosure to ensure that you see the solution.</p> <p>Reboot the host and the storage system.</p>
I received the following message: "Lost communication with server. The server maybe down."	During heavy host operations and/or data I/O, the system may become too busy to complete CGI requests from the GUI in the time allocated.	<p>After several updated attempts have failed you will see this message. At this time you can try to use the Browsers' refresh function to reload the StorView GUI. If that is unsuccessful, you may need to stop and then restart the StorView Server service. If you continue to receive that message, close the browser and wait until I/O traffic has settle down before opening the StorView GUI back up.</p> <p>You will still continue to receive email notifications and Event logging.</p>
During heavy data I/O, when I try to make a configuration change I get a failure saying that the controller is busy.	The controller's onboard resources are consumed.	Configuration changes during heavy I/O are not recommended. You can either wait until there is less data traffic or keep re-trying the command until it is successful.
Enclosure image on the Main screen is dimmed or greyed out.	Enclosure Support option has been disabled.	Access the Controller Information window by clicking the Controller icon. Verify the option "Enclosure Support" is check and click APPLY. Close the window.
Cannot start SSL StorView while using a Microsoft Proxy Server.	Issue with allowing the secure port on the proxy.	Refer to section <a href="#">4.1. Starting StorView</a> for details on issues with a proxy server.

## 17. More on SnapShot

### 17.1. Overview

This chapter is provided to give a more in depth look at Snapshot, in order to provide a better understanding of the technology used.

The Snapshot engine has been designed to give the user considerable flexibility when setting up and performing snapshot and snapback operations. A snapshot allows the user to recreate a point in time image of a logical drive that contains exactly the same data at the point the snapshot was taken.

To recreate the data of the snapshot, a snapback is initiated. This will cause the controller to copy data from the Overwrite Data Area (ODA) back to the snapshotted logical drive. Once completed, the logical drive will now contain exactly the same data as it did at the time of the snapshot.

A snapshot image can also be presented back to the host server as a snapshot (virtual) logical drive, in which the server will see this virtual LUN and its data just like any other logical drive.

One of the most important components of Snapshot is the Overwrite Data Area or ODA. The ODA serves as a storage area for disk data that needs to be protected from being overwritten by new data after a snapshot was issued.

### 17.2. Terminology

The following describes some of the terminology used when working with Snapshot.

Table 35 - Snapshot Terminology

Term	Description
Internal Logical Drive	An internal logical drive is identical to a regular logical drive except that it is NOT made visible to a host adapter as a LUN. Instead internal logical drives are used for setting up snapshot ODA that are only accessed internally by the RAID controller.
Overwrite Data Area, ODA	An internal storage area on an array that is dedicated to storing data from a snapshotted logical drive. The data stored on the ODA is the data from the logical drive that needed to be overwritten after a snapshot was initiated. The ODAs are mapped on top of internal logical drives. An ODA cannot be accessed externally through a host LUN. They are only accessed internally.

Term	Description
Snapback	The process of restoring a logical drive from a selected snapshot. This process takes place internally in the RAID controller firmware and needs no support from any backup utility.
Snapshot	A method for producing a point-in-time image of a logical drive. In the process of initiating a snapshot, no data is actually copied from the snapshotted logical drive. However as new writes are made to a snapshotted logical drive, existing data blocks are copied to the ODA before the new data is written to the logical drive.
Snapshot Number	Identifier that references one of several snapshots of the same logical drive.
Snapshotted LD	A logical drive that has one or more snapshots initiated on it.
ODA Size	The size of the Overwrite Data Area.
ODA Stripe Size	The read/write block size that the system will use when copying data from the original logical drive to the ODA.
Virtual LUN or Snapshot LUN	A special LUN created from a combination of the snapshotted logical drives' data and the data contained in the ODA.

---

## 17.3. Overwrite Data Area

The Overwrite Data Area (ODA) is a reserved area on the disk array used for storing data from logical drives that are operating under a snapshot condition. They cannot be made visible as LUNs like other logical drives.

Up to 256 ODAs can be configured. Each ODA can have any capacity assigned to them and they can be mapped onto any disk array with any RAID level (except RAID 6), chunk size and stripe size.

Once the ODA is created, snapshots can be taken on a logical drive at any time. Each logical drive can have up to 64 snapshots issued.

ODA initialization is performed automatically by the RAID Controller during configuration. The user will only need to provide one parameter during initialization and that is the ODA strip size. This defines how much data that will be protected per copy on write operation. When the snapshot engine receives a write command and has determined the underlying data must be protected, the size of the read and write commands will be determined by the ODA stripe size. If the ODA stripe size is set to 128Kbytes, then the protective read write will be done on 256 disk blocks at a time.

---

## 17.4. Overwrite Data Area Location

The overwrite data areas can be configured in two different ways. On the same disk array as a regular logical drive or on one or more separate disk arrays dedicated for the use of ODAs.

If the ODA area is located on the same physical disk array as the logical drive being snapshot, the performance of both reads and writes to the logical drive will be affected. Instead if the ODA area is located on its own dedicated disk array, it results in much better performance characteristics.

It is always preferred to dedicate a disk array to be used exclusively by one or more ODAs.

---

## 17.5. Selecting the ODA Stripe Size

During ODA initialization, the user is asked to select a stripe size for the ODAs. Five different options are available: 64 Kbytes, 128 Kbytes, 256 Kbytes, 512 Kbytes, and 1 Mbyte. In an array with multiple ODAs defined, each ODA can be set up with its own stripe size different from the other ODAs.

The ODA stripe size represents the size of the data area that is always protected during a snapshot copy-on-write operation. Choosing the best ODA stripe size is not possible without some knowledge about

snapshot. The ODA saves data to be overwritten in multiples of its stripe size. A large stripe size will offer a smaller number of data areas while a smaller stripe size will offer more data areas.

If the server application issues a high percentage of small writes in a random address pattern, then it will be better to select a small ODA stripe size such as 64 Kbytes. That way the ODA has a lot of smaller data areas to use for a large amount of small random writes. A disadvantage will be that the probability for follow on hits in an already protected area of the logical drive is smaller compared with using a larger ODA stripe size.

If the server application issues a high percentage of small random writes clustered together in groups on the logical drive, then it would be preferred to use a large ODA stripe size. This increases the probability that other small writes will be addressed to a disk area that is already protected by a larger stripe of data.

This will minimize the amount of copy-on-write sequences needed and therefore minimize the impact on the disk array's performance. The disadvantage would be that the total number of ODA stripes would be reduced if the write access pattern suddenly became very random in its nature.

For applications that perform a large percentage of sequential writes it would be best to select a larger ODA stripe size to minimize the amount of copy-on-write operations needed. The results will be better performance.

---

the intended write data access pattern on the logical drive to be

## 17.6. Overwrite Data Area Storage Capacity

The maximum storage capacity of an ODA is dependant on the ODA stripe size.

Table 36 - Available ODA Strip Sizes

ODA Stripe Size	ODA Storage Capacity
64 Kbytes	512 GBs
128 Kbytes	1 Terabyte
256 Kbytes	2 Terabytes
512 Kbytes	4 Terabytes
1 Mbyte	8 Terabytes

If the Overwrite Data Area storage space gets close to running out of storage space, it will alert the user. If the Overwrite Data Area does get filled to its limit and can no longer store copy-on-write data, the Snapshot will become out of phase with the underlying snapshotted logical drive resulting in all Snapshots for this logical drive being invalidated by the RAID Controller.

## 17.7. ODA Assignment

Each logical drive can be snapshotted up to 64 times. Each one of these 64 snapshots represents a full virtual copy of the snapshotted logical drive at the point in time of the snapshot.

Before snapshots can be issued on a logical drive an ODA area must be assigned for its use. Once the ODA has been assigned it will remain assigned as long as it is needed.

## 17.8. Snapshots and Caching

It is of the highest importance that disk data caching be handled properly when using the RAID Controller snapshot engine.

The purpose of a snapshot is to generate a point-in-time reference marker that allows a system administrator restoration of a logical drive data to the state it was at the time the snapshot was taken. Due to disk data caching, it is required to keep track of the status of the various caches present in the server installation environment. A server might believe a block of data has been committed to the disk when it is actually kept waiting in the RAID Controller cache. Likewise, an application such as a database might believe its data has been saved on the disk but in reality

the data still resides in the server OS cache and has not yet been sent to the disk array.

### 17.8.1. Application Caches

Advanced applications, such as databases, might implement their own data cache to speed up execution. Before a snapshot can be issued this application cached data must be flushed. If the application does not provide a means to flush its cache, close down the application before the snapshot is issued.

Refer to your application user documentation for more information.

### 17.8.2. Operating System Cache

The second highest level of caching disk data is provided by the operating system (OS). The OS will dedicate a large percentage of its DRAM space to be used for disk data caching. It is required to ask the OS to flush its entire disk cache just before the snapshot is issued.

Refer to your operating system user guide for more information.

### 17.8.3. RAID Controller Cache

The RAID Controller(s) implement large caches for disk data. A large portion of these caches are used as delayed writeback caches for holding data blocks not yet written to the disk drives. It is very important to have the RAID Controller(s) perform a flush of its cache before a snapshot is started.

## 17.9. Virtual LUNs

Virtual LUNs are also known as snapshot LUNs and can be made available for access through a dynamically created LUN. Reads issued to a virtual LUN will be serviced in a similar way to reads issued to a normal LUN. The data will be the exact data as it were at the time of the snapshot. This allows an application, such as disk backup software to perform a backup operation of the Virtual LUN.

A virtual LUN will respond to any type of incoming SCSI command in the same way as the original logical drive. The exception is with write commands, since the snapshots represent frozen data in time, it is not allowed for the host to modify this data.

### 17.9.1. Virtual LUN Numbering

Each Virtual LUN will be presented to the host side as a regular logical drive (LUN). A logical unit number will be assigned and used by the accessing host. These LUN numbers are automatically assigned.

Up to 256 ODAs can be defined within the array. Therefore the ODAs are numbered 0-255. Each ODA can store up to 64 snapshots, so the snapshot number will range between 0-63. The possible range for the Virtual LUNs are 0-511.

The side effect of LUN numbering is a conflict with a normal LUN is high. Therefore, the virtual LUNs are not automatically made visible to the host system. This is accomplished using LUN mappings and each Virtual LUN is manually assigned a LUN number that will then become visible to the host.

Example: a disk array with three normal data LUNs configured as LUN 0-2 and one ODA, will have its first Virtual LUN numbered 0. To avoid a conflict between the original data LUN 0, the Virtual LUN will have to be LUN mapped to a non-conflicting LUN number, e.g., LUN 10. All accesses to LUN 10 will then be translated into Virtual LUN 0 internally by the RAID Controller.

---

## 18. Contacting Customer Support

### 18.1. LaCie Technical Support Contacts

LaCie Asia, Singapore, and Hong Kong Contact us at: <a href="http://www.lacie.com/asia/contact/">http://www.lacie.com/asia/contact/</a>	LaCie Australia Contact us at: <a href="http://www.lacie.com/au/contact/">http://www.lacie.com/au/contact/</a>
LaCie Belgium Contact us at: <a href="http://www.lacie.com/be/contact/">http://www.lacie.com/be/contact/</a> (Français)	LaCie Canada Contact us at: <a href="http://www.lacie.com/caen/contact/">http://www.lacie.com/caen/contact/</a> (English)
LaCie Denmark Contact us at: <a href="http://www.lacie.com/dk/contact">http://www.lacie.com/dk/contact</a>	LaCie Finland Contact us at: <a href="http://www.lacie.com/fi/contact/">http://www.lacie.com/fi/contact/</a>
LaCie France Contact us at: <a href="http://www.lacie.com/fr/contact/">http://www.lacie.com/fr/contact/</a>	LaCie Germany Contact us at: <a href="http://www.lacie.com/de/contact/">http://www.lacie.com/de/contact/</a>
LaCie Italy Contact us at: <a href="http://www.lacie.com/it/contact/">http://www.lacie.com/it/contact/</a>	LaCie Japan Contact us at: <a href="http://www.lacie.com/jp/contact/">http://www.lacie.com/jp/contact/</a>
LaCie Netherlands Contact us at: <a href="http://www.lacie.com/nl/contact/">http://www.lacie.com/nl/contact/</a>	LaCie Norway Contact us at: <a href="http://www.lacie.com/no/contact/">http://www.lacie.com/no/contact/</a>
LaCie Spain Contact us at: <a href="http://www.lacie.com/es/contact/">http://www.lacie.com/es/contact/</a>	LaCie Sweden Contact us at: <a href="http://www.lacie.com/se/contact">http://www.lacie.com/se/contact</a>
LaCie Switzerland Contact us at: <a href="http://www.lacie.com/chfr/contact/">http://www.lacie.com/chfr/contact/</a> (Français)	LaCie United Kingdom Contact us at: <a href="http://www.lacie.com/uk/contact/">http://www.lacie.com/uk/contact/</a>
LaCie Ireland Contact us at: <a href="http://www.lacie.com/ie/contact/">http://www.lacie.com/ie/contact/</a>	LaCie USA Contact us at: <a href="http://www.lacie.com/contact/">http://www.lacie.com/contact/</a>
LaCie International Contact us at: <a href="http://www.lacie.com/intl/contact/">http://www.lacie.com/intl/contact/</a>	

## 19. Warranty Information

LaCie warrants your drive against any defect in material and workmanship, under normal use, for the period designated on your warranty certificate. In the event this product is found to be defective within the warranty period, LaCie will, at its option, repair or replace the defective drive. This warranty is void if:

- ◆ The drive was operated/stored in abnormal use or maintenance conditions;
- ◆ The drive is repaired, modified or altered, unless such repair, modification or alteration is expressly authorized in writing by LaCie;
- ◆ The drive was subjected to abuse, neglect, lightning strike, electrical fault, improper packaging or accident;
- ◆ The drive was installed improperly;
- ◆ The serial number of the drive is defaced or missing;
- ◆ The broken part is a replacement part such as a pickup tray, etc.
- ◆ The tamper seal on the drive casing is broken.

LaCie and its suppliers accept no liability for any loss of data during the use of this device, or for any of the problems caused as a result.

LaCie will not, under any circumstances, be liable for direct, special or consequential damages such as, but not limited to, damage or loss of property or equipment, loss of profits or revenues, cost of replacement goods, or expense or inconvenience caused by service interruptions.

Any loss, corruption or destruction of data while using a LaCie drive is the sole responsibility of the user, and under no circumstances will LaCie be held liable for the recovery or restoration of this data.

Under no circumstances will any person be entitled to any sum greater than the purchase price paid for the drive.

To obtain warranty service, call LaCie Technical Support. You will be asked to provide your LaCie product's serial number, and you may be asked to furnish proof of purchase to confirm that the drive is still under warranty.

All drives returned to LaCie must be securely packaged in their original box and shipped with postage prepaid.

---

**IMPORTANT INFO:** Register online for free technical support:  
[www.lacie.com/register](http://www.lacie.com/register)

---