



Sun StorEdge™ 3000 Family Best Practices Manual

Sun StorEdge 3510 FC Array

Sun StorEdge 3511 FC Array with SATA

Sun Microsystems, Inc.
www.sun.com

Part No. 816-7325-15
June 2004, Revision A

Submit comments about this document at: <http://www.sun.com/hwdocs/feedback>

Copyright © 2004 Dot Hill Systems Corporation, 6305 El Camino Real, Carlsbad, California 92009, USA. All rights reserved.

Sun Microsystems, Inc. and Dot Hill Systems Corporation may have intellectual property rights relating to technology embodied in this product or document. In particular, and without limitation, these intellectual property rights may include one or more of the U.S. patents listed at <http://www.sun.com/patents> and one or more additional patents or pending patent applications in the U.S. and other countries.

This product or document is distributed under licenses restricting its use, copying distribution, and decompilation. No part of this product or document may be reproduced in any form by any means without prior written authorization of Sun and its licensors, if any.

Third-party software is copyrighted and licensed from Sun suppliers.

Parts of the product may be derived from Berkeley BSD systems, licensed from the University of California. UNIX is a registered trademark in the U.S. and in other countries, exclusively licensed through X/Open Company, Ltd.

Sun, Sun Microsystems, the Sun logo, Sun StorEdge, AnswerBook2, docs.sun.com, and Solaris are trademarks or registered trademarks of Sun Microsystems, Inc. in the U.S. and in other countries.

U.S. Government Rights—Commercial use. Government users are subject to the Sun Microsystems, Inc. standard license agreement and applicable provisions of the FAR and its supplements.

DOCUMENTATION IS PROVIDED "AS IS" AND ALL EXPRESS OR IMPLIED CONDITIONS, REPRESENTATIONS AND WARRANTIES, INCLUDING ANY IMPLIED WARRANTY OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE OR NON-INFRINGEMENT, ARE DISCLAIMED, EXCEPT TO THE EXTENT THAT SUCH DISCLAIMERS ARE HELD TO BE LEGALLY INVALID.

Copyright © 2004 Dot Hill Systems Corporation, 6305 El Camino Real, Carlsbad, Californie 92009, Etats-Unis. Tous droits réservés.

Sun Microsystems, Inc. et Dot Hill Systems Corporation peuvent avoir les droits de propriété intellectuels relatants à la technologie incorporée dans le produit qui est décrit dans ce document. En particulier, et sans la limitation, ces droits de propriété intellectuels peuvent inclure un ou plus des brevets américains énumérés à <http://www.sun.com/patents> et un ou les brevets plus supplémentaires ou les applications de brevet en attente dans les Etats-Unis et dans les autres pays.

Ce produit ou document est protégé par un copyright et distribué avec des licences qui en restreignent l'utilisation, la copie, la distribution, et la décompilation. Aucune partie de ce produit ou document ne peut être reproduite sous aucune forme, par quelque moyen que ce soit, sans l'autorisation préalable et écrite de Sun et de ses bailleurs de licence, s'il y en a.

Le logiciel détenu par des tiers, et qui comprend la technologie relative aux polices de caractères, est protégé par un copyright et licencié par des fournisseurs de Sun.

Des parties de ce produit pourront être dérivées des systèmes Berkeley BSD licenciés par l'Université de Californie. UNIX est une marque déposée aux Etats-Unis et dans d'autres pays et licenciée exclusivement par X/Open Company, Ltd.

Sun, Sun Microsystems, le logo Sun, Sun StorEdge, AnswerBook2, docs.sun.com, et Solaris sont des marques de fabrique ou des marques déposées de Sun Microsystems, Inc. aux Etats-Unis et dans d'autres pays.

LA DOCUMENTATION EST FOURNIE "EN L'ÉTAT" ET TOUTES AUTRES CONDITIONS, DECLARATIONS ET GARANTIES EXPRESSES OU TACITES SONT FORMELLEMENT EXCLUES, DANS LA MESURE AUTORISEE PAR LA LOI APPLICABLE, Y COMPRIS NOTAMMENT TOUTE GARANTIE IMPLICITE RELATIVE A LA QUALITE MARCHANDE, A L'APTITUDE A UNE UTILISATION PARTICULIERE OU A L'ABSENCE DE CONTREFAÇON.



Please
Recycle



Adobe PostScript

Contents

Overview	1
Comparison of the Sun StorEdge 3510 FC Array and the Sun StorEdge 3511 SATA Array	3
Sun StorEdge 3500 Series Architecture	4
DAS and SAN Storage Architectures	5
Fibre Channel Protocols	6
Supported RAID Levels	7
Logical Drives	7
Cache Optimization	9
Array Management Tools	11
Saving and Restoring Configuration Information	12
Direct-Attached Storage	12
Storage Area Networking	13
Scaling Capacity	14
First Steps	16
General Configuration Considerations	17
Configuring an Array's RCOM Channel	18
Cautions	20

Non-Redundant DAS Configurations	21
Non-Redundant DAS Tips and Techniques	22
Non-Redundant DAS Setup Details	23
High-Availability DAS Configurations	26
High-Availability DAS Tips and Techniques	27
Standard High-Availability DAS Setup Details	28
Six Node High-Availability DAS Setup Details	31
Full-Fabric SAN Configurations	33
Full-Fabric SAN Tips and Techniques	34
Full-Fabric SAN Setup Details	35
High-Performance SAN Configurations	38
High-Performance SAN Tips and Techniques	39
High-Performance SAN Setup Details	40
Scaling Arrays Into High Capacity Configurations	43
Limitations	43
Sun StorEdge 3510 FC Array High Capacity Configurations	44
Sun StorEdge 3511 SATA Array High Capacity Configurations	53
Summary	59

Best Practices for the Sun StorEdge 3510 and 3511 FC Arrays

This Best Practices Manual describes both the Sun StorEdge™ 3510 FC Array and Sun StorEdge 3511 FC Array with SATA, and complements the *Sun StorEdge 3000 Family Installation, Operation, and Service Manual* for these products.

The *Sun StorEdge 3511 FC Array with SATA* is also known as the *Sun StorEdge 3511 SATA array* which will be the term used throughout most of this document. Procedures for configuration of the Sun StorEdge 3500 series arrays are identical except where specifically noted in this manual.

Note – The primary use of the two arrays is very different due to different disk technologies. Be sure to read the following Overview section and the [“Comparison of the Sun StorEdge 3510 FC Array and the Sun StorEdge 3511 SATA Array”](#) on page 3 for optimum use of each array.

Overview

Today’s business, educational and consumer environments continue to create a seemingly inexhaustible amount of information. In addition, there is a constant strain on the trade-offs between performance and cost when architecting and deploying a storage solution. Storage technology has provided a multitude of options to appropriately store information based upon retrieval characteristics such as: frequency, random/sequential type of access, criticality, time-sensitivity and cost basis.

Both the Sun StorEdge 3500 series arrays feature full 2Gb host Fibre Channel (FC) connections and are next-generation storage systems designed to provide direct attached storage (DAS) to entry-level, mid-range, and enterprise servers, or to serve as the disk storage within a storage area network (SAN).

Modern businesses employ servers of all sizes requiring storage to match the type of data storage and processing being used by the applications involved, as well as consider the cost of the server. The Sun StorEdge 3500 product family is designed to address network application priorities and cost considerations by applying the appropriate disk technology as follows:

- Applications demanding transaction-based random access patterns with high duty cycles are better suited by the FC-disk solution of the StorEdge 3510 FC array.

The StorEdge 3510 FC array is designed for business-critical applications and their high-value data. These critical applications require higher availability (not a trait of SATA-disk based solutions).

- Applications requiring streaming or sequentially accessed data or those with cost-sensitive duty cycles such as backup and archival are better served by the Serial ATA-disk solution of the StorEdge 3511 SATA array.

The StorEdge 3511 SATA array is designed for applications requiring lower-cost bulk storage (not a trait of FC-disk based solutions), where application or data availability is not a primary concern.

The StorEdge 3511 SATA array is only used for archival, backup, or static storage; it is not used to store single instances of data, and cannot serve as a boot device.

Designed for smaller direct-attached storage (DAS) needs or for consolidated storage applications using a Storage Area Network (SAN), matching the array reliability and disk performance to the application availability requirements has never been easier.

This document gives a high level overview of the Sun StorEdge 3500 series arrays and outlines five sample storage solutions for entry-level, mid-range, and enterprise servers. Use these solutions as-is or tailor them to fit your exact needs. Examples of customization opportunities include adding disks, enclosures and software or even combining configurations. Choosing the solution that best matches your particular environment will provide the best results.

The benefits and purpose of each array varies according to network requirements:

- The Sun StorEdge 3510 FC array uses Fibre Channel (FC) disk technology and is ideal for performance-sensitive, intense I/O applications.
- The Sun StorEdge 3511 SATA array uses Serial ATA (SATA) disk technology, and is ideal for cost-sensitive archival, static storage, and backup applications.

Comparison of the Sun StorEdge 3510 FC Array and the Sun StorEdge 3511 SATA Array

Before installing and configuring your array, please review the key differences between the Sun StorEdge 3510 FC array and the Sun StorEdge 3511 SATA array. The following list provides a comparison of the Sun StorEdge 3510 FC array and the Sun StorEdge 3511 SATA array features.



Caution – Although the two products are very similar in appearance and setup, the configurations have very important differences. While the StorEdge 3510 FC Array can be used for all applications, its higher cost makes it a poor choice for certain applications. The StorEdge 3511 SATA array should only be used where indicated below. Its features are a subset of the StorEdge 3510 FC array’s and therefore may not be used everywhere the StorEdge 3510 FC array can be used. Using the StorEdge 3511 SATA array where a FC disk-based solution should be used may result in the loss of data access or even, in extreme cases, data corruption.

TABLE 1 Comparison of Sun StorEdge 3510 FC Array and the Sun StorEdge 3511 SATA Array Features

	Sun StorEdge 3510 FC Array	Sun StorEdge 3511 SATA Array
Applications	Best suited for production applications where the superior features of FC technical characteristics and performance is a requirement. This includes the below online applications: <ul style="list-style-type: none"> • Database • Decision Support • Data warehousing • Electronic commerce • Enterprise resource planning • Messaging, file and print 	Best suited for an inexpensive secondary storage application where higher capacity drives are needed which are not mission-critical, and where lower performance and less than 7/24 availability is an option. This includes the below near-line applications: <ul style="list-style-type: none"> • Information lifecycle management • Content addressable storage • Backup and restore • Secondary SAN storage • Near-line DAS storage • Static reference data storage
Disks	Fibre Channel disks: 36, 73, or 146 GB at 10k rpm 36 or 73 GB at 15k rpm	SATA disks: 250 GB at 7200 rpm
Maximum FC Host Ports per I/O Controller Module	4 (one SFP port each for channels 0, 1, 4, and 5)	6 (two SFP ports each for channels 1 and 0; one SFP port each for channels 4 and 5)
Maximum Number of Expansion Units Connected to a RAID Array	8	5

TABLE 1 Comparison of Sun StorEdge 3510 FC Array and the Sun StorEdge 3511 SATA Array Features

	Sun StorEdge 3510 FC Array	Sun StorEdge 3511 SATA Array
Maximum Number of Disks per Configuration	108 (1 RAID array + 8 expansion units)	72 (1 RAID array + 5 expansion units)
Maximum Number of Logical Drives	8 logical drives	8 logical drives
Maximum Total Storage Capacity	15.75 TB	18.0 TB
Maximum Usable Storage Capacity	14.0 TB RAID 5 15.1 TB RAID 0	14.0 TB RAID 5 16.0 TB Raid 0
JBOD Array Support	One JBOD per server	Not supported

Sun StorEdge 3500 Series Architecture

All members of the StorEdge 3500 series arrays share the same architecture elements. This section discusses those elements, making note of the few differences in how the architecture was implemented differently in the StorEdge 3510 FC and 3511 SATA products.

The Sun StorEdge 3500 series RAID controllers have six FC channels. RAID controller channels 0, 1, 4, and 5 are normally designated for connection to hosts or Fibre Channel switches. RAID controller channels 2 and 3 are dedicated drive channels that connect to disks. Each channel has a single port connection except the Sun StorEdge 3511 SATA array which has two extra ports (two connections for channels 0 and 1).

In a dual RAID controller configuration, the architecture of the loops within the chassis provides both RAID controllers the same host channel designators. Each host channel of the top RAID controller shares a loop with the matching host channel on the bottom RAID controller. For example, channel 0 of the top RAID controller shares the same loop as channel 0 of the bottom RAID controller. This provides four distinct loops for connectivity. The individual loops provide LUN failover without causing HBA path failover in the event of a controller failure.

A single RAID controller configuration is slightly different. The lower I/O board has drive channels but does not have host channels. Overall, the same number of loops are available, but with only half as many host channel ports.

Sun StorEdge 3510 FC Array. RAID controller channels 0, 1, 4, and 5 are normally designated host channels. Any of the host channels can be configured as a drive channel. In a dual controller configuration, each host loop includes two ports per loop, one port on the top controller and one port on the bottom controller.

All six fibre channels in an I/O controller module support 1-Gbit or 2-Gbit data transfer speeds.

RAID controller channels 2 and 3 are dedicated drive channels that connect to expansion units. Each I/O board has two ports designated as disk drive loops. These ports connect to the internal dual-ported FC disk drives and are used to add expansion units to the configuration. The two drive loop ports on the upper I/O board form FC loop 2 (channel 2) while the two drive ports on the lower I/O board form FC loop 3 (channel 3). FC loop 2 provides a data path from both RAID controllers to the A loop of the internal disk drives, while FC loop 3 provides a data path from both RAID controllers to the B loop of the internal disk drives.

Sun StorEdge 3511 SATA Array. RAID controller channels 0 and 1 are dedicated host channels. Channels 4 and 5 are normally designated host channels but can be configured as drive channels. RAID controller channels 2 and 3 are dedicated drive channels that connect to expansion units.

How the channels are exposed on the StorEdge 3511 RAID controller is one of the key differences between the StorEdge 3510 and 3511. Host channels 0 and 1 include four ports per loop (two ports on the upper controller, and two ports on the lower controller). Channels 0 and 1 support 1 Gbit or 2 Gbit data transfer speed.

Channels 4 and 5 provide two ports per loop (one port on each controller). Channels 4 and 5 only support 2 Gbit data transfer speed.

Each RAID controller has two ports designated as disk drive loops. The drive ports support only 2 Gbit data transfer speed. These ports connect to the internal Serial ATA disk drives using internal FC-SATA routing technology. Drive ports are used to add expansion chassis to the configuration. Like the host channels, each drive channel of the top RAID controller shares a loop with the matching drive channel on the bottom RAID controller. For example, drive channel 2 of the top RAID controller shares the same loop as channel 2 of the bottom RAID controller.

DAS and SAN Storage Architectures

There are two popular methods for connecting storage to servers.

- Direct attached storage (DAS) uses a direct connection between a server and its storage system. The DAS solution of connecting each server to its own dedicated storage system is straight-forward, and the absence of storage switches can minimize costs in some instances.
- A storage area network (SAN) places a Fibre Channel (FC) storage switch between network servers and storage systems. A SAN solution shares a storage system among multiple servers using FC storage switches, and reduces the total number of storage systems required for a particular environment at the cost of additional element management (FC switches) and FC path complexity.




FIGURE 1 DAS and SAN Storage Architectures

Selecting the best storage architecture for a particular environment can be a confusing exercise. In general, some environments are well-suited for DAS while others will benefit greatly from SAN.

The challenge of selecting between DAS and SAN is often further complicated by the need to choose between different storage systems, one designed for DAS or another intended for SAN. Fortunately, Sun StorEdge 3500 series arrays inherently supports both DAS and SAN.

Fibre Channel Protocols

The Sun StorEdge 3500 series arrays support point-to-point and Fibre Channel–Arbitrated Loops (FC–AL) protocols. Using the point-to-point protocol with the Sun StorEdge 3500 series arrays requires a switched fabric network (SAN), whereas selecting FC–AL mode enables the arrays to be used in either DAS or SAN environments. Using point-to-point protocol enables full-duplex use of the available channel bandwidth, whereas selecting FC–AL mode limits host channels to half-duplex communication.

In a point-to-point configuration, only one ID can be assigned to each host channel. If more than one ID is assigned, the point-to-point protocol rules are violated. Any host channel with more than one ID will not be able to log in to an FC switch in fabric mode. This “one-ID-per-channel” requirement is true in both single-controller and dual-controller configurations. Thus, in dual-controller configurations, either the primary or the secondary controller can have an ID assigned, but not both. This yields:

4 (host channels) × 1 (ID per channel) × 32 (LUNs per ID) = 128 maximum addressable LUNs in a fabric point-to-point environment. If dual paths are desired for each logical device, a maximum of 64 dual-pathed LUNs are available.

In an FC-AL configuration, multiple IDs can be assigned to any given host channel. The maximum number of storage partitions that can be mapped to a RAID array is 1024. There are several ways that 1024 LUNs can be configured. For example:

4 (host channels) × 8 (IDs per channel) × 32 (LUNs per ID) =
1024 maximum addressable LUNs in a FC-AL environment.

However, configuring the maximum number of LUNs increases overhead and can have a negative impact on performance.

The FC-AL protocol should be selected for environments needing more than 128 LUNs, or where a switched fabric network is not available.

Supported RAID Levels

There are several RAID level choices: RAID 0, 1, 3, 5, 1+0 (10), 3+0 (30), and 5+0 (50). RAID levels 1, 3, and 5 are the most commonly used. The Sun StorEdge 3500 series arrays support the use of both global and local spare drives in the unlikely event of disk failure. It is good practice to use spare drives when configuring RAID devices. Refer to the *Sun StorEdge 3000 Family RAID Firmware User's Guide* for detailed information about how RAID levels and spare drives are implemented.

Logical Drives

A logical drive (LD) is a group of physical drives configured with a RAID level. Each LD can be configured for a different RAID level.

Sun StorEdge 3500 series arrays support a maximum of eight LDs. A LD can be managed by either the primary or secondary controller. The best practice for creating LDs is to add them evenly across the primary and secondary controllers. The most efficient maximum configuration would have four LDs assigned to each controller. With at least one LD assigned to each controller, both controllers are active. This configuration is known as an active-active controller configuration and allows maximum use of a dual controller array's resources.

Supporting large storage capacities requires using maximum-sized logical drives. See [TABLE 2](#) and [TABLE 4](#) to determine the maximum number of disk drives you can include in a logical drive. The largest supported logical drive configuration is determined by the size of your disk drives, cache optimization and the logical drive's RAID level (see [TABLE 3](#) and [TABLE 5](#)). Maximizing existing logical drives

before building a new logical drive is recommended. In the Sun StorEdge 3511 SATA array, for example, this can result in a supported storage capacity of 16 Tbytes (8 logical drives x 2 Tbytes each = 16 Tbytes total capacity).

TABLE 2 Maximum Number of Disks per Sun StorEdge 3510 FC Array Logical Drive

Disk Capacity (GB)	RAID 5 Random	RAID 5 Sequential	RAID 3 Random	RAID 3 Sequential	RAID 1 Random	RAID 1 Sequential	RAID 0 Random	RAID 0 Sequential
36.2	14	31	14	31	28	36	14	36
73.4	7	28	7	28	12	30	6	27
146.8	4	14	4	14	6	26	3	13

TABLE 3 Maximum Usable Capacity (Gbyte) per Sun StorEdge 3510 FC Array Logical Drive

Disk Capacity	RAID 5 Random	RAID 5 Sequential	RAID 3 Random	RAID 3 Sequential	RAID 1 Random	RAID 1 Sequential	RAID 0 Random	RAID 0 Sequential
36.2	471	1086	471	1086	507	543	507	1122
73.4	440	1982	440	1982	440	1101	440	1982
146.8	440	1908	440	1908	440	1908	440	1908

Note – For optimum configuration of the Sun StorEdge 3511 SATA array, only sequential optimization is used.

TABLE 4 Maximum Number of Disks per Sun StorEdge 3511 SATA Array Logical Drive

Disk Capacity (GB)	RAID 5 Sequential	RAID 3 Sequential	RAID 1 Sequential	RAID 0 Sequential
250.0	8	8	8	8

TABLE 5 Maximum Usable Capacity (Gbyte) per Sun StorEdge 3511 SATA Array Logical Drive

Disk Capacity	RAID 5 Sequential	RAID 3 Sequential	RAID 1 Sequential	RAID 0 Sequential
250.0	1908	1908	2000	2000

Note – You cannot use all disks for data when using 108 146-Gbyte or 72 250-Gbyte disks. The remaining disks are usually used as spares.

Be sure to check the latest product Release Notes for additional guidelines or limitations for large configurations.

Each LD can be partitioned into up to 128 separate partitions or used as a single partition. Partitions are presented to hosts as LUNs.

Once the LDs have been created, assigned to a controller, and partitioned, the partitions must be mapped to host channels as LUNs in order for them to be seen by a host. It is usually desirable to map each partition to two host channels for redundant pathing.

A partition can only be mapped to a host channel where its controller has an assigned ID. For example, if LD 0 is assigned to the primary controller, all partitions on LD 0 will need to be mapped to a host channel ID on the primary controller (PID). Any LDs assigned to the secondary controller will need to have all partitions mapped to a host channel ID on the secondary controller (SID).

When attaching FC cables for LUNs configured with redundant paths, make sure one cable is connected to a channel on the upper controller and the other cable is connected to a different channel on the lower controller. Then, if multipathing software is configured on the host, a controller can be hot-swapped in the event of failure without losing access to the LUN.

For example, suppose partition 0 of LD0 is mapped to Channel 0 PID 42 and Channel 5 PID 47. To ensure that there is no single point of failure (SPOF), connect a cable from the host HBA or a switch port to the upper board port FC0, and connect a second cable from the lower board port FC5 to a different host HBA or switch.

Cache Optimization

The Sun StorEdge 3500 series provides settings for both sequential I/O and random I/O. Sequential I/O is the default setting. As previously discussed, the StorEdge 3511 SATA array provides a subset of the StorEdge 3510's function. Given the characteristics of SATA disks and the application performance requirements for bulk storage, the StorEdge 3511 SATA array should never be used with the random I/O setting. Poor performance and loss of data availability may occur. Sequential I/O is the default setting. See [TABLE 6](#) for a brief overview of the optimization modes.

The logical drive, cache memory, and other controller internal parameters are adjusted to transfer information most efficiently for each optimization mode; they also control the size of logical drives and number of disks per logical drive. For the maximum capacity and number of disks in each optimization mode, refer to:

- “Maximum Number of Disks per Sun StorEdge 3510 FC Array Logical Drive” on page 8
- “Maximum Usable Capacity (Gbyte) per Sun StorEdge 3510 FC Array Logical Drive” on page 8.
- “Maximum Number of Disks per Sun StorEdge 3511 SATA Array Logical Drive” on page 8
- “Maximum Usable Capacity (Gbyte) per Sun StorEdge 3511 SATA Array Logical Drive” on page 8

Note – On the Sun StorEdge 3511 SATA array, use the sequential cache optimization option only (default). Changing the cache optimization mode to random significantly reduces the supported size of logical drives from 2 TB to 512 GB.

TABLE 6 Sequential and Random Optimization Modes

	Sequential Optimization	Random Optimization
Applicable products	Sun StorEdge 3510 FC arrays Sun StorEdge 3511 SATA arrays	Sun StorEdge 3510 FC arrays
Read/Write Data Throughput	Large 128K blocks	Small 32K blocks
Maximum size of logical drive	2 Tbyte	512 Gbyte
Example applications *	Video and imaging applications	Database/transaction processing applications

* See also “[Comparison of Sun StorEdge 3510 FC Array and the Sun StorEdge 3511 SATA Array Features](#)” on page 3.

Note – Sequential or random optimization must be set *prior* to creating logical drives. The change takes effect after the controller resets.

There are two limitations that apply to the optimization modes.

- One optimization mode must be applied to all logical drives in a RAID array.
- Once the optimization mode is selected and logical drives are created, the optimization mode of those logical drives cannot be changed. The only way to change the optimization mode is to delete all logical drives, select the new optimization mode, reboot the array, and create new logical drives. Any existing data on the logical drives is lost during this procedure, which is why it is important to correctly select the optimization mode early in the planning process.

Array Management Tools

The Sun StorEdge 3500 series arrays use the same management interfaces and techniques. The StorEdge 3500 series arrays can be configured and monitored using the following methods as appropriate:

Sun StorEdge 3500 series arrays can be configured and monitored through any of the following methods:

- Using the out-of-band serial port connection (RAID only), a Solaris `tip` session or terminal emulation program for other supported operating systems can be used to access an array's internal firmware application. All procedures can be performed by using the firmware's terminal interface via the COM port.
- Using the out-of-band Ethernet port connection, Sun StorEdge Configuration Service software or the Sun StorEdge CLI can configure and manage an array from a host system.
 - To set up and use the Configuration Service software package, refer to the *Sun StorEdge 3000 Family Configuration Service User's Guide*.
 - The CLI is installed as part of the `SUNWsccli` package. The main advantages of the CLI are that commands can be scripted and information can be passed to other programs. Information about CLI functionality can be found in the *Sun StorEdge 3000 Family CLI User's Guide*, and in the `sccli` man page once the package is installed.
 - The StorEdge 3511 SATA array will respond more slowly than the StorEdge 3510 FC Array when being managed via either CLI or Configuration Service. Due to this, it is not recommended to use either in-band management method with the StorEdge 3511 SATA array.
- Using the out-of-band Ethernet port connection, `telnet` can be used to access the firmware application. All procedures except the initial assignment of an IP address can be done through an Ethernet port connection. Refer to the *Sun StorEdge 3000 Family Installation, Operation and Service Manual* for your array for more information.

Note – The StorEdge 3511 SATA array should only be managed with the out-of-band methods which access the firmware application via Ethernet or serial connections.



Caution – In-band configuration options from a host system include the Sun StorEdge Configuration Service software or the command-line interface (CLI). If you assign an IP address to an array in order to manage it out-of-band, for security reasons make sure that the IP address is on a private network, rather than a publicly routable network.

Saving and Restoring Configuration Information

An important feature of these management tools is the ability to save and restore configuration information in a number of ways. Using the array's firmware application, the configuration information (NVRAM) can be saved to disk. This provides a backup of the controller-dependent configuration information such as channel settings, host IDs, FC protocol, and cache configuration. It does not save LUN mapping information. The NVRAM configuration file can restore all configuration settings but does not rebuild logical drives.

The Sun StorEdge Configuration Service software can be used to save and restore all configuration data, including LUN mapping information. It can also be used to rebuild all logical drives and therefore can be used to completely duplicate an array's configuration to another array.


Direct-Attached Storage

One powerful feature of Sun StorEdge 3500 series arrays is their ability to support multiple direct-attached servers without requiring storage switches. They accomplish this by using intelligent internal Fibre Channel networks. Servers can be directly connected using built-in external Fibre Channel ports, if available, or add-in Fibre Channel host adapter cards.

- **Sun StorEdge 3510 FC Array.** The array automatically configures its ports to match the transfer speed and communication method of each connection.
- **Sun StorEdge 3511 SATA Array.** Channels 0 and 1 automatically configure their ports to match the transfer speed and communication method of each connection. Channels 4 and 5 only support 2 Gbyte transfer speed.



Standard DAS Configuration



High Availability DAS Configuration

FIGURE 2 Two DAS Configurations

The actual number of servers that can be connected varies according to the number of controllers in the array. It also depends on the number of Fibre Channel connections used for each server and the total number of small form-factor pluggable (SFP) interface modules installed. DAS configurations often include single or dual servers only, though a dual-controller array can support up to:

- **Sun StorEdge 3510 FC Array.** Four servers with redundant connections, or eight servers in non-redundant DAS configurations.
- **Sun StorEdge 3511 SATA Array.** Six servers with redundant connections, or 12 servers in non-redundant DAS configurations.

Note – The Sun StorEdge 3511 SATA array cannot be used to store single instances of data but can be used in multi-path and multi-host configurations where the array has a backup or archival role.

Additional SFP modules are required to support more than two servers with redundant connections or four servers in non-redundant configurations.

Note – Except in some clustering configurations, in a DAS Loop configuration when you connect two hosts to channel 0 (both FC0 ports of either controller), or channel 1 (both FC1 ports of either controller) on a Sun StorEdge 3511 SATA array, you must use host filtering if you want to control host access to storage. Refer to the user documentation for your clustering software to determine whether the clustering software can manage host access in this configuration.

Storage Area Networking

Combining storage switches with a Sun StorEdge 3500 series array configuration creates a SAN, increasing the number of servers that can be connected. Essentially, the maximum number of servers that can be connected to the SAN becomes equal to the number of available storage switch ports. Storage switches generally include the ability to manage and monitor the Fibre Channel networks they create, which can reduce storage management workloads in multiple server environments.

Sun StorEdge 3500 series arrays are designed to be deployed in SANs based on switched Fibre Channel fabrics. In a SAN scenario, the server HBAs are connected to one side of the fabric and storage is connected to the other. A SAN fabric automatically routes Fibre Channel packets between ports on one or many Fibre Channel switches.

SAN deployment enables Sun StorEdge 3500 series arrays to be used by a larger number of hosts. This storage strategy tends to utilize storage resources more effectively and is commonly referred to as *storage consolidation*.

The number of hosts that can effectively share one Sun StorEdge 3500 series array depends on several factors, such as the type of host application, bandwidth requirements, and the need for concurrent IOPs. Since most applications have moderate performance needs, it is quite feasible to have several hosts sharing the same Sun StorEdge 3510 FC or 3511 SATA array controller, with the following network characteristics:

- The applications used with the StorEdge 3511 SATA array are characteristically higher bandwidth and lower IOPS. Sharing a StorEdge 3511 SATA array amongst several servers will depend upon the division of performance among the applications installed on each server actively accessing the StorEdge 3511 SATA array.
- The StorEdge 3510 FC Array is better suited for larger configurations with broader application requirements and higher IOPS.

A SAN can also support multiple Sun StorEdge 3500 series arrays. Increasing the number of arrays makes more performance and capacity available within a storage network for sharing among the servers connected to the SAN. A SAN also provides great flexibility in how storage capacity can be allocated among servers and eliminates cabling changes when reallocation of storage becomes necessary.

When a Sun StorEdge 3500 series array is deployed in a SAN, both point-to-point (full fabric) and arbitrated loop (public loop) modes are supported. Point-to-point mode allows for slightly better full duplex performance but limits the total number of addressable LUNs to 128, or to 64 when redundant pathing is used.

Scaling Capacity

Sun StorEdge 3500 series arrays are available in a number of configurations to address a broad range of storage capacities.

- The StorEdge 3510 FC array is available with 15k rpm FC disks to address the needs of high IOPS applications.
- The StorEdge 3511 SATA array is available in larger capacities to better address the needs of the bulk storage market. The additional storage capacity of the StorEdge 3511 SATA array provides more storage for each server attached. Larger capacity disks do a poor job of supporting larger groups of servers, so using the extra capacity to spread among more servers would result in extremely poor performance and possible loss of data.

Base systems include single or redundant controllers and a choice of five or twelve disks. This results in storage capacities as small as:

- **Sun StorEdge 3510 FC Array.** 0.73 Tbyte with five 146-Gbyte disks
- **Sun StorEdge 3511 SATA Array.** 1.25 Tbyte with five 250-Gbyte disks

Capacities can be as large as:

- **Sun StorEdge 3510 FC Array.** 1.75 Tbyte with twelve 146-Gbyte disks
- **Sun StorEdge 3511 SATA Array.** 3.00 Tbyte with twelve 250-Gbyte disks

Additional storage capacity can be dynamically created, starting with a system with five disks and then adding one or more disks. Expansion units can be dynamically added to base systems when more storage capacity is required than a single Sun StorEdge 3500 series array can provide.

Note – Sun StorEdge 3510 FC Arrays can only be connected with Sun StorEdge 3510 FC Expansion Units. Similarly, Sun StorEdge 3511 SATA arrays can only be connected with Sun StorEdge 3511 SATA expansion units. Arrays and expansion units for the two systems can not be mixed.

Sun StorEdge 3500 series arrays remain single storage systems as expansion units are added, even though there are multiple interconnected physical units. Expansion units simply add bays to base units to increase the total number of disks that can be supported. A fully configured system can support:

- **Sun StorEdge 3510 FC Array.** Using one base unit and eight expansion units (108 disks) with 146-Gbyte disks, a maximum of storage capacity of 15.668 Tbyte.
- **Sun StorEdge 3511 SATA Array.** Using one base unit and five expansion units (72 disks) with 250-Gbyte disks, a maximum storage capacity of 16 Tbyte.

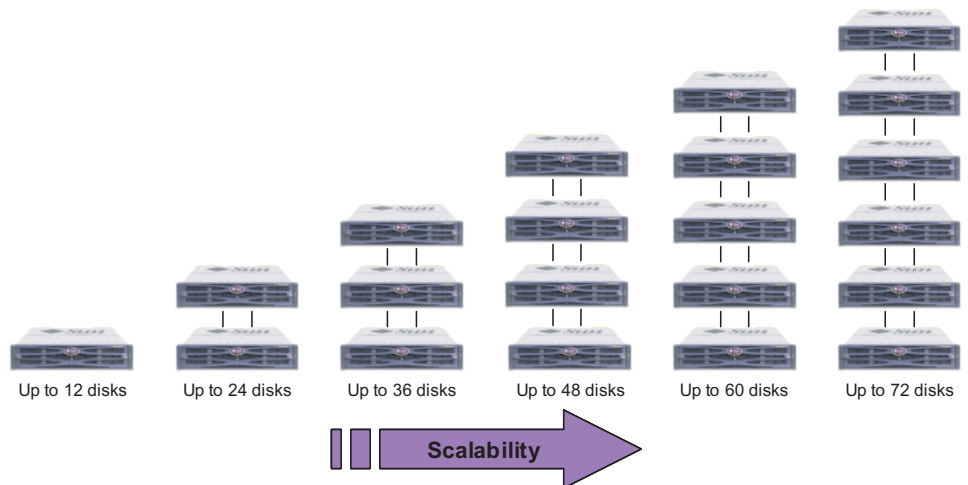


FIGURE 3 Increasing Capacity



Caution – For optimum and successful use of each array, check that you are using the array with the correct applications (see “[Overview](#)” on page 1) and according to the best practices configurations outlined in this manual.

First Steps

There are two simple yet effective approaches for designing a Sun StorEdge 3500 series array solution into your environment. Both methods allow for the rapid estimation of an appropriate DAS or SAN solution. Regardless of which method is used, the storage needs of each application and server involved must be identified to establish the total amount of storage capacity required.

When Designing a Storage Solution for an Existing Environment

The first method works well for existing environments. Start by identifying the number of servers that can immediately benefit from the storage the Sun StorEdge 3500 series array provides. A SAN solution can provide the necessary connectivity to support:

- **Sun StorEdge 3510 FC Array.** Five or more servers. If there are four or fewer servers, a DAS solution is sufficient.
- **Sun StorEdge 3511 SATA Array.** Seven or more servers. If there are six or fewer servers, a DAS solution is sufficient. If you connect two servers to channel 0 or to channel 1, use host filtering if you want to control host access to logical drives.

With either the Sun StorEdge 3510 FC array or the Sun StorEdge 3511 SATA array, a SAN solution can be a powerful option, even when the array is connected to a number of servers that can otherwise be supported in a DAS solution. Combining both the Sun StorEdge 3510 FC array and Sun StorEdge 3511 SATA array on the same SAN allows for a tiered storage strategy, using StorEdge SAM-FS as the data mover amongst tiers. Determine how much storage is currently accessible to these servers and plan for that total capacity as the minimum amount of Sun StorEdge 3500 series disk capacity needed.

When Designing a New Storage Solution

Another technique involves matching a particular environment to one of the best practices solutions described in this document. This approach works particularly well with new deployments, but it can be used for existing environments as well. Take note of special features, such as the number of connections between servers and

storage. While these solutions do not match every environment exactly, use the closest one as a design blueprint that can be customized to suite your particular environment. For environments with different server configurations, choose the solution that best matches the servers whose applications are mission-critical or most important.

General Configuration Considerations

The entry-level configuration for an FC array uses only one RAID controller. If this configuration is used, two single-controller arrays should use host-based mirroring to ensure high reliability, availability, and serviceability (RAS).

Note – Please refer to Veritas Volume Manager or an equivalent host mirroring application to set up the optimum configuration with single-controller arrays.

It is a best practice to use dual-controller arrays to avoid a single point of failure. A dual-controller FC array features a default active-to-active controller configuration. This configuration improved application availability because, in the unlikely event of a controller failure, the array automatically fails over to a second controller, resulting in no interruption of data flow. Single controller arrays are provided for small configurations requiring fast, scratch disk, as in EDA environments.

Sun StorEdge 3500 series arrays are extremely flexible, but when designing storage solutions remember to keep them as simple as possible. Keep the following suggestions in mind when designing the configuration of a Fibre Channel storage system:

- To ensure power redundancy, connect the two power modules to two separate circuits, such as one commercial circuit and one UPS.
- In a single-controller configuration, disable the Write-Back Cache feature to avoid the possibility of data corruption in the event of a controller failure. This will have a negative effect on performance. To avoid either issue, use dual controllers.
- Using two single-controllers in a clustering environment with host-based mirroring provides some of the advantages of using a dual controller. However, you still need to disable the Write-Back Cache in case one of the single controllers fails and you want to avoid data corruption. For this reason, a dual controller configuration is preferable.
- Prior to creating logical drives and mapping them to host channels, set the appropriate cache optimization, Fibre Channel protocol, and controller channel IDs. Reset the controller after these configuration parameters have been set.
- For best performance and RAS, create logical drives across expansion units.

- To avoid disruptions to other hosts sharing the same array, do not share a logical drive among multiple hosts.
- Use either local or global spare drives when creating logical drives. Any free drive can be designated as a spare and more than one drive can be used as a spare.
- Use dual pathing for each LUN and use Sun StorEdge Traffic Manager software to provide load balancing across controller ports for increased performance.
- The maximum number of LUNs when using point-to-point protocol is 128 for single-path configurations and 64 for dual-path configurations.
- Power-up the equipment in the following order:
 - a. Expansion units
 - b. RAID array
 - c. Host computers
- Install Sun SAN Foundation software along with the latest patches and firmware.
- Install Sun Storage Automated Diagnostic Environment (StorADE) 2.3 software. (See FIN# I0959-1).
- Connect the Ethernet management ports to a private Ethernet network (Sun Alert# 26464).
- After completing the configuration of a Sun StorEdge 3500 series array, the configuration should be saved using the firmware "Save nvram to disks" menu option and the Sun StorEdge Configuration Service Console's "save configuration" utility.

Configuring an Array's RCCOM Channel

Redundant controller communication (RCCOM) provides the communication channels by which two controllers in a redundant RAID array communicate with one another. This communication allows the controllers to monitor each other, and includes configuration updates, and control of cache. By default, channels 2 and 3 are configured as DRV + RCCOM (Drive and RCCOM). In this configuration, RCCOM is distributed over all DRV + RCCOM channels. However, when host channels remain unused two alternative configurations are available.

Using Four DRV + RCCOM Channels

If only channels 0 and 1 are used for communication with servers, channels 4 and 5 can be configured as DRV + RCCOM, thus providing four DRV + RCCOM channels (channels 2, 3, 4 and 5). An advantage of this configuration is channels 4 and 5 are still available for connection of expansion units. The impact of RCCOM is reduced

because it is now distributed over four channels instead of two. If at a later time you choose to add an expansion unit, it will not be necessary to interrupt service by resetting the controller after reconfiguring a channel.

To configure channels 4 and 5 as additional DRV + RCCOM channels:

1. **Access the firmware application of the array.**
2. **On the Main Menu, choose “view and edit Scsi channels.”**
3. **Select channel 4.**
4. **Choose “channel Mode → Drive + RCCOM.”**
5. **Choose Yes to confirm.**
6. **Choose No decline the controller reset.**
7. **Select channel 5.**
8. **Choose “channel Mode → Drive + RCCOM.”**
9. **Choose Yes to confirm.**
10. **Choose Yes to reset the controller.**

Using Channels 4 and 5 as RCCOM Channels

When only channels 0 and 1 are used for communication with servers, another option is to assign channels 4 and 5 as dedicated RCCOM channels. This reduces the impact of RCCOM on the drive channels by removing RCCOM from drive channels 2 and 3. In this configuration, however, channels 4 and 5 cannot be used to communicate with hosts or to attach expansion modules. To configure channels 4 and 5 as dedicated RCCOM channels:



Caution – If you later reconfigure channels 4 and 5 as host or drive channels, restore channels 2 and 3 as DRV + RCCOM channels or the RAID array will no longer operate.

1. **Access the firmware application of the array.**
2. **On the Main Menu, choose “view and edit Scsi channels.”**
3. **Select channel 4.**
4. **Choose “channel Mode → RCCOM.”**
5. **Choose Yes to confirm.**

6. Choose **No** to decline the controller reset.
7. Select channel 5.
8. Choose “channel Mode → RCCOM.”
9. Choose **Yes** to confirm.
10. Choose **No** to decline the controller reset.
11. Select channel 2.
12. Choose “channel Mode → drive.”
13. Choose **Yes** to confirm.
14. Choose **No** to decline the controller reset.
15. Select channel 3.
16. Choose “channel Mode → drive.”
17. Choose **Yes** to confirm.
18. Choose **Yes** to reset the controller.

Cautions



Caution – The UNIX and Solaris `format` and Solaris `probe-scsi-all` commands do not display all mapped LUNs if there is not a partition or logical drive mapped to LUN 0.



Caution – Do not use both in-band and out-of-band connections at the same time to manage the array. Otherwise conflicts between multiple operations might occur.



Caution – Suspend all I/O activities and unmount filesystems before issuing download, reset, or shutdown commands.



Caution – Do not connect HBAs of different speeds to the same FC channel. Otherwise data corruption might occur.

Non-Redundant DAS Configurations

Note – Using single connections between Fibre Channel arrays and servers creates single points of failure (SPOF) that can cause interruptions in the event a connection becomes unreliable or fails. This is not a recommended configuration unless host-based mirroring is utilized to protect against single points of failure.




FIGURE 4 Three Non-Redundant DAS Configurations

Note – The Sun StorEdge 3511 SATA array cannot be used to store single instances of data but can be used in multi-path and multi-host configurations where the array has a backup or archival role.

TABLE 7 Configuration Overview for Non-Redundant DAS

	Single Server Configurations	Dual Server Configurations	Quad Server Configurations
Number of Servers	1	2	4
RAID Enclosures	1	1	1
Expansion Units	As needed	As needed	One or more
Number of Controllers	1	1	1
Number of Disks	5 or more	12 or more	24 or more
Cache Optimization (Sun StorEdge 3510 FC array)	Random or sequential	Random or sequential	Random or sequential
Cache Optimization (Sun StorEdge 3511 SATA array)	Sequential	Sequential	Sequential
RAID Levels	Application-dependent	Application-dependent	Application-dependent
Fibre connection option	Loop only	Loop only	Loop only
Drive Configuration	One or more logical drives plus one or more global spare ¹	Two or more logical drives plus two or more global spares ¹	Four or more logical drives plus four or more global spares ¹
Host Adapters per server	Single-Port 2-Gbit FC	Single-Port 2-Gbit FC	Single-Port 2-Gbit FC
Traffic Manager	Not required	Not required	Not required
Storage Switches	Not required	Not required	Not required

¹ Refer to Table 1 to determine how many physical drives can be included in one logical drive.

Non-Redundant DAS Tips and Techniques

- A Sun StorEdge Fibre Channel array with a single controller can be configured to support up to four host connections (Sun StorEdge 3510 FC array) or 6 host connections (Sun StorEdge 3511 SATA array). These connections can be used in pairs, individually, or in any combination of both.
- You will need to add SFP modules to support more than two host connections to the Sun StorEdge 3500 series array. Add:
 - One SFP module to support three connections
 - Two SFP modules to support four connections

- Three SFP modules to support five connections (Sun StorEdge 3511 SATA array only)
- Four SFP modules to support six connections (Sun StorEdge 3511 SATA array only)
- Using two single-port FC host bus adapters (HBAs) or a dual-port 2-Gbit FC HBA in single-server or dual-server configurations makes optimum use of the Fibre Channel array's performance. Mapping logical drive partitions to two paths while using multipathing software and load balancing provides the best performance.

Non-Redundant DAS Setup Details

FIGURE 5 and FIGURE 6 show a Sun StorEdge 3500 series arrays in a non-redundant configurations.




FIGURE 5 Sun StorEdge 3510 FC Array Non-Redundant DAS Connections




FIGURE 6 Sun StorEdge 3511 SATA Array Non-Redundant DAS Connections

TABLE 8 Setup Summary for Non-Redundant DAS

Channel Number	Primary ID Number	Secondary ID Number
0	40	N/A
1	43	N/A
2	14	N/A
3	14	N/A
4	44	N/A
5	47	N/A

The general procedure for creating this configuration follows.

1. Check the position of installed SFP modules. Move or add SFP modules as necessary to support the connections needed.
2. Connect expansion units if needed.
 - Sun StorEdge 3510 FC Array. See [“Sun StorEdge 3510 FC Array High Capacity Configurations” on page 44.](#)
 - Sun StorEdge 3511 SATA Array. See [“Sun StorEdge 3511 SATA Array High Capacity Configurations” on page 53](#)

3. **Configure cache optimization. If you are configuring a Sun StorEdge 3511 SATA array, confirm that cache optimization is sequential.**
4. **Create one Logical Drive for each server and configure spare disks.**
5. **Map Logical Drive 0 to controller channel 0.**
6. **Map Logical Drive 1 (if created) to controller channel 5.**
7. **Map Logical Drive 2 (if created) to controller channel 1.**
8. **Map Logical Drive 3 (if created) to controller channel 4.**
9. **Connect the first server to controller port FC0.**
10. **Connect the second server (if needed) to controller port FC5.**
11. **Connect the third server (if needed) to controller port FC1.**
12. **Connect the fourth server (if needed) to controller port FC4.**
13. **After completing the configuration, save the configuration using “Save nvram to disks” menu option in the firmware application and the “save configuration” command in Sun StorEdge Configuration Service.**

High-Availability DAS Configurations




FIGURE 7 High-Availability DAS Configurations

Note – The Sun StorEdge 3511 SATA array cannot be used to store single instances of data but can be used in multi-path and multi-host configurations where the array has a backup or archival role.

TABLE 9 Configuration Overview for High Availability DAS

	Single Server Configurations	Dual Server Configurations	Quad Server Configurations
Number of Servers	1	2	4
RAID Enclosures	1	1	1
Expansion Units	As needed	As needed	One or more
Number of Controllers	2	2	2

TABLE 9 Configuration Overview for High Availability DAS (*Continued*)

	Single Server Configurations	Dual Server Configurations	Quad Server Configurations
Number of Disks	5 or more	12 or more	24 or more
Cache Optimization (Sun StorEdge 3510 FC Array)	Random or sequential	Random or sequential	Random or sequential
Cache Optimization (Sun StorEdge 3511 SATA Array)	Sequential	Sequential	Sequential
RAID Levels	Application-dependent	Application-dependent	Application-dependent
Fibre connection option	Loop only	Loop only	Loop only
Drive Configuration	One or more logical drives plus one or more global spares ¹	Two or more logical drives plus two or more global spares ¹	Four or more logical drives plus four or more global spares ¹
Host Adapters per server	Two Single-Port 2-Gbit FC	Two Single-Port 2-Gbit FC	Two Single-Port 2-Gbit FC
Traffic Manager	Required	Required	Required
Storage Switches	Not required	Not required	Not required

¹ Refer to Table 1 to determine how many physical drives can be included in one logical drive.

High-Availability DAS Tips and Techniques

- A Sun StorEdge 3510 FC array with two controllers can be configured to support up to eight host connections. The Sun StorEdge 3511 SATA array can be configured to support up to 12 host connections. These connections can be used in pairs for redundancy, individually or in any combination of both.
- You will need to add SFP modules to support more than four host connections to the array. For example, add two SFP modules to support six connections and add four SFP modules to support eight connections.
- Using two single-port 2-Gbit FC host adapters in a high-availability configuration will make optimum use of a Sun StorEdge Fibre Channel array's redundancy. Mapping logical drive partitions to two paths while using multipathing software provides the best redundancy.
- For complete redundancy and high availability, use multipathing software such as Sun StorEdge Traffic Manager. To configure multipathing:
 - Establish two connections between a server and a Sun StorEdge Fibre Channel array.
 - Install and enable the software on the server.

- Map the logical drive to both of the controller channels to which the server is connected.

Standard High-Availability DAS Setup Details

FIGURE 8 and FIGURE 9 show a Sun StorEdge 3510 FC array and a Sun StorEdge 3511 SATA array in a high-availability DAS configuration.




FIGURE 8 Sun StorEdge 3510 FC Array High-Availability DAS Connections

Note – The Sun StorEdge 3511 SATA array cannot be used to store single instances of data but can be used in multi-path and multi-host configurations where the array has a backup or archival role.




FIGURE 9 Sun StorEdge 3511 SATA Array High-Availability DAS Connections

TABLE 10 Setup Summary for High-Availability DAS

Channel Number	Primary ID Number	Secondary ID Number
0	40	N/A
1	N/A	42
2	14	15
3	14	15
4	44	N/A
5	N/A	46

The general procedure for creating this configuration follows.

1. Check the position of installed SFP modules. Move them as necessary to support the connections needed.
2. Connect expansion units if needed.
 - Sun StorEdge 3510 FC Array. See [“Sun StorEdge 3510 FC Array High Capacity Configurations”](#) on page 44.

- Sun StorEdge 3511 SATA Array. See [“Sun StorEdge 3511 SATA Array High Capacity Configurations”](#) on page 53
3. Configure cache optimization. If you are configuring a Sun StorEdge 3511 SATA array, confirm that cache optimization is sequential.
 4. Ensure that the fibre connection is set to loop mode.
 5. Configure target IDs.
 6. Create one Logical Drive for each server and configure spare disks.
 7. Map Logical Drive 0 to channels 0 and 5 of the primary controller.
 8. Map Logical Drive 1 (if created) to channels 1 and 4 of the secondary controller.
 9. Map Logical Drive 2 to channels 0 and 5 of the primary controller.
 10. Map Logical Drive 3 (if created) to channels 1 and 4 of the secondary controller.
 11. Connect the first server (Server 1) to port FC0 of the upper controller and port FC5 of the lower controller.
 12. Connect the second server (Server 2), if needed to port FC1 of the lower controller and port FC4 of the upper controller.
 13. Connect the third server (Server 3), if needed to port FC0 of the lower controller and port FC5 of the upper controller.
 14. Connect the fourth server (Server 4), if needed to port FC1 of the upper controller and port FC4 of the lower controller.
 15. Install and enable multipathing software on each connected server.
 16. After completing the configuration, save the configuration using “Save nvram to disks” menu option in the firmware application and the “save configuration” command in Sun StorEdge Configuration Service.

Six Node High-Availability DAS Setup Details

FIGURE 10 show a Sun StorEdge 3511 SATA array connected to six servers in a high-availability DAS configuration.




FIGURE 10 Sun StorEdge 3511 SATA Array Six Node High-Availability Connections

Note – The Sun StorEdge 3511 SATA array cannot be used to store single instances of data but can be used in multi-path and multi-host configurations where the array has a backup or archival role.

The general procedure for creating this configuration follows.

1. Add SFPs to support twelve connections.
2. Connect expansion units if needed. See [“Sun StorEdge 3511 SATA Array High Capacity Configurations”](#) on page 53
3. Confirm that cache optimization is sequential.

4. Ensure that the fibre connection is set to loop mode.
5. Configure target IDs.
6. Create one logical drive for each server and configure spare disks.
7. Map Logical Drive 0 to channels 0 and 5 of the primary controller.
8. Map Logical Drive 1 to channels 1 and 4 of the secondary controller.
9. Map remaining even-numbered logical drives to channels 0 and 5 of the primary controller.
10. Map remaining odd-numbered logical drives to channels 1 and 4 of the secondary controller.
11. Connect the first server (Server 1) to the left-most port FC0 of the upper controller and port FC4 of the lower controller.
12. Connect the second server (Server 2) to the right-most port FC0 of the upper controller and port FC5 of the lower controller.
13. Connect the third server (Server 3) to port FC4 of the upper controller and the left-most FC1 port of the lower controller.
14. Connect the fourth server (Server 4) to port FC5 of the upper controller and the right-most FC1 port of the lower controller.
15. Connect the fifth server (Server 5) to the left-most FC1 port of the upper controller and the left-most FC0 port of the lower controller.
16. Connect the sixth server (Server 6) to the right-most FC1 port of the upper controller and the right-most FC0 port of the lower controller.
17. Install and enable multipathing software on each connected server.
18. After completing the configuration, save the configuration using "Save nvram to disks" menu option in the firmware application and the "save configuration" command in Sun StorEdge Configuration Service.

Full-Fabric SAN Configurations




FIGURE 11 Typical Full-Fabric SAN Configuration

Note – The Sun StorEdge 3511 SATA array cannot be used to store single instances of data but can be used in multi-path and multi-host configurations where the array has a backup or archival role.

TABLE 11 Configuration Overview for a Full-Fabric SAN

	Small Configuration	Medium Configuration	Large Configuration
Sun StorEdge 3511 SATA array			
Number of Servers	2 to 4	2 to 8	2 to 16
Cache Optimization	Sequential	Sequential	Sequential
Sun StorEdge 3510 FC array			
Number of Servers	2 to 4	2 to 14	2 to 30
Cache Optimization	Random or sequential	Random or sequential	Random or sequential
Common Parameters			
RAID Enclosures	1	1	1
Expansion Units	As needed	As needed	As needed
Number of Controllers	2	2	2
Number of Disks	12 or more	12 or more	12 or more
RAID Levels	Application-dependent	Application-dependent	Application-dependent
Fibre connection option	Point-to-point	Point-to-point	Point-to-point
Drive Configuration	Two or more LUNs with two or more global spares ¹	Two or more LUNs with two or more global spares ¹	Two or more LUNs with two or more global spares ¹
Host Adapters per server	Dual-port 2-Gbit FC	Dual-port 2-Gbit FC	Dual-port 2-Gbit FC
Traffic Manager	Required	Required	Required
Storage Switches	Two 8-port 2-Gbit FC fabric switches	Two 16-port 2-Gbit FC fabric switches	Two 32-port 2-Gbit FC fabric switches

¹ Refer to Table 1 to determine how many physical drives can be included in one logical drive.

Full-Fabric SAN Tips and Techniques

- In the fabric SAN configuration, the switches communicate with the array host ports using a fabric point-to-point (F_port) mode. This enables transparent controller fail-over and fail-back without server-resident software. However, supporting hot-swap servicing of a failed controller requires the use of multipathing software, such as the Sun StorEdge Traffic Manager, on the connected servers.

- Use of fabric point-to-point (F_port) connections between an array and fabric switches limits to 128 the total number of LUNs that can be presented. Fibre channel standards allow only one ID per port when operating point-to-point protocols, resulting in a maximum of four IDs, with a maximum of 32 LUNs each, supporting up to 128 LUNs.
- When a fabric switch is connected to one port of channel 0 or channel 1 of a Sun StorEdge 3511 SATA array, no connections can be made with the other three ports of that channel. If channel 0 (port FC0) is connected to a fabric switch, for example, the second port for channel 0 on that controller, and the FC0 ports on a redundant controller, cannot be used. Similarly, if channel 1 (port FC1) is connected to a fabric switch, the second FC1 port on that controller and the FC1 ports on a redundant controller, cannot be used.

Full-Fabric SAN Setup Details

FIGURE 12 and FIGURE 13 show a Sun StorEdge 3510 FC array and a Sun StorEdge 3511 SATA array in a full-fabric SAN configuration.




FIGURE 12 Sun StorEdge 3510 FC Array Full-fabric SAN Connections




FIGURE 13 Sun StorEdge 3511 SATA Array Full-fabric SAN Connections

TABLE 12 Setup Summary for Full Fabric SAN

Channel Number	Primary ID Number	Secondary ID Number
0	40	N/A
1	N/A	42
2	14	15
3	14	15
4	44	N/A
5	N/A	46

The general procedure for creating this configuration follows.

1. Check the position of installed SFP modules. Move them as necessary to support the connections needed.
2. Connect expansion units if needed.
 - Sun StorEdge 3510 FC Array. See [“Sun StorEdge 3510 FC Array High Capacity Configurations”](#) on page 44.

- Sun StorEdge 3511 SATA Array. See [“Sun StorEdge 3511 SATA Array High Capacity Configurations”](#) on page 53
3. Configure cache optimization. If you are configuring a Sun StorEdge 3511 SATA array, confirm that cache optimization is sequential.
 4. Ensure that the fibre connection is set to point-to-point.
 5. Ensure only one target ID per channel is configured.
 6. Create at least two logical drives and configure spare disks.
 7. Create one or more logical drive partitions for each server.
 8. Map Logical Drive 0 to channels 0 and 4 of the primary controller
 9. Map Logical Drive 1 to channels 1 and 5 of the secondary controller.
 10. If more than two logical drives were created, map even-numbered logical drives to channels 0 and 4 of the primary controller and odd-numbered logical drives to channels 1 and 5 of the secondary controller.
 11. Connect the first switch to port FC0 of the upper controller and port FC1 of the lower controller.
 12. Connect the second switch to port FC4 of the lower controller and port FC5 of the upper controller.
 13. Connect each server to each switch.
 14. Install and enable multipathing software on each connected server.
 15. After completing the configuration, save the configuration using “Save nvram to disks” menu option in the firmware application and the “save configuration” command in Sun StorEdge Configuration Service.

High-Performance SAN Configurations




FIGURE 14 Typical High-Performance SAN Configurations

Note – The Sun StorEdge 3511 SATA array cannot be used to store single instances of data but can be used in multi-path and multi-host configurations where the array has a backup or archival role.

TABLE 13 Configuration Overview for High-Performance SAN

	Small Configuration	Medium Configuration	Large Configuration
Number of Servers (Sun StorEdge 3511 SATA array)	2 to 4	2 to 4	2 to 8
Number of Servers (Sun StorEdge 3510 FC array)	2 to 4	2 to 14	2 to 62
RAID Enclosures	1	1	1
Expansion Units	As needed	As needed	As needed
Number of Controllers	2	2	2

TABLE 13 Configuration Overview for High-Performance SAN (*Continued*)

	Small Configuration	Medium Configuration	Large Configuration
Number of Disks	12 or more	12 or more	12 or more
Cache Optimization (Sun StorEdge 3510 FC array)	Random or sequential	Random or sequential	Random or sequential
Cache Optimization (Sun StorEdge 3511 SATA array)	Sequential	Sequential	Sequential
RAID Levels	Application-dependent	Application-dependent	Application-dependent
Fibre connection option	Point-to-point only	Point-to-point only	Point-to-point only
Drive configuration	Two or more LUNs with two or more global spares ¹	Two or more LUNs with two or more global spares ¹	Two or more LUNs with two or more global spares ¹
Host Adapters per server	Dual-port 2-Gbit FC	Dual-port 2-Gbit FC	Dual-port 2-Gbit FC
Traffic Manager	Required	Required	Required
Storage Switches	Two 8-port 2-Gbit FC fabric switches	Two 16-port 2-Gbit FC fabric switches	Two 32-port 2-Gbit FC fabric switches

¹ Refer to Table 1 to determine how many physical drives can be included in one logical drive.

High-Performance SAN Tips and Techniques

- In the fabric SAN configuration, the switches communicate with the array host ports using a fabric point-to-point (F_port) mode. This enables transparent controller fail-over and fail-back without server-resident software. However, supporting hot-swap servicing of a failed controller requires the use of multipathing software, such as the Sun StorEdge Traffic Manager, on the connected servers.
- Use of fabric point-to-point (F_port) connections between an array and fabric switches limits to 128 the total number of LUNs that can be presented. Fibre channel standards allow only one ID per port when operating point-to-point protocols, resulting in a maximum of four IDs, with a maximum of 32 LUNs each, supporting up to 128 LUNs.
- When a fabric switch is connected to one port of channel 0 or channel 1 of a Sun StorEdge 3511 SATA array, no connections can be made with the other three ports of that channel. If channel 0 (port FC0) is connected to a fabric switch, for example, the second port for channel 0 on that controller, and the FC0 ports on a

redundant controller, cannot be used. Similarly, if channel 1 (port FC1) is connected to a fabric switch, the second port for channel 1 on that controller and both FC1 ports on a redundant controller, cannot be used.

High-Performance SAN Setup Details

FIGURE 15 and FIGURE 16 shows a Sun StorEdge 3510 FC array and a Sun StorEdge 3511 SATA array in a high-performance SAN configuration.




FIGURE 15 Sun StorEdge 3510 FC Array High-Performance SAN Connections




FIGURE 16 Sun StorEdge 3511 SATA Array High-Performance SAN Connections

TABLE 14 Setup Summary for a High-Performance SAN

Channel Number	Primary ID Number	Secondary ID Number
0	40	N/A
1	N/A	42
2	14	15
3	14	15
4	44	N/A
5	N/A	46

The general procedure for creating this configuration follows.

1. Check the position of installed SFP modules. Move them as necessary to support the connections needed.
2. Connect expansion units if needed.
 - Sun StorEdge 3510 FC Array. See [“Sun StorEdge 3510 FC Array High Capacity Configurations”](#) on page 44.

- Sun StorEdge 3511 SATA Array. See [“Sun StorEdge 3511 SATA Array High Capacity Configurations”](#) on page 53
3. Configure cache optimization. If you are configuring a Sun StorEdge 3511 SATA array, confirm that cache optimization is sequential.
 4. Ensure that the fibre connection set to point-to-point mode.
 5. Configure target IDs.
 6. Create at least two logical drives and configure spare disks.
 7. Create one or more logical drive partitions for each server.
 8. Map Logical Drive 0 to channels 0, 1, 4 and 5 of the primary controller
 9. Map Logical Drive 1 to channels 0, 1, 4 and 5 of the secondary controller.
 10. If more than two logical drives were created, map even-numbered logical drives to channels 0, 1, 4 and 5 of the primary controller and odd-numbered logical drives to channels 0, 1, 4 and 5 of the secondary controller.
 11. Connect the first switch to port FC0 of the upper controller and port FC1 of the lower controller.
 12. Connect the second switch to port FC4 of the lower controller and port FC5 of the upper controller.
 13. Connect each server to each switch.
 14. Install and enable multipathing software on each connected server.
 15. After completing the configuration, save the configuration using “Save nvram to disks” menu option in the firmware application and the “save configuration” command in Sun StorEdge Configuration Service.

Scaling Arrays Into High Capacity Configurations

Note – High capacity Sun StorEdge 3500 series configurations are supported, with certain limitations. For example, only the sequential cache optimization mode can be used. Other important limitations may exist. Consult the Sun StorEdge 3510 and 3511 FC array documentation for additional information. A Sun StorEdge 3510 FC array with more than two expansion units is a high capacity configuration. A Sun StorEdge 3511 SATA array with one expansion unit or more is a high capacity configuration.

Carefully consider the limitations of these special high-capacity configurations. Using multiple arrays connected to the same SAN normally provides significantly better performance than one high-capacity configuration.

Limitations

- Random optimization should not be used with high-capacity configurations. Use sequential optimization instead. Random cache optimization greatly reduces the maximum number of supported disks.
- Only the cabling configurations detailed in this section are supported when configuring high-capacity Sun StorEdge 3500 series arrays.
- Maximize the size of each logical drive (up to 1,908 Gbyte for RAID 5 logical drives) before creating another logical drive. Otherwise the maximum number of eight logical drives might limit future expansion.
- Limiting the maximum number of expansion units on a Sun StorEdge 3510 FC array to 7 provides much greater configuration flexibility. Doing so enables the use of up to 96 disks.
- If you stack a tabletop array and expansion units on top of each other, you can stack up to five physical units. Do not stack more than five physical units.
- Many of these configurations require the use of optional extended-length cables, part number X9732A. Other items may also be required. Refer to the *Sun StorEdge 3000 Family FRU Installation Guide* for information about supported cables, SFPs, and other user-replaceable items.

Sun StorEdge 3510 FC Array High Capacity Configurations

Sun StorEdge 3510 FC arrays typically allow the connection of up to two expansion units to support a maximum of 36 disks. However, you can create larger configurations that support as many as eight expansion units and up to 108 disks if you use the guidelines in this section.

See [“Sun StorEdge 3511 SATA Array High Capacity Configurations” on page 53](#) for information about adding expansion units to a Sun StorEdge 3511 SATA array.

Recommendations

- These configurations work best when you select the highest-capacity disk available without exceeding maximum limits set by the firmware. Refer to the *Sun StorEdge 3000 Family FRU Installation Guide* for information about supported disk drives, cables, SFPs, and other user-replaceable items.
- The use of a full-fabric SAN configuration is highly recommended for host connections, as shown in [“Full-Fabric SAN Configurations” on page 33.](#)

Note – The following high-capacity configurations illustrate cabling techniques for drive ports. They must be combined with a standard configuration shown in this document, such as shown in [“High-Performance SAN Configurations” on page 38.](#)

Connecting One Expansion Unit




FIGURE 17 Sun StorEdge 3510 FC Array Configuration with One Expansion Unit

Connecting Two Expansion Units




FIGURE 18 Sun StorEdge 3510 FC Array Configuration with Two Expansion Units

Connecting Three Expansion Units




FIGURE 19 Sun StorEdge 3510 FC Array Configuration with Three Expansion Units

Connecting Four Expansion Units




FIGURE 20 Sun StorEdge 3510 FC Array Configuration with Four Expansion Units

Connecting Five Expansion Units

This high-capacity configuration using five expansion units requires the configuration of RAID channels 4 and 5 as drive channels. This restricts host connections to RAID channels 0 and 1, affecting supported host configurations. Use four expansion units or fewer when possible.




FIGURE 21 Sun StorEdge 3510 FC Array Configuration with Five Expansion Units

Connecting Six Expansion Units

This high-capacity configuration using six expansion units requires the configuration of RAID channels 4 and 5 as drive channels. This restricts host connections to RAID channels 0 and 1, affecting supported host configurations. Use four expansion units or fewer when possible.




FIGURE 22 Sun StorEdge 3510 FC Array Configuration with Six Expansion Units

Connecting Seven Expansion Units

This high-capacity configuration using seven expansion units requires the configuration of RAID channels 4 and 5 as drive channels. This restricts host connections to RAID channels 0 and 1, affecting supported host configurations. Use four expansion units or fewer when possible.




FIGURE 23 Sun StorEdge 3510 FC Array Configuration with Seven Expansion Units

Connecting Eight Expansion Units

This high-capacity configuration using eight expansion units requires the configuration of RAID channels 4 and 5 as drive channels. This restricts host connections to RAID channels 0 and 1, affecting supported host configurations. Use four expansion units or fewer when possible.

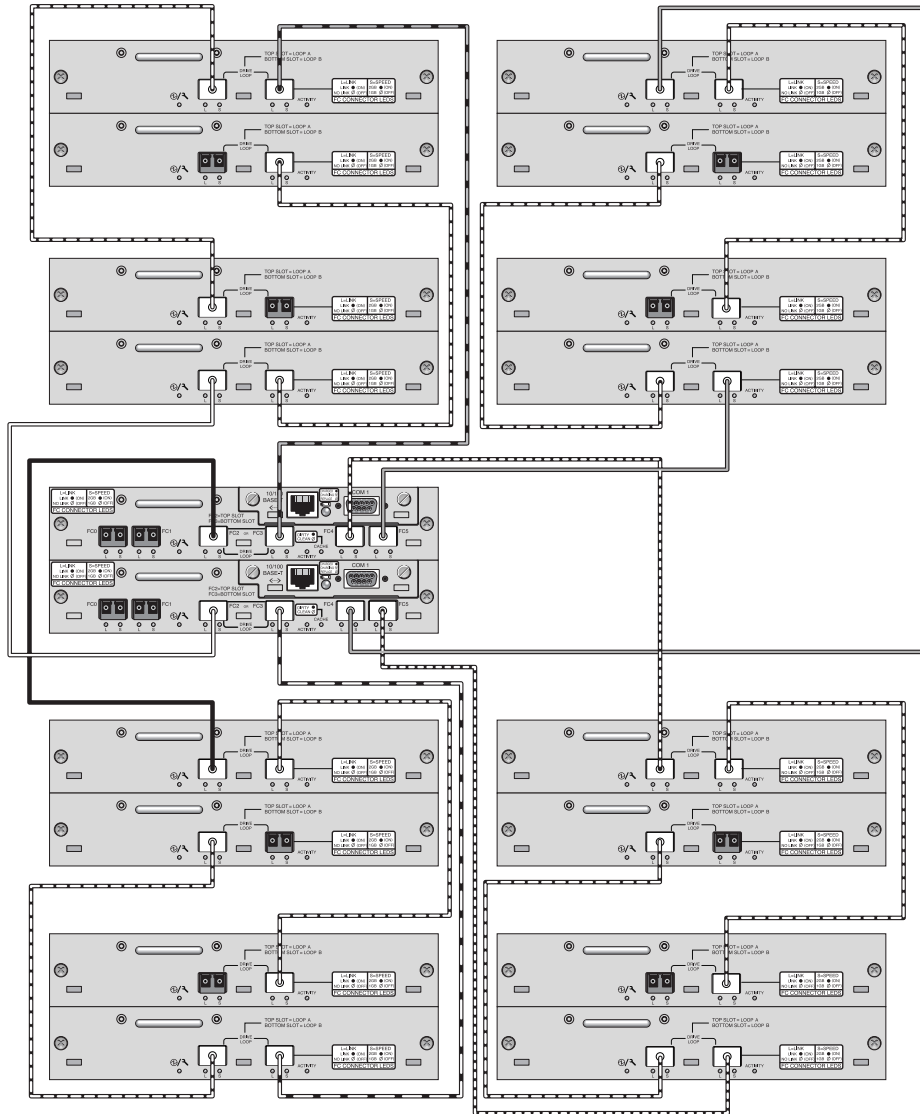


FIGURE 24 Sun StorEdge 3510 FC Array Configuration with Eight Expansion Units

Sun StorEdge 3511 SATA Array High Capacity Configurations

Sun StorEdge 3511 SATA arrays allow the connection of up to five expansion units to support a maximum of 72 disks. You can create large configurations beyond 12 disks if you use the guidelines in this section.

Note – The Sun StorEdge 3511 SATA array cannot be used to store single instances of data but can be used in multi-path and multi-host configurations where the array has a backup or archival role.

See [“Sun StorEdge 3510 FC Array High Capacity Configurations”](#) on page 44 for information about adding more than two expansion units to a Sun StorEdge 3510 FC array.

For additional late-breaking technical and configuration details about the new SATA disk technology and the Sun StorEdge 3511 SATA array, be sure to check the *Sun StorEdge 3511 FC Array with SATA Release Notes* on the Sun product web sites.

Recommendations

- Refer to the *Sun StorEdge 3000 Family FRU Installation Guide* for information about supported disk drives, cables, SFPs, and other user-replaceable items.
- The use of a full-fabric SAN configuration is highly recommended for host connections, as shown in [“Full-Fabric SAN Configurations”](#) on page 33.

Note – The following high-capacity configurations illustrate cabling techniques for drive ports. They must be combined with a standard configuration shown in this document, such as shown in [“High-Performance SAN Configurations”](#) on page 38 and must be used with the network applications appropriate to the array, described in [“Overview”](#) on page 1 and summarized in the [“Comparison of the Sun StorEdge 3510 FC Array and the Sun StorEdge 3511 SATA Array”](#) on page 3.

Connecting One Expansion Unit




FIGURE 25 Sun StorEdge 3511 SATA Array Configuration with One Expansion Unit

Connecting Two Expansion Units




FIGURE 26 Sun StorEdge 3511 SATA Array Configuration with Two Expansion Units

Connecting Three Expansion Units




FIGURE 27 Sun StorEdge 3511 SATA Array Configuration with Three Expansion Units

Connecting Four Expansion Units




FIGURE 28 Sun StorEdge 3511 SATA Array Configuration with Four Expansion Units

Connecting Five Expansion Units

This high-capacity configuration using five expansion units requires the configuration of RAID channels 4 and 5 as drive channels. This restricts host connections to RAID channels 0 and 1, affecting supported host configurations. Use four expansion units or fewer when possible.




FIGURE 29 Sun StorEdge 3511 SATA Array Configuration with Five Expansion Units

Summary

Entry-level, mid-range, and enterprise servers are used for a wide range of applications with distinct storage requirements, so the Sun StorEdge 3500 series arrays feature a modular architecture with flexible configurations. For example, a solution can be deployed as direct-attached storage (DAS) or as part of a storage area network (SAN). Configuration preferences include RAID protection levels, single or redundant controllers, total storage capacity, multipathing and more.

Modularity and flexibility allow Sun StorEdge 3500 series array storage solutions to be quickly and easily adapted to a particular environment.

