

# Sun StorageTek™ 6540 Array

## Just the Facts

### Copyrights

©2006 Sun Microsystems, Inc. All Rights Reserved.

Sun, Sun Microsystems, the Sun logo, Sun StorageTek, Solaris, Solstice, SunSpectrum, Sun Spectrum Bronze, SunSpectrum Silver, SunSpectrum Gold, and SunSpectrum Platinum are trademarks or registered trademarks of Sun Microsystems, Inc. in the United States and other countries.

All SPARC trademarks are used under license and are trademarks or registered trademarks of SPARC International, Inc. in the United States and other countries. Products bearing SPARC trademarks are based upon an architecture developed by Sun Microsystems, Inc.

AMD, the AMD Arrow logo, AMD Opteron, combinations thereof, are a trademark of Advanced Micro Devices, Inc.

Oracle, Netscape, and Mozilla are registered trademark 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: 09/08/2006

# Table of Contents

<b>Positioning</b> .....	<b>5</b>
Introduction.....	5
Robust Data Protection.....	6
Simplified, Application-Oriented Management.....	6
Round-the-Clock Availability.....	6
Key Features.....	7
1. Award-Winning Quality.....	7
Key Messages for the Sun StorageTek 6540 Array.....	9
4 Gbit/s – End-to-End infrastructure .....	9
Best Choice for Commercial, Business-Critical Applications.....	9
Near-Line Storage, Integrated Into a Backup or Disaster Recovery Solution.....	9
Consolidating Storage for Many Heterogeneous Server Operating Systems.....	9
Managing Large Amounts of Fixed Content with Larger Drives.....	9
Investment Protection .....	10
Product Family Placement.....	10
Positioning With the Sun StoragTek™ 6140 Array.....	10
Positioning With the Sun StorageTek™ 6940 System.....	10
Summary Comparison.....	11
Product Availability.....	11
Target Users.....	12
Target Markets.....	12
Large-Scale Configurations (Up to 224 Drives, Requiring High IOPS Performance).....	13
Bandwidth-Intensive Applications (Requiring High Throughput).....	13
Affordable Entry-Point (Without Sacrificing Availability).....	13
Low Acquisition Costs (Without Sacrificing Future Scalability).....	14
<b>Selling Highlights</b> .....	<b>15</b>
Market Value Proposition.....	15
Target Applications.....	15
Compatibility.....	15
<b>Enabling Technology</b> .....	<b>16</b>
Technology Overview.....	16
4 Gb/s Fibre Channel Technology.....	16
Sun StorageTek Common Array Manager Software.....	17
<b>System Architecture</b> .....	<b>19</b>
Overview.....	19
Sun StorageTek 6150 System Controller Tray .....	19
Sun StorageTek 6540 Array Controller.....	20
Cache Usage.....	22
Write Policy.....	22
Write Cache.....	22
Cache Protection.....	22
Cache Flushing.....	23
Immediate Volume Availability.....	23
Sun StorageTek 6540 Array 4 Gb/s Expansion Tray .....	23
Switched Architecture.....	24
Improved Diagnostics for Less Time in Degraded Mode.....	24
Lower Latency and Linear Scaling for More Responsive Applications.....	25
Preventive Maintenance.....	25
DACstore.....	25
SATA Drives.....	25
Battery.....	26
<b>Configuration</b> .....	<b>27</b>
Cabling.....	27
Configuration Flexibility .....	28

Management Software .....	30
Volumes.....	31
Storage Profiles.....	32
Virtual Disks.....	32
Storage Pools.....	33
Remote CLI Client .....	33
Configuration Metadata .....	33
Online Administration.....	34
Performance Tracking Utility.....	34
Global Hot Spares.....	35
Multipathing .....	35
Investment Protection / Data-In-place Migration .....	36
Heterogeneous Hosts.....	37
Sun StorageTek Data Snapshot Software — Optional.....	37
Key Features and Benefits.....	39
Sun StorageTek Integration Module for VERITAS NetBackup With Oracle® — Optional.....	39
Sun StorageTek Data Volume Copy Software — Optional.....	39
Key Features and Benefits.....	40
Target Applications for Both Sun StorageTek Data Snapshot and Data Volume Copy Software.....	41
Sun StorageTek Data Replicator Software — Optional.....	41
Key Features and Benefits.....	43
Target Applications for Sun StorageTek Data Replicator Software .....	43
<b>Features, Functions, and Benefits.....</b>	<b>45</b>
<b>Reliability, Availability, and Serviceability (RAS).....</b>	<b>48</b>
Reliability.....	48
Availability.....	48
Serviceability.....	49
<b>Best Practices.....</b>	<b>50</b>
Planning and Acquisition.....	50
Best Practices Require That the Application(s) Are Well Understood.....	50
How Many Drives are Required?.....	50
Which RAID Level is Best for the Application?.....	51
How Many Disk Drives Should Be In a Volume Group?.....	52
RAID Topology.....	53
Controller Cache Parameters.....	53
Host Striping of Data on an Array.....	53
Striping for Better Performance.....	54
Server Hardware — Bus Bandwidth and HBAs.....	56
Command Tag Queuing (CTQ).....	57
Fibre Channel Switch Fabric.....	57
Data Services Performance Enhancement.....	57
Sun StorageTek Data Snapshot Software.....	57
Sun StorageTek Data Volume Copy Software.....	58
Sun StorageTek Data Replicator Software.....	58
SANtricity Performance Monitor (Optional With SANtricity Only).....	58
In Summary.....	58
<b>Specifications.....</b>	<b>60</b>
<b>System Requirements .....</b>	<b>61</b>
System Requirements.....	61
Supported Operating Systems.....	61
New Sun StorageTek Interop Tool.....	62
Supported Sun Software.....	62
Third-Party Software.....	62
<b>System Configuration and Management.....</b>	<b>63</b>

Sun StorageTek Common Array Manager Software.....	63
Centralized Administration.....	63
Array Summary Window.....	63
Storage Window.....	63
Sun Storage Automated Diagnostic Environment.....	63
Management Host System Requirement:.....	64
<b>Ordering Information.....</b>	<b>65</b>
Shipping Configurations.....	65
<b>Marketing Part Numbering Scheme.....</b>	<b>66</b>
Sun StorageTek 6540 Array (Pre-Built, Pre-Racked Systems).....	66
Sun StorageTek 6540 Controller Tray.....	66
Configuration Matrix / Marketing Part Numbers.....	68
Ordering Details and Marketing Part Numbers .....	68
Webdesk Configurator.....	68
Sweet Spot Configurations.....	69
Customized ATO Options.....	70
Optional — Sun StorageTek Traffic Manager Software (SSTM) for Non-Solaris Operating Systems.....	76
Optional — Sun Services.....	76
SunSpectrumSM Instant Upgrades (Warranty Upgrades).....	77
<b>Upgrades.....</b>	<b>79</b>
Upgrade Paths.....	79
Sun Storage Consolidation Trade-in Program.....	79
<b>Service and Support.....</b>	<b>80</b>
Warranty.....	80
Education.....	81
Sun Enterprise Services.....	81
Professional Services.....	81
Storage Migration Service.....	82
Backup and Restore Services.....	82
Point-In-Time Copy Services.....	82
SAN Architecture and Implementation.....	82
Consulting Services.....	82
<b>Glossary.....</b>	<b>83</b>
<b>Materials Abstract.....</b>	<b>89</b>
<b>FAQs.....</b>	<b>91</b>
<b>Internal Information.....</b>	<b>93</b>
Competitive Information.....	93
More Performance Than the EMC CX3-80.....	94
More Flexibility and Scalability Than the HDS AMS1000.....	95
More Expansibility Features Than the HP EVA8000.....	95

# Positioning

---



**Figure 1.** Sun StorageTek™ 6540 Array

## Introduction

In high-performance technical computing (HPTC), success depends on the speed at which data is acquired, processed and distributed. These environments store vast amounts of data that is utilized for high-bandwidth programs and complex application processing such as sophisticated, data-intensive research, rich media, 3-D computer modeling, seismic processing, data mining, and large-scale simulation. The Sun StorageTek™ 6540 array powers storage systems designed for markets with high volumes of data processing and computing demands, providing the performance needed to excel in these applications. Despite recent competitive product announcements, the Sun StorageTek 6540 array remains the best performing array on the market in its class.

Customers trying to improve application performance, centralize primary (Fibre Channel) and secondary (SATA) storage, keep up with demanding data availability and data protection requirements, as well as unyielding data growth can appreciate the built-in data services, online administration, and industry-leading performance of the Sun StorageTek 6540 array. The array's modular, pay-as-you-grow scalability and broad flexibility enables the storage system to grow with the application, lowering acquisition and expansion costs. This array also enables disk-to-disk backup and archiving solutions though its ability to mix SATA and Fibre Channel (FC) storage media. It also enables IT administrators to deploy tiered storage scenarios with its centralized management and data migration tools. And, with multiple cache options and online scalability of up to 224 disk drives, the Sun StorageTek 6540 array easily satisfies demanding application requirements.

The Sun StorageTek 6540 array rebrands the StorageTek FlexLine FLX380 and integrates this array into the broader Sun StorageTek™ Modular Line family. The Sun StorageTek 6540 array builds on the performance leadership of StorageTek™ FLX380 and adds the following features:

- RoHs compliant
- New Common Storage Module 200 (CSM200) — the common building block for the entire Sun StorageTek Modular Line family
- Sun StorageTek™ Common Array Manager software — a single management interface for the entire Sun StorageTek Modular Line
- Sun™ Rack 1000-38 cabinet and new cosmetics — appearance aligned with the rest of the Sun StorageTek Modular Line
- Release of 4 Gb/s FC disk drives deliver end-to-end 4 Gb/s performance
- New larger 500 GB SATA drives for greater capacity and density
- Full suite of data services — Sun StorageTek™ Data SnapShot, Sun StorageTek™ Data Volume Copy, Sun StorageTek™ Data Replicator software

Plus all the great features of the StorageTek FLX380:

- Still the best price/performance system in its class
- Flexibility to support multiple drive types simultaneously
- Backwards compatibility — auto-negotiate 4, 2, and 1 Gb/s storage area networks (SANs)
- Investment protection — supports new expansion trays as well as previous trays (post General Availability or GA)

## **Robust Data Protection**

Integrated data services maximize application availability. Sun StorageTek Data Snapshot software is ideal for backups because it enables the administrator to create real-time copies of critical volumes with minimal disruption to applications. Sun StorageTek Data Volume Copy software provides the ability to quickly and securely duplicate full volumes, making it ideal for data mining, analysis, or enhanced backup operations. Sun StorageTek Data Replicator software enables organizations to maintain up-to-date copies of critical data at multiple sites, which is particularly useful in creating business continuity solutions and disaster planning scenarios.

## **Simplified, Application-Oriented Management**

Sun StorageTek Common Array Manager software makes it easy to manage multiple arrays, and storage profiles provide the ability to quickly and predictably provision storage to precise application demands. Built-in performance tracking helps optimize performance by enabling granular visibility of critical array statistics. And, the Sun StorageTek 6540 array's health and diagnostics utility minimizes downtime through predictive, preemptive maintenance and automated diagnostics.

## **Round-the-Clock Availability**

Robust, fully redundant architectures are required to keep businesses up and running around the clock. Extensive reliability, availability, and serviceability (RAS) features, including online reconfiguration and performance tuning, hot-swappable components, automated path failover, and non-disruptive firmware upgrades help ensure availability. Individually isolated drives help speed problem resolution for improved diagnostics. Point-to-point switched bunch of disks (SBOD) drive connectivity enables seamless scalability and pin-point control over back-end operations as well as lower latencies for more responsive applications.

The Sun StorageTek 6540 array is a modular, rack mounted and scalable array designed specifically to address the needs of the open systems environment. The Sun StorageTek 6540 array consists of a

minimum of *one* controller tray and up to 14 Common Storage Module 200 (CSM200) expansion trays. The Sun StorageTek 6540 controller tray (1 x 1) has three cache options — 4 GB, 8 GB, or 16 GB of cache (cache is not field upgradable). Expansion trays with FC drives (73 GB 15K RPM – 4 Gb/s, 146 GB 15K RPM – 4 Gb/s, 146 GB 10K RPM – 2 Gb/s, or 300 GB 10K RPM – 2 Gb/s) can be added to improve performance and increase capacity. For secondary storage requirements, expansion trays with SATA drives (SATA-II 500 GB 7.2K RPM) can be integrated to create a cost-effective solution.

The Sun StorageTek Common Array Manager software for the Sun StorageTek 6540 array lowers storage management costs with centralized management, run-anywhere flexibility, a single interface across all Sun StorageTek 6540 arrays with optional fully integrated data services such as storage domains, data snapshot, data volume copy, and data mirror features. The Sun StorageTek Common Array Manager software improves storage utilization with increased configuration flexibility, incremental capacity and performance scaling, storage consolidation, and capacity-efficient data services.

## Key Features

The Sun StorageTek 6540 array offers a number of innovative features to address real-world IT problems.

### 1. Award-Winning Quality

The Sun StorageTek 6540 array (originally part of the FlexLine 200/300 Series) dominated the competition with a number one ranking in the Diogenes Labs — Storage magazine Quality Awards (see “Quality Awards: Midrange Arrays,” Storage Magazine, November 2005).

### 2. Performance Leadership

When introduced in May, 2005 as the FlexLine FLX380, this was the fastest system in its class. Now, more than a year later and even with new products from key competitors like EMC, HP, and HDS, the Sun StorageTek 6540 array is still the fastest modular system in the mid-range. With 575,000K of cache-based I/Os per second (IOPS), enterprises can rev up their infrastructure, transact more business, increase productivity, and grow their top-line.

### 3. Robust Data Protection

The Sun StorageTek 6540 array delivers integrated data services that maximize application availability and provide advanced data protection.

- *Sun StorageTek Data Snapshot software* provides the ability to create real-time copies of critical volumes, which can be used to eliminate backup windows and rapidly recover data.
- *Sun StorageTek Data Volume Copy software* provides the ability to quickly and securely duplicate a volume for data mining, analysis, or enhanced backup operations.
- *Sun StorageTek Data Replicator software* enables real-time synchronous or asynchronous data replication to either local campus or metro or remote data centers to protect mission-critical information. Up-to-date copies of critical data can be maintained at multiple sites, which is particularly useful in creating business continuity solutions and disaster planning scenarios.

### 4. Simplified, Application-Oriented Management

The Sun StorageTek 6540 array's application-oriented management helps save time, while protecting the integrity and performance of business-critical applications. The Sun StorageTek Common Array Manager software makes it easy to manage numerous arrays, and storage profiles enable administrators to quickly and predictably provision storage to precise application demands. In addition, the built-in performance tracking utility helps optimize application performance and minimize downtime through predictive, preemptive maintenance and automated diagnostics for proactive health monitoring.

### 5. Round-the-Clock Availability

The Sun StorageTek 6540 array's robust, highly available, fully redundant architecture is designed to keep applications up and running around the clock, thanks to extensive RAS features, such as online administration and reconfiguration, redundant hot-swappable components, automated path failover, and non-disruptive controller firmware upgrades. Individually isolated drives enable improved diagnostics and faster problem resolution. Fewer drive-loop nodes and faster communication lowers loop latency and provides linearly scalable performance for more responsive applications. This is accomplished with switching technology (SBOD) that provides point-to-point connections of the drives in the expansion trays.

To maximize the power of these robust data services, Sun StorageTek 6540 array provides multiple classes of storage, from high-performance Fibre Channel to high-density, low-cost SATA, enabling tiered data-protection solutions. In addition, the modular, scalable and high-performance Sun StorageTek 6540 array offers the following features:

- Dual active, hot-swappable RAID controllers in a 4 RU chassis
- Backwards compatibility — auto-negotiates 4, 2, and 1 Gb/s SANs
- 4 GB, 8 GB, or 16 GB of battery-backed cache
- Scales from 5 drives to 224 drives (365 GB to 112 TB)
- Eight 4 Gb/s FC host ports
- Eight 4 Gb/s FC drive ports
- Up to 14 CSM200 expansion trays per system
- Five to 16 hard disk drives per tray
- Support for legacy expansion trays (available post GA)
- Optional 4, 8, 16, and 64 storage domains enable storage consolidation in multi-platform environments
- Storage profiles, built-in best practices wizard enables the administrator to quickly provision storage by application
- Performance tracking tool provides real-time information on over 30 of the most critical performance measurements
- Centralized administration, manages many Sun StorageTek 6540 arrays from anywhere via a browser or command-line interface (CLI) interface
- Dynamic reconfiguration enables: capacity and volume expansion, RAID and segment size migration, virtual disk reconfiguration, and firmware upgrades
- Support RAID 0, 1, 3, 5, and 1+0
- Support for point-to-point, arbitrated loop, and switched fabric
- Maximum 30 drives per virtual disk
- Support for server-attach and SAN-attach
- High-speed hardware XOR engine generates RAID parity calculations
- Modular *pay-as-you-grow* scalability enables optimal just-in-time purchasing
- Available rack-mounted in the Sun Rack 1000 or as loose components for mounting into an existing rack

## **Key Messages for the Sun StorageTek 6540 Array**

### **4 Gbit/s – End-to End infrastructure**

Sun StorageTek 6540 array, with its integrated 4 Gb/s technology, has the unique ability to match bandwidth to application needs and concurrently reduce risk with its performance. Large data sets load



faster and data protection methods such as remote replication, snapshots, and volume copies can happen in less time.

### **Best Choice for Commercial, Business-Critical Applications**

- Designed to handle large data sets — The Sun StorageTek Modular Line's most powerful dual controllers, combined with a 4 Gb/s end-to-end design, meets the rigors of commercial applications and high performance computing requirements.
- Enables better application responsiveness — These same features can also provide better application responsiveness by moving data faster than in 2 Gb/s environments.

### **Near-Line Storage, Integrated Into a Backup or Disaster Recovery Solution**

The Sun StorageTek 6540 array can be employed to implement a disk-to-disk back-up solution using FC as primary storage and SATA as secondary storage. This solution offers:

- Reduced backup windows
- Improved backup and recovery reliability
- Built-in RAID redundancy
- Easier data management
- Eliminates the need to monitor mechanical issues of tape for backup
- Uses existing backup software tools and processes

Use in concert with tape archiving for the following benefits:

- Application servers are off-line an average of 14 hours with traditional server-to-tape back-up — with SATA disk between the primary storage and tape, applications are available during back-up
- Supports rapid data restoration at disk speeds
- Functions as near-line storage
- Acts as a cache between primary storage and tape
- Lengthy backups no longer tie up primary storage

### **Consolidating Storage for Many Heterogeneous Server Operating Systems**

With support for the Solaris™ Operating System (OS), Windows, Linux, IBM-AIX, HP-UX, Novell, and IRIX, the Sun StorageTek 6540 array can support diverse computing environments with heterogeneous servers. Consolidate onto a single storage platform for ease of management.

### **Managing Large Amounts of Fixed Content with Larger Drives**

The *8-second rule* applies when businesses find that their customers are waiting too long to get what they want online, so they simply click elsewhere. The increase in Internet users and online content can overtax even a well designed Web site's performance. The combination of Sun StorageTek 6540 array's best in class performance and the largest, low-cost SATA drives available today enables this array to act as *caching services* to serve up Web content faster. Additionally, the Sun StorageTek Data Snapshot software enables disk-to-disk backup and disk-to-disk-to tape backup, enabling content management strategies according to the type and value of the data.

### **Investment Protection**

- Sun StorageTek 6540 array delivers flexibility and investment protection with data-in-place, on-line upgradeability across the entire Sun StorageTek Modular Line.
- Support for previous generations of SAN infrastructure and expansion trays let enterprises reuse and repurpose storage investments as their business needs change.

- Sun StorageTek Common Array Manager software provides common, application-oriented management across the entire Sun StorageTek Modular Line for ease-of-use, granular control, and business advantage driven by simplicity.
- IT staff learns the Sun StorageTek Common Array Manager software once, and uses the same software to manage across the Sun StorageTek Modular Line and through multiple upgrades.

## Product Family Placement

The Sun StorageTek 6540 array belongs to the Sun StorageTek Modular Line. The core family traits are value, flexibility, and data protection. Value is about getting the most for your money with a cost-effective balance of performance, availability, and scalability. It's about delivering high-end features previously only available on expensive, high-end systems and it's about affordable application and storage consolidation. Flexibility helps organizations to easily keep up with ever-evolving applications and to non-disruptively adapt to changing application requirements. World-class investment protection provides data-in-place upgrade-ability across the line and across generations, preserving investment. Data protection is second to none for business-critical applications with a highly-available system design, a broad suite of data protection software and 24x7 service and support.

## Positioning With the Sun StorageTek™ 6140 Array

The Sun StorageTek 6140 array is positioned as an entry-level array in the mid-range that offers enterprise-level data protection for cost-conscious companies. It is the Sun StorageTek Modular Line's most affordable array, enabling companies to start small and grow.

- Enterprise-class data protection
  - ☞ Protect data while conserving costly disk space with Sun StorageTek Data Snapshot software
  - ☞ Easily perform offline data analysis and test with full volume copies
- Comprehensive management
  - ☞ Application-oriented storage provisioning and pooling
  - ☞ Built-in performance tracking and automated diagnostics enable better utilization
- High availability
  - ☞ Redundant design with non-disruptive controller firmware upgrades

## Positioning With the Sun StorageTek™ 6940 System

With heterogeneous storage consolidation and migration, multi-dimensional scalability, and affordable high-end data protection, the Sun StorageTek 6940 system is intended for:

- Storage consolidation, with the ability to unify disparate servers and storage and simplify storage management
- Business continuity, with the ability to reduce unplanned downtime and easily duplicate and restore data to protect assets from loss
- Investment protection, with the ability to integrate multiple classes of storage for longer usable life, and the ability to help eliminate the complexity of migrating data

Positioned against the Sun StorageTek 6540 array, the Sun StorageTek 6940 system has:

- Enhanced data services features and functions, with dedicated CPUs to serve those data services
- The ability to scale capacity up to 512 drives and 28 4Gb/s FC host ports
- Symmetric performance with up to 32GB of cache and a cross-bar architecture
- Targeted for data center environments by incorporating heterogeneous storage resources

- Consolation features provide the ability to physically isolate multiple applications into storage domains

## Summary Comparison

	Sun StorageTek 6140 Array	Sun StorageTek 6540 Array	Sun StorageTek 6940 System
Maximum Drives	112	224	512
Maximum Cache	4 GB	16 GB	28 GB
Drive Options	FC and SATA	FC and SATA	FC and SATA
Management	Application-oriented	Application-oriented	Application-oriented
Data Snapshot	Yes	Yes	Yes
Data Mirror	No	No	Yes
Data Volume Copy	Yes	Yes	No
Data Replicator	Yes	Yes	Yes
Embedded Storage Service Processor	No	No	Yes
Dedicated Virtualization Engine	No	No	Yes
Designed for HPC	No	Yes	No
Designed for consolidating multiple applications	No	No	Yes
Warranty	3 years	3 years	3 years

## Product Availability

Key schedule dates for the Sun StorageTek 6540 array are:

Event	Date
Sun product Introduction (Presto)	07/11/06
WebDesk	07/11/06
Public Announcement	08/10/06
Revenue Release (RR)	08/30/06
General Availability (GA)	October 06

## Target Users

Most sales occur at the high end of the organization — often with CIOs and CFOs — who are concerned with top line growth, mitigating risk, and reducing costs. Business managers or IT directors have the most acute understanding of unrelenting storage growth and limited budgets. They have the expertise to recommend and then drive the acceptance of technical solutions within the organization.

New initiatives in the enterprise must be cost-justified with management teams who are demanding that IT department do more with less. It is with this group that the features of the Sun StorageTek 6540 array are going to resonate the most strongly. The following table provides tips on how to approach each type of user.

Target User	Buying Influence Needs
CEOs and CFOs	Discuss how the Sun StorageTek 6540 array can reduce total cost of ownership and mitigate the risk of inadequate storage and protection strategies, and how Sun can provide consultative strategies to help the company achieve its goals for growth.
Line of Business Manager	Discuss the impact of more responsive applications and a more productive staff on a company's ability to generate more revenue.
CIOs and IT directors	Discuss best-in-class price/performance and investment protection as a solution for tight budgets, limited resources, and the rising cost of managed storage.
Technical staff	Discuss ease of management, enterprise-class reliability, and full data protection as a solution for heterogeneous environments, devices not keeping up with high-performing applications, long backups, and unreliable recoveries.

## Target Markets

This newest member of the Sun storage family is for enterprises looking for a fully-featured, high performance solution for deploying business-critical applications. In high performance technical computing, success depends on the speed at which data is acquired, processed, and distributed. These environments store vast amounts of data that is utilized for high-bandwidth programs and complex application processing such as sophisticated, data-intensive research, rich media, 3-D computer modeling, seismic processing, data mining, and large scale simulation. This system is designed to meet the rigors of these environments. Organizations trying to improve application performance, centralize primary (FC) and secondary (SATA) storage, keep up with demanding data availability and data protection requirements, and unyielding data growth can appreciate the built-in data services, online administration, and industry-leading performance of the Sun StorageTek 6540 array.

The target markets are:

- Business-critical commercial applications like databases, CRM, SRM
- High-performance technical computing (HPTC)
- Enterprise and data center application (online transaction processing or OLTP, data warehouse, e-commerce)
- Computer generated animation (CGA)
- Service providers (e-mail, v-mail, static/dynamic Web servers)
- Image capture/retrieval applications (medical imaging, high-performance data acquisition)
- E-commerce (interactive online, e-commerce organizations, digital media, financial services organizations, and Internet-centric organizations)
- Government and Education
- Healthcare

The Sun StorageTek 6540 array is a great fit for Global 2000 data centers and government/scientific accounts with high-performance computing requirements, infrastructure bottlenecks, or backup/recovery challenges. Here are some markers for high-quality prospects:

Industry/Customer	Key Features to Highlight
Operating in open systems environments	<ul style="list-style-type: none"> <li>• Eight 4 Gb/s FC host-side ports</li> <li>• Broad interoperability matrix supports complex data center environments</li> </ul>
Running large, mission-critical databases or OLTP applications	<ul style="list-style-type: none"> <li>• Powerful dual controllers, up to 16 GB cache, and 4 Gb/s end-to-end architecture</li> </ul>

Industry/Customer	Key Features to Highlight
Urgent need for speed and availability	<ul style="list-style-type: none"> <li>• Best performing storage array in its class, audited by third party benchmark council (SPC)</li> <li>• Storage profiles enable predictable and precise application-oriented provisioning to quickly provision storage for mission-critical applications</li> <li>• Only midrange array that enables applications to continue running while changing every critical system component, including the midplane</li> </ul>
Upgrading SAN infrastructure to 4 Gb/s FC	<ul style="list-style-type: none"> <li>• Support for 4 Gb/s end-to-end, as well as the ability to auto-negotiate to 1, 2, and 4 Gb/s</li> <li>• Support for mixed speed loops</li> <li>• Ability to migrate to 4 Gb/s with data intact</li> </ul>

Almost all businesses today require continuous uptime. With advanced data services, the Sun StorageTek 6540 array enables businesses to cost-effectively meet these critical requirements while enhancing their backup operations. The Sun StorageTek 6540 array is ideally suited for businesses seeking a high-performance array with enterprise-level data protection and centralized, application-oriented management.

### **Large-Scale Configurations (Up to 224 Drives, Requiring High IOPS Performance)**

In large-scale configurations, the Sun StorageTek 6540 array provides exceptional IOPS performance value to easily handle compute-intensive, transaction-oriented applications requiring random, small-block I/O — such as database, enterprise resource planning (ERP), OLTP, Web servers, and Microsoft Exchange servers. Its industry-leading IOPS — including 575,000 dual-controller burst IOPS from cache and 86,000 sustained IOPS from disk — easily handles the most transaction-heavy applications. With the availability of SBOD (switched bunch of disks) expansion trays, the Sun StorageTek 6540 array can easily scale as the demand of applications increase.

### **Bandwidth-Intensive Applications (Requiring High Throughput)**

The Sun StorageTek 6540 array's 4 Gb/s host connectivity, advanced cache algorithm, and switched SBOD drive connectivity make it an exceptional solution for high-bandwidth applications, such as seismic processing, video streaming, and multi-media servers. Its industry-leading performance loads large data sets in less time and handles multiple I/O streams for rich media applications. With sustained throughput of 1600 MB/s for disk reads and 1400 MB/s for disk writes, it is ideal for HPC environments that require the highest levels of throughput for bandwidth-intensive applications. Additionally, the Sun StorageTek 6540 array provides the extra horsepower to efficiently handle the additional performance demands of data replication operations.

### **Affordable Entry-Point (Without Sacrificing Availability)**

The Sun StorageTek 6540 array's modular design enables lower acquisition and expansion costs while maintaining maximum flexibility. Its modular design helps avoid over-configuration — creating an affordable entry-point without sacrificing future scalability. And, the modular pay-as-you-grow scalability keeps expansion costs down through optimal just-in-time purchasing.

### **Low Acquisition Costs (Without Sacrificing Future Scalability)**

The dynamic, online reconfiguration and expansion capabilities lower acquisition and expansion costs. Companies can purchase what they need, when they need it, match the storage to the application, add storage without impacting data availability, and create an affordable system without sacrificing future scalability. The array's support for up to 224 disk drives in FC or SATA make it a great choice for performance- and capacity-oriented storage requirements. In addition, the Sun StorageTek 6540 array (as

well as the Sun StorageTek 6140 array) is the only storage system that enables the IT operator to entirely change the volume characteristics on the fly, reducing costs and improving operational efficiencies by scaling with the application.

# Selling Highlights

---

## Market Value Proposition

- Improved storage productivity — tiered storage reduces the cost of keeping data available and improves efficiency across the storage infrastructure.
- Ongoing investment protection — The Sun StorageTek 6540 array integrates easily into existing infrastructures. It auto-negotiates 1, 2, and 4 Gb/s FC speeds, and supports both 2 and 4 Gb/s drives simultaneously.
- High data availability — high-performance I/O, redundant controllers, path failover, and simple upgrades, and robust data protection software help keep data available at all times.
- Fast restores — high-speed restores with enhanced data protection minimize business interruptions, helping to keep data available around the clock for the always-open enterprise.

## Target Applications

The most appropriate environments for the Sun StorageTek 6540 array are those that need extreme OLTP and database performance. Such environments particularly need the higher availability, performance and improvement in data protection processes that the Sun StorageTek 6540 array offers.

Large databases that have heavy workloads, OLTP, data mining, etc. — the Sun StorageTek 6540 array provides the performance demanded by HPC environments that require the highest levels of throughput for bandwidth-intensive applications. It is ideally suited to markets with high volumes of data processing and computing demands including energy, entertainment, biosciences, supercomputing, and scientific research and development. The Sun StorageTek 6540 array's controller is the culmination of more than 20 years of design knowledge focused entirely on open systems storage. This end-to-end, 4 Gb/s Fibre Channel RAID controller delivers the performance demanded by HPC environments that require the highest levels of throughput for bandwidth intensive applications.

- Any sort of data replication or mirroring for back-up are great matches.
- Any applications that involve streaming data, such as media companies creating video clips and movies on demand, surveillance systems and satellite feeds, or hospitals storing medical images.
- Given these applications — it's the usual suspects for early adopters — federal accounts like Department of Defense, Home Land Security, government labs, big banks, etc.

## Compatibility

For the most updated compatibility matrix for the Sun StorageTek 6540 array, refer to the online Interop Tool at: <https://extranet.stortek.com/interop/interop?cmd=report>

Or contact Sales Support (sales.support@sun.com) for more information.

# Enabling Technology

---

## Technology Overview

The Sun StorageTek 6540 array is a modular, rack mountable, highly scalable, high performance system with integrated data services. The standard configuration consists of an integrated controller tray (1 x 1) with dual-active FC RAID controllers in a 4 RU enclosure. Up to 14 Common Storage Module 200 (CSM200) expansion trays (0 x 1) can be dynamically added to the base controller tray to create a maximum configuration with 224 drives. Expansion trays (0 x 1) can support both FC and SATA drives. Fibre Channel drives, either 2Gb/s (146 GB, 10K RPM or 300 GB, 10K RPM) or 4 Gb/s ( 73 GB, 15K RPM, or 146 GB, 15K RPM) can be added to improve performance and increase capacity, while SATA drives (500 GB, 7K RPM) can be utilized for secondary storage requirements. The Sun StorageTek 6540 array's ability to concurrently support FC and SATA drives allows for implementation of a single storage solution for both near-line and primary storage. Also, 2 Gb/s and 4 Gb/s FC drives can be mixed in the same system (not in the same tray and not on the same loop) to provide maximum flexibility when architecting for performance, price, and reliability.

## 4 Gb/s Fibre Channel Technology

4 Gb/s FC systems are ideally suited for applications that need to quickly transfer large amounts of data, such as remote replication across a SAN, streaming video on demand, medical imaging, data mining, data warehousing, and large databases supporting OLTP. Additionally, large increases in online data, driven by applications that utilize radio frequency ID (RFID), can leverage the benefits of 4 Gb/s FC technology.

The benefits of 4 Gb/s technology include:

- Faster communication between servers and storage
- The ability to more rapidly transfer data from storage to server or restore data from online backup media
- Higher reliability as a result of a reduced number of connections
- Investment protection due to backward compatibility with 1 Gb/s and 2 Gb/s components

Faster communications between servers and storage, and within the storage system itself can have benefits in the following markets:

- **Tiered storage** — more organizations are starting to deploy storage tiers to reduce storage costs, which means that data needs to be routinely migrated from primary (tier 1) storage to secondary (tier 2) storage across the storage network. Depending on the amount of data, this can be time-consuming. 4 Gb/s FC provides a massive pipe that allows data to be quickly moved from one storage system to another, enabling transfers to happen at twice the speed of 2 Gb/s FC technology.
- **Data replication** — data replication is an ideal mechanism to provide data protection against disasters, whether deployed on a campus or over a wide area network (WAN). Replicated data is the same at the remote site and can be used in a variety of ways such as data restoration, mining, analysis, testing, and remote backup. Depending on the application, synchronizing the local and remote sites can require a large portion of bandwidth on the SAN or WAN. 4 Gb/s FC offers a high-bandwidth infrastructure that is ideally suited for campus replication over a SAN, enabling data to be synchronized and restored at twice the speed of 2 Gb/s FC.
- **Streaming video** — large block, sequential I/O applications that include content creation and delivery, modeling, rendering, and publishing can benefit from the additional bandwidth provided by 4 Gb/s FC.



- **Large data analysis** — a 4 Gb/s FC storage can enable companies to accelerate and scale simulation, visualization, modeling, and rendering applications.
- **Database in memory** — enterprises running database-in-memory applications need to have very large data sets in memory. Loading and refreshing these data sets is a time-critical process. With 4 Gb/s storage systems, data can be loaded from storage in half the time.
- **Faster data recovery** — as most IT managers know, recovering data takes longer than backing it up. With 7x24x365 operations and stringent service level agreements (SLAs), the time it takes to recover is becoming more important every day. Storage based on 4 Gb/s FC can deliver the performance required to shorten recovery time.
- **Data warehousing** — data warehouses require both high bandwidth and high random I/O performance. The balanced performance of 4 Gb/s FC offers the highest level of performance in both bandwidth and IOPS.

## Sun StorageTek Common Array Manager Software

The Sun StorageTek 6540 array comes with the Sun StorageTek Common Array Manager software, which is common across the Sun StorageTek Modular Line. The Sun StorageTek Common Array Manager software helps reduce complexity and lowers storage management costs with centralized management, run-anywhere flexibility, a single interface across many Sun StorageTek 6540 arrays and fully integrated Storage Automated Diagnostic Environment software plus integrated management for all of the optional data services software. The Sun StorageTek Common Array Manager software improves storage utilization through configuration flexibility, incremental capacity and performance scaling, storage consolidation, and capacity-efficient data services. This robust array management software helps enable an organization to easily scale storage resources without increasing administration costs or complexity.

The Sun StorageTek Common Array Manager software controls the Sun StorageTek 6540 array and provides administrators a powerful, yet easy to use management interface. With the Sun StorageTek Common Array Manager software, all administrative tasks, including configuration, re-configuration, expansion, maintenance, and performance tracking can be performed with no system downtime and no interruption to system I/O. The Sun StorageTek Common Array Manager software's configuration flexibility includes the ability to mix drive types (capacities and RPM speeds), RAID levels, segment sizes, virtual disk sizes, volume characteristics, and cache policies all within a single storage system. And, its centralized administration and run-anywhere GUI enables management of all Sun StorageTek 6540 arrays from any location on the network — regardless of host platforms.

The Sun StorageTek Common Array Manager software's robust data services extend the functionality of Sun StorageTek Common Array Manager software — creating even more powerful storage. Storage domains allows a single storage system to function as multiple, logical storage systems — enabling storage consolidation in heterogeneous environments. Sun StorageTek Data Snapshot software instantaneously creates capacity-efficient, point-in-time volume images, providing a logical volume for such uses as file restoration and backup. Sun StorageTek Data Volume Copy software creates a complete physical copy, or a clone, of a volume within a storage system. This unique entity can be assigned to any host and used for application testing, development, information analysis, or data mining. The Sun StorageTek Data Replicator software enables real-time synchronous or asynchronous data replication to either local campus, metro, or remote data centers to protect a company's mission-critical information. Organizations can maintain up-to-date copies of critical data at multiple sites, which is particularly useful in creating business continuity solutions and disaster planning scenarios.

The Storage Automated Diagnostic Environment component of the Sun StorageTek Common Array Manager software offers proactive health checking, intelligent diagnosis, fault isolation, event notification, and fault management reporting for the Sun StorageTek 6540 array from a single management console. This software helps improve the time to recover and increases infrastructure uptime, thus contributing to overall improved application service levels.

In addition to integrated data services, the Sun StorageTek 6540 array also offers a wide support of various operating platforms for heterogeneous data center environments. The supported operating platforms include: Solaris 8, 9, and 10 Operating Systems, Microsoft Windows 2000, Windows Server 2003, RedHat Linux, SUSE Linux, IBM AIX, HP-UX, Novell NetWare, and SGI IRIX (single path only).

# System Architecture


## Overview

The Sun StorageTek 6540 array is a high-performance, enterprise-class, end-to-end 4 Gb/s Fibre Channel solution that combines massive performance with the high reliability, availability, flexibility, and manageability. The Sun StorageTek 6540 array consists of two primary components, the controller tray (1 x 1), and expansion trays (0 x 1). Custom configurations can be built to address specific performance or capacity requirements. Each of the primary components is fully fault tolerant with automated I/O path failover and redundant, hot-swappable components to help ensure the highest availability.

## Sun StorageTek 6540 Array Controller Tray

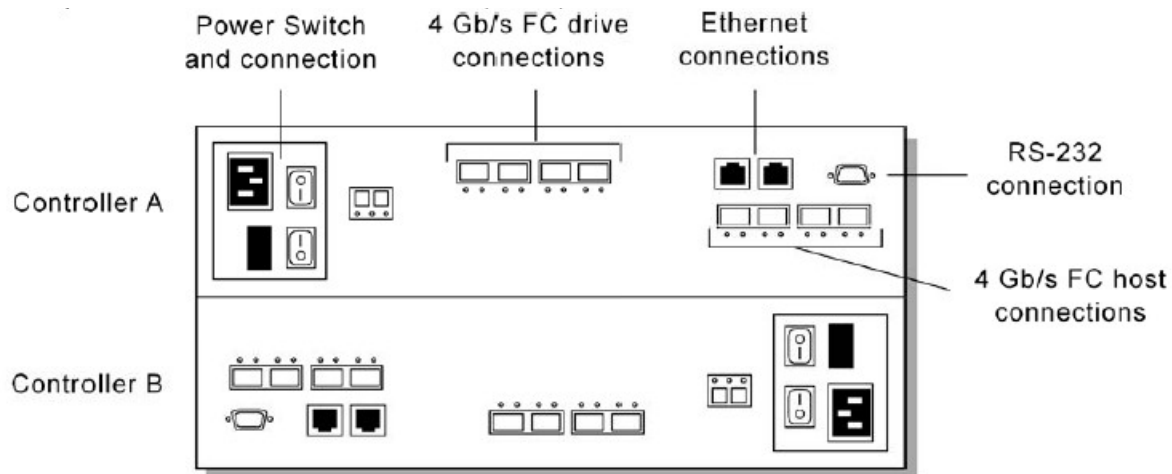
The Sun StorageTek 6540 array controller tray houses the dual-active, intelligent array controllers, provides eight host or SAN connections and supports up to 224 disk drives. Redundant controllers and a OS-dependent failover driver provides automated I/O path protection and helps ensure continued data access. In the event of a component failure anywhere in the data path between server and disk drives, I/O is re-routed down the surviving path for uninterrupted access to the data and application availability.

The controller tray has five total components, as illustrated in Figure 2 — all of which are hot-swappable, field-replaceable units or FRUs (the battery is classified as a customer replaceable unit, CRU). Each FRU can be accessed either from the front or back of the command module and easily removed and/or replaced. The controllers insert from the rear of the module and support all external connections — including host and drive connections, Ethernet connections, RS-232 connections, and power supply connections. Each controller module houses a 375-watt power supply (of which only 233 watts are used), two fans, and a battery charger. The interconnect module houses the cache batteries and functions as the midplane. This implementation enables the midplane to be removed and/or replaced without interrupting I/O.



**Figure 2.** Sun StorageTek 6540 Array Controllers

The Sun StorageTek 6540 array has eight 4 Gb/s Fibre Channel *host interfaces* (four per controller) supporting eight fabric connections (FC-SW) for storage area networks or eight arbitrated loop (FC-AL) connections for direct-attached servers. These ports are illustrated in Figure 3. Its native 4 Gb/s interfaces auto-negotiate 4, 2, and 1 Gb/s Fibre Channel link speeds, enabling the system to easily integrate into existing infrastructures.



**Figure 3.** Sun StorageTek 6540 Array Controllers (Front View)

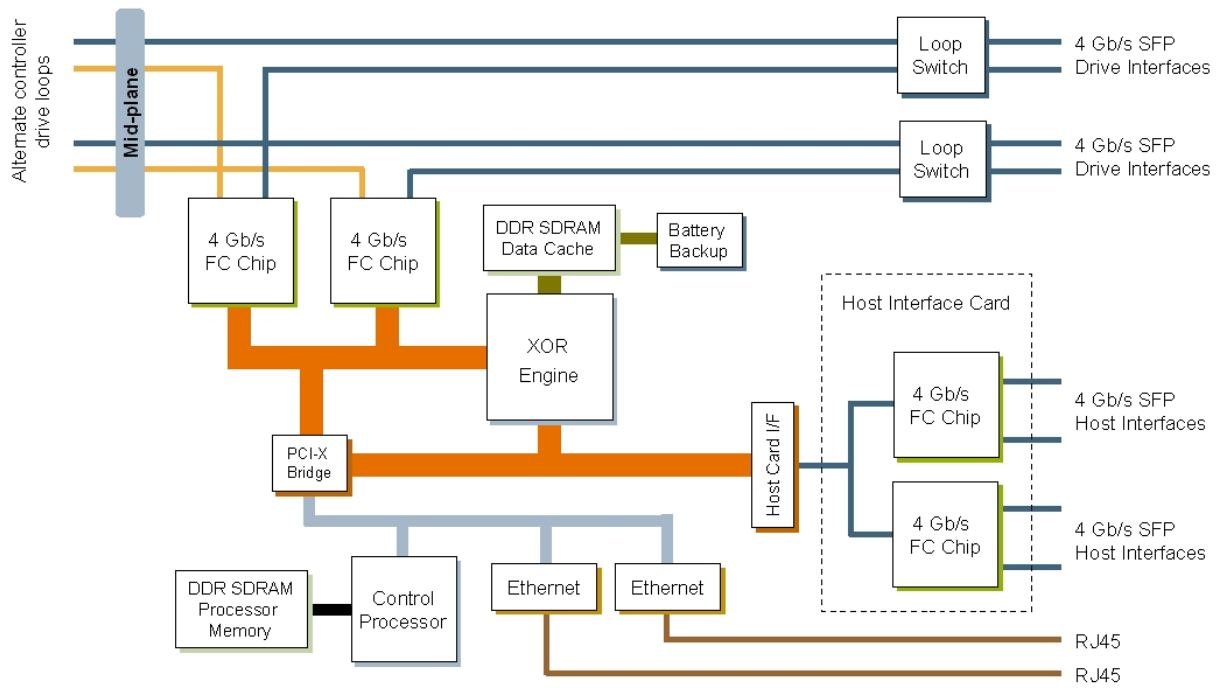
The Sun StorageTek 6540 array has eight 4 Gb/s *drive-side interfaces*. The controller tray supports up to 224 low-profile Fibre Channel or SATA disk drives in 14 16-drive Common Storage Module 200 with a maximum of 2048 addressable volumes. The drive side interfaces auto-negotiate 4 and 2 Gb/s FC link speeds, enabling the system to concurrently support current 2 Gb/s drive modules and 4 Gb/s drive module at their native speeds.

The controller tray has four Ethernet ports — two per controller. One port per controller is for managing the Sun StorageTek 6540 array over the network. A second port per controller is available for advanced diagnostics by service personnel and is isolated to prevent exposure to the LAN. Additionally, each controller has an RS-232 interface connector for troubleshooting and running diagnostics.

### Sun StorageTek 6540 Array Controller

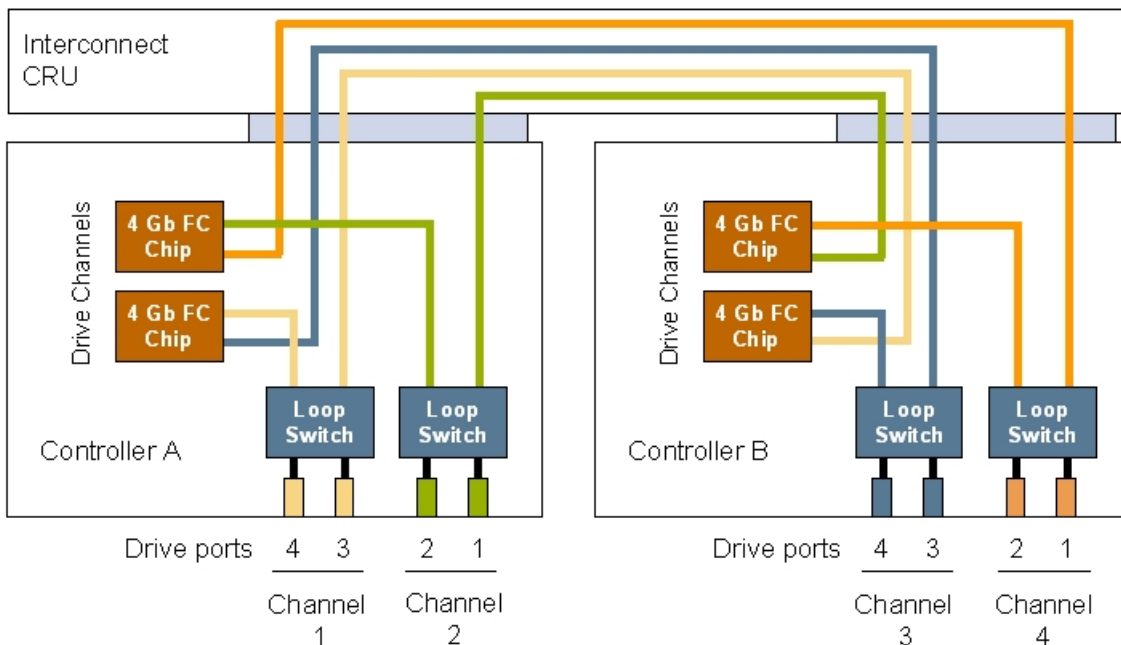
The Sun StorageTek 6540 array controller utilizes multiple, specialized processing elements, each with its own memory to optimize processing power. This approach reduces contention and enables each element to focus on its specific job. The next generation, high-speed XOR engine generates RAID parity with no performance penalty, enabling this compute intensive task of to be handled efficiently and effortlessly. A separate processor focuses on data movement control, allowing setup and control instructions to be processed and dispatched independent of data.

Each controller has four 4 Gb/s FC host-side interfaces and four 4 Gb/s FC drive channels. On the drive side, two of the channels (one per chip) connect to local loop switches and the other two connect to the alternate controller's loop switches through the interconnect module (midplane). (See Figure 4 below.) This implementation allows each controller to access all eight external drive loops for maximum performance and availability.



**Figure 4. Sun StorageTek 6540 Array Controller Diagram**

Each loop switch is capable of delivering 400 MB/s of bandwidth through its ports; however, both ports of a loop switch run at the same link speed, shown in Figure 5. This ingenious design enables the controller to simultaneously support both 2 Gb/s and 4 Gb/s drive enclosures while allowing each to run at its native speed on separate loop switches.



**Figure 5. Sun StorageTek™ 6540 Array Loop Switch Diagram**

## Cache Usage

The Sun StorageTek 6540 array dynamically utilizes cache for both read and write caching. This allows for more efficient use of available cache by automatically adapting its utilization to match the I/O requests from hosts as the needs of the applications change during the computing day. This adaptability is one reason that the Sun StorageTek 6540 array can sustain higher levels of real-world performance than many monolithic and modular systems with larger caches.

The Sun StorageTek 6540 array can provide up to 16 GB of cache for those applications with a high *locality of reference* that can derive the most benefit from larger cache. For applications that infrequently access the same data and do not benefit from larger cache, smaller 4 and 8 GB options are available. Larger cache may also benefit environments that utilize data services such as sun StorageTek Data Snapshot or Data Replicator software, or that are streaming large, fixed content files like video.

## Write Policy

The Sun StorageTek 6540 array offers both write-back and write-through caching options to optimize performance or protection. The write policy determines whether I/O completion is signaled when data transfers to the cache or when data writes to disk. With write-back caching, the write request is signaled complete as soon as the data is in cache, and actual writing to disk occurs at a later time. Write-through caching signals the completion of a write request when the data is safely stored on disk. Write-back caching creates significantly higher write performance for most applications, and any potential risk is eliminated as the Sun StorageTek 6540 array cache is mirrored and battery backed up — helping to ensure that it is protected at all times.

## Write Cache

The Sun StorageTek 6540 array's cache algorithms are designed to provide the lowest latency I/O access possible when securing the user's data in storage system cache. And, the Sun StorageTek 6540 array makes its cache ready to receive the next write request from the host by efficiently writing cached data to disk in a manner that minimizes disk I/O. By identifying high locality of reference for write data in its cache, the Sun StorageTek 6540 array is able to group all data in a RAID stripe and minimize the amount of time required to write data to disk, thus eliminating the need for disk read operations in computing parity in RAID 5. This helps ensure that disk utilization is minimized for a particular usage pattern, allowing maximum I/O throughput per spindle available.

## Cache Protection

Cache mirroring protects unwritten data during a controller or cache memory failure. When cache mirroring is enabled, data that writes to the cache memory of one controller also writes to the cache memory of the other controller. Therefore, if one controller fails, the other can complete all outstanding write operations. This cache parameter is set on a per volume basis. Data in the controller cache memory is also protected in case of power outages. The Sun StorageTek 6540 array controllers contain separate, independent batteries that protect the data in cache by keeping a level of power until the data can write to the drive media. The Sun StorageTek 6540 array contains enough battery-power to easily protect a 1 GB cache memory for up to 7 days of continuous system-wide power outage. (It might take up to 24 hours to charge the battery after the initial set up.) Battery performance is continuously monitored. If batteries are ever removed, or non-functional, the Sun StorageTek 6540 array retains total integrity by automatically switching to the write-through caching mode.

## Cache Flushing

To prevent data loss or corruption, the controller periodically writes unwritten data residing in its cache to disk (flushes the cache) based on one of two factors:

- A time-based flush writes cache data to disk after a user-defined period of time (in seconds). This value can range from immediate to infinite, with a value of 10 seconds or less recommended for optimal data protection.
- A threshold-based flush occurs when the amount of unwritten data in cache reaches a certain level, called a start percentage. The controller writes cache data to disk until the amount of data in cache drops to a stop percentage level. For example, the controller can start flushing the cache when the cache reaches 80 percent full, and stop flushing the cache when the cache reaches 20 percent full. The start and stop percentage levels are user-defined to maximize data protection and/or performance.

## Immediate Volume Availability

Immediate Volume Availability allows reads and writes immediately after a *Create Volume* command is issued for RAID 1, 3, 5, and 1+0. The logical unit is available for normal I/O requests immediately after a *good status* is returned to the host in response to the *Create Volume* command. Immediate Volume Availability eliminates the need for formats (write zeroes), both within the controller and from the host. This feature initializes the first and the last 10 MB of the volume to zero. This is done to help ensure that any file-system information, that might have been stored on this logical unit disk previously, is no longer valid. A read from another area on the volume that was not previously written returns un-initialized data (not zeroes). Thus, the host cannot expect unwritten data to equate to zeroes. If this data must be zeroes then Immediate Volume Availability must be disabled or a new *Format* command must be issued after the redundancy initialization completes. After initializing the two 10 MB regions of the volume, the controller returns *good status* to the *Create Volume* command. The controller then initiates background redundancy initialization to establish good parity for the volume. The rate of the initialization is controlled by the reconstruction parameters supplied in the *Create Volume* command. If the host sends a *Format* command to a volume on which redundancy initialization is either pending or in progress, a *good status* is returned in response to that *Format* command without initiating a write zeroes task. Should a *Format* command arrive after redundancy initialization has completed, a write zeroes task is performed.

## Sun StorageTek 6540 Array 4 Gb/s Expansion Tray

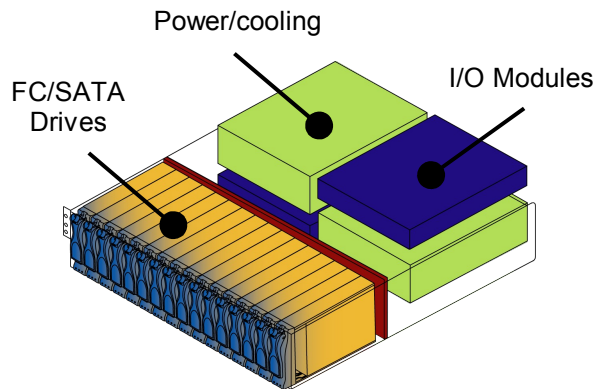
The Sun StorageTek 6540 array 4 Gb/s expansion tray is also known as the Common Storage Module 200 (CSM200). With an imbedded *loop switch*, redundant dual-active drive loop access, and dual-ported FC or SATA drives, the Sun StorageTek 6540 array expansion tray is more than *just-a-bunch-of-disks*. Combined, these advanced features deliver maximum performance and help to ensure complete accessibility to all drives in the event of a loop or cable failure. For maximum investment protection, these CSM200 expansion modules are the same trays that can be used in the entire Sun StorageTek Modular Line (Sun StorageTek 6140 array, and Sun StorageTek 6940 system).

Like the Sun StorageTek 6540 controller tray, all of the primary components in the 4 Gb/s expansion tray are hot-swappable. This includes the FC or SATA disk drives, power supplies, cooling fans, and I/O modules (IOMs), shown in Figure 6. The IOMs deliver online information about component status, such as temperature, fan speed, or faults, to the storage management software and provide the ability to cascade multiple drive modules to a single array module.

The expansion tray accommodates up to 16 low-profile FC or SATA disk drives in a 3U enclosure, enabling the Sun StorageTek 6540 array to maximize spindle density and capacity in a minimum amount of space. A single 19-inch, 38 RU cabinet can support a Sun StorageTek 6540 array controller module and up to 11 CSM200 expansion trays for 176 spindles in a single data center floor tile. And, as the Sun StorageTek 6540 array supports a range of drive capacities and rotational speeds, both spindle density (dollars per IOPS per square foot) and storage capacity (dollars per MB per square foot) can be maximized in a single system. As environmental costs continue to grow, the reduced rack space and low

power and cooling demands of the Sun StorageTek 6540 array can significantly lower total cost of ownership.

The highest levels of performance and availability associated with FC drive enclosures are assured when intermixing FC and SATA drives. Each SATA drive in the 4 Gb/s expansion tray includes a SATA Interface Card, which provides these drives with the dual-link capability and switching functionality needed to achieve redundant I/O paths — a requirement for enterprise-class storage systems. Additionally, the SATA Interface Card's speed-matching technology enables the 3 Gb/s SATA II drive to function effectively as a 4 Gb/s drive — ensuring each 4 Gb/s expansion tray achieves maximum performance when intermixed with 4 Gb/s FC and 3 Gb/s SATA II drives.



**Figure 6.** Sun StorageTek 6540 Array Expansion Tray

## Switched Architecture

The 4 Gb/s expansion tray has an imbedded loop switch that isolates each drive on a private loop with the I/O module, enabling direct and detailed communication with each individual drive. Each tray has two loop switches; one on each I/O module. This new, advanced technology introduces a number of customer and Sun Service advantages, such as improved diagnostics, lower latency and linear scalability, and preventative maintenance.

## Improved Diagnostics for Less Time in Degraded Mode

The embedded loop switch creates what is essentially a private FC arbitrated loop with each individual drive. Drive isolation eliminates the risk of a single drive disrupting the loop and potentially causing other drives to fail on the loop (referred to as fault propagation). Drive isolation also enables diagnostics to be performed on an individual drive level, simplifying analysis of drive and loop problems. The longer it takes to identify and diagnose a problem drive, the longer the overall storage system performance is degraded. In traditional loop-based drive enclosures, where multiple drives are operating on a given loop, it can be extremely difficult to identify which drive is faulty. By isolating each drive on a private loop, the CSM200 expansion tray greatly simplifies analysis of drive and loop problems. More detailed and isolated diagnostic information can be gathered and problem drives can be located and remedied faster.



## Lower Latency and Linear Scaling for More Responsive Applications

With traditional loop-based enclosures, each drive is a node on the drive loop. Large loops with multiple nodes create a longer and slower I/O path from a controller to a drive and back to the controller. The Sun StorageTek 6540 array's switched drive module adds only one node at the I/O module to each redundant drive loop, and its drives are just one point-to-point step away from the now shorter and quicker loop. Fewer nodes and faster communication significantly lowers the latency and provides linearly scalable performance for more responsive applications. This means that there is no performance “knee” for the 6540 as in previous generations. Performance scales with capacity all the way to a maximum 224 drive configuration. Additionally, the Sun StorageTek 6540 array's switched enclosure's ability to acquire detailed diagnostic information enables more effective performance tuning.

## Preventive Maintenance

The Sun StorageTek 6540 array's switched enclosures put all drives through an insertion test (port test before insert, or PTBI) before they are added to the FC loop. The intent is to identify faulty drives before they can disrupt the loop. PTBI causes a drive that is added to the loop to be loop initialization primatived (LIP'd) and then looped back to itself. The switched drive module monitors the drive for 5 seconds, and if the loop remains up, the drive is LIP'd again and allowed to join the loop.

## DACstore

The Sun StorageTek 6540 array stores configuration metadata in a private 512 MB region on every configured drive called the DACstore. The DACstore area contains drive state and status information, volume state and status information, and controller and subsystem information. The DACstore on each drive stores that drive's state and status, the worldwide name of its volume group, the volumes it contains, and the definitions for those volumes. Additionally, one drive in each volume group (with a minimum of three in each storage system) stores the controller- and subsystem-level information.

Storing metadata in the nonvolatile DACstore provides higher availability and enables easier re-configuration and migration. As the system configuration data resides on every configured drive, controllers and/or multiple drives can be removed or swapped without losing the system configuration. Drives can be relocated within the storage system to improve channel utilization/protection or even migrated as a complete volume group into another storage system. And in both instances, all configuration metadata and user data remains intact on the drives.

## SATA Drives

All advanced technology attachment (ATA) drives are *not* the same, and all storage systems based on ATA drives are *not* the same. The Sun StorageTek 6540 array with SATA drives are designed for midrange and enterprise application environments that need lower cost storage, but cannot give up the availability, manageability, data integrity, and software/hardware features achieved with FC based storage systems. The robust RAID controller firmware and enclosure technology enable higher performance, availability, and manageability — creating robust SATA storage solutions that are truly enterprise class.

SATA and FC drives have very different characteristics — from performance, to reliability, to price. FC drives are still the best choice for applications that are high use, high workload, require high reliability, and high performance. However, SATA drives are optimal for a range of throughput-intensive applications where the activity consists of large block writes and occasional reads, as opposed to heavy block-level activity. SATA drives are best suited for the following applications:

- Bandwidth streaming applications (such as seismic processing)
- Certain business applications with low or limited IOPS performance
- Near-line secondary storage

- Disaster recovery typical site
- Reference data
- Tape cache
- Fixed content to manage retention data
- Hierarchical storage management

SATA drives deliver about one-quarter the performance of FC with the same number of drives and are unacceptable for applications requiring anywhere from adequate to maximum IOPS performance. SATA drives are a good fit for bandwidth applications as they can provide enough throughput to maximize a controller's MB/s performance. To achieve this, however, it does take about a third more SATA drives as Fibre Channel.

The SATA drive bracket is designed with a unique SATA interface card, which enables SATA drives to emulate the FC interface, providing the following key features and functionality:

- **Availability** — SATA disk drives are single-ported and lack true multipathing and multi-initiator capabilities. The SATA interface card provides the dual-link capability and switching functionality needed to achieve dual-active redundant I/O paths.
- **Connectivity** — Enables a SATA drive bracket, with the same form factor as the FC drive bracket, to be connected to the Sun StorageTek 6540 array controller enclosure and 4 Gb/s expansion tray midplane.
- **Usability** — The SATA interface card provides the SATA drives with basic FC-like functionality not otherwise available, enabling them to be managed like FC drives. These drive services include command queuing, a drive activity LED (external to the drive), unique worldwide name, drive signing, controlled drive spin up and down, and redundant drive loop diagnostics.

## Battery

The controller tray includes a lithium ion battery for long battery life. The battery requires no scheduled maintenance during the 3 year warranty.


# Configuration

Additional hallmark features of the Sun StorageTek 6540 array are flexibility and pay-as-you-grow scalability. Its unequalled flexibility includes the ability to mix drive types (FC and SATA, capacities, and/or RPM speeds), RAID levels, segment sizes, virtual disk sizes, volume characteristics, and cache policies all within a single storage system. The Sun StorageTek 6540 array allows dynamic re-configuration of the following parameters: capacity expansion, volume expansion, RAID migration, segment size migration, virtual disk re-configuration, and firmware upgrades. This configuration flexibility enables a single storage system to support multiple, heterogeneous servers with vastly different performance and/or capacity demands and meet varying workloads. And, the Sun StorageTek 6540 array delivers scaling flexibility as it can provide online expansion from 365 GB to over 112 TB — one drive at a time or a fully-populated 16-drive expansion tray at a time.

## Cabling

Each Sun StorageTek 6540 array controller tray and expansion tray fiber port and all fiber interface cables support full-duplex serial communications, meaning data is transmitted in two directions simultaneously. Therefore, a single cable from one component to another creates a complete fiber loop.

Drive loops are configured as redundant pairs utilizing one port from each controller, as illustrated in Figure 7. This helps ensure data access in the event of a path/loop or controller failure. Both ports of the channel run at the same speed.



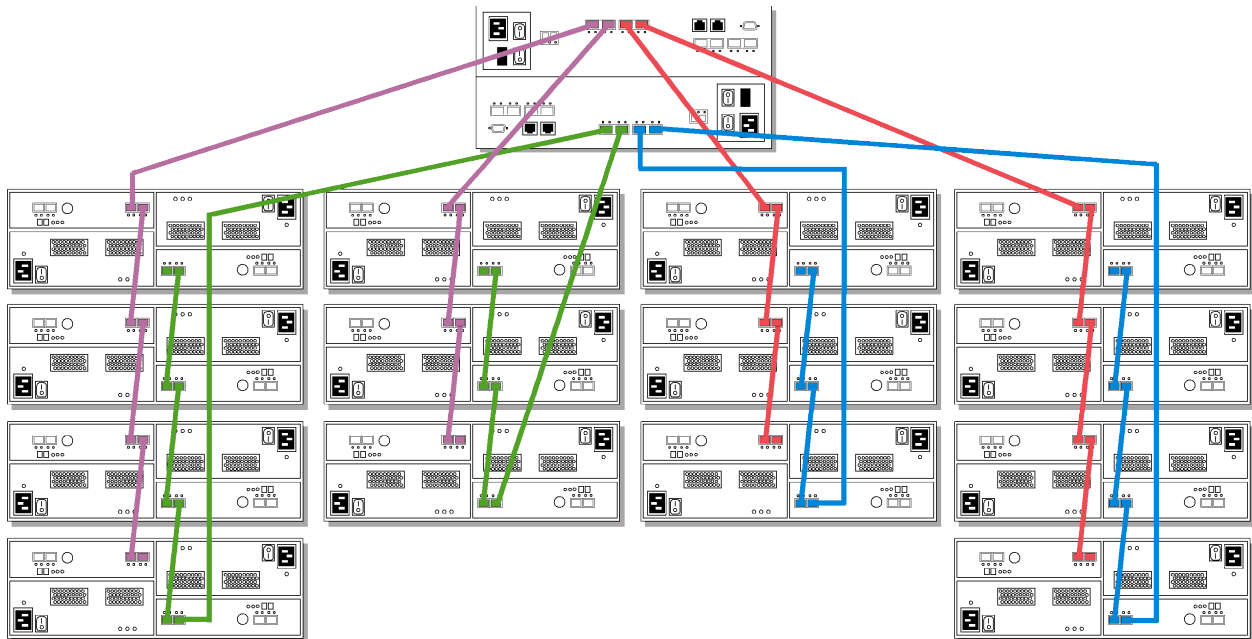
## Drive Connections

**Figure 7.** Sun StorageTek 6540 Array Loop Pairs

Expansion trays have redundant, dual-active drive loops running to the dual-ported FC disk drives or SATA disk drives. Each 4 Gb/s drive loop comes *in* through one port on the I/O module and goes *out* through the other — if additional expansion trays are attached.

Expansion trays are configured with redundant, dual-active loops from the Sun StorageTek 6540 array controller tray. This again helps to ensure data access in the event of a path loop or power failure. Each FC RAID controller card has access to all four drive channels for the highest performance and availability.

The Sun StorageTek 6540 array controller trays are configured with a *top-down/bottom-up* cabling scheme. This helps to ensure data access to surviving drive modules even in the unlikely event that a full drive module is unavailable. The diagram in Figure 8 shows a maximum configuration of 16 drive modules cabled utilizing all four drive loop pairs (eight total loops) and top-down/bottom-up cabling for full redundancy.



**Figure 8.** Sun StorageTek 6540 Array Full Configuration

All controller tray optical SFP and I/O Modules copper interconnect connections are designed with port bypass technology. Port bypass technology maintains fiber loop integrity when new FC connections are made — or if existing connections are removed — by automatically opening and closing ports as needed. This means a second host can be added, or SAN connections can be re-configured, without affecting I/O of the existing attached host. And, full drive modules can be added or removed without interrupting access to data.

HotScale technology combines the port bypass features of the Sun StorageTek 6540 array with Sun StorageTek Common Array Manager software's online administrative capabilities to deliver unsurpassed storage system availability. HotScale technology enables storage capacity to be added, configured, and assigned to a host without system downtime, without interrupting existing I/O, and without requiring a server or storage re-boot.

## Configuration Flexibility

The Sun StorageTek Common Array Manager software's unequalled configuration flexibility enables customized performance tuning, maximum capacity utilization, and the highest data protection. With Sun StorageTek Common Array Manager software, a single Sun StorageTek 6540 array can concurrently support multiple RAID levels (0, 1, 3, 5), mixing drive technologies, capacities and rotational speeds, multiple and various virtual disk sizes (up to 30 drives), and one or more volumes per virtual disk. Additionally, each volume supports individualized configuration settings, including controller ownership,

segment size, modification priority, and all cache policies. And, all options can be dynamically configured and re-configured to help ensure that data is always accessible.

This industry-leading flexibility enables the Sun StorageTek Common Array Manager software to best match application data availability, performance, or capacity requirements. This is especially important in consolidated environments where multiple servers, each with varying and potentially different performance demands, are sharing a single storage system. While every environment is different and has different priorities in regards to performance, data availability, and capacity utilization, it is possible to outline some best practice guidelines centered on drive types and RAID levels.

Disk drive type — every Sun StorageTek 6540 array supports both fat (high capacity) and fast (high rotational speed) disk drives. The Sun StorageTek Common Array Manager software enables these different drive types to concurrently reside in a single storage system enabling high-capacity and high-performance applications to co-exist in the same installation.

Applications such as databases, OLTP, messaging servers, and Web servers with random, small-block I/O requiring high IOPS performance are best served by a large number of high-spin drives. The Sun StorageTek 6540 array has two fast drive offerings:

- 73 GB 15K RPM — FC
- 146 GB 15K RPM — FC

Applications such as video streaming, multi-media, seismic processing, and high performance computing with sequential, large-block I/O requiring high throughput (MB/s) are satisfied with a smaller number of high-capacity disk drives. The Sun StorageTek 6540 array has three fat drive offerings:

- 146 GB 10K RPM — FC
- 300 GB 10K RPM — FC
- 500 GB 7.2K RPM — SATA

Selecting RAID configurations is equally critical for each application to meet desired data availability, performance, or capacity requirements. And, like drive types and configuration settings, the Sun StorageTek Common Array Manager software supports intermixing RAID configurations to provide maximum flexibility and utilization. RAID 5 is generally considered the best balance of cost, performance, and availability. The table below outlines the trade-offs of the various supported RAID configurations.

	RAID-0	RAID-1 and 1+0	RAID-3	RAID-5
<b>Description</b>	Data is striped across multiple drives.	Data is mirrored to another drive.	Data is distributed across multiple drives. Parity information is written to one disk in the group.	Drives operated independently with data and parity blocks distributed across all drives in the group.
<b>Minimum Number of Drives</b>	1	2	3	3
<b>Maximum Number of Drives</b>	30	30	30	30
<b>Usable capacity as a percentage of raw capacity</b>	100%	50%	66.67% to 96.67%	66.67% to 96.67%
<b>Application</b>	IOPS   MB/s	IOPS	MB/s	IOPS   MB/s

	<b>RAID-0</b>	<b>RAID-1 and 1+0</b>	<b>RAID-3</b>	<b>RAID-5</b>
<b>Advantages</b>	Performance due to parallel operation of the access.	Performance as multiple requests can be fulfilled simultaneously.	High performance for large, sequentially accessed files. Parity utilizes small portion of raw capacity.	Good for reads, small IOPS, many concurrent IOPS and random I/Os. Parity utilizes small portion of raw capacity.
<b>Disadvantages</b>	No redundancy. One drive fails, data is lost.	Storage costs are doubled	Degraded performance with 8-9 I/O threads, random IOPS, smaller more numerous IOPS.	Writes are particularly demanding.

## Management Software

The Sun StorageTek Common Array Manager software is the primary interface for configuring and managing the array. The management software consists of a suite of tools that are installed on an external management host. The management software runs on the following platforms:

- Windows XP Professional SP2
- Solaris 10 OS, x86 platforms
- Windows Server 2003, SP1 Enterprise Edition
- Windows 2000 Server and Advanced Server SP4
- Solaris 8, 9, or 10, SPARC platforms
- Windows Server 2003, SP1 Standard Edition
- Solaris 10 OS, AMD Opteron™ platforms

The management software enables the storage administrator to manage the Sun StorageTek 6540 array from any system with a Web browser that is on the same network as the management host. The supported browsers are:

- Netscape Navigator™ 7.0
- Microsoft Internet Explorer 5.0
- Mozilla™ 1.2.1

For latest supported browsers, see the *Sun StorageTek 6540 Array Release Notes*.

The Sun StorageTek Common Array Manager software provisions and maintains the storage for data hosts:

- Common arrays are collections of volumes that share a profile defining the common configuration of the volumes. Storage profiles define the characteristics of a common array. The administrator can choose one from the set of pre-configured profiles or create a new one.
- Volumes are divisions of a common array, consisting of virtual disks, representing the storage space that is used by the data hosts in the environment.
- Virtual disks, also called RAID sets, are collections of locations in the memory of more than one physical disk. The storage array handles a virtual disk as if it were an actual disk.
- Host groups are a collection of hosts that share access to the same volumes.
- Snapshots are copies of the data in a volume at a specific moment. Snapshots offer a high-availability alternative to backups because the the array can remain online to create the snapshot, and the snapshots take less less space than the original data.
- Volume copies are copies of the complete contents of one data volume that are located on another data volume on the same array.

## Volumes

The Sun StorageTek 6540 array's physical disks are managed as a pool of storage space for creating volumes. A volume is a logical structure on a storage system and is created by slicing a virtual disk into a stripe set with a defined capacity. Volumes are containers into which applications, databases, and file systems can put data. Volumes are created from the virtual disks, based on the characteristics of the common array associated with the virtual disks. Based on user's specifications, the array automatically allocates a virtual disk that can satisfy the volume configuration requirements.

With the Sun StorageTek Common Array Manager software, each virtual disk supports up to 256 volumes, with a maximum of 1024 volumes per Sun StorageTek 6540 array. During or after standard volume creation, a host or host group can be mapped to the volume in order to give the host or host group read and write privileges to the volume. Each host, including any host that is a member of a host group, must be assigned one or more initiators before the host or host group can be mapped to the volume.

There are several different types of volumes:

- *Standard volume* — A standard volume is a logical structure created on a storage array for data storage. When a volume is created, initially it is a standard volume. Standard volumes are the typical volumes that users access from data hosts.
- *Source volume* — A standard volume becomes a source volume when it participates in a volume copy operation as the source of the data to be copied to a target volume. The source and target volumes maintain their association through a copy pair. When the copy pair is removed, the source volume reverts back to a standard volume.
- *Target volume* — A standard volume becomes a target volume when it participates in a volume copy operation as the recipient of the data from a source volume. The source and target volumes maintain their association through a copy pair. When the copy pair is removed, the target volume reverts back to a standard volume.
- *Snapshot volume* — A snapshot volume is a point-in-time image of a standard volume. The management software creates a snapshot volume when the snapshot feature is utilized. The standard volume on which a snapshot is based is also known as the base or primary volume.
- *Reserve volume* — A snapshot reserve volume is automatically created a snapshot is created. The reserve volume stores information about the data that has changed since the volume snapshot was created. When a snapshot is deleted, the management software also deletes its associated reserve volume.

A volume can be created on a virtual disk as long as the RAID level, the number of disks, and the disk type (FC or SATA) of the virtual disk matches the storage profile associated with the volume's pool. The virtual disk must also have enough capacity for the volume. The administrator chooses the method of determining which virtual disk is used to create the volume. The following options are available:

- *Automatic* — The Sun StorageTek Common Array Manager software automatically searches for and selects a virtual disk that matches the necessary criteria. If none are available, it creates a new virtual disk if enough space is available.
- *Create volume on an existing virtual disk* — The administrator manually selects the virtual disks on which to create the volume from the list of all available virtual disks. It is important to ensure that the number of disks selected have enough capacity for the volume.
- *Create a new virtual disk* — The administrator creates a new virtual disk on which to create the volume. It is important to make sure that the number of disks selected have enough capacity for the volume.

A volume can be added to an existing storage domain, including the default storage domain, or a new volume can be created by mapping the volume to a host or host group. A storage domain is a logical entity used to partition storage that allows a host or host group to have read and write access to the

volume. The default storage domain contains all hosts and host groups without explicit mappings and enables them to share access to all volumes that are not explicitly mapped.

## Storage Profiles

A storage profile is set of attributes that can be applied to a common array to allocate storage, instead of having to set each attribute individually. The system has a predefined set of storage profiles. The administrator can choose a profile suitable for the application that is using the storage or create a custom profile. The Sun StorageTek 6540 array has a default storage profile with RAID-5 storage characteristics that is suitable for many storage applications. The default pool uses the default profile. If none of these profiles is suitable, a new storage profile can be created. Each storage profile has the following settings.

Parameter	Value or Variable Type	Description
<b>Name</b>	Up to 32 characters	Unique identifier for the storage profile.
<b>RAID Level</b>	0, 1, 3, 5 RAID level	Configured across all disks within a virtual disk.
<b>Segment Size</b>	8 KB, 16 KB, 32 KB, 64 KB, 128 KB, 256 KB, 512 KB	Segment size is the portion of a disk allocated to a virtual disk stripe.
<b>Readahead</b>	Enabled or Disabled	Read-ahead mode of the array. Cache read-ahead enables the controller to copy additional data blocks into cache while the controller reads and copies host requested data blocks from disk into cache.
<b>Disk Type</b>	ANY, FC, or SATA	Disk type.
<b>Number of Disks</b>	Up to 30 drives	The number of disks to be grouped together in a virtual disk. For example, if a common array is created with a profile that has the number of disks parameter set to a number, all virtual disks that are part of that common array must have the same number of disks. If the number of disks parameter is set to the Variable value, the administrator is prompted for the number of disks when storage is added to the pool.

## Virtual Disks

A virtual disk is a set of drives that the controller logically groups together to provide one or more volumes to an application host. Each virtual disk has an assigned RAID level, and typically all the drives in a virtual disk are the same capacity. Virtual disks can be created using drives of different capacities, however, the usable capacities of all the drives in that virtual disk are adjusted down to the disk drive of the smallest capacity.

The drives that make up a virtual disk can be selected manually or automatically, and can reside in a single drive module or across multiple drive modules. When using the automatic configuration option, virtual disks are configured to balance the load across as many drive channels and drive modules as possible. This helps ensure maximum protection and optimum performance.

While multiple RAID levels can be intermixed in a single Sun StorageTek 6540 array, each virtual disk has a single, assigned RAID level and can comprise of up to 30 drives. This RAID level can be dynamically migrated through the Sun StorageTek Common Array Manager software. The software also provides the ability to dynamically add additional drives, up to two at a time, to existing virtual disks. Existing volumes are then re-striped across all drives in the newly expanded virtual disk. And, like RAID level migration, this occurs while the virtual disk and all its volumes remain online and accessible for I/O.

Virtual disks are created and removed indirectly through the process of creating or deleting volumes or snapshots. RAID systems provide storage by making the data on many small disks readily available to



file servers, hosts, or the network as a single array. RAID systems use two or more drives in combination for fault tolerance and performance. One of the factors in data throughput and availability is how the data is stored within the array, that is, the array's RAID level. In the Sun StorageTek 6540 array, disk drives within a tray are grouped together into RAID sets, also called virtual disks, according to RAID level.

The Sun StorageTek 6540 array supports the following RAID levels

- **RAID-0** — Stripes data across multiple disks, but without redundancy. This improves performance but does not deliver fault tolerance.
- **RAID-1** — Mirrors a disk so that all data is copied to a separate disk.
- **RAID-3** — Stripes data at the byte level across multiple disks while writing the stripe parity to a parity disk. Provides high throughput for a single streamed file. Checks parity on reads.
- **RAID-5** — Stripes data at the byte level, and provides stripe error correction (parity checking) information. For this level, the minimum number of drives is three. RAID-5 results in excellent performance and good fault tolerance. Parity checking specifies that when the RAID controller writes information onto disks, it also writes redundant information, called parity bits. If a disk fails, the parity information enables the RAID controller to re-create the lost information as it is requested. Since the parity information is spread across multiple disks, only a percentage of the disks is used for parity information, which improves the efficiency of available storage space.

## Storage Pools

A storage environment can be divided into storage pools. Each pool is associated to a profile that allows the storage pool to meet specific input/output (I/O) requirements. Each system has a default pool that uses the default profile, which implements RAID-5 storage characteristics. The default pool satisfies most common storage requirements. The system also provides a set of storage profiles that satisfy certain I/O requirements that are optimal for the type of application to which they refer. If none of the factory profiles are suitable for an application's needs, a custom storage profile can be created. When a new storage pool is created, a specific profile is assigned to it. Removing a storage pool destroys all stored data in the pool and removes all volumes that are members of the pool. The data can be restored from backup after new storage pools are added, but it is far easier to avoid the difficulty in the first place.

## Remote CLI Client

The Sun StorageTek 6540 array can also be configured and managed using the remote CLI client. The CLI provides the same control and monitoring capability as the Web browser, and it is also scriptable for running frequently performed tasks. The remote CLI client is available for the Solaris OS, Windows, and several other operating systems. See the *Sun StorageTek 6540 Array Release Notes* for details on the supported operating system platforms.

## Configuration Metadata

The Sun StorageTek 6540 array controllers store configuration metadata in a private 512 MB region on every configured drive. This private metadata area contains drive state and status information, volume state and status information, and controller and subsystem information. In addition, it also stores a drive's state and status, the worldwide name (WWN) of its virtual disk, the volumes it contains, and the definitions for those volumes. Additionally, one drive in each virtual disk (with a minimum of three in each storage system) stores the controller- and subsystem-level information.

Storing metadata in the nonvolatile private metadata area provides the highest availability and enables easier re-configurations and migrations. As system configuration data resides on every configured drive, controllers and/or multiple drives can be removed or swapped without losing the system configuration. Drives can be relocated within the storage system to improve channel utilization and protection or even

migrated as a complete virtual disk into another storage system. And, in both instances, all configuration metadata and user data remains intact on the drives.

## Online Administration

The Sun StorageTek 6540 array enables nearly all storage management tasks to be performed while the storage remains online with complete read/write data access. This allows storage administrators to make configuration changes, conduct maintenance, or expand the storage capacity without disrupting I/O to its attached hosts. The online capabilities include:

- **Dynamic Capacity Expansion (DCE)** — DCE enables the capacity and number of drives to be increased for a given virtual disk. Up to two drives at a time can be added to existing virtual disks, with the existing volumes then being striped across the full set of drives now comprising the virtual disk.
- **Dynamic Volume Expansion (DVE)** — This feature allows the capacity of an existing volume to be expanded by using the free capacity on an existing virtual disk. **Note:** Increasing the capacity of a standard volume is only supported on certain operating systems.
- **Dynamic RAID level migration (DRM)** — Changes the RAID level of a virtual disk. **Note:** The virtual disk must have enough capacity and the appropriate number of drives to support the new RAID level.
- **Dynamic segment size migration (DSS)** — A default segment size is set during volume creation, based on the virtual disk RAID level and the volume usage specified. These two parameters should optimize the segment size appropriately for the environment. When monitoring the storage system reflects less than optimal performance, however, the segment size of a given volume can be changed.
- **Dynamic defragmentation** — This feature rearranges volumes and consolidates free capacity within a virtual disk, resulting in optimized access patterns for existing and newly created volumes.
- **Virtual Disk/Volume configuration** — HotScale technology enables new drive modules to be added to the Sun StorageTek 6540 array, virtual disks to be configured, and volumes to be created without disrupting access to existing data. Once a newly created volume is defined, Sun StorageTek Common Array Manager software's immediate LUN availability and *hot add* features enable it to instantly be mapped and accessed by a host(s).
- **Performance monitoring** — The Sun StorageTek Common Array Manager software for the Sun StorageTek 6540 array enables the storage system performance to be monitored in real-time to assist in performance tuning decisions.
- **Non-disruptive controller firmware upgrades** — The controller firmware can be upgraded with no interruption to data access. The new firmware is transferred to the first controller, its flash memory is updated, and the controller is rebooted. The firmware is then passed from the first controller to the second, and the process is repeated. With a multi-path driver installed, access to storage system volumes is maintained throughout the process.

## Performance Tracking Utility

The Sun StorageTek Common Array Manager software provides quick visualize performance health statistics, which allows fine tuning of attributes to reach the optional configuration for precise application needs. Since performance degradation is often a precursor to component or system downtime, performance tracking enables predictive, preemptive maintenance. The Performance Monitor data table tracks the following parameters by device (Sun StorageTek 6540 array controllers, volumes, and storage system totals):

<b>Performance Metric</b>	<b>Definition</b>
<b>Total IOPS (I/Os per second)</b>	The total number of I/Os handled by the component per second
<b>Run Average Total IOPS</b>	The average total number of I/Os handled by the component over the entire performance sampling period
<b>Peak Total IOPS</b>	Highest total IOPS over polling period
<b>Read Percentage</b>	The percentage of read operations handled by the component between two sampling periods (T1, T2)
<b>Write Percentage</b>	The percentage of write operations handled by the component between two sampling periods (T1, T2)
<b>Cache Read Hit Percentage</b>	The percentage of cache read hit operations handled by the component between two sampling periods (T1, T2)
<b>Total Data Transferred (KB/s)</b>	Total data transferred by component (reads/writes)
<b>KB/s Read</b>	Total data read in KB per second
<b>Run Average KB/s Read</b>	The average total data read in KB per second by the component over the entire performance sampling period
<b>Peak KB/s Read</b>	Highest KB/s read over polling period
<b>KB/s Written</b>	Total data written in KB per second
<b>Run Average KB/s Written</b>	The average total data written in KB per second by the component over the entire performance sampling period
<b>Peak KB/s Written</b>	Highest KB/s written over polling period
<b>Average Read Size (Bytes)</b>	The average read size for a component for a single sampling period
<b>Average Write Size (Bytes)</b>	The average write size for a component for a single sampling period

## Global Hot Spares

If a drive fails in the Sun StorageTek 6540 array, the controller uses redundancy data to reconstruct the data onto a hot spare drive. The hot spare is automatically substituted for the failed drive without requiring user intervention. And, when the failed drive is physically replaced, the data from the hot spare is automatically copied back to the replacement drive. The system supports up to 15 global hot spare drives, and each can be a spare for any disk in the array.

A hot-spare is a drive, containing no data, that acts as a standby in the storage array in case a drive fails in a RAID-1, RAID-3, or RAID-5 volume. The hot-spare adds another level of redundancy to the storage array. Generally, the drive assigned as a hot spare should have a capacity that is equal to or greater than the capacity of the largest drive on the storage array. If a hot spare is available when a drive fails, the hot spare is automatically substituted for the failed drive, without user intervention. Upon physical replacement of the failed drive, the data from the hot spare is copied back to the replacement drive. This is called copyback. If a hot spare drive is not available, a failed disk drive can still be replaced while the storage array is operating. If the drive is part of a RAID-1, RAID-3, or RAID-5 virtual disk, the controller uses redundancy data to automatically reconstruct the data onto the replacement drive. This is called reconstruction.

## Multipathing

With multipathing, also called multipath failover, an array or network can detect when an adapter has failed and automatically switch access to an alternate adapter. Multipathing enables high-availability configuration because it ensures that the data path remains active. Multipathing also helps increase performance to multi-controller disk arrays by spreading I/O between multiple paths into the array.

Within the array, common arrays use multipathing by default. To complete the data path, data hosts also need the ability to multipath. Therefore, all data hosts need one of the following software products:

- Sun StorageTek™ Traffic Manager software,
- MPxIO, part of the Sun StorageTek SAN Foundation software and integrated into the Solaris 10 OS
- VERITAS Volume Manager with Dynamic Multipathing (DMP)
- RDAC is available for Windows 2000, Windows 2003 and Linux on the Sun Download Center (<http://www.sun.com/download/>)
- RDAC and DPF drivers are also available for AIX through a Get To Yes review
- See the online Interoperability Tool (<https://extranet.storitek.com/interop/interop>) for the latest information.

## Investment Protection / Data-In-place Migration

To maximize investments in the Sun StorageTek Modular Line, the CSM200 drive trays can be utilized between the Sun StorageTek 6140, 6540 arrays, and 6940 systems. This flexibility not only allows the same drive trays to be redeployed as enterprises move up within the Sun StorageTek Modular Line, in most cases, even the data can be preserved during the migration.

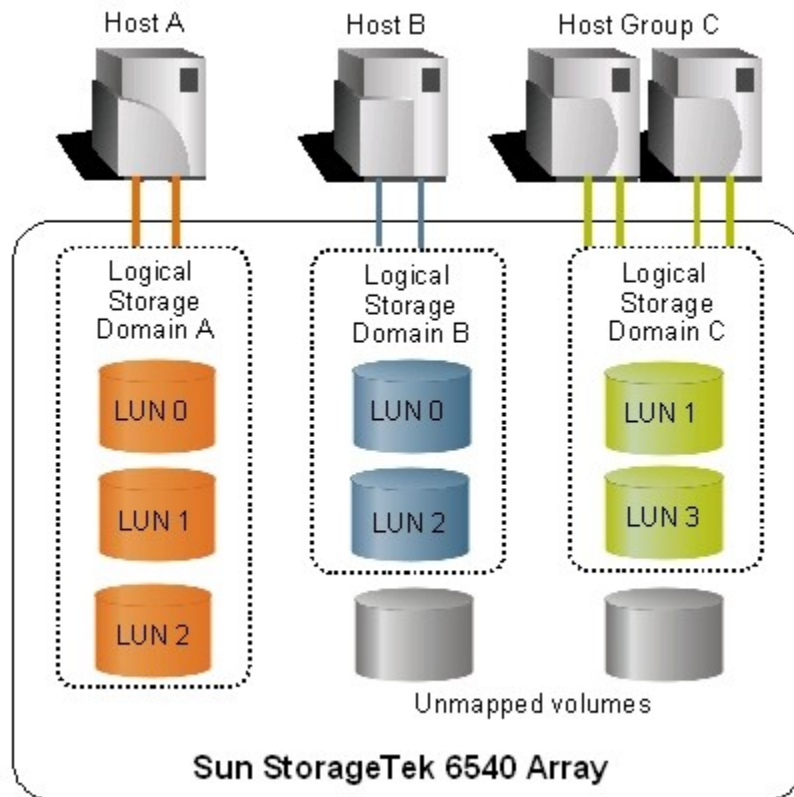
1. CSM200 expansion modules can be used across Sun StorageTek 6140, 6540 arrays, and 6940 systems
2. CSM200 expansion modules are expected to be supported behind the Sun StorageTek™ 6130 array (Q3CY06)
3. CSM200 expansion modules are expected to be supported behind the StorageTek Flexline FLX380 system (Q3CY06)
4. CSM200 expansion modules are expected to be supported behind the StorageTek™ Flexline FLX280/FLX240 arrays (Q3CY06)
5. CSM100 expansion trays are expected to be supported behind the Sun StorageTek 6540 array (Q3CY06)
6. FLA200/FLC200/FLA300 expansion modules are expected to be supported behind the Sun StorageTek 6140 and 6540 arrays (Q3CY06)

## Storage Domains — Optional

The Sun StorageTek Common Array Manager software's Storage Domain feature enables a single Sun StorageTek 6540 array to be logically partitioned and function as up to 64 virtual arrays. A storage domain is a logical entity consisting of one or more storage system volumes that are accessed by a single host or shared among hosts that are part of a host group. A storage partition is created when the administrator defines a single host or a collection of hosts, called a host group, and then defines a volume-to-LUN mapping. This mapping defines what host or host group have access to a particular volume in the storage system. Hosts and host groups can only access data through assigned volume-to-LUN mappings. Partition access is maintained at the controller level, for complete data integrity in multi-host, multi-OS environments.

Volume-to-LUN mapping creates valuable flexibility for the storage administrator as any available volume can be mapped to any attached server. So, while the individual servers see a virtual array that consists of only their LUNs/volumes, the volumes can be intermixed throughout the storage system within one or more virtual disks. Logical partitioning combined with the Sun StorageTek Common Array Manager software's configuration capabilities enables administrators to choose from a range of volumes with different characteristics to meet a server's exact needs for a given LUN. Each volume can

have unique configuration settings and reside on different drive types with different RAID levels. This flexibility enables a range of hosts with different capacity, performance, or data protection demands to effectively share a single Sun StorageTek 6540 array. All storage domain licenses are available as optional upgrades.



**Figure 9:** Storage Domains

## Heterogeneous Hosts

The heterogeneous hosts feature allows the firmware on each controller in the storage system to tailor its behavior (such as LUN reporting and error conditions) to the needs of the host operating system. This provides each individual host the view of the storage system that it would experience if it had exclusive access to the array.

In a heterogeneous environment, each host type must be set to the appropriate operating system during host port definition. A host type can be a completely different operating systems, such as the Solaris OS and Windows 2000, or variants of the same operating system, e.g., clustered and non-clustered. Each storage host port can be configured for multiple host types enabling complete flexibility for heterogeneous consolidation. Heterogeneous host settings are only available when storage domains are enabled.

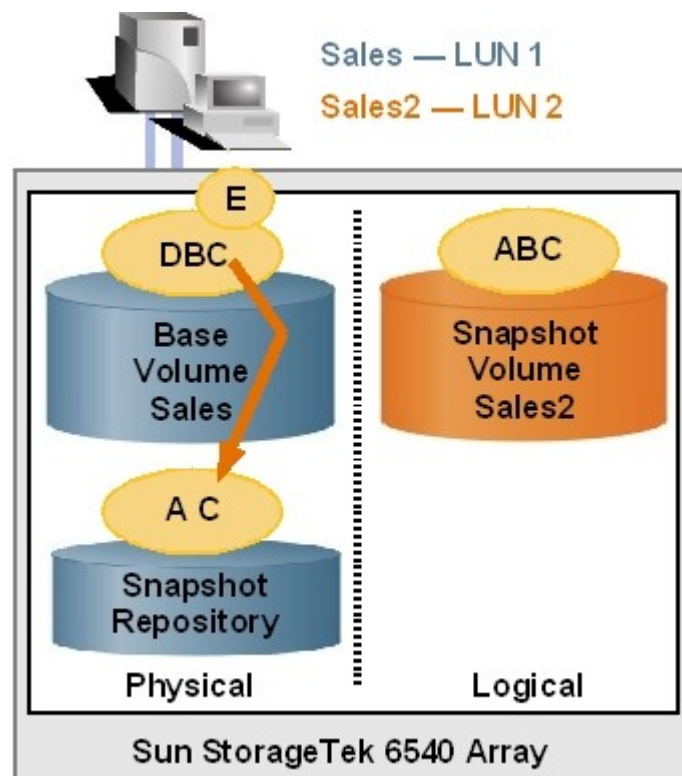
## Sun StorageTek Data Snapshot Software — Optional

Sun StorageTek Data Snapshot software provides an additional level of data protection and the means to improve production data utilization. It enables non-production servers to access an up-to-date copy of production data for a variety of applications — including backup, application testing, or data mining — while the production data remains online and user-accessible.

A snapshot volume is a point-in-time image of a volume that provides an immediate, space-efficient, logical copy of a volume. The snapshot volume appears and functions as a standard volume, enabling it to be used as a backup source, restoration point, information analysis / manipulation base, development environment, and for other applications that benefit from a volume copy. Snapshots are created instantaneously and require only a small amount of disk space. Snapshot volumes appear and function as normal volumes. A snapshot volume is a point-in-time (PiT) image of a volume with all of the read, write, and copy capabilities of the original volume.

Sun StorageTek Data Snapshot software uses an innovative copy-on-first-write technology to maintain the logical snapshot volume while minimizing disk utilization. When the snapshot is *taken* the controller suspends I/O to the base volume for a few seconds while it creates a physical volume — called the repository volume — to store snapshot metadata and copy-on-first-write data. The first time a data block is written to after the snapshot is taken, the old data is copied to a reserved area called a repository. Pointers into the repository are maintained and used to present the snapshot as a virtual LUN. This repository volume combined with the original base volume creates the logical snapshot volume.

Since the only data blocks that are physically stored in the repository volume are those that have changed since the time of the snapshot, the snapshot technology uses less disk space than a full physical copy. The repository volume is typically 20 percent of the base volume, but varies depending on the amount of changes to the data. The longer a snapshot is active, the larger the repository is needed. The Sun StorageTek Data Snapshot software provides notification when the repository volume nears a user-specified threshold (a percentage of its full capacity). And at any time, the Sun StorageTek Common Array Manager software's Dynamic Volume Expansion feature can be used to dynamically expand the repository volume. Sun StorageTek Data Snapshot software allows up to four snaps per volume and up to 1024 copies per Sun StorageTek 6540 array.



**Figure 10.** Sun StorageTek Data Snapshot Software

## Key Features and Benefits

- **Up to four snapshots per volume**

- ☞ Better protection — Support for multiple copies of data volumes allows frequent and regular snapshots to be used to protect against data loss from an operational problem.
- ☞ Protects data by providing up to four readily available online copies that reduces restore time.

- **Rapid availability**

- ☞ Snapshot creation within seconds — Sun StorageTek Data Snapshot software can create a snapshot volume in seconds, which avoids the lengthy time required to do a full disk-to-disk copy or tape backup and then restore. Fast creation of snapshots means less waiting to use them.
- ☞ Improves employee productivity by having an immediate copy. No more waiting for large volumes of data to copy, snapshot is nearly instantaneous.

- **Read and write support for snapshots**

- ☞ Support all application needs — read and write access to snapshots means that virtually any type of testing or analysis can be applied to using snapshots without jeopardizing primary production data. While the snapshot version of data can be modified by a secondary application, the primary data volume continues to be used and modified by the original application server.
- ☞ Provides more rapid application development by immediately creating a test environment and capitalizing on the ability to write to the snapshot image.

- **Space saving design**

- ☞ Maintain multiple copies with minimal extra disk expense — the copy-on-first-write design allows multiple versions of data to be protected with minimal disk space consumption. As a result, Sun StorageTek Data Snapshot software may require only 5 to 100 percent of additional capacity for four snapshots instead of the 400 percent (for four copies) of additional capacity that would be required with full volume copies.

## Sun StorageTek Integration Module for VERITAS NetBackup With Oracle® — Optional

Sun StorageTek Integration Module for Netbackup with Oracle is a solutions guide comprised of a media kit with CD and hard copy documentation that includes scripts and guidelines for integrating Sun StorageTek Data Snapshot software with VERITAS NetBackup and Oracle 10g for the Sun StorageTek 6140 and 6540 arrays. Although the backup software market is a mature market space, clearly backup implementations continue to be a long-standing, serious problem for enterprises today. More specifically, administrators attempting to protect and recover data with traditional backup methods are meeting several challenges that include ensuring recoveries are performed expeditiously in light of massive data volumes, growing data volumes that produce longer backup windows, and longer backup windows creating less application availability.

Sun StorageTek Integration Module for VERITAS NetBackup with Oracle solves problems such as application availability due to expanding backup windows, speed of recovery (especially of mission critical applications) from crashes, errors, or disasters, and low productivity of the system administrator. Solving these problems results in greater application availability, a more efficient backup and recovery solution, and a higher state of error or disaster recovery preparedness.

## Sun StorageTek Data Volume Copy Software — Optional


Sun StorageTek Data Volume Copy software provides administrators with another tool to effectively manage information growth and maximize the utilization of Sun StorageTek 6540 arrays. It creates a complete physical copy, or a clone, of a volume within a storage system. The clone volume is a unique



entity that can be assigned to any host and used by applications requiring a point-in-time (PiT) copy of production data — such as backup, application testing or development, information analysis, or data mining — without affecting the performance of the production volume.

The software is configured and accessed via easy-to-use wizards or command line interface. It supports up to eight concurrent copies within a single storage system. Sun StorageTek Data Volume Copy software is a background operation with five user-defined priority settings, enabling administrators to minimize either copy time or the overall I/O impact to the storage system. And as the software is controller-based, it requires no host interaction or CPU cycles, minimizing the impact to applications and the storage infrastructure.

Sun StorageTek Data Volume Copy software can be used in conjunction with Sun StorageTek Data Snapshot software — which creates a PIT image of a volume while maintaining read and write access — enabling a complete PIT clone to be created without interrupting the I/O activity of the production volume. Additionally, Sun StorageTek Data Volume Copy software can be used to redistribute data — moving volumes from older, slower disk drives to newer, faster, or higher capacity drives — to optimize application performance and/or capacity utilization.



**Figure 11:** Data Volume Copy

## Key Features and Benefits

- **Redistribute data**

- ☞ Move volumes from older, slower disk drives to faster, higher capacity disk drives to optimize application performance and capacity utilization.

- **Support up to 8 active copies per volume**

- ☞ Each clone volume is an independent entity that can be assigned to any host and used by any application requiring a point-in-time complete copy of data such as backup, application testing and development, or information analysis and data mining without affecting production volume performance.



- **User -defined priority settings**

☞ Sun StorageTek Data Volume Copy software is a background operation with five priority settings enabling administrators to minimize either copy time or I/O impact to the system as business conditions and requirements dictate.

- **Controller-based**

☞ Sun StorageTek Data Volume Copy software requires no host interaction or CPU cycles – minimizing impact to applications and the storage infrastructure.

## **Target Applications for Both Sun StorageTek Data Snapshot and Data Volume Copy Software**

Sun StorageTek Data Snapshot and Data Volume Copy software enable companies to improve business operations by copying data sets to allow them to run more applications in parallel and create more frequent recovery points. The result can be an improved return on information. The value realized depends on how they are using it and is summarized in the following example applications.

- *Reducing Backup Windows* — Backups are a critical aspect of any business continuity strategy. However, backup windows are shrinking while the amount of information that needs to be protected is rapidly expanding. Backing up snapshots and/or clones, rather than online data, enables critical online transactions to keep running during the backup process.
- *New Application Development and Testing* — Many enterprise systems are being re-engineered today to respond more efficiently to the networked business environment. Sun StorageTek Data Snapshot and Data Volume Copy software can enhance these efforts by enabling business applications to be tested on real, up-to-date data. This technology helps to accelerate new application development and helps to promote faster time-to-market.
- *Analyzing Up-to-Date Information* — Businesses are constantly seeking new approaches to gain a competitive marketplace advantage. The faster they can convert data to information, the quicker they can make critical decisions. Sun StorageTek Data Snapshot and Data Volume Copy software can be used to enable businesses to analyze more current information, leading to faster and better business decisions. One way to frequently analyze up-to-date information is to take a snapshot of an online database and then use that snapshot as the data set for a decision support application.
- *Reducing the Window for Data Loss* — Sun StorageTek Data Snapshot and Data Volume Copy software can be used to reduce the window for data loss in the event of a database or file system corruption. For example, while a typical backup policy might provide for a nightly backup, creating two snapshots a day can reduce the potential window for data loss from 24 hours to 12 hours. The lighter impact of snapshots can enable more frequent copies and result in a better level of data protection, availability, and recovery point objective (RPO).

## **Sun StorageTek Data Replicator Software — Optional**

The StorageTek Data Replicator software enables real-time synchronous or asynchronous data replication to either local campus or metro or remote data centers to protect a company's mission critical information. The StorageTek Data Replicator software provides volume-level replication, which enables volume replication between physically separate Sun StorageTek 6540 arrays in real time. Organizations can keep up-to-date copies of critical data at multiple sites, which is particularly useful in creating business continuity solutions and disaster planning scenarios.


A basic replication configuration consists of two peer sites (primary and secondary) and their associated volumes. The primary volume is the volume at the main or primary site that is being replicated while the volume at the destination site is referred to as the secondary volume. A replication set consists of a group of primary volumes and their associated secondary volumes at the remote site. A volume can only be replicated to a single remote site.

The software transports data between the two Sun StorageTek 6540 arrays by means of synchronous or asynchronous replication mode using a dedicated FC connection. The replication network used can be either FC or TCP/IP with conversion equipment or over public or private telecommunications infrastructures. These features provide the capability to replicate between sites located throughout the world and enables data to be written transparently to both primary and secondary sites either simultaneously (synchronous) or with a managed delay (asynchronous).

If there is a break in the network or if the secondary volume is unavailable, the software automatically switches to suspended mode, in which it ceases replication and tracks changes to the primary volume in a separate volume known as a replication repository. When the link is restored, the software uses the information in the replication repository volume to re-synchronize the volumes and returns to replication the data.

Sun StorageTek Data Replicator software supports write order consistency, where multiple volumes are combined to preserve write order, shown in Figure 10. Each Sun StorageTek 6540 array supports one consistency group. The value of write order consistency is that the remote volumes, replicated using asynchronous mode, are in a re-startable consistent state regardless of how far behind the data transfer to the remote volumes lags.

Sun StorageTek Data Replicator software can also be used to restore data from a secondary volume to a primary volume by reversing the roles of the primary and secondary volumes. Role reversal is a failover technique in which a primary volume failure causes the secondary volume to assume the role for the primary volume. The application software accesses the secondary volume directly until the failure at the primary volume is corrected. Once the primary volume has been brought online again, any data changes that occurred can be written back to the primary volume.



**Figure 12.** Sun StorageTek Data Replicator Software

## Key Features and Benefits

- **Enterprise-wide data protection**
  - ☞ Supports data center, campus, Metropolitan Area Network (MAN), and Wide Area Network (WAN) replication
- **Synchronous and asynchronous replication modes**
  - ☞ Synchronous provides zero data loss remote replication
  - ☞ Asynchronous enables economical and long distance remote replication
  - ☞ Supports synchronous and asynchronous modes over both IP and Fibre Channel links
  - ☞ Replication distances can range from building to building within a campus or between sites over long distances
- **Write-order consistency across volumes**
  - ☞ Preserves write transaction order across remote volumes
  - ☞ Protects against data corruption
  - ☞ Enables remote volumes to be immediately used as re-startable volumes in the event of a primary site failure
  - ☞ Replicated data can be used for application testing, data mining, or back-up to tape.
- **Role reversal**
  - ☞ Use role reversal to make the secondary site the new primary site
  - ☞ Allows a secondary site to be assigned as primary while the primary site is rolled back to a known, corruption-free point in time
- **Dynamic mode switching**
  - ☞ Ability to switch from synchronous to asynchronous replication without suspending the remotely replicated volumes accommodates changing application and bandwidth requirements without sacrificing protection
- **Many-to-one support**
  - ☞ Multiple remote systems can replicate to a single system for centralized data protection, mining, or backups

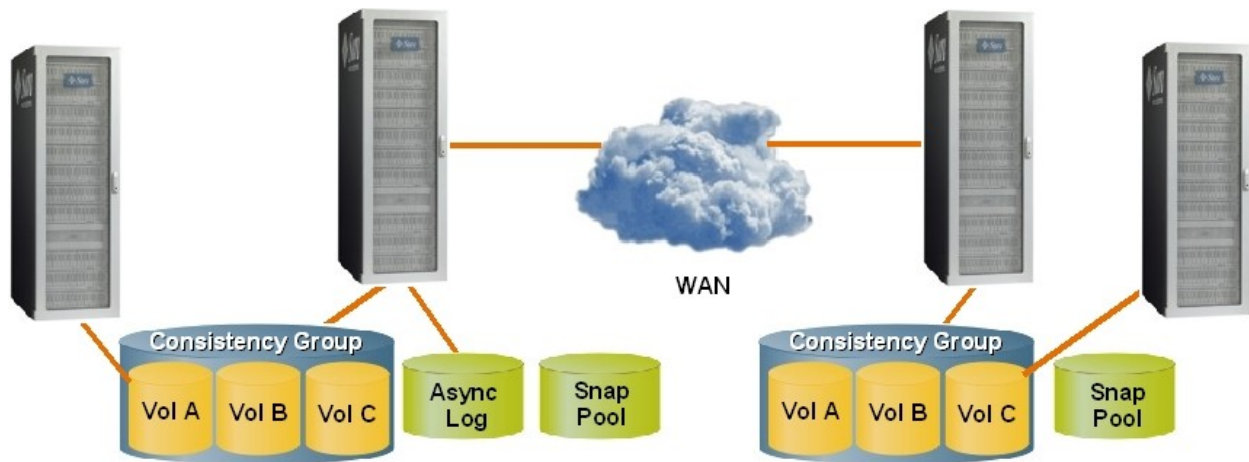
## Target Applications for Sun StorageTek Data Replicator Software

Sun StorageTek Data Replicator enables companies to improve their business operations by replicating data sets to a remote and physically separate locations. The result can be protected or improved business operations. The value realized is summarized in the following example applications.

- *Business and Data Continuanace* — Business/data continuance uses remote replication to make a secondary site available as a primary site in the event of a disaster or unplanned outage at the primary site. Both synchronous and asynchronous replication are used in business/data continuance implementations, selected by the distance between the replication sites and the amount of data loss an enterprise is willing to sustain in the event of an outage.
- *Content Distribution* — Many enterprises duplicate data from a central core site to another remote site for the purpose of updating the information repository at that site to reflect current changes to the data. An example would be product price lists that are distributed from a central corporate site to a geographically remote site on a daily basis.
- *New Application Development and Testing* — Many enterprise systems are being re-engineered today to respond more efficiently to the networked business environment . Sun StorageTek Data Replicator software can enhance these efforts by enabling business applications to be tested on real, up-to-date

data. This technology helps to accelerate new application development and helps to promote faster time-to-market.


- *Analyzing Up-to-Date Information* — Businesses constantly seek new approaches to gain a competitive marketplace advantage. The faster they can convert data to information, the quicker they can make critical decisions. Sun StorageTek Data Replicator software can be used to enable businesses to analyze more current information, leading to faster and better business decisions. One way to frequently analyze up-to-date information is to take a snapshot of an online database and then use that snapshot as the data set for a decision support application.



**Figure 13:** Data Replication

## Features, Functions, and Benefits

Feature	Function	Benefit
<b>Application-oriented management utility, including common array and storage profile capabilities</b>	<ul style="list-style-type: none"> <li>• Simplify storage provisioning.</li> </ul>	<ul style="list-style-type: none"> <li>• Quickly and predictably provision storage to precise application demands.</li> </ul>
<b>Online storage management, redundant hot-swappable components, and automated path failover</b>	<ul style="list-style-type: none"> <li>• Provide round-the-clock availability.</li> </ul>	<ul style="list-style-type: none"> <li>• Confidence that business-critical data is online and accessible all of the time .</li> </ul>
<b>Support for high-performance FC and high-capacity SATA disk drives</b>	<ul style="list-style-type: none"> <li>• Expansion trays with FC drives can be added to improve performance and increase capacity. For secondary storage requirements, expansion trays with SATA drives can be integrated to create a cost-effective solution.</li> </ul>	<ul style="list-style-type: none"> <li>• Enables a single Sun StorageTek 6540 array to satisfy primary and secondary storage requirements.</li> </ul>
<b>Fully-integrated data services</b>	<ul style="list-style-type: none"> <li>• Sun StorageTek Data Snapshot software enables administrators to create real-time copies of critical volumes, which can be used to eliminate the backup window and rapidly recover data.</li> <li>• Sun StorageTek Data Volume Copy software enables administrators to quickly and securely duplicate a volume for data mining, analysis, or enhanced backup operations.</li> <li>• Sun StorageTek Data Replicator software enables real-time synchronous or asynchronous data replication to either local campus, metro, or remote data centers to protect mission critical information.</li> </ul>	<ul style="list-style-type: none"> <li>• Maximize application availability and protection.</li> </ul>
<b>Switching technology (SBOD) in expansion storage modules</b>	<ul style="list-style-type: none"> <li>• Provides point-to-point connections of the drives in the expansion trays.</li> </ul>	<ul style="list-style-type: none"> <li>• Individual drive isolation enables improved diagnostics and faster problem resolution.</li> <li>• Fewer drive-loop nodes and faster communication lowers loop latency and provides linearly scalable performance for more responsive applications.</li> </ul>

Feature	Function	Benefit
<b>Modular design and pay-as-you-grow scalability</b>	<ul style="list-style-type: none"> <li>• Online capacity expansion.</li> </ul>	<ul style="list-style-type: none"> <li>• Lowers acquisition and expansion costs by enabling organizations to purchase what they need, when they need it.</li> <li>• Match the storage to the application.</li> <li>• Add storage without impacting data availability.</li> <li>• Create an affordable entry-point without sacrificing future scalability.</li> </ul>
<b>Robust, yet intuitive, storage management software</b>	<ul style="list-style-type: none"> <li>• Simplify storage management.</li> </ul>	<ul style="list-style-type: none"> <li>• Helps ensure maximum utilization of storage capacity and complete control over rapidly growing storage environments.</li> </ul>
<b>Storage domains</b>	<ul style="list-style-type: none"> <li>• Virtual arrays — segment storage capacity by host and applications.</li> </ul>	<ul style="list-style-type: none"> <li>• Lower total cost of ownership through storage consolidation in heterogeneous host environments.</li> </ul>
<b>Metadata is stored on every configured drive in the array</b>	<ul style="list-style-type: none"> <li>• Centralized administration from anywhere with Sun StorageTek Common Array Manager software.</li> </ul>	<ul style="list-style-type: none"> <li>• Data-intact drive portability supports changing requirements.</li> <li>• Redundant protection of critical configuration metadata.</li> <li>• Loss of a management station does not cause loss of any configuration data.</li> </ul>
<b>Online configuration and re-configuration capabilities</b>	<ul style="list-style-type: none"> <li>• Re-configure the array's attributes on-the-fly.</li> </ul>	<ul style="list-style-type: none"> <li>• Dynamically tune volume attributes to meet ever-changing requirements.</li> </ul>
<b>Dual 4 Gb/s cache FC RAID controller cards in a 4 RU chassis, eight 4 Gb/s FC Host ports per Sun StorageTek 6540 array controller tray</b>	<ul style="list-style-type: none"> <li>• Redundant, hotswappable components</li> </ul>	<ul style="list-style-type: none"> <li>• Robust design with massive scalability.</li> </ul>
<b>Up to 14 expansion trays per controller tray, scales from 5 drives to 2  drives (365 GB to 112 TB)</b>	<ul style="list-style-type: none"> <li>• Creates an affordable entry-point, avoids over-configuration without sacrificing future scalability.</li> </ul>	<ul style="list-style-type: none"> <li>• The Sun StorageTek 6450 system's modular, building block design enables lower acquisition and expansion costs while maintaining maximum flexibility.</li> </ul>
<b>Dynamic re-configuration functionality</b>	<ul style="list-style-type: none"> <li>• Dynamic volume expansion, RAID migration, capacity expansion, segment size migration, volume and virtual disk re-configuration, controller firmware upgrades.</li> </ul>	<ul style="list-style-type: none"> <li>• Round-the-clock availability.</li> </ul>

Feature	Function	Benefit
<b>Performance monitoring and tuning</b>	<ul style="list-style-type: none"> <li>• Helps optimize application performance.</li> </ul>	<ul style="list-style-type: none"> <li>• Minimize downtime through predictive, preemptive maintenance and automated diagnostics for proactive health monitoring.</li> </ul>
<b>Supports RAID 0, 1, 3, 5, and 1+0</b>	<ul style="list-style-type: none"> <li>• Ability to cost-effectively match performance and availability to application requirements.</li> </ul>	<ul style="list-style-type: none"> <li>• The Sun StorageTek 6540 array's extremely high-speed disk performance enables faster, more responsive applications that can improve transaction rates and customer satisfaction.</li> </ul>

# Reliability, Availability, and Serviceability (RAS)

---

## Reliability

Reliability features of the Sun StorageTek 6540 array include the following:

- No Single point of failure
- Passive midplane (except for FRU-ID EEPROM)
- Automatic sector reallocation on RAID controller
- I2C for components management and fault detection
- Path fail over (PFO) protection — coverage for I/O module pulls, hardware, or firmware failure
- Path fail back protection — PFB — after a PFO the ability to add a new I/O module back into the tray
- Global hot spare (up to 15)
- Link redundancy chip and 8- to 10-bit encoding on FC-AL loops
- ECC on data cache
- Hardware support for parity for RISC RAM
- Cooling fan module pro-actively adjusts fan speed based on temperature monitoring
- Temperature sensors located closer to heat spots to provide accurate temperature measurements
- RAID set / disk scrubbing
- Global hot spare (array or dedicated hot spare)
- I/O module loop back diagnostic
- I/O module FC loop down detection (backend loop down detection)
- I/O statistics

## Availability

Availability features of the Sun StorageTek 6540 array include the following:

- Dual-active, hot-swappable FC RAID controllers with cache mirroring
- Dual-active drive channels
- Hot-swappable disk drives
- Dual hot-swappable redundant 400W power supplies and fan modules with individual power cord
- Automatic drive failure detection and rebuild using global hot spare drives and cache battery backup
- All FRUs are hot swappable
- Write through / write back with mirroring
- Up to 7 day battery backup
- Fast on-the-fly background volume (RAID Set) initialization
- Dynamic re-configuration: virtual disk strip size, RAID migration, volume create/delete, defragmentation, online controller firmware upgrade, expansion trays addition
- Persistent group reservation (supporting industry-standard SCSI-3 persistent reserve commands)
- Array-based CLI, host-based GUI/CLI management software



## **Serviceability**

Serviceability features of the Sun StorageTek 6540 array include the following:

- Low FRU count — the backplane, FC RAID controllers, support modules with battery, cooling fan and power supplies — can be hot-swapped with no tools required
- Failed FRU ID
- Host-based phone home capability
- Auto disk recognition
- Advanced run time diagnostics
- Background media scan
- Quick snap locking mechanisms for easy insertion and extraction of disks and other FRUs

# Best Practices

---

## Planning and Acquisition

It is key to get a proper understanding of the applications and the expectations for them before the purchase of storage resources in order to best assure a satisfactory outcome. This is especially true of the Sun StorageTek 6540 array, which offers such broad configuration flexibility and adaptability. It is often not possible to gain a deep understanding of an application's requirements either because of severe budgetary limitations or other factors, such as management mandates for storage consolidation and/or the pre-assignment of redeployed hardware that *will just have to do the job*. In the latter cases, planning may not help much when there are few options. The former situation, however, does allow business and technical trade-offs to be made that can provide better results when the available resources are selected on the basis of what is needed for the application.

## Best Practices Require That the Application(s) Are Well Understood

Two typical classifications methods are based on the access patterns for the data: transaction processing and data streaming.

- Transaction processing is the interactive online communication with record-oriented data.
  - ☞ Sample applications might be order entry or an airline reservation system.
  - ☞ Transaction processing characteristically employs short records or *blocks*, asynchronously and randomly accessed.
  - ☞ Performance is measured in terms of IOPS.
- Data streaming reflects the continuous ingest, processing, or output of data.
  - ☞ Sample applications might be satellite data collection or media serving.
  - ☞ Data streaming more classically utilizes very large data blocks that are sequentially accessed.
  - ☞ Performance is measured in terms of Mbytes transferred per second, or MB/sec.
    - Some applications require a combination of the above two patterns, and the Sun StorageTek 6540 array is well suited to both transactional and streaming data patterns, as long as the storage system is configured to best address the performance characteristics of each.

Performance enhancement of intelligent RAID-based systems like the Sun StorageTek 6540 array is best achieved by taking advantage of parallel and overlapped operations that overall appear to exceed the basic I/O operating specifications of the actual hardware. This is normally achieved by taking advantage of dual controllers, each with large, flexible cache buffers, multiple independent channels, numerous drive spindles across which the data is striped, and dynamically managed I/O queues that utilize access optimization algorithms.

## How Many Drives are Required?

The design of the Sun StorageTek 6540 array provides the user with a very scalable storage system that can be predictably configured for most applications. This is because in addition to its explicit objective of maximum availability, the RAID architecture provides exceptional parallelism and extensive overlap of I/O operations, as described above, in addition to the use of intrinsically fast components. The Sun StorageTek 6540 array also boasts an extremely low access time to data. A transaction processing application designed around a relational database normally requires several random accesses for modest amounts of data with each access. Overall performance of this application is said to be IOPS bound, because input/output overhead is a function of transaction rate, but could possibly be enhanced through performance optimization techniques described in this section.

While numerous factors come into play, the most common is that spreading a database across more spindles results in faster access time. For planning purposes, knowing how many IOPS the target application generates today should help estimate the future IOPS requirement based on predicted business growth.

Although the typical 15K RPM Fibre Channel disk drive is capable of approximately 350 IOPS per spindle in a benchmark environment, 200 to 250 IOPS per spindle is a better sizing parameter for *real world* environments. This figure is based on SPC-1 benchmark testing ([www.storageperformance.org](http://www.storageperformance.org)). The optimized IOPS rate divided by 200 or 250 should provide a first cut estimate of the number of disk drives needed to achieve top target transactional performance.

This drive count should then be bumped up to address the additional requirements for parity, growth, and spare disks in a RAID system. 200 to 250 is merely an estimate: 15K RPM drives are about 15 percent faster than traditional 10K RPM drives, providing a higher IOPS rate.

Plan a storage solution around a key application. If the calculations result in a large number of drives that seem difficult to manage, just remember that the notion of buying the largest drive size because it offers more storage capacity can be false economy. If the application demands a maximum level of performance, select the higher spindle count but with lower capacities. (In those situations where one of several mechanisms for data/business protection is deployed, such as snapshot copy or remote volume mirroring, the incremental drive requirement for those functions should be considered separately.)

Other applications, such as data acquisition for signal processing are bandwidth-oriented, access very large blocks of data with sequential access pattern, rather than few random data accesses. IOPS are not as significant for these applications. Instead, they require continuous streaming of large volumes of data, referred to as bandwidth or throughput-oriented. Performance requirements in this case are measured in terms of MB per second. For planning purposes with these applications, 4 Gb/s Fibre Channel technology can support 50 MB/s bandwidth. Divide the expected bandwidth by that number and it provides the minimum base number of disks required. Round up for parity, growth, and disks that are slower than 15K RPM spindle speeds.

## Which RAID Level is Best for the Application?

The RAID level for one or all volume groups can have a major impact on cost, performance, and data recovery.

- **RAID 0** uses the least amount of storage, requiring no capacity for redundancy or parity protection, since it does not offer either.
  - ☞ It appears to be the cheapest solution because it requires no extra hardware for data protection or recovery, not taking into account the cost of recovery from data loss due to operator error, hardware malfunction, or other backup recovery.
  - ☞ RAID 0 also appears to be a high performance solution, since there is no lost overhead due to managing redundancy, although the cost of data recovery includes some lost production time associated with lost data recovery.
  - ☞ RAID 0 is rarely selected as the level of choice for important business or technical data.
- **RAID levels 1 or 1+0** require the most storage, which is more than twice the minimum capacity of the actual data requirements, and therefore the most expensive.
  - ☞ Using two complete copies of a volume, data recovery is not only thorough, but data access operates very quickly, since the controller has the choice of accessing data records from either set, it is more likely to select the closer one.
  - ☞ Recovery from data or hardware losses is relatively fast, since the failed set can be rapidly recopied from the surviving mirror twin without requiring mathematical recalculations.

- **RAID 5 or 3** is some capacity compromise among the RAID levels mentioned above, since it requires a parity drive for each volume group of between 2 and 29 data drives.
  - ☞ Failed devices or data errors can be recovered by mathematically analyzing the surviving data.
  - ☞ Low parity ratios (such as 4+1) result in lower savings, but faster recovery times. Higher parity ratios (such as 15+1) are less costly than lower ratio solutions, but require a substantially greater error recovery time.
  - ☞ Bear in mind the earlier discussion about higher drive spindle counts having a positive influence on IOPS performance. At the LUN or RAID group level, the greater the spindle count, the higher the expected level of IOPS performance.
  - ☞ RAID groups within the same system or application not only do not need to use the same RAID level, they do not need the same spindle-counts, or parity ratios, or block sizes.
    - I/O for most applications is skewed much more heavily toward reads than writes. Applications that have more of a write bias, such as greater than 20 to 25 percent write operations, might be better suited to the use of RAID 1 or 1+0.

The process of rationalizing the business decision of which RAID level and parity level to choose for the various logical storage volumes should be well-thought out and based on the requirements of the application and the organization. The advantage of Sun StorageTek 6540 array management software is that decisions are not irrevocable. Changing RAID levels, drive counts, parity ratios, volume group size, etc., can be dynamically performed from one value to another, without interrupting normal application processing or database access — a clear advantage over alternative designs.

## How Many Disk Drives Should Be In a Volume Group?

For transaction processing or IOPS heavy applications, use some of the guidelines already discussed:

1. Pick a RAID level that is reasonable for the expected activity.  
For this example, assume RAID 5.
2. Pick a drive type and capacity.  
For this example, assume 146 GB 10K RPM Fibre Channel disk drives.
3. Pick a parity drive ratio.  
For this example, seven data drives to one parity drive is a good ratio.
4. Quick check:  
Database requirement = about 1 TB, including planned growth.

Seven times 146GB drives = 1022 GB = ~1 TB .

Plus one parity drive adds up to eight drives, which equals one half of a drive rack .

The volume group must be defined as eight drives (seven data and one parity).

This volume group should be capable of sustaining at least 2000 random IOPS.

If this sample system behaved more like a streaming (high bandwidth) application, the 7+1P volume group could sustain over 300 MB/s sustained throughput.

Another guideline to keep in mind for volume groups is that the number of data spindles in the data volume group should be able to accommodate the maximum size of a data stripe. That is to say, the maximum size of the data stripe should be the equal to or greater than the segment size multiplied by the number of data disk spindles reserved for that volume group.

## RAID Topology

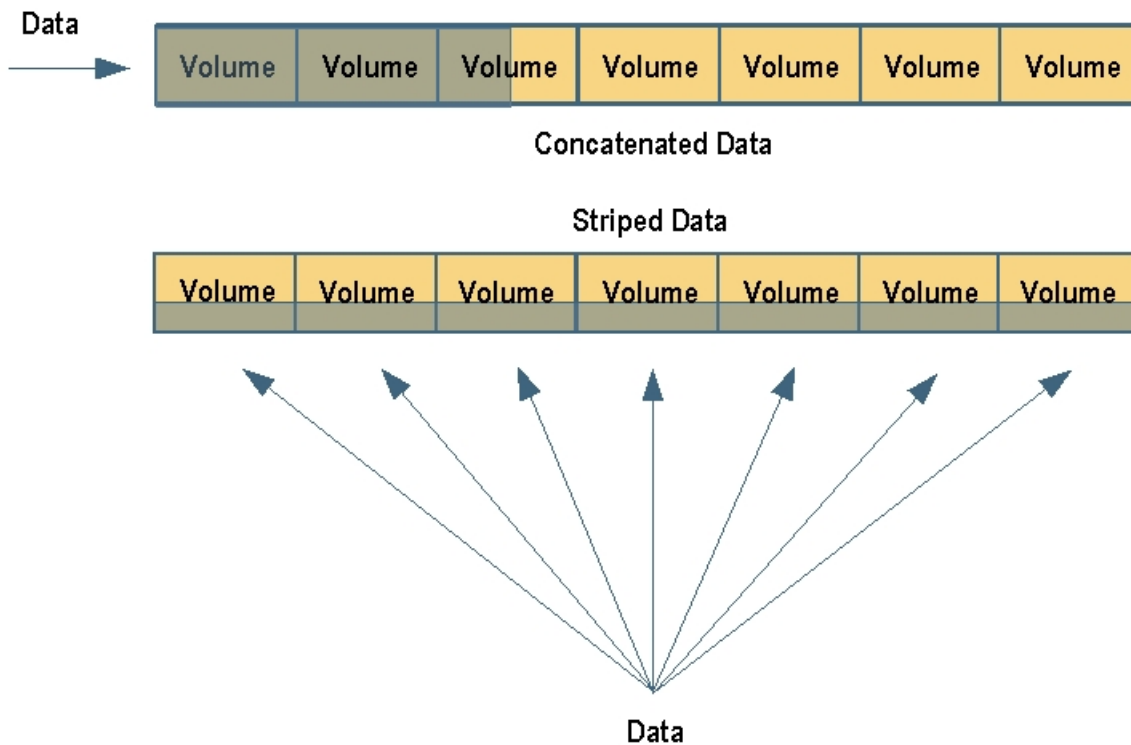
Most applications require several to numerous data segments. The segment can be either a small volume or, if the volume is significantly large based on the mathematical rules provided in the above section, each and every portion of a volume group distributed among the different spindles can also be referred to as a segment. When segments for the same application require access concurrently during normal operation, some performance benefit can be achieved if those segments are allocated to spindles in as many different drive enclosures as possible. This is sometimes referred to as *vertical striping*, as contrasted with segments striped to or from spindles within the same drive enclosure. The simple explanation for this is that striping benefits from overlapped I/O, but cannot benefit when there are conflicts for the same resource. Therefore, reduce all obvious points of conflict for the same resource.

## Controller Cache Parameters

Better I/O efficiency is achieved if the I/O block size for controller cache is equal to the stripe size of the I/O record. IOPS applications typically have smaller block sizes than bandwidth data. If this is the case, try to select the block size as a simple multiple or factor of the stripe size.

## Host Striping of Data on an Array

There are two primary ways that data can be striped across volumes in an array or arrays. RAID 0 (called striped) or concatenated. Both are illustrated in Figure 14. It is important to be careful when laying out several volumes in one volume group.



**Figure 14.** Concatenated and Striped Data

Concatenated data can fill up one volume before filling up the next volume. This is useful when using several volumes in one volume group, but performance is limited to one controller path at a time.

Striping writes to all volumes at the same time. This is useful when using volumes in different volume groups or arrays. This method increase performance. However, it is not useful with several volumes in the same volume group, as this degrades performance. Simultaneous reads or writes across different areas (stripes) on the same disk or volume can cause disk contention.


## Striping for Better Performance

Some administrators try to control how data is laid out on the array. They often create several volumes on a volume group and use a logical volume manager on the host to stripe the volumes a one volume. Although half of the volumes might be on one controller and the other half on the other controller, certain configurations can still cause drive head contention in the volume group. For example, the array in Figure 15 is configured with an 8+1 RAID group with 10 volumes. The 10 volumes are striped as one volume in the host using a logical volume manager.

Disk 1	Disk 2	Disk 3	Disk 4	Disk 5	Disk 6	Disk 7	Disk 8	Disk 9	
									Volume 1
									Volume 2
									Volume 3
									Volume 4
									Volume 5
									Volume 6
									Volume 7
									Volume 8
									Volume 9
									Volume 10

Volume Group 1 comprised of 8+1


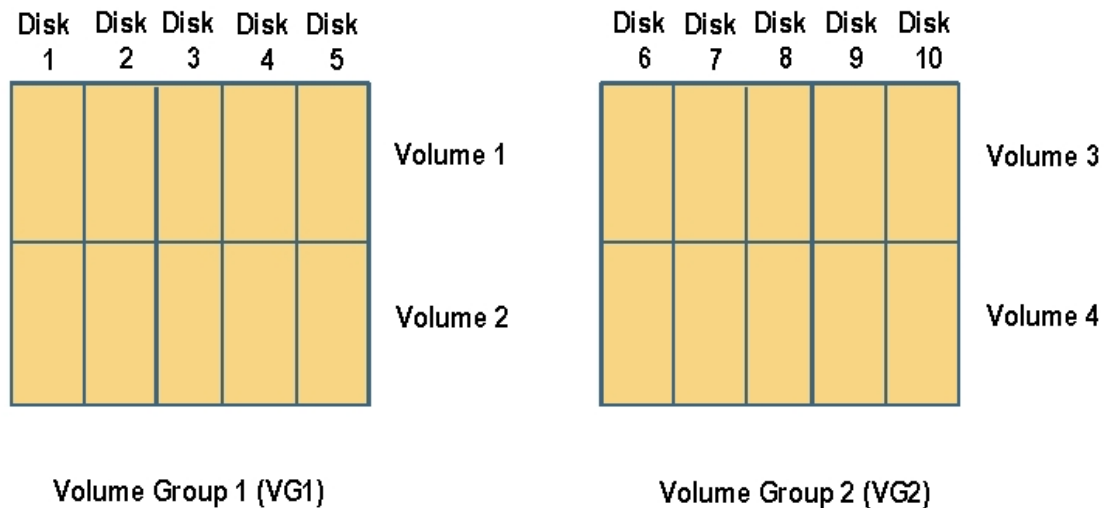
Host sees 10 drives.  
LVM uses 1 stripe  
across all 10 volumes.  
Reads and writes are  
spread across all 10  
volumes at the same  
time, causing drive  
contention



**Figure 15.** Volumes Striped as One Volume

When reads or writes occur in this configuration, the operations try to read or write to 10 different places on each of the 10 drives at the same time. This causes a performance problem, especially as I/O increases.

The example in Figure 16 uses larger volume sizes and fewer volumes, as well as an additional volume group. With the data spread across two volume groups, head contention is less of an issue.



**Figure 16.** Volumes Striped as Four Volumes and Two Volume Groups

If using a logical volume manager to stripe the volumes at the host, there are now two larger volumes striped across two volume groups. The other two volumes in the same volume group can be assigned to a different host. There is now one volume in one volume group on controller A, and one volume in the other volume group assigned to controller B. This configuration provides better performance than the previous scenario.

## Server Hardware — Bus Bandwidth and HBAs

An important limiting factor in I/O performance is the I/O capability of the server that hosts the application. The aggregate bandwidth of the server to the storage subsystem is measured in MB/s and consists of the total capability of the buses to which the storage system is connected. For example, a 64-bit PCI bus clocked at 133 MHz has a maximum bandwidth of  $133 \times 10^6 \text{ cycles/second} \times 8 \text{ bytes/cycle} = 1064 \times 10^6 \text{ bytes/s} = \text{approximately } 1 \text{ GB/s}$ .

Multiple HBAs on this bus share this single source of I/O bandwidth and each HBA can have multiple FC ports, which typically operate at 1 Gb/s, 2 Gb/s, or 4 Gb/s (100, 200, or 400 MB/s, respectively). Therefore, the ability to drive a storage subsystem can be throttled by either the server bus or the HBAs.

When a server is configured, or whenever I/O performance is analyzed, it is important to understand how much server bandwidth is available and which devices share it. If the aggregate maximum bandwidth of



the HBAs exceed that of the server, then the server can become a bottleneck to throughput. If it is lower, however, it might be possible to increase the throughput of an application by adding additional HBAs to the server, or replacing lower performance HBAs with faster ones, or reassigning the fastest HBAs to the applications that can take advantage of their higher performance.

See the online Interop Tool or Sun StorageTek 6540 Array Release Notes for supported HBAs.

## Command Tag Queuing (CTQ)

Command Tag Queuing refers to the controller's ability to line up (queue) multiple SCSI commands for a single LUN and execute them in an optimized order that minimizes rotational and seek latencies.

Although CTQ might not help in some instances, e.g., single-threaded I/O, it does not hurt performance and therefore is generally recommended. CTQ is enabled by default on Sun StorageTek 6540 arrays, but it must also be enabled on the HBA (see HBA vendor documentation) and in the host OS. The capability of a single host varies by OS type, but a general formula for calculating CTQ is:

$$\text{OS CTQ depth setting} = \text{maximum OS queue depth (<255)} / \text{total \# LUNs}$$

If the HBA has a lower CTQ capacity than the above calculation result, it limits the actual setting. The method for setting CTQ varies by OS type. Refer to the OS man page for the following commands and parameters for detailed information.

- Solaris: `sd_max_throttle` in `/etc/system`.
- HP-UX (11.0): `scsictl` command.
- HP-UX (11i): `scsi_max_depth` dynamic parameter added; use `kmtune` command.
- AIX: use `lsattr -E -l hdiskn` to view LUN setting and `chdev -l hdiskn -a q_type=simple -a queue_depth=NewValue` to change queue depth for a LUN. The `-T` and `-P` flags control when the change is effective and its permanence.
- IRIX: for each LUN, use the `fx` command (e.g., `fx -x "dksc 6,2,2"`) following the menus down to `/label/set/param`, where options are provided for Enable/Disable and for CTQ depth.
- Linux: OS default is viewed/set is set in generic `sg` driver at `/proc/scsi/sg` HBA parameter is set in HBA driver configuration file, e.g., `lpfc.conf`.
- Windows Edit Registry settings per HBA vendor documentation.

## Fibre Channel Switch Fabric

Use any FC switch settings recommended for your particular storage system model and release level, which are available from the supplier of your storage system. For example, on Brocade switches, verify the *In-Order Delivery* parameter is enabled. Understand the paths from servers to storage and in a multi-switch SAN fabric, where I/O traverses ISLs (Inter-switch link), ensure that sufficient ISL bandwidth is configured.

## Data Services Performance Enhancement

### Sun StorageTek Data Snapshot Software

For optimal performance when using the Sun StorageTek Data Snapshot software:

- Locate repository volumes on disks that are separate from production LUNs, to isolate repository writes and minimize copy-on-write penalty.
- Try to schedule read I/Os to the snapshot volume at off-peak times when I/O activity on the source LUN is lower.

## Sun StorageTek Data Volume Copy Software

The Sun StorageTek Data Volume Copy software uses optimized large blocks to complete the volume copy as quickly as possible, so little tuning is required other than setting the copy priority to the highest level.

Best practices for using volume copy include:

- Disable all snapshot volumes associated with a base volume before selecting it as a volume copy target volume.

## Sun StorageTek Data Replicator Software

For optimal performance when using the Sun StorageTek Data Replicator software:

- Upgrade the storage subsystems at both sites to the latest firmware levels available.
- Locate repository volumes on RAID-1 volumes that are separated from production volumes to isolate writes and help optimize performance.
- In general, use more disks in the target LUN volume group, e.g., 8+1 instead of 4+1.
- For the target LUN, enable *write caching*, but disable *write cache mirroring*.
- For the source LUN, enable *read caching*, *write caching* and *write cache mirroring*.
- Use the highest priority level for synchronization for optimal remote volume mirroring performance, assuming that the impact on host I/O performance is acceptable.
- For optimal performance, enable the *In Order Delivery* option on Brocade switches.

## SANtricity Performance Monitor (With SANtricity Only)

SANtricity provides an integrated Performance Monitor that reports the following statistics for each volume in the storage system:

- Total I/Os (since start of this monitoring session)
- Read percentage (percent of read I/Os)
- Cache hit percentage (percent of reads satisfied from cache)
- Current KB/s (since last polling interval or requested update)
- Maximum KB/s (highest value since last start)
- Current I/Os (since last polling interval or requested update)
- Maximum I/O/ (highest value since last start)

This tool adds the storage system view of performance to those provided by other host-based or fabric-based monitoring tools. Refer to SANtricity online help for detailed usage information about Performance Monitor.

An important value of the performance monitor is to objectively and quantitatively identify the performance hot-spots, pointing out those areas requiring greatest attention, as well as performance bottlenecks, which must be addressed before most kinds of performance optimization can be attempted.

## In Summary

Optimizing a storage system is much like tuning any other kind of system comprised of multiple individual and interrelated sub-functions and subsystems. The storage system can be optimized with a good understanding of how all of the pieces fit together to the point that it appears to be an optimum solution.

A new storage system is much like starting with a new desk: it's clean, free of clutter with lots of space to store new things and easily find them again. Over time the clear, empty spaces start to fill up, and stacks of similar things become fragmented into multiple batches that can become easily separated, causing the efficiency in accessing data to consistently diminish. Implementing these best practices can set the stage for a more effective and efficient storage management paradigm for longer term efficiency.

# Specifications

## Specifications for the Sun StorageTek 6540 array controller tray:

Feature	Specification
<b>Physical Planning</b>	
Dimensions	6.9 inches high (17.5 centimeters) 19.0 inches wide (48.2 centimeters) 25.6 inches deep (65.1 centimeters)
Maximum Weight	80.2 pounds (36.5 kilograms)
<b>Environmental (Operating)</b>	
Temperature	10° to 40° Celsius (50° to 104° Fahrenheit)
Relative Humidity	20% to 80% noncondensing
Altitude	100 feet (30.5 meters) below sea level to 10,000 feet (3048 meters)
Shock	25g @ 3.75ms, triangular
Vibration	Random vibration to specified power spectral density
Heat Output	235 Watts (800 BTU/hour)
<b>Environmental (Non-Operating)</b>	
Temperature (Storage)	-10° to 65° Celsius (-14° to 149° Fahrenheit)
Temperature (Transit)	-40° to 65° Celsius (-40° to 149° Fahrenheit)
Humidity (Storage)	10% — 90%, maximum dew point is 26° Celsius (79° Fahrenheit), 10% per hour Gradient
Humidity (Transit)	5% — 95%, maximum dew point is 26° Celsius (79° Fahrenheit), 10% per hour Gradient
Altitude (Storage)	100 feet (30.5 meters) below sea level to 10,000 feet (3048 meters)
Altitude (Transit)	100 feet (30.5 meters) below sea level to 40,000 feet (12,000 meters)
Shock	Unit will be able to withstand a free-fall drop from 2.5 inches onto a hard surface
Vibration	5-150-5 Hz, Sinusoidal wave form 0.5g
<b>Power Requirements</b>	
AC Power (CU)	3.0A maximum operating @ 100 VAC (90 — 136 VAC range), 50/60 Hz 1.5A maximum operating @ 220 VAC (198 — 264 VAC range), 50/60 Hz
AC Power (EXP)	3.5A maximum operating @ 100 VAC (90 — 136 VAC range), 50/60 Hz 1.4A maximum operating @ 240 VAC (198 — 264 VAC range), 50/60 Hz
<b>Standards Compliance</b>	
Safety regulations	IEC 60950-1, UL/CSA 60950-1, EN 60950-1, CB Report with all national differences
Emissions regulations	FCC Class A (47 CFR 15, Subpart B), ICES-003 Class A, EN55022 Class A, VCCI Class A, CNS13438 Class A, CISPR22 Class A
Immunity	EN55024
Certifications: safety/EMC	Safety: cULus, CE and GOST R (S-mark and CCC exempt as “Storage” product EMC: Class A for FCC, CE, VCCI, C-Tick, GOST R and BSMI (CCC exempt as “Storage” product)

# System Requirements

## System Requirements

### Supported Operating Systems

Operating System	Multi-Path Failover Driver	Cluster
Solaris 8,9, and 10 on SPARC® processor-based platforms	MPxIO	Sun™ Cluster 3.X software
	VERITAS Dynamic Multipathing (DMP) 4.X	VERITAS Cluster Service
	Redundant Disk Array Controller (RDAC) (Solaris 8 & 9 only)	Sun Cluster 3.X
Solaris 10 on AMD Opteron™ processor-based platforms	MPxIO	Sun Cluster 3.X
Windows 2003 R2 (Web, Standard., Enterprise, and Data Center) 32- and 64-bit	MPIO (Q2FY2007)	
	RDAC	Microsoft Cluster Server (MSCS)
	DMP 4.X (32-bit only)	VERITAS Cluster Service
Windows 2000 (Server and Advanced Server)	RDAC	MSCS
Redhat Linux A S/W S/ES Enterprise Edition 3.X/4.X (32- and 64-bit)	RDAC	Life Keeper Steeleye
	DMP 4.X (Q2FY2007)	
SuSe Linux LES 8.X, 9.X (32- and 64-bit)	RDAC	Life Keeper Steeleye
	DMP 4.X (Q2FY2007)	
AIX 5.2 / 5.3	RDAC	High Availability Cluster Multi-Processing (HACMP) (GTY required)
	DPF	HACMP (GTY required)
	DMP 4.X	
HP-UX 11.11 and 11.23	Logical Volume Manager (LVM)	MC/Service Guard (MCSG)
	DMP 4.X	VERITAS Cluster Service
Netware 6 SP4, 6.5 SP1	Novell NPE	Novell Cluster Services (NCS)
SGI IRIS 6.5.28	none	none
VMWARE for Solaris	MPxIO	(Q2FY2007)
VMWARE for Windows	RDAC	(Q2FY2007)

### New Sun StorageTek Interop Tool

The Sun StorageTek 6540 array is fully supported with major host operating systems and multi-path drivers. For details, please refer to Sun StorageTek Interop Tool at:

<https://extranet.stortek.com/interop/interop>

## Supported Sun Software

- Sun StorageTek Storage Domain (4, 8, 16, 64)
- Sun StorageTek Data Snapshot software for the Sun StorageTek 6540 array
- Sun StorageTek Data Volume Copy software for the Sun StorageTek 6540 array
- Sun StorageTek Data Replicator software for the Sun StorageTek 6540 array
- Sun StorageTek™ Performance Suite with Sun StorageTek™ QFS Software 4.0, or later
- Sun StorageTek™ Utilization Suite with Sun StorageTek™ SAM-FS Software 4.0, or later
- Sun StorageTek Traffic Manager software 4.6 for multipathing management (integrated into the Solaris 10 OS)
- Sun Cluster 3.0, update 3, or higher
- Sun StorageTek™ Enterprise Backup software 7.1 and later
- Solstice™ DiskSuite 4.2.1 (in conjunction with Solaris 8 Operating System)
- Storage Automated Diagnostic Environment 2.4
- Solaris™ Volume Manager (embedded in Solaris 9 Operating System)

## Third-Party Software

- VERITAS NetBackup (NBU) 5.X/6.X
- VERITAS Volume Manager with DMP (VxVM) 3.5\* or later
- VERITAS File System (VxFS) 3.5, or higher
- VERITAS Storage Foundation 4.X

# System Configuration and Management

---

## Sun StorageTek Common Array Manager Software

Sun StorageTek Common Array Manager software for the Sun StorageTek 6540 array provides a powerful, yet easy to use, Java™ platform based GUI for administering the Sun StorageTek 6540 arrays. Sun StorageTek Common Array Manager software enables online administration, a consistent interface across all operating systems, and the ability to monitor and manage one or all Sun StorageTek 6540 arrays from any location on the network.

### Centralized Administration

Sun StorageTek Common Array Manager software enables all Sun StorageTek 6540 arrays to be managed from a single interface at one or more locations on the network. Its browser user interface (BUI) provides a comprehensive view of all storage systems in the management domain. From the BUI, new storage systems can be manually or automatically detected and added. Each storage system is managed through a *Storage Window*. The console is specific to an individual storage system, however multiple Storage Windows can be launched from the *Array Summary Window* to simultaneously manage multiple storage systems.

### Array Summary Window

The Array Summary Window is a browser based GUI used to configure and maintain the Sun StorageTek 6540 arrays. The Array Summary Window allows auto discovery and manual registration of the different Sun StorageTek 6540 arrays. In addition, it displays detailed information for the array such as: name, health, type, firmware version, total capacity, available capacity, and network address.

### Storage Window

The Storage Window displays the application-oriented storage profiles, logical components (volumes and virtual disk, pools, and data services), and physical components (initiators, host groups, hosts, trays, and drives) for each Sun StorageTek 6540 array. All storage management operations for a selected storage system or for selected components within a storage system are launched from the appropriate Storage Window menu.

## Sun Storage Automated Diagnostic Environment

The Sun StorageTek Common Array Manager software offers proactive health checking, intelligent diagnosis, fault isolation event notification, and fault management reporting for the Sun StorageTek 6540 array from a single management console. This software helps improve recoverability and increase infrastructure uptime, thus contributing to overall improved application service levels. More specifically, it provides the following functionality:

- Collects health, configuration, and other non-customer-related data
- Evaluates statistical error reports
- Notifies designated parties about events, when action is required
- Step-by-step instructions to add expansion trays to the Sun StorageTek 6540 array
- Monitors host message files for errors in order to obtain status information about the Sun StorageTek 6540 arrays
- Makes decisions on actionable service issues
- Troubleshooting and fault isolation of the Sun StorageTek 6540 array

- Provides mechanisms for service personnel to remotely access the system to gather additional data, perform maintenance, perform upgrades, and invoke diagnostics
- Guides service personnel through FRU isolation, replacement, and validation
- Device revision checking for firmware of the Sun StorageTek 6540 array

## **Management Host System Requirement:**

The external management host where the management software resides has the following system requirements:

- Platform: SPARC processor-based server or workstation
- Operating system: Solaris 8 OS Update 5, Solaris 9 OS or later
- Disk space: 500 MB (includes 300 MB in the /opt directory and 200 MB in the /var directory)
- Minimum system memory (two arrays, two users): 512 MB
- Recommended system memory: 1 GB
- Client memory: 256 KB

Supported browsers:

- Netscape™ Navigator 7.0
- Microsoft Internet Explorer 5.0
- Mozilla™ 1.2.1



# Ordering Information

---

Ordering information and part numbers for the Sun StorageTek 6540 are provided in this section. Also, more detailed instructions for using Webdesk can be found in the *Sun StorageTek 6540 Array Ordering Guide* (part number 819-6139).

## Shipping Configurations

- Controller tray (1 x 1) includes the following:
  - ☐ Eight shortwave SFPs for host connections
  - ☐ Two 5-meter LC-LC Fibre Channel cables
  - ☐ Sun StorageTek Common Array Manager software, CD
  - ☐ Getting Started Guide
  - ☐ Please note: a server required to host Common Array Management software and the cables needed to connect the management server to the 6540 controller are not included.
- Expansion tray (0 x 1) includes the following:
  - Dual I/O modules per tray
  - Two 2-meter optical interconnect cables
  - Accessing documentation pointer
- Minimum of 5 drives per tray, maximum of 16 drives per tray
  - ☐ FC 4 Gb/s hard disk drives: 73 GB 15K RPM, 146 GB 15K RPM
  - ☐ FC 2 Gb/s hard disk drives: 146 GB 10K RPM, 300 GB 10K RPM
  - ☐ SATA-II 3 Gb/s hard disk drives: 500 GB 7.2K RPM
  - ☐ Dual FC RAID controller cards per tray
- Available as rack ready or rack-mounted configurations
- A country kit with the appropriate power cords is added by the WEBDESK depending on the destination of the shipment
- Optional data services and premium features:
  - ☐ Sun StorageTek Storage Domain — 4 right-to-use (RTU) license
  - ☐ Sun StorageTek Storage Domain — 8 RTU license
  - ☐ Sun StorageTek Storage Domain — 16 RTU license
  - ☐ Sun StorageTek Storage Domain — 64 RTU license
  - ☐ Sun StorageTek Storage Domain — 4 to 8 upgrade RTU license
  - ☐ Sun StorageTek Storage Domain — 8 to 16 upgrade RTU license
  - ☐ Sun StorageTek Storage Domain — 16 to 64 upgrade RTU license
  - ☐ Sun StorageTek Data Snapshot software (DSS) RTU license
  - ☐ Sun StorageTek Data Volume Copy software (DVC) RTU license
  - ☐ Sun StorageTek Data Replicator software (DRS) RTU license
  - ☐ Sun StorageTek DSS and DVC combo RTU license
  - ☐ Sun StorageTek DSS and 8-storage domain combo RTU license
  - ☐ Sun StorageTek DSS, DVC, and DRS combo RTU license
  - ☐ Sun StorageTek DSS, DVC, DRS, and 64-storage domain combo RTU license

- Optional 3U rack rail kits for:
  - Optional 3U rack rail kits for Sun™ Rack 900/1000 and Sun Fire™ system cabinets
- Optional SANtricity Management Software

# Marketing Part Numbering Scheme

---

## Sun StorageTek 6540 Array (Pre-Built, Pre-Racked Systems)

Example part number = XTA6540-01-214U09Z

X = X-option or no X for factory configured

T = Storage family product

A = Product revision number: A — first release

6540 = Product class/model number: 6540 — Sun StorageTek 6540 array

-01- = Type of controller tray in the system

01 = 4 GB cache, Fibre Channel

02 = 8 GB cache, Fibre Channel

03 = 16 GB cache, Fibre Channel

2 = Quantity of controller boards per control unit

0 = Expansion unit

1 = Single board per control unit

2 = Dual boards per control unit

14 = Quantity of control unit/component group

First digit is the total number of controller trays in the standard configuration

Second digit is the total number of trays in the standard configuration

11 = 1 controller unit, 1 tray total, 1 x 1

12 = 1 controller unit, 2 tray total, 1 x 2

14 = 1 controller unit, 4 tray total, 1 x 4

U = Drive form factor/size/RPM

U = FC, 2 Gb/s 3.5 inch 10K RPM 146 GB

X = FC, 2 Gb/s 3.5 inch 10K RPM 300 GB

A = SATA-II, 3 Gb/s 3.5 inch 7.2K RPM 500 GB

B = FC, 4 Gb/s 3.5 inch 15K RPM 73 GB

C = FC, 4 Gb/s 3.5 inch 15K RPM 146 GB

09 = Raw, unformatted capacity in TB

Z = RoHS compliant

## Sun StorageTek 6540 Controller Tray

Example part number = XTA6540R11A2—08FCZ

X = X-option or no X for factory configured

T = Storage family product

A = Product revision number: A — first release

6540 = Product class/model number: 6540 — Sun StorageTek 6540 array

**R** = Physical configuration\*

R = Rack ready (field installed)

M = Rack mounted (factory installed)

**11** = Quantity of control unit/component group

First digit is the total number of controller trays in the standard configuration

Second digit is the total number of trays in the standard configuration

11 = 1 controller tray, 1 tray total, 1 x 1 (controller tray)

01 = 0 controller unit, 1 tray total, 0 x 1 (expansion tray)

14 = 1 controller unit 4 tray total, 1 x 4

**A** = Power

A = AC power

D = DC power

**2** = Quantity of controller boards per control unit

0 = Expansion unit

1 = Single board per control unit

2 = Dual boards per control unit

— = Reserved character

**08** = Cache size: 4 GB, 8 GB, 16 GB of cache

**FC** = Interface

FC = Fibre Channel

**Z** = RoHS compliant

*Note: \* When ordering the “R” (rack ready) configurations for the Sun Rack 900/1000 racks, order the appropriate accessories (additional Fibre Channel cables, rack mount rail kits, etc.) separately. If ordered via the Webdesk Configurator “M” (rack mounted) configurations come in a rack and rail kits are already included.*

# Configuration Matrix / Marketing Part Numbers

## Ordering Details and Marketing Part Numbers

Orders for the Sun StorageTek 6540 arrays **must be placed** through either the :

- WebDesk *Configurator* — for rack mounted trays (factory installed in rack/cabinet).
- If non-racked systems are required, order x-option parts and do not use the Configurator.

### There are three ways to order the Sun StorageTek 6540 array:

1. **Sweet Spot configurations:** Limited selection of pre-built, pre-racked standard configurations. Each one of these systems has a 4 GB cache controller with four fully populated expansion trays. Rack, rack rails, and all cabling internal to the system are included in a single part number
2. **Customized ATO options:** Build-to-order configurations installed in rack and cabled at the factory. Trays with a minimum of 5 or maximum of 16 drives are available at Revenue Release. Trays with single drive increments will be available later. Also, systems with 12 or more trays that require two racks will be available by General Availability.
3. **X-option Loose trays:** Individual components for upgrades or to install in an existing rack. Must be racked and cabled in the field.

### Option 1: Sweet Spot Configurations

Do not use the Configurator for these parts, use the Quote tab. See the *Sun StorageTek 6540 Array Ordering Guide* (part number 819-6139) for more information.

Based on configuration specified, WebDesk picks from the following marketing part numbers:

Marketing Part Number	Description
<b>Sun StorageTek 6540 Array Full System Configurations — 4 GB cache Sun StorageTek 6540 array with four expansion trays — x-option pre-racked with 16 drives per tray</b>	
<b>XTA6540-01-214B04Z</b>	Rack-Mount Sun StorageTek 6540 storage system with 4 Gb/s control module (controller tray), 4 GB cache memory and 4672 GB raw capacity. System includes: 2 * 2 GB cache memory FC RAID controllers; 2 * redundant 375 W power supplies, each with 2 * redundant cooling fans; 8 * FC ports for expansion trays and 8 * optical FC host ports with shortwave SFPs; 1* interconnect battery canister, with 2 battery packs; 4 CSM200 trays each with 16 x 73 GB 15K RPM 4Gb/s FC-AL drives. Installed in a 38 RU Sun StorageTek cabinet; fully cabled w/rails for CSM200 arrays; 3 year on-site warranty. RoHS-5 compliant.
<b>XTA6540-01-214U09Z</b>	Rack-Mount Sun StorageTek 6540 storage system with 4 Gb/s control module (controller tray), 4 GB cache memory and 9344 GB raw capacity. System includes: 2 * 2 GB cache memory FC RAID controllers; 2 * redundant 375 W power supplies, each with 2 * redundant cooling fans; 8 * FC ports for expansion trays and 8 * optical FC host ports with shortwave SFPs; 1* interconnect battery canister, with 2 battery packs; 4 CSM200 trays each with 16 x 146 GB 15K RPM 4 Gb/s FC-AL drives. Installed in a 38 RU Sun StorageTek cabinet; fully cabled w/rails for CSM200 arrays; 3 year on-site. RoHS-5 compliant.

Marketing Part Number	Description
<b>XTA6540-01-214C09Z</b>	Rack-Mount Sun StorageTek 6540 storage system with 4 Gb/s control module (controller tray), 4 GB cache memory and 9344 GB raw capacity. System includes: 2 * 2 GB cache memory FC RAID controllers; 2 * redundant 375 W power supplies, each with 2 * redundant cooling fans; 8 * FC ports for expansion trays and 8 * optical FC host ports with shortwave SFPs; 1* interconnect battery canister, with 2 battery packs; 4 CSM200 trays each with 16 x 146 GB 15K RPM 4 Gb/s FC-AL drives. Installed in a 38 RU Sun StorageTek cabinet; fully cabled w/rails for CSM200 arrays; 3 year on-site warranty. RoHS-5 compliant.
<b>XTA6540-01-214X19Z</b>	Rack-Mount Sun StorageTek 6540 storage system with 4 Gb/s control module (controller tray), 4 GB cache memory and 19,200 GB raw capacity. System includes: 2 * 2 GB cache memory FC RAID controllers; 2 * redundant 375 W power supplies, each with 2 * redundant cooling fans; 8 * FC ports for expansion trays and 8 * optical FC host ports with shortwave SFPs; 1* interconnect battery canister, with 2 battery packs; 4 CSM200 trays each with 16 x 300 GB 10K RPM 4 Gb/s FC-AL drives. Installed in a 38 RU Sun StorageTek cabinet; fully cabled w/rails for CSM200 arrays; 3 year on-site warranty. RoHS-5 compliant.
<b>XTA6540-01-214A32Z</b>	Rack-Mount Sun StorageTek 6540 storage system with 4 Gb/s control module (controller tray), 4 GB cache memory and 32,000 GB raw capacity. System includes: 2 * 2 GB cache memory FC RAID controllers; 2 * redundant 375 W power supplies, each with 2 * redundant cooling fans; 8 * FC ports for expansion trays and 8 * optical FC host ports with shortwave SFPs; 1* interconnect battery canister, with 2 battery packs; 4 CSM200 trays each with 16 x 500 GB 10K RPM 4 Gb/s FC-AL drives. Installed in a 38 RU Sun StorageTek cabinet; fully cabled w/rails for CSM200 arrays; 3 year on-site warranty. RoHS-5 compliant.

## Option 2: Customized ATO Options

If the specified configuration cannot be met with one of the x-option part numbers above, the Webdesk Configurator generates a list of assemble-to-order (ATO) marketing part numbers as outlined below. This option results in *multiple* marketing part numbers (controller, expansion, drive type, number of drives per tray, and optional data services). At Revenue Release, only 5-drive and 16-drive tray options are available. Single increment drive trays will be available later. Also, Sun StorageTek 6540 arrays with 12 or more drive trays requiring more than one rack will be available by GA.

### Step 1. Select Controller Cache

Marketing Part Number	Description
<b>Sun StorageTek 6540 Controller Options</b>	
<b>TA6540M11A2-04FCZ</b>	Rack-Mount Sun StorageTek 6540 4 Gb/s control module (controller tray) with 4 GB cache memory. It includes: 2 * 2 GB cache memory FC RAID controllers; 2 * redundant 375 W power supplies, each with 2 * redundant cooling fans; 8 * FC ports for expansion trays and 8 * optical FC host ports with shortwave SFPs; 1* interconnect battery canister, with 2 battery packs; at least one CSM200 drive tray required. RoHS-5 compliant.
<b>TA6540M11A2-08FCZ</b>	Rack-Mount Sun StorageTek 6540 4 Gb/s control module (controller tray) with 8 GB cache memory. It includes: 2 * 4 GB cache memory FC RAID controllers; 2 * redundant 375 W power supplies, each with 2 * redundant cooling fans; 8 * FC ports for expansion trays and 8 * optical FC host ports with shortwave SFPs; 1* interconnect battery canister, with 2 battery packs; at least one CSM200 drive tray required. RoHS-5 compliant.

Marketing Part Number	Description
<b>TA6540M11A2-16FCZ</b>	Rack-Mount Sun StorageTek 6540 4 Gb/s control module (controller tray) with 16 GB cache memory. It includes: 2 * 8GB-cache memory FC RAID Controllers; 2 * redundant 375 W power supplies, each with 2 * redundant cooling fans; 8 * FC ports for expansion trays and 8 * optical FC host ports with shortwave SFPs; 1* interconnect battery canister, with 2 battery packs; at least one CSM200 drive tray required. RoHS-5 compliant.

**Step 2. Add capacity with expansion trays, drive type, and number of drives :**

<b>Expansion trays, 0 x 1 — ATO rack-mounted with 5 drives per tray (minimum configurations)</b>	
<b>TACSM2M01A0B365</b>	RoHS-5 compliant, Sun StorageTek CSM200, rack-mounted expansion tray, 365 GB, 5 * 73 GB 15K RPM 4 Gb/s FC-AL drives, 2 * I/O modules, 2 * redundant power supplies and cooling fans, 2 * copper FC ports for expansion, 2 * 2M copper cables, (for factory integration only).
<b>TACSM2M01A0U730</b>	RoHS-5 compliant, Sun StorageTek CSM200, rack-mounted expansion tray, 730 GB, 5 * 146 GB 10K RPM 2 Gb/s FC-AL drives, 2 * I/O modules, 2 * redundant power supplies and cooling fans, 2 * copper FC ports for expansion, 2 * 2M copper cables, (for factory integration only).
<b>TACSM2M01A0C730</b>	RoHS-5 compliant, Sun StorageTek CSM200, rack-mounted expansion tray, 730 GB, 5 * 146 GB 15K RPM 4 Gb/s FC-AL drives, 2 * I/O modules, 2 * redundant power supplies and cooling fans, 2 * copper FC ports for expansion, 2 * 2M copper cables, (for factory integration only).
<b>TACSM2M01A0X1500</b>	RoHS-5 compliant, Sun StorageTek CSM200, rack-mounted expansion tray, 1500 GB, 5 * 300 GB 10K RPM 2 Gb/s FC-AL drives, 2 * I/O modules, 2 * redundant power supplies and cooling fans, 2 * copper FC ports for expansion, 2 * 2M copper cables, (for factory integration only).
<b>TACSM2M01A0A2500</b>	RoHS-5 compliant, Sun StorageTek CSM200, rack-mounted expansion tray, 2500 GB, 5 * 500 GB 7.2K RPM SATA drives, 2 * I/O modules, 2 * redundant power supplies and cooling fans, 2 * copper FC ports for expansion, 2 * 2M copper cables, (for factory integration only).
<b>Expansion trays, 0x1 — ATO rack-mounted with 16 drives per tray (maximum configurations)</b>	
<b>TACSM2M01A0B1168</b>	RoHS-5 compliant, Sun StorageTek CSM200, rack-mounted expansion tray, 1168 GB, 16 * 73 GB 15K RPM 4 Gb/s FC-AL drives, 2 * I/O modules, 2 * redundant power supplies and cooling fans, 2 * copper FC ports for expansion, 2 * 2M copper cables, (for factory integration only).
<b>TACSM2M01A0U2336</b>	RoHS-5 compliant, Sun StorageTek CSM200, rack-mounted expansion tray, 2336 GB, 16 * 146 GB 10K RPM 2 Gb/s FC-AL drives, 2 * I/O modules, 2 * redundant power supplies and cooling fans, 2 * copper FC ports for expansion, 2 * 2M copper cables, (for factory integration only).
<b>TACSM2M01A0C2336</b>	RoHS-5 compliant, Sun StorageTek CSM200, rack-mounted expansion tray, 2336 GB, 16 * 146 GB 15K RPM 4 Gb/s FC-AL drives, 2 * I/O modules, 2 * redundant power supplies and cooling fans, 2 * copper FC ports for expansion, 2 * 2M copper cables, (for factory integration only).
<b>TACSM2M01A0X4800</b>	RoHS-5 compliant, Sun StorageTek CSM200, rack-mounted expansion tray, 4800 GB, 16 * 300 GB 10K RPM 2 Gb/s FC-AL drives, 2 * I/O modules, 2 * redundant power supplies and cooling fans, 2 * copper FC ports for expansion, 2 * 2M copper cables, (for factory integration only).
<b>TACSM2M01A0A8000</b>	RoHS-5 compliant, Sun StorageTek CSM200, rack-mounted expansion tray, 8000 GB, 16 * 500 GB 7.2K RPM SATA drives, 2 * I/O modules, 2 * redundant power supplies and cooling fans, 2 * copper FC ports for expansion, 2 * 2M copper cables, (for factory integration only).

**Step 3 – Racks are added when ordering via the Configurator. If you do not want racks in your order, see Option 3: X-options below. Expansion cabinet is automatically added for systems with 12 or more expansion trays. SunStorageTek 6540 cabinets already include all rack rails.**

<b>Cabinets</b>	
<b>TA6540-BCAB-38-OZ</b>	Sun StorageTek 6540 base cabinet, ATO
<b>TA6540-ECAB-38-OZ</b>	Sun StorageTek 6540 expansion cabinet, ATO

The base cabinet includes a modem plus rack rails and optical cables needed for a controller tray and 11 expansion trays. Therefore, there is no need to purchase additional rails or cables when adding capacity upgrades to this rack in the future. The expansion cabinet has no modem; also, rails and cables for only three expansion trays. Power cords for both cabinets are added based on shipping geography.

**Step 4. Add Optional Data Services: (licensed per Sun StorageTek 6540 array).**

<b>Marketing Part Number</b>	<b>Description</b>
<b>XTA6540-DSS-ARY</b>	Sun StorageTek Data Snapshot software right-to-use (RTU) license key for Sun StorageTek 6540 array.
<b>XTA6540-DVC-ARY</b>	Sun StorageTek Data Volume Copy software right-to-use (RTU) license key for Sun StorageTek 6540 array.
<b>XTA6540-DRS-ARY</b>	Sun StorageTek Data Replicator software right-to-use (RTU) license key for Sun StorageTek 6540 array.
<b>XTA6540-DOM4-ARY</b>	Four domain right-to-use (RTU) license key for Sun StorageTek 6540 array.
<b>XTA6540-DOM8-UPG</b>	Eight domain right-to-use (RTU) license key, upgrade from 4 domain for Sun StorageTek 6540 array.
<b>XTA6540-DOM8-ARY</b>	Upgrade to 8 domain right-to-use (RTU) license key for Sun StorageTek 6540 array.
<b>XTA6540-DOM16-UPG</b>	Sixteen domain right-to-use (RTU) license key, upgrade from 8 domain for Sun StorageTek 6540 array.
<b>XTA6540-DOM16-ARY</b>	Upgrade to 16 domain right-to-use (RTU) license key for Sun StorageTek 6540 array.
<b>XTA6540-DOM64-UPG</b>	Sixty-four domain right-to-use (RTU) license key, upgrade from 16 domain for Sun StorageTek 6540 array.
<b>XTA6540-DOM64-ARY</b>	Upgrade to 64 domain right-to-use (RTU) license key for Sun StorageTek 6540 array.
<b>XTA6540-CMBO-DOM8</b>	Sun StorageTek Data Services software Combo right-to-use (RTU) license key for Data Snapshot and 8 domains. No capacity limit.
<b>XTA6540-CMBO-DOM64</b>	Sun StorageTek Data Services software Combo right-to-use (RTU) license key for Data Snapshot, Data Volume Copy, Data Replicator, and 64 domains. No capacity limit.
<b>XTA6540-CMBO-ARY-2</b>	Sun StorageTek Data Services software Combo right-to-use (RTU) license key for Data Snapshot and Data Volume Copy software. No capacity limit.
<b>XTA6540-CMBO-ARY-3</b>	Sun StorageTek Data Services software Combo right-to-use (RTU) license key for Data Snapshot, Data Volume Copy, and Data Replicator software. No capacity limit.



<b>OPTIONAL — BY APPROVAL ONLY: SANtricity Management Software (only for customers with existing SANtricity implementation)</b>	
<b>SN599-LCO-99H9</b>	SANTRICITY 9.16, license for Sun StorageTek 6540 array. Need to order a media kit for the appropriate platform. Electronic software distribution.
<b>SN59B-916-E9M9</b>	SANTRICITY 9.16, AIX-HOST 6540, media and documentation. Need to order license P/N: SN199-LCO-99H9 if this is the first media kit for the Sun StorageTek 6140 array. No license required for second media kit.
<b>SN59H-916-E9M9</b>	SANTRICITY 9.16, HP-HOST 6540, media and documentation. Need to order license P/N: SN199-LCO-99H9 if this is the first media kit for the Sun StorageTek 6140 array. No license required for second media kit.
<b>SN59K-916-E9M9</b>	SANTRICITY 9.16, W2K-HOST 6540, media and documentation. Need to order license P/N: SN199-LCO-99H9 if this is the first media kit for the Sun StorageTek 6140 array. No license required for second media kit.
<b>SN59L-916-E9M9</b>	SANTRICITY 9.16, LINUX-HOST 6540, media and documentation. Need to order license P/N: SN199-LCO-99H9 if this is the first media kit for the Sun StorageTek 6140 array. No license required for second media kit.
<b>SN59R-916-E9M9</b>	SANTRICITY 9.16, IRIX-HOST 6540, media and documentation. Need to order license P/N: SN199-LCO-99H9 if this is the first media kit for the Sun StorageTek 6140 array. No license required for second media kit.
<b>SN59S-916-E9M9</b>	SANTRICITY 9.16, SUN SOLARIS-HOST 6540, media and documentation. Need to order license P/N: SN199-LCO-99H9 if this is the first media kit for the Sun StorageTek 6140 array. No license required for second media kit.
<b>SN59V-916-E9M9</b>	SANTRICITY 9.16, NETWARE-HOST 6540, media and documentation. Need to order license P/N: SN199-LCO-99H9 if this is the first media kit for the Sun StorageTek 6140 array. No license required for second media kit.
<b>SN59W-916-E9M9</b>	SANTRICITY 9.16, W2K3-HOST 6540, media and documentation. Need to order license P/N: SN199-LCO-99H9 if this is the first media kit for the Sun StorageTek 6140 array. No license required for second media kit.

### **Option 3: X-options — Components for upgrades or to add to existing customer rack.**

<b>Marketing Part Number</b>	<b>Description</b>
<b>Controller tray, 1 x 1 — rack ready</b>	
<b>XTA6540R11A2-04FCZ</b>	RoHS-5 compliant, Sun StorageTek 6540, Fibre Channel RAID controller with 4 GB cache and 8 host ports.
<b>XTA6540R11A2-08FCZ</b>	RoHS-5 compliant, Sun StorageTek 6540, Fibre Channel RAID controller with 8 GB cache and 8 host ports.
<b>XTA6540R11A2-16FCZ</b>	RoHS-5 compliant, Sun StorageTek 6540, Fibre Channel RAID controller with 16 GB cache and 8 host ports.

<b>Expansion trays, 0 x 1 — x-options rack-ready with 5 drives per tray (minimum configurations)</b>	
<b>XTACSM2R01A0B365</b>	RoHS-5 compliant, Sun StorageTek CSM200, rack-ready expansion tray, 365 GB, 5 * 73 GB 15K RPM 4 Gb/s FC-AL drives, 2 * I/O modules, 2 * redundant power supplies and cooling fans, 2 * copper FC ports for expansion, 2 * 2M copper cables.

<b>Expansion trays, 0 x 1 — x-options rack-ready with 5 drives per tray (minimum configurations)</b>	
<b>XTACSM2R01A0U730</b>	RoHS-5 compliant, Sun StorageTek CSM200, rack-ready expansion tray, 730 GB, 5 * 146 GB 10K RPM 2 Gb/s FC-AL drives, 2 * I/O modules, 2 * redundant power supplies and cooling fans, 2 * copper FC ports for expansion, 2 * 2M copper cables.
<b>XTACSM2R01A0C730</b>	RoHS-5 compliant, Sun StorageTek CSM200, rack-ready expansion tray, 730 GB, 5 * 146 GB 15K RPM 4 Gb/s FC-AL drives, 2 * I/O modules, 2 * redundant power supplies and cooling fans, 2 * copper FC ports for expansion, 2 * 2M copper cables.
<b>XTACSM2R01A0X1500</b>	RoHS-5 compliant, Sun StorageTek CSM200, rack-ready expansion tray, 1500 GB, 5 * 300 GB 10K RPM 2 Gb/s FC-AL drives, 2 * I/O modules, 2 * redundant power supplies and cooling fans, 2 * copper FC ports for expansion, 2 * 2M copper cables.
<b>XTACSM2R01A0A2500</b>	RoHS-5 compliant, Sun StorageTek CSM200, rack-ready expansion tray, 2500 GB, 5 * 500 GB 7.2K RPM SATA drives, 2 * I/O modules, 2 * redundant power supplies and cooling fans, 2 * copper FC ports for expansion, 2 * 2M copper cables.
<b>Expansion trays, 0 x 1 — x-options rack-ready with 16 drives per tray (maximum configurations)</b>	
<b>XTACSM2R01A0B1168</b>	RoHS-5 compliant, Sun StorageTek CSM200, rack-ready expansion tray, 1168 GB, 16 * 73 GB 15K RPM 4 Gb/s FC-AL drives, 2 * I/O modules, 2 * redundant power supplies and cooling fans, 2 * copper FC ports for expansion, 2 * 2M copper cables (standard configuration).
<b>XTACSM2R01A0U2336</b>	RoHS-5 compliant, Sun StorageTek CSM200, rack-ready expansion tray, 2336 GB, 16 * 146 GB 10K RPM 2 Gb/s FC-AL drives, 2 * I/O modules, 2 * redundant power supplies and cooling fans, 2 * copper FC ports for expansion, 2 * 2M copper cables (standard configuration).
<b>XTACSM2R01A0C2336</b>	RoHS-5 compliant, Sun StorageTek CSM200, rack-ready expansion tray, 2336 GB, 16 * 146 GB 15K RPM 4 Gb/s FC-AL drives, 2 * I/O modules, 2 * redundant power supplies and cooling fans, 2 * copper FC ports for expansion, 2 * 2M copper cables (standard configuration).
<b>XTACSM2R01A0X4800</b>	RoHS-5 compliant, Sun StorageTek CSM200, rack-ready expansion tray, 4800 GB, 16 * 300 GB 10K RPM 2 Gb/s FC-AL drives, 2 * I/O modules, 2 * redundant power supplies and cooling fans, 2 * copper FC ports for expansion, 2 * 2M copper cables (standard configuration).
<b>XTACSM2R01A0A8000</b>	RoHS-5 compliant, Sun StorageTek CSM200, rack-ready expansion tray, 8000 GB, 16 * 500 GB 7.2K RPM SATA drives, 2 * I/O modules, 2 * redundant power supplies and cooling fans, 2 * copper FC ports for expansion, 2 * 2M copper cables (standard configuration).

Expansion trays with custom number of drives can be built using the following parts. Drives will be integrated into the expansion tray at the factory but the expansion tray will have to be racked and cabled in the field. Minimum of 5 drives is required. Mixing drive types in a tray is not factory configurable but can be accomplished in the field.

<b>Expansion trays, 0x1 — diskless chassis — rack ready</b>	
<b>TACSM2R01A0-0</b>	RoHS-5, Sun StorageTek CSM200, rack-ready expansion tray — diskless chassis, 0 GB, 0 drives; must order minimum 5 drives, 2 * I/O modules, 2 * redundant power supplies and cooling fans, 4 * copper FC ports for expansion, 2 * 2M copper cables, (for factory integration only)

<b>Drives for diskless chassis — rack ready</b>	
<b>TA-FC1CF-73G15K</b>	RoHS-6, Sun StorageTek 6140 array / CSM200, 73 GB 15K rpm FC-AL drive (for factory integration only)
<b>TA-FC1CF-146G10K</b>	RoHS-6, Sun StorageTek 6140 array / CSM200, 146 GB 10K rpm FC-AL drive (for factory integration only)
<b>TA-FC1CF-146G15K</b>	RoHS-6, Sun StorageTek 6140 array / CSM200, 146 GB 10Krpm FC-AL drive (for factory integration only)
<b>TA-FC1CF-300G10K</b>	RoHS-6, Sun StorageTek 6140 array / CSM200, 300 GB 10K rpm FC-AL drive (for factory integration only)
<b>TA-ST1CF-500G7K</b>	RoHS-6, Sun StorageTek 6140 array / CSM200, 500 GB 7.2K rpm SATA drive (for factory integration only)

<b>Cabinets – x-options</b>	
<b>XTA6540-BCAB-38-0Z</b>	Sun StorageTek 6540 base cabinet, x-option
<b>XTA6540-ECAB-38-0Z</b>	Sun StorageTek 6540 expansion cabinet, x-option

<b>Part Number</b>	<b>Description</b>	<b>Category</b>
<b>Drives and cables</b>		
<b>XTA-FC1CF-73G15K</b>	RoHS-6 compliant, CSM200, 73 GB 15K RPM FC-AL drive	Hard disk drives
<b>XTA-FC1CF-146G10K</b>	RoHS-6 compliant, CSM200, 146 GB 10K RPM FC-AL drive	
<b>XTA-FC1CF-146G15K</b>	RoHS-6 compliant, CSM200, 146 GB 10K RPM FC-AL drive	
<b>XTA-FC1CF-300G10K</b>	RoHS-6 compliant, CSM200, 300 GB 10K RPM FC-AL drive	
<b>XTA-ST1CF-500G7K</b>	RoHS-6 compliant, CSM200, 500 GB 7.2K RPM SATA drive	
<b>SG-XPCI1FC-EM4</b>	4 Gb/s, Emulex, PCI-X, single port	Host Bus Adaptors
<b>SG-XPCI2FC-EM4</b>	4 Gb/s, Emulex, PCI-X, dual port	
<b>SG-XPCI1FC-QF4</b>	4 Gb/s, Qlogic, PCI-X, single port	
<b>SG-XPCI2FC-QF4</b>	4 Gb/s, Qlogic, PCI-X, dual port	
<b>SG-XPCIE1FC-EM4</b>	4 Gb/s, Emulex, PCI-E, single port	
<b>SG-XPCIE2FC-EM4</b>	4 Gb/s, Emulex, PCI-E, dual port	
<b>SG-XPCI1EFC-QF4</b>	4 Gb/s, Qlogic, PCI-E, single port	
<b>SG-XPCI2EFC-QF4</b>	4 Gb/s, Qlogic, PCI-E, dual port	
<b>X9730A-Z</b>	0.8m LC to LC FC optical cable, RoHS-5 compliant.	LC – LC Fibre Channel cables
<b>X9732A-Z</b>	2M LC to LC FC optical cable, RoHS-5 compliant	
<b>X9733A-Z</b>	5M LC to LC FC optical cable, RoHS-5 compliant	
<b>X9734A-Z</b>	15M LC to LC FC optical cable, RoHS-5 Compliant.	
<b>X9736A-Z</b>	25M LC to LC FC optical cable, RoHS-5 Compliant.	
<b>X9738A-Z</b>	50M LC to LC FC optical cable, RoHS-5 Compliant.	
<b>X9740A-Z</b>	100M LC to LC FC optical cable, RoHS-5 compliant	

Part Number	Description	Category
<b>X9722A-Z</b>	2M LC to SC FC optical cable, RoHS-5 compliant.	LC – SC Fibre Channel cables
<b>X9723A-Z</b>	5M LC to SC FC optical cable, RoHS-5 compliant.	
<b>X9724A-Z</b>	15M LC to SC FC optical cable, RoHS-5 compliant	
<b>X9735A-Z</b>	25M SC to LC FC optical cable. RoHS-5 compliant.	
<b>X9737A-Z</b>	50M SC to LC FC optical cable, RoHS-5 compliant.	
<b>XTACSM2-RK-3RU-19U</b>	Sun StorageTek 6540 array / CSM200, 3U rack rail kit for Sun Rack 900/1000 racks	Rack rail kits

## Optional — Sun StorageTek Traffic Manager Software (SSTM) for Non-Solaris Operating Systems

Part Number	Title and Shipping Configuration
<b>MPDIB-450-99YS</b>	Sun StorageTek Traffic Manager 4.5 Media Kit and documentation for IBM AIX5.X
<b>MPDIX-450-99YS</b>	Sun StorageTek Traffic Manager 4.5 Media Kit and documentation for MS Windows
<b>MPDIH-450-99YS</b>	Sun StorageTek Traffic Manager 4.5 Media Kit and documentation for HP-UX
<b>MPDIL-450-99YS</b>	Sun StorageTek Traffic Manager 4.5 Media Kit and documentation for Red Hat Linux Enterprise Editions 2.1 & 3.0 (WS, AS, ES)

The SSTM 4.6 SAN software release will be distributed via patches posted to SunSolve.

## Optional — Sun Services

The Sun StorageTek 6540 array models are available with optional enhanced services packages, which allow organizations to rapidly implement complex data storage environments. These extended service packages help ensure the use of sound storage installation and configuration practices, thereby allowing Sun to implement the support infrastructure that is required to maintain the most demanding enterprise and data center environments.

Recommended services available for the Sun StorageTek 6540 arrays are shown in the table below.

Installation Service	Americas Part Number	EMEA/APAC Part Number
Sun StorageTek Array Installation, base charge for the Sun StorageTek 6540 array/CSM200, 3U Tel	ARRAY-INS-BAS2	EIS-ARRAY
Sun StorageTek Tray Installation, per tray charge for Sun StorageTek 6540 array	ARRAY-INS-PER-TRAY	EIS-ARRAY-TRAY
Sun StorageTek Array Installation, after hours base charge for the Sun StorageTek 6540 array	ARRAY-INS-BAS2-AH	EIS-ARRAY-AH
Sun StorageTek Tray Installation, after hours per tray charge for the Sun StorageTek 6540 array	ARRAY-INS-PRTRY-AH	EIS-ARRAY-TRAY-AH
Various storage services are offered by Sun Professional Services (SunPS <sup>SM</sup> )	Various	Available now, please refer to the following URL for more details: <a href="http://sunps.central/services/storage/index.html">http://sunps.central/services/storage/index.html</a>

## SunSpectrum<sup>SM</sup> Instant Upgrades (Warranty Upgrades)

The SunSpectrum<sup>SM</sup> program is a service offering that allows customers to choose the level of service best suited to their needs. 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 Solaris Operating System, and telephone support for Sun software packages. Customers should check with their local Sun Services representative for program and feature availability in their areas.

For information specific to the Sun StorageTek 6540 array, refer to:

<http://www.sun.com/service/support/products/storage/>

SunSpectrum program support contracts are available both during and after the warranty program. Customers can 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 range from SunSpectrum Bronze<sup>SM</sup> level to SunSpectrum Platinum<sup>SM</sup> level. Contact a Sun Services representative for further details. The following part numbers can be used to upgrade system warranty to the given level of SunSpectrum service for each specified product:

Part Numbers	Service
W9D-SE6140-1-1G	Sun StorageTek 6540-1 array expansion tray upgrade to 1 year of Gold support.
W9D-SE6140-1-1P	Sun StorageTek 6540-1 array expansion tray upgrade to 1 year of Platinum support.
W9D-SE6140-1-24-1G	Sun StorageTek 6540-1 array expansion tray upgrade to Gold support + 7X24 on-site support for 1 year.
W9D-SE6140-1-24-2G	Sun StorageTek 6540-1 upgrade to Gold support + 7X24 on-site support for 2 years.
W9D-SE6140-1-24-3G	Sun StorageTek 6540-1 array expansion tray upgrade to Gold support + 7X24 on-site support for 3 years.
W9D-SE6140-1-2G	Sun StorageTek 6540-1 array expansion tray upgrade to 2 years of Gold support.
W9D-SE6140-1-2P	Sun StorageTek 6540-1 array expansion tray upgrade to 2 years of Platinum support.
W9D-SE6140-1-3G	Sun StorageTek 6540-1 array expansion tray upgrade to 3 years of Gold support.
W9D-SE6140-1-3P	Sun StorageTek 6540-1 array expansion tray upgrade to 3 years of Platinum support.
W9D-SE6140-1G	Sun StorageTek 6540 controller tray upgrade to 1 year of Gold support.
W9D-SE6140-1P	Sun StorageTek 6540 controller tray upgrade to 1 year of Platinum support.
W9D-SE6140-24-1G	Sun StorageTek 6540 controller tray upgrade to Gold support + 7X24 on-site support for 1 year.
W9D-SE6140-24-2G	Sun StorageTek 6540 controller tray upgrade to Gold support + 7X24 on-site support for 2 years.
W9D-SE6140-24-3G	Sun StorageTek 6540 controller tray upgrade to Gold support + 7X24 on-site support for 3 years.
W9D-SE6140-2G	Sun StorageTek 6540 controller tray upgrade to 2 years of Gold support.
W9D-SE6140-2P	Sun StorageTek 6540 controller tray upgrade to 2 years of Platinum support.

<b>Part Numbers</b>	<b>Service</b>
<b>W9D-SE6140-3G</b>	Sun StorageTek 6540 controller tray upgrade to 3 years of Gold support.
<b>W9D-SE6140-3P</b>	Sun StorageTek 6540 controller tray upgrade to 3 years of Platinum support.
<b>WPD-SE6140-1-24-3G</b>	Sun StorageTek Remote Response 6540-1 array expansion tray upgrade to Gold support + 7x24 on-site support for 3 years for the price of 2 years. Promotional Part Number.
<b>WPD-SE6140-1-3P</b>	Sun StorageTek Remote Response 6540-1 array expansion tray upgrade to Platinum support for 3 years for the price of 2 years. Promotional Part Number.
<b>WPD-SE6140-24-3G</b>	Sun StorageTek Remote Response 6540 controller tray upgrade to Gold support + 7x24 on-site support for 3 years for the price of 2 years. Promotional Part Number. Requires SSRR activation.
<b>WPD-SE6140-3P</b>	Sun StorageTek Remote Response 6540 controller tray upgrade to Platinum support for 3 years for the price of 2 years. Promotional Part Number. Requires SSRR activation.



# Upgrades

---

## Upgrade Paths

Customers can protect their investment in Sun StorageTek storage by trading up to a Sun StorageTek 6540 array. Customers can trade-in Sun StorageTek™ T3 arrays, Sun StorageTek™ A5200 arrays, Sun StorageTek™ 3X00 arrays, Sun StorageTek 6120 arrays, or Sun StorageTek 6320 arrays on a one for one basis or trade in many systems utilizing the Sun storage consolidation program. Non Sun storage systems can also be traded-in.

Trade-ins to the Sun StorageTek 6540 array are available as full system swaps through the Sun Upgrade Advantage Program. Please note, disk drives from a Sun StorageTek 6120 array, Sun StorageTek 6320 array, or other legacy Sun systems cannot migrate to the Sun StorageTek 6540 array. Customers must return a full system cabinet with all drives. Sun handles the disposition of older assets for the customer at no charge. Please order one RMA Kit (part number UG-RMA) for one for one trade-ins. When consolidating, order one RMA Kit (part number CU-CONSOL-RMA) when multiple systems are to be traded-in. RMA kits provide customers instructions on where to return the used residual equipment.

The array to be traded-in must be owned by, used by, and in the possession of the customer at least (90) days prior to trading-in. To qualify for the trade-in allowance, the customer must return within 90 days, a fully functional system with all drives. If trade-in is not returned within 90 days, Sun can bill back and customer agrees to pay promptly the full amount of the trade-in allowance. No credit or return of trade-in equipment is allowed after expiration of 90 days. For customers who want to use the online returns Web tool or are requesting pickup of traded equipment must do so before their return date. More information on Sun returns can be found at: [www.sun.com/ibb/upgrades/sunreturns.html](http://www.sun.com/ibb/upgrades/sunreturns.html).

## Sun Storage Consolidation Trade-in Program

Sun Storage Consolidation Trade-in Program enables customers to trade-in many older Sun or non Sun storage systems either for a new Sun storage system or to enrich an existing Sun storage system. An allowance code is used when trading up to a Sun StorageTek 6540 array. Storage UAP product matrices containing standard upgrade allowance codes are included in the Sun Configuration guide, at <http://ibb.eng/upgrades> or on SunWin:

<u>One for one storage trade-ins</u>	<u>Sun Win Token #94719</u>
<u>Storage Consolidation</u>	<u>Sun Win Token #112445</u>
<u>Component trade-ins</u>	<u>Sun Win Token #108142</u>

To determine the value of the trade-in, apply the allowance to the list price of the Sun StorageTek 6540 array. The value of the trade-in can be reinvested towards the purchase of professional services, support, or additional components.

For additional help on upgrades, IBB representatives are trained and ready to assist:

<http://ibb.eng/org/> or send an email to [ibbrvalues@sun.com](mailto:ibbrvalues@sun.com).

Resources:

- Internal IBB Upgrades Home Page: [ibb.eng/upgrades](http://ibb.eng/upgrades)
- To find your local IBB sales representative: <http://ibb.eng/org/>
- External IBB Upgrades Home Page: [www.sun.com/ibb/upgrades](http://www.sun.com/ibb/upgrades)

☞ Storage Consolidation Calculator: <http://ibb.eng/>

This tool helps calculate what is to be traded-in and what is to be purchased and provides a trade-in allowance code.

☞ Sun Returns: [www.sun.com/ibb/upgrades/sunreturns.html](http://www.sun.com/ibb/upgrades/sunreturns.html).

# Service and Support

The SunSpectrum<sup>SM</sup> 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 Solaris Operating System, and telephone support for Sun software packages. Customers should check with their local Sun Services representatives for program and feature availability in their areas.

SunSpectrum program support contracts are available both during and after the warranty program. Customers can choose to uplift the service and support agreement to meet their business needs by purchasing a SunSpectrum contract. For more information on the SunSpectrum program offerings refer to the following URL: [http://service.central/TS/ESP/SunSpectrum/Feature\\_Matrix/index.html](http://service.central/TS/ESP/SunSpectrum/Feature_Matrix/index.html).

The four levels of SunSpectrum support contracts are outlined below.

Program	Description
<b>Mission-Critical SunSpectrum Platinum<sup>SM</sup> Support</b>	Designed for mission-critical storage systems, the Platinum Service Plan provides complete 24x7 coverage, our highest priority response, interoperability assistance, and additional specialized services to help customers proactively maintain high availability.
<b>Business-Critical SunSpectrum Gold<sup>SM</sup> Support</b>	Key storage systems need to be available. With SunSpectrum Gold, customers receive business-critical support that includes extended hardware service coverage hours and interoperability assistance. Telephone support is provided 24x7. On-site support is provided from 8 a.m. To 8 p.m. Mon. through Fri. with 4 hour response.
<b>System Coverage SunSpectrum Silver<sup>SM</sup> Support</b>	For basic support at a great value, SunSpectrum Silver is the answer. Customers get all of the essentials, including hardware service coverage during normal business hours.
<b>Self-Directed SunSpectrum Bronze<sup>SM</sup> Support</b>	Provided for customers who rely primarily upon their own in-house service capabilities. Allows customers to deliver high quality service by giving them access to UNIX <sup>®</sup> expertise, Sun certified replacement parts, software releases and technical tools. Support is provided 8 a.m. to 5 p.m. Mon. through Fri.

For information specific to the Sun StorageTek 6540 array models, refer to:

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

## Warranty

The following table indicates warranty details for the Sun StorageTek 6540 array:

Repair Support Duration	Software Support Duration	Phone Coverage		Hardware Coverage		
		Hours of Coverage	Call-Back Response Time	Hours of Coverage	Response Time <sup>1</sup>	Delivery Method
3 years	90 days	8 a.m. – 8 p.m.	Customer-defined priority (4 hr, 8 hr, NBD)	Local business hours, M-F, 8 a.m. - 5 p.m.	<ul style="list-style-type: none"> <li>Next Business Day</li> <li>8 Business Hours</li> </ul>	On-site

For the latest warranty information for the Sun StorageTek 6540 array, refer to:

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

<sup>1</sup> Average response times may vary by country.



## Education

For further information on courses visit Sun Ed Web site at <http://www.sun.com/service/suned>, or to order, call: 1-800-422-8020.

## Sun Enterprise Services

Sun Enterprise Services provides first call support and escalates to Sun StorageTek DMG Global Services. Services offered through Sun Enterprise Services are:

- Sun Installation Services — provides installation and limited implementation of Sun StorageTek 6540 array
  - ☞ Site audit
  - ☞ Installation planning
  - ☞ System installation specification
  - ☞ Statement of installation
  - ☞ Installation and customization of Sun hardware and software
  - ☞ Installation verification
  - ☞ System turnover
- Sun Integration Services
  - ☞ Project management
  - ☞ Requirements validation
  - ☞ Solution design validation
  - ☞ Build specification report creation
  - ☞ Test plan development and execution
  - ☞ Documentation
  - ☞ Knowledge transfer
- Sun Storage Consolidation Services
  - ☞ TCO review services
  - ☞ SAN assessment, architecture, and implementation
  - ☞ Hardware and software installation
  - ☞ Storage migration services
  - ☞ Education services
  - ☞ Support service

## Professional Services

Sun Client Solutions is updating existing Storage Migration, Backup and Recovery, SAN, and Point-in-Time Copy Services to include support for Sun StorageTek 6540 array. A fixed price Sun StorageTek 6540 Array Implementation service is not being released for this product.

## **Storage Migration Service**

This service can help customers safely transfer data from one storage system to another storage system, without pulling their internal resources from other critical business tasks.

## **Backup and Restore Services**

**Assessment** — The Backup and Restore Assessment Service delivers an assessment of an existing Sun StorageTek™ Enterprise NetBackup or Solstice Backup environment. The assessment is to help ensure that the existing configuration and operational environment meets customer requirements. It reveals weaknesses or shortcomings in the areas of server/client configurations. Additionally, it reviews the historical operations of the backup and restore environment to ensure the problems are not reoccurring or are recurring occasionally, the service reviews the system management processes and personnel to help ensure operational continuity of the environment.

**Architecture** — Sun's Backup and Restore Architecture service develops an open architecture for backup and restore which meets the customer's needs and fits their computing environment. This architecture service becomes a foundation for implementation of a comprehensive backup and restore solution.

## **Point-In-Time Copy Services**

This service delivers a customized unassisted backup and recovery solution for the Oracle RDBMS Instances installed on a specific set of servers. This service includes assessment of a customer's business and Service Level Agreements (SLA) requirements; current architecture study of existing storage, backup strategies, and system availability to determine an appropriate archival solution utilizing existing hardware whenever possible; and design and implementation of a Automated Point In Time BackUp Logistics and Logic solution.

## **SAN Architecture and Implementation**

This service complements the Sun StorageTek Array Architecture and Implementation Service that addresses directly attached storage configurations.

The objective of the service is to facilitate the use of Sun StorageTek products in a SAN environment. While some solutions may include non-Sun branded products, one Sun server and one Sun StorageTek array must be included in the proposed solution.

## **Consulting Services**

Sun Client Services can offer multiple levels of consulting services to ensure a smooth data migration from other storage products to the Sun StorageTek 6540 array. Sun Client Services offers a wide range of data management and migration services to accommodate most customer environments and circumstances. Contact local Sun Client Services for further details and quotation for the specific customer environment.

Sun also offers a wide variety of consultative services that can help the customer architect their Sun StorageTek 6540 array into existing storage infrastructures. These services are generally custom-priced engagements that can assist with the design and implementation of larger storage architectures. These services can also assist with analysis of total cost of ownership (TCO), storage migration, comprehensive review of backup and recovery procedures, data replication design and implementation, and security issues.

# Glossary

---

<b>Array</b>	Storage system consisting of trays and controller units. Storage system consisting of a minimum of one controller unit. Also includes one or more slots to house disks drives.
<b>Array hot-spare</b>	Disk that serves as a hot-spare within an array. A reserve disk that can be made available to all virtual disks within an array.
<b>Bandwidth</b>	A measure of the capacity of a communication channel, usually specified in MB/second.
<b>Block or block size</b>	Amount of data sent or received by the host per I/O operation. Atomic read/write operation to/from a disk. Size of data unit that is striped across disks.
<b>Cache</b>	Staging area used to provide higher performance to applications for reads and writes. During reads, the controller unit tries to keep the latest and most often accessed data in cache and also tries to pre-stage cache with future data during sequential accesses. For writes, cache is used to provide delayed writes to the disks. This delays the parity calculations and disk writes during RAID 5 operations. More optimization and advanced staging algorithms thus provide better performance.
<b>Cache hit</b>	Read or write request for data that is already in cache. Therefore, a request can be serviced without needing to go to disk.
<b>CLI</b>	Command Line Interface.
<b>Controller tray</b>	A tray with one or more installed controller units.
<b>Controller unit</b>	The intelligence card that manages RAID functions and failover characteristics for an array or tray, or group of trays.
<b>Copy-On-Write (COW)</b>	The process that Sun StorageTek Data Snapshot software uses to preserve point-in-time data when new data is written to disk. With each write, the system preserves the old data in snapshot reserve space, so that it can re-create the volume as it existed at the time of the snapshot.
<b>DAS</b>	Direct attach storage. Storage directly attached to servers or hosts (as opposed to SAN storage where storage is attached to a network of storage devices).
<b>Data path</b>	The path traveled by data packet — between the host processor and the disk.
<b>Disk</b>	Physical entity that stores data (as compared to a virtual disk, which is a logical grouping of disks or storage extents).
<b>Disk Slot</b>	Slots on trays that house physical disks.
<b>DMP</b>	VERITAS dynamic multipathing.
<b>Drive depopulation</b>	Drive depopulation allows additional spindles to be added to drive trays that are not fully populated with 14 drives.
<b>ECC</b>	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.

<b>Fabric</b>	A group of interconnections between ports that includes a fabric element. A collection of switches and the connections between them.
<b>FC-AL</b>	Fibre Channel arbitrated loop, a loop topology used with Fibre Channel.
<b>Fiber</b>	A wire or optical strand. Spelled <i>fib</i> re in the context of Fibre Channel.
<b>Fibre Channel</b>	A set of standards for a serial I/O bus capable of transferring data between two ports up to 100 MB per second. Fibre Channel supports point-to-point, arbitrated loop, and switched topologies. Fibre Channel can be implemented with either optical fiber (note spelling) or copper.
<b>Fiber-optic cable</b>	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.
<b>F_Port</b>	On a Fibre Channel switch, a port that supports an N_port.
<b>HBA</b>	Host bus adapter. A card that connects a peripheral device to the computer system's I/O bus.
<b>host</b>	Represents a data host and is comprised of 0 or more initiators. A host has a name that is unique to the storage array on which it lives (the uniqueness is enforced by the storage array firmware).
<b>host group</b>	A collection of host objects on a particular storage array. Host groups have names that are unique to the storage array on which they live (the uniqueness is enforced by the storage array firmware).
<b>Host ports</b>	DSP ports attached to hosts or host-facing SAN ports. Any DSP port is capable of being either a host port or a storage port.
<b>Hot-spare</b>	Disk used by a controller unit to replace a failed disk.
<b>Hot-swappable</b>	A hot-swappable 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 configures 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.
<b>In-band</b>	Using the data path between a host(s) and a storage device to transport system management traffic.
<b>Initiator</b>	A port on a Fibre Channel card of a data host. The initiator has a world-wide name (WWN) that is globally unique and a name. Note that because the array allows the use of duplicate names for initiators, the management software can allow duplicate names to be used on the same storage array.
<b>I/O</b>	Input/output.
<b>I/O rate</b>	A measure of a device's capacity to transfer data to and from another device within a given time period, typically as I/O operations per second.
<b>IOPS</b>	Input/output operations per second. A measure of I/O performance, this is commonly used to quote random I/O performance.
<b>IP</b>	Internet protocol. A set of protocols developed by the United States Department of Defense to communicate between dissimilar computers across networks.
<b>LED</b>	Light emitting diode.

<b>LUN</b>	Logical Unit as defined by SNIA. Defines a volume as it is mapped to particular host(s) or initiator(s). Distinguished from a volume in a sense that the same volume can represent a different LUN to different host(s) or initiator(s).
<b>LUN mapping</b>	Assigning volume permissions (read-only, read/write, or none) to a host or initiator.
<b>LUN masking</b>	A technique that prevents all but certain initiators from gaining access to a volume.
<b>Management path</b>	The out-of-band path that connects components of the system to the storage service processor.
<b>Mapping</b>	An association between a storage volume and either a host or a host group that has a Logical Unit Number (LUN). All mappings are implicitly read/write mappings between the volume and all initiators contained inside the host or host group.
<b>Mirrored cache</b>	Redundant copies of data residing in cache — the (write) data residing in cache that has not yet been written to the hard disks is duplicated for failover operation.
<b>Mirroring (RAID)</b>	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.
<b>MTBF</b>	Mean time between failures. A measure of reliability, this is the average expected time between failures of equipment, usually measured in operating hours.
<b>MTBDL</b>	Mean time between data loss. In a RAID system, this is the average expected time between two rapid disk failures that would cause irreparable data loss.
<b>MTTR</b>	Mean time to repair. A measure of availability, this is the average time the system is out of commission to complete a repair process.
<b>Multipathing</b>	Providing two or more physical paths to a given target or device.
<b>Non-disruptive</b>	Does not prevent system or data access at any time during operation. Data path access is not sacrificed, and the host does not see any I/O failure (unless dual points of failure). Availability of management services not included.
<b>N_Port</b>	A Fibre Channel port in a point-to-point or fabric topology.
<b>NL_Port</b>	A port attached to a node for use in all three topologies (point-to-point, arbitrated loop, or fabric).
<b>NVRAM cache</b>	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.
<b>OLTP</b>	On-line transaction processing.
<b>Optical fiber</b>	Any filament of fiber, made of dielectric material, that guides light.
<b>Out-of-band</b>	Using a path other than the data path to transport system management information. Connecting to a management port using an IP network, for example.

<b>Parity</b>	Additional information stored along with the data that allows the controller to reconstruct lost data on RAID 3 or 5 LUNs if a single drive fails.
<b>Partitions</b>	See Storage Domains.
<b>Path failure</b>	The loss of a data or management path.
<b>Path failover and recovery</b>	The process of automatically moving traffic to a different path when a path either fails, or is re-instated.
<b>Point-in-time copy</b>	A frozen copy of a volume's data, as created when taking a snapshot.
<b>Port</b>	An access point on a device for attaching a link.
<b>Primary volume</b>	Used in the context of snapshots, the primary volume is the live data set from which the snapshot copy is made.
<b>Profile</b>	A set of attributes applied to a set of storage in a Common Array designed to help optimize that pool for a particular access pattern and/or level of data protection. Profiles are associated with Common Arrays to define the attributes of the pool.
<b>Protocol</b>	A convention for data transmission that defines timing, control, format, and data representation.
<b>RAID</b>	Redundant array of independent disks. A set of disk drives that 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, availability, and cost characteristics.
<b>RAS</b>	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.
<b>Read-ahead</b>	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.
<b>Reconstruction</b>	The process of rebuilding lost data on a replacement disk after a disk failure.
<b>Redundancy</b>	Duplication for the purpose of achieving fault tolerance. Refers to duplication or addition of components, data, and functions within the array.
<b>Repository Volume</b>	A special volume created as a resource for a snapshot volume.
<b>SAN</b>	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 simplifying management across servers.
<b>Snapshot Volumes</b>	A point-in-time image of a standard volume. Only possible if the premium feature for snapshots is enabled.
<b>SNMP</b>	Simple network management protocol. A simple protocol designed to allow networked entities (for example, hosts, routers) to exchange monitoring information.

<b>Source Volume</b>	A source volume used in a volume copy operation. Can be a standard volume or a snapshot volume. It is a logical entity that has an associated storage profile, and has a collection of storage volumes that all adhere to the storage profile. A Common Array is always scoped by a storage array and has a unique name relative to that storage array.
<b>Storage domain</b>	A logical domain with its own storage, and its own management environment. It is one or more mappings between a host or host group, and volumes. A Storage Domain may contain up to 256 volumes, each volume having a unique LUN id with a value between 0 and 255. Storage Domain is synonymous to the term “Partitioning”.
<b>Storage pool</b>	A collection of disks, virtual disks or storage extents, generally with common configuration, availability, and performance characteristics, that can be carved into volumes.
<b>Storage Profile</b>	A collection of attributes that identifies the storage properties of a storage volume. The attributes will contain values that are applicable at the level of Storage Array, Vdisk, and Volume.
<b>Storage Volume</b>	A volume on a particular Vdisk on a particular storage array.
<b>Striping</b>	Laying data out over a series of disks or virtual disks, allows multiple disk controllers to simultaneously access data, thus improving performance. Spreading or interleaving logical contiguous blocks of data across multiple independent disk spindles. Striping allows multiple disk controllers to simultaneously access data, improving performance.
<b>Stripe size</b>	Total amount of data in a disk stripe, that is, the block size multiplied by number of data disks in the stripe.
<b>Stripe width</b>	Total number of disks in a disk stripe.
<b>Sun StorageTek Data Snapshot software</b>	Dependent style point-in-time copies, using copy on first write technology.
<b>Sun StorageTek Common Array Manager software</b>	Uses virtualization capabilities of the Sun StorageTek 6540 array to simplify storage management, using Common Arrays and application-oriented storage profiles.
<b>Switch</b>	The name of an implementation of the fabric topology. A fabric element that implements a fabric. The fabric element that allows each port of a switch to be connected to any other port on that switch. A collection of switches implement a fabric and provide the network through which any device can communicate with any other device.
<b>Syslog</b>	The internal log file maintained by Sun StorageTek 6540 arrays to track events and alerts as well as informational and notice messages. This log file can be sent periodically to a host server for evaluation using the <code>syslogd(1M)</code> function.
<b>Target volume</b>	A standard volume to which the data on the source volume will be copied during a volume copy operation. Note all data on the target volume is destroyed during the copy operation and when the copy operation is complete the target volume will be read-only.
<b>Telemetry stream</b>	Stream of data generated by monitoring agents.
<b>Throughput</b>	A measure of sequential I/O performance, quoted as megabytes per second (MB/sec.). See <i>IOPS and I/O rate</i> .

<b>Topology</b>	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.
<b>Transfer rate</b>	The rate at which data is transferred, usually measured in megabytes per second (MB/sec.).
<b>Tray</b>	An enclosure containing disks.
<b>Tray depopulation</b>	Trays delivered without the full compliment of disks installed. Allows additional disks to be added to trays that are not fully populated.
<b>Vdisk</b>	See Virtual Disk.
<b>Virtual disk</b>	Any abstraction or collection of disks that appears as a single disk to the device mounting it. is a volume group (a RAIDset) on a particular storage array comprised of similar disk types (either fiber channel or SATA).
<b>Volume</b>	A logical structure created on a Storage Array for data storage. Can be a single physical disk or a virtual disk mapped from one or more underlying extents. These are the typical volumes that users will access from data hosts.
<b>Warm boot device</b>	Bootable on all supported HBAs with storage booted before server booting.
<b>XOR</b>	eXclusive OR. A binary mathematical operation performed on data to produce parity information. In RAID levels 3 and 5, parity is generated from the user data, stored, and used to regenerate lost data if a drive failure occurs.
<b>Zone or zoning</b>	Provided by fabric switches, a function that allows segmentation of node by physical port, name, or address.



# Materials Abstract

All materials are available on SunWIN except where noted otherwise.

Collateral	Description	Purpose	Distribution	Token # or COMAC Order #
<b>Product Literature</b>				
– <i>Sun StorageTek 6540 Array, Just the Facts</i>	Reference Guide (this document)	Training Sales Tool	SunWIN, Reseller Web	SunWIN #480929
– <i>Sun StorageTek 6540 Array, Customer Presentation</i>	Customer Presentation with Notes	Sales Tool	SunWIN, Reseller Web	SunWIN #480927
– <i>Sun StorageTek 6540 Array, Technical Presentation</i>	Technical Presentation with Notes	Sales Tool	SunWIN, Reseller Web	SunWIN #480928
– <i>Sun StorageTek Modular Line Customer Presentation</i>	Customer Presentation	Sales Tool	SunWIN, Reseller Web	SunWIN #429305
<b>References</b>				
– <i>Data Sheet</i>	Data Sheet	Sales Tool	SunWIN, Reseller Web, COMAC	SunWIN #480926
– <i>Sun StorageTek 6540 Array Ordering Guide</i>	Internal Ordering Guide	Sales Support	Sun Documentation Center	819-6139
– <i>Data Replication Strategies</i>	Whitepaper	Sales Tool	SunWIN, Reseller Web	SunWIN #446863
– <i>Storage Performance Benchmarking 101</i>	Whitepaper	Sales Tool	SunWIN, Reseller Web	SunWIN #418102
– <i>Sun StorageTek SAN 4.X WWW Support Matrix</i>	WWW Compatibility Matrix	Sales Tool	SunWIN, Reseller Web	SunWIN #397802

Collateral	Location Information
<b>External Web Sites</b>	
– <i>Sun Web Site</i>	<a href="http://www.sun.com">http://www.sun.com</a>
– <i>Sun Network Storage Main Page</i>	<a href="http://www.sun.com/storage/">http://www.sun.com/storage/</a>
– <i>Sun StorageTek 6540 Array Main Page</i>	<a href="http://www.sun.com/storage/midrange/6000/6500/6540/index.xml">http://www.sun.com/storage/midrange/6000/6500/6540/index.xml</a>
– <i>Sun SAN WWW</i>	<a href="http://www.sun.com/storage/san/multiplatform_support.html">http://www.sun.com/storage/san/multiplatform_support.html</a>
– <i>Sun Customer Ready Systems (CRS) Program</i>	<a href="http://www.sun.com/integration">http://www.sun.com/integration</a>

Collateral	Location Information
– <i>Sun Upgrade Advantage Program</i>	www.sun.com/ibb/upgrades
<b>Internal Web Sites</b>	
– <i>Sun StorageTek 6540 Array Main Page</i>	<a href="http://mysales.central/public/storage/products/midrange/6540">http://mysales.central/public/storage/products/midrange/6540</a>
– <i>Sun Customer Ready Systems (CRS) Program</i>	http://systems.corp/integration
– <i>Sun Upgrade Advantage Program</i>	http://ibb.eng/upgrades
<b>Further Assistance</b>	
– <i>Sales Support</i>	800-ASK-4STK sales.support@sun.com
– <i>Sun StorageTek 6540 Hotlines</i>	800-381-9958 (US) / 650-331-3283 (outside the US)
– <i>Sun StorageTek 6540 Expert Alias</i>	STK6540deals@sun.com

# FAQs

---

## **Q1. What is the Sun StorageTek 6540 array?**

The Sun StorageTek 6540 array is the newest addition to Sun's line of storage arrays, providing unprecedented performance in a family of industry-leading technology. The array controller integrates 4 gigabit per second (Gb/s) Fibre Channel (FC) connectivity, high-performance FC or high-capacity serial ATA (SATA) disk drives, a next-generation XOR engine, and robust storage management to power this high performing solution. With eight interfaces for host or SAN connectivity, support for 4, 2, or 1 Gb/s FC infrastructures, eight redundant drive loops, and flexible cache offerings, the Sun StorageTek 6540 array is a great fit for enterprise environments requiring the highest levels of performance, connectivity, flexibility, scalability and manageability.

## **Q2. What are its key value propositions?**

4 Gb/s FC technology leadership provides investment protection for today's infrastructure and future protection for tomorrow's. Additionally, next-generation 4 Gb/s FC technology enables faster communication between servers and storage devices resulting in faster applications and quicker access to data.

Excelling at IOPS and MB/s, the Sun StorageTek 6540 array is a great fit for both transaction-oriented and bandwidth-intensive applications. Its industry-leading performance drives faster transactional applications, loads large data sets in less time, handles multiple I/O streams for rich media applications, supports larger consolidations, and replicates data in less time.

Multiple cache sizes provide customers with options to best meet their requirements. Larger cache sizes (up to 16 GB per system) are available for applications with a high locality of reference that frequently access the same data and derive the most benefit from larger cache. For general purpose usage and maximum performance value, smaller options are available.

## **Q3. What is the Sun StorageTek 6540 array's performance?**

Transactional — random, small-block I/O

Burst I/O from cache: 575,000 IOPS

Sustained I/O from FC media (reads): 86,000 IOPS

Sustained I/O to FC media (writes): 22,000 IOPS

Bandwidth — sequential, large-block I/O

Sustained throughput from FC media (reads): 1600 MB/s

Sustained throughput from FC SATA (reads): 1550 MB/s

Sustained throughput to FC media (writes): 1300 MB/s

Sustained throughput to SATA media (writes): 1200 MB/s

## **Q4. How much cache does the Sun StorageTek 6540 array offer?**

From 4 GB to 16 GB of cache. Offering multiple cache sizes provides customers with options to best meet their requirements. Cache is not field upgradeable.

## **Q5. What's the target market?**

The Sun StoragTek 6540 array's 4 Gb/s interfaces make it a great fit for all environments — providing current and future infrastructure investment protection. Additionally, the 6998 storage system's improved performance, connectivity, cache offerings and 2/1 Gb/s FC backwards compatibility makes it a great fit for enterprise environments that require the highest levels of performance, connectivity, flexibility, scalability, and manageability.

The Sun StoragTek 6540 array's 4 Gb/s interfaces and outstanding throughput capabilities are intended to appeal to high-performance computing (HPC) and technical computing environments — such as energy, entertainment, biosciences, supercomputing, and scientific research and development — that store and utilize vast amounts of data for high-bandwidth programs and complex application processing. By delivering twice the bandwidth of 2 Gb/s Fibre Channel, 4 Gb/s Fibre Channel can significantly reduce the time it takes to transfer data from one location to another. These customers are expected to be early adapters of 4 Gb/s technology and ideal markets for the Sun StorageTek 6540 array.

Applications that benefit from the 4 Gb/s FC interfaces and high bandwidth can be found in general purpose computing and HPC environments. A few examples include:

- Campus area replication — Compared with 2 Gb/s technology, 4 Gb/s FC reduces the time, exposure and performance impact of creating and updating mirrors within FC distances. With a 4 Gb/s SAN, the remote site is synchronized in roughly half the time.
- Quick recovery — 4 Gb/s Fibre Channel reduces the time it takes to restore backup images from disk-to-disk backups. With 4 Gb/s technology's improved bandwidth performance, restores complete in half the time of 2 Gb/s.
- Moving data between tiers of storage — With 4 Gb/s FC, the time to copy large blocks of data between primary (FC) and secondary (SATA) storage systems can be reduced by half.
- Data set in memory — 4 Gb/s FC reduces the time it takes to load large, memory-resident data sets by up to one half.

Additionally, 4 Gb/s FC simplifies striping for large data analysis — such as simulations, seismic, or research — by reducing the number of HBAs and switch ports.

## **Q6. Why do I need 4 Gb/s when my infrastructure is 2 Gb/s?**

There are several reasons why buying a 4 Gb/s FC system today makes sense. The information explosion continues to drive the need for higher-performance communication. Next-generation 4 Gb/s FC enables faster communication between servers and storage devices resulting in faster applications and quicker access to data.

4 Gb/s FC is the next logical step in technology and is expected to be the dominant interface by the end of 2006. When purchasing a new product, it needs to be based on a technology that is going to be around throughout its useful life. 2 Gb/s FC is going to be phased out long before many of today's purchases are rotated out of production. Buying 4 Gb/s today removes pre-purchase obsolescence — dramatically reducing buyer remorse and setting the stage for end user satisfaction.

4 Gb/s Fibre Channel's backwards compatibility enables it to be seamlessly integrated into existing 1 Gb/s and 2 Gb/s FC infrastructures. This allows users to immediately benefit from improved functionality in their new 4 Gb/s Fibre Channel products, provides them with investment protection of their current environments, and provides future protection when their infrastructure eventually makes the transition to 4 Gb/s.

# Internal Information

## Sun Proprietary — Confidential: Internal Use Only

Sun Microsystems Inc. proudly announces the Sun StorageTek 6540 array with data services! This array is positioned to meet the needs of commercial, business-critical applications that require enterprise-level data protection, and centralized, application-oriented management. This is a storage solution is designed to handle large data sets and enables better application responsiveness with end-to-end 4 Gb/s technology. The Sun StorageTek 6540 array is ideal for customers who require enterprise-level data protection and who value quality and price/performance. These arrays can consolidate storage classes without disrupting critical applications, and manage massive growth in static content without compromising system availability.

Feel free to contact the [Competitive Intelligence](#) group for more up to date information.

## Competitive Information

The Sun StorageTek 6540 array is targeted at upper half of the midrange storage market with demanding performance requirements. Primary competitors are the EMC CX3-80, HDS AMS1000, and HP EVA8000.

<i>Specification</i>	<i>Sun StorageTek 6540 Array</i>	<i>EMC CX3-80</i>	<i>HDS AMS1000</i>	<i>HP EVA8000</i>
<b>Controllers per system</b>	2	2	2	2
<b>No. of host ports per system</b>	8	8	8	8
<b>Host Port speed</b>	4 Gb/s	4 Gb/s	4 Gb/s FC, or 4 Fcand 4 iSCSI @ 1 Gb/s, or 8 NAS @ 1 Gb/s	4 Gb/s
<b>Drive Port speed</b>	4 Gb/s	4 Gb/s	2 Gb/s	2 Gb/s
<b>No. of drive ports per system</b>	8	8		8
<b>Drive types supported</b>	4 Gb/s FC 2 Gb/s FC 3 Gb/s SATA	4 Gb/s FC 2 Gb/s FC 2 Gb/s LC-FC		
<b>Maximum no. of drives</b>	224	480	450 FC, 420 SATA	240
<b>Maximum capacity</b>	112 TB	239 TB	129.4 TB FC 173.8 TB SATA	72 TB
<b>Maximum cache</b>	16 GB	16 GB	16 GB	8 GB
<b>Burst cache reads</b>	575,000 IOPS	275,000 IOPS	231,000 IOPS	200,000 IOPS
<b>Sustained disk reads</b>	86,000 IOPS	Unknown	44,200 IOPS	Unknown
<b>Sustained throughput disk read</b>	1600 MB/s	Unknown	1400 MB/s	1300 MB/s

<i>Specification</i>	<i>Sun StorageTek 6540 Array</i>	<i>EMC CX3-80</i>	<i>HDS AMS1000</i>	<i>HP EVA8000</i>
<b>Total LUNs supported</b>	2048	2048	4096	512

Positioned against these products the Sun StorageTek 6540 array delivers substantial competitive advantages:

### **More Performance Than the EMC CX3-80**

The EMC end-to-end 4 Gb/s FC disk arrays - CX3-20, CX3-40 & CX3-80 were announced in May 2006 as replacements for the CX300, CX500, and CX700 systems. EMC claims these new units have twice the capacity, twice the memory, and twice the processing power of previous generation CX systems. Their key messages are: new architecture; performance, scalability, and availability; ease of use and management; economical; investment protection.

- EMC is showing competitive performance graphs based on an internal EMC test that is not documented or explained, EMC must test with an industry standard such as the SPC-1 and SPC-2 benchmarks to prove performance leadership.
- EMC is telling customers they can offer inexpensive HPC leadership with up to 2700 MB/s from cache. HPC applications typically use inexpensive, high capacity SATA drives. For EMC, this means using their 2 Gb/s LC-FC (FATA-like) drives, which have a maximum data rate of only 1600 MB/s.
- The EMC CX3-80 can support up to 480 drives (>200 TB), yet IDC reports that the average EMC CX700 storage system ships with only 10 TB of storage. If customers are using 300 GB FC disks, this requires only 33 disks. If using SATA, as few as 20 drives are needed. EMC's own sales indicates their new scaling capability is interesting, but irrelevant to most of their customers.
- EMC claims that they can create an inexpensive, fast, and flexible subsystem using LC-FC drives (low-cost FC FATA-like drives). However, the first five slots of the first drive tray must use standard FC drives. In addition, the FLARE OS requires 33 GB of space on each of the first five drives (for a total of 165 GB). Sun StorageTek 6540 array uses only 10 GB of DACstore. EMC uses 16 times more space.
- EMC is touting 16 GB of cache as a big selling point for performance. However, they only provide between 25% and 43% of the memory as actual data cache. What EMC refers to as cache is actually a block of memory that is used for CPUs, Windows XP, FLARE code, management software, and data cache. In actuality, only 6.7 GB of their 16 GB of cache is utilized for data cache. The Sun StorageTek 6540 array's cache is dedicated data cache.
- EMC claims their new UltraScale controller technology — built on PCI Express — expands bandwidth by three times. Going to PCI Express does offer potential increases in internal bandwidth, but the bottom line is that the only measurements that count are throughput (MB/s) and IOPS. The only way EMC's claims of increase performance can be verified is to provide standardized, reproducible benchmarks such as SPC-1 and SPC-2.
- EMC also claims increased redundancy and fault tolerance with n+1+1 redundant power and cooling. This seems like overkill as a double failure of power supplies rarely occurs. In addition, the Sun StorageTek 6540 array's power/cooling units are designed to easily operate both controllers in the system if one unit should fail.
- The Sun StorageTek 6540 system (FLX380) dominated the competition with a number one ranking in the Diagnoses Labs — Storage magazine Quality Awards (see “Quality Awards: Enterprise arrays,” Storage Magazine, August 2005). EMC's CLARiiON systems placed in the bottom half of every category and came in last for features, quality, and technical support. The CX3 systems might give EMC more features, but it's hard to tell if their quality or technical support will show any improvement.

## **More Flexibility and Scalability Than the HDS AMS1000**

Substantially more RAID setting options:

- HDS1000 has one segment size option (64K), prohibiting performance tuning.
- The Sun StorageTek 6540 array offers five segment size options and the ability to non-disruptive change many settings with active I/O.

More flexible volume management:

- HDS1000 only supports horizontal vdisks comprised of contiguous drives.
- The Sun StorageTek 6540 array vdisks can use drives from any location and supports vertical striping.

More flexible data service deployments:

- HDS1000 ShadowImage needs identical source and target RAID level.
- The Sun StorageTek 6540 array can have different target and source settings, ideal for tiered storage requirements.

Advanced drive connectivity:

- HDS1000 uses loop-based drive connectivity.
- The Sun StorageTek 6540 array uses switching-based (SBOD).

## **More Expansibility Features Than the HP EVA8000**

Better system redundancy:

- EVA8000 uses a single cable to mirror cache.
- The Sun StorageTek 6540 array uses dual, redundant cabling.

Highest built-in availability:

- EVA8000 (VRAID) is a volume manager with no control over the underlying vdisks, require planning and maintenance.
- The Sun StorageTek 6540 array provides complete control for vdisks and supports data intact reconfigurations and automatic rebuilds.

Smaller footprint:

- EVA8000 scales up to 240 drives in two 42 RU cabinets.
- The Sun StorageTek 6540 array scales up to 224 drives with 16 drives per 3 RU tray instead of HP's 14 drives per 3 RU tray.

Are HP's new 4 Gb/s controllers capable of auto-negotiating to 1 and 2 Gb/s? The 2 and 4 Gb/s controllers are not supported in the same system, so both controllers must be upgraded.

Disk groups cannot contain both FC and FATA disks. Disk groups must contain only one type of disk.

Up to 200K IOPS and up to 1300 MB/s throughput per controller pair with four FC loop switches.

Better provisioning and management:

- EVA8000 relies solely on RAID settings to provision storage.
- The Sun StorageTek 6540 array delivers provisioning by RAID type and/or application type.