

Sun StorEdge™ T3 Array with 1-GB Cache Controller

Just the Facts



Copyrights

©2001 Sun Microsystems, Inc. All Rights Reserved.

Sun, Sun Microsystems, the Sun logo, Sun StorEdge, Solaris, SunSpectrum, SunSpectrum Gold, ArrayStart, Sun StorEdge ArrayStart, Jiro, Java, Solstice, Solstice DiskSuite, Solstice Backup, StorTools, Sun Enterprise, Ultra, Sun Fire, Sun Blade, SunSolve, SunPS, SunSpectrum Platinum, SunSpectrum Silver, SunSpectrum Bronze, and Gigaplane are trademarks or registered trademarks of Sun Microsystems, Inc. in the United States and other countries.

UNIX is a registered trademark in the United States and other countries, exclusively licensed through X/Open Company, Ltd.

Last update: 9/5/2001



Table of Contents

Positioning.....	5
Introduction.....	5
Key Features.....	6
Features and Benefits.....	7
Product Family Placement.....	9
Key Messages.....	9
Storage Management Software.....	11
Target Markets.....	13
Target Applications.....	13
Selling Highlights.....	14
Market Value Proposition.....	14
Compatibility.....	14
Enabling Technology.....	16
Features Overview.....	16
Fibre Channel Loop-Switching Architecture	16
Three-Dimensional Scalability: Capacity, Performance, and Availability.....	17
Enhanced RAID 5 Performance.....	17
ECC Protection in the 1-GB Cache Controller.....	17
System Architecture.....	18
Basic Architecture.....	18
External Interfaces.....	19
Administration Path.....	20
3-D Scalability.....	20
Availability, Reliability, and Serviceability via FRUs.....	20
Path Failure.....	24
Implicit LUN Failover vs. Explicit LUN Failover.....	25
Controller Failure.....	25
LUN Configurations per Tray.....	28
Reliability, Availability, and Serviceability (RAS).....	30
Reliability.....	30
Availability.....	30
Serviceability.....	30
Specifications.....	32
Performance.....	34
Benchmarks.....	34
Key Findings.....	34
Cache Memory and Mirroring.....	34
Seek Range.....	34
Predictable Performance.....	34
Block Size.....	35
Adaptive Cache Optimizations.....	35
Tuning.....	36
RAID Level.....	36
Stripe Unit Size.....	37
Reconstruction Rate (RAID 5).....	37
Capacity Planning.....	37
Requirements and Configuration.....	38
System Requirements.....	38



Solaris Operating Environment-Based Software Compatibility.....	39
System Configuration.....	39
Front Components.....	40
Interconnect.....	42
Configuration Options.....	42
Configuration Diagrams.....	43
System Administration.....	45
System Maintenance.....	45
Localization and Internationalization.....	45
Software Administration.....	46
Compatible Software.....	46
Host Operating Environments.....	46
StorTools Diagnostic Software.....	47
VERITAS Volume Manager Software.....	47
Sun StorEdge Component Manager Software.....	47
Ordering Information.....	51
Sun StorEdge T3 Array Ordering.....	51
Part Numbering Scheme.....	51
Configuration Matrix.....	52
Shipping Configurations.....	53
Sun StorEdge T3 Array for the Workgroup, 1-Array Configurations	54
Sun StorEdge T3 Array for the Enterprise, 2-Array Configurations	55
Sun StorEdge T3 Array for the Enterprise, 2-Array Configurations	56
Sun StorEdge T3 Array for the Enterprise, 8-Array Configurations	58
Options.....	59
Services Available.....	61
Warranty.....	61
Service and Support.....	62
Warranty Duration and Terms.....	63
Education.....	63
Glossary.....	64
Materials Abstract.....	70
FAQs.....	73



Positioning

Introduction



Figure 1. The Sun StorEdge™ T3 array for the enterprise

Simple. Scalable. Flexible. To meet the growing demands for storage caused by the "Net Effect" — the explosive growth of users, amount of data, and type of data — Sun developed the Sun StorEdge™ T3 array. The Sun StorEdge T3 array represents an advanced generation of modular, reliable storage systems that is designed to allow businesses to scale storage capacity, performance, and availability seamlessly as business needs grow.

The advanced system architecture of the Sun StorEdge T3 array meets the demands of a wide range of applications but is particularly well suited for database, NFS, web serving, and e-mail applications, as well as a variety of workloads. Its modular design combines advanced hardware RAID controller technology, hot-swappable systems components, and industry-standard Fibre Channel technology. These features provide high-performance, redundant data paths between disks and servers, and create a scalable storage system that can address stringent availability requirements.

Customers are experiencing dramatic growth in their storage requirements, and are changing the way they view, evaluate, and purchase their storage. Today, storage purchases are based upon application and business needs, which are independent of the type of host system platform. Increasingly, storage systems are re-purposed to different uses or used as host platforms within a department or enterprise. Many customers simply prefer to purchase their storage from a single-source vendor.

To be suitable for an IT environment with different open systems host platforms, a storage system must be compatible with the various host platform types found in a typical enterprise, and the high-quality Sun StorEdge T3 array storage system is no exception. Supporting multiple host platforms while providing versatile three-way scalability across capacity expansion, increases in performance, and enterprise-class availability features, the Sun StorEdge T3 array solves the need for a simple, scalable, flexible storage system which can accommodate a variety of open systems platforms. Available in configurations such as the completely redundant, fault-tolerant failover storage system, complete with redundant failover switches, this partner-pair approach provides enterprise-class availability and performance at an exceptional price.



The Sun StorEdge T3 arrays are offered in two configurations, each providing different levels of availability, software, and services. The two Sun StorEdge T3 storage array systems are:

- The Sun StorEdge T3 array for the enterprise (partner pair models) offers redundant RAID controllers for high-availability enterprise and data center environments. These models are fully redundant, with hot-swappable hardware components and with capacities ranging from 660 GB (36.4-GB HDDs) up to 5.2 TB (73.4-GB HDDs) per cabinet. This pair of arrays is available in tabletop, rack-ready, or rack-installed configurations with two or eight arrays, with or without dual hubs or dual redundant switches.
- The Sun StorEdge T3 array for the workgroup with a single RAID controller offers a reliable and cost-effective storage system with 327-GB (36.4-GB HDDs), or 660-GB (73.4-GB HDDs) individual drive tray capacities. These models are available in tabletop or rack-ready versions.

The Sun StorEdge T3 array includes array management software that allows large numbers of arrays to be centrally administered. Management functions include the Sun StorEdge Component Manager software, which allows administrators to configure, monitor, control, and diagnose the Sun StorEdge T3 storage system from a single graphical interface. This capability allows an organization to grow its storage capacity without increasing administration costs and complexity.

In addition to providing advanced hardware and software components, the Sun StorEdge T3 array for the enterprise models include an enhanced services package, which allows customers to rapidly implement complex, reliable, and monitored storage environments. This extended service package helps ensure the use of sound storage installation and configuration practices, thereby allowing Sun to put the support infrastructure in place that is required to maintain the most demanding enterprise and data center environments.

The Sun StorEdge T3 array is available in configurations using 36.4-GB or 73.4-GB, bi-directional, dual-ported 10000-rpm FC-AL disk drives. Each drive tray includes a hardware RAID controller with battery-backed cache, dual hot-swap/redundant power/cooling units with built-in redundant UPS batteries, four hot-swap/redundant electrically independent cooling fans, hot-swap/redundant interconnect cards, and nine drives. The Sun StorEdge T3 array is massively scalable to 168 TB (256 drive trays across 32 rack cabinets using 73.4-GB HDDs) on a single server.

Key Features

- Fibre Channel architecture from front to back
- Hardware RAID storage for RAID 0, RAID 1 (1+0), and RAID 5 levels, or a combination of two of these
- High-availability, hot-swap/redundant failover dual-controller design (partner pair models)
- Available in configurations using 36.4-GB or 73.4-GB, bi-directional, dual-ported 10000-rpm FC-AL disk drives
- Each drive tray includes a hardware RAID controller with battery-backed cache, dual hot-swap/redundant power/cooling units with built-in redundant UPS batteries, four hot-swap/redundant electrically independent cooling fans, hot-swap/redundant interconnect cards, and nine drives for predictable linear performance scalability, and consistent reliability/availability
- The Sun StorEdge T3 array is massively scalable to 168 TB (256 drive trays across 32 rack cabinets using 73.4-GB HDDs) on a single server
- Easy-to-use Sun StorEdge Component Manager 2.2 software for array configuration and central administration
- Installation, configuration, and support service are optional



- VERITAS Volume Manager with dynamic multipathing (DMP) 3.0.4, 3.1, 3.1.1, or 3.2 included (Sun StorEdge T3 array for the enterprise models — Tier 2 license)
- Compatible with Sun StorEdge Instant Image and Sun StorEdge Network Data Replicator (SNDR) storage software programs for enterprise-critical applications
- Rackmounting options for Sun StorEdge expansion cabinets, Sun Enterprise™ server cabinets, third-party cabinets, and standard 2- or 4-post Telco racks
- Sun StorEdge arrays for the enterprise models are available in configurations with dual 7-port hubs or dual redundant failover 8-port switches

Features and Benefits

The Sun StorEdge T3 storage arrays have attractive features that make these units the products of choice for forward-moving customers.

Features

- 73.4-GB bi-directional dual-ported FC-AL HDDs
- Individual servers can support up to 168 TB of attached Sun StorEdge T3 array storage
- Three-way scalability — linear performance, massive capacity, and consistent reliability/ availability
- Full Fibre Channel architecture — Sun StorEdge T3 array architecture takes advantage of 100 MB/sec. FC-AL host interface and dual 100 MB/sec. back-end FC-AL drive loops
- Sun's Solaris™ Operating Environment, Microsoft Windows NT, Linux kernel, HP-UX, and IBM AIX host platform support
- Highly scalable tabletop, rack-ready, or factory rackmounted configurations

Benefits

- High-speed, high-capacity drives allows the Sun StorEdge T3 array to scale to 168 TB on a single server and reduces the price per GB
- Massive scalability for demanding storage requirements
- Helps lowers cost per GB; performance increases steadily as capacity is added; reliability/availability remains constant as capacity is added
- High performance that does not bog down for demanding applications; intelligent FC-AL diagnostic circuitry for fast diagnosis and resolution
- Save money and time by standardizing on one centrally managed high-performance, high-availability storage platform
- Flexible installation allows capacities to scale from 327 GB for to over 5.2 TB for single-cabinet configuration, allowing the modular Sun StorEdge T3 array to cost-effectively grow along with each business



Features

- Fully redundant, hot-swappable hardware components (all major FRUs)
- Online reconfiguration
- Dual redundant load-sharing/load-balancing power supplies
- Four electrically independent cooling fans
- Mirrored cache (partner pair models)
- Battery-backed cache with built-in hot-swap/redundant UPS batteries to allow immediate destage to disk upon sense of power loss
- Integrated battery backup for cache protection
- No internal cabling
- Dual power cords on each array
- Data cache and processor buffers are ECC protected
- Field-replaceable units (FRUs) are easy to identify, access, and hot-swap
- Hardware RAID controller with dedicated XOR parity engine
- Sun StorEdge T3 arrays are Sun Remote Service (SRS) Event Monitoring and Management Service 2.0 ready
- 2-year (24 x 7 phone support, same business day on-site repair response) (Sun StorEdge T3 array for the enterprise models)
- Installation and configuration through the Sun StorEdge ArrayStartSM service is optional
- Fully redundant failover configurations available
- VERITAS dynamic multipathing (partner pair models)

Benefits

- Continuous data availability and reliability; easy, quick, and non-disruptive serviceability
- Provides transparent and automatic error detection and correction.
- Easy serviceability, decreased downtime, and reduction in potential errors
- Provides data protection, high-speed low-latency RAID 5 performance, increased availability, and ease of configuration
- In Sun Enterprise Services, service orders are generated when systems problems are first detected, decreasing time to resolve problems or potential problems
- 24 x 7 remote system monitoring with 4-hour, on-site response time
- Experienced systems engineers help ensure proper installation and configuration
- Excellent value — available redundant, failover factory configurations for the enterprise include switches at very aggressive pricing for quick, easy, error-free installation
- Increased availability through automatic host loop failover



Product Family Placement

The Sun StorEdge T3 array is an open, modular network system for simple, reliable, and scalable storage. This array is an economical, yet powerful and scalable, storage system which delivers advanced data services to the storage network. The Sun StorEdge T3 array helps maximize performance in online transaction, decision support, messaging, web hosting, and high-performance computing environments. Its design is based on a simple building block concept, combining advanced RAID technology with industry-standard, Fibre Channel architecture in a modular package. Concatenating these modular building blocks creates a high-performance, highly available, centrally administered network storage system. The scalability of the Sun StorEdge T3 array allows it to grow along with businesses.

The Sun StorEdge T3 array also offers competitive and flexible performance and availability features at a low cost per GB of storage. The Sun StorEdge T3 array is ideal for companies facing rapidly growing storage capacity on disparate host platforms.

Key Messages

The Sun StorEdge T3 array delivers these key benefits:

- Three-way scalability — As capacity is added, performance is linearly increased and availability/reliability remain consistent
- Simplicity — Easy to use, manage, monitor, and control from a single central administration console
- World-class Sun service and support
- Shipped with Jiro™ technology-compliant storage management software

Three-Way Scalability

The Sun StorEdge T3 array design provides *three-way scalability* for flexible capacity, performance, and availability. Customers can scale up in the following ways:

- **Increase capacity** — Each drive tray has nine disk drives; simply add more drive trays to increase storage capacity using 36.4-GB or 73.4-GB disks.
- **Increase performance** — Each hardware RAID controller with pipelined XOR and dual backend loops is responsible for handling exactly nine hard drives so bandwidth and I/O are additive; adding more controller units increases both overall system MB/sec. and IOPS performance in a predictable, linear fashion. The Sun StorEdge T3 array provides performance strengths exceeding many disk systems currently on the market with respect to transaction-intensive online services, data-intensive file services, high-performance computing applications, and data warehousing/data mart data analysis applications. Here are some performance highlights:
 - 96 MB/sec. reads from disk per controller unit with RAID 5 (8 + 1 volumes)
 - 95 MB/sec. writes to disk per controller unit with RAID 5 (8 + 1 volumes)
 - 7,250 IOPS (100-percent reads from cache, 8-KB block size) *per controller unit* — **Sun StorEdge T3 array for the workgroup, 74% increase from the earlier model with 256-MB cache.**
 - 14,500 IOPS (100-percent reads from cache, 8-KB block size) *per controller unit* — **Sun StorEdge T3 array for the enterprise, 78% increase from the model with 256-MB cache.**
 - 264 MB/sec. sustained internal throughput *per controller unit*
- **Increase availability** — Each drive tray consists of two hot-swap/redundant load-sharing, load-balancing power supplies and four hot-swap/redundant electrically independent fans (if a power



supply fails, all four fans continue to spin) to power and cool exactly nine hard drives. Every time another drive tray with nine hard disks is added, another two power supplies and four cooling fans are added so the system does not exceed its power or temperature boundaries. Moreover, each time a drive tray is added, another data path is added for nine hard disks, so each additional tray means additional data paths for increased overall system resiliency. Cache is battery-backed by two built-in hot-swap redundant UPS batteries which are dedicated to exactly nine hard drives, so if power loss is detected, the data staged in the write cache is immediately written (destaged) to the hard disks and the unit is then gracefully shut down, with no data loss. Drive tray availability can be scaled in two ways:

- In partnered pairs with redundant automatic failover hot-swap RAID controllers with fully duplexed cache (Sun StorEdge T3 array for the enterprise) for host failover environments
- Drive trays can be added individually (Sun StorEdge T3 array for the workgroup models) for host mirroring environments

Rackmounted configurations of up to eight controller units per cabinet are possible. A second unit provides full redundancy and provides increased performance to 14,500 IOPS (100-percent reads, 8-KB blocks) with an additional effective bandwidth of 90 MB/sec. Two and eight partner pairs in rack configurations are currently available, with or without dual hubs or switches.

Simplicity, Ease of Use

The open, modular architecture of the Sun StorEdge T3 array makes it easy to install, configure, service, and administer:

- All major components are easily accessible either at the front or rear of the unit.
- All components except the system chassis (which includes the cableless passive midplane board) can be hot-swapped. Hot-swappable, redundant-capable field-replaceable units (FRUs) include nine preconfigured RAID disk drives, two power/cooling units, two unit interconnect cards, and the RAID controller card.
- Interconnectivity between units is easily accomplished using the interconnect cables on each unit interconnect card; no recabling within the units is required.
- The RAID controller card has built-in network services, Telnet, FTP, and HTTP through a 10/100BASE-T Ethernet network connection for an out-of-band administrative interface which can be centrally monitored and controlled from a single administrative console.
- The administrative console is host platform independent, so all Sun StorEdge T3 arrays connected on the same Ethernet LAN can be centrally administered from one console, independent of the host platform operating system.
- Sun StorEdge Component Manager software provides a friendly graphical user interface for managing the Sun StorEdge T3 array controller unit physical enclosure.
- Software support for VERITAS Volume Manager software allows enhanced management of disk volumes (VERITAS Volume Manager software is included with all Sun StorEdge T3 array for the enterprise configurations).

World-Class Sun-Backed Service and Support

Each Sun StorEdge T3 array is backed by Sun's Global Warranty Program. Sun StorEdge T3 array for the enterprise models include a 2-year enhanced warranty with two years of 24 x 7 phone support and same business day hardware repair support. Sun StorEdge T3 array for the workgroup models include a 2-year standard warranty.



Each Sun StorEdge T3 array is backed by Sun's world-class service and support. Sun StorEdge T3 arrays are serviced worldwide by trained and certified Sun personnel so customers received high-quality service and highly reliable support.

Serviceability features include SNMP monitoring; failure detection, reporting, and recovery for all FRUs; hot-swap capability for all FRUs; and support for online reconfiguration through an independent Ethernet interface. The Sun StorEdge T3 array can be remotely monitored using the Sun Remote Services (SRS) 2.0 Event Monitoring and Management service.

Storage Management Software

Sun StorEdge Component Manager software is included with the Sun StorEdge T3 array. Sun StorEdge Component Manager software is a server-installed, Java™ technology-based application for managing attached Sun StorEdge arrays and their hardware components. Sun StorEdge Component Manager software is compliant with Jiro technology.

This software allows the user to perform control directives on some of the components (for example powering off a disk). In addition, Sun StorEdge Component Manager software affords constant monitoring of the Sun StorEdge T3 arrays or Sun StorEdge A5X00 array enclosures and is capable of providing alarm notification and remote reporting (via e-mail, files, and system logging) upon detection of abnormal activities or conditions within a designated storage enclosure.

Note: *Sun StorEdge Component Manager software does not manage the logical organization of stored data.*

The Sun StorEdge Component Manager software plug-in is accessible from the Sun StorEdge Management Console software. Sun StorEdge Management Console software serves as a storage management framework for Java technology-based management plug-ins, making it easier to operate, administer, and maintain storage area networks. The management GUI gives system administrators and service personnel a centralized point for storage administration. Sun StorEdge Management Console software enhances the reliability, availability, and serviceability (RAS) of storage assets. The Sun StorEdge Management Console software implements elements of the specification for the Jiro open storage management platform to enable automated management services and interoperability across heterogeneous storage area networks.

Administrators and service personnel of supported arrays can use Sun StorEdge Component Manager software's GUI to:

- Monitor enclosures by using hardware polling, alarm notification, event logging, and remote reporting for abnormal conditions and activities
- Display the status of enclosures, enclosure components, and their associated properties
- Exercise control directives on some enclosure components, such as powering off a disk, or enabling or disabling a RAID controller

Features of Sun StorEdge Component Manager 2.2 Software

The 2.2 release of Sun StorEdge Component Manager software includes support for the Sun StorEdge T3 array, along with enhanced support for the Sun StorEdge A5X00 arrays. In addition, a diagnostics interface has been added for the Sun StorEdge T3 array.



Administrators and service personnel of Sun StorEdge T3 arrays may use Sun StorEdge Component Manager 2.2 software on the Solaris Operating Environment and on the Microsoft Windows NT operating environment to:

- Monitor enclosures by using hardware polling, alarm notification, event logging, and remote reporting for abnormal conditions and activities
- Display the status of enclosures, enclosure components, and their associated properties through photorealistic views
- Exercise control directives on some enclosure components, for example, disabling a RAID controller
- Configure controller units for one or two LUNs
- Obtain block I/O statistics for array controllers

Support for the Sun StorEdge A5X00 arrays on the Solaris Operating Environment has also been enhanced to include these features:

- Faster enclosure discovery
- The system monitors for bad Fibre Channel loops; if one is discovered and there is another path to the disks, the system has the capability to disable the Fibre Channel port on the loop
- Global setting for polling interval

This release also forwards alarm notifications to Sun Management Center software via the Sun StorEdge Management Console 2.2 software.

Other features include:

- Threading discovery and dynamic report/update
- The system can discover enclosures asynchronously and also during the start of the Sun StorEdge Component Manager software
- All polling intervals are now in the Sun StorEdge Component Manager software's configuration panel instead of in each enclosure configuration panel for ease of use

Target Users

The primary customers for the Sun StorEdge T3 array are Sun network storage direct accounts, storage-only resellers, solution resellers, OEMs, distributors, and system remarketers. Sun StorEdge T3 arrays meet user needs, as shown in the table below.

Individual User	Buying Influence Needs
MIS manager	FC-AL technology
Procurement	Investment protection in existing and future products
Developer	Standards compliance for implementation of FC-AL products
Systems administrator	Flexible management in both software and hot-swap components
Operations	High availability, allowing efficient system operation
End user	High performance, resulting in quick I/O response



Target Markets

The Sun StorEdge T3 array is well suited for the capacity and performance requirements of application servers, network data services, and performance-oriented systems. Strategic uses for this array within the manufacturing, government, finance, and entertainment industries include those shown in the table below.

Industry/Customer	Key Features to Highlight
Disaster recovery	Future support for remote mirroring to 10 kilometers (using FC switches)
Data center storage	Enterprise-class redundancy and mission-critical availability features
Technical computing	High-performance data storage for engineering design projects
Scientific computing	High bandwidth for data capture, retrieval, and storage
Workgroup storage	Scalable design with proper capacity to support workgroup sizes

The Sun StorEdge T3 array is ideal for environments in these price ranges (in US dollars):

- \$100,000 and up for enterprises (Sun StorEdge T3 array for the enterprise)
- \$25,000 to \$100,000 for workgroups and departments (Sun StorEdge T3 array for the workgroup)

Targeted groups within these market segments are:

- Global 1000 customers moving to storage area networks (SANs)
- NFS system customers
- Data modeling and data-intensive customers
- Technical and scientific customers
- High-performance computing (HPC) customers
- Internet service providers (ISPs)
- E-commerce customers (interactive online, e-commerce organizations, digital media, financial services organizations, and Internet-centric organizations)

Target Applications

The Sun StorEdge T3 array is well suited for customers who desire scalable hardware RAID Fibre Channel storage. Target applications include the following:

- Service providers (e-mail, v-mail, static/dynamic web servers, e-commerce)
- Enterprise and data center (OLTP, data warehouse, e-commerce)
- Technical and scientific applications (high-performance computing)
- Workgroup (NFS, e-mail, file/print services)
- Computer generated animation (CGA)
- Image capture and retrieval applications such as medical imaging and high-performance data acquisition
- Video streaming applications



Selling Highlights

Market Value Proposition

The Sun StorEdge™ T3 array is an open, modular network storage solution for simple, reliable, and scalable data storage.

Organizations of all sizes, in every industry, want a simple, reliable, and manageable approach to storage. Customers seek solutions based on open standards that are capable of solving their needs for access to information across their networks. The Sun StorEdge T3 array is an economical, yet powerful and scalable, storage system that delivers advanced data services to the storage network. This system helps maximize performance in NFS, web serving, e-mail, and a variety of other applications. The array's design is based on a simple building block concept, combining advanced RAID technology with industry-standard Fibre Channel technology in a modular package. Combining building blocks creates a high-performance, highly available network storage system that can be centrally administered. The scalability of the Sun StorEdge T3 array allows each customers storage system to grow as their storage needs grow.

Compatibility

The following table lists the supported features and product attributes for both the Sun StorEdge T3 array for the workgroup and the Sun StorEdge T3 array for the enterprise.

Feature	Description
Sun StorEdge T3 array for the enterprise	<ul style="list-style-type: none">• Two- and eight-array configurations• Fully populated with nine 36.4-GB or 73.4-GB, 10000-rpm FC-AL drives• Preconfigured for RAID 5 (8 + 1 no hot-spare)• Dual hot-swap redundant RAID controllers with 1 GB mirrored cache• RAID 0, 1 (1+0), or 5• Two to four LUNs per partner pair• Rackmounted options (two and eight-array configurations, with or without hubs or switches)• Rackmount-ready or tabletop options (two-array configurations)• Sun StorEdge Component Manager 2.2, or higher, software, license, media, and documentation• VERITAS Volume Manager 3.0.4, 3.1, or 3.2 software for soft partitioning (included with all Sun StorEdge T3 array for the enterprise; Tier 2 server license)• Installation/configuration via Sun StorEdge ArrayStartSM program (optional)
Sun StorEdge T3 array for the workgroup	<ul style="list-style-type: none">• One-array configuration• Fully populated with nine 36.4-GB or 73.4-GB, 10000-rpm FC-AL drives• Preconfigured for RAID 5 (8 + 1 no hot-spare)• Single RAID controller with 1-GB cache• RAID 0, 1 (1+0), or 5• Supports up to two LUNs per array• Rackmount-ready or tabletop options• Sun StorEdge Component Manager 2.2 software, license, media, and documentation



Feature	Description
Software features	<ul style="list-style-type: none"> • GUI • CLI management interface
Supported Solaris™ Operating Environment-based software	<ul style="list-style-type: none"> • Sun StorEdge Instant Image 2.0, or higher • Sun StorEdge Network Data Replicator (SNDR) 2.0, or higher • Sun Cluster 3.0 • Solstice Backup™ 5.5.1, or higher • Solstice DiskSuite™ 4.2.1, or higher • Sun Enterprise Server Alternate Pathing (AP) 2.3.1, or higher • Sun StorEdge Data Management Center 3.0, or higher • Sun StorEdge Component Manager (CM) 2.1, or higher • StorTools™ 3.3 diagnostics, or higher • VERITAS NetBackup (NBU) 3.4, or higher • VERITAS Volume Manager 3.0.4, 3.1 or 3.2 (or later) with DMP (VxVM) • VERITAS File System (VxFS) 3.3.3, or higher • VERITAS Cluster Server 1.3, or higher



Enabling Technology

Features Overview

In today's Net economy, fast response and continuous availability are key to any company's survival. The competition is a mere click away, and, in general, if Internet users have to wait much longer than five seconds for requested data, they may be gone and may not return. Sun understands this market dynamic. As a result, Sun built cache mirroring and fully redundant, hot-swappable hardware components into the Sun StorEdge™ T3 array, which provide customers with the IT muscle to deliver information continuously, where and when they need it.

The Sun StorEdge T3 array employs impressive technologies to build break-away solutions:

- Advanced architecture that uses full Fibre Channel connectivity, loop-switching design, and failover security
- Simple and capable storage unit design that is highly scalable and easily accessed for adding, upgrading, and swapping components
- Supported by Sun StorEdge Component Manager software, which provides server-based, single-point administration for attached arrays and their hardware components
- Sun Remote Services program support for continuous remote systems monitoring

Fibre Channel Loop-Switching Architecture

Sun loop-switching architecture is at the heart of the Sun StorEdge T3 array controller units ability to scale performance linearly as capacity grows. Traditional SCSI and Fibre Channel arrays typically consist of one or more controllers connected to one or more disk trays by dual Fibre Channel loops. In a bandwidth-intensive application, for example, a traditional arrays two internal disk loops could be completely saturated with just two disk trays delivering data concurrently. Additional capacity does not increase bandwidth.

The Sun StorEdge T3 array employs a radically different architecture to connect its disks to the controller functions. Sun has allowed segmentation of the disk interconnect loops on a tray-by-tray basis into smaller dedicated sub-loops which can all operate independently and concurrently without interference. This function is under control of the Sun StorEdge T3 array controllers and is managed dynamically in order to help optimize performance and deliver complete failover data access.

A small Sun StorEdge T3 array system consisting of two controller units typically segments one set of redundant loops into two dedicated data loops — one per controller — so that each controllers data accesses can proceed without interference from the other controller. The second loop is typically not segmented and is used to transfer mirrored cache writes between the controllers, again without interfering with either controllers disk accesses.

A full rack of Sun StorEdge T3 array controller units organized into four failover partner groups delivers 800 MB/second of dedicated data access internal loop bandwidth and 400 MB/second of dedicated cache mirror bandwidth for a total of 1.2 GB/second of internal data transfer bandwidth. Compared to the more traditional array design using just two internal loops delivering only 200 MB/second of total bandwidth, the Sun StorEdge T3 arrays ability to scale is virtually unparalleled in the industry.



Three-Dimensional Scalability: Capacity, Performance, and Availability

The ability to adapt to rapid and unexpected change in I/O workloads, capacity requirements, and system-delivered bandwidth while delivering continuous availability is another design requirement that is becoming absolutely necessary in the dot-com world. The Sun StorEdge T3 array addresses this requirement by delivering some of the best overall scalability in the market. The system scales capacity, performance, and availability in a linear and predictable fashion. As capacity increases, performance increases by a predictable amount. Availability scales as additional data paths and redundancy are added with each array.

In reality, capacity, bandwidth and I/O per second (IOPS) can change independently of each other. For example, an increase in streaming audio or visual content may drive bandwidth, while a SAP application may drive IOPS and capacity. The Sun StorEdge T3 arrays revolutionary switched-loop architecture allows customers to scale on any combination of capacity, IOPS, or bandwidth without being forced into "over buy" situations which are common with many of the large traditional storage solutions that have been offered in the past.

To simply add capacity, a Sun StorEdge T3 array unit can be added to a configuration. If I/O increases, a controller can be added to the unit to effectively double the cache and potential IOPS into the storage subsystem. The same flexibility applies to bandwidth-sensitive applications, allowing customers to address ever-changing and dynamic application demands.

Enhanced RAID 5 Performance

The Sun StorEdge T3 array is designed for outstanding Fibre Channel RAID 5 performance. The system employs a revolutionary hardware-pipelined, XOR engine (pXOR) that operates, and can calculate parity data, at the full internal data bus speed of the array. This helps reduce the requirement for cache memory accesses, which serve as a bottleneck in most RAID 5 arrays.

The Sun StorEdge T3 arrays pXOR unit helps improve performance by reducing mechanical disk seek and latency periods. The pXOR engine does this by accumulating partial XOR sums (up to 128), allowing the array to accumulate stored write-behind data in its cache memory and deliver it to the disk back-store in large chunks that can take advantage of the command sorting feature of the disks. Reducing mechanical wait states is among the best ways to improve application performance.

In addition, the Sun StorEdge T3 array's ability to optimize physical disk I/O access patterns directly delivers improved application response time and improved overall system performance.

ECC Protection in the 1-GB Cache Controller

Sun has added ECC circuitry on all critical data paths. All of the data paths within the XOR array are ECC protected, including the data cache and the accumulator memories. ECC protection provides the transparent and automatic detection and correction of internal errors to the Sun StorEdge T3 array.



System Architecture

This section provides an overview of the system architecture of the Sun StorEdge™ T3 array. For more detailed information on the architecture of this array, refer to the Sun StorEdge T3 Array with 1-GB Cache Controller Technical White Paper (SunWIN number 311986).

Basic Architecture

The Sun StorEdge T3 arrays architecture begins with a basic *controller unit*. The standalone controller unit is the smallest possible array configuration. The architecture integrates disks, data cache, hardware RAID, power, cooling, uninterrupted power supply (UPS), diagnostic capabilities, and administration into a versatile, standalone component. The controller unit includes external connections to a data host (or hub or switch), and to a management network.




Figure 2. Logical view of the Sun StorEdge T3 arrays controller unit, with separate data path connected to an application host and administration path connected to a management host

Two units may be paired in a *partner group* to create a configuration with redundant controllers and redundant data and management paths, allowing for cache mirroring, controller failover, and path failover capability. The partner group is thus the minimum storage configuration for enterprise environments that call for high availability.




Figure 3. Logical view of the Sun StorEdge T3 controller units

External Interfaces

The Sun StorEdge T3 array controller module includes three external interfaces. A Fibre Channel arbitrated loop (FC-AL) port transports data to the application host. An Ethernet port handles administrative (configuration, monitoring) communication with the management host. And, finally, an RJ-45 serial port (RS-232) is used for advanced service procedures, such as boot diagnostics.

Only application data travels across the FC-AL channel, and only administrative information moves across the network channel. This separation of responsibilities has several advantages. It enables greater reliability, because diagnostic reporting is preserved even when the host channel is down. It provides greater performance, because administrative traffic does not interfere with application I/O. It also provides greater security: a junior system administrator may be granted access to monitor and service the unit without access to the application server or even application data on the Sun StorEdge T3 array.

In addition, the separate administrative path enables greater efficiency and productivity in the data center by allowing for centralization of administration. A site may have multiple, heterogeneous,

geographically distributed application servers with local Sun StorEdge T3 arrays. All Sun StorEdge T3 arrays can be connected via Ethernet and TCP/IP to a single management server, which provides centralized administration with a single user interface.

Administration Path

On the Sun StorEdge T3 array, the administration path provides connectivity from the controller card in the master controller unit (and alternate master controller unit) to all FRUs throughout the entire partner group. It also connects the master controller unit (and alternate master controller unit) to the external Ethernet and serial ports. It is over this path that configuration, diagnostics, and monitoring takes place.

The Sun StorEdge T3 array has two internal serial lines connecting the CPU with all non-disk FRUs. These lines are used as redundant internal administration paths, communicating configuration, control, monitoring, and diagnostic information. The serial lines extend through the daisy-chain unit interconnect cables to all units in a partner group, creating a single administrative domain.

The external Ethernet administration path is used to exchange configuration and monitoring information between the Sun StorEdge T3 array unit's CPU and the management host(s). The CPU has no access to the application data, and no application data is available via the administration path. This separation of data and administration paths provides security by shielding application data from those individuals performing everyday service and administration. The path disunion also means a path for communicating with the Sun StorEdge T3 array remains available even if the data path or application host has failed.

In a partner pair configuration, one controller unit is always designated as the master controller unit, and its partner is the alternate master controller unit. Both controllers are actually FC active, but only the master controller is active for management.

All administration and all external communication, on behalf of both units in a partner pair configuration, are conducted via the master controller. Only in case of master controller failure does the alternate master controller take over administration.

3-D Scalability

The Sun StorEdge T3 arrays architectural design provides for 3-D scalability. Controller unit configurations can be added to meet requirements in capacity, bandwidth, and transaction rate as business requirements grow. This scalability and flexibility protects the original investment, and it allows the customer to "pay as you grow." Sun StorEdge T3 arrays simply can be added to the existing storage infrastructure as business requirements change. In addition, the centralized administration capability of the Sun StorEdge T3 array answers increasing storage needs without adding management complexity.

Availability, Reliability, and Serviceability via FRUs

In the Sun StorEdge T3 array, all active components are designed to be N+1 redundant, including disks, power supply, fans, and UPS. On the back end, loops, loop switching, diagnostics, and administration channels are also redundant. When configured in a partner group, even controllers, host channels, and external administration channels are redundant.

Active components are consolidated into four types of FRUs: disk drive, power/cooling unit (PCU), unit interconnect card (UIC), and controller. All FRUs are hot-swappable to prevent servicing downtime and to help minimize mean time to repair (MTTR).



Disk Drives

Every Sun StorEdge T3 array controller drive tray unit is configured with nine bi-directional dual-ported 10000-rpm FC-AL disk drives of either 36.4-GB or 73.4-GB capacity (as disks of larger capacity become available, the Sun StorEdge T3 array will support those drives). The disk drives are concealed by a removable front bezel which provides electromagnetic interference (EMI) shielding. Each disk FRU consists of a custom enclosure which holds 1.6-inch high disk drive. The drive plus custom drive enclosure constitutes the disk FRU. Individual disk drives are not visible to the application host; rather, they are configured into one or two RAID 5, RAID 0, or RAID 1 (1+0) logical volumes. The ninth disk may optionally be configured as a non-floating hot-spare.

Each drive has a private region of 200 MB reserved for system use. All remaining capacity is available for use by the application host. On the master and alternate master controller units, the system area is used on all drives as a nine-way mirror, containing a copy of the operating system, file system, and firmware. Multiple versions of firmware may be saved, allowing the flexibility to back out or revert to an earlier version if necessary. On the master controller unit, the system area also includes configuration information, system log, and other assorted files for internal use.

Power Cooling Unit (PCU)

Each Sun StorEdge T3 arrays tray includes two redundant power and cooling units (PCUs). Each PCU has an external power connection, allowing for connection to two independent power grids or circuits. There is one internal 325-Watt auto-switching power supply per PCU. In case of external power failure or PCU failure, one power supply is sufficient to indefinitely power the nine drives contained within a Sun StorEdge T3 array controller unit drive tray.

The power supplies share the voltage load (load sharing) and the current load is evenly distributed across both power supplies (load balancing), both of which combined greatly increase the life of the power supplies compared to those systems which either do not load share or load balance or both. Virtually all power supplies must operate within specific boundary conditions, which are above 15 percent and below 85 percent of the average load rating — any constant load outside those boundaries significantly shortens the life of the power supplies which can result in frequent failures. To prevent failures caused by overload, the power supply circuitry inside the Sun StorEdge T3 array is designed to evenly split both the voltage and the current loads across both power supplies, thereby insuring approximately a 50/50 load on each supply, which is well within the 15/85 boundary condition.

The four cooling fans in the Sun StorEdge T3 array are fully redundant — if any one or even two fail, the controller unit remains operational indefinitely, as the nine hard drives continue to be cooled within the predetermined operating temperature range. Moreover, these cooling fans are electrically independent from the power supplies, so that if one power supply should fail, all four cooling fans continue to spin indefinitely.

Unit Interconnect Card (UIC)

The unit interconnect card (UIC) has three interrelated functions: join units in a daisy chain, perform back-end loop-switching, and to maintain diagnostic state registers.

Loop resiliency circuitry (LRC) provides internal Fibre Channel switching and bypass capability, connecting all back-end Fibre Channel components. This includes all nine drives, one of the back-end QLogic 2200 Fibre Channel interfaces on the controller board, and one pair of ports on the Unit Interconnect Card, used to join units together in a daisy chain. The LRC also provides Fibre Channel clock regeneration, which is crucial to preventing signal jitter.

The remainder of the UIC consists of an 8051 CPU connected to FLASH PROM FRU ID, control/sense registers, and a dual UART. The CPU is connected to the LRC through the registers, which collect and



maintain component status and diagnostic states. The administration serial line from the controller CPU connects through the midplane to the UIC CPU, and then to the dual UART, which extends one serial line to each of the two back-end external ports.

The LRC switching capability gives the Sun StorEdge T3 array back-end reliability, availability, and serviceability (RAS) capabilities. In normal operation, each UIC selectively enables a subset of its components for purposes of load balancing, cache mirroring, and redundancy.

When two controller units (Unit 1 and Unit 2) are configured as a partner pair, Unit 1, Loop 1 UIC enables the Loop 1 back-end controller ASIC and drives 4 through 9. Any I/O requests to drives 4 through 9 on Unit 1 travel through this loop; the Loop 1 back-end daisy-chain ports are not enabled on the loop. Therefore, Loop 1 in Unit 1 and Loop 1 in Unit 2 are maintained as two independent loops.

The Loop 2 UIC in Unit 1 enables the Loop 2 back-end controller ASIC, drives 1 through 3, and the Loop 2 back-end daisy chain ports; the same applies to Unit 2. Because the daisy-chain ports are enabled, there is just one Loop 2 for the entire partner group. So, a single Loop 2 spans both units: from the Unit 1 Loop 2 ASIC to its Loop 2 UIC (including drives 1-3), to the Unit 1 Loop 2 back-end daisy-chain ports, to the Unit 2 Loop 2 back-end daisy-chain ports, to its Loop 2 UIC (including drives 1 through 3), to the Unit 2 Loop 2 controller ASIC.

This back-end configuration of three loops for the partner group provides static load balancing in normal operation. Two of the loops, one in each unit, each carry the load for six drives. The third loop, which spans both units, also carries the load for six drives (three in each unit), plus any mirrored cache writes.

The LRC switching function also provides powerful diagnostic capabilities. In case of failure on a loop, a diagnostic routine can be run in which the UIC systematically switches components in and out of the loop until the offending component has been identified. Then, the UIC can switch off, or bypass, that component until it is replaced. By fencing off the failed component from the loop, the loop can be restored to active use.

In addition to the Fibre Channel port used to link the data path between units, the UIC includes a serial port to link the administrative path between units. The Fibre Channel and serial ports are combined into a single non-standard physical connector. A single non-standard cable, called the unit interconnect cable, combines the Fibre Channel and serial lines that link the Sun StorEdge T3 array units.

The application data path is used to process I/O between the application host and disks exclusively. No configuration or monitoring is performed over the data path to the host, other than normal SCSI inquiry requests. No configuration or monitoring is handled over the internal data paths, other than to store/retrieve configuration and monitoring data to/from the reserved system area on the disk drives.

Data movement is by DMA to and from the QLogic 2200 ASICs. All data goes through the cache and the inline XOR engine. Because all XOR operations are completed as data moves in and out of cache, there is virtually no performance penalty for calculating RAID 5 parity on the Sun StorEdge T3 array. When write-behind mode is enabled, host writes are acknowledged when they reach cache, and are later destaged to disk. When cache mirroring is enabled in a partner group configuration, host writes are acknowledged only after they both reach cache and are copied to the partner controller's cache. Write data is later destaged to disk according to cache destage rules, based on idle time, utilization, and error conditions.

Within the Sun StorEdge T3 array, all data travels through cache. The system uses static load balancing to spread I/O across the two back-end loops. Data destined for drives 4 through 9 is sent through Loop 1, while cache mirroring data, plus data destined for drives 1 through 3, is sent through Loop 2. Should one loop become disabled, the surviving loop handles the full back-end load.



RAID Controller Card

The RAID controller card provides cache, RAID management, administration, diagnostics, and external interfaces. Controller units include one controller FRU. Two controller units are paired in a partner group for cache mirroring and controller redundancy.

The controller is both the data processing and administrative "brain" of the Sun StorEdge T3 array. It provides all the Sun StorEdge T3 arrays external interfaces and controls all back-end activities, whether they be related to data management or administration. The controller's data host interface is a QLogic 2200 FC-AL interface ASIC. It connects to a 64-bit, 33-MHz PCI bus, which functions as the backbone of the Sun StorEdge T3 array. Also residing on the PCI bus is 1 GB of SDRAM cache, with a custom inline FPGA XOR engine that has 2 MB of VRAM. Two more QLogic 2200s provide the interfaces to two back-end FC-AL loops. Finally, there is a bridge chip on the backbone, providing a transition to a 32-bit, 33-MHz PCI administration bus to the controller CPU.




Figure 4. Controller architecture

The controller has an administration bus, which connects FLASH PROM/FRU ID, external 10/100BASE-T Ethernet port, and serial line interface to an external RS-232 (RJ-45). The external ports use standard connectors: Native LC, SFF – FC-AC connectors or FC-AL, RJ45 for Ethernet, and RJ45 for the serial port (RS-232). Connector pinout specifications are provided in the product documentation. The administration bus also connects to two internal serial lines to the unit interconnect cards.

The controller CPU is a PPC750 running at 400 MHz. Note that the CPU is not on the system bus; its involvement with application data is limited to managing the data, not manipulating it. Although the CPU controls DMA transfer of data between host interface FC-AL ASIC and cache, and between cache and back-end FC-AL ASICs, data never travels through the CPU itself. Even XOR parity calculations are performed not by the CPU but rather by the inline XOR engine, as data moves in and out of cache. This

efficient data flow is a key factor in achieving superior RAID 5 performance in the Sun StorEdge T3 array.

Because the Ethernet and serial lines are also isolated from the controller backbone, it is not possible to transfer data through the external Ethernet or serial ports. They are available exclusively for administration, just as the external FC-AL loop is available exclusively for application I/O.

To achieve high availability using the Sun StorEdge T3 array, two controller units are configured in a single partner group. Although united in a partner group, each controller unit processes data to its disks independently. However, write data that is placed in cache to be destaged at a later time is mirrored to the partner controller unit cache prior to returning an ACK through the host interface. The mirroring is accomplished over one of the two back-end FC-AL loops. Under normal operation, Loop 1 in each controller unit remains independent of the partner's Loop 1. Loop 2, which provides the cache mirroring path, is a continuous loop between the two partners.

Each controller has its own data path to the application host, hub, or switch. In normal operation, each path to a given controller carries only data for the volumes contained within the same StorEdge T3 array as its controller — the path serves as the active or primary path for those volumes. However, each controller/path is also capable of carrying data intended for its partner controller/path — the path also serves as the secondary or passive path for its partner if necessary.

The failover scheme for the Sun StorEdge T3 array distinguishes between two types of failures: path failures and controller failures. The failover approaches for both failure types have much in common, but they also exhibit some distinct differences.

Path Failure

Path failure occurs when I/O to a unit is interrupted for any reason other than for a controller failure — the failure could be in a cable, in a host adapter, or could even occur because of a non-I/O root cause such as removal of an application host system I/O board. Regardless of the cause of the interruption, I/O requests targeted at a LUN eventually times out. The I/Os are then redirected to the alternate path for that LUN — the path to the other Sun StorEdge T3 array in the partner group. The redirection is managed on the application host by the alternate pathing software appropriate to that host. On Solaris™ Operating Environment platforms, the user can choose between Solaris Operating Environment MPXIO (Multi-path I/O), Solaris Operating Environment alternate pathing (AP) and VERITAS dynamic multipathing (DMP); on other platforms, the user can choose between VERITAS DMP and the Sun-supplied failover driver available with the Sun StorEdge T3 array.

When a controller receives an I/O request targeted at a LUN belonging to its partner controller, it verifies its partner controller is healthy, and then takes over control of the LUN. This procedure is called a LUN failover. The back-end connection between Loop 1 of Unit 1 and Loop 1 of Unit 2 is healed, so that there is now a single Loop 1 for the entire partner group, as well a single Loop 2 (Recall that there is already a single Loop 2 in normal operation, which provides a path for cache mirroring). I/O targeted at the partner's LUN is directed across the back end (Loop 1 or Loop 2, as appropriate) and into the proper drives.

In a path failure scenario, both controllers remain healthy. So, if write-behind cache was enabled, it remains enabled. Writes go into the controller with the live path, are written into the local cache, copied via Loop 2 into the partner's cache, acknowledged to the application host, and, in due course, destaged to disk across the back-end channels.

Even when a failure occurs on the data path to the master controller, administrative procedures continue unchanged. The controller and administrative path remain healthy, so the master controller continues administration even though it may have temporarily ceased performing data I/O operations.



Also note that no special communication, or "heartbeat" is needed between the application host and partner group to initiate failover or fail-back.

With the exceptions of SCSI inquiry or read/write of block 0 of any LUN, the occurrence of I/O down the alternate path automatically triggers failover. Likewise, resumption of I/O down the primary path (except for SCSI inquiry or read/write of block 0 of any LUN) automatically triggers fail-back. It is the responsibility of the alternate pathing software to "ping" the primary path periodically (by sending a SCSI inquiry or read/write of block 0 of any LUN), to see if it has been restored.

Implicit LUN Failover vs. Explicit LUN Failover

The LUN failover mechanism just described is more specifically called "implicit LUN failover" (ILF) because there is no explicit communication between host and Sun StorEdge T3 array that anything has gone wrong; rather the need to perform a LUN failover is implicit in the fact that I/O has started coming down the alternate path. This mechanism works well when the Sun StorEdge T3 array partner pair is connected to a single host, and eliminates the complexity and overhead of explicit communications between host and Sun StorEdge T3 array when a path fails.

However, when a Sun StorEdge T3 array partner pair is connected to two hosts which share LUNs, it is possible for a path failure to occur that affects only one of the two hosts. For example, an HBA could fail in one of the hosts, or the path between switch and one of the hosts could fail. The host which sees the failure sends I/O down the alternate path, while the second host, unaware of any failure, continues to send I/O down the primary path. This causes a "ping-pong" effect: a series of failovers and failbacks that thrash the partner pair and bring I/O to a halt.

In this case, explicit communication regarding the failure is required between Sun StorEdge T3 array partner pair and the two hosts, in order to prevent the failure. This mechanism, combining LUN failover with Sun StorEdge T3 array communication to the hosts, is called, not surprisingly, "explicit LUN failover" (ELF). The Sun StorEdge T3 array partner pair coordinates with ELF-aware alternate pathing software on the host to ensure that if one host has suffered a path failure causing LUN failover, the primary path to the other host is locked out, forcing it to switch to its alternate path as well. Examples of ELF-aware alternate pathing software on Solaris Operating Environment platforms include Solaris Operating Environment MPXIO and VERITAS Volume Manager (VxVM) 3.2 (or later) DMP. On other platforms, the user can choose between VERITAS DMP and the native failover driver supplied with the Sun StorEdge T3 array.

Controller Failure

To the application host, a controller failure appears identical to a path failure, and the response and recovery procedures are also identical. I/O requests down the channel of the failed controller time out. The host-based alternate pathing software re-routes I/O down the failover channel. The software periodically pings the primary channel, and when it gets a response, it then re-routes I/O back to the primary channel. The time it takes to effect a controller failover is slightly longer than the time needed to effect a LUN failover.

On the Sun StorEdge T3 array, the path failover resulting from controller failure causes LUN failover, as with any other path failover. But there the similarity ends — loss of communication heartbeat informs the surviving controller its partner controller has failed. The surviving controller then takes significant additional recovery actions on both the data and administration fronts.

Write-behind cache and cache mirroring are both disabled. Both back-end Loop 1 loops on both partners are joined into a single loop, just like Loop 2. Any uncommitted write data in the surviving cache is then flushed to disk, including mirrored uncommitted write data destined for LUNs of the failed controller.



If the alternate master controller has failed, then no administrative changes are needed. If the master controller has failed, then the alternate master controller must take over the role of the master controller. The alternate master controller takes on the MAC and IP addresses of the master controller, as well as the host name, activates its dormant Ethernet connection, and resumes IP activity on its administration path. As far as the network and any management consoles are concerned, nothing about the network topology has changed — the alternate master controller, for all intents and appearances, is now the master controller.

Because the MAC and IP addresses for the partner group have not changed, there is no need to change routing tables and maps. When a failed Sun StorEdge T3 array controller is replaced, insertion of the replacement controller is automatically detected, the controller is booted, and the unit's heartbeat is restored. LUN fail-back is achieved the same manner as non-controller path failover: when the host pings the primary path and receives a response, it re-routes I/O back to the primary path. However, if the controller failure resulted in administrative controller failover, controller replacement does not cause administrative fail-back. The former alternate master controller continues to act as the new master controller, rather than suffer the overhead of the administrative fail-back. The former alternate master Administrative Network controller continues to act as master controller until a system reset or power cycle, or failure of that controller. In any of these cases, the controller in Unit 1 is restored as master controller.

Note that even when the alternate master controller takes over as master controller, the populated system area continues to be a nine-way mirror on Unit 1. This means the system boots from a firmware image on the drives of Unit 1, and the syslog continues to be written on Unit 1 drives.

Data Flow

The application data path is used to process I/O between the application host and disks exclusively. No configuration or monitoring is performed over the data path to the host, other than normal SCSI inquiry requests. No configuration or monitoring is handled over the internal data paths, other than to store/retrieve configuration and monitoring data to/from the reserved system area on the disk drives.

Data movement is by DMA to and from the QLogic 2200 ASICs. All data goes through the cache and inline XOR engine. Because all XOR operations are completed as data moves in and out of cache, there is virtually no performance penalty for calculating RAID 5 parity on the Sun StorEdge T3 array. When write-behind mode is enabled, host writes are acknowledged when they reach cache, and are later destaged to disk.

When cache mirroring is enabled in a partner group configuration, host writes are acknowledged only after they both reach cache and are copied to the partner controllers cache. Write data is later destaged to disk according to cache destage rules, based on idle time, utilization, and error conditions. Within the Sun StorEdge T3 array, all data travels through cache. The system uses static load balancing to spread I/O across the two back-end loops. Data destined for drives 4 through 9 is sent through Loop 1, while cache mirroring data, plus data destined for drives 1 through 3, is sent through Loop 2. Should one loop become disabled, the surviving loop handles the full back-end load.

Data Cache

The primary purpose of the data cache in the Sun StorEdge T3 array is to provide a low latency buffer for write data, allowing writes to be quickly acknowledged to the application host. The cache is especially crucial to RAID 5 write performance, because it can coalesce several partial-stripe writes into a single read/modify/write operation. A secondary benefit of the cache is to buffer read data, allowing for low latency on repeated reads of the same data.

Adaptive cache is a key feature of the Sun StorEdge T3 array. The algorithms used for allocating, coalescing, and flushing data are automatically and dynamically adjusted based on I/O patterns. This



limits the amount of cache configuration needed to be performed by the user, thus greatly simplifying administration, improving ease of use, and enabling optimal cache behavior for current I/O patterns.

Each Sun StorEdge T3 array controller includes 1-GB SDRAM data cache. Cache organization and behavior are tightly coupled with LUN stripe width and Sun StorEdge T3 array block size (the amount of data in the stripe that goes on each disk). The Sun StorEdge T3 array block size is a system configuration parameter set by the user to be 16, 32, or 64 KB. The cache buffer size equals the block size, so the block size configuration parameter defines both the size of the cache buffers and the unit of data written to each disk in a RAID stripe.

Because the cache size is fixed at 1 GB, the number of cache buffers varies. There are 65536, 32768, or 16384 cache buffers, depending on block size of 16, 32, or 64 KB, respectively. Each cache block is composed of eight segments. This means that segment size is 2, 4, or 8 KB, for a block size of 16, 32, or 64 KB, respectively. Segmentation of the cache block is crucial to performance of the adaptive cache, because the segment size defines the Sun StorEdge T3 array controller units atomic I/O to disk. This means not only is it possible to optimize for partial stripe reads and writes, but it is also possible to optimize for partial block reads and writes.

Note the host I/O size is not necessarily the same as the Sun StorEdge T3 array block size. There are cases where optimal performance can be achieved when the segment size matches the host I/O size or where the stripe size matches the host I/O size.

All data travels through cache, and all data in the cache is read cache. Therefore, all data — whether read or written into cache; whether dirty or clean — is available for subsequent host read requests as a cache hit. Dirty write data is limited to 204.8 MB to help ensure that, in case of power failure, there is sufficient battery power to ensure all dirty data is safely written to disk. Writes are cached only when write-behind mode is enabled.

Dirty write data is flushed to disk under the following conditions:

- **Demand flushing** — When the dirty write data 80-percent threshold has been reached, the CPU causes as many as 20 stripes of dirty data to be flushed to disk. The least recently used stripes are chosen to be flushed.
- **Idle time flushing** — If no host requests are received for a full second, one stripe is flushed. After 10-ms with no host requests, two stripes are flushed. After another 10 ms, four stripes are flushed, and so on, continuing up to a maximum of 128 stripes per flush, until either the cache is emptied of dirty write data or a host command is received. The stripes chosen for each flush are those least recently used.
- **LUN flushing** — If a volume is unmounted by the user, any dirty cache associated with that volume is flushed. All host commands are queued while the LUN data is flushed.
- **Controller flushing** — All dirty data in the cache is flushed when the system is shut down, when the user manually forces a "sync" operation, when there is a controller failure, or when a power failure has occurred.

Configuration Details

In each Sun StorEdge T3 array, the nine disk drives are configured into one or two logical volumes or LUNs, which are the atomic units presented to the application host. In other words, the application host does not see the individual disk drives.

Sun StorEdge T3 array configuration tools employ the following configuration rules:

- LUNs must consist of whole disks.
- A disk may not be partitioned into different LUNs.
- LUNs may not span physical drive trays.



- One LUN is equal to one array group. An array group, or drive group, is the exact equivalent of one RAID level; a RAID level is either a stripe (RAID 0 or RAID 5) or striped and mirrored (RAID 1+0); so, an array group using RAID 5 can consist of either three, four, five, six, seven, eight, or nine drives; an array group using RAID 1 (1+0) can consist of exactly two drives.
- There is a minimum of one LUN per tray and a maximum of two LUNs per controller drive tray unit.
- There are three RAID levels to choose from, as follows:
 - **RAID 0** — Striping without parity or mirroring; if a drive fails, all the data is lost and non-recoverable unless it is available somewhere else within the organization. *Note that this RAID level, due to its lack of redundancy, is not usually recommended.* The minimum size for a RAID 0 LUN is two disks and the maximum size is nine disks, and there can be up to two RAID 0 LUNs per drive tray if no other RAID level is used. Use this level for high-speed streaming of large file reads (for example, video) of non-critical data which is easily available elsewhere within the organization.
 - **RAID 1 (1+0)** — Each data block in a RAID 1 (1+0) volume is mirrored on two drives, and the blocks are striped across all the drives in the volume. If one of the mirrored pair fails, the data from the other drive is used. Because the data is mirrored in a RAID 1 configuration, the volume has only half the capacity of the assigned drives. For example, if you create a four-drive RAID 1 volume with 18-GB drives, the resulting data capacity is $2 \times 18 = 36$ GB.
 - **RAID 5** — Striping with rotated parity (both file data and parity data are evenly distributed across all drives in the array). The minimum size for a RAID 5 LUN is three disks and the maximum size is nine disks, and there can be up to two RAID 5 LUNs per drive tray if no other RAID level is used. Use this level for most applications which do not require the special characteristics of the above RAID levels.
- If a hot-spare is used, it must be used with all LUNs in a given tray, and it must be declared when the first LUN on a tray is created.

These rules may appear restrictive, but they provide for a much greater degree of simplicity of configuration compared to competitive systems on the market today. The user just makes three basic decisions:

- Will there be a hot-spare drive or not?
- How many LUNs, one or two? If two, how many disks for each LUN?
- What RAID level is required for each LUN?

LUN Configurations per Tray

Choosing between one or two LUNs per tray requires balancing considerations of capacity, performance, and availability.

A configuration of a single LUN per tray provides better performance over dual LUNs per tray because of the additional administrative overhead required for two managing LUNs. However, one LUN per tray requires a RAID 5 stripe of at least eight disks, and a failure of one of those disks can result in double the reconstruct time — which is the time the RAID is in a vulnerable state because the subsequent failure of another drive within the same stripe causes data loss — compared to a RAID 5 stripe of only four disks. And, using smaller drives in a RAID 5 stripe further decreases the reconstruct time.

Note: *If small 1-GB or 2-GB LUNs are desired, 14 partitions (two LUNs x seven Solaris Operating Environment partitions) still do not create enough volumes to utilize all disk capacity. In this case, the use of host-based volume management software (such as VERITAS Volume Manager) is needed to create*



the required quantity of sub-volumes, regardless of whether there are one or two native volumes on the Sun StorEdge T3 array.

LUN configurations can be created as follows:

- **RAID 1 (1+0)** — Minimum two disks, maximum nine disks or 8 + 1 hot spare.
- **RAID 5** — Single LUN: either nine disks (8+1) and no hot-spare or eight disks (7+1) with hot-spare. Dual LUN: Either one 4-disk (3+1) LUN and one 5-disk (4+1) LUN and no hot-spare or two 4-disk LUNs (3+1 each) with hot-spare. Sun StorEdge T3 array hardware and firmware have been optimized for RAID 5. In most cases, RAID 5 outperforms RAID 1. If read/write ratio is 1:1 or higher, use RAID 5.
- **RAID 0** — Single LUN: nine disks. Dual LUN: nine disks. Use of RAID 0 is advised only in conjunction with an external form of data protection, such as using host-based mirroring across two Sun StorEdge T3 arrays, or in the rare case of using host-based RAID 5 stripes across multiple trays.

# of LUNs	Hot Spare?	RAID Level for LUN 1	# Drives Used for LUN 1	RAID Level for LUN 2	# Drives Used for LUN 2
1	NO	5	9	-----	-----
1	YES	5	8	-----	-----
2	NO	5	3	5	6
2	YES	5	3	5	5
2	NO	5	4	5	5
2	YES	5	4	5	4
2	NO	5	5	5	4
2	YES	5	5	5	3
2	NO	5	6	5	3

Note: RAID 0 not shown because of no redundancy for RAID level 0.



Reliability, Availability, and Serviceability (RAS)

Reliability

Reliability features of the Sun StorEdge™ T3 arrays include the following:

- Error checking and correction on disk drives
- Skip sectors and spare cylinders on disk drives
- Automatic sector reallocation on RAID controller
- Link redundancy chip and 8- to 10-bit encoding on FC-AL loops
- ECC on data cache
- Passive midplane (except ID signature) and temperature sensor

Availability

Availability features of the Sun StorEdge T3 array include the following:

- Hot-swap redundant load-sharing/load-balancing auto-sensing 110VAC/220VAC power supplies with dual power cords
- Built-in hot-swap redundant UPS batteries for cache backup which power the controller unit and its nine disks so the contents in cache can be destaged to the disks upon sense of power loss (then graceful shut down), helping to ensure no data is lost no matter how long the power is out
- Four hot-swap redundant electrically independent cooling fans
- Hot-swap redundant unit interconnect cards
- Hot-swap redundant dual-ported FC-AL drives (RAID 5), non-floating hot-sparing capability, and dual backend drive loops per controller
- Hot-swap redundant RAID controllers for automatic failover and cache mirroring (Sun StorEdge T3 array for the enterprise models)
- Redundant host interfaces (Sun StorEdge T3 array for the enterprise models)

Serviceability

The system administrator may selectively replace unit components as needed. Serviceability features of the Sun StorEdge T3 array include the following:

- Low FRU count (four excluding cables)
- The four basic FRUs — the drives, the PCUs, the interconnect cards, and the RAID controller — can be hot-swapped with no tools required
- Status/failure LED on each FRU
- Each FRU is electronically identifiable
- Fibre Channel world-wide name support
- Online installation, scaling, and service
- Upgradeable drive firmware (with only the associated volume off-line during upgrade)



- Detection and reporting for incorrect drive position
- Automatic drive-ID selection

Specifications

Specifications for the Sun StorEdge™ T3 array controller unit array include those shown in the following table.

Performance	Specifications
Bandwidth	93 MB/second writes, 96 MB/second reads
IOPS	<ul style="list-style-type: none"> • Sun StorEdge T3 array for the workgroup = 7,250 per controller unit • Sun StorEdge T3 array for the enterprise = 14,500 per controller unit (cache to cache)
RAID disk capacity	Nine x 36.4-GB (327-GB total), or 73.4-GB disks (660 GB total) raw
Controller read/write cache	1 GB
Physical Planning	
Tabletop unit dimensions	5.5 inches high 3U (13.97 cm) 17.5 inches wide (44.45 cm) 18.5 inches deep (46.99 cm)
Weight	67 pounds (30.15 kilograms) maximum with half-height drives
Input power	450W maximum
Auto-ranging	100/240 VAC at 47 to 63 Hz single-phase
FRU access at front	Disk drives
FRU access at rear	Controller card, unit interconnect cards, and power/cooling units
Environmental (operating)	
Temperature	5 to 35° C, maximum gradient 20° C per hour
Relative humidity	20 to 80% noncondensing, maximum gradient 10% per hour
Effective altitude	-1,000 to +10,000 feet (-305 to +3,048 meters)
Shock (from any axis X, Y, Z)	4.0 g for maximum duration of 11 ms (half sinewave)
Vibration (from any axis X, Y, Z)	5 to 500 Hz at 0.25 g
Environmental (nonoperating)	
Temperature	-20 to 60° C, maximum gradient 20° C per hour
Relative humidity	5 to 93% noncondensing, maximum gradient 10% per hour
Effective altitude	-1,000 to +40,000 feet (-305 to +12,192 meters)
Shock (from any axis X, Y, Z)	10.0 g for maximum duration of 11 ms (half sinewave)
Vibration (from any axis X, Y, Z)	5 to 500 Hz at 1.0 g
Connectors	
Disk drives	3.5-inch, 36.4-GB or 73.4-GB 10000-rpm, bi-directional dual-ported FC-AL drives
Controller card	Native FC connectors RJ45 10/100BASE-T (100 meters max. cable) RJ45 RS-232 (25 meters max. cable)



Performance	Specifications
Native FC	LC/SSF small form factor connections
Unit interconnect cards	DB-9 (proprietary)
Power/cooling unit	CEE 22-V (latching)
Power plug	NEMA 5-15P
Standards compliance	
Safety and emissions	FCC Class A, UL/CSA, CE
Interfaces and protocols	FC-AL, SCSI, HTTP, HTML, Telnet, and FTP
Policies	Y2K, L10n, I18n (except RAID controller firmware)
Other	RS232C Ethernet 802.3

Performance

Benchmarks

The Sun StorEdge™ T3 array performance test results were run using 18.2-GB disks. The goal of the benchmark activities was to measure the Sun StorEdge T3 array performance characteristics, understand the performance, and tune performance if any parameters can be tuned. Numerous configurations were tested, including different RAID levels.

The Sun StorEdge T3 array disk trays default parameter settings provide a basis for good performance. Optimal performance for a given workload can be achieved with minor changes to these default settings.

The Sun StorEdge T3 array disk tray performs predictably, so minor variations in workload result in small changes in performance. Any parameters can be tuned.

Additional benchmarks are planned. For the latest results, refer to the following URLs:

<http://www.sun.com/storage/t3es>

<http://www.sun.com/storage/t3wg>

<http://webhome.ebay/networkstorage/products/t3/performance/index.html>

Key Findings

The Sun StorEdge T3 array exhibits the greatest — and most noticeable — performance improvement when I/O requests are greater than or equal to 64 KB, for example, in DSS (analytical) applications. With I/O-intensive applications such as online transaction processing (OLTP), the Sun StorEdge T3 array performs comparably to existing Sun StorEdge arrays.

Cache Memory and Mirroring

Enabling cache mirroring reduces performance for small, sequential write operations by about 50 percent, since only about half the cache memory is available per controller, and the controllers have the overhead of copying writes synchronously. Read and large write operations, and small, random write operations perform roughly the same as without cache mirroring enabled.

Seek Range

The seek range has a significant impact on IOPS and response time. In general, as the seek range decreases, performance increases, which is true for all network storage systems.

Predictable Performance

The Sun StorEdge T3 array performs consistently, so minor variations in workload result in small changes in performance. As a result of this predictability, customers can extrapolate what kind of performance to expect from the Sun StorEdge T3 array disk tray in different workload environments, based on the data available in this report.

The equipment and configurations used in the Sun StorEdge T3 array performance benchmark are shown below.



Block Size

Block size, which is also known as stripe unit size, is the size of the data unit being "striped" across the disks. Block size affects performance as expected. For applications with a large number of sequential I/O requests, such as decision support systems (DSS) and high performance computing (HPC), a block size of 64 KB is beneficial. For latency-sensitive applications such as OLTP applications, using a 16-KB block size is best. Other applications that are between the two extremes may perform best with the 32-KB block size. The block size setting for the Sun StorEdge T3 array also determines internal settings in the controller (such as number of parallel XOR operations the controller can perform) which affect the self-tuning behavior of the array. This is the single most important "tuning knob" for the array.

Reconstruction Time

Degraded mode occurs when a disk drive fails with no hot-spare available in the configuration. The reconstruction mode occurs during one of the following scenarios:

- When the Sun StorEdge T3 unit is reconstructing all of the failed drive information to the hot-spare drive from parity or mirror drive information
- When a failed disk drive is replaced with new disk drive

In degraded mode, a Sun StorEdge T3 array with a logical volume containing a failed disk drive showed a significant decrease in performance. However, the second logical volume in the partner group operated normally without any performance changes. The reconstruction time varied, depending on how the "reconstruction rate" parameter is set and the number of I/O requests from the host. The higher the reconstruction rate parameter is set, the faster the reconstruction rate and the slower the I/O to the host.

Adaptive Cache Optimizations

For small-block random writes (OLTP), the minimum write size from host into the cache is a segment (1/8 of a Sun StorEdge T3 array units block). If a host write is smaller than one segment, then the entire segment must be read from disk and modified by the write. The block is held in cache as long as possible to allow subsequent random writes of additional segments in the same block to occur. When the entire block is filled, it can be written as a single atomic write to disk, thus consolidating eight host writes into a single disk write.

Even when less than an entire block must be written, if the segments are contiguous, they can be written as a single atomic write to disk, without having to read the remainder of the block from disk into cache. If, for some reason, one or more non-contiguous segments in a block must be written (for example, because the 80-percent write threshold was reached), then a read/modify/write sequence of the entire block must be performed.

The algorithm of segment write with block writes from host to disk is especially crucial to RAID 5 partial stripe write performance. In addition, it benefits RAID 1 (1 +0) write performance.

• Large-block sequential writes

For large-block sequential writes, the ideal host I/O size for large-block sequential writes is equal to either the Sun StorEdge T3 array unit's block size or to an integer multiple of Sun StorEdge T3 array unit's block size. This allows full block atomic writes to be performed. If RAID 5 is being used, this further allows for parity to be calculated in atomic units of segments.

Another feature of Sun StorEdge T3 array unit's adaptive cache is that even with write-behind mode in effect, large-block sequential writes are treated as write-through data. There is little advantage to holding large-block sequential writes in cache because they are unlikely to be read again soon as cache hits.



Furthermore, large sequential I/O tends to saturate cache, eventually resulting in the same effect as write-through mode, but meanwhile monopolizing the cache at the expense of other random I/O which might also be occurring. So when the Sun StorEdge T3 array detects large-block sequential I/O, it writes to disk before sending an acknowledgment to the host, thus freeing up the same block for the next sequential write, and keeping the remainder of the cache available for random writes.

- **Small-block sequential reads**

For small-block sequential reads, the cache read-ahead parameter is configurable. If more than two host I/O blocks (*Note: Not Sun StorEdge T3 array blocks*) are read consecutively, then the entire array block which holds those I/O blocks is read into cache. The default setting is "on," indicating read-ahead mode is enabled. The read-ahead parameter may be viewed and set from any administrative CLI or GUI tool which has write access to the Sun StorEdge T3 array.

Examples of small-block sequential reads are:

- Example 1: Sun StorEdge T3 array with a 64-KB block size, 8-KB host I/O block size, read-ahead enabled. Two consecutive 8-KB blocks are read by the host, causing two 8-KB array segments of a 64-KB block to be read into cache. Because read-ahead is enabled, the Sun StorEdge T3 array reads the remainder of its 64-KB block (that is, six more 8-KB segments) into cache.
- Example 2: Sun StorEdge T3 array with a 64-KB block size, 2-KB host I/O block size, read-ahead enabled. Two consecutive 2-KB blocks are read by the host, causing one 8-KB T3 segment, four host I/O blocks, to be read into cache. Because read-ahead is enabled, the array reads the remainder of its 64-KB block (that is, seven more 8-KB segments, 28 more host I/O blocks) into cache. Thus, a total of 32 host I/O blocks is read into cache: the two requested blocks plus 30 more.
- Example 3: Sun StorEdge T3 array with a 64-KB block size, 64-KB host I/O block size, read-ahead enabled. Two consecutive 64-KB blocks are read by the host, causing two entire 64-KB blocks to be read into cache. Even though read-ahead is enabled, the array does not read any additional data into cache, because there is no remaining portion of a Sun StorEdge T3 array block to read.

Note that if the host I/O block size multiplied by two is as large or larger than the Sun StorEdge T3 array block size, then the read-ahead parameter has no effect, and read ahead never occurs, even if the read-ahead parameter is enabled.

Tuning

Unlike most enterprise-class RAID arrays on the market today, the Sun StorEdge T3 array avoids the myriad of tuning parameters that making tuning performance an art. There are three primary factors to tune in the Sun StorEdge T3 array, and each has only one or two settings that make sense for a particular situation. These parameters are RAID level, stripe unit size, and reconstruction rate.

RAID Level

The Sun StorEdge T3 array offers the choice for RAID 0, RAID 1 (1+0), or RAID 5. For almost all applications, RAID 5 is the appropriate choice. RAID 0 has great performance when transferring large block files, but no data redundancy (neither parity or mirroring). It may be appropriate for limited applications such as storing temporary or reproducible data, but is not appropriate for general-purpose use unless combined with some sort of software RAID layer to add reliability.



Stripe Unit Size

There are three possible stripe unit size (or block size) settings for the Sun StorEdge T3 array: 16, 32, or 64 KB, and should be matched to the I/O size of the application. In general, transaction-processing environments (such as most database, ISP, and EPR applications) use a small I/O size and tend to benefit from using a 16-KB stripe unit size. Applications using large I/O sizes (such as HPC, data warehousing, and imaging) tends to benefit from using a 64-KB stripe unit size. Applications using a medium I/O size (such as NFS file servers, especially in a data-intensive environment) should fare best with a stripe unit size in the middle, so a 32-KB setting is recommended. The choice of stripe unit size intuitively follows the I/O size of the application, and should be easy to select once the application environment is understood.

Reconstruction Rate (RAID 5)

The reconstruction rate is number between 1 and 16, and by default is set to 8. The setting controls how much reconstruction I/O the controller can perform between servicing host I/Os. The default setting of 8 is a good compromise between acceptable performance to the host and time to completely reconstruct the data (which is also the window of vulnerability to a catastrophic double-failure). If the reconstruction time needs to be as short as possible and severely degraded performance during that time is acceptable, the reconstruction rate should be set to 16; if host performance needs to be as high as possible and reconstruction time is not important, the reconstruction rate should be decreased closer to 1. At the extremes of one and 16, the controller is almost entirely dedicated to serving host I/O in degraded mode or reconstructing lost data onto a properly working disk. In general, use the default setting of 8 unless the application has a strong compelling reason to use another setting.

Capacity Planning

Capacity planning involves understanding the overall system demand today and in the future, and configuring sufficient resources to have a high level of confidence in meeting those demands. Using the Sun StorEdge T3 array, capacity planning is easy. Characteristics such as linear scalability, predictable performance, and simplified tuning help minimize many of the traditional capacity planning challenges.

Version (v2.0) of the Configuration Rules for Mission-Critical Storage incorporates the Sun StorEdge T3 array and provides more information on capacity planning. Using that document even a novice engineer can arrive at a well-designed and sized configuration. Contact a Sun sales representative for more information about this process.



Requirements and Configuration

System Requirements

Multi-Platform Support

Please refer to the following URLs for the latest information regarding Multi-Platform Support:

<http://www.sun.com/storage/t3es>

<http://www.sun.com/storage/t3wg>

<http://webhome.ebay/networkstorage/products/t3/compatibility/index.html>

Solaris™ Operating Environment Platform Support

The Sun StorEdge T3 array supports the following hardware platforms:

- Sun Enterprise™ 220R, 250, 420R, 450, 3X00–6X00, and 10000 servers
- Sun Fire™ 280R, 3800, 4800, 4810, and 6800 servers
- Netra™ t 1125, 1400, and 1405 servers
- Ultra™ 60 and 80 workstations

Supported Host Platforms

Refer to the following URLs for the latest data on supported host platforms for the Sun StorEdge T3 array:

<http://www.sun.com/storage/t3es>

<http://www.sun.com/storage/t3wg>

<http://webhome.ebay/networkstorage/products/t3/compatibility/index.html>

- **Sun StorEdge T3 for the enterprise (TT, RR, RK, RH, and RS models)**

Customers who wish to install the Sun StorEdge T3 array for the enterprise under the Microsoft Windows NT, Linux, HP-UX (with or without HP MC/Serviceguard HA clustering), or IBM AIX (with or without IBM HACMP HA clustering) operating system running on a supported host server order any Sun StorEdge T3 array for the enterprise TT, RR, RK, or RH model. The Sun-supplied failover driver needs to be downloaded (please refer to the matrix above, this document, for the supported host server/HBA/ driver versions). Use the native system commands to create volumes on Sun StorEdge T3 array LUNs.

- **Sun StorEdge T3 for the workgroup (TT and RR models)**

Customers who wish to install the Sun StorEdge T3 array for the workgroup under the Microsoft Windows NT, Linux, HP-UX, or IBM AIX operating system running on a supported host server simply order any Sun StorEdge T3 array for the workgroup TT or RR model. This array works with the native OS driver for that platform (please refer to the matrix above, this document, for the supported host server/HBA/driver versions). Use the native system commands to create volumes on Sun StorEdge T3 array LUN(s). Nothing else is required. Solaris Operating Environment support has not changed for these models. For HA clustering with these independent controller units, host-based mirroring is required.



The URL for linking to each download site for each Sun-supplied non-Solaris Operating Environment failover driver is:

http://www.sun.com/storage/t3es/multi_platform.html

Then, simply go to the bottom of the page and click on the appropriate download driver link.

(HINT: Since the page is a long one, simply click once anywhere on the page and then use the keyboard combination [Ctrl][End] to quickly go to the bottom where the links are. Users can also hit the [PgDn] key five times in quick succession.)

Storage Network Interconnect

A maximum 50/125 micron multimode fiber-optic cable length of 500 meters is supported.

Maximum Number of Sun StorEdge T3 Arrays Support Matrix

Refer to the following URL for the latest support data:

<http://www.sun.com/storage/t3es>

<http://www.sun.com/storage/t3wg>

<http://webhome.ebay.com/networkstorage/products/t3/compatibility/index.html>

Solaris Operating Environment-Based Software Compatibility

Sun has extensively tested the Sun StorEdge T3 arrays for compatibility with Sun server, storage, and storage software product lines. This includes the Sun Ultra 60 workstation, the Sun Enterprise 10000 storage environments, and all currently shipping Sun StorEdge array products (Sun StorEdge A1000/D1000 and A5X00 arrays). Numerous software packages have also been rigorously tested across these platforms, including the following:

- Sun StorEdge Instant Image 2.0, or higher
- Sun StorEdge Network Data Replicator (SNDR) 2.0, or higher
- Sun Cluster 3.0
- Solstice Backup™ 5.5.1, or higher
- Solstice DiskSuite™ 4.2 and 4.2.1, or higher
- Sun Enterprise Server Alternate Pathing (AP) 2.3.1, or higher
- Sun StorEdge Data Management Center 3.0, or higher
- Sun StorEdge Component Manager (CM) 2.2
- StorTools™ 3.3 diagnostics, or higher
- VERITAS NetBackup (NBU) 3.4, or higher
- VERITAS Volume Manager with DMP (VxVM) 3.0.4, 3.1 or 3.2
- VERITAS File System (VxFS) 3.3.3, 3.4, or higher
- VERITAS Cluster Server 1.3, or higher

System Configuration

The Sun StorEdge T3 array is offered in two configurations, each providing different levels of availability, software, and services.



- The Sun StorEdge T3 array for the enterprise (partner pair) offers redundant RAID controllers for high-availability enterprise and data center environments. The enterprise models have fully redundant, hot-swappable hardware components, and capacities ranging from 660 GB up to 5.2 TB per rack cabinet. Sun StorEdge T3 array for the enterprise models are available in tabletop, rack-ready, or rack-installed configurations with two, or eight arrays, with or without dual 7-port hubs or dual 8-port switches:
 - Dual high-availability hardware RAID controller for RAID 0, 1 (1+0), or 5
 - Battery-backed 1-GB read/write mirrored cache
 - Dual dedicated hardware RAID 5 XOR parity engines
 - Dual hot-swap redundant power/cooling modules
 - Non-floating hot-spare disk drive capability
 - Dual host data paths
 - 2-year enhanced warranty, 24 x 7 phone support and same business day hardware repair response
 - Sun Remote Services (SRS) Event Monitoring and Management service ready
 - Sun StorEdge Component Manager 2.2 software for centralized array monitoring, configuration and control
- The Sun StorEdge T3 array for the workgroup with a single RAID controller offers a reliable and cost-effective storage system with 327-GB or 660-GB tray capacities. T3WG models are available in tabletop or rack-ready versions. Features include:
 - Hardware RAID controller for RAID 0, 1 (1+0), or 5
 - Battery-backed 1 GB read/write cache
 - Dedicated hardware RAID 5 XOR parity engine
 - Reliable single RAID controller design
 - Dual hot-swap redundant power/cooling modules
 - Non-floating hot-spare disk drive capability
 - Sun Remote Services (SRS) Event Monitoring and Management service ready
 - Sun StorEdge Component Manager 2.2 for centralized array monitoring, configuration, and control
 - 2-year standard warranty

Front Components

The front of the Sun StorEdge T3 unit provides user access to the units nine Fibre Channel RAID disk drives. Each disk drive has an 36.4-GB or 73.4-GB capacity, its own dual Fibre Channel interface and circuitry, and front-panel LEDs for drive activity (green) and status (amber).

To open the front cover (or bezel), the user depresses the latches on each side of the front panel and pulls forward to completely remove the bezel. The units nine disk drives are positioned sideways with their activity and status LEDs at top.






Figure 5. Front view of the Sun StorEdge T3 controller unit

Rear Components

The rear of the Sun StorEdge T3 controller unit provides user access to redundant Fibre Channel unit interconnect cards (lower left), one RAID controller card (lower right), and redundant power/cooling units (upper left and right):

- Each of the two unit interconnect cards includes interface circuitry and two Fibre Channel connectors for interconnecting units.
- The RAID controller card includes RAID controller hardware and firmware, one host Fibre Channel interface, a 10BASE-T/100BASE-T Ethernet host interface, and an RS-232 COM service port.
- Each of the two power/cooling unit contains a power supply, two cooling fans, an integrated UPS battery, and status indicators for AC LED (green/amber) and POWER OK LED (green).




Figure 6. Rear view of the Sun StorEdge T3 controller unit

Interconnect




Figure 7. Sun StorEdge T3 interconnect

The Sun StorEdge T3 array uses two unit interconnect cables for connecting units together. The unit interconnect cables use a proprietary connector. In addition to FC-AL signals, these cables also carry a serial bus that functions as the nervous system of the Sun StorEdge T3 array. The serial bus carries FRU state information to the RAID controller card that then processes the information. This data is used to create `syslog` entries.

The unit interconnect cables may be replaced without taking the partner group off-line. One cable **MUST** be in place at all times for the partner group to remain functional.


Native LC Fiber Channel Connectors supplied with each controller unit, is used to attach a multimode fiber cable. This cable is then attached to a supported FC-AL HBA.

Configuration Options


Refer to the diagrams in this section and to the product descriptions in the Ordering section for more details.

- There are three configuration options: tabletop, rack-ready, and rack-installed.
- There are two drive capacities (36.4 GB and 73.4 GB).


Configuration Diagrams




2 controllers, direct host




8 controllers, direct host




2 controllers, 2 hubs/switches, 1 host



8 controllers, 2 hubs/switches, 1 host



4 controllers, 2 hubs/switches, 2 hosts



8 controllers, 2 hubs/switches, 2 hosts

System Administration

System Maintenance

The modular design of the Sun StorEdge™ T3 array makes it easy to replace failed components. Field replaceable units (FRUs) include the following:

- RAID controller card* (one per controller tray)
- Nine 10000-rpm 36.4-GB or 73.4-GB bi-directional dual-ported FC-AL disk drives*
- Two FC-AL unit interconnect cards*
- Two power/cooling units* (including UPS batteries)
- Empty chassis (with center-plane installed)

*All major FRUs can be hot-swapped to ease system maintenance and help prevent system downtime.

The controller units nine disk drives can be accessed simply by opening the front cover of the enclosure. The rear of the Sun StorEdge T3 unit provides access to the other FRUs. See the "Requirements and Configuration" section for front and rear diagrams and additional information.

Localization and Internationalization

The Sun StorEdge T3 array is fully compliant with localization and internationalization policies, with exception of the controller firmware and `syslog` entries.



Software Administration

This section describes the following aspects of software administration for the Sun StorEdge™ T3 array:

- Compatible Sun™ and third-party software products
- Supported host operating environments
- StorTools™ system diagnostic software
- Storage management software
 - Optional VERITAS Volume Manager software
 - Sun StorEdge Component Manager software

Compatible Software

- **Sun software**
 - Sun StorEdge Instant Image 2.0, or higher
 - Sun StorEdge Network Data Replicator (SNDR) 2.0, or higher
 - Sun StorEdge Data Management Center 3.0, or higher
 - Solstice Backup™ 5.5.1, or higher
 - Solstice DiskSuite™ 4.2 and 4.2.1, or higher
 - Sun Enterprise Server Alternate Pathing (AP) 2.3.1, or higher
 - Sun StorEdge Component Manager (CM) 2.2
 - StorTools 3.3 diagnostics, or higher
- **Third-party software**
 - VERITAS NetBackup (NBU) 3.2, 3.4, or higher
 - VERITAS Volume Manager with DMP (VxVM) 3.0.4, 3.1, 3.1.1 or 3.2
 - VERITAS File System (VxFS) 3.3.3, 3.4, or higher

Host Operating Environments

The Sun StorEdge T3 array supports the following host operating environments:

- Solaris™ 2.6, Solaris 7, and Solaris 8 Operating Environment
 - Solaris Operating Environment PatchPro patches are available on the SunSolve™ program web site at <http://sunsolve.sun.com>.
- Microsoft Windows NT 4.0 (SP6) Operating Environment (requires version 1.14, or higher, Sun StorEdge T3 array firmware)
- HP-UX (11.0) Operating Environment (requires version 2.0, or higher, Sun StorEdge T3 array firmware)
- IBM AIX (4.3.3) Operating Environment (requires version 2.0, or higher, Sun StorEdge T3 array firmware)



- Linux kernel (2.2.15—GNU/Debian, Caldera Open Linux, E-Server v2.3, Red Hat v6.2, SuSE v6.4, Turbo Linux Server v6.0) Operating Environments (requires version 2.0, or higher, Sun StorEdge T3 array firmware)

StorTools Diagnostic Software

The Sun StorEdge T3 array supports the StorTools 3.3, or higher, utility for diagnosing host-loop integrity. The RAID controller in the Sun StorEdge T3 array manages disk-loop integrity.

VERITAS Volume Manager Software

Optional VERITAS Volume Manager software supports and adds functionality to the Sun StorEdge T3 array RAID controller hardware.

Features

- Host-based RAID levels 0, 1 (1+0) and 5
- Dynamic multipathing
- Hot relocation
- Disk groups
- Volume resizing
- File system expansion
- On-line snapshots
- Graphical user interface
- Performance analysis tool

Benefits

- Helps improve performance and data availability
- Balances performance, adds reliability, and provides path failover
- Helps increase data availability
- Facilitates movement of data between hosts
- Helps enable volumes to change as needs dictate
- Helps enable file systems to grow dynamically
- Facilitates on-line backups
- Eases administration
- Helps enable problem (bottleneck) isolation and tuning

Sun StorEdge Component Manager Software

Sun StorEdge Component Manager software can be ordered with the Sun StorEdge T3 array as a no charge item. Sun StorEdge Component Manager software is a Jiro™ open storage management platform-compliant, Java™ technology-based application for managing attached Sun StorEdge arrays regardless of host platform connection and their hardware components. Administrators and service personnel who install, monitor, and maintain supported Sun StorEdge arrays can use Sun StorEdge Component Manager software to perform a variety of subsystem tasks:

- Monitor enclosures by using hardware polling, alarm notification, event logging, and remote reporting for abnormal conditions and activities.
- Display the status of enclosures, enclosure components, and their associated properties.
- Exercise control directives on some enclosure components, for example, powering off a disk.

Sun StorEdge Component Manager software provides out-of-band management through the Ethernet connection on the RAID controller of the Sun StorEdge T3 controller unit array.

Note: *Sun StorEdge Component Manager software does not manage the logical organization of stored data.*



Sun StorEdge Component Manager 2.2 software runs on Sun Fire™ 280R, 3800, 4800, 4810, and 6800 servers, Sun Enterprise™ 250, 220R, 450, 420R, 3X00, 4X00, 5X00, 6X00, and 10000 servers, and also on Sun Blade™ 1000 and Sun Ultra™ 5, 60, and 80 workstations. Sun StorEdge Component Manager 2.2 software runs under the Solaris 2.6, 7, or 8 Operating Environment.

Use the Sun StorEdge Component Manager software when a graphical representation of array and enclosure attributes is preferred over selected command-line actions. Sun StorEdge Component Manager software may not be suitable, on server platforms with small memory configurations. On such platforms, Sun StorEdge Component Manager software's memory consumption of 70 MB and up for multiple arrays may outweigh the advantages of its graphical services.

Market Value Proposition

Sun StorEdge Component Manager software is GUI-based software for managing the physical attributes of one or more attached Sun StorEdge enclosures. Sun StorEdge Component Manager software helps enhance the reliability, availability, and serviceability of storage assets under its custody through the following:

- Enclosure discovery
- Intuitive health displays
- Alarm generation
- Log monitoring
- Hardware control

Product Positioning

Sun StorEdge Component Manager software is positioned as an operating system-independent Java technology-based storage management software plug-in that is accessible from the Sun StorEdge Management Console software. Sun StorEdge Component Manager 2.2 software is fundamentally a tool for controlling and monitoring the physical attributes of the Sun StorEdge T3 arrays. Sun StorEdge Component Manager software is *not* intended to address the logical organization of the data stored in those arrays.

In the larger picture, Sun StorEdge Component Manager software takes on the role of configuration, control, monitoring, and the diagnostic entry point for physical components. This is planned as an evolving effort over the next few years, and is intended to encompass new Sun StorEdge arrays, switches, and tape subsystems.

Sun StorEdge Management Console Software

Sun StorEdge Management Console software implements elements of the Jiro open storage management platform through a collection of fundamental Java technology-based storage management functions. Sun StorEdge Component Manager software and other Java technology-based Sun StorEdge management services rely on the Sun StorEdge Management Console software for these functions:

- GUI navigation
- Common event notification
- Common logging
- Alarm generation
- Remote notification (phone home)



- Distributed managed object discovery and communications
- Persistent state

The Sun StorEdge Management Console software launcher window presents a GUI for navigating through Sun StorEdge Component Manager software features. The main window includes a toolbar that the administrator can use to launch a new Sun StorEdge Management Console window, the Alarm Viewer, the Log Viewer, or online help. This window also includes management application tabs for Health, Configuration, and Control, plus alarm status buttons.

Sun StorEdge Component Manager Software GUI

Sun StorEdge Component Manager software's graphical user interface (GUI) and underlying management services run as plug-ins under the Sun StorEdge Management Console. Sun StorEdge Component Manager software's look-and-feel is consistent with other Sun StorEdge management services that are under development.

Sun StorEdge Component Manager software allows administrators and service personnel to monitor, display, and control devices and activities within supported Sun StorEdge T3 arrays. Descriptions for Sun StorEdge Component Manager software's basic functions follow.

Function	Description
Enclosure discovery	Identify the set of storage arrays and associated enclosures managed by a host.
Enclosure monitoring — health	Monitor and log the operational status of one or more storage arrays and associated enclosures managed by a host.
Enclosure monitoring — alarm generation and viewer	Create, display, and manage alarms generated by abnormal events detected by Sun StorEdge Component Manager software. Notification can include a combination of visual alarms to the GUI, e-mail messages, or SNMP traps for remote reporting.
Enclosure management and control	Control storage array and associated enclosure components; enable and disable RAID controllers
Log viewing	View log entries by administrator. Logging provides auditing results of tasks executed by Sun StorEdge Component Manager software, or provides information about a sequence of events.
Log file monitoring	Monitor a log file continuously for particular string patterns. Once a pattern is identified, rules that are preset in the application can trigger alarms or events.

Architecture

Sun StorEdge Component Manager software is a client-agent model application. The client part is a lightweight Java technology-based application GUI. The agent part is also written in the Java programming language, and runs on a Solaris Operating Environment server (host) to which the storage enclosure is attached.

Sun StorEdge Component Manager software is comprised of three layers:

- The top layer is the GUI-client layer that accepts user commands and reports results.
- The middle layer is the application layer that contains the logic to initiate data acquisition, process information, and schedule and generate sets of actions.
- The bottom layer is the managed-objects layer that performs hardware commands to satisfy requests through the underlying Java Native Interface (JNI).



The three layers communicate through distributed object management facilities in the Sun StorEdge Management Console software.



Ordering Information

Sun StorEdge™ T3 Array Ordering

The following items are included with all Sun StorEdge™ T3 arrays:

- Sun StorEdge Component Manager 2.2 software (order part number SCMMS-220-99Y9; No charge line item)
- Sun Remote Services and Event Monitoring and Management 2.0, or higher software ready
- VERITAS Volume Manager software (promotion, no charge line item)

Part Numbering Scheme

In the table below is a description of the basic part numbering scheme. Here is a breakdown of the marketing names:

- T3 = The Sun StorEdge T3 array
- B = The second generation of this array, with 1-GB cache controller
- WG = Sun StorEdge T3 array for the workgroup model
- ES = Sun StorEdge T3 array for the enterprise model

Marketing Name	Configuration	Arrays/Controllers	Total Capacity	Rack/Revenue Rec.
(T3BWG) T3 array for the workgroup	(TT) Tabletop	(11) 1 array, 1 controller	Given in Gigabytes (GB)	(GR4) For factory installation in 68-inch Sun Enterprise™ 5500/6500 cabinet
(T3BES) T3 array for the enterprise	(RR) Rack ready (Rack kits sold separately)	(22) 2 arrays, 2 controllers		(R4) For factory installation in Sun Fire™ cabinet
	(RK) Rackmounted in Sun StorEdge 72-inch cabinet	(88) 8 arrays, 8 controllers		(R5) For factory installation in Sun StorEdge 72-inch cabinet
	(RH) Rackmounted in Sun StorEdge 72-inch cabinet			
	(RS) Sun cabinet with two (8) port switches			



Configuration Matrix

		Sun StorEdge T3 array for the workgroup	Sun StorEdge T3 array for the enterprise	
Configuration	Drive	1 Array	2 Arrays	8 Arrays
(TT) Table Top	36-GB	XT3BWG-TT-11-327	XT3BES-TT-22-655	
	73-GB	XT3BWG-TT-11-660	XTB3ES-TT-22-1321	
(RR) Customer rack-ready Requires rackmounting kit	36-GB	XT3BWG-RR-11-327	XT3BES-RR22-655	
	73-GB	XT3BWG-RR-11-660	XT3BES-RR22-1321	
(RR) Rack-ready, mountable in Sun system cabinet or 72-inch Sun StorEdge cabinet. Rack must be ordered separately. GR4 - 68-inch Sun Enterprise 5500/6500 cabinet R5 - 72-inch Sun StorEdge cabinet R4 - Sun Fire cabinet	36-GB	T3BWG-RR-11-327GR4 T3BWG-RR-11-327R5 T3BWG-RR-11-327R4	T3BES-RR22-655GR4 T3BES-RR22-655R5 T3BES-RR22-655R4	
	73-GB	T3BWG-RR-11-660GR4 T3BWG-RR-11-660R5 T3BWG-RR-11-660R4	T3BES-RR22-1321GR4 T3BES-RR22-1321R5 T3BES-RR22-1321R4	
(RK) Sun rackmounted in 72-inch Sun StorEdge cabinet	36-GB		XT3BES-RK-22-655 T3BES-RK-22-655	XT3BES-RK-88-2620 T3BES-RK-88-2620
	73-GB		XT3BES-RK-22-1321 T3BES-RK-22-1321	XT3BES-RK-88-5284 T3BES-RK-88-5284
(RH) Sun rackmounted with two 7-port hubs in 72-inch Sun StorEdge cabinet	36-GB		XT3BES-RH-22-655 T3BES-RH-22-655	XTB3ES-RH-88-2620 T3BES-RH-88-2620
	73-GB		XT3BES-RH-22-1321 T3BES-RHS-22-1321	XT3BES-RH-88-5284 T3BES-RH-88-5284
(RS) Sun rackmounted with two 8-port switches in 72-inch Sun StorEdge cabinet	36-GB		XT3BES-RS-22-655 T3BES-RS-22-655	XT3BES-RS-88-2620 T3BES-RS-88-2620
	73-GB		XT3BES-RS-22-1321 T3BES-RS-22-1321	XT3BES-RS-88-5284 T3BES-RS-88-5284
		Included with Sun StorEdge T3 array for the workgroup	Included with Sun StorEdge T3 array for the enterprise	
Order SCMMS-220-99Y9 no charge part number		Sun StorEdge Component Manager 2.2 software	Sun StorEdge Component Manager 2.2 software	
Included during Mission Critical Sales Process. Line charges are additional.		Sun Remote Services Event Monitoring & Management 2.0 Ready	Sun Remote Services Event Monitoring & Management 2.0, or higher Ready	
		Standard 2-yr. warranty	2 year 24 x 7 enhanced warranty	
			VERITAS Volume Manager (VxVM) 3.0.4/ 3.1/ 3.1.1/ 3.2 No charge part number (One Tier 2 Server License) VxVM 3.0.4 Order VVMT3-304-9999 VxVM 3.1.0 Order VVMT3-310-9999 VxVM 3.1.1 Order VVMT3-311-9999 VxVM 3.2 Order VVMT3-320-9999	

NOTES:

1. Bold part numbers are the "Standard Configurations" with shorter lead times.
(for example, **XT3BES-RK-88-2620**)
2. All part numbers beginning with "X" are factory assemble to order (ATO).
(for example, XT3BES-RS-88-5384)



Shipping Configurations

- **Sun StorEdge T3 array for the enterprise (T3BES models)**

Includes two or eight array configurations with each tray containing one RAID controller, nine 10000-rpm 36.4-GB, or 73.4-GB FC-AL disk drives, one multimode fiber-optic cable, two AC power-line cords, and Installation Guide.

- Preconfigured for RAID 5 (8+1 no hot-spare)
- Rackmounted in 72-inch cabinet (two or eight array systems)
- Tabletop or rackmount ready (one or two array systems)
- Sun StorEdge Component Manager 2.2 software media, documentation
- VERITAS Volume Manager (VxVM) 3.0.4, 3.1, 3.1.1 or 3.2 (Tier 2 departmental server license)
- 2-year enhanced warranty
- Sun StorEdge ArrayStart™ installation and configuration (optional)
- Sun Remote Services (SRS) 2.0 ready

- **Sun StorEdge T3 array for the workgroup (T3BWG models)**

Includes one tray containing one RAID controller, nine 10000-rpm 36.4-GB or 73.4-GB FC-AL disk drives, one multimode fiber-optic cable, two AC power-line cords, and Installation Guide.

- Preconfigured for RAID 5 (8+1 no hot-spare)
- Rackmount-ready or tabletop versions offered
- Sun StorEdge Component Manager 2.2 software license, media, and documentation
- 2-year standard warranty
- Sun Remote Services (SRS) 2.0 ready



Sun StorEdge T3 Array for the Workgroup, 1-Array Configurations

Order Number	Title and Description
XT3BWG-TT-11-327 (<i>X = Field Install</i>)	327-GB Sun StorEdge T3 tabletop array for the workgroup (286-GB usable RAID 5 storage) 9 x 36.4-GB, 10000-rpm, FC-AL drives, one RAID controller with 1-GB cache, preconfigured as one RAID 5 LUN (8+1), one FC-AL port, one 5-meter fiber-optic cable, two power supplies w/internal battery backup, two fans, two power cords, SRS ready
XT3BWG-TT-11-660 (<i>X = Field Install</i>)	660-GB Sun StorEdge T3 tabletop array for the workgroup 9 x 73.4-GB, 10000-rpm, FC-AL drives, one RAID controller with 1GB cache, preconfigured as one RAID 5 LUN (8+1), one FC-AL port, one 5-meter fiber-optic cable, two power supplies w/internal battery backup, two fans, two power cords, SRS ready
XT3BWG-RR-11-327 (<i>X = Field Install</i>)	327-GB Sun StorEdge T3 array for the workgroup (286-GB usable RAID 5 storage) 9 x 36.4-GB, 10000-rpm, FC-AL drives, one RAID controller with 1-GB cache, preconfigured as one RAID 5 LUN (8+1), one FC-AL port, one 5-meter fiber-optic cable, two power supplies w/internal battery backup, two fans, two power cords, requires separately ordered rackmount rail kit (X9673, X9674), SRS ready See Ordering Notes 1, 2, and 3 below.
XT3BWG-RR-11-660 (<i>X = Field Install</i>)	660-GB Sun StorEdge T3 array for the workgroup rack-ready 9 x 73.4-GB, 10000-rpm, FC-AL drives, one RAID controller with 1-GB cache, preconfigured as one RAID 5 LUN (8+1), one FC-AL port, one 5-meter fiber-optic cable, two power supplies w/internal battery backup, two fans, two power cords, requires separately ordered rackmount rail kit (X9673, X9674), SRS ready See Ordering Notes 1 and 3 below.
T3BWG-RR-11-327GR4 (<i>No X = Factory Install</i>)	327-GB Sun StorEdge T3 array for the workgroup (9 x 36.4-GB, 10000-rpm FC-AL drives); rackmounted in 68-inch Sun Enterprise 5500/6500 cabinet
T3BWG-RR-11-327R5 (<i>No X = Factory Install</i>)	327-GB Sun StorEdge T3 array for the workgroup (9 x 36.4-GB, 10000-rpm FC-AL drives); rackmounted in R5 72-inch Sun StorEdge cabinet
T3BWG-RR-11-327R4 (<i>No X = Factory Install</i>)	327-GB Sun StorEdge T3 array for the workgroup (9 x 36.4-GB, 10000-rpm FC-AL drives); rackmounted in Sun Fire cabinet for use with Sun Fire servers
T3BWG-RR-11-660GR4 (<i>No X = Factory Install</i>)	660-GB Sun StorEdge T3 array for the workgroup (9 x 73.4-GB, 10000-rpm FC-AL drives); rackmounted in 68-inch Sun Enterprise 5500/6500 cabinet



Order Number	Title and Description
T3BWG-RR-11-660R5 <i>(No X = Factory Install)</i>	660-GB Sun StorEdge T3 array for the workgroup (9 x 73.4-GB, 10000-rpm FC-AL drives); rackmounted in R5 72-inch Sun StorEdge cabinet
T3BWG-RR-11-660R4 <i>(No X = Factory Install)</i>	660-GB Sun StorEdge T3 array for the workgroup (9 x 73.4-GB, 10000-rpm FC-AL drives); rackmounted in Sun Fire cabinet for use with Sun Fire servers

Sun StorEdge T3 Array for the Enterprise, 2-Array Configurations

Order Number	Title and Description
XT3BES-TT-22-655 <i>(X = Field Install)</i>	655-GB Sun StorEdge T3 tabletop array for the enterprise, includes two arrays configured in one partner group, (572-GB usable RAID 5 storage) preconfigured as two RAID 5 LUNs (8+1) two RAID controllers, 1-GB mirrored cache, 18 x 36.4-GB, 10000-rpm, FC-AL drives, two 5-meter fiber-optic cables, two unit interconnect cables, two power cords, 2-year enhanced warranty included, SRS ready
XT3BES-TT-22-1321 <i>(X = Field Install)</i>	1321-GB Sun StorEdge T3 tabletop array for the enterprise, includes two arrays configured in one partner group, preconfigured as two RAID 5 LUNs (8+1) two RAID controllers, 1-GB mirrored cache, 18 x 73.4-GB, 10000-rpm, FC-AL drives, two 5-meter fiber-optic cables, two unit interconnect cables, two power cords, 2-year enhanced warranty, SRS ready
XT3BES-RR-22-655 <i>(X = Field Install)</i>	655-GB Sun StorEdge T3 array for the enterprise, includes two arrays configured in one partner group, (572-GB usable RAID 5 storage) preconfigured as two RAID 5 LUNs (8+1) two RAID controllers, 1-GB mirrored cache, 18 x 36.4-GB, 10000-rpm, FC-AL drives, two 5-meter fiber-optic cables, two unit interconnect cables, two power cords, ArrayStart installation and configuration (T3BES Model only), 2-year enhanced warranty included, SRS ready, requires separately ordered rackmount rail kit (X9673, X9674) See Ordering Notes 1, 2, and 4 below.
XT3BES-RR-22-1321 <i>(X = Field Install)</i>	1321-GB Sun StorEdge T3 array for the enterprise, includes two arrays configured in one partner group, preconfigured as two RAID 5 LUNs (8+1) two RAID controllers, 1-GB mirrored cache, 18 x 73.4-GB, 10000-rpm, FC-AL drives, two 5-meter fiber-optic cables, two unit interconnect cables, two power cords, 2-year enhanced warranty included, SRS ready, requires separately ordered rackmount rail kit (X9673, X9674) See Ordering Notes 1 and 4 below.



Order Number	Title and Description
T3BES-RR-22-655GR4 <i>(No X = Factory Install)</i>	655-GB Sun StorEdge T3 array for the enterprise (two arrays, each 9 x 36.4-GB, 10000-rpm FC-AL drives); rackmounted in 68-inch Sun Enterprise 5500/6500 cabinet
T3BES-RR22-655R5 <i>(No X = Factory Install)</i>	655-GB Sun StorEdge T3 array for the enterprise (two arrays, each 9 x 36.4-GB, 10000-rpm FC-AL drives); rackmounted in R5 72-inch Sun StorEdge cabinet
T3BES-RR22-655R4 <i>(No X = Factory Install)</i>	655-GB Sun StorEdge T3 array for the enterprise (two arrays, each 9 x 36.4-GB, 10000-rpm FC-AL drives); rackmounted in Sun Fire StorEdge cabinet for use with Sun Fire servers
T3BES-RR22-1321GR4 <i>(No X = Factory Install)</i>	1321-GB Sun StorEdge T3 array for the enterprise (two arrays, each 9 x 73.4-GB, 10000-rpm FC-AL drives); rackmounted in 68-inch Sun Enterprise 5500/6500 cabinet
T3BES-RR22-1321R5 <i>(No X = Factory Install)</i>	1321-GB Sun StorEdge T3 array for the enterprise (two arrays, each 9 x 73.4-GB, 10000-rpm FC-AL drives); rackmounted in R5 72-inch Sun StorEdge cabinet
T3BES-RR22-1321R4 <i>(No X = Factory Install)</i>	1321-GB Sun StorEdge T3 array for the enterprise (two arrays, each 9 x 73.4-GB, 10000-rpm FC-AL drives); rackmounted in Sun Fire StorEdge cabinet for use with Sun Fire servers

Sun StorEdge T3 Array for the Enterprise, 2-Array Configurations

Order Number	Title and Description
XT3BES-RK-22-655 <i>(X = Field Install)</i> T3BES-RK-22-655 <i>(No X = Factory Install)</i>	655-GB Sun StorEdge T3 array for the enterprise includes one array configured as a partner group, 572-GB usable RAID 5 storage preconfigured as four RAID 5 LUNs (8+1), 18 x 36.4-GB, 10000-rpm, FC-AL drives, four 15-meter fiber-optic cables, four unit interconnect cables, installed in one 72-inch Sun StorEdge expansion cabinet, fans and door included, 2-year enhanced warranty included, SRS ready See Ordering Note 5 below.
XT3BES-RK-22-1321 <i>(X = Field Install)</i> T3BES-RK-22-1321 <i>(No X = Factory Install)</i>	1321-GB Sun StorEdge T3 array for the enterprise includes one array configured as a partner group, 1.1-TB usable RAID 5 storage preconfigured as four RAID 5 LUNs (8+1), 18 x 73.4-GB, 10000-rpm, FC-AL drives, four 15-meter fiber-optic cables, four unit interconnect cables, installed in one 72-inch Sun StorEdge expansion cabinet, fans and door included, 2-year enhanced warranty included, SRS ready See Ordering Note 5 below.



Order Number	Title and Description
XT3BES-RH-22-655 <i>(X = Field Install)</i> T3BES-RH-22-655 <i>(No X = Factory Install)</i>	655-GB Sun StorEdge T3 array for the enterprise includes one array configured as a partner group, 572-GB usable RAID 5 storage preconfigured as four RAID 5 LUNs (8+1), 18 x 36.4-GB, 10000-rpm, FC-AL drives, four 15-meter fiber-optic cables, four unit interconnect cables, installed in one 72-inch Sun StorEdge expansion cabinet (fans and door included), dual 7-port FC hubs with five GBICs each, enhanced warranty, SRS ready See Ordering Note 5 below.
XT3BES-RH-22-1321 <i>(X = Field Install)</i> T3BES-RH-22-1321 <i>(No X = Factory Install)</i>	1321-GB Sun StorEdge T3 array for the enterprise includes one array configured as a partner group, 1.1-TB usable RAID 5 storage preconfigured as four RAID 5 LUNs (8+1), 18 x 73.4-GB, 10000-rpm, FC-AL drives, four 15-meter fiber-optic cables, four unit interconnect cables, installed in one 72-inch Sun StorEdge expansion cabinet (fans and door included), dual-port FC hubs with five GBICs each, 2-year enhanced warranty included, SRS ready See Ordering Note 5 below.
XT3BES-RS-22-655 <i>(X = Field Install)</i> T3BES-RS-22-655 <i>(No X = Factory Install)</i>	655-GB Sun StorEdge T3 array for the enterprise, includes one array configured as a partner group, pre-configured as two RAID 5 LUNs (8+1), 18 x 36.4-GB 10000-rpm dual-ported bi-directional FC-AL drives, two 15-meter fiber-optic cables, two unit interconnect cables, installed in a 72-inch Sun StorEdge expansion cabinet (fans and door included), dual redundant 8-port FC switches with five GBICs each, 2-year enhanced warranty included, SRS Ready
XT3BES-RS-22-1321 <i>(X = Field Install)</i> T3BES-RS-22-1321 <i>(No X = Factory Install)</i>	1321-GB Sun StorEdge T3 array for the enterprise, includes one array configured as a partner group, pre-configured as two RAID 5 LUNs (8+1), 18 x 73.4-GB 10000-rpm dual-ported bi-directional FC-AL drives, two 15-meter fiber-optic cables, two unit interconnect cables, installed in a 72-inch Sun StorEdge expansion cabinet (fans and door included), dual redundant 8-port FC switches with five GBICs each, 2-year enhanced warranty included, SRS Ready



Sun StorEdge T3 Array for the Enterprise, 8-Array Configurations

Order Number	Title and Description
XT3BES-RK-88-2620 <i>(X = Field Install)</i> T3BES-RK-88-2620 <i>(No X = Factory Install)</i>	2620-GB StorEdge T3 array for the enterprise, includes eight arrays configured in four partner groups, 2312-GB usable RAID 5 storage preconfigured as eight RAID 5 LUNs (8+1), 72 x 36.4-GB 10000-rpm FC-AL drives, eight 15-meter fiber-optic cables, eight unit interconnect cables, installed in one 72-inch Sun StorEdge expansion cabinet, fans and door included, ArrayStart installation and configuration (T3BES models only), 2-year enhanced warranty included, SRS ready See Ordering Note 5 below.
XT3BES-RK-88-5284 <i>(X = Field Install)</i> T3BES-RK-88-5284 <i>(No X = Factory Install)</i>	5284-GB StorEdge T3 array for the enterprise, includes eight arrays configured in four partner groups, preconfigured as eight RAID 5 LUNs (8+1), 72 x 73.4-GB 10000-rpm FC-AL drives, eight 15-meter fiber-optic cables, eight unit interconnect cables, installed in one 72-inch Sun StorEdge expansion cabinet, fans and door included, ArrayStart installation and configuration (T3BES models only), 2-year enhanced warranty included, SRS ready See Ordering Note 3 below.
XT3BES-RH-88-2620 <i>(X = Field Install)</i> T3BES-RH-88-2620 <i>(No X = Factory Install)</i>	2620-GB Sun StorEdge T3 array for the enterprise, includes eight arrays configured in four partner groups, 2312-GB usable RAID 5 storage preconfigured as eight RAID 5 LUNs (8+1), 72 x 36.4-GB 10000-rpm FC-AL drives, eight copper-to-optic media interface adapter, two 15-meter fiber-optic cables, eight unit interconnect cables, installed in one 72-inch Sun StorEdge expansion cabinet, fans and door included, dual 7-port hubs with five GBICs each, ArrayStart installation and configuration (T3BES models only), 2-year enhanced warranty included, SRS ready See Ordering Note 5 below.
XT3BES-RH-88-5284 <i>(X = Field Install)</i> T3BES-RH-88-5284 <i>(No X = Factory Install)</i>	5284-GB Sun StorEdge T3 array for the enterprise, includes eight arrays configured in four partner groups, preconfigured as eight RAID 5 LUNs (8+1), 72 x 73.4-GB 10000-rpm FC-AL drives, two 15-meter fiber-optic cables, eight unit interconnect cables, installed in one 72-inch Sun StorEdge expansion cabinet, fans and door included, dual 7-port hubs with five GBICs each, ArrayStart installation and configuration (T3BES models only), 2-year enhanced warranty included, SRS ready See Ordering Note 3 below.
XT3BES-RS-88-2620 <i>(X = Field Install)</i> T3BES-RS-88-2620 <i>(No X = Factory Install)</i>	2620-GB Sun StorEdge T3 array for the enterprise, includes four arrays configured in four partner groups, pre-configured as eight RAID 5 LUNs (8+1), 72 x 36.4-GB 10000-rpm dual-ported bi-directional FC-AL drives, two 15-meter fiber-optic cables, eight unit interconnect cables, installed in a 72-inch Sun StorEdge expansion cabinet (fans and door included), dual redundant 8-port FC switches with five GBICs each, 2-year enhanced warranty included, SRS Ready



Ordering Notes

1. Delete "X" and append "R5" to the order number (for example, T3BWG-RR-11-327R5) for factory installation in an 72-inch Sun StorEdge expansion cabinet rack or a Sun Enterprise 4502-R cabinet. Rackmounting rails are included with the non X-option part number.
2. Delete "X" and append "GR4" to the order number (for example, T3BWG-RR-11-327GR4) for factory installation in a Sun Enterprise 5500 or 6500 system cabinet. Rackmounting rails are included with the non X-option part number.
3. When ordering the X-option part number (for example, XT3BWG-RR-11-327) include the appropriate 4U single rackmount rail kit, X9673A OR X9674A.
4. When ordering the X-option part number (for example, XT3BES-RR-22-660) include the appropriate 7U dual rackmount rail kit, X9663A OR X9664A.
5. Delete "X" (for example, T3BES-RH-88-2620) for factory configuration with the Sun Enterprise 10000 cabinet.
6. Delete "X" for factory configured options (for example, 9629A).

Options

Part Number	Description	Category
X9647A	Upgrade from Sun StorEdge T3 to Sun StorEdge T3 array with 1-GB cache hardware RAID controller. Includes 0.4-meter LC FC-AL cable, SC-SC coupler, and documentation	Upgrade Kit to Sun StorEdge T3 array with 1-GB cache
X9720A	SC to SC Fibre Channel cable coupler with female to female SC connections	
X9720A	SC to SC Fibre Channel cable coupler with female to female SC connections	
X9721A	0.4-meter LC to SC Fibre Channel optic cable.	LC/SFF (small form factor) to SC connectors FC-AL cables and accessories for Sun StorEdge T3 array with 1-GB cache
X9722A	2-meter LC to SC Fibre Channel optic cable	
X9723A	5-meter LC to SC Fibre Channel optic cable	
X9724A	15-meter LC to SC Fibre Channel optic cable	
X973A	2-meter SC to SC Fibre Channel optic cable	SC connectors FC-AL cables for connection between HBAs to hubs/switches
X9715A	5-meter SC to SC Fibre Channel optic cable	
X978A	15-meter SC to SC Fibre Channel optic cable	
X9673A	4U single rackmount rail kit for 72-inch Sun StorEdge expansion cabinet or RETMA rack	Rack rail kits for the Sun StorEdge T3 array for the workgroup
X9674A	4U single rackmount rail kit for Sun Enterprise 5500/6500 cabinet	
X9677A	4U single rackmount universal rail kit for mounting T3WG single brick models into Sun expansion cabinets (72-inch Sun StorEdge, Sun Fire, Sun Enterprise 5500/6500, or RETMA cabinets)	



Part Number	Description	Category
X9663A	7U dual rackmount rail kit for 72-inch Sun StorEdge expansion or RETMA rack	Rack Rail Kits for the Sun StorEdge T3 array for the enterprise
X9664A	7U dual rackmount rail kit for Sun Enterprise 5500 /6500 cabinet	
X9667A	7U dual rackmount universal rail kit for mounting Sun StorEdge T3ES and T3AES partner pair models into Sun expansion cabinets (72-inch Sun StorEdge, Sun Fire, Sun Enterprise 5500/6500, and RETMA cabinets)	
X9665A	7U dual rackmount kit for 2- or 4-post Telco rack	Rack accessories
X9676A	Sun StorEdge T3 mounting tray kit (tabletop to rackmount upgrade); one required for each tray	
(X)9629A	1 rack unit (RU) air baffle for Sun StorEdge T3 array inside Sun Enterprise 5500 and 6500 system cabinets (one required in each system cabinet)	
SG-ARY040A	Sun StorEdge 72-inch cabinet with front door, fan tray and documentation for StorEdge T3 Array.	Rack
(X)6731A	FC-AL 100 MB/sec. short-wave GBIC module	GBIC module
(X)6729A	FC-100 PCI FC-AL host bus adapter single channel	Host bus adapters
(X)6730A	SOC+ SBus FC-AL host bus adapter dual channel	
(X)6799A	FC-100 PCI FC-AL host bus adapter single channel (required for use with switches)	
(X)6748A	FC-100 (compact) cPCI FC-AL host bus adapter dual channel (required for use with switches)	
(X)6732A	FC-AL 100 MB/sec. 7-slot hub with no GBIC	Hubs and accessories
(X)6735A	Dual hub rackmount kit for Sun StorEdge expansion rack	
X6746A	Redundant Fibre Channel 8-port switch pair includes two 8-port Fibre Channel switches, documentation (Release Notes, Installation Guide, Sun Switch Configuration Guide, Sun Switch Trouble Shooting Guide and Admin GUI), a CD containing Fibre Channel switch administration GUI, and 19-inch rackmount hardware for both the switches and rack. Warranty: 2-year SunSpectrum Gold warranty included.	Switches
SG-XSW16-32P	Redundant Fibre Channel 16-port switch pair includes two 16-port Fibre Channel switches, documentation (Release Notes, Installation Guide, Sun Switch Configuration Guide, Sun Switch Trouble Shooting Guide and Admin GUI), a CD containing Fibre channel switch Administration GUI, and 19-inch rackmount hardware for both the switches and rack. Warranty: 2-year SunSpectrum Gold warranty included.	
X6714A	FC-AL 36.4-GB, 10000-rpm disk drive in mounting tray (spares use only)	Drive spares for the Sun StorEdge T3 array
X6713A	FC-AL 73.4-GB, 10000-rpm disk drive in mounting tray (spares use only)	



Services Available

The Sun StorEdge T3 array models have optional enhanced services package, which allows customers to rapidly implement complex, reliable, and monitor complex data storage environments. This extended service package helps ensure the use of sound storage installation and configuration practices, thereby allowing Sun to put the support infrastructure in place that is required to maintain the most demanding enterprise and data center environments.

Recommended services available for the Sun StorEdge T3 array for the workgroup and enterprise systems are shown in the table below. These services are available at an additional cost.

Service	Part Number/Availability
Sun StorEdge ArrayStart installation service Base Charge.	ARRAY-INS-BAS2
Sun StorEdge ArrayStart installation service Per-Tray Charge.	ARRAY-INS-PER-TRAY
Sun StorEdge Remote Services (SRS) 2.0.1, or higher	Line charges, GEO dependent
Sun Professional Services (SunPS SM) Performance Analysis and Capacity Planning	Available now
Sun StorEdge Upgrade Installation service <ul style="list-style-type: none">• Installation Service Base charge• Installation Service Per-Tray Charge	<ul style="list-style-type: none">• T3_UPGR_BASE• T3_UPGR_PER_TRAY

Warranty

The warranty information is described in the Service and Support section of this document.



Service and Support

The SunSpectrumSM program is an innovative and flexible service offering that allows customers to choose the level of service best suited to their needs, ranging from mission-critical support for maximum solution availability to backup assistance for self-support customers. The SunSpectrum program provides a simple pricing structure in which a single fee covers support for an entire system, including related hardware and peripherals, the SolarisTM Operating Environment software, and telephone support for SunTM software packages. The majority of Sun customers today take advantage of the SunSpectrum program, underscoring the value that it represents. Customers should check with their local Sun Enterprise Services representatives for program and feature availability in their areas.

For information specific to the Sun StorEdgeTM T3 array models, refer to:

http://www.sun.com/service/support/products/storage/t3_workgroup.html

For information specific to the Sun StorEdge T3 array for the enterprise models, refer to:

http://www.sun.com/service/support/products/storage/t3_enterprise.html

SunSpectrum program support contracts are available both during and after the warranty program. Customers may choose to uplift the service and support agreement to meet their business needs by purchasing a SunSpectrum contract.

The four levels of SunSpectrum support contracts are outlined below.

Program	Description
Mission-Critical SunSpectrum PlatinumSM Support	Designed to support client-server, mission critical solutions by focusing on failure prevention, rapid recovery and year round technical services planning. Support is provided 24 x 365.
Business-Critical SunSpectrum GoldSM Support	Includes a complete package of proactive and responsive services for customers who require maximum uptime for their strategic business-critical systems. Support is provided 24 x 365.
System Coverage SunSpectrum SilverSM Support	Combines the service expertise, responsive on-site support and technical support by telephone and SunSolve TM CD/on-line services. Support is provided 8 a.m. to 8 p.m. Mon. through Fri.
Self-Directed SunSpectrum BronzeSM Support	Provided for customers who rely primarily upon their own in-house service capabilities. Enables customers to deliver high quality service by giving them access to UNIX [®] expertise, Sun certified replacement parts, software releases and technical tools. Support is provided 8 a.m. to 5 p.m. Mon. through Fri.

Warranty Duration and Terms

For the latest warranty information for the Sun StorEdge T3 array models refer to:

<http://www.sun.com/service/support/warranty/features.html>

- **Sun StorEdge T3 array for the enterprise models**

Two years enhanced warranty:

- Same business day average response time
- On-site parts delivery
- Four business hours average response time
- On-site hardware service response
- Local business hours coverage
- 24 x 7 phone support

- **Sun StorEdge T3 array for the workgroup models**

Two years standard warranty:

- First year:
 - Second business day average response time
 - On-site parts delivery
 - 16 business hours on-site hardware service response
 - Local business hours coverage
 - Local business hours phone support
- Second year:
 - No on-site parts delivery
 - 15-day return to Sun parts exchange
 - Local business hours phone support

Education

The following SunU and SunED classes are available for this product:

- SunU course numbers 41805 and 41806
- SunED course number SM-255



Glossary

Array group	An arrangement of disk drives and the data on them which conform to a RAID level that facilitates organization and the movement of disks between systems. <i>See also</i> Disk group.
Block	An overly used term. Often used to describe the amount of data sent or received by the host per I/O operation. Also used to describe the size of an atomic read/write operation to/from a disk. In the context of the Sun StorEdge™ T3 array, represents the size of each cache buffer, and also the disk interleave factor (also known as stripe unit, chunk, interlace factor). Sun StorEdge T3 array block size can be 16, 32, or 64 KB.
Bus	A point-to-point network component. Used by Sun™ Management Center software to represent a network link to which many other hosts may be connected.
Cache hit	A read or write request for data that is already in cache. Therefore, a request can be serviced without needing to go to disk.
Channel	An interface directed toward high-speed transfer of large amounts of information.
Chunk	A quantity of information that is handled as a unit by the host and disk device.
Circuit-switched bus	A bus in which a transaction is normally implemented in an automatic fashion. Simple and easy to construct, a circuit-switched bus is often less efficient than a comparable packet-switched bus. An SBus is a circuit-switched bus.
Clean data	Any read data or write data that has been committed to disk. In other words, a copy of data that is safely on disk.
Concatenation	A volume created by sequentially mapping blocks on disks to a logical device. Two or more partitions can be concatenated and accessed as a single device.
Controller unit	The standalone controller unit is the smallest possible array configuration. The architecture integrates disks, data cache, hardware RAID, power, cooling, uninterrupted power supply (UPS), diagnostic capabilities, and administration into a versatile, standalone component. The controller unit includes external connections to a data host (or hub or switch), and to a management network.
Dirty data	Write data that is in cache and has been acknowledged to the application host, but which has not yet been committed to disk.
Disk array	A storage subsystem containing an arrangement or arrangements of multiple disk drives, designed to provide performance, high availability, serviceability, or other benefits.
Disk group	An arrangement of disk drives and the data on them which conform to a RAID level that facilitates organization and the movement of disks between systems. <i>See also</i> Array group.



Duplexed cache	Cache mirroring with duplicate data paths to and from the cache. <i>See</i> Mirrored cache.
ECC	Error Correction Code: Extra bits added to Words, or Double Words, that correct all single-bit errors, and detect all double-bit errors. A superior technology to Parity, which detects, but does not correct, single-bit errors, and cannot detect double-bit errors.
Event	A change in the state of a managed object.
Fabric	A group of interconnections between ports that includes a fabric element.
FC-AL	Fibre Channel arbitrated loop, a loop topology used with Fibre Channel.
Fiber	A wire or optical strand. Spelled <i>fib</i> re in the context of Fibre Channel.
Fiber-optic cable	Jacketed cable made from thin strands of glass, through which pulses of light transmit data. Used for high-speed transmission over medium to long distances.
Frame	An indivisible unit for transfer of information in Fibre Channel.
FRU	Field replaceable unit. A component which can be removed and replaced during service in the field.
GBIC	Gigabit interface converter.
Heterogeneous hosts	Application servers running different (disparate) operating systems which are attached to the same storage system.
Hot-plug	A hot-plug component means that it is electrically safe to remove or add that component while the machine is still running. Typically, the system must be rebooted before the hot-plug component is configured into the system.
Hot-spare	Drive(s) within a storage system held in reserve to replace any other drive when it fails. hot-spares are continuously powered up and spinning, but are not actually part of the array because they contain no data. This allows the array processor to have immediate access to a functioning drive for possible reconstruction of lost data when a disk fails.
Hot-swap	A hot-swap component can be installed or removed by simply pulling the component out and putting the new one in. The system either automatically recognizes the component change and configure itself as necessary or requires user interaction to configure the system; however, in neither case is a reboot required. All hot-swappable components are hot-pluggable, but not all hot-pluggable components are hot-swappable.
Hub	A dedicated bandwidth device for connecting fiber cables.

Interleaved memory	Helps reduce memory access time by permitting multiple memory components to operate in parallel. Memory is divided into n banks arranged so that every n th byte is supplied by a different memory bank. In a two-way interleaved system, the first double word is supplied by bank 0 while the second is supplied by bank 1. Normally, the size and extent of interleave is arranged so that a single typical request is satisfied by as many banks as possible. This arrangement permits a single memory request to be fulfilled without waiting for memory recycle time.
I/O rate	A measure of a devices capacity to transfer data to and from another device within a given time period, typically as I/O operations per second.
IOPS	Input/output operations per second. A measure of I/O performance, this is commonly used to quote random I/O performance.
IP	Internet protocol. A set of protocols developed by the United States Department of Defense to communicate between dissimilar computers across networks.
Jiro™	An open platform initiative that simplifies storage management by providing interoperability between storage vendor products.
LED	Light emitting diode.
Link	One inbound fiber and one outbound fiber connected to a port.
LRC	Loop redundancy circuit
Micron	One millionth of a meter. Also called <i>micrometer</i> .
Mirror synchronization	The process by which VERITAS Volume Manager software keeps two or more copies of data identical.
Mirrored cache	Redundant copies of data residing in cache — the (write) data residing in cache which has not yet been written to the hard disks is duplicated for failover operation. <i>See</i> Duplexed cache.
Mirroring (RAID)	Redundant storage of data, achieved by duplicating files (so, there is always a primary file and a copy of the primary file) onto separate disks. Duplicate files are either copied in whole (RAID 1).
Module	A software component that may be loaded dynamically to monitor data resources of systems, applications, and network devices.
Multimode fiber	An optical wave guide which allows more than one mode (rays of light) to be guided.
Network	An arrangement of nodes and connecting branches, or a configuration of data processing devices and software connected for information exchange.
N_Port	A port attached to a node for use with point-to-point or fabric topology.
NL_Port	A port attached to a node for use in all three topologies (point-to-point, arbitrated loop or fabric).
Node	A device that has at least one N_Port or NL_Port.



NVRAM cache	A non-volatile (battery-backed) random access memory area used as an intermediate store for data between a host computer system and disk drives to achieve faster writes and in some cases, faster reads.
Optical Fiber	Any filament of fiber, made of dielectric material, that guides light.
Packet-switched bus	A bus in which information is transmitted in fixed-sized units. This type of bus is often associated with the use of split transactions. Gigaplane™ and UPA are packet-switched buses.
Parity	In an array environment, data that is generated from user data and is used to regenerate user data lost due to a drive failure. Used in RAID 5.
Partner group (or partner pair)	Two controller units paired together to create a configuration with redundant controllers with redundant data and management paths, allowing for mirrored cache duplexing which provide controller failover, and path failover capability. The partner group is thus the minimum storage configuration for enterprise environments that call for high availability. Applies to T3ES model only.
Point-to-point	A topology where exactly two ports communicate.
Port	An access point on a device for attaching a link.
Protocol	A convention for data transmission that defines timing, control, format, and data representation.
RAID	Redundant array of independent disks. A set of disk drives which appear to be a single logical disk drive to an application such as a database or file system. Different RAID levels provide different capacity, performance, high availability, and cost characteristics.
RAID 0	RAID level 0, or striping, without parity or mirroring protection. Data is distributed evenly at the block level among disks for performance. No redundancy is provided, and the loss of a single disk causes the loss of data on all disks. Use this level for high-speed streaming of large file reads (for example, video) of non-critical data which is easily available elsewhere within the organization.
RAID 1 (1+0)	RAID level 1 (1+0), or mirroring with striping. Data is stored at the file level. Files reside on separate disks, and two copies of the data are kept. Use this level for mirroring the host operating system and/or application programs or for creating a high-traffic log volume.
RAID 5	RAID level 5, or striping with distributed parity. Both data and parity are distributed evenly across all the disks in the array at the block level. No single disk can compromise the integrity of the data. RAID 5 balances the optimization of performance, reliability and cost. Use this level for most applications which do not require the special characteristics of the above RAID levels.
RAS	Reliability, availability, and serviceability. Reliability is a measure of the likelihood that problems will occur. A highly reliable system has few problems. Once a problem occurs, availability is the measure of how the system protects the user from being adversely affected by the problem. Serviceability is a measure of how easy it is to repair the problem.



Read-ahead	Sequential data read from disk into cache without having actually been requested by the application host, in anticipation that it will be requested by the host. When the request occurs, it can be serviced as a low latency cache hit, thus improving host application performance.
Receiver	The circuitry that receives signals on a fiber, and the ultimate destination of data transmission.
Reconstruction	The process of rebuilding lost data on a replacement drive after a drive failure.
Redundancy	Duplication for the purpose of achieving fault tolerance. Refers to duplication or addition of components, data and functions within the array.
Responder	The logical function in an N_Port responsible for supporting the exchange initiated by the originator in another N_Port.
SCSI	Small computer systems interface. An ANSI standard for controlling peripheral devices by one or more host computers.
SAN	Storage area network. SAN architecture uses high-performance, high-capacity Fibre Channel switches to connect storage islands to servers. This approach provides physical connectivity, facilitating information sharing or simplify management across servers.
Segment	An overly used term; in the context of the Sun StorEdge T3 array, 1/8 of a cache buffer. In the Sun StorEdge T3 array, a segment is the smallest size of I/O possible between cache and disk. Segment size is 2, 4, or 8 KB, depending on block size.
Serial transmission	Data communication mode where bits are sent in sequence in a single fiber.
Single-mode fiber	A step index fiber wave guide in which only one mode (ray of light) propagates above the cutoff wavelength.
SNMP	Simple network management protocol. A simple protocol designed to allow networked entities (for example, hosts, routers) to exchange monitoring information.
Stripe size	Total amount of data in a disk stripe; that is, the block size multiplied by number of data disks in the stripe.
Stripe width	Total number of disks in a disk stripe.
Striping	Spreading or interleaving logical contiguous blocks of data across multiple independent disk spindles. Striping allows multiple disk controllers to simultaneously access data, improving performance.
Switch	The name of an implementation of the fabric topology.
Switched-loop architecture	Splits the drive interface into multiple, independent loops so that the RAID controller has its own drive loop, plus access to other drive loops. Improves performance and expansion flexibility for enterprise networks.
Throughput	A measure of sequential I/O performance, quoted as megabytes per second (MB/second). <i>See</i> IOPS and I/O rate.



Topology	The components used to connect two or more ports together. Also, a specific way of connecting those components, as in point-to-point, fabric, or arbitrated loop.
Transfer rate	The rate at which data is transferred, usually measured in Megabytes (MB) per second.
Volume	A volume is a virtual disk into which a file system, DBMS, or other application can place data. A volume can physically be a single disk partition or multiple disk partitions on one or more physical disk drives. Applications that use volumes do not need to be aware of their underlying physical structure. Software handles the mapping of virtual partition addresses to physical addresses.
Write-behind mode	A data write is acknowledged to the application host as soon as it is in cache, without having yet been committed to disk, in order to reduce write latency. Also known as write-back or fast-write mode.
Write-through mode	A data write is acknowledged only after data has been fully committed to disk.

Materials Abstract

All materials are available on SunWIN except where noted otherwise.

Collateral	Description	Purpose	Distribution	Token # or COMAC Order #
PowerPack				
– <i>Sun StorEdge™ T3 Array with 1-GB Cache Controller, Just the Facts</i>	Reference Guide (this document)	Sales Tool, Training	SunWIN, Reseller Web	311985
– <i>Sun StorEdge T3 Array Customer Presentation</i>	Presentation Overview; Slide Notes for Presentation	Sales Tool	SunWIN, Reseller Web	120838
– <i>Sun StorEdge T3 Array with 1-GB Cache Controller, Addendum to Customer Presentation</i>	Presentation Addendum	Sales Tool	SunWIN, Reseller Web	312720
– <i>Sun StorEdge T3 Array Technical Presentation</i>	Presentation with Slide Notes	Sales Tool	SunWIN, Reseller Web	120839
– <i>Sun StorEdge T3 Array Multi-Platform Presentation</i>	Presentation with Slide Notes	Sales Tool	SunWIN, Reseller Web	125114
Product Literature				
– <i>Sun StorEdge T3 Array Quick Reference Card</i>	Quick Reference Card	Sales Tool	SunWIN, Reseller Web	73691
– <i>Literature: Sun StorEdge T3 Array for the Workgroup Data Sheet</i>	Data Sheet	Sales Tool	SunWIN, Reseller Web, COMAC	108576, DE1074-3
– <i>Literature: Sun StorEdge T3 Array for the Enterprise Data Sheet</i>	Data Sheet	Sales Tool	SunWIN, Reseller Web, COMAC	117451, DE1165-3
– <i>Sun StorEdge T3 Array Elevator Pitch</i>	Presentation with Notes	Sales Tool	SunWIN, Reseller Web	120363
– <i>Sun StorEdge T3 Array FastFacts</i>	Fast Facts	Sales Tool	SunWIN, Reseller Web	120364
– <i>Sun StorEdge T3 Array, Just the Facts</i>	Comprehensive Guide to Sun StorEdge T3 Array Product Line	Sales Tool, Training	SunWIN, Reseller Web	112864



Collateral	Description	Purpose	Distribution	Token # or COMAC Order #
– <i>Sun StorEdge™ T3 Array Multiplatform Support, Just the Facts</i>	Reference Guide for Multiplatform Support	Sales Tool, Training	SunWIN, Reseller Web	127758
– <i>Sun StorEdge™ T3 Array for the Enterprise with Redundant Fibre Channel Switch Configurations, Just the Facts</i>	Reference Guide for Switch Support	Sales Tool, Training	SunWIN, Reseller Web	130189
White Papers				
– <i>Sun StorEdge T3 Array Performance Tuning White Paper</i>	Technical Brief	Training	SunWIN, Reseller Web	119879
– <i>Sun StorEdge T3 Array with 1-GB Cache Controller Architecture White Paper</i>	Technical Brief	Training	SunWIN, Reseller Web	311986
– <i>Fibre Channel Technology from Sun Microsystems</i>	Technical Brief	Training	SunWIN, Reseller Web	65659
– <i>Fibre Channel versus Alternative Storage Interfaces: An Overview</i>	Technical Brief	Training	SunWIN, Reseller Web	65663
Quote Sheets				
– <i>Sun StorEdge T3 Array Customer Quote Sheet</i>	Quote Sheet	Sales Tool	SunWIN, Reseller Web	119896, FE1270-0
– <i>Sun StorEdge T3 Array Partner Quote Sheet</i>	Quote Sheet	Sales Tool	SunWIN, Reseller Web	119934
Success Stories				
– <i>AB Watley Success Story</i>	Success Story	Sales Tool	SunWIN, Reseller Web	120147
– <i>Network Commerce Inc. Success Story</i>	Success Story	Sales Tool	SunWIN, Reseller Web	120005
– <i>Bluelight Success Story</i>	Success Story	Sales Tool	SunWIN, Reseller Web	120003
Competitive				
– <i>CLARiiON Beat Sheet Competitive White Paper</i>	Competitive White Paper	Training	SunWIN	112069
– <i>EMC Beat Sheet Competitive White Paper</i>	Competitive White Paper	Training	SunWIN	109825



Collateral	Description	Purpose	Distribution	Token # or COMAC Order #
– <i>Sun StorEdge T3 Array Competitive Presentation</i>	Competitive Presentation	Sales Tool	SunWIN, Reseller Web	120840
– <i>Sun StorEdge v. EMC Pocketfacts</i>	Pocket Facts	Training	SunWIN	117277, BE962-0
– <i>Competitive Edge Sun StorEdge T3 Solution vs. CLARiiON FC4500</i>	Competitive White Paper	Training	SunWIN	120367
– <i>Competitive Edge Sun StorEdge T3 Solution vs. Compaq RA8000</i>	Competitive White Paper	Training	SunWIN	120368
– <i>Competitive Edge Sun StorEdge T3 Solution vs. EMC 8430</i>	Competitive White Paper	Training	SunWIN	120369
External Web Sites				
– <i>Sun StorEdge Array Main Page</i>	http://www.sun.com/storage/disk.html			
– <i>Sun StorEdge T3 Array for the Enterprise Information</i>	http://www.sun.com/storage/t3es			
– <i>Sun StorEdge T3 Array for the Workgroup Information</i>	http://www.sun.com/storage/t3wg			
– <i>Fibre Channel Association</i>	http://www.fibrechannel.com			
– <i>Fibre Channel Loop Community</i>	http://www.fcloop.org			
Internal Web Sites				
– <i>Storage Products Internal Site for the Sun StorEdge T3 Array for the Enterprise</i>	http://webhome.ebay/networkstorage/products/T3ES			
– <i>Storage Products Internal for the Sun StorEdge T3 Array for the Workgroup</i>	http://webhome.ebay/networkstorage/products/T3WG			
– <i>Configuration Rules Page</i>	http://webhome.ebay/networkstorage/performance/confrules			
– <i>Resources Web Site</i>	http://webhome.ebay/networkstorage/contacts/			
– <i>Network Storage Sales Center (Help Desk)</i>	http://webhome.ebay/networkstorage/salesupportctr			
– <i>Sun StorEdge T3 Array Software Resources</i>	http://icode.ebay			
– <i>Onestop SE Information</i>	http://onestop.eng/hw/t3.shtml			
– <i>Product Documentation</i>	http://thedance.ebay/hardware/arrays/purple/hardware.html			
– <i>SunSpectrumSM Program Information</i>	http://service.central/TS/ESP/SunSpectrum/Feature_Matrix/index.html			



FAQs

Sun StorEdge™ T3 array Frequently Asked Questions list is no longer included in this document.

There are now two separate Sun StorEdge T3 array FAQ documents

- Internal use only = Sun StorEdge T3 Array INTERNAL FAQs, found on <http://webhome.ebay/networkstorage/products/> under Sun StorEdge T3 array
- External use = Sun StorEdge T3 Array EXTERNAL FAQs, found on <http://www.sun.com/storage/disk.html>

The FAQ lists are now kept as separate documents because of the frequency of updating required. This way, customers can find them in the same place and can know that they are up to date. Both FAQ lists are formatted using Courier 10 font for the purpose of maintaining formatting for copy and pasting into text-driven applications such as UNIX® platform-based e-mail messages.

