



Sun StorEdge™ 3000 Family Best Practices Manual

Sun StorEdge 3510 FC Array

Sun StorEdge 3511 SATA Array

Sun Microsystems, Inc.
www.sun.com

Part No. 816-7325-16
November 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.

Contents

Preface	vii
1. Overview	1
Fibre Channel and SATA Array Architectures	2
Sun StorEdge 3510 FC Array Features	2
Sun StorEdge 3511 SATA Array Features	3
How Sun StorEdge 3511 SATA Arrays Differ From Sun StorEdge 3510 FC Arrays	4
2. Underlying Concepts and Practices	7
Fibre Channel Protocols	7
Supported RAID Levels	8
Logical Drives	8
Cache Optimization	10
Array Management Tools	12
Configuring an Array's RCCOM Channel	13
Using Four DRV + RCCOM Channels	13
Using Channels 4 and 5 as RCCOM Channels	14
Saving and Restoring Configuration Information	15
3. Planning Your Storage Architecture	17

Direct-Attached Storage	18
Storage Area Networking	19
Scaling Capacity	20
First Steps in Designing a Solution	22
Designing a Storage Solution for an Existing Environment	22
Designing a New Storage Solution	23
General Configuration Considerations	23
4. DAS Configurations	27
Single-Controller DAS Configurations	27
Single-Controller DAS Tips and Techniques	28
Single-Controller DAS Setup Details	29
Dual-Controller Multipath DAS Configurations	32
Dual-Controller Multipath DAS Tips and Techniques	33
Sample Dual-Controller Multipath DAS Setup Details	33
5. SAN Configurations	37
Full-Duplex SAN Configurations	37
Full-Duplex SAN Tips and Techniques	38
Sample Full-Duplex SAN Setup Details	38
Redundant Point-to-Point SAN Configurations	42
Redundant Point-to-Point Tips and Techniques	42
Redundant Point-to-Point Setup Details	43
6. Using Multiple Expansion Units in High Capacity Configurations	47
Sun StorEdge 3510 FC Array High-Capacity Configurations	47
Limitations	48
Connecting One Expansion Unit	49
Connecting Two Expansion Units	50
Connecting Three Expansion Units	51

Connecting Four Expansion Units	52
Connecting Five Expansion Units	53
Connecting Six Expansion Units	55
Connecting Seven Expansion Units	57
Connecting Eight Expansion Units	59
7. Sun StorEdge 3511 SATA Array High-Capacity Configurations	61
Connecting One Expansion Unit	62
Connecting Two Expansion Units	63
Connecting Three Expansion Units	64
Connecting Four Expansion Units	65
Connecting Five Expansion Units	66

Preface

This manual gives an overview of Sun StorEdge™ 3510 FC arrays and Sun StorEdge 3511 SATA arrays, and presents several sample storage solutions for entry-level, mid-range, and enterprise servers.



Caution – You should read the *Sun StorEdge 3000 Family Safety, Regulatory, and Compliance Manual* before beginning any procedure in this manual.

How This Book Is Organized

This book includes the following topics:

[Chapter 1](#) provides an overview of the array's features.

[Chapter 2](#) provides an overview of important concepts and practices that underly the configurations you can use.

[Chapter 3](#) helps you determine which direct-attached storage (DAS) and storage area network (SAN) configurations are most appropriate to your requirements.

[Chapter 4](#) presents several sample DAS configurations you can use.

[Chapter 5](#) presents several sample SAN configurations you can use.

[Chapter 6](#) describes high-capacity configurations involving multiple expansion units and presents sample Sun StorEdge 3510 FC configurations that maximize reliability, availability, and serviceability (RAS).

[Chapter 7](#) presents sample high-capacity Sun StorEdge 3511 SATA configurations that maximize reliability, availability, and serviceability (RAS).

Using UNIX Commands

This document might not contain information on basic UNIX® commands and procedures such as shutting down the system, booting the system, and configuring devices. Refer to the following for this information:

- Software documentation that you received with your system
- Solaris™ operating system documentation, which is at

<http://docs.sun.com>

Shell Prompts

Shell	Prompt
C shell	<i>machine-name%</i>
C shell superuser	<i>machine-name#</i>
Bourne shell and Korn shell	\$
Bourne shell and Korn shell superuser	#

Typographic Conventions

Typeface*	Meaning	Examples
AaBbCc123	The names of commands, files, and directories; on-screen computer output	Edit your <code>.login.login</code> file. Use <code>ls -a</code> to list all files. % You have mail.
AaBbCc123	What you type, when contrasted with on-screen computer output	% su Password:
<i>AaBbCc123</i>	Book titles, new words or terms, words to be emphasized. Replace command-line variables with real names or values.	Read Chapter 6 in the <i>User's Guide</i> . These are called <i>class</i> options. You <i>must</i> be superuser to do this. To delete a file, type <code>rm filename</code> .

* The settings on your browser might differ from these settings.

Related Documentation

Title	Part Number
<i>Sun StorEdge 3510 FC Array and Sun StorEdge 3511 SATA Array Release Notes</i>	817-6597
<i>Sun StorEdge 3000 Family Installation, Operation, and Service Manual for the Sun StorEdge 3510 FC Array and Sun StorEdge 3511 SATA Array</i>	816-7300
<i>Sun StorEdge 3000 RAID Firmware 3.25 and 3.27 User's Guide</i>	817-3711
<i>Sun StorEdge 3000 Family Configuration Service 1.5 User's Guide</i>	817-3337
<i>Sun StorEdge 3000 Family Diagnostic Reporter 1.5 User's Guide</i>	817-3338
<i>Sun StorEdge 3000 Family Software Installation Manual (1.5/1.6)</i>	817-3764
<i>Sun StorEdge 3000 Family CLI 1.6 User's Guide</i>	817-4951
<i>Sun StorEdge 3000 Family Rack Installation Guide for 2U Arrays</i>	817-3629
<i>Sun StorEdge 3000 Family FRU Installation Guide</i>	816-7326
<i>Sun StorEdge 3000 Family Safety, Regulatory, and Compliance Manual</i>	816-7930

Accessing Sun Documentation

All *Sun StorEdge 3000 Family Best Practices Manual* array documentation is available online at the following locations:

http://www.sun.com/products-n-solutions/hardware/docs/Network_Storage_Solutions/Workgroup/3510

http://www.sun.com/products-n-solutions/hardware/docs/Network_Storage_Solutions/Workgroup/3511

<http://docs.sun.com/db/coll/3510FCarray>

<http://docs.sun.com/db/coll/3511FCarray>

Contacting Sun Technical Support

For late-breaking news and troubleshooting tips, review the *Release Notes* at the locations shown in “[Accessing Sun Documentation](#)” on page x.

If you have technical questions about this product that are not answered in the documentation, go to:

<http://www.sun.com/service/contacting>

To initiate or check on a USA-only service request, contact Sun support at:

800-USA-4SUN

To obtain international technical support, contact the sales office of each country at:

<http://www.sun.com/service/contacting/sales.html>

508 Accessibility Features

The Sun StorEdge documentation is available in Section 508-compliant HTML files that can be used with assistive technology programs for visually impaired personnel. These files are provided on the Documentation CD for your product as well as on the websites identified in the previous “Accessing Sun Documentation” section. Additionally, the software and firmware applications provide keyboard navigation and shortcuts that are documented in the user's guides.

Sun Welcomes Your Comments

Sun is interested in improving its documentation and welcomes your comments and suggestions. You can submit your comments by going to:

<http://www.sun.com/hwdocs/feedback>

Please include the title and part number of your document with your feedback:

Sun StorEdge 3000 Family Best Practices Manual, part number 816-7325-16.

Overview

This Best Practices Manual describes the use of Sun StorEdge™ 3510 FC arrays and Sun StorEdge 3511 Serial ATA (SATA) arrays, as well as the use of their corresponding expansion units. It complements the *Sun StorEdge 3000 Family Installation, Operation, and Service Manual* for these products.

This manual gives a high level overview of Sun StorEdge 3510 FC arrays and Sun StorEdge 3511 SATA arrays, and presents several 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. Choosing the solution that best matches your particular environment will provide the best results.

The Sun StorEdge 3510 FC array is a next-generation Fibre Channel storage system 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). This solution features powerful performance and reliability, availability and serviceability (RAS) features using modern FC technology. As a result, the Sun StorEdge 3510 FC array is ideal for performance-sensitive applications and for environments with many entry-level, mid-range, and enterprise servers, such as:

- Internet
- Messaging
- Database
- Technical
- Imaging

The Sun StorEdge 3511 SATA array is best suited for inexpensive secondary storage applications that are not mission-critical where higher capacity drives are needed, and where lower performance and less than 7/24 availability is an option. This includes near-line applications such as:

- Information lifecycle management
- Content addressable storage
- Backup and restore
- Secondary SAN storage

- Near-line DAS storage
- Static reference data storage

Note – Procedures described for configuring the Sun StorEdge 3510 and 3511 arrays are the same except where noted in this manual.

Fibre Channel and SATA Array Architectures

The Sun StorEdge 3510 FC array and Sun StorEdge 3511 SATA array share many architectural elements. This section discusses those elements, making note of the few differences in how the architecture is implemented differently in the two arrays.

Sun StorEdge 3510 FC and Sun StorEdge 3511 SATA 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.

In a single RAID controller configuration, 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. All six fibre channels in a Sun StorEdge 3510 FC array's I/O controller module support 1-Gbit or 2-Gbit data transfer speeds.

Sun StorEdge 3510 FC Array Features

On the 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.

Sun StorEdge 3510 FC 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 Features

On the Sun StorEdge 3511 SATA array, RAID controller channels 0 and 1 are dedicated host channels. Channels 4 and 5 are host channels by default but can be configured as drive channels. RAID controller channels 2 and 3 are dedicated drive channels that connect to expansion units.

Unlike the Sun StorEdge 3510 FC array, on the Sun StorEdge 3511 RAID controller 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 rates.

Sun StorEdge 3511 SATA RAID controller channels 4 and 5 provide two ports per loop (one port on each controller). Channels 4 and 5 support only a 2-Gbit data transfer rate.

Each Sun StorEdge 3511 SATA RAID controller has two ports designated as disk drive loops. The drive ports support only a 2-Gbit data transfer rate. These ports connect to the internal SATA disk drives using internal FC-SATA routing technology. These drive ports are also used to add expansion units 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.

How Sun StorEdge 3511 SATA Arrays Differ From Sun StorEdge 3510 FC Arrays

Sun StorEdge 3511 SATA arrays use Serial ATA (SATA) disk technology, and are recommended for cost-sensitive archival, static storage, and backup applications. Primary applications where the Sun StorEdge 3511 SATA array might be used rather than the Sun StorEdge 3510 FC array include:

- Applications requiring streaming or sequentially accessed data or those with cost-sensitive duty cycles such as backup and archival storage are best served by the Serial ATA-disk solution of the Sun 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 should be used only for archival, backup, or static storage; do not use it to store single instances of data. Do not use it as a boot device.

The Sun StorEdge 3511 SATA array can be used in multi-path and multi-host configurations where the array has a backup or archival role.

- Applications demanding transaction-based random access patterns with high duty cycles are best suited to a Fibre Channel disk-based solution such as the Sun StorEdge 3510 FC array. The Sun 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).

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 in the following table.

Note – Although the two products are very similar in appearance and setup, the configurations have very important differences. While the Sun StorEdge 3510 FC array can be used for all applications, its higher cost makes it a less desirable choice for certain applications. The Sun StorEdge 3511 SATA array's features are a subset of the Sun StorEdge 3510 FC array's functionality, and it should only be used as indicated below. Using the Sun StorEdge 3511 SATA array where a Fibre Channel disk-based solution is appropriate may result in the loss of data access or even, in extreme cases, in data corruption.

TABLE 1-1 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 are required. This includes online applications such as: <ul style="list-style-type: none"> • Database • Decision support • Data warehousing • Electronic commerce • Enterprise resource planning • Messaging, file and print 	Best suited for inexpensive secondary storage applications that are not mission-critical where higher capacity drives are needed, and where lower performance and less than 7/24 availability is an option. This includes near-line applications such as: <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
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
Direct host-attached JBOD Support	One JBOD per server	Not supported

Underlying Concepts and Practices

This chapter provides a brief overview of important concepts and practices that underly the configurations you can use. These concepts and practices are described in greater detail in other books in the Sun StorEdge 3000 Family documentation set. Refer to [“Related Documentation” on page ix](#) for a list of those books.

Fibre Channel Protocols

The Sun StorEdge 3510 FC array and Sun StorEdge 3511 SATA array support point-to-point and Fibre Channel–Arbitrated Loops (FC–AL) protocols. Using the point-to-point protocol with the Sun StorEdge 3510 FC arrays and Sun StorEdge 3511 SATA arrays require a switched fabric network (SAN), whereas selecting the FC-AL protocol enables the arrays to be used in either DAS or SAN environments. Using the point-to-point protocol enables full-duplex use of the available channel bandwidth, whereas using the FC-AL protocol 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 into 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:

$4 \text{ (host channels)} \times 1 \text{ (ID per channel)} \times 32 \text{ (LUNs per ID)} = 128$ 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) x 8 (IDs per channel) x 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

Several RAID levels are available: 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. Sun StorEdge 3000 family 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 logical drive can be configured for a different RAID level.

Sun StorEdge 3510 FC arrays and Sun StorEdge 3511 SATA arrays support a maximum of eight logical drives. A logical drive can be managed by either the primary or secondary controller. The best practice for creating logical drives is to add them evenly across the primary and secondary controllers. The most efficient maximum configuration would have four logical drives assigned to each controller. With at least one logical drive 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 2-1](#) and [TABLE 2-3](#) 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 2-2](#) and [TABLE 2-4](#)). 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 2-1 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 2-2 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 2-3 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 2-4 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 logical drive can be partitioned into up to 128 separate partitions or used as a single partition. Partitions are presented to hosts as LUNs.

Once the logical drives 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 logical drives 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 LD 0 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 FC 0, and connect a second cable from the lower board port FC 5 to a different host HBA or switch.

Cache Optimization

Sun StorEdge 3000 family arrays provide settings for both sequential I/O and random I/O. Sequential I/O is the default setting. 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 9.](#)

- “Maximum Usable Capacity (Gbyte) per Sun StorEdge 3510 FC Array Logical Drive” on page 9.
- “Maximum Number of Disks per Sun StorEdge 3511 SATA Array Logical Drive” on page 9.
- “Maximum Usable Capacity (Gbyte) per Sun StorEdge 3511 SATA Array Logical Drive” on page 9.

TABLE 2-5 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 “How Sun StorEdge 3511 SATA Arrays Differ From Sun StorEdge 3510 FC Arrays” on page 4.

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

Optimization mode limitations include the following:

- 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, making it necessary to determine the appropriate optimization mode prior to creating logical drives.
- Considering the characteristics of SATA disks and the application performance requirements for near-line storage, use Sun StorEdge 3511 SATA arrays and expansion units only with sequential I/O optimization.
- 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 has 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 the risk of data corruption. For this reason, a dual controller configuration is preferable.

Array Management Tools

Sun StorEdge 3000 family arrays use the same management interfaces and techniques. They 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, 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.
- Using the out-of-band Ethernet port connection or in-band FC connection, Sun StorEdge Configuration Service software or the Sun StorEdge CLI can configure and manage an array from a host system. Sun StorEdge Configuration Service provides a graphical user interface (GUI) that displays information about multiple aspects of the system at a glance. The main advantages of the CLI are that commands can be scripted and information can be passed to other programs.

Note – To set up and use the Sun StorEdge 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. 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.

SATA drives respond more slowly than FC drives when being managed by either Sun StorEdge Configuration Service or the CLI. From a performance standpoint, it is preferable to use these applications out-of-band to monitor and manage a Sun StorEdge 3511 SATA array. However, security considerations may take precedence over performance considerations. If you assign an IP address to an array in order to manage it out-of-band, for security reasons consider keeping the IP address on a private network rather than a publicly routable network.

Note – 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.

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.

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.

Planning Your Storage Architecture

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 straightforward, 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 3-1 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 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 3510 FC arrays and Sun StorEdge 3511 SATA arrays inherently support both DAS and SAN.

Direct-Attached Storage

One powerful feature of Sun StorEdge 3510 FC arrays and Sun StorEdge 3511 SATA 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 arrays automatically configures ports to match the transfer speed and communication method of each connection.

Note – Some older 1-Gbyte FC HBAs do not correctly support current auto-negotiation. In such configurations, set the transfer speed to 1-Gbyte rather than Auto. Refer to the release notes for your array to see the HBAs supported for your host and any limitations. Refer to the *Sun StorEdge 3000 Family RAID Firmware User's Guide* for information about how to set transfer speed.

- For Sun StorEdge 3511 SATA arrays, 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 a 2-Gbyte transfer rate.



Standard DAS Configuration



High Availability DAS Configuration

FIGURE 3-2 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:

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

Note – Do not use a Sun StorEdge 3511 SATA array to store single instances of data. It is more suitable for use in multipath 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. For information about obtaining and relocating SFP modules, refer to the *Sun StorEdge 3000 Family Installation, Operation, and Service Manual* for the Sun StorEdge 3510 FC Array and Sun StorEdge 3511 SATA Array.

Note – Except in some clustering configurations, in a DAS Loop configuration when you connect two hosts to channel 0 (both FC 0 ports of either controller), or channel 1 (both FC 1 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 3000 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 3510 FC arrays and Sun StorEdge 3511 SATA 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 3510 FC arrays and Sun StorEdge 3511 SATA 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 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 Sun StorEdge 3511 SATA array are characteristically higher bandwidth and lower IOPS. Sharing a StorEdge 3511 SATA array among several servers will depend upon the division of performance among the applications installed on each server actively accessing the Sun StorEdge 3511 SATA array.
- The Sun StorEdge 3510 FC Array is better suited for larger configurations with broader application requirements and higher IOPS.

A SAN can also support multiple 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 an array is deployed in a SAN, both point-to-point (full fabric) and arbitrated loop (public loop) modes are supported. Point-to-point mode enables superior 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 3510 FC arrays and Sun StorEdge 3511 SATA arrays are available in a number of configurations to address a broad range of storage capacities.

- The Sun StorEdge 3510 FC array is available with 15,000 RPM FC disks to address high-performance requirements.

- The Sun 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 Sun StorEdge 3511 SATA array provides more storage for each server attached. Larger capacity disks do not effectively support large groups of servers or users, so spreading the extra capacity among a larger number of servers results in poor performance and, in extreme cases, 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:

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

Capacities can be as large as:

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

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 array can provide.

Sun StorEdge 3510 FC arrays and Sun StorEdge 3511 SATA 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 Sun StorEdge 3510 FC array can support eight expansion units for a total of 108 146-Gbyte disks, providing a maximum of storage capacity of 15.668 Tbyte.
- A Sun StorEdge 3511 SATA array can support five expansion units for a total of 72 250-Gbyte disks, providing a maximum storage capacity of 16 Tbyte.

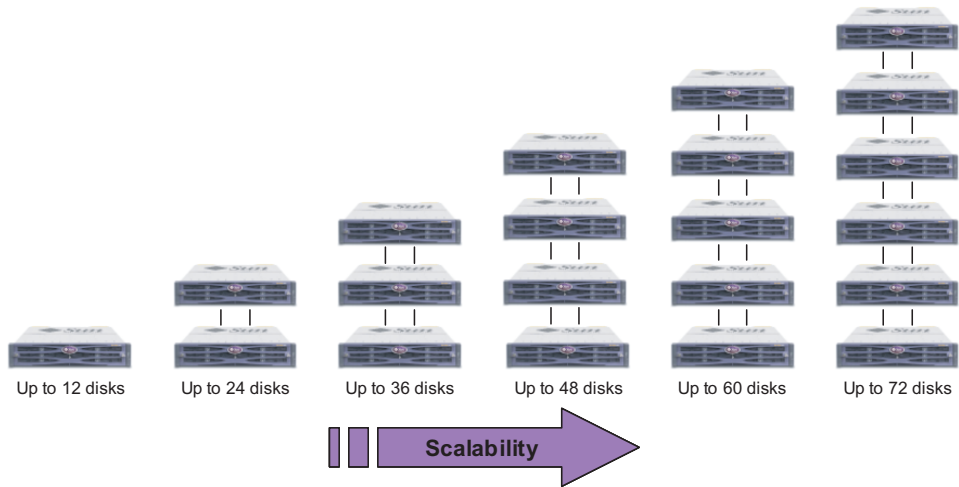


FIGURE 3-3 Increasing Capacity



Caution – For optimum and successful use of each array, check that you are using the array with the correct applications.

First Steps in Designing a Solution

There are two simple yet effective approaches for designing a Sun StorEdge 3510 FC arrays and Sun StorEdge 3511 SATA 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.

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 that the array provides.

- A Sun StorEdge 3510 FC array can support five or more servers. If there are four or fewer servers, a DAS solution is sufficient.
- A Sun StorEdge 3511 SATA array can support 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 disk capacity needed.

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 or SATA 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.

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 3510 FC arrays and Sun StorEdge 3511 SATA 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 3510 FC array or Sun StorEdge 3511 SATA array, the configuration should be saved using the firmware “Save nvram to disks” menu option or the Sun StorEdge Configuration Service Console’s “save configuration” utility.

DAS Configurations

This chapter presents several DAS configurations you can use.

Single-Controller 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. Similarly, using only a single controller creates a single point of failure, unless single controllers are used in pairs and mirrored. Using a dual-controller configuration, as shown in [“Dual-Controller Multipath DAS Configurations” on page 32](#), is preferable to using a single controller or a pair of single controllers.




FIGURE 4-1 Three Single-Controller DAS Configurations

Note – Do not use a Sun StorEdge 3511 SATA array to store single instances of data. It is more suitable for use in multipath and multi-host configurations where the array has a backup or archival role.

Single-Controller 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 six host connections (Sun StorEdge 3511 SATA array). These connections can be used in pairs, individually, or in any combination of both.
- You need to add SFP modules to support more than two host connections to the Sun StorEdge 3510 FC arrays or to the Sun StorEdge 3511 SATA array. For information about adding SFP modules, refer to the *Sun StorEdge 3000 Family Installation, Operation, and Service Manual* for the Sun StorEdge 3510 FC array and the Sun StorEdge 3511 SATA array.

- This single-controller configuration offers no redundancy. Consider mirroring single controllers or using dual controllers instead to achieve redundancy and increase reliability, availability, and serviceability. See [“General Configuration Considerations”](#) on page 23

Single-Controller DAS Setup Details

FIGURE 4-2 shows a single-controller Sun StorEdge 3510 FC array in a DAS configuration. FIGURE 4-3 shows a single-controller Sun StorEdge 3511 SATA array in a DAS configuration.




FIGURE 4-2 Sample Single-Controller Sun StorEdge 3510 FC DAS Connections




FIGURE 4-3 Sample Single-Controller Sun StorEdge 3511 SATA DAS Connections

TABLE 4-1 Sample Setup Summary for Single-Controller DAS Configurations

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.

For information about connecting multiple expansion units, see [“Sun StorEdge 3510 FC Array High-Capacity Configurations”](#) on page 47 or [“Sun StorEdge 3511 SATA Array High-Capacity Configurations”](#) on page 61.

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 FC 0.**
10. **Connect the second server (if needed) to controller port FC5.**
11. **Connect the third server (if needed) to controller port FC 1.**
12. **Connect the fourth server (if needed) to controller port FC4.**
13. **After completing the configuration, save the configuration using the “Save nvram to disks” menu option in the firmware application or the “save configuration” command in Sun StorEdge Configuration Service.**

Dual-Controller Multipath DAS Configurations




FIGURE 4-4 Sample Dual-Controller Multipath DAS Configurations

Note – Do not use a Sun StorEdge 3511 SATA array to store single instances of data. It is more suitable for use in multipath and multi-host configurations where the array has a backup or archival role.

Dual-Controller Multipath 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 need to add SFP modules to support more than four host connections to an array. For information about adding SFP modules, refer to the *Sun StorEdge 3000 Family Installation, Operation, and Service Manual* for the Sun StorEdge 3510 FC array and the Sun StorEdge 3511 SATA array.
- Using two single-port 2-Gbit FC host adapters in a high-availability configuration makes optimum use of a 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 host-based 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.

Sample Dual-Controller Multipath DAS Setup Details

FIGURE 4-5 shows a Sun StorEdge 3510 FC array in a dual-controller multipath DAS configuration. **FIGURE 4-6** shows a Sun StorEdge 3511 SATA array in a dual-controller multipath DAS configuration.




FIGURE 4-5 Sample Dual-Controller Multipath Sun StorEdge 3510 DAS Connections

Note – Do not use a Sun StorEdge 3511 SATA array to store single instances of data. It is more suitable for use in multipath and multi-host configurations where the array has a backup or archival role.




FIGURE 4-6 Sample Dual-Controller Multipath Sun StorEdge 3511 DAS Connections

TABLE 4-2 Sample Setup Summary for Dual-Controller Multipath Sun StorEdge 3511 DAS Configurations

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.

For information about connecting multiple expansion units, see [“Sun StorEdge 3510 FC Array High-Capacity Configurations” on page 47](#) or [“Sun StorEdge 3511 SATA Array High-Capacity Configurations” on page 61](#).

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 4 of the primary controller.

8. Map Logical Drive 1 (if created) to channels 1 and 5 of the secondary controller.

9. Map Logical Drive 2 to channels 0 and 4 of the primary controller.

10. Map Logical Drive 3 (if created) to channels 1 and 5 of the secondary controller.

11. Connect the first server (Server 1) to port FC 0 of the upper controller and port FC5 of the lower controller.

12. Connect the second server (Server 2), if needed to port FC 1 of the lower controller and port FC4 of the upper controller.

13. Connect the third server (Server 3), if needed to port FC 0 of the lower controller and port FC5 of the upper controller.

14. Connect the fourth server (Server 4), if needed to port FC 1 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 the “Save nvram to disks” menu option in the firmware application or the “save configuration” command in Sun StorEdge Configuration Service.

SAN Configurations

This chapter presents several SAN configurations you can use.

Full-Duplex SAN Configurations




FIGURE 5-1 Typical Full-Fabric SAN Configuration

Note – Do not use a Sun StorEdge 3511 SATA array to store single instances of data. It is more suitable for use in multipath and multi-host configurations where the array has a backup or archival role.

Full-Duplex SAN Tips and Techniques

- In a full-duplex 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 FC 0) is connected to a fabric switch, for example, the second port for channel 0 on that controller, and the FC 0 ports on a redundant controller, cannot be used. Similarly, if channel 1 (port FC 1) is connected to a fabric switch, the second FC 1 port on that controller and the FC 1 ports on a redundant controller, cannot be used.
- For Sun StorEdge 3511 SATA arrays, ensure that the optimization mode is set to Sequential.

Sample Full-Duplex SAN Setup Details

[FIGURE 5-2](#) shows a Sun StorEdge 3510 FC array in a full-duplex SAN configuration. [FIGURE 5-3](#) shows a Sun StorEdge 3511 SATA array in a full-duplex SAN configuration.




FIGURE 5-2 Sample Sun StorEdge 3510 FC Full-Duplex SAN Connections




FIGURE 5-3 Sample Sun StorEdge 3511 SATA Full-Duplex SAN Connections

TABLE 5-1 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.

For information about connecting multiple expansion units, see [“Sun StorEdge 3510 FC Array High-Capacity Configurations”](#) on page 47 or [“Sun StorEdge 3511 SATA Array High-Capacity Configurations”](#) on page 61.

3. Configure cache optimization. If you are configuring a Sun StorEdge 3511 SATA array, confirm that the cache optimization policy 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 FC 0 of the upper controller and port FC 1 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 the "Save nvram to disks" menu option in the firmware application or the "save configuration" command in Sun StorEdge Configuration Service.

Redundant Point-to-Point SAN Configurations




FIGURE 5-4 Sample Redundant Point-to-Point SAN Configuration

Note – Do not use a Sun StorEdge 3511 SATA array to store single instances of data. It is more suitable for use in multipath and multi-host configurations where the array has a backup or archival role.

Redundant Point-to-Point Tips and Techniques

- In the redundant point-to-point 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 FC 0) is connected to a fabric switch, for example, the second port for channel 0 on that controller, and the FC 0 ports on a redundant controller, cannot be used. Similarly, if channel 1 (port FC 1) is connected to a fabric switch, the second port for channel 1 on that controller and both FC 1 ports on a redundant controller cannot be used.

Redundant Point-to-Point Setup Details

FIGURE 5-5 shows a Sun StorEdge 3510 FC array in a redundant point-to-point SAN configuration. FIGURE 5-6 shows a Sun StorEdge 3511 SATA array in a redundant point-to-point SAN configuration.




FIGURE 5-5 Sample Redundant Point-to-Point Sun StorEdge 3510 SAN Connections




FIGURE 5-6 Sample Redundant Point-to-Point Sun StorEdge 3511 SAN Connections

TABLE 5-2 Setup Summary for a Redundant Point-to-Point Configuration

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.**

For information about connecting multiple expansion units, see [“Sun StorEdge 3510 FC Array High-Capacity Configurations” on page 47](#) or [“Sun StorEdge 3511 SATA Array High-Capacity Configurations” on page 61](#).

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 mode.**

5. **Configure target IDs.**

6. **Create at least two logical drives and configure spare disks.**

7. **Assign even-numbered logical drives beginning with Logical Drive 0 to the primary controller.**

8. **Assign odd-numbered logical drives beginning with Logical Drive 1 to the secondary controller.**

9. **Create one or more logical drive partitions for each server.**

10. **Map the LUNs from even-numbered logical drives beginning with Logical Drive 0 to channels 0 and 4 of the primary controller.**

11. **Map the LUNs from even-numbered logical drives beginning with Logical Drive 1 to channels 1 and 5 of the secondary controller.**

Note – See [“Fibre Channel Protocols” on page 7](#) for a description of the maximum number of devices you can map in various configurations.

12. **Connect the first switch to port FC 0 of the upper controller and port FC 1 of the lower controller.**

13. **Connect the second switch to port FC4 of the lower controller and port FC5 of the upper controller.**

14. **Connect each server to each switch.**

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.**

Using Multiple Expansion Units in High Capacity Configurations

High-capacity configurations using multiple expansion units are appropriate for some situations, although using multiple arrays connected to the same SAN generally offers significantly better performance than a single high-capacity configuration. High-capacity configurations require considerable planning in order to connect expansion units in a way that ensures maximum reliability, availability, and serviceability (RAS), and avoids single points of failure.

This chapter presents a few limitations on high-capacity configurations you should consider, and shows you sample configurations for one-to-eight Sun StorEdge 3510 FC expansion units. The following chapter shows you sample configurations for one-to-five Sun StorEdge 3511 SATA expansion units when connected to a Sun StorEdge 3511 SATA array.

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 61](#) for information about adding expansion units to a Sun StorEdge 3511 SATA array.

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-Duplex SAN Configurations” on page 37.](#)”

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 [“Redundant Point-to-Point SAN Configurations” on page 42.](#)”

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.
- 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.

Connecting One Expansion Unit




FIGURE 6-1 Sun StorEdge 3510 FC Array Configuration with One Expansion Unit

Connecting Two Expansion Units




FIGURE 6-2 Sun StorEdge 3510 FC Array Configuration with Two Expansion Units

Connecting Three Expansion Units




FIGURE 6-3 Sun StorEdge 3510 FC Array Configuration with Three Expansion Units

Connecting Four Expansion Units




FIGURE 6-4 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 6-5 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 6-6 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 6-7 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.

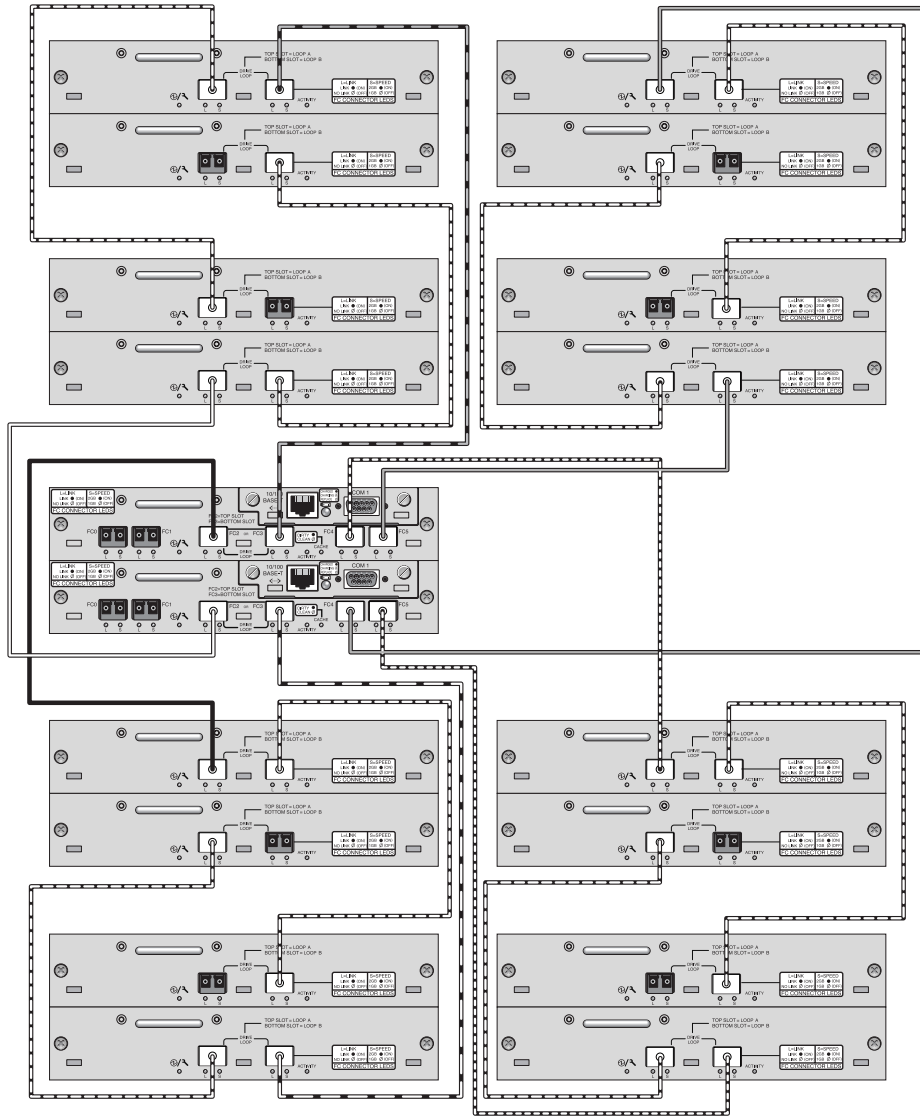


FIGURE 6-8 Sun StorEdge 3510 FC Array Configuration with Eight Expansion Units

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 47](#) for information about adding more than two expansion units to a Sun StorEdge 3510 FC array.

The use of a full-fabric SAN configuration is highly recommended for host connections, as shown in [“Full-Duplex SAN Configurations” on page 37](#).

Refer to the *Sun StorEdge 3000 Family FRU Installation Guide* for information about supported disk drives, cables, SFPs, and other user-replaceable items.

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 [“Redundant Point-to-Point SAN Configurations” on page 42](#) and must be used with the network applications appropriate to the array, as summarized in the [“How Sun StorEdge 3511 SATA Arrays Differ From Sun StorEdge 3510 FC Arrays” on page 4](#).

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 SATA array Release Notes* on the Sun product web sites.

Connecting One Expansion Unit




FIGURE 7-1 Sun StorEdge 3511 SATA Array Configuration with One Expansion Unit

Connecting Two Expansion Units




FIGURE 7-2 Sun StorEdge 3511 SATA Array Configuration with Two Expansion Units

Connecting Three Expansion Units




FIGURE 7-3 Sun StorEdge 3511 SATA Array Configuration with Three Expansion Units

Connecting Four Expansion Units




FIGURE 7-4 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 7-5 Sun StorEdge 3511 SATA Array Configuration with Five Expansion Units

