



SANtricity ES Storage Manager Concepts Guide for Version 10.77

Version 10.77

May 2011

51328-00, Rev. A



Revision History

Version and Date	Description of Changes
51328-00, Rev. A May 2011	Initial release of the document.

LSI and the LSI & Design logo, StorageTek, SANtricity, HotScale, and SANshare are trademarks or registered trademarks of LSI Corporation or its subsidiaries or Sun Microsystems, Inc. All other brand and product names may be trademarks of their respective companies.

This document contains proprietary information of LSI Corporation and Sun Microsystems, Inc. The information contained herein is not to be used by or disclosed to third parties without the express written permission of an officer of LSI or Sun.

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

LSI products are not intended for use in life-support appliances, devices, or systems. Use of any LSI product in such applications without written consent of the appropriate LSI officer is prohibited.

LSI Corporate Headquarters
Milpitas, CA
800-372-2447

Email
globalsupport@lsi.com

Website
www.lsi.com

Document Number: 51328-00, Rev. A
Copyright © 2011 LSI Corporation. All rights reserved.
Copyright © 2011 Sun Microsystems, Inc. All rights reserved.

Table of Contents

Chapter 1: Storing Your Data	1
Storage Arrays	1
Storage Area Networks	1
Management Methods	1
Out-of-Band Management	2
In-Band Management	2
RAID Levels and Data Redundancy	2
Dynamic RAID-Level Migration	4
Hardware Redundancy	4
Controller Cache Memory	5
Module Loss Protection	5
Hot Spare Drives	6
Channel Protection	6
I/O Data Path Protection	6
Multi-Path Driver with AVT Enabled	7
Multi-Path Driver with AVT Disabled	7
Target Port Group Support	7
Load Balancing	8
Round Robin with Subset	8
Least Queue Depth with Subset	8
Least Path Weight with Subset	8
Chapter 2: Introducing the Storage Management Software	9
Enterprise Management Window	9
Parts of the Enterprise Management Window	10
EMW Devices Tab	10
EMW Setup Tab	12
Adding and Removing a Storage Array	12
Array Management Window	13
Starting the Array Management Window	13
Summary Tab	13
Logical Tab	14
Physical Tab	15
Mappings Tab	17
AMW Setup Tab	18
Support Tab	19
Managing Multiple Software Versions	19
Chapter 3: Configuring the Storage Arrays	20
Volumes and Pools	20
Standard Volumes	20
Pools	21
Pool Creation	21
Dynamic Capacity Expansion	23
Register the Volume with the Operating System	24
Premium Features	24
Storage Domains	24
Snapshot Volume Premium Feature	25

Data Replicator Software Premium Feature	27
Volume Copy Premium Feature	30
Drive Security and Enterprise Key Manager	35
T10 Protection Information Premium Feature	39
Solid State Disks	39
Heterogeneous Hosts	40
Password Protection	40
Persistent Reservations Management	41
HotScale Technology	41
Chapter 4: Maintaining and Monitoring Storage Arrays	42
Storage Array Health	42
Background Media Scan	42
Event Monitor	42
Alert Notifications	43
Performance Monitor	44
Viewing Operations in Progress	46
Retrieving Trace Buffers	47
Upgrading the Controller Firmware	47
Monitoring the Status of the Download	49
Problem Notification	50
Event Log Viewer	51
Storage Array Problem Recovery	51
Recovery Guru	51
Glossary	52
A	52
C	52
D	52
F	53
H	53
I	54
L	54
M	54
N	54
O	55
P	55
R	55
S	57
T	58
U	59
V	59
W	59

Chapter 1: Storing Your Data

The topics in this section describe the basic storage concepts, methods for managing storage arrays, including data-protection strategies, and multi-path failover drivers.

For additional information and detailed procedures for the options described in this section, refer to the online help topics for your version of the storage management software.

Storage Arrays

A storage array has redundant components, including drives, controllers, power supplies, and fans. These redundant components keep the storage array operational if a single component fails.

The storage array configuration provides a secure and robust system with which to store large amounts of data and allows for a variety of backup and retrieval scenarios. Administrators can set up the storage management software to maintain a specific level of security and configuration on the storage area network, such that the network requires little human interaction to perform its daily functions.

Storage Area Networks

A storage area network (SAN) transfers data between computers and storage systems. A SAN is comprised of many hardware components. Each hardware component might have a device manager or third-party management software.

A SAN includes one or more storage arrays that are managed by one or more servers or hosts running the SANtricity ES Storage Manager.

NOTE The SANtricity ES Storage Manager software is also referred to as the storage management software.

You can use the storage management software to add, monitor, manage, and remove the storage arrays on your SAN. Within the storage management software, you can configure the data to be stored in a particular configuration over a series of physical storage components and logical (virtual) storage components.

The I/O data and management instructions are sent from a host to the controllers in the storage array. When the I/O data reaches the controllers, they are distributed across a series of drives, which are mounted in modules.

The SAN can also include storage management stations, which also run the storage management software. A storage management station manages the storage arrays but does not send I/O data to them. Although physical storage array configurations vary, all SANs work using these basic principles.

Management Methods

Depending on your system configuration, you can use an out-of-band management method, an in-band management method, or both to manage a storage array controller from a storage management station or host.

NOTE A maximum of eight storage management stations can concurrently monitor an out-of-band managed storage array. This limit does not apply to systems that manage the storage array through the in-band management method.

Out-of-Band Management

You can use the out-of-band management method to manage a storage array directly over the network through an Ethernet connection, from a storage management station to the ethernet port on the controllers. This management method lets you manage all of the functions in the storage array.

NOTE Storage management stations require Transmission Control Protocol/Internet Protocol (TCP/IP) to support the out-of-band management of storage arrays.

In-Band Management

You can use the in-band management method to manage a storage array in which the controllers are managed through an I/O connection from a storage management station to a host that is running host-agent software. The I/O connection can be Serial Attached SCSI (SAS), Fibre Channel (FC), or internet SCSI (iSCSI). The host-agent software receives communication from the storage management client software and passes it to the storage array controllers along an I/O connection. The controllers also use the I/O connections to send event information back to the storage management station through the host.

When you add storage arrays by using this management method, you must specify only the host name or IP address of the host. After you add the specific host name or IP address, the host-agent software automatically detects any storage arrays that are connected to that host.

NOTE Systems running desktop (non-server) Windows operating systems and desktop Linux operating systems can be used only as storage management stations. You cannot use systems running desktop operating systems to perform I/O to the storage array and to run the host-agent software.

RAID Levels and Data Redundancy

RAID is an acronym for Redundant Array of Independent Disks. The storage solution stores the same data or information about the data (parity) in different places on multiple hard drives. Data can be written in parallel to multiple drives, which can improve performance. If a drive fails, the redundant data or parity data is used to regenerate the data on the replacement drive.

RAID relies on a series of configurations, called levels, to determine how user data and redundancy data are written to and retrieved from the drives. Each level provides different performance features and protection features. The storage management software offers six formal RAID level configurations: RAID Level 0, RAID Level 1, RAID Level 3, RAID Level 5, RAID Level 6, and RAID Level 10.

RAID Level 1, RAID Level 3, RAID Level 5, RAID Level 6, and RAID Level 10 write redundancy data to the drive media for fault tolerance. The redundancy data might be a copy of the data or parity data. Parity data is derived through a logical operation on the data, and is used for reconstruction of lost data. The parity data might exist on only one drive, or the parity data might be distributed among all of the drives in a pool.

The controller logically groups a set of drives together to create a pool. Each pool can contain one or more volumes. You can configure only *one* RAID level across each pool. Each pool stores its own redundancy data. The capacity of the pool is the aggregate capacity of the member drives, minus the capacity that is reserved for redundancy data. The amount of capacity needed for redundancy data depends on the RAID level used.

Table 1 RAID Level Configuration Table

RAID Level	Short Description	Detailed Description
RAID Level 0	No protection against loss of a drive (non-redundant), striping mode	<ul style="list-style-type: none"> ■ A minimum of one drive is required for RAID Level 0. ■ RAID Level 0 can use the maximum number of drives in a storage array. ■ You can use RAID Level 0 for high-performance needs, but it does not provide <i>data redundancy</i>. ■ Data is striped across all of the drives in the pool. ■ Do not use this RAID level for high data-availability needs. RAID Level 0 is better for non-critical data. ■ A single drive failure in a pool causes all of the volumes associated with the pool to fail, and data loss will occur
RAID Level 1 or RAID Level 10	Striping and mirroring mode	<ul style="list-style-type: none"> ■ A minimum of two drives are required for RAID Level 1: one for the user data and one for the mirrored data. If you select four or more drives, RAID Level 10 is automatically configured across the pool: two drives for the user data, and two drives for the mirrored data. ■ RAID Level 1 and RAID Level 10 can use the maximum number of drives in a storage array. ■ RAID Level 1 and RAID Level 10 typically provide the best write performance, but not in all cases. On a RAID Level 1 volume, data is written to a duplicate drive. On a RAID Level 10 volume, data is striped across mirrored pairs. ■ If one of the drives in a drive-pair fails, the system can instantly switch to the other drive without any loss of data or service. ■ RAID Level 1 and RAID Level 10 use drive mirroring to make an exact copy from one drive to another. ■ A single drive failure causes associated volumes to become degraded, but the mirror drive allows access to the data. ■ Two or more drive failures in a pool causes the volumes associated with the pool to fail, and data loss will occur.

RAID Level	Short Description	Detailed Description
RAID Level 3	High-bandwidth mode	<ul style="list-style-type: none"> ■ A minimum of three drives is required for RAID Level 3. ■ RAID Level 3 is limited to a maximum of 30 drives in a pool. ■ RAID Level 3 stripes both user data and redundancy data (parity) across the drives. ■ RAID Level 3 uses the equivalent of the capacity of one drive (in a pool) for redundancy data. ■ RAID Level 3 is used for applications with large data transfers, such as multimedia or medical imaging that write and read large sequential chunks of data. ■ A single drive failure in a pool causes the associated volumes to become degraded, but the redundancy data allows access to the data. ■ Two or more drive failures in a pool causes the volumes associated with the pool to fail, and data loss will occur.
RAID Level 5	High I/O mode	<ul style="list-style-type: none"> ■ A minimum of three drives is required for RAID Level 5. ■ RAID Level 5 is limited to a maximum of 30 drives in a pool. ■ RAID Level 5 stripes both user data and redundancy data (parity) across the drives. ■ RAID Level 5 uses the equivalent of the capacity of one drive (in a pool) for redundancy data. ■ A single drive failure in a pool causes associated volumes to become degraded, but the redundancy data allows access to the data. ■ Two or more drive failures in a pool causes the volumes associated with the pool to fail, and data loss will occur.
RAID Level 6	High I/O mode with simultaneous drive failure protection	<ul style="list-style-type: none"> ■ A minimum of five drives is required for RAID Level 6. ■ RAID Level 6 is limited to a maximum of 30 drives in a pool. ■ RAID Level 6 stripes both user data and redundancy data (parity) across the drives. ■ RAID Level 6 uses the equivalent of the capacity of two drives (in a pool) for redundancy data. ■ RAID Level 6 provides the best data availability. RAID Level 6 protects against the simultaneous failure of two pool member drives by using two independent error-correction schemes.

Dynamic RAID-Level Migration

Dynamic RAID-Level Migration (DRM) is a modification operation that lets you change the RAID level on a selected pool without impacting the I/O. You can continue to access data on pools, volumes, and drives during the migration process.

The pool must contain sufficient free space and the required number of drives, or the DRM request is rejected. You cannot cancel the DRM operation after the process begins.

NOTE If RAID Level 6 is a premium feature on your storage array, you must enable RAID Level 6 with a feature key file before migrating a pool to RAID Level 6.

Hardware Redundancy

Data-protection strategies provided by the storage array hardware include controller cache memory, hot spare drives, background media scans, and channel protection.

Controller Cache Memory

Write caching, or caching a drive segment to a memory buffer before writing to the drive, can increase I/O performance during data transfers.

Write-cache mirroring protects data during a controller-memory failure or a cache-memory failure. When you enable write cache, cached data is mirrored across two redundant controllers with the same cache size. Therefore, if one controller fails, the alternate controller can complete all outstanding write operations.

To prevent data loss or corruption, the controller periodically writes cache data to a drive (flushes the cache) when the amount of unwritten data in the cache reaches a certain level, called a start percentage, or when data has been in the cache for a predetermined amount of time. The controller continues to write data to a drive until the amount of data in the cache drops to a stop percentage level. You can configure the start percentage and the stop percentage to suit your own storage requirements. For example, you can specify that the controller start flushing the cache when it reaches 80-percent full and stop flushing the cache when it reaches 16-percent full.

In case of power outages, data in the controller cache memory is protected. Controller modules and array modules contain batteries that protect the data in the cache by maintaining a level of power until the data can be written to the drive media or a flash memory card.

If the controller supports a flash memory card, the cache data can be written to the flash memory card when a power outage occurs. For example, the ST2500 M2 array module supports a flash memory card to write the cache data. The battery is only needed to maintain power while the data in the cache is written to the flash memory card. The flash memory card provides nonvolatile backup of the cache data in case of long power outages. When power is restored to the controllers, the cache data can be read from the flash memory card.

If a power outage occurs when there is no UPS, and there is no battery or the battery is damaged, the data in the cache that has not been written to the drive media is lost. This situation occurs even if the data is mirrored to the cache memory of both controllers. It is, therefore, important to change the batteries in the controller module and the array module at the recommended time intervals.

Module Loss Protection

When you create a pool using the module loss protection feature, all of the drives in the pool are found in different drive modules. Module loss protection provides more data protection if access to the module is lost. This feature is used by default when you choose the automatic configuration option.

Module loss protection depends on the number of modules that are available, the value set for the Redundant Array of Independent Disks (RAID) level, and the number of drives in the pool. For example, module loss protection cannot be achieved if a RAID Level 5 pool is comprised of eight drives, but there are only three modules. Configuring your pools to have module loss protection is recommended. If your configuration supports the minimum number of drive modules for your RAID level, create your pools to have module loss protection.

RAID Level	Criteria for Module Loss Protection
RAID Level 0	No module loss protection (RAID Level 0 does <i>not</i> provide redundancy).
RAID Level 1 or RAID Level 10	For RAID Level 1, the pool must use a minimum of two drives found in separate modules. For RAID Level 10, the pool must use a minimum of four drives found in separate modules.
RAID Level 3	The pool must use a minimum of three drives found in separate modules.
RAID Level 5	The pool must use a minimum of three drives found in separate modules.
RAID Level 6	The pool must use a minimum of five drives, with a maximum of two drives in any module.

Hot Spare Drives

A valuable strategy to protect data is to assign available drives in the storage array as hot spare drives. A hot spare is a drive, containing no data, that acts as a standby in the storage array in case a drive fails in a RAID Level 1, RAID Level 3, RAID Level 5, RAID Level 6, or RAID Level 10 pool. The hot spare adds another level of redundancy to the storage array. Generally, hot spare drives must have capacities that are equal to or greater than the used capacity on the drives that they are protecting. Hot spare drives must be of the same media type and same interface type as the drives that they are protecting.

If a drive fails in the storage array, the hot spare can be substituted automatically for the failed drive without requiring your intervention. If a hot spare is available when a drive fails, the controller uses redundancy data to reconstruct the data onto the hot spare. After the failed drive is physically replaced, you can use either of the following options to restore the data:

- When you have replaced the failed drive, the data from the hot spare is copied back to the replacement drive. This action is called *copyback*.
- You can assign the hot spare as a permanent member of the pool. Performing the copyback function is not required for this option.

The availability of module loss protection for a pool depends on the location of the drives that comprise the pool. Module loss protection might be lost because of a failed drive and the location of the hot spare drive. To make sure that module loss protection is not affected, you must replace a failed drive to initiate the copyback process.

The storage array automatically selects T10 Protection Information (PI) capable drives for hot spare coverage of PI-enabled volumes. Make sure to have PI-capable drives in the storage array for hot spare coverage of PI-enabled volumes.

Security capable drives provide coverage for both security capable and non-security capable drives. Non-security capable drives can provide coverage only for other non-security capable drives.

If you do not have a hot spare, you can still replace a failed drive while the storage array is operating. If the drive is part of a RAID Level 1, RAID Level 3, RAID Level 5, RAID Level 6, or RAID Level 10 pool, the controller uses redundancy data to automatically reconstruct the data onto the replacement drive. This action is called *reconstruction*.

Channel Protection

In a Fibre Channel environment, channel protection is usually present for any storage array. When the storage array is cabled correctly, two redundant arbitrated loops (ALs) exist for each drive.

I/O Data Path Protection

Input/output (I/O) data path protection to a redundant *controller* in a *storage array* is accomplished with these multi-path drivers:

- The *Auto-Volume Transfer (AVT)* feature and the *Multi-Path I/O (MPIO)* driver in the Windows operating system (OS).
- The Multi-Path Proxy (MPP) -based *Redundant Dual Active Controller (RDAC)* multi-path driver in the Linux OS.
- The Multi-Plexed I/O (MPxIO) driver in the Solaris OS.
- The native failover driver using Target Port Group Support (TPGS) in the HP-UX OS version 11.31.
- The native failover driver in the VMware OS.
- The native failover driver using TPGS in the Mac OS X.

AVT is a built-in feature of the controller *firmware* that permits ownership of a volume to be transferred to a second controller if the preferred controller fails. When you use AVT with a multi-path driver, AVT helps to make sure that an I/O data path is always available for the *volumes* in the storage array.

If a component, such as a controller, a cable, or an *environmental services monitor (ESM)*, fails, or an error occurs on the data path to the preferred controller, AVT and the multi-path driver automatically transfer the *pools* and volumes to the alternate “non-preferred” controller for processing. This failure or error is called a *failover*.

Multi-path drivers, such as MPIO, RDAC, and MPxIO, are installed on host computers that access the storage array and provide I/O path failover. The AVT feature is used specifically for single-port cluster failover. The AVT feature mode is automatically selected by the host type.

Multi-Path Driver with AVT Enabled

Enabling AVT in your storage array and using it with a host multi-path driver helps to make sure that an I/O data path is always available for the storage array volumes.

When you create a volume in a storage array where AVT is enabled, a controller must be assigned to own the volume, called the *preferred owner*. The preferred controller normally receives the I/O requests to the volume. If a problem along the data path, such as a component failure, causes an I/O request to fail, the multi-path driver sends the I/O to the alternate controller.

NOTE You should have the multi-path driver installed at all times. You should always enable the AVT mode. Set the AVT mode to a single port cluster host type.

After the I/O data path problem is corrected, the preferred controller automatically re-establishes ownership of the volume as soon as the multi-path driver detects that the path is normal again.

Multi-Path Driver with AVT Disabled

When you disable AVT in your storage array, the I/O data path is still protected as long as a multi-path driver is installed on each host that is connected to the storage array. However, when an I/O request is sent to a specified volume, and a problem occurs along the data path to its preferred controller, all volumes on the preferred controller are transferred to the alternate controller, not just the specified volume.

Target Port Group Support

Target Port Group Support (TPGS) is another multi-path driver that is available on specific combinations of operating systems and failover drivers that can be present on a host. TPGS provides failover for a storage array. Failover is an automatic operation that switches the data path for a volume from the preferred controller to the alternate controller in the case of a hardware failure.

TPGS is part of the ANSI T10 SPC-3 specification. It is implemented in the controller firmware. TPGS is similar to other multi-pathing options, such as Auto-Volume Transfer (AVT) and Redundant Dual Active Controller (RDAC), which were developed prior to defining a multi-pathing standard. The advantage of TPGS is that it is based on the current standard, which allows interoperability with multi-pathing solutions from other vendors. Interoperability with other multi-pathing solutions simplifies administration of the host.

Each host type uses only one of the multi-path methods: RDAC, AVT, or TPGS.

Load Balancing

Load balancing is the redistribution of read/write requests to maximize throughput between the server and the storage array. Load balancing is very important in high workload settings or other settings where consistent service levels are critical. The multi-path driver transparently balances I/O workload without administrator intervention. Without multi-path software, a server sending I/O requests down several paths might operate with very heavy workloads on some paths, while other paths are not used efficiently.

The multi-path driver determines which paths to a device are in an active state and can be used for load balancing. The load-balancing policy uses one of three algorithms: round robin, least queue depth, or least path weight. Multiple options for setting the load-balancing policies let you optimize I/O performance when mixed host interfaces are configured. The load-balancing policies that you can choose depend on your operating system. Load balancing is performed on multiple paths to the same controller, but not across both controllers.

Operating System	Multi-Path Driver	Load Balancing Policy
Windows	MPIO DSM	Round robin, least queue depth, least path weight
Red Hat Enterprise Linux (RHEL)	RDAC	Round robin, least queue depth
SUSE Linux Enterprise (SLES)	RDAC	Round robin, least queue depth
Solaris	MPxIO	Round robin

Round Robin with Subset

The round-robin with subset I/O load-balancing policy routes I/O requests, in rotation, to each available data path to the controller that owns the volumes. This policy treats all paths to the controller that owns the volume equally for I/O activity. Paths to the secondary controller are ignored until ownership changes. The basic assumption for the round-robin policy is that the data paths are equal. With mixed-host support, the data paths might have different bandwidths or different data transfer speeds.

Least Queue Depth with Subset

The least queue depth with subset policy is also known as the least I/Os policy or the least requests policy. This policy routes the next I/O request to the data path on the controller that owns the volume that has the least outstanding I/O requests queued. For this policy, an I/O request is a command in the queue. The type of command or the number of blocks that are associated with the command is not considered. The least queue depth with subset policy treats large block requests and small block requests equally. The data path selected is one of the paths in the path group of the controller that owns the volume.

Least Path Weight with Subset

The least path weight with subset policy assigns a weight factor to each data path to a volume. An I/O request is routed to the path with the lowest weight value to the controller that owns the volume. If more than one data path to the volume has the same weight value, the round-robin with subset path selection policy is used to route I/O requests between the paths with the same weight value.

Chapter 2: Introducing the Storage Management Software

The topics in this section describe the basic layout of the SANtricity ES Storage Manager software. The SANtricity ES Storage Manager software has two windows that provide management functionality and a graphical representation of your storage array: the Enterprise Management Window (EMW) and the Array Management Window (AMW).

NOTE The SANtricity ES Storage Manager software is also referred to as the storage management software.

In general, you will use the following process when using the storage management software. You use the EMW to add the storage arrays that you want to manage and monitor. Through the EMW, you also receive alert notifications of errors that affect the storage arrays. If you are notified in the EMW that a storage array has a non-Optimal status, you can start the AMW for the affected storage array to show detailed information about the storage array condition.

NOTE Depending on your version of storage management software, the views, menu options, and functionality might differ from the information presented in this section. For information about available functionality, refer to the online help topics that are supplied with your version of the storage management software.

Enterprise Management Window

The Enterprise Management Window (EMW) is the first window to appear when you start the storage management software. The EMW lets you perform these management tasks:

- Discover hosts and storage arrays automatically on your local sub-network.
- Manually add and remove hosts and storage arrays.
- Monitor the health of the storage arrays and report a high-level status by using the applicable icon.
- Configure alert notifications through email or Simple Network Management Protocol (SNMP) and report events to the configured alert destinations.
- Launch the applicable Array Management Window (AMW) for a selected storage array to perform detailed configuration and management operations.
- Run scripts to perform batch management tasks on a particular storage array. For example, scripts might be run to create new volumes or to download new controller firmware. For more information on running scripts, refer to the online help topics in the EMW.
- Upgrade the controller firmware.

A local configuration file stores all of the information about storage arrays that you have added and any email destinations or SNMP traps that you have configured.

Parts of the Enterprise Management Window

The Enterprise Management Window (EMW) has these areas that provide options for managing your storage array.

Part	Description
Title bar	"Enterprise Management" in the title bar text indicates that this is the EMW.
Menu bar	The menu bar contains various options to manage the storage arrays. For more information about menu bar options, refer to the EMW Menu Bar Options online help topic in the Enterprise Management Window of SANtricity ES Storage Manager.
Toolbar	The toolbar contains icons that are shortcuts to common commands. To show the toolbar, select View >> Toolbar .
Tabs	The EMW contains two tabs: <ul style="list-style-type: none"> ■ Devices – Shows the discovered storage arrays and their status and also shows unidentified storage arrays. ■ Setup – Allows you to perform initial setup tasks with the storage management software.
Status bar	The Status bar shows a summary of the health of your storage arrays, messages, and a progress bar. To show the Status bar, select View >> Status Bar .

EMW Devices Tab

The **Devices** tab in the EMW presents two views of the storage arrays that are managed by the storage management station:

- Tree view
- Table view

Tree View

The Tree view provides a tree-structured view of the nodes in the storage system. The Tree view shows two types of nodes:

- Discovered Storage Arrays
- Unidentified Storage Arrays

Both the Discovered Storage Arrays node and the Unidentified Storage Arrays node are child nodes of the storage management station node.

The Discovered Storage Arrays node has child nodes that represent the storage arrays that are currently managed by the storage management station. Each storage array is labeled with its machine name and is always present in the Tree view. When storage arrays and hosts with attached storage arrays are added to the EMW, the storage arrays become child nodes of the Discovered Storage Arrays node.

NOTE If you move the mouse over the storage array node, a tooltip shows the controller's IP address.

The Unidentified Storage Arrays node shows storage arrays that the storage management station cannot access because the name or IP address does not exist.

You can perform these actions on the nodes in the Tree view:

- Double-click the storage management station node and the Discovered Storage Arrays node to expand or collapse the view of the child nodes.
- Double-click a storage array node to launch the Array Management Window for that storage array.
- Right-click a node to open a pop-up menu that contains the applicable actions for that node.

The right-click menu for the Discovered Storage Arrays node contains these options:

- **Add Storage Array**
- **Automatic Discovery**
- **Refresh**

These options are the same as the options in the **Tools** menu. For more information, refer to the online help topics in the Enterprise Management Window.

Table View

Each managed storage array is represented by a single row in the Table view. The columns in the Table view show data about the managed storage array.

Column	Description
Name	The name of the managed storage array. If the managed storage array is unnamed, the default name is <code>Unnamed</code> .
Type	The type of managed storage array. This type is represented by an icon.
Status	An icon and a text label that report the status of the managed storage array.
Management Connections	<p>Out-of-Band – This storage array is an out-of-band storage array.</p> <p>In-Band – This storage array is an in-band storage array that is managed through a single host.</p> <p>Out-of-Band, In-Band – This storage array is a storage array that is both out-of-band and in-band. Click Details to see more information about any of these connections.</p>
Comment	Any comments that you have entered about the specific managed storage array.

Sort the rows in the Table view in ascending order or descending order by either clicking a column heading or by selecting one of these commands:

- **View >> By Name**
- **View >> By Status**
- **View >> By Management Connection**
- **View >> By Comment**

Showing Managed Storage Arrays in the Table View

You can change the way that *managed storage arrays* appear in the Table view.

- Select the storage management station node to show all of the known managed storage arrays in the Table view.
- Select a Discovered Storage Array node or an Undiscovered Storage Array node in the Tree view to show any storage arrays that are attached to that specific host in the Table view.

NOTE If you have not added any storage arrays, the Table view is empty.

- Select a storage array node in the Tree view to show only that storage array in the Table view.

NOTE Selecting an Unidentified node in the Tree view shows an empty Table view.


EMW Setup Tab

The EMW Setup tab is a gateway to tasks that you can perform when you set up a storage array. Using the EMW Setup tab, you can perform these tasks:


- Add a storage array
- Name or rename a storage array
- Configure an alert
- Manage a storage array by launching the Array Management Window (AMW)
- Upgrade the controller firmware
- Open the Inherit Systems Settings window

Adding and Removing a Storage Array

You can add a storage array by using these methods in the storage management software.

Location	Procedure
Tree view	Right-click the root node from the Tree view, and select Add Storage Array from the pop-up menu.
Toolbar	Click the icon to add the storage array. 
Edit menu	Select Edit >> Add Storage Array .
Setup tab	Select Add Storage Array .

You can remove a storage array by using these methods, which remove only the icon from the view without physically deleting the storage array. You can select more than one storage array to delete at a time.

Location	Procedure
Tree view	Right-click the storage array that you want to remove from the Tree view, and select Remove >> Storage Array from the pop-up menu.
Toolbar	Select the storage array that you want to remove from the Tree view or Table view, and click the icon to remove the storage array. 
Edit menu	Select the storage array that you want to remove from the Tree view or Table view, and select Edit >> Remove >> Storage Array .

Array Management Window

The Array Management Window (AMW) is a Java™ technology-based software that is launched from the Enterprise Management Window (EMW). The AMW provides management functions for a single storage array. You can have more than one AMW open at the same time to manage different storage arrays. The AMW includes these management functions for a *storage array*:

- Provides storage array options, such as locating a storage array, configuring a storage array, renaming a storage array, or changing a password.
- Provides the ability to configure *volumes* from your storage array capacity, define *hosts* and *host groups*, and grant host or host group access to sets of volumes called *storage domains*.
- Monitors the health of storage array components and reports a detailed status using applicable icons.
- Provides you with the applicable recovery procedures for a failed logical component or a failed hardware component.
- Presents a view of the *Event Log* for the storage array.
- Presents profile information about hardware components, such as *controllers* and *drives*.
- Provides controller management options, such as changing ownership of volumes or placing a controller online or offline.
- Provides drive management options, such as assigning hot spares and locating the drive.
- Monitors storage array performance.

Starting the Array Management Window

To start the Array Management Window (AMW) from the Enterprise Management Window (EMW), perform one of these tasks:

- Click the **Devices** tab, and double-click the name of the storage array that you want to manage.
- Click the **Devices** tab, right-click the name of the storage array you want to manage, and select **Manage Storage Array**.
- Click the **Devices** tab, and select **Tools >> Manage Storage Array**.
- Click the **Setup** tab, and select **Manage a Storage Array**. In the Select Storage Array dialog, select the name of the storage array that you want to manage, and click **OK**.

Summary Tab

The **Summary** tab in the AMW shows information about the storage array. Links to the Storage Array Profile dialog, relevant online help topics, and the storage concepts tutorial also appear. Additionally, the link to the Recovery Guru dialog appears when the storage array needs attention.

In the **Summary** tab, you can view this information:

- The status of the storage array
- The hardware components in the storage array
- The capacity of the storage array
- The hosts, the mappings, and the storage domains in the storage array
- The pools and volumes in the storage array

Logical Tab

The **Logical** tab in the AMW contains two panes: the Logical pane and the Properties pane.

NOTE You can resize either pane by dragging the splitter bar, located between the two panes, to the right or to the left.

Logical Pane

The Logical pane provides a tree-structured view of the logical nodes. Click the plus (+) sign or the minus (-) sign adjacent to a node to expand or collapse the view. You can right-click a node to open a pop-up menu that contains the applicable actions for that node.

Nodes in the Logical Pane

The *storage array*, or root node, has three types of child nodes.

Child Nodes of the Root Node	Description of the Child Nodes
Unconfigured Capacity	This node represents the storage array capacity that is not configured into a pool.
Pool	This node has two types of child nodes: <ul style="list-style-type: none"> ■ Volume – This node represents a configured and defined volume. Multiple Volume nodes can exist under a Pool node. ■ Free Capacity – This node represents a region of capacity that you can use to create one or more new volumes within the pool. Multiple Free Capacity nodes can exist under a Pool node.

NOTE Multiple Unconfigured Capacity nodes appear if your storage array contains drives with different media types (hard drive or Solid State Disk [SSD]) and different interface types. Each drive type has an associated Unconfigured Capacity node shown under the Total Unconfigured Capacity node if *unassigned drives* are available in the drive module.

Types of Volumes

These types of volumes appear under the Pool node:

- *Standard volumes* is the basic structure that you create in the storage array to store data. A volume is configured from a pool with a specific RAID level to meet the software application's needs for data availability and I/O performance. The operating system sees a volume as one drive.
- *Primary volumes* that participate in a mirror relationship in the primary role. Primary volumes are standard volumes with a synchronized mirror relationship. The remote *secondary volume* that is associated with the primary volume appears as a child node.
- *Secondary volumes* appear directly under the Pool node when the local storage array contains this volume.
- *Mirror repository volumes* are special volumes in the storage array that are created as a resource for each controller in both local storage arrays and remote storage arrays. The controller stores duplicate information on the mirror repository volume, including information about remote writes that are not yet written to the secondary volume.
- *Snapshot repository volumes* is a volume in the storage array that is used as a resource for a snapshot volume.
- *Snapshot volumes* are child nodes of their associated *base volume*.
- *Source volumes* are standard volumes that participate in a volume copy relationship. Source volumes are used as the copy source for a target volume. Source volumes accept host I/O requests and store application data. A source volume can be a standard volume, a snapshot volume, a snapshot base volume, or a Data Replicator Software primary volume.

- *Target volumes* are standard volumes that participate in a volume copy relationship and contain a copy of the data from the source volume. Target volumes are read-only and do not accept write requests. A target volume can be created from a standard volume, the base volume of a snapshot volume, or a Remote Volume Mirror primary volume. The volume copy overwrites any existing volume data if an existing volume is used as a target.

Properties Pane

The Properties pane provides detailed information about the component selected in the Logical pane. The information varies depending on what type of component is selected.

You can view the physical components that are associated with a logical component by selecting the **Logical** tab, right-clicking a component, and selecting **View Associated Physical Components**.

Physical Tab

The **Physical** tab in the AMW contains two panes: the Physical pane and the Properties pane.

NOTE You can resize either pane by dragging the splitter bar, located between the two panes, to the right or to the left.




The Physical pane provides a view of the hardware components in a storage array, including their status. You can right-click a hardware component to open a pop-up menu that contains the applicable actions for that component.

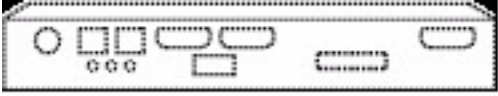


NOTE The orientation of the Physical pane is determined by the actual layout of the storage array. For example, if the storage array has horizontal drive modules, the storage management software shows horizontal drive modules in the Physical pane.

The Properties pane provides information for the hardware component that is selected in the Physical pane. The information in the Properties pane is specific to each hardware component. If you select a controller icon in the Physical pane, a list of properties for that controller is shown in the Properties pane. If you select drive icon in the Physical pane, a list of properties for that drive is shown in the Properties pane.

Controller Status

The status of each controller is indicated by an icon in the Physical pane. This table describes the various controller icons. Depending on your hardware model, the icons might differ from the icons shown in this table.

Icon	Status
	Online, Optimal
	Offline
	Service Mode



Icon	Status
	Slot Empty
	Needs Attention (if applicable for your hardware model)
	Suspended (if applicable for your hardware model)

View Module Components


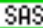



The **View Module Components** command on each module shows the status of the secondary components within the module, such as power supplies, fans, and temperature sensors.

Drive Modules

For each drive module that is attached to the storage array, a drive module appears in the Physical pane. If your storage array contains different media types or different interface types, a drive type icon appears to indicate the type of drives in the drive module. This table describes the different drive type icons that might appear.

Icon	Status
	This drive module contains only hard drives.
	This drive module contains only Solid State Disks (SSDs).

This table describes the different drive interface type icons that might appear.

Icon	Status
	This drive module contains only Encryption Services (ES) security capable drives.
	This drive module contains only Serial Attached SCSI (SAS) drives.
	This drive module contains only Fibre Channel (FC) drives.
	This drive module contains only Serial ATA (SATA) drives.
	This drive module contains only T10 Protection Information (PI) capable drives.

You can click **Show** in the Physical pane to view where a specific drive type is located in the drive module.

Mappings Tab

The **Mappings** tab in the AMW contains two panes: the Topology pane and the Defined Mappings pane.

NOTE You can resize either pane by dragging the splitter bar, located between the two panes, to the right or to the left.

Topology Pane

The Topology pane shows a tree-structured view of logical nodes that are related to *storage domains*. Click the plus (+) sign or the minus (-) sign adjacent to a node to expand or collapse the view. You can right-click a node to open a pop-up menu that contains the applicable actions for that node.

Nodes in the Topology Pane

The storage array, or the root node, has these types of child nodes.

Child Nodes of the Root Node	Description of the Child Nodes
Undefined Mappings	The Undefined Mapping node has one type of child node. Individual Undefined Mapping – Represents a <i>volume</i> with an undefined mapping. Multiple Volume nodes can exist under an <i>Undefined Mappings</i> node.
Default Group	NOTE If Storage Domains is disabled, all of the created volumes are in the Default Group. A Default Group node has two types of child nodes: <ul style="list-style-type: none"> ■ Host Group – Defined host groups that are not participating in specific mappings are listed. This node can have host child nodes, which can have child host port nodes. ■ Host – Defined hosts that are not part of a specific host group but are part of the Default Group and are not participating in specific mappings are listed. This node can have child host port nodes.
Unassociated Host Port Identifier	An Unassociated Host Port Identifier node has one type of child node. Host Port Identifier – Host port identifier that has not been associated with any host.
Host Group	A Host Group node has one type of child node. Host – Defined hosts that belong to this defined host group are listed. This node can have child host port nodes. NOTE The host nodes that are child nodes of this host group can also participate in mappings specific to the individual host rather than the host group.
Host	A Host node has one type of child node. Host Port – This node has child nodes that represent all of the host ports or single ports on a host adapter that are associated with this host.

Storage Domain Icon

The storage domain icon, when present in the Topology pane, indicates that a storage domain has been defined for a host group, or a host. This icon also appears in the status bar when storage domains have been defined.



Defined Mappings Pane

The Defined Mappings pane shows the mappings associated with a node selected in the Topology pane.

The information in the table appears for a selected node.

Column Name	Description
Volume Name	The user-supplied volume name. The factory-configured <i>access volume</i> also appears in this column. NOTE An access volume mapping is <i>not</i> required for a storage array with an <i>in-band</i> connection and can be removed.
Accessible By	Shows the Default Group, a defined host group, or a defined host that has been granted access to the volume in the mapping.
LUN	The LUN assigned to the specific volume that the host or hosts use to access the volume.
Volume Capacity	Shows the volume capacity in units of GB.
Type	Indicates whether the volume is a standard volume or a snapshot volume.

You can right-click a volume name in the Defined Mappings pane to open a pop-up menu. The pop-up menu contains options to change and remove the mappings.

The information shown in the Defined Mappings pane varies according to what node you select in the Topology pane, as shown in this table.

Node Selected	Information That Appears in the Defined Mappings Pane
Root (storage array) node	All defined mappings.
Default Group node or any child node of the Default Group	All mappings that are currently defined for the Default Group (if any).
Host Group node (outside of Default Group)	All mappings that are currently defined for the Host Group.
Host node that is a child node of a Host Group node	All mappings that are currently defined for the Host Group, plus any mappings specifically defined for a specific host.
Host Port node or individual host port node outside of the Default Group	All mappings that are currently defined for the host port's associated host.

AMW Setup Tab

The AMW **Setup** tab provides links to these tasks:

- Locating the storage array
- Renaming the storage array
- Setting a storage array password
- Configuring the network parameters for the iSCSI host ports
- Configuring the storage array
- Mapping volumes to hosts
- Saving configuration parameters in a file
- Defining the hosts and host ports
- Configuring the Ethernet management ports
- Viewing and enabling the premium features

- Managing the additional iSCSI settings for authentication, identification, and discovery

The iSCSI options are shown in the AMW **Setup** tab only when the controllers contain iSCSI host ports.

Support Tab

The **Support** tab in the AMW provides links to these tasks:

- Recovering from a storage array failure by using the Recovery Guru
- Gathering support information, such as the *Event Log* and a description of the storage array, to send to your Sun Customer Care Center representative
- Viewing the description of all components and properties of the storage array
- Downloading the controller *firmware*, the *NVSRAM*, the drive firmware, the *ESM* firmware, and the ESM configuration settings
- Viewing the Event Log of the storage array
- Viewing the online help topics
- Viewing the version and copyright information of the storage management software

You can click a link to open the corresponding dialog.

Managing Multiple Software Versions

When you open the Array Management Window (AMW) to manage a storage array, the version of software that is appropriate for the version of firmware that the storage array uses is opened. For example, you manage two storage arrays using this software; one storage array has firmware version 6.14, and the other has firmware version 7.7x, where x represents a number. When you open the AMW for a particular storage array, the correct AMW version is used. The storage array with firmware version 6.14 uses version 9.14 of the storage management software, and the storage array with firmware version 7.7x uses version 10.7x of the storage management software. You can verify the version that you are currently using by selecting **Help >> About** in the AMW.

This bundling of previous versions of the AMW provides the flexibility of upgrading the firmware only on selected storage arrays instead of having to perform an upgrade on all of the storage arrays at one time.

Chapter 3: Configuring the Storage Arrays

The topics in this section describe the methods for configuring storage arrays, including managing security, and premium features.

For additional information and detailed procedures for the options described in this section, refer to the online help topics in SANtricity ES Storage Manager.

Volumes and Pools

When you configure a storage array for the first time, you must consider which data protection strategy is most appropriate for your storage array, together with how the total storage capacity must be organized into volumes and shared among hosts.

The storage management software identifies several distinct volumes:

- Standard volumes
- Snapshot volumes
- Snapshot repository volumes
- Primary volumes
- Secondary volumes
- Mirror repository volumes
- Source volumes
- Target volumes

Standard Volumes

A standard volume is a logical structure that is created on a storage array for data storage. A standard volume is defined from a set of drives called a pool, which has a defined RAID level and capacity. You can create a volume from unconfigured capacity, unassigned drives, or Free Capacity nodes on the storage array. If you have not configured any volumes on the storage array, the only node that is available is the Unconfigured Capacity node.

Use the Create Volume Wizard to create one or more volumes on the storage array. During the volume creation process, the wizard prompts you to select the capacity to allocate for the volumes and to define basic volume parameters and optional advanced volume parameters for the volume.

NOTE The host operating system might have specific limits about how many volumes that the host can access. You must consider these limits when you create volumes that are used by a particular host.

Storage Array	Maximum Number of Volumes per Storage Array	Maximum Number of Volumes per Storage Domain
6580/6780 controller module	Up to 2048	Up to 256
ST2500 M2 array module	Up to 512	Up to 256
6180 array module	Up to 1024	Up to 256

Pools

A pool is a set of drives that the controller logically groups together to provide one or more volumes to an application host. All of the drives in a pool must have the same media type and interface type.

To create a pool, you must specify two key parameters: the RAID level and the capacity (how large you want the pool to be). You can either select the automatic choices provided by the software or select the manual method to indicate the specific drives to include in the pool. Whenever possible, use the automatic method because the storage management software provides the best selections for drive groupings.


Pool Creation

The Create Pool Wizard guides you through the steps to create one or more pools in a storage array and to configure basic pool parameters and optional pool parameters.

NOTE The storage management software determines the default initial capacity selections based on whether you select free capacity, unconfigured capacity, or unassigned drives in the Create Pool Wizard. After the wizard begins, you can change the capacity by defining a new volume capacity.

You can organize available capacity on a storage array by using these types of storage spaces:

- **Free capacity** – Free capacity is unassigned space in a pool that you can use to create a volume. When you create a volume from free capacity, an additional volume is created on an existing pool.
- **Unconfigured capacity** – Unconfigured capacity is available space on drives of a storage array that has not been assigned to a pool. One unconfigured capacity node exists for each type of drive media and drive interface.
- **Unassigned drive** – An unassigned drive is a drive that is not being used in a pool or is not assigned as a hot spare.



- 1. Free Capacity
- 2. Pool
- 3. Volume
- 4. Volume
- 5. Volume
- 6. Hot Spare Drive
- 7. Unconfigured Capacity

Specifying Volume Parameters

Parameter	Free Capacity	Unconfigured Capacity	Unassigned Drive
Pool Creation	The pool is predefined.	You must create a pool before configuring a new volume.	You must create a pool before configuring a new volume.
Specify Capacity/Name Dialog	Assign a name to the volume. Change the default capacity.	Assign a name to the volume. Change the default capacity.	Assign a name to the volume. Change the default capacity.
Storage Domains will be used	Select the Map Later using the Mappings View option. This option specifies that a LUN not be assigned to the volume during volume creation. This option defines specific mappings and creates storage domains.	Select the Map Later using the Mappings View option. This option specifies that a LUN not be assigned to the volume during volume creation. This option defines specific mappings and creates storage domains.	Select the Map Later using the Mappings View option. This option specifies that a LUN not be assigned to the volume during volume creation. This option defines specific mappings and creates storage domains.
Storage Domains will not be used	Select the Default Mapping option. This option automatically assigns the next available LUN in the Default Group to the volume. The option grants volume access to host groups or hosts that have no specific mappings, which are shown under the Default Group node in the Topology pane.	Select the Default Mapping option. This option automatically assigns the next available LUN in the Default Group to the volume. The option grants volume access to host groups or hosts that have no specific mappings, which are shown under the Default Group node in the Topology pane.	Select the Default Mapping option. This option automatically assigns the next available LUN in the Default Group to the volume. The option grants volume access to host groups or hosts that have no specific mappings, which are shown under the Default Group node in the Topology pane.
Advanced Volume Parameters	You can customize these advanced volume parameters: <ul style="list-style-type: none"> ■ Volume I/O characteristics ■ Preferred controller owner 	You can customize these advanced volume parameters: <ul style="list-style-type: none"> ■ Volume I/O characteristics ■ Preferred controller owner 	You can customize these advanced volume parameters: <ul style="list-style-type: none"> ■ Volume I/O characteristics ■ Preferred controller owner

Dynamic Capacity Expansion

Dynamic Capacity Expansion (DCE) is a modification operation in the storage management software that increases the capacity of a pool. This modification operation allows you to add unassigned drives to a pool. Adding unassigned drives increases the free capacity in the pool. You can use this free capacity to create additional volumes.

This operation is considered to be dynamic because you have the ability to continually access data in the pool throughout the entire operation.

Keep these guidelines in mind when you add unassigned drives to a pool:

- The number of unassigned drives that you can select for a DCE modification operation is limited by the controller firmware. You can add two unassigned drives at a time. However, after you have completed a DCE operation, you can add more drives again until the desired capacity is reached.
- The existing volumes in the pool do not increase in size when you add unassigned drives to expand the free capacity. This operation redistributes existing volume capacity over the larger number of drives in the pool.
- The unassigned drives that you are adding to the pool must be of the same media type and interface type. Mixing different drive types within a single pool is not permitted. Whenever possible, select drives that have a capacity equal to the capacities of the current drives in the pool.
- In a RAID Level 1 pool, you must add two drives to make sure that data redundancy is configured.
- Only security capable drives can be added to a security enabled pool or a security capable pool.
- In a pool that is T10 Protection Information (PI) capable and contains a PI-enabled volume, you can add only PI-capable drives.

Register the Volume with the Operating System

After you have created all of your volumes and have assigned mappings, use a volume registration utility, such as the hot_add utility when using RDAC, to scan the mapped volumes and register the volumes with the operating system.

You can run the hot_add utility to make sure that the operating system is aware of the newly created volumes.

If available for your operating system, you can run the host-based SMdevices utility to associate the physical storage array name and the volume name.

Premium Features

The storage management software has the following premium features that provide data-protection strategies:

- Storage Domains
- Snapshot Volume
- Data Replicator Software (this premium feature is supported only in storage arrays with the Fibre Channel [FC] host ports)
- Volume Copy
- Drive Security and Enterprise Key Manager
- T10 Protection Information (PI)
- Solid State Disks (SSDs)

Storage Domains

Storage Domains lets hosts with different operating systems share access to a storage array. Hosts with different operating systems that share access to a storage array are called heterogeneous hosts.

A storage domain is a logical entity that consists of one or more storage array volumes that can be shared among hosts. To create a storage domain after the total storage capacity has been configured into volumes, you must define a single host or collection of hosts (or host group) that will access the storage array. Then you must define a mapping, which lets you specify the host group or the host that will have access to a particular volume in your storage array.

Based on the premium feature key file purchased, the storage management software can support the maximum storage domains shown in this table.

Storage Array	Maximum Number of Storage Domains Supported
6580/6780 controller module	Up to 512
ST2500 M2 array module	Up to 128
6180 array module	Up to 128

You can define a maximum of 256 volumes per partition (except for the HP-UX 11.23 operating system); this number is limited to the total number of volumes on your storage array.

Snapshot Volume Premium Feature

The Snapshot Volume premium feature creates a logical point-in-time image of another volume. Snapshot Volume is a premium feature of the storage management software. You or your storage vendor must enable this premium feature.

Because the only data blocks that are physically stored in the snapshot repository volume are those that have changed since the time that the snapshot volume was created, the snapshot volume uses less drive space than a full physical copy.

Typically, you create a snapshot so that an application (for example, a backup application) can access the snapshot and read the data; meanwhile, the base volume stays online and is user accessible. When the backup is completed, the snapshot volume is no longer needed.

You can also create snapshots of a base volume and write data to the snapshot volumes to perform testing and analysis. Before upgrading your database management system, for example, you can use snapshot volumes to test different configurations. Then you can use the performance data that is provided by the storage management software to help you decide how to configure your live database system. The maximum number of snapshots supported by the storage array is shown in this table.

Storage Array	Maximum Number of Snapshots per Volume	Maximum Number of Snapshots per Storage Array
6580/6780 controller module	Up to 16	Up to 1024
ST2500 M2 array module	Up to 16	Up to 256
6180 array module	Up to 8	Up to 512

Creating Snapshot Volumes

When a snapshot volume is created, the controller suspends I/O activity to the base volume for a few seconds while it creates a physical volume, called the snapshot repository volume. The snapshot repository volume stores the snapshot volume metadata and the copy-on-write data.

You can create snapshot volumes by using the Create Snapshot Volume Wizard in the Array Management Window. The first dialog of the Create Snapshot Volume Wizard lets you select either the simple path or the advanced path to be followed through the wizard. You can choose the simple path to create a snapshot volume if the pool of the base volume has the required amount of free capacity. The simple path lets you specify the basic parameters for the snapshot volume. The simple path accepts the default settings for the advanced parameters.

NOTE If sufficient free capacity is not available in the pool of the base volume, the Create Snapshot Volume Wizard uses the advanced path by default.

In the advanced path, either you can choose to place the snapshot repository volume in another pool, or you can use unconfigured capacity in the storage array to create a new pool. The advanced path lets you customize the advanced settings for the snapshot volume, such as the full conditions of the snapshot repository volume and the notification settings.

If you want to create a snapshot volume that performs snapshot operations at a later time or at regularly occurring intervals, specify a schedule. If you do not specify a schedule, the snapshot operation occurs immediately.

Scheduling Snapshots

If you want to create a snapshot volume that performs snapshot operations at a later time or at regularly occurring intervals, add a schedule to the snapshot volume. If you do not add a schedule to the snapshot volume, the snapshot operation occurs immediately. You can add a schedule when you create a snapshot volume, or you can add a schedule to an existing snapshot volume. Each snapshot volume can have only one schedule.

Typical Uses of Scheduling Snapshots

Scheduled backups – For example, an application stores business-critical data in two volumes in the storage array. You back up this data every work day at 11:00 p.m. To accomplish this type of backup, select the first volume. Create a schedule that runs once a day on Monday, Tuesday, Wednesday, Thursday, and Friday. Choose a time between the end of your work day and 11:00 p.m. Select a starting date of today and no end date. Apply this schedule to the second volume, also. Map the two snapshot volumes to your backup host, and perform the regular backup procedures. Unmap the two snapshot volumes before the next scheduled snapshot operation time. If you do not unmap the snapshot volumes, the storage array skips the next snapshot operation to avoid data corruption.

Rapid recovery – In this example, you back up your data at the end of every work day and keep hourly snapshots from 8:00 a.m. to 5:00 p.m. If data loss or corruption occurs during the work day, you can recover the data from the snapshots so that the data loss window is smaller than one hour. To accomplish this type of recovery, create a schedule that contains a start time of 8:00 a.m. and an end time of 5:00 p.m. Select 10 snapshots per day on Monday, Tuesday, Wednesday, Thursday, and Friday. Select a start date of today and no end date. Create an end-of-day backup as described in the "Scheduled backups" example.



Guidelines for Creating Schedules

Keep the following guidelines in mind when creating schedules for snapshot volumes:

- Either you can create a schedule when you create a snapshot volume, or you can add a schedule to an existing snapshot volume.
- Scheduled snapshot operations do not take place when these conditions occur:
 - The snapshot volume is mapped.
 - The storage array is offline or powered off.
 - The snapshot volume is used as a source volume in a Volume Copy operation, and the status of the copy operation is Pending or In progress.
- If you delete a snapshot volume that has a schedule, the schedule is also deleted.
- Schedules are stored in the configuration database in the storage array. The management station does not need to be running the Enterprise Management Window (EMW) or the Array Management Window (AMW) for the scheduled snapshot operation to occur.

Enabling and Disabling Schedules

You temporarily can suspend scheduled snapshot operations by disabling the schedule. When a schedule is disabled, the schedule's timer continues to run, but the scheduled snapshot operations do not occur. This table shows the icons for scheduled snapshots.

Icon	Description
	The schedule is enabled. Scheduled snapshots will occur.
	The schedule is disabled. Scheduled snapshots will not occur.

Discontinuing the Use of a Snapshot Volume

As long as a snapshot volume is enabled, storage array performance is affected by the copy-on-write activity to the associated snapshot repository volume. When you no longer need a snapshot volume, you can disable it, reuse it, or delete it.

- **Disable** – Stops copy-on-write activity. This option keeps the snapshot volume and snapshot repository volume intact.
- **Reuse** – Creates a different point-in-time image of the same base volume. This action takes less time to configure than re-creating the snapshot volume.

- **Delete** – Completely removes the snapshot volume and the associated snapshot repository volume. If you want to re-enable a snapshot volume, you must re-create it.

Disabling and Restarting Multiple Snapshots

If multiple volumes require regular snapshots for backup purposes, keeping the snapshots enabled might significantly affect storage array performance. In this situation, you can disable the snapshot function for multiple volumes and then restart the snapshots for all of the volumes before the next backup is scheduled.

The list of snapshots to be restarted is treated as a single operation. The new point-in-time snapshot images are created from the previously defined parameters. If an error is encountered on any of the listed snapshots, none of the snapshots on the list are re-created.

Dynamic Volume Expansion

NOTE Increasing the capacity of a standard volume is only supported on certain operating systems. If volume capacity is increased on a host operating system that is not supported, the expanded capacity is unusable, and you cannot restore the original volume capacity.

Dynamic Volume Expansion (DVE) is a modification operation that increases the capacity of standard volumes or snapshot repository volumes. The increase in capacity can be achieved by using any free capacity available on the pool of the standard volume or the snapshot repository volume. Data is accessible on pools, volumes, and drives throughout the entire modification operation.

If you receive a warning that the snapshot repository volume is in danger of becoming full, you can use the DVE modification operation to increase the capacity of the snapshot repository volume.

Increasing the capacity of a snapshot repository volume does not increase the capacity of the associated snapshot volume. The capacity of the snapshot volume is always based on the capacity of the base volume at the time that the snapshot volume was created.

Data Replicator Software Premium Feature

The Data Replicator Software premium feature is used for online, real-time data replication between storage arrays over a remote distance. Storage array controllers manage the mirroring, which is transparent to host machines and software applications. You create one or more mirrored volume pairs that consist of a primary volume at the primary site and a secondary volume at a secondary, remote site. After you create the mirror relationship between the two volumes, the current owner of the primary volume copies all of the data from the primary volume to the secondary volume. This process is called a full synchronization.

There is a base number of defined mirrors that are allowed for each storage array. You can increase the number of defined mirrors that are allowed per model with the purchase of an optional feature pack upgrade key. This table shows the maximum number of defined mirrors to which you can upgrade with a feature pack upgrade key.

Storage Array	Maximum Number of Defined Mirrors
6580/6780 controller module	Up to 128
ST2500 M2 array module	Up to 16
6180 array module	Up to 64

The Data Replicator Software premium feature is not supported in a simplex configuration. You must disable the Data Replicator Software premium feature before converting a storage array from a duplex configuration to a simplex configuration. The Data Replicator Software premium feature is supported only in storage arrays with the Fibre Channel (FC) host ports. The Data Replicator Software premium feature also requires a Fibre Channel network switch.

ATTENTION Possible loss of data access – You cannot create a mirror relationship if the primary volume contains unreadable sectors. Furthermore, if an unreadable sector is discovered during a mirroring operation, the mirror relationship fails.

NOTE Because replication is managed on a per-volume basis, you can mirror individual volumes in a primary storage array to appropriate secondary volumes in several *different* remote storage arrays.

Disaster Recovery

The secondary, remote volume is unavailable to secondary host applications while mirroring is in progress. In the event of a disaster at the primary site, you can fail over to the secondary site. To fail over, perform a role reversal to promote the secondary volume to a primary volume. Then the recovery host is able to access the newly promoted volume, and business operations can continue.

Data Replication

When the current owner of the primary volume receives a write request from a host, the controller first logs information about the write to a special volume. This volume is called a mirror repository volume. It writes the data to the primary volume. Next, the controller initiates a remote write operation to copy the affected data blocks to the secondary volume at the remote site.

Finally, the controller sends an I/O completion indication back to the host system to confirm that the data was copied successfully to the secondary storage array. The write mode that you selected when you first created a remote volume mirror determines when the I/O completion indication is sent to the host system.

The storage management software provides two write modes:

- **Synchronous** – When you select this write mode, any host write requests are written to the primary volume and then copied to the secondary storage volume. The controller sends an I/O completion indication to the host system *after* the copy has been successfully completed.
- **Asynchronous** – When you select this write mode, host write requests are written to the primary volume. Then the controller sends an I/O completion indication back to the host system *before* the data has been successfully copied to the secondary storage array.

When write caching is enabled on either the primary volume or the secondary volume, the I/O completion is sent when data is in the cache on the side (primary or secondary) where write caching is enabled. When write caching is disabled on either the primary volume or the secondary volume, the I/O completion is not sent until the data has been stored to physical media on that side.

Host write requests received by the controller are handled normally. No communication takes place between the primary storage array and the secondary storage array.

Link Interruptions or Secondary Volume Errors

When processing write requests, the primary controller might be able to write to the primary volume, but a link interruption prevents communication with the remote secondary controller.

In this case, the remote write cannot complete to the secondary volume. The primary volume and the secondary volume are no longer appropriately mirrored. The primary controller changes the mirrored pair into Unsynchronized status and sends an I/O completion to the primary host. The primary host can continue to write to the primary volume, but remote writes do not take place.

When connectivity is restored between the current owner of the primary volume and the current owner of the secondary volume, a full synchronization takes place. Only the blocks of data that have changed on the primary volume during the link interruption are copied to the secondary volume. The mirrored pair changes from an Unsynchronized state to Mirror Synchronization in Progress status.

The primary controller also marks the mirrored pair as Unsynchronized when a volume error on the secondary side prevents the remote write from completing. For example, an offline secondary volume or a failed secondary volume can cause the remote mirror to become unsynchronized. When the volume error is corrected (the secondary volume is placed online or is recovered to Optimal status), a full synchronization automatically begins. The mirrored pair then changes to Synchronization in Progress status.

Connectivity and Volume Ownership

A primary controller attempts to communicate only with its matching controller in the secondary storage array. For example, controller A in the primary storage array attempts communication only with controller A in the secondary storage array. The controller (A or B) that owns the primary volume determines the current owner of the secondary volume. If the primary volume is owned by controller A on the primary side, the secondary volume is owned by controller A on the secondary side. If primary controller A cannot communicate with secondary controller A, controller ownership changes do not take place.

The next remote write processed automatically triggers a matching ownership change on the secondary side if one of these conditions exists:

- When an I/O path error causes a volume ownership change on the primary side
- If the storage administrator changes the current owner of the primary volume

For example, a primary volume is owned by controller A, and then you change the controller owner to controller B. In this case, the next remote write changes the controller owner of the secondary volume from controller A to controller B. Because controller ownership changes on the secondary side are controlled by the primary side, they do not require any special intervention by the storage administrator.

Controller Resets and Storage Array Power Cycles

Sometimes a remote write is interrupted by a controller reset or a storage array power cycle before it can be written to the secondary volume. The storage array controller does not need to perform a full synchronization of the mirrored volume pair in this case. A controller reset causes a controller ownership change on the primary side from the preferred controller owner to the alternate controller in the storage array. When a remote write has been interrupted during a controller reset, the new controller owner on the primary side reads information stored in a log file in the mirror repository volume of the preferred controller owner. It then copies the affected data blocks from the primary volume to the secondary volume, eliminating the need for a full synchronization of the mirrored volumes.

Data Replicator Software Premium Feature Activation

Like other premium features, you enable the Data Replicator Software premium feature by purchasing a feature key file from your storage supplier. You must enable the premium feature on both the primary storage array and the secondary storage array.

Unlike other premium features, you also must *activate* the premium feature after you enable it. To activate the premium feature, use the Activate Data Replicator Software Wizard in the Array Management Window (AMW). Each controller in the storage array must have its own mirror repository volume for logging write information to recover from controller resets and other temporary interruptions. The Activate Data Replicator Software Wizard guides you to specify the placement of the two mirror repository volumes (on newly created free capacity or existing free capacity in the storage array).

After you activate the premium feature, one Fibre Channel (FC) host side I/O port on each controller is solely dedicated to Data Replicator Software operations. Host-initiated I/O operations are not accepted by the dedicated port. I/O requests received on this port are accepted only from remote controllers that are participating in Data Replicator Software operations with the controller.

Connectivity Requirements

You must attach dedicated Data Replicator Software ports to a Fibre Channel fabric environment. In addition, these ports must support the Directory Service interface and the Name Service.

You can use a fabric configuration that is dedicated solely to the Data Replicator Software ports on each controller. In this case, host systems can connect to the storage arrays using fabric, Fibre Channel Arbitrated Loop (FC-AL), or point-to-point configurations. These configurations are totally independent of the dedicated Data Replicator Software fabric.

Alternatively, you can use a single Fibre Channel fabric configuration for both the Data Replicator Software connectivity and for the host I/O paths to the controllers.

The maximum distance between the primary site and the secondary site is 10 km (6.2 miles), using single-mode fiber gigabit interface converters (GBICs) and optical long-wave GBICs.

Restrictions

These restrictions apply to mirrored volume candidates and storage array mirroring:

- RAID level, caching parameters, and segment size can be different on the two mirrored volumes.
- The secondary volume must be at least as large as the primary volume.
- The only type of volume that can participate in a mirroring relationship is a standard volume. Snapshot volumes cannot participate.
- You can create a snapshot volume by using either a primary volume or a secondary volume as the base volume.
- A primary volume can be a source volume or a target volume in a volume copy. A secondary volume cannot be a source volume or a target volume unless a role reversal was initiated after the copy has completed. If a role reversal is initiated during a Copy in Progress status, the copy fails and cannot be restarted.
- A given volume might participate in only *one* mirror relationship.

Volume Copy Premium Feature

ATTENTION Possible loss of data access – The volume copy operation overwrites existing data on the target volume and renders the volume read-only to hosts. This option fails all snapshot volumes that are associated with the target volume, if any exist.

The Volume Copy premium feature copies data from one volume (the source) to another volume (the target) in a single storage array.

Use the Volume Copy premium feature to perform these tasks:

- Copy data from pools that use smaller capacity drives to pools that use larger capacity drives.
- Create an online copy of data from a volume within a storage array, while still being able to write to the volume with the copy in progress.
- Back up data or restore snapshot volume data to the base volume.

Volume Copy is a premium feature of the storage management software and must be enabled either by you or your storage vendor.

Storage Array	Maximum Number of Volume Copies per Storage Array
6580/6780 controller module	Up to 2047
ST2500 M2 array module	Up to 511
6180 array module	Up to 1023

Volume Copy Features

Data Copying for Greater Access

As your storage requirements for a volume change, use the Volume Copy premium feature to copy data to a volume in a pool that uses larger capacity drives within the same storage array. This premium feature lets you perform these functions:

- Move data to larger drives; for example, 73 GB to 146 GB
- Change to drives with a higher data transfer rate; for example, 2 Gb/s to 4 Gb/s
- Change to drives using new technologies for higher performance

Data Backup

The Volume Copy premium feature lets you back up a volume by copying data from one volume to another volume in the same storage array. You can use the target volume as a backup for the source volume, for system testing, or to back up to another device, such as a tape drive.

Snapshot Volume Data Restoration to the Base Volume

If you need to restore data to the base volume from its associated snapshot volume, use the Volume Copy premium feature to copy data from the snapshot volume to the base volume. You can create a volume copy of the data on the snapshot volume, and then copy the data to the base volume.

ATTENTION Possible loss of data – If you are using the Windows 2000 operating system or the Linux operating system, use the Volume Copy premium feature with the Snapshot Volume premium feature to restore snapshot volume data to the base volume. Otherwise, the source volume and the target volume can become inaccessible to the host.

Types of Volume Copies

You can perform either an *offline* volume copy or an *online* volume copy. To ensure data integrity, all I/O to the target volume is suspended during either volume copy operation. This suspension occurs because the state of data on the target volume is inconsistent until the procedure is complete. After the volume copy operation is complete, the target volume automatically becomes read-only to the hosts.

The offline and online volume copy operations are described as follows.

Offline Copy

An offline copy reads data from the source volume and copies it to a target volume, while suspending all updates to the source volume with the copy in progress. All updates to the source volume are suspended to prevent chronological inconsistencies from being created on the target volume. The offline volume copy relationship is between a source volume and a target volume.

Source volumes that are participating in an offline copy are available for read requests only while a volume copy has a status of In Progress or Pending. Write requests are allowed after the offline copy has completed. If the source volume has been formatted with a journaling file system, any attempt to issue a read request to the source volume might be rejected by the storage array controllers, and an error message might appear. The journaling file system driver issues a write request before it attempts to issue the read request. The controller rejects the write request, and the read request might not be issued due to the rejected write request. This condition might result in an error message appearing, which indicates that the source volume is write protected. To prevent this issue from occurring, do not attempt to access a source volume that is participating in an offline copy while the volume copy has a status of In Progress. Also, make sure that the Read-Only attribute for the target volume is disabled after the volume copy has completed to prevent error messages from appearing.

Online Copy

An online copy creates a point-in-time snapshot copy of a volume within a storage array, while still being able to write to the volume with the copy in progress. This function is achieved by creating a snapshot of the volume and using the snapshot as the actual source volume for the copy. The online volume copy relationship is between a snapshot volume and a target volume. The volume for which the point-in-time image is created is known as the base volume and must be a standard volume in the storage array.

A snapshot volume and a snapshot repository volume are created during the online copy operation. The snapshot volume is not an actual volume containing data; rather, it is a reference to the data that was contained on a volume at a specific time. For each snapshot that is taken, a snapshot repository volume is created to hold the copy-on-write data for the snapshot. The snapshot repository volume is used only to manage the snapshot image.

Before a data block on the source volume is modified, the contents of the block to be modified are copied to the snapshot repository volume for safekeeping. Because the snapshot repository volume stores copies of the original data in those data blocks, further changes to those data blocks write only to the source volume.

NOTE If the snapshot volume that is used as the copy source is active, the base volume performance is degraded due to copy-on-write operations. When the copy is complete, the snapshot is disabled, and the base volume performance is restored. Although the snapshot is disabled, the repository infrastructure and copy relationship remain intact.

The online copy function is enabled with the Snapshot Volume premium feature. To use the online copy function, you must enable the Snapshot Volume premium feature by purchasing a feature key file from your storage vendor.

Components of the Volume Copy Premium Feature

The Volume Copy premium feature includes these components:

Create Copy Wizard, which assists in creating a volume copy.

You can use the Create Copy Wizard to guide you through the following steps in creating a Volume Copy:

- Selecting a source volume from a list of available volumes and the type of copy you want to perform (offline or online)
- Selecting a target volume from a list of available volumes
- Allocating capacity for the snapshot repository volume for online copy types
- Setting the copy priority for the volume copy

When you have completed the wizard dialogs, the volume copy starts, and data is read from the source volume and written to the target volume. Operation in Progress icons appear on the source volume and the target volume while the volume copy has a status of In Progress or Pending.

Copy Manager, which monitors volume copies after they have been created.

After you create a volume copy with the Create Copy Wizard, you can monitor the volume copy through the Copy Manager. You can use the Copy Manager to perform the following actions:

- Monitor the progress of a volume copy
- Stop a volume copy
- Re-copy a volume copy
- Remove copy pairs
- Change target volume permissions
- Change copy priority

Keep these guidelines in mind when you create a volume copy.

Failed Controller	<p>You must manually change controller ownership to the alternate controller to allow the volume copy to complete under all of these conditions:</p> <ul style="list-style-type: none"> ■ The preferred controller of the source volume fails. ■ The ownership transfer does not occur automatically in the failover.
Volume Failover for Online Copy Types	<p>Ownership changes affect the base volume and all of its snapshots. The same controller should own the base volume, the snapshot volume, and the snapshot repository volume. The rules that apply to the base volume for host-driver-based or controller-based failover modes also apply to the associated snapshots and snapshot repository volumes. If a failover situation occurs, all related volumes change controller ownership as a group.</p>
Volume Copy and Modification Operations for Offline Copy Types	<p>For offline copy operations, if a modification operation is running on a source volume or a target volume, and the volume copy has a status of In Progress, Pending, or Failed, the volume copy does not take place. If a modification operation is running on a source volume or a target volume after a volume copy has been created, the modification operation must complete before the volume copy can start. If a volume copy has a status of In Progress, any modification operation does not take place.</p>
Preferred Controller Ownership	<p>During a volume copy, the same controller must own both the source volume and the target volume. If both volumes do not have the same preferred controller when the volume copy starts, the ownership of the target volume is automatically transferred to the preferred controller of the source volume. When the volume copy is completed or is stopped, ownership of the target volume is restored to its preferred controller. If ownership of the source volume is changed during the volume copy, ownership of the target volume is also changed.</p>
Failed Volume Copy	<p>A volume copy can fail due to these conditions:</p> <ul style="list-style-type: none"> ■ A read error from the source volume ■ A write error to the target volume ■ A failure in the storage array that affects the source volume or the target volume, such as a remote volume mirror role reversal <p>When the volume copy fails, a Needs Attention icon appears in the Array Management Window. While a volume copy has this status, the host has read-only access to the source volume. Read requests from and write requests to the target volume do not take place until the failure is corrected by using the Recovery Guru.</p>
Volume Copy Status	<p>If eight volume copies with a status of In Progress exist, any subsequent volume copy will have a status of Pending, which remains until one of the eight volume copies completes.</p>
Snapshot Volume	<p>A volume copy fails all snapshot volumes that are associated with the target volume, if any exist. If you select a base volume of a snapshot volume, you must disable all of the snapshot volumes that are associated with the base volume before you can select it as a target volume. Otherwise, the base volume cannot be used as a target volume.</p> <p>A volume copy overwrites data on the target volume and automatically makes the target volume read-only to hosts.</p>

Snapshot Failure	If a snapshot volume that is serving as an online copy fails, the volume copy relationship is still maintained between the snapshot volume and the target volume. If the snapshot failure occurs when the physical copy is in progress, the status of "Failed" is displayed in the Copy Manager.
Volume Consistency	When using the online volume copy operation, make sure that the source volume is in a consistent state. If the source volume is not consistent, the online volume copy is also inconsistent. An inconsistent volume might be unusable for its purpose, such as backup.
Copy Failure for Online Copy Types	A copy failure terminates the copy-on-write process for the snapshot volume. If a copy failure occurs due to a snapshot failure because of snapshot repository volume overflow, you can correct the failure by deleting the copy relationship and re-creating it.

Restrictions on Volume Copy

These restrictions apply to the source volume, the target volume, and the storage array when performing volume copy operations.

For an offline volume copy, the source volume is available for read requests only while a volume copy has a status of In Progress or Pending. Write requests are allowed after the volume copy is completed.

- You can use a volume as a target volume in only *one* volume copy at a time.
- The maximum allowable number of volume copies per storage array depends on the number of target volumes that are available in your storage array.
- A storage array can have up to *eight* volume copies running at any given time.
- The capacity of the target volume must be equal to or greater than the capacity of the source volume.

For an offline volume copy, a source volume can be one of the following volumes:

- A standard volume
- A snapshot volume
- A snapshot base volume
- A remote volume mirror primary volume

For an online volume copy, a source volume can only be a *standard volume*.

- If the source volume is a primary volume, the capacity of the target volume must be equal to or greater than the usable capacity of the source volume.
- You cannot use the snapshot volume copy until after the online copy operation completes.
- You cannot use any of the Snapshot Volume options (**Disable**, **Re-create**, **Create Copy**, **Delete**, and **Rename**) or perform host mapping on a snapshot volume that was created using the online copy operation in the Create Copy Wizard.

A target volume can be one of these volumes:

- A standard volume
- A base volume of a disabled snapshot volume or a failed snapshot volume
- A remote volume mirror primary volume

NOTE If you choose a base volume of a snapshot volume as your target volume, you must disable all snapshot volumes that are associated with the base volume before you can select it as a target volume. Otherwise, you cannot use the base volume as a target volume.

Volumes that have these statuses cannot be used as a source volume or a target volume:

- A volume that is reserved by the host cannot be selected as a source volume or a target volume
- A volume that is in a modification operation

- A volume that is the source volume or a target volume in another volume copy operation with a status of Failed, In Progress, or Pending
- A volume with a status of Failed
- A volume with a status of Degraded

For detailed information about this premium feature, refer to the online help topics in the Array Management Window.

Drive Security and Enterprise Key Manager

Drive Security is a premium feature that prevents unauthorized access to the data on a Encryption Services (ES) drive that is physically removed from the storage array. Controllers in the storage array have a security key. Secure drives provide access to data only through a controller that has the correct security key. Drive Security is a premium feature of the storage management software and must be enabled either by you or your storage vendor.

The Drive Security premium feature requires security capable ES drives. A security capable ES drive encrypts data during writes and decrypts data during reads. Each security capable ES drive has a unique drive encryption key.

When you create a secure pool from security capable drives, the drives in that pool become security enabled. When a security capable drive has been security enabled, the drive requires the correct security key from a controller to read or write the data. All of the drives and controllers in a storage array share the same security key. The shared security key provides read and write access to the drives, while the drive encryption key on each drive is used to encrypt the data. A security capable drive works like any other drive until it is security enabled.

Whenever the power is turned off and turned on again, all of the security enabled drives change to a *security locked* state. In this state, the data is inaccessible until the correct security key is provided by a controller.

The Enterprise Key Manager premium feature integrates external key management products.

You can view the Drive Security status of any drive in the storage array. The status information reports whether the drive is in one of these states:

- Security Capable
- Secure – Security enabled or security disabled
- Read/Write Accessible – Security locked or security unlocked

You can view the Drive Security status of any pool in the storage array. The status information reports whether the storage array is in one of these states:

- Security Capable
- Secure

This table interprets the security properties status of a pool.

Table 2 Pool Security Properties

	Security Capable – yes	Security Capable – no
Secure – yes	The pool is composed of all ES drives and is in a Secure state.	Not applicable. Only ES drives can be in a Secure state.
Secure – no	The pool is composed of all ES drives and is in a Non-Secure state.	The pool is not entirely composed of ES drives.

When the Drive Security premium feature has been enabled, the **Drive Security** menu appears in the **Storage Array** menu. The **Drive Security** menu has these options:

- **Security Key Management**
- **Create Security Key**
- **Change Security Key**
- **Save Security Key**
- **Validate Security Key**
- **Import Security Key File**

The **Security Key Management** option lets you specify how to manage the security key. By default, the security key is managed locally by the controllers. The controllers generate the security key and save the security key in the nonvolatile static random access memory (NVS RAM) of the controllers. You can use the Enterprise Key Manager to have an external key management server generate the security key.

NOTE If you have not created a security key for the storage array, the **Create Security Key** option is active. If you have created a security key for the storage array, the **Create Security Key** option is inactive with a check mark to the left. The **Change Security Key** option, the **Save Security Key** option, and the **Validate Security Key** option are now active.

The **Import Security Key File** option is active if there are any security locked drives in the storage array.

When the Drive Security premium feature has been enabled, the **Secure Drives** option appears in the **Pool** menu. The **Secure Drives** option is active if these conditions are true:

- The selected storage array is not security enabled but is comprised entirely of security capable drives.
- The storage array does not contain any snapshot base volumes or snapshot repository volumes.
- The pool is in an Optimal state.
- A security key is set up for the storage array.

The **Secure Drives** option is inactive if the conditions are not true.

The **Secure Drives** option is inactive with a check mark to the left if the pool is already security enabled.

You can erase security enabled drives so that you can reuse the drives in another pool, in another storage array, or if you are decommissioning the drives. When you erase security enabled drives, you make sure that the data cannot be read. When all of the drives that you have selected in the Physical pane are security enabled, and none of the selected drives are part of a pool, the **Secure Erase** option appears in the **Drive** menu.

The storage array password protects a storage array from potentially destructive operations by unauthorized users. The storage array password is independent from the Drive Security premium feature and should not be confused with the pass phrase that is used to protect copies of a security key. However, it is good practice to set a storage array password before you create, change, or save a security key or unlock secure drives.

Using Enterprise Key Manager

The Enterprise Key Manager premium feature lets you specify how to manage the security key. You can choose to manage the security key locally by the controllers or externally by an external key management server. By default, the security key is managed locally by the controllers. The controllers generate the security key and save the security key in the nonvolatile static random access memory (NVS RAM) of the controllers. You can also use the Enterprise Key Manager to have an external key management server generate the security key. To change the management method, select **Storage Array >> Drive Security >> Security Key Management**.

ATTENTION Changing the method of managing the security key from local to external requires creating and saving a new security key. This action makes any previously saved security key for the storage array invalid.

NOTE External key management must be enabled for both the source storage array, from which the key is saved, and any target storage array that imports the key. The key management server used by the source storage array must be accessible to the target storage array.

A copy of the security key must be kept on some other storage medium for backup, in case of controller failure or for transfer to another storage array. A *pass phrase* that you provide is used to encrypt and decrypt the security key for storage on other media. The storage array password protects a storage array from potentially destructive operations by unauthorized users. The storage array password is independent from the Drive Security premium feature and should not be confused with the pass phrase that is used to protect copies of a security key. However, it is good practice to set a storage array password before you change a security key.

Creating a Security Key

Drives with the Encryption Services technology are *security capable*. This capability enables the controller to apply security to every security capable drive in the storage array. The controller firmware creates a key and activates the drive's security function, which encrypts data as it enters, and decrypts data as it is read. Without the key, the data written on a drive is inaccessible and unreadable. A security enabled drive can also be configured to require a password, PIN, or certificate; however, this function is separate from the encryption and decryption processes.

The storage array password protects a storage array from potentially destructive operations by unauthorized users. The storage array password is independent from the Drive Security premium feature and should not be confused with the pass phrase that is used to protect copies of a Drive Security key. However, it is good practice to set a storage array password before you create a Drive Security key.

After the controller creates the key, the storage array moves from a state of *security capable* to a state of *security enabled*. The security enabled condition requires the drives to obtain a key to access their media. As an added security measure, when power is applied to the storage array, the drives are all placed in a *security locked* state. They are only unlocked during drive initialization with the controller's key. The *security unlocked* state allows the drives to be accessible so that read and write activities can be performed.

Changing a Security Key

A new security key is generated by the controller firmware for these reasons:

- You need to change the security key.
- You need to change the method of managing the security key from local to external.

ATTENTION Changing the method of managing the security key makes any previously saved security keys invalid.

The new security key is stored in the *nonvolatile static random access memory (NVS RAM)* of the controllers. The new key replaces the previous key. You cannot see the security key directly. A copy of the security key must be kept on some other storage medium for backup, in case of controller failure or for transfer to another storage array. A *pass phrase* that you provide is used to encrypt and decrypt the security key for storage on other media.

The storage array password protects a storage array from potentially destructive operations by unauthorized users. The storage array password is independent from the Drive Security feature and should not be confused with the pass phrase that is used to protect copies of a Drive Security key. However, it is good practice to set a storage array password before you change a Drive Security key.

Saving a Security Key

You save an externally storable copy of the security key when the security key is first created and each time it is changed. You can create additional storable copies at any time. To save a new copy of the security key, you must provide a pass phrase. The pass phrase that you choose does not need to match the pass phrase that was used when the security key was created or last changed. The pass phrase is applied to the particular copy of the security key that you are saving.

Keep these guidelines in mind when you create a pass phrase:

- The pass phrase must be between eight and 32 characters long.
- The pass phrase must contain at least one uppercase letter.
- The pass phrase must contain at least one lowercase letter.
- The pass phrase must contain at least one number.
- The pass phrase must contain at least one non-alphanumeric character, for example, <, >, @, or +.

The characters you enter are not readable in the **Pass phrase** text box.

The storage array password protects a storage array from potentially destructive operations by unauthorized users. The storage array password is independent from the Drive Security feature and should not be confused with the pass phrase that is used to protect copies of a security key. However, it is good practice to set a storage array password before you save a security key.

Unlocking Secure Drives

You can export a security enabled pool to move the associated drives to a different storage array. After you install those drives in the new storage array, you must unlock the drives before data can be read from or written to the drives. To unlock the drives, you must supply the security key from the original storage array. The security key on the new storage array will be different and will not be able to unlock the drives.

You must supply the security key from a security key file that was saved on the original storage array. You must provide the pass phrase that was used to encrypt the security key file to extract the security key from this file.

The storage array password protects a storage array from potentially destructive operations by unauthorized users. The storage array password is independent from the Drive Security feature and should not be confused with the pass phrase that is used to protect copies of a security key. However, it is good practice to set a storage array password before you unlock secure drives.

Validating the Security Key

You validate a file in which a security key is stored through the Validate Security Key dialog. To transfer, archive, or back up the security key, the controller firmware encrypts (or wraps) the security key and stores it in a file. You must provide a pass phrase and identify the corresponding file to decrypt the file and recover the security key.

NOTE You can also install the security key from an external key management server. External key management must be enabled for both the source storage array and the target storage array. The key management server used by the source storage array must be accessible by the target storage array.

Data can be read from a security enabled drive only if a controller in the storage array provides the correct security key. If you move security enabled drives from one storage array to another, you must also import the appropriate security key to the new storage array. Otherwise, the data on the security enabled drives that were moved is inaccessible.

NOTE After 20 consecutive unsuccessful attempts to validate a security key, you might be blocked from making further attempts at validation. The Recovery Guru guides you to reset the limit and make additional attempts. Data on the drives is temporarily inaccessible during the reset procedure.

T10 Protection Information Premium Feature

The T10 Protection Information (PI) premium feature checks for and corrects errors that might occur as data is moved within the controller, such as from cache to the drive. This checking leads to correction of write errors and increases data integrity across the entire storage system. PI is implemented using the SCSI direct-access block-device protection information model. PI creates error-checking information, such as a cyclic redundancy check (CRC) and appends that information to each block of data. Any errors that might occur when a block of data is transmitted or stored is then detected and corrected by checking the data with its error-checking information.

Only certain configurations of hardware, including PI-capable drives, controllers, and host interface cards (HICs), support the PI premium feature. When you install the PI premium feature on a storage array, SANtricity ES Storage Manager provides options to use PI with certain operations. For example, you can create a pool that includes PI-capable drives and then create a volume within that pool that is PI-enabled. Other operations that use a PI-enabled volume have options to support the PI premium feature.

For detailed information about this premium feature, refer to the online help topics in the Array Management Window.

Solid State Disks

Some controllers and drive modules now support Solid State Disks (SSDs). SSDs are data storage devices that use solid state memory (flash) to store data persistently. An SSD emulates a conventional hard drive, thus easily replacing it in any application. SSDs are available with the same interfaces used by hard drives.

The advantages of SSDs over hard drives are:

- Faster start up (no spin up)
- Faster access to data (no rotational latency or seek time)
- Higher I/O operations per second (IOPS)
- Higher reliability with fewer moving parts
- Lower power usage
- Less heat produced and less cooling required

SSD support is a premium feature of the storage management software that must be enabled by either you or your storage vendor.

Identifying SSDs

You can identify SSDs in the storage management software either by the label "SSD" or this icon.



In addition to drive firmware, SSDs have field-programmable gate array (FPGA) code that might be updated periodically. An FPGA version is listed in the drive properties, which you can see in the storage management software by selecting a drive on the **Physical** tab. Also, SSDs do not have a speed listed in the drive properties like hard drives do.

Creating Pools

All of the drives in a pool must have the same media type (hard drive or SSD) and the same interface type. Hot spare drives must also be of the same drive type as the drives they are protecting.

Wear Life

A flash-based SSD has a limited wear life before individual memory locations can no longer reliably persist data. The drive continuously monitors itself and reports its wear life status to the controller. Two mechanisms exist to alert you that an SSD is nearing the end of its useful life: average erase count and spare blocks remaining. You can find these two pieces of information in the drive properties, which you can see in the storage management software by selecting a drive on the **Physical** tab.

The average erase count is reported as a percentage of the rated lifetime. When the average erase count reaches 80 percent, an event is logged to the Major Event Log (MEL). At this time, you should schedule the replacement of the SSD. When the average erase count reaches 90 percent, a Needs Attention condition occurs. At this time, you should replace the SSD as soon as possible.

The spare blocks remaining are reported as a percentage of the total blocks. When the number of spare blocks remaining falls below 20 percent, an event is logged to the MEL. At this time, you should schedule the replacement of the SSD. When the number of spare blocks remaining falls below 10 percent, a Needs Attention condition occurs. At this time, you should replace the SSD as soon as possible.

Write Caching

Write caching will always be enabled for SSDs. Write caching improves performance and extends the life of the SSD.

Background Media Scans

Background media scans are not necessary for SSDs because of the high reliability of SSDs.

Heterogeneous Hosts

Heterogeneous hosts are hosts with different operating systems that share access to the same storage array. When you change a host type, you are changing the *operating system (OS)* for the host adapter's host port.

To specify different operating systems for attached hosts, you must specify the appropriate *host* type when you define the host ports for each host. Host types can be completely different operating systems, or can be variants of the same operating system. By specifying a host type, you define how the controllers in the storage array will work with the particular operating system on the hosts that are connected to it.

Password Protection

NOTE Running operations that alter the configuration of your storage array can cause serious damage, including data loss. Configuring a password for each storage array that you manage prevents unauthorized access to destructive commands.

For added security, you can configure each storage array with a password to protect it from unauthorized access. A password protects any options that the controller firmware deems destructive. These options include any functions that change the state of the storage array, such as creating a volume or modifying the cache setting.

NOTE If you forget the password, contact your Sun Customer Care Center representative.

After the password has been set on the storage array, you are prompted for that password the first time you attempt an operation in the Array Management Window that can change the state of the storage array, such as modifying the cache settings. You are asked for the password only once during a single management session.

For storage arrays with a password and alert notifications configured, any attempts to access the storage array without the correct password are reported.

The storage management software provides other security features to protect data, including generation numbering to prevent replay attacks and hashing and encryption to guard against client spoofing and snooping.

Persistent Reservations Management

ATTENTION Sun Customer Care Center representative supervision required – Do not perform this procedure unless you are supervised by your Sun Customer Care Center representative.

Persistent reservation management lets you view and clear volume reservations and associated registrations. Persistent reservations are configured and managed through the cluster server software and prevent other hosts from accessing particular volumes.

Unlike other types of reservations, a persistent reservation performs these functions:

- Reserves access across multiple host ports
- Provides various levels of access control
- Offers the ability to query the storage array about registered ports and reservations
- Optionally, provides for persistence of reservations in the event of a storage array power loss

The *storage management software* lets you manage persistent reservations by performing these tasks:

- Viewing registration and reservation information for all of the volumes in the storage array
- Saving detailed information on volume reservations and registrations
- Clearing all registrations and reservations for a single volume or for all of the volumes in the storage array.

HotScale Technology

HotScale™ technology lets you configure, reconfigure, add, or relocate storage array capacity without interrupting user access to data.

Port bypass technology automatically opens ports and closes ports when drive modules are added to or removed from your storage array. Fibre Channel loops stay intact so that system integrity is maintained throughout the process of adding and reconfiguring your storage array.

For more information about using the HotScale technology, contact your Sun Customer Care Center representative.

Chapter 4: Maintaining and Monitoring Storage Arrays

The topics in this section describe the methods for maintaining storage arrays, including troubleshooting storage array problems, recovering from a storage array problem using the Recovery Guru, and configuring alert notifications using the Event Monitor.

For additional conceptual information and detailed procedures for the options described in this section, refer to the Learn About Monitoring Storage Arrays online help topic in the Enterprise Management Window.

Storage Array Health

NOTE To receive notification of events for the storage arrays, you must configure alert notifications in the Enterprise Management Window, and the Event Monitor must be running.

The Enterprise Management Window summarizes the conditions of all of the known storage arrays being managed. Appropriate status indicators appear in the Tree view on the **Devices** tab, the Table view on the **Devices** tab, and the Health Summary Status area in the lower-left corner of the window. To show the status bar, select **View >> Status Bar**.

Background Media Scan

A background media scan is a background process that is performed by the controllers to provide error detection on the drive media. A background media scan can find media errors before they disrupt normal drive reads and writes. The background media scan process scans all volume data to make sure that it can be accessed. The errors are reported to the Event Log.

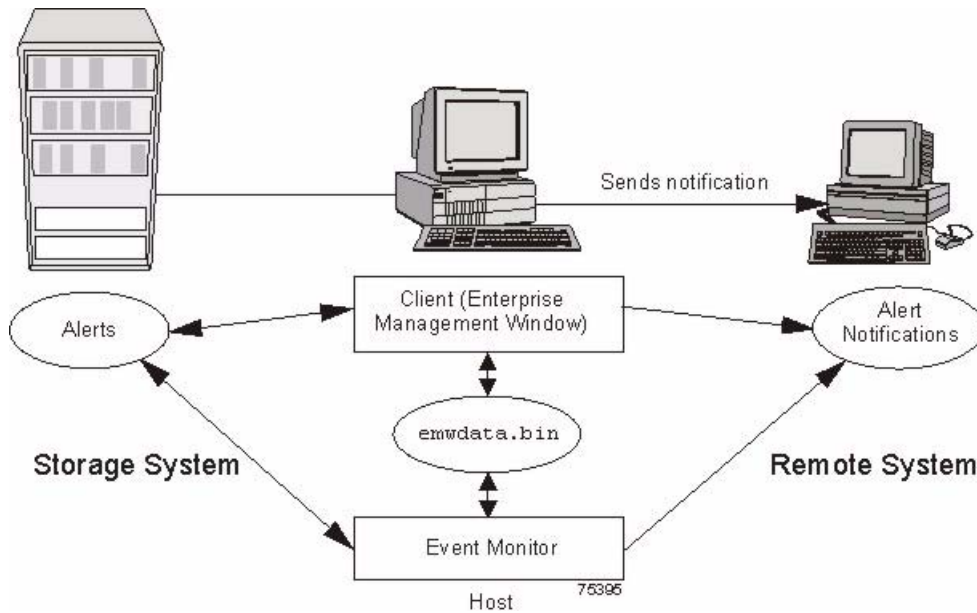
A background media scan runs on all volumes in the storage array for which it has been enabled. You must enable the media scan for the entire storage array, and for individual volumes. If you enable a redundancy check, the background media scan also scans the redundancy data on a RAID Level 1 volume, a RAID Level 3 volume, a RAID Level 5 volume, or a RAID Level 6 volume.

Event Monitor

The Event Monitor runs continuously in the background, monitoring activity on a storage array and checking for problems. Examples of problems include impending drive failures or failed controllers. If the Event Monitor detects any problems, it can notify a remote system by using email notifications, *Simple Network Management Protocol (SNMP)* trap messages, or both, if the Enterprise Management Window is not running.

The Event Monitor is a client that is bundled with the client software. Install the Event Monitor on a computer that runs 24 hours a day. The client and the Event Monitor are installed on a *storage management station* or a *host* that is connected to the storage arrays. Even if you choose not to install the Event Monitor, you can still configure alert notifications on the computer on which the client software is installed.

The following figure shows how the Event Monitor and the Enterprise Management Window client software send alerts to a remote system. The storage management station contains a file with the name of the storage array being monitored and the address to which alerts will be sent. The alerts and errors that occur on the storage array are continuously being monitored by the client software and the Event Monitor. The Event Monitor continues to monitor the client, even after the client software package is shut down. When an event is detected, a notification is sent to the remote system.



Because the Event Monitor and the Enterprise Management Window share the information to send alert messages, the Enterprise Management Window has some visual cues to assist in the installation and synchronization of the Event Monitor.

Using the Event Monitor involves these three key steps:

1. Installing the client software
2. Setting up the alert destinations for the storage arrays that you want to monitor from the Enterprise Management Window
3. Synchronizing the Enterprise Management Window and the Event Monitor

Alert Notifications

You can configure alert notifications by using the storage management software.

Configuring Alert Notifications

You must configure alert notification settings to receive email notifications or SNMP notifications when an event occurs in a storage array. The notification summarizes the event and details about the affected storage array, including these items:

- The name of the affected storage array
- The host IP address (for an in-band managed storage array)
- The host name and ID (shown as out-of-band if the storage array is managed through the Ethernet connection of each controller)
- The event error type related to an Event Log entry
- The date and the time when the event occurred
- A brief description of the event

NOTE To set up alert notifications using SNMP traps, you must copy and compile a management information base (MIB) file on the designated network management station.

Three key steps are involved in configuring alert notifications:

1. Select a node in the Enterprise Management Window that shows alert notifications for the storage arrays that you want to monitor. You can select every storage array being managed, every storage array attached to and managed through a particular host, and individual storage arrays.
2. Configure email destinations, if desired.
3. Configure SNMP trap destinations, if desired. The SNMP trap destination is the IP address or the host name of a station running an SNMP service, such as a network management station.

Customer Support Alert Notifications

If an event occurs in a storage array, the Enterprise Management Window contains options to configure the system to send email notifications to a specified customer support group. After the alert notification option is configured, the email alert notification summarizes the event, provides details about the affected storage array, and provides customer contact information. For more information about setting up this file, contact your Sun Customer Care Center representative.

Performance Monitor

The Performance Monitor provides visibility into performance activity across your monitored storage devices. You can use the Performance Monitor to perform these tasks:

- View in real time the values of the data collected for a monitored device. This capability helps you to determine if the device is experiencing any problems.
- See a historical view of a monitored device to identify when a problem started or what caused a problem.
- Specify various reporting attributes, such as time increments and filtering criteria, to examine performance trends and to pinpoint the cause of availability and performance issues.
- Display data in tabular format (actual values of the collected metrics) or graphical format (primarily as line-graphs), or export the data to a file.

About Metrics

Metrics are measurements of the data that the Performance Monitor collects from the storage devices that you monitor. Metrics help to pinpoint problems and define their cause. Metrics define the types of data that you collect as well as the type of data source from which you collect the data.

Performance Metric Data

You can collect the following metric data:

- **Total I/Os** – Total I/Os performed by this device since the beginning of the polling session.
- **Read Percentage** – The percentage of total I/Os that are read operations for this device. Write percentage can be calculated as 100 minus this value.
- **Cache Hit Percentage** – The percentage of total I/Os that are processed with data from the cache rather than requiring a read from drive.
- **I/O per second** – The number of I/O requests serviced per second during the current polling interval (also called an I/O request rate).
- **KBs or MBs per second** – The transfer rate during the current polling interval. The transfer rate is the amount of data in kilobytes (Table view) or megabytes (Graphical view) that can be moved through the I/O data connection in a second (also called throughput).

NOTE A kilobyte is equal to 1024 bytes, and a megabyte is equal to 1024 x 1024 bytes (1,048,576 bytes).

Metric Sources

Metrics define how the Performance Monitor collects data from supported data sources called metric sources. Metric sources are the aspects of a storage array or a controller that provide data. You can configure the Performance Monitor to report data from the following metric sources:

- Volume
- Pool
- Controller
- Storage array

You can use the data to create reports, and make tuning decisions based on the data values. If a value is outside of the desired range or is in an undesired state, you can take action to correct the problem.

NOTE The Performance Monitor reports volume metrics and pool metrics at the storage array level, regardless of volume controller ownership changes that might occur during monitoring.

Viewing Performance Data

The Performance Monitor provides both real-time analysis and historical context of performance metrics. The metrics are available in either of two views:

- **Table view** – In the Table view, the data is displayed in a tabular format. The actual numeric values of the collected metrics are displayed in a data table.
- **Graphical view** – In the Graphical view, the data is presented with a single x-axis and a single y-axis. The x-axis represents the time for which you selected to view performance data. The y-axis represents the metric you selected to view for a particular metric source.

Performance Tuning

The Performance Monitor provides you with data about devices. You use this data to make storage array tuning decisions, as described in the following table. When performance issues are encountered, tuning is required to alleviate the issues.

Performance Metric Data	Implications for Performance Tuning
Total I/Os	<p>This data is useful for monitoring the I/O activity of a specific controller and a specific volume, which can help identify possible high-traffic I/O areas.</p> <p>If the I/O rate is slow on a volume, try increasing the pool size by selecting Pool >> Add Free Capacity (Drives).</p> <p>You might notice a disparity in the total I/Os (workload) of controllers. For example, the workload of one controller is heavy or is increasing over time while that of the other controller is lighter or more stable. In this case, you might want to change the controller ownership of one or more volumes to the controller with the lighter workload. Use the volume total I/O statistics to determine which volumes to move.</p> <p>You might want to monitor the workload across the storage array. Look at the Total I/Os column of the Storage Array Totals row in the Performance Monitor window. If the workload continues to increase over time while application performance decreases, you might need to add additional storage arrays. By adding storage arrays to your enterprise, you can continue to meet application needs at an acceptable performance level.</p>
Read Percentage	<p>Use the Read Percentage for a volume to determine actual application behavior. If a low percentage of read activity exists relative to write activity, you might want to change the RAID level of a pool from RAID Level 5 to RAID Level 1 to obtain faster performance.</p>
Cache Hit Percentage	<p>A higher cache hit percentage is desirable for optimal application performance. A positive correlation exists between the cache hit percentage and the I/O rates.</p> <p>The cache hit percentage of all of the volumes might be low or trending downward. This trend might indicate inherent randomness in access patterns. In addition, at the storage array level or the controller level, this trend might indicate the need to install more controller cache memory if you do not have the maximum amount of memory installed.</p> <p>If an individual volume is experiencing a low cache hit percentage, consider enabling dynamic cache read prefetch for that volume. Dynamic cache read prefetch can increase the cache hit percentage for a sequential I/O workload.</p>
KB/s or MB/s	<p>The transfer rates of the controller are determined by the application I/O size and the I/O rate. Generally, small application I/O requests result in a lower transfer rate but provide a faster I/O rate and shorter response time. With larger application I/O requests, higher throughput rates are possible. Understanding your typical application I/O patterns can help you determine the maximum I/O transfer rates for a specific storage array.</p>
IOPS	<p>Factors that affect input/output operations per second (IOPS) include these items:</p> <ul style="list-style-type: none"> ■ Access pattern (random or sequential) ■ I/O size ■ RAID level ■ Segment size ■ The number of drives in the pools or storage array <p>The higher the cache hit rate, the higher I/O rates will be.</p> <p>You can see performance improvements caused by changing the segment size in the IOPS statistics for a volume. Experiment to determine the optimal segment size, or use the file system size or database block size.</p> <p>Higher write I/O rates are experienced with write caching enabled compared to disabled. In deciding whether to enable write caching for an individual volume, look at the current IOPS and the maximum IOPS. You should see higher rates for sequential I/O patterns than for random I/O patterns. Regardless of your I/O pattern, enable write caching to maximize the I/O rate and to shorten the application response time.</p>

For detailed information about the Performance Monitor, refer to the online help topics in the Array Management Window.

Viewing Operations in Progress

The Operations in Progress dialog displays all of the long-running operations that are currently running in the storage array. From this dialog, you cannot interact with the operations. You can only view their progress.

The Operations in Progress dialog remains open until you close it or until you close the Array Management Window (AMW). You can do other tasks in the AMW while the Operations in Progress dialog is open.

You can view the progress for the following long-running operations:

- Dynamic Capacity Expansion (DCE) – Adding capacity to a pool
- Dynamic RAID Migration (DRM) – Changing the RAID level of a pool
- Checking the data redundancy of a pool
- Defragmenting a pool
- Initializing a volume
- Dynamic Volume Expansion (DVE) – Adding capacity to a volume
- Dynamic Segment Size (DSS) – Changing the segment size of a volume
- Reconstruction – Reconstructing data from parity because of unreadable sectors or a failed drive
- Copyback – Copying data from a hot spare drive to a new replacement drive
- Volume copy
- Synchronizing a remote mirror

For detailed information about this feature, refer to the online help topics in the Array Management Window.

Retrieving Trace Buffers

NOTE Use this option only under the guidance of your Sun Customer Care Center representative.

You can save trace information to a compressed file. The firmware uses the trace buffers to record processing, including exception conditions, that might be useful for debugging. Trace information is stored in the current buffer. You have the option to move the trace information to the flushed buffer after you retrieve the information. You can retrieve trace buffers without interrupting the operation of the storage array and with minimal effect on performance.

A zip-compressed archive file is stored at the location you specify on the host. The archive contains trace files from one or both of the controllers in the storage array along with a descriptor file named `trace_description.xml`. Each trace file includes a header that identifies the file format to the analysis software used by the Sun Customer Care Center representative. The descriptor file has the following information:

- The World Wide Identifier (WWID) for the storage array.
- The serial number of each controller.
- A time stamp.
- The version number for the controller firmware.
- The version number for the management application programming interface (API).
- The model ID for the controller board.
- The collection status (success or failure) for each controller. If the status is Failed, the reason for failure is noted, and there is no trace file for the failed controller.

For detailed information about this feature, refer to the online help topics in the Array Management Window.

Upgrading the Controller Firmware

You can upgrade the firmware of the controllers in the storage array by using the storage management software.

In the process of upgrading the firmware, the firmware file is downloaded from the host to the controller. After downloading the firmware file, you can upgrade the controllers in the storage array to the new firmware immediately. Optionally, you can download the firmware file to the controller and upgrade the firmware later at a more convenient time.

The process of upgrading the firmware after downloading the firmware file is known as *activation*. During activation, the existing firmware file in the memory of the controller is replaced with the new firmware file.

The firmware upgrade process requires that the controllers have enough free memory space in which the firmware file resides until activation.

A version number exists for each firmware file. For example, 06.60.08.00 is a version number for a firmware file. The first two digits indicate the major revision of the firmware file. The remaining digits indicate the minor revision of the firmware file. You can view the version number of a firmware file in the Upgrade Controller Firmware window and the Download Firmware dialog. For more information, refer to the Downloading the Firmware online help topic in the Enterprise Management Window.

The process of upgrading the firmware can be either a major upgrade or a minor upgrade depending on the version of the firmware. For example, the process of upgrading the firmware is major if the version of the current firmware is 06.60.08.00, and you want to upgrade the firmware to version 07.36.12.00. In this example, the first two digits of the version numbers are different and indicate a major upgrade. In a minor upgrade, the first two digits of the version numbers are the same. For example, the process of upgrading the firmware is minor if the version of the current firmware is 06.60.08.00, and you want to upgrade the firmware to version 06.60.18.00 or any other minor revision of the firmware.

You can use the Enterprise Management Window to perform both major upgrades and minor upgrades. You can use the Array Management Window to perform minor upgrades only.

The storage management software checks for existing conditions in the storage array before upgrading the firmware. Any of these conditions in the storage array can prevent the firmware upgrade:

- An unsupported controller type or controllers of different types that are in the storage array that cannot be upgraded
- One or more failed drives
- One or more hot spare drives that are in use
- One or more pools that are incomplete
- Operations, such as defragmenting a pool, downloading of drive firmware, and others, that are in progress
- Missing volumes that are in the storage array
- Controllers that have a status other than Optimal
- The storage domains database is corrupt
- A data validation error occurred in the storage array
- The storage array has a Needs Attention status
- The storage array is unresponsive, and the storage management software cannot communicate with the storage array
- The Event Log entries are not cleared

You can correct some of these conditions by using the Array Management Window. However, for some of the conditions, you might need to contact your Sun Customer Care Center representative. The storage management software saves the information about the firmware upgrade process in log files. This action helps the Sun Customer Care Center representative to understand the conditions that prevented the firmware upgrade.

You can view the status of a storage array in the Status area of the Upgrade Controller Firmware window. Based on the status, you can select one or more storage arrays for which you want to upgrade the firmware.





You also can use the command line interface (CLI) to download and activate firmware to several storage arrays. For more information, refer to the About the Command Line Interface online help topic in the Enterprise Management Window.



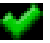
Monitoring the Status of the Download

Monitor the progress and completion status of the *firmware* and *NVSRAM* download to the controllers to make sure that errors did not occur. After the Confirm Download dialog is dismissed, the file is transferred to the storage array. Each controller is sent the new file one at a time. If the file transfer to the first controller succeeds, then the file is transferred to the second controller. The status of the file transfer and the update to each participating controller appear in the Upgrade Controller Firmware window.

NOTE When the firmware download successfully completes, a dialog might appear stating that the current version of the Array Management Window (AMW) is not compatible with the new firmware just downloaded. If you see this message, dismiss the AMW for the storage array, and open it again after selecting the storage array in the Enterprise Management Window (EMW) and selecting **Tools >> Manage Storage Array**. This action launches a new version of the AMW that is compatible with the new firmware.

The progress and status of optimal controllers that are participating in the download appear. Controllers with statuses other than Optimal are not represented.

Status	Description
During Firmware or NVSRAM Download	
Progress bar	Transferring the firmware or the NVSRAM and the completed percentage
During Firmware or NVSRAM Activation	
Progress bar	Activating the firmware or the NVSRAM and the completed percentage of firmware activation
After Download and Results	
Firmware Pending 	The storage array has pending firmware that is ready for activation.
Refreshing	The storage array status is refreshing.
Error 	An error occurred during the operation.
Unresponsive 	The storage array cannot be contacted.
Not-upgradeable 	The storage array cannot be upgraded for one or more reasons. For more information, refer to the Upgrading the Controller Firmware online help topic.

Status	Description
Health Check Passed 	No problems were detected, and you can upgrade the storage array.
Upgradeable: Needs Attention 	One or more problems were detected, but you can still upgrade the storage array.
Firmware Upgraded 	The firmware is successfully upgraded in the storage array.

During firmware downloads, the storage management software periodically polls the controller to see if the download has completed successfully. Sometimes, controller problems occur that keep the download from occurring. This table shows the results of firmware downloads if a controller is failed.

Task	Result
You download new firmware to a storage array. A controller in the storage array fails, and you replace the failed controller with a new one.	After the new controller is installed, the storage array detects the controller replacement and synchronizes the firmware on both controllers.
You download new firmware to a storage array. A controller in the storage array fails, but you place the controller back online (assuming the problem was with something other than the controller).	The firmware synchronization does <i>not</i> occur.

Problem Notification

NOTE To receive notification of events for the storage arrays, the Enterprise Management Window (EMW) or the Event Monitor must be running. In addition, you must have configured the alert notifications in the Enterprise Management Window.

Typically, storage array problems are indicated by using these status notifications:

- A Needs Attention status icon appears in several locations:
 - In the Status bar of the EMW
 - In the Tree view and the Table view on the **Devices** tab of the EMW
 - In the title bar of the Array Management Window (AMW)
 - In the storage array name and status area above the tabs in the AMW
- On the **Summary** tab, the **Logical** tab, and the **Physical** tab in the AMW

Event Log Viewer

The Event Log is a detailed record of events that occur in the storage array. You can use the Event Log as a supplementary diagnostic tool to the Recovery Guru for tracing storage array events. Always refer to the Recovery Guru first when you attempt to recover from component failures in the storage array.

The Event Log is stored in reserved areas on the disks in the storage array.

You can perform these actions in the Event Log window:

- View and filter the events that are displayed in the Event Log.
- Update the display to retrieve any new events.
- View detailed information about a selected event.
- Save selected Event Log data to a file.
- Clear the events in the Event Log.

The Event Log displays three levels of events: Critical, Informational, and Warning. To configure the destination addresses for delivery of email and SNMP trap messages that contain event details affecting managed storage arrays, select **Edit >> Configure Alerts** in the Enterprise Management Window. For more information about SMTP notification, refer to the online help topics in the Enterprise Management Window.

Viewing the Event Log

From the Array Management Window (AMW), select **Advanced >> Troubleshooting >> View Event Log**.

Several minutes might elapse for an event to be logged and to become visible in the Event Log window.

Storage Array Problem Recovery

When you see a storage array **Needs Attention** icon or link, launch the Recovery Guru. The Recovery Guru is a component of the Array Management Window that diagnoses the problem and provides the appropriate procedure to use for troubleshooting.

Recovery Guru

The Recovery Guru window is divided into three panes:

- **Summary** - This pane lists storage array problems.
- **Details** - This pane shows information about the selected problem in the Summary pane.
- **Recovery Procedure** - This pane lists the appropriate steps to resolve the selected problem in the Summary pane.

For detailed information about the Recovery Guru, refer to the online help topics in the Array Management Window.

Glossary

A

Auto-Volume Transfer (AVT)

A feature of the controller firmware that helps to manage each volume in a storage array. When used with a multi-path driver, AVT helps to make sure that an I/O data path always is available for the volumes in the storage array.

C

configured capacity

Space on drives in a storage array that has been designated for use in a pool.

controller

A circuit board and firmware that is located within a controller module or an array module. A controller manages the input/output (I/O) between the host system and data volumes.

copyback

The process of copying data from a hot spare drive to a replacement drive. When a failed drive has been physically replaced, a copyback operation automatically occurs from the hot spare drive to the replacement drive.

D

Default Group

A standard node to which all host groups, hosts, and host ports that do not have any specific mappings are assigned. The standard node shares access to any volumes that were automatically assigned default logical unit numbers (LUNs) by the controller firmware during volume creation.

duplex

A disk array system with two active controllers handling host input/output (I/O) requests, referred to as dual-active controllers.

Dynamic RAID-Level Migration (DRM)

A modification operation that changes the Redundant Array of Independent Disks (RAID) level on a selected pool. During the entire modification process, the user can access data on pools, volumes, and drives in the storage management software. The user cannot cancel this operation after it starts.

Dynamic Volume Expansion (DVE)

A modification operation in the storage management software that increases the capacity of a standard volume or a snapshot repository volume. The operation uses the free capacity available on the pool of the standard volume or the snapshot repository volume. This operation is considered to be dynamic because the user has the ability to continually access data on pools, volumes, and drives throughout the entire operation.

F

Fibre Channel (FC)

A high-speed, serial, storage and networking interface that offers higher performance and greater capacity and cabling distance. FC offers increased flexibility and scalability for system configurations and simplified cabling. FC is a host interface that is a channel-network hybrid using an active, intelligent interconnection scheme (topology) to connect devices over a serial bus. The storage management software uses this connection between the host (where it is installed) and each controller in the storage array to communicate with the controllers.

firmware

Low-level program code that is installed into programmable read-only memory (PROM), where it becomes a permanent part of a computing device. The firmware contains the programming needed for boot and to implement storage management tasks.

Free Capacity node

A contiguous region of unassigned capacity on a defined pool. The user assigns free capacity space to create volumes.

full disk encryption (FDE)

A type of drive technology that can encrypt all data being written to its disk media.

H

HBA host port

The physical and electrical interface on the host bus adapter (HBA) that provides for the connection between the host and the controller. Most HBAs will have either one or two host ports. The HBA has a unique World Wide Identifier (WWID) and each HBA host port has a unique WWID.

heterogeneous hosts

Hosts with different operating systems that share access to the same storage array.

host

A computer that is attached to a storage array. A host accesses volumes assigned to it on the storage array. The access is through the HBA host ports or through the iSCSI host ports on the storage array.

host group

A logical entity that identifies a collection of hosts that share access to the same volumes.

hot spare drive

A spare drive that contains no data and that acts as a standby in case a drive fails in a Redundant Array of Independent Disks (RAID) Level 1, RAID Level 3, RAID Level 5, or RAID Level 6 volume. The hot spare drive can replace the failed drive in the volume.

I**in-band management**

A method to manage a storage array in which a storage management station sends commands to the storage array through the host input/output (I/O) connection to the controller.

L**logical unit number (LUN)**

The number assigned to the address space that a host uses to access a volume. Each host has its own LUN address space. Therefore, the same LUN can be used by different hosts to access different volumes.

M**media scan**

A background process that runs on all volumes in the storage array for which it has been enabled. A media scan provides error detection on the drive media. The media scan process scans all volume data to verify that it can be accessed. Optionally, the media scan process also scans the volume redundancy data.

mirror repository volume

A special volume on the storage array that is created as a resource for each controller in both local storage arrays and remote storage arrays. The controller stores duplicate information on the mirror repository volume, including information about remote writes that are not yet written to the secondary volume. The controller uses the mirrored information to recover from controller resets and from accidental powering-down of storage arrays.

N**network management station (NMS)**

A console with installed network management software that is Simple Network Management Protocol (SNMP) compliant. The NMS receives and processes information about managed network devices in a form that is supported by the Management Information Base (MIB) that the NMS uses.

SANtricity ES Storage Manager provides information about critical events, using SNMP trap messages, to the configured NMS.

node

CONTEXT [Network] [Storage System] An addressable entity connected to an input/output (I/O) bus or network. Used primarily to refer to computers, storage devices, and storage subsystems. The component of a node that connects to the bus or network is a port. (*The Dictionary of Storage Networking Terminology*)

O

out-of-band management

A method to manage a storage array in which a storage management station sends commands to the storage array through the Ethernet connections on the controller.

P

parity

A method that provides complete data redundancy while requiring that only a fraction of the storage capacity of mirroring. The data and parity blocks are divided between the drives so that if any single drive is removed (or fails), the data on the drive can be reconstructed. Data is reconstructed by using the data on the remaining drives. The parity data might exist on only one drive, or the parity data might be distributed between all of the drives in the Redundant Array of Independent Disks (RAID) group.

pool

A set of drives that is logically grouped and assigned a RAID level. Each pool created provides the overall capacity needed to create one or more volumes.

premium feature

A feature that is not available in the standard configuration of the storage management software.

primary volume

A standard volume in a mirror relationship that accepts host input/output (I/O) and stores application data. When the mirror relationship is first created, data from the primary volume is copied in its entirety to the associated secondary volume. The primary volume contains the original user data in a mirroring relationship.

protocol

CONTEXT [Fibre Channel] [Network] [SCSI] A set of rules for using an interconnect or a network so that information conveyed on the interconnect can be correctly interpreted by all parties to the communication. Protocols include such aspects of communication as data representation, data item ordering, message formats, message and response sequencing rules, block data transmission conventions, timing requirements, and so forth. (*The Dictionary of Storage Networking Terminology*, 2004)

R

RAID Level 0

A level of non-redundant Redundant Array of Independent Disks (RAID) in which data is striped across a volume or pool. RAID Level 0 provides high input/output (I/O) performance and works well for non-critical data. All drives are available for storing user data; however, data redundancy does not exist. Data availability is more at risk than with other RAID levels, because any single drive failure causes data loss and a volume status of Failed.

RAID Level 0 is not actually RAID unless it is combined with other features to provide data and functional redundancy, regeneration, and reconstruction, such as RAID Level 1+0 or RAID Level 5+0.

RAID Level 1

A redundant Redundant Array of Independent Disks (RAID) level in which identical copies of data are maintained on pairs of drives, also known as mirrored pairs. RAID Level 1 uses disk mirroring to make an exact copy from one drive to another drive.

RAID Level 1 offers the best data availability, but only half of the drives in the pool are available for user data. If a single drive fails in a RAID Level 1 pool, all associated volumes become degraded, but the mirrored drive allows access to the data. RAID Level 1 can survive multiple drive failures as long as no more than one failure exists per mirrored pair. If a drive pair fails in a RAID Level 1 pool, all associated volumes fail, and all data is lost.

RAID Level 3

A high-bandwidth mode Redundant Array of Independent Disks (RAID) level in which both user data and redundancy data (parity) are striped across the drives. The equivalent of one drive's capacity is used for redundancy data. RAID Level 3 is good for large data transfers in applications, such as multimedia or medical imaging, that read and write large sequential blocks of data.

If a single drive fails in a RAID Level 3 pool, all associated volumes become degraded, but the redundancy data allows access to the data. If two or more drives fail in a RAID Level 3 pool, all associated volumes fail, and all data is lost.

RAID Level 5

A high input/output (I/O) Redundant Array of Independent Disks (RAID) level in which data and redundancy are striped across a pool or volume. The equivalent of one drive's capacity is used for redundancy data. RAID Level 5 is good for multiuser environments, such as database or file system storage, where typical I/O size is small, and there is a high proportion of read activity.

If a single drive fails in a RAID Level 5 pool, then all associated volumes become degraded, but the redundancy data allows access to the data. If two or more drives fail in a RAID Level 5 pool, then all associated volumes fail, and all data is lost.

RAID Level 6

A further development of Redundant Array of Independent Disks (RAID) Level 5. RAID Level 6 protects against simultaneous failure of two member drives by using two independent error correction schemes. Although RAID Level 6 provides ultra-high data reliability, its write penalty is even more severe than that of RAID Level 5 because redundant information must be generated and written twice for each application update. As with RAID Level 4 and RAID Level 5, the write penalty in RAID Level 6 is often mitigated by other storage technologies, such as caching.

RAID Level 10

A striping and mirroring mode used for high performance.

redundancy (data)

Additional information stored along with user data that enables a controller to reconstruct lost data. Redundant Array of Independent Disks (RAID) Level 1 uses mirroring for redundancy. RAID Level 3, RAID Level 5, and RAID Level 6 use redundancy information, sometimes called parity, that is constructed from the data bytes and is striped along with the data on each drive.

redundancy (hardware)

The use of some hardware components that take over operation when the original hardware component fails. For example, if one power-fan CRU fails in a module, the second power-fan CRU can take over the power and cooling requirements for the module.

redundancy check

A scan of volume redundancy data, performed as a part of a background media scan.

Redundant Array of Independent Disks (RAID)

CONTEXT [Storage System] A disk array in which part of the physical storage capacity is used to store redundant information about user data stored on the remainder of the storage capacity. The redundant information enables regeneration of user data in the event that one of the array's member disks or the access path to it fails.

Although it does not conform to this definition, disk striping is often referred to as RAID (RAID Level 0). (*The Dictionary of Storage Networking Terminology*)

remote mirror

A mirrored volume pair that consists of a primary volume at the primary site and a secondary volume at a secondary, remote site.

The secondary, remote volume is unavailable to secondary host applications while mirroring is underway. In the event of disaster at the primary site, the user can fail over to the secondary site. The failover is done by performing a role reversal to promote the secondary volume to a primary volume. Then the recovery host will be able to access the newly promoted volume, and business operations can continue.

redundancy check

A scan of volume redundancy data, performed as a part of a background media scan.

remote mirroring

A configuration in which data on one storage array (the primary storage array) is mirrored across a fabric storage area network (SAN) to a second storage array (the secondary storage array). In the event that the primary storage array fails, mirrored data at the secondary site is used to reconstruct the data in the volumes.

S**secondary volume**

A standard volume in a mirror relationship that maintains a mirror (or copy) of the data from its associated primary volume. The secondary volume is available for host read requests only. Write requests to the secondary volume are not permitted. In the event of a disaster or catastrophic failure of the primary site, the secondary volume can be promoted to a primary role.

Simple Network Management Protocol (SNMP)

CONTEXT [Network] [Standards] An IETF protocol for monitoring and managing systems and devices in a network. The data being monitored and managed is defined by a Management Information Base (MIB). The functions supported by the protocol are the request and retrieval of data, the setting or writing of data, and traps that signal the occurrence of events. (*The Dictionary of Storage Networking Terminology*)

simplex

A one-way transmission of data. In simplex communication, communication can only flow in one direction and cannot flow back the other way.

snapshot repository volume

A volume in the storage array that is made as a resource for a snapshot volume. A snapshot repository volume holds snapshot volume metadata and copy-on-write data for a specified snapshot volume.

snapshot volume

A point-in-time image of a standard volume. A snapshot is the logical equivalent of a complete physical copy, but a snapshot is created much more quickly than a physical copy. In addition, a snapshot requires less unconfigured capacity.

SNMP trap

A notification event issued by a managed device to the network management station when a significant event occurs. A significant event is not limited to an outage, a fault, or a security violation.

Solid State Disk (SSD)

[Storage System] A disk whose storage capability is provided by solid-state random access or flash memory rather than magnetic or optical media.

A solid state disk generally offers very high access performance compared to that of rotating magnetic disks, because it eliminates mechanical seek and rotation time. It may also offer very high data transfer capacity. Cost per byte of storage, however, is typically higher. (*The Dictionary of Storage Networking Terminology*)

source volume

A standard volume in a volume copy that accepts host input/output (I/O) and stores application data. When the volume copy is started, data from the source volume is copied in its entirety to the target volume.

standard volume

A logical component created on a storage array for data storage. Standard volumes are also used when creating snapshot volumes and remote mirrors.

storage management station

A computer running storage management software that adds, monitors, and manages the storage arrays on a network.

striping

CONTEXT [Storage System] Short for data striping; also known as Redundant Array of Independent Disks (RAID) Level 0 or RAID 0. A mapping technique in which fixed-size consecutive ranges of virtual disk data addresses are mapped to successive array members in a cyclic pattern. (*The Dictionary of Storage Networking Terminology*)

T**target volume**

A standard volume in a volume copy that contains a copy of the data from the source volume.

topology

The logical layout of the components of a computer system or network and their interconnections. Topology deals with questions of what components are directly connected to other components from the standpoint of being able to communicate. It does not deal with questions of physical location of components or interconnecting cables. (*The Dictionary of Storage Networking Terminology*)

U

Unconfigured Capacity node

The capacity present in the storage array from drives that have not been assigned to a pool.

V

volume

The logical component created for the host to access storage on the storage array. A volume is created from the capacity available on a pool. Although a volume might consist of more than one drive, a volume appears as one logical component to the host.

Volume Copy

A premium feature that copies data from one volume (the source volume) to another volume (the target volume) within a single storage array.

W

write caching

An operation in which data is moved from the host to the cache memory on the controllers. This operation allows the controllers to copy the data to the drives that comprise a volume. Write caching helps improve data throughput by storing the data from the host until the controller can access the volume and move the data.

Copyright © 2011 LSI Corporation. All rights reserved.
Copyright © 2011 Sun Microsystems, Inc. All rights reserved.
Printed in U.S.A.

