

Sun™ Storage F5100 Flash Array Just the Facts



Version 1.2 - 10/04/2009



Sun Microsystems, Inc. 4150 Network Circle, Santa Clara, CA 95054 USA **Phone** 1-650-960-1300 or 1-800-555-9SUN **Web**
sun.com

© 2008 Sun Microsystems, Inc. All rights reserved. Sun, Sun Microsystems, the Sun logo, StorageTek, IPX, JVM, ONC+, NFS, WebNFS, Java, Netra, Sun N1, ONC, Solaris, Sun Fire, Sun StorEdge, SunLink, Sun Global Services, SunSpectrum, SunSpectrum Silver, SunSpectrum Gold, SunSpectrum Platinum, Sun Enterprise, and Netra 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. UNIX is a registered trademark in the United States and other countries, exclusively licensed through X/Open Company Ltd. AMD, Opteron, the AMD logo, the AMD Opteron logo are trademarks or registered trademarks of Advanced Micro Devices.

SunWIN: 569042

Sun Storage F5100 Flash Array : Just the Facts
Sun Confidential: Internal and Sun Channel Partners Use Only



Table of Contents

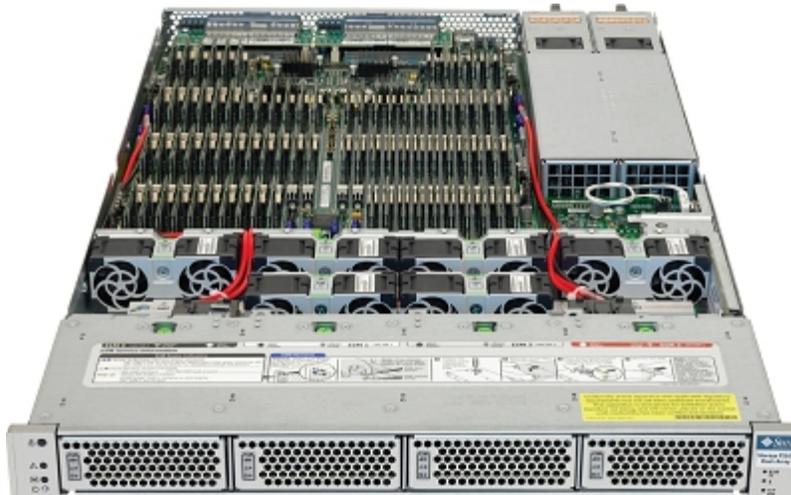
Table of Contents

Positioning: Sun™ Storage F5100 Flash Array	6
Introduction.....	6
Sun FlashFire Product Portfolio.....	7
Sun™ Storage F5100 Flash Array	8
Sun StorageTek™ Common Array Manager software	9
Key Features, Functions, and Benefits: F5100 Flash Array.....	9
Key Messages: Sun™ Storage F5100 Flash Array.....	13
Product Highlights and Key Take Aways.....	16
Target Customers.....	17
Target Markets.....	19
Target Applications.....	20
Product Availability.....	21
Enabling Technology for Sun™ Storage F5100 Flash Array	22
Solid-state Flash Technology.....	22
Interface Technologies Overview.....	22
Sun StorageTek™ Common Array Manager (CAM) software.....	23
Sun Solaris™ ZFS.....	23
System Architecture: Sun™ Storage F5100 Flash Array	24
Architecture.....	24
External features and connections.....	27
Sun Flash Module.....	27
F5100 management — Sun StorageTek CAM Software.....	29
Performance Sun™ Storage F5100 Flash Array	30
Performance Data: Four Corners Test.....	30
Performance Considerations: 4K Block Alignment	32
Performance Considerations: Software.....	34
Reliability, Availability, and Serviceability (RAS): F5100	35
RAS features.....	35
Solid-state Flash technology	36
Super Capacitors - ESMS.....	36
Specifications: Sun™ Storage F5100 Flash Array	38
System Requirements: Sun™ Storage F5100 Flash Array	39
Supported Host Platforms	39
Supported SAS HBAs	40
Supported Operating Systems	40
Supported Sun Enterprise and Application Software	41
Management Host System Requirements.....	42
System Configuration and Deployments:	
Sun™ Storage F5100 Flash Array	43
Flash Module Upgrades and Configurations.....	43
Min and Maximum Host Connect/HBA s.....	43
Configuring F5100 for Best Performance.....	44
Configuring the F5100 for Highest Capacity.....	45
Configuring the F5100 for High Availability	46
Deployment and Best Practices	49

Deployment Considerations.....	49
Configuration examples.....	51
Performance tuning.....	53
Index mirroring, migration, operation, and validation.....	54
Summary.....	55
Ordering Information: Sun™ Storage F5100 Flash Array.....	56
WebDesk Configurator	56
Configuration Options.....	56
Configuration Matrix / Marketing Part Numbers.....	57
Supported HBAs.....	58
SAS Cables.....	58
Service and Support: Sun™ Storage F5100 Flash Array.....	60
Warranty.....	60
Support.....	60
Services.....	60
SunSpectrumSM Service Plans.....	61
Professional Services.....	61
Glossary.....	63
Resources: Sun™ Storage F5100 Flash Array.....	64
Internal Information: Sun™ Storage F5100 Flash Array.....	66
Partner Information.....	66
F5100 Flash Array Sales Program Overviews.....	69
Competitive Information	70
Key Competitors: Systems and Storage Vendors	71
Competitive Comparisons: Solid State Arrays and Flash Cards.....	72
Frequently Asked Questions:	
Sun™ Storage F5100 Flash Array	74

Positioning: Sun™ Storage F5100 Flash Array

World's fastest and eco-efficient solid-state Flash Array with unprecedented 1.6 million IOPS performance in just 1.75 inches of rack space



Introduction

Today more and more applications, especially databases, are being choked by disk drives that can no longer keep up with CPU performance, causing latencies, I/O bottlenecks. Disk access speeds remain one of the major bottlenecks. While CPU performance has been doubling every year and servers continue accelerating their performance capability with multicore CPUs, disk drive performance has not kept pace. Today, even the high performance 15K rpm mechanical disk drives are over 260 times slower than what the CPUs are capable of delivering causing application performance to suffer from storage latencies and I/O bottlenecks.

The traditional approach of using large numbers of mechanical disk drives to address growing storage performance needs can greatly increase power, cooling, and space costs. IT managers are looking for more cost-effective and highly scalable storage solutions that can quickly accelerate application performance while also reducing operating costs.

The Sun™ Storage F5100 Flash Array with its blazing record performance offers a new approach to solving the storage performance issues that slow down applications while reducing operating cost.. With well over 1 million IOPS in just 1RU of space and 300 watts of power, the system can help eliminate storage bottlenecks, decrease latencies, improve response times and allow you to increase transactional scalability with up to 100X less power and space versus traditional disk drive solutions.

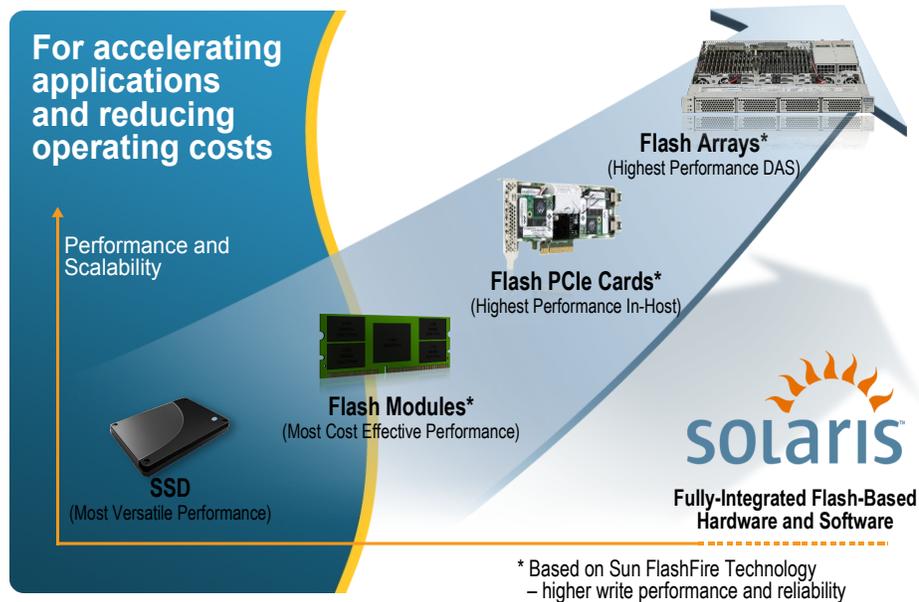
Sun Storage F5100 Flash Array : Just the Facts
Sun Confidential: Internal and Sun Channel Partners Use Only



Sun FlashFire Product Portfolio

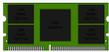
The Sun Storage F5100 Flash Array is one many solid-state storage products in the Sun's FlashFire portfolio of solid-state flash products that can help you meet the challenges of data-intensive computing and faster application response times with greater speed and efficiency.

The Sun Flash Product Portfolio



Sun offers the broadest portfolio of highly reliable enterprise-grade solid state products, including SSDs, Flash Modules, Flash PCIe cards and ultra high performance Flash Array.

Sun Flash Products for Turbocharging your Applications

 1.6M IOPS, 1.92TB	<ul style="list-style-type: none"> • Large Databases (Oracle, MySQL, Sybase, DB2, etc) • I/O intensive HPC applications • Financial modeling • Scale Out Databases • Storage grids • Hybrid Storage Pools • Data logs • Use in Servers & Blades • Local cacheing & logs • Hybrid Storage Pools <p>Use in Servers and Storage Systems Tier "0" storage, read oriented Hybrid Storage Pools</p>
 Over 100K IOPS, 96GB	
 Over 25K IOPS, 24GB	
 35K IOPS, 32GB	

Sun Storage F5100 Flash Array : Just the Facts
 Sun Confidential: Internal and Sun Channel Partners Use Only



Sun™ Storage F5100 Flash Array

The Sun Storage F5100 Flash Array is the world's fastest and solid-state flash array delivering a record 1.6 million read and 1.2 million write IOPS performance with only 300 watts of power and 1 rack unit (1.75 inches) of space. That equals the I/O performance of over 3,000 high performance disk drives with over 40,000 watts of power.

The array reduces latency and eliminates storage-related I/O bottlenecks that slow down application response times. It's optimized for heavy read and write I/O workload applications, such as Oracle OLTP databases, indexes, hot tables, and data structures, where I/O performance, high transactional throughput, and low latency are the most critical to improve business productivity, efficiency and cost.

The F5100 offers customers the best, most efficient way to accelerate Oracle and other database performance. It's a simple, high-performance, eco-efficient solid-state storage solution that doesn't require expensive tuning or space. At a glance:

- World's fastest Flash array with over 1M IOPS in 1RU
- Equals the I/O performance of over 3,000 disk drives
- Accelerates database applications and eliminates bottlenecks
- Improves application response times up to 10x
- Slashes energy consumption by up to 100x compared to spinning disks
- Provides up to 2 TB of solid-state Flash storage
- SAS domains are provided to match performance to manageable chunks to servers.

The Sun F5100 Flash Array comes in a streamlined 1.75 enclosure that contains up to 80 solid-state Flash Modules (solid-state drives) with Common Array Manager software that provides unified management and monitoring via a single storage management window across a wide range of operating systems.

Basic Data: F5100 Flash Array (full capacity)	
Storage Type	Solid-State Enterprise-Grade SLC Flash
Random I/O Performance (4K)	1.6 million read 1.2 million write
Sequential Bandwidth (1M)	12.8 GB/s read 9.7 GB/s write
Capacity	Up to 1.92TB solid state
Domain	4 x SAS Domains
Interface/Ports	SAS-1 / 16 x 4-wide ports (64 paths)
Size	1 RU (1.75 inches)
Power (average r/w)	300 watts active 200 watts idle
Power backup	Integrated, super capacitor based Energy Storage Modules
Management	Common Array Manager software (included)
Operating Systems	Solaris, Open Solaris, Windows , Linux (Red Hat)

Basic Data: F5100 Flash Array (full capacity)	
Warranty	1 year

Sun StorageTek™ Common Array Manager software

Sun StorageTek™ Common Array Manager software (CAM), included with the F5100 Flash Array, is a browser-based, application-oriented management console. Through it, all of the array's functionality is accessible from almost anywhere. This intuitive, centralized management interface needs only a secure Internet connection and a browser window to be at the user's fingertips. The interface provides powerful, granular controls for expert users, as well as indispensable ease-of-use for those less familiar with storage management. Whether the task is replicating between sites, running snapshots, or configuring for specific application workloads, quick wizards and menus get users there fast.

CAM provides both a browser interface and a command-line interface (CLI) for configuring and managing arrays on an external management host or a host directly attached to storage. In addition, CAM provides a small relay agent (also known as a proxy) which allows the aggregation of many hosts directly attached to the F5100 to be managed by an instance of CAM software installed on a dedicated management host. The Common Array Manager software enables online administration, a consistent interface across all operating systems, and the ability to monitor and manage one or all F5100 arrays from any location on the network. CAM also provides SAS Zoning for the F5100.

Key Advantages over Disk: F5100 Flash Array

In addition to greater reliability and solid-state robustness, the Sun Storage F5100 delivers massive amount of I/O performance with a fraction of the power, space and cost of what it would take otherwise with high performance disk drives.

F5100 Advantages

1.2 Million IOPS Comparison			
	F5100 Flash Array	Disk Drives (15K rpm)	F5100 Advantage
Devices	1 unit	3,450 drives	
Power	300 watts	58,650 watts	200X
Space	1 RU	604 RU	600X
\$/IOPS	\$ 0.13	\$ 1.30	10X
Power Cost (3 yrs)	\$ 788	\$ 158,200	200X
Space Cost (3 yrs)	\$ 112	\$ 65,920	600X

Unprecedented Performance and Eco-efficiency

- Accelerate applications, improve business response
- Reduce power, cooling and space costs
- Improve productivity and operational efficiency
- Reduce TCO

Key Features, Functions, and Benefits: F5100 Flash Array

The Sun Storage F5100 Flash Array delivers the performance of over 3,000 disk drives while consuming 100x less power and space than traditional disk solutions. Plus, it provides up to 2 TB of solid-state Flash storage and easily scales to meet even the most demanding needs.

Feature	Function	Benefit
Unprecedented Record Performance	Over 1.6 million read and over 1.2 million write IOPS (4K) performance in only 1 RU of space - equaling that of over 3000 disk drives	Accelerates application performance, such as databases, instantly and efficiently without costly tuning. Reduces latency and eliminates storage bottlenecks. Improves business response and transactional scalability. Improved productivity and efficiency. Reduces cost, eliminate disk "short stroking" and unnecessary over purchase to address growing performance needs as in the past.
Unmatched Power and Space Efficiency	Over 1.6 million maximum delivered IOPS (4K) with only 300 watts of power and 1RU of space - Less than 1/100 the space and power of disk drives for same performance level	Lower power and cooling costs. Less space. Greater eco-efficiency and much lower operating costs.
High Reliability	Robust enterprise-class non-volatile solid-state SLC Flash technology, advanced ECC, advanced wear-leveling, high write endurance, redundant power and cooling.	Years of trouble free operation and durability exceeding that of hard disk drives. Less susceptibility to shock, vibration and other environmental conditions than mechanical disk drives.
Integrated backup protection	Integrated super capacitor based Energy Storage Modules for write cache and meta-data write-through protection	Protection of critical applications against data loss in the event of sudden power loss.

Feature	Function	Benefit
Easy to Use and Deploy	Compatibility with existing applications and operating environments with ultra-high performance disk like emulation without need of special software of drivers	Easy to install and use within existing applications and computing environments. You can turbocharge your IO intensive database applications quickly and efficiently without expensive tuning, training or software changes
Ease of management	CAM (Common Array Management) Software provides a common, simple-to-use interface for configuring and managing all your Sun Storage arrays.	The software takes the complexity out of storage management, enabling you to accelerate deployment, simplify your environment, and improve utilization. The Common Array Manager Service Advisor also offers proactive health checking and supports quick time to service. System error messages are connected directly to specific repair procedures in Sun's knowledge base, making common repairs simple and easy
Configuration flexibility	Storage Domains and SAS Zoning	Four internal domains with up to 20 Flash Modules each provide configuration flexibility and enables mirroring protection. Zoning (CAM Access Configurations) enables up to 16 host servers (HBA) to access Flash Modules individually

<i>Feature</i>	<i>Function</i>	<i>Benefit</i>
Scalability to meet all performance and capacity needs	Three convenient performance / capacity configurations (0.48/0.96/1.92TB) with upgrades in a single unit. Multiple units in a rack for most demanding needs performance and availability needs	Flexible configurations and upgrades to satisfy all performance and capacity needs for today and in the future
High bandwidth and server connectivity	16 x 4-wide SAS ports, or 64 channels	Up to 16 direct server connections for maximum throughput
Sun StorageTek™ Common Array Manager (CAM) Software	Application-aware storage pools: fast, easy application deployment Proactive monitoring: dramatically reduced service complexity Centralized support for a range of Sun Storage offerings	Less training, reduced operations costs Quicker deployment time Quick service for maximum uptime Saves time; no need to go from one management interface to another when managing multiple Sun modular arrays; no need to re-train staff from one array model to the other or when upgrading Ensures maximum utilization of storage capacity and complete control over a rapidly growing storage environment

Key Messages: Sun™ Storage F5100 Flash Array

Sun Storage F5100 Flash Array, is the world's fastest, most power and space efficient solid-state flash array for accelerating database applications and reducing operating costs.

The F5100 delivers unprecedented performance, power and space efficiency. Based on enterprise-class SLC Flash technology, the F5100 delivers well over 1 Million IOPS performance and up to 1.92TB of solid-state capacity with less than 300 watts of power and 1U (1.75 inches) of rack space. - equaling the IO performance of over 3,000 disk drives with 1/100th the space and power. By reducing latency and eliminating I/O bottlenecks, the F5100 quickly and efficiently accelerates application performance and increases transactional scalability for improved business response, SLAs, and lower TCO. The F5100 is ideal for I/O intensive applications, such as Oracle OLTP with heavy IO reads and 4K block aligned writes.

It allows customers to quickly accelerate their I/O intensive applications/workloads and optimize their overall storage infrastructures for the best performance, eco-efficiency, reliability, and cost savings.

- **The Sun Storage F5100 Flash Array can accelerate your databases, such as Oracle or MySQL by up to 4 times or more.** It delivers world-record performance with over 1M IOPS, optimized for heavy workload applications such as Oracle OLTP and SAS databases where I/O performance and low latency is most critical. Users can instantly speed up database performance for best I/O and highest transactional throughput, and keep up with rapid growth and scalability with non-disruptive adds, moves, and updates.
- **Reduce costs by 99% over conventional storage systems while realizing significant power, cooling, and space efficiencies.** With the best price/performance in its class, this flash array uses only 1/100 of the space and power of traditional servers and storage for the same performance level, chalking up industry-leading \$/IOPS, watts/IOPS, and space/IOPS. It offers over 40X better IOPS/\$ value than disk.
- **Reduced risk, increase uptime and service levels make this the most reliable flash array product on the market.** Based on advanced enterprise SLC Flash Module technology, the F5100 Flash Array has a unique enterprise system design for maximum performance, endurance and reliability with over 2M MTBF hours Flash Modules –2x that of most disk drives. A health check monitoring and reporting feature is included: Common Array Manager (CAM) software.
- **Sun makes it easy to deploy the Sun Storage F5100 Flash Array** with services that can help mitigate the risk of downtime, data loss, and costly delays.

KEY MESSAGES AND EVIDENCE:

The Sun Storage F5100 Flash Array is the world's first enterprise flash array and delivers unprecedented performance with well over 1 million IOPS in just 1RU of space with 100X better power and space efficiency than traditional disk. It allows customers to quickly accelerate their Oracle and other database applications and optimize their overall storage infrastructure for the best performance, eco efficiency, reliability, and cost savings.

- The Sun Storage F5100 Flash Array is the best and easiest way to accelerate your applications, allowing you to reduce Oracle database I/O service times by up to 15X, thus improving throughput and transaction times

- The Sun Storage F5100 Flash Array delivers world record performance with over 1M IOPs – equaling the IOPS of the largest multi-rack enterprise storage systems at a fraction of their power, space, and cost.
- The Sun Storage F5100 Flash Array is optimized for IO intensive workload applications such as Oracle OLTP where high performance and low latency is most critical.
- The Sun Storage 5100 Flash Array provides shorter read service times (1-2ms vs 5-15 ms)-- 1/15 the latency of hard disk drives-- to accelerate database performance and productivity, improve application response, and reduce costs associated with database tuning and deployment.
- Customers need to do less query and database tuning to achieve the same client-side response time
- Reducing IO service time and increasing read throughput from media also allows a system to scale to a larger number of clients and transactions
- Performance optimized for high reads and 4K block aligned writes
 - Instantly speed up database performance for best IO and highest transactional throughput
 - Sun Storage F5100 Flash Array as a database accelerator allows you to put your indexes and hot files on solid state flash for maximum performance and minimum latency, and cold files on mechanical disk drives for lowest storage cost in order to optimize your database response times and storage efficiency.
 - Provides much greater transactional scalability and efficiency with over 400x IOPS per GB throughput advantage over disk drives
 - Works with native OS SVM, Oracle VSM, and Veritas VxVM for host based mirroring and file management
 - Faster database performance by accelerating their indexes, hot tables and data structures for best performance and lowest latency
 - Unleashes system and applications previously constrained by high storage latencies and IO bottlenecks to allows customers do do more work with increased productivity and efficiency.
 - Reduces computing costs and increases server efficiency by minimizing storage related IO bottlenecks.
 - Keep up with rapid growth and scalability, minimize business impacts and maintain availability with non-disruption adds, moves, and updates.
 - Easy to add to to existing and new storage infrastructures for instant performance boost and scalability
 - Take advantage of Solaris ZFS for automating data management, enhancing data protection and optimizing the storage hierarchy for maximum performance and lowest cost
 - Quickly optimize your storage infrastructure for maximum performance and efficiency, accelerating your application and improving your productivity while reducing storage and operating cost
 - Unleash your system capability with quick and easy performance boost and scalability

- Manage multiple storage and flash arrays using a single management framework (CAM)
 - Easily scale performance and capacity without the cost or inconvenience of adding DRAM and disk.
 - Improve storage performance, increase operational productivity and efficiency and reduce cost
 - Performance scales from 300K IOPS to 1.6M IOPS (4K) in a single 1RU unit. Capacity scales from .48TB to 1.92TB in a single 1RU. In theory, 42 F5100 Flash Array units can fit into a single rack for up to 50M IOPS and 80TB of solid-state capacity to better meet your most demanding performance scalability and business growth needs
 - Reduce or all together eliminate storage tuning with partial or complete data sets
-
- Gain up to a 80% cost savings over conventional storage systems while realizing significant power, cooling, and space efficiencies with the Sun Storage F5100 Flash Array
 - Gain up to 1.6M read and 1.2M write IOPS performance (equaling that of ofer 3000 disk drives) with just 300W of power and 1RU of space
 - The most power and space-efficient solid state array – Scaling from 1million IOPS performance and 2TB solid state capacity in 1RU to multiples of that in a single rack
 - Unique design for maximum performance and 2M MTBF hour Flash Module reliability
 - Reduce operational costs associated with database tuning and deployment
 - Increase data center efficiencies (Cost/TCO) with greater storage performance efficiency with less power, cooling, space and cost.
 - Best price for performance in its class
 - Scale performance and capacity while lowering cost of servers and storage
 - The Sun Storage F5100 Flash Array uses only 1/100 the space and power (200W idle, 300W active) of traditional servers and storage for the same performance level
 - Industry leading \$/IOPS, Watts/IOPS and Space/IOPS (Over 40X better IOPS/\$ than disk drives)
 - Reduce total storage costs by replacing expensive and inefficient disk drives with a more cost effective Hybrid of 5100 Flash Array and high-capacity disk drives that maximizes storage performance and minimizes storage cost
 - Immediate benefits in new or existing environments with no data migration or extensive tuning needs
 - Reduce risk, increase uptime and service levels with the Sun Storage F5100 Flash Array, the most reliable flash array product on the market
 - Based on advanced enterprise SLC Flash Module technology designed to maximum longevity and reliability
 - A solution designed from ground up for highest level of reliability with

advanced wear leveling, bad block mapping, data protection, RAID, mirroring, solid state robustness with Flash Module reliability exceeding that of HDD with over 2 million hours MTBF

- Unique system design for maximum reliability, exceeding that of traditional enterprise storage
- Reduce risk and increase uptime and service levels
- Unique super capacitor backup protection of write cache in case of power failure; no need for batteries
- Unique enterprise system design for maximum performance, endurance and reliability with over 2M MTBF hours (2x that of most disk drives) and more than 6 years of continuous use (at 50%/50% write/read cycles)
- Health check monitor and reporting via CAM software (Common Array Manager), included for simple, secure, centralized configuration management and monitoring
- Automatically and silently detect, diagnosis, and resolve problems leading to increased uptime and service levels when used with Solaris
 - Quickly detect & diagnose problems with an extensible set of agents which can automatically take faulty components offline (Solaris Fault Management Architecture - FMA)
 - Provides diagnostic message links to articles in Sun's knowledge base which guide administrators through easy to understand corrective tasks when human intervention is required
 - Automatic data integrity checking & correction with block level checksums. When a corrupt block is identified, Solaris ZFS™ software automatically repairs it with another copy, providing for self-healing capabilities at the block level and preventing silent data corruption.

Product Highlights and Key Take Aways

The Sun Storage F5100 is a new class of storage products based on non-volatile solid-state Flash technology designed for maximum performance, reliability and eco-efficiency.

F5100 Highlights

- Unprecedented performance - over 1.6M read and 1.2M write IOPS (4K)
- Over 12.8GB/sec sequential read, 9.67GB/s sequential write (1M)
- Low power - 300 watts of active power
- Small Space - 1U (1.75-inches) of rack space
- Up to 1.92TB of solid state capacity in just 1RU
- Three configurations: 0.48TB, 0.96TB and 1.92TB, upgradable to full
- Capacity upgradable up to full 1.92TB with easy to install 24GB Flash modules
- Disk emulation - high compatibility, no special drivers required
- Direct SAS connectivity with up to 16 servers (16 HBAs with zoning)
- 16 x 4-wide SAS ports for maximum throughput
- 4 storage domains for protection and flexibility
- SAS Zoning

Sun Storage F5100 Flash Array : Just the Facts
Sun Confidential: Internal and Sun Channel Partners Use Only



- CAM (Common Array Manager) Software - Comprehensive management tool for simple, secure, centralized configuration management and monitoring
- High reliability - enterprise-class SLC Flash, advanced wear-leveling, high write endurance, redundant power and cooling
- Integrated cache protection with super capacitor based 4 energy storage modules
- Supports ZRAID and host mirroring (between units or internal domains)

Redefining storage performance

The Sun Storage F5100 Flash Array redefines storage performance and efficiency with an incredible 1.6 million read and 1.2 million write IOPS in just 1.75 inches of rack space. It sets new benchmark records for IOPS/\$, IOPS/watt and IOPS/space. I/O intensive database applications with heavy I/O read and 4K block aligned write workloads can take advantage of the Sun Storage F5100 Flash Array to accelerate their performance, improve response times and increase transactional scalability while reducing power and space costs.

Breakthrough economics and eco-efficiency

The Sun Storage F5100 Flash Array offers a breakthrough in storage economics with industry leading price/performance, power/performance and space/performance efficiency. By using the Sun Storage F5100 Flash Array to address your performance needs, you can supplant costly and inefficient 15K RPM disk drives, which drive up power and storage costs. This can help you significantly reduce both storage costs and operating costs, resulting in much lower TCO.

High reliability

The Sun Storage F5100 Flash Array provides a very high level of reliability with an all solid state durable design consisting of 80 non-volatile enterprise-class SLC Flash Modules with MTBF exceeding that of enterprise disk drives. The 24GB flash modules internally have 32G of pages for advanced wear leveling, bad block mapping, performance pipelining and page erase management for highest level of write endurance, reliability and longevity. The system includes redundant power and cooling, with integrated super-capacitors based data protection and to help prevent interruptions in write operations should a power failure occur. Host based mirroring can be used for mirroring internal domains or flash arrays.

The flash array also takes advantage of data integrity features in the ZFS™ file system, including ZFS RAID and automatic data integrity checking and correction with block level checksums. When a corrupt block is identified, self-healing features in ZFS software will direct another copy to be written as an automatic repair, thus helping to prevent silent data corruption

Easy to deploy and manage

The Storage 5100 Flash Array appears as a normal storage device enabling easy deployment with instant performance benefits into new or existing environments, whether Solaris, Windows or Linux. The flash array can be managed together with other Sun storage arrays using Sun StorageTek™ Common Array Manager software, which provides a common, simple-to-use management interface. Common Array Manager Service Advisor technology also offers proactive health checking and supports quick time to service. System error messages are connected directly to specific repair procedures in Sun's knowledge base, making common repairs simple and easy.

Target Customers

The primary customers for the Sun Storage F5100 Flash Array are Sun network storage direct accounts, storage-only resellers, solution resellers, server resellers, OEMs, distributors, and system integrators. The F5100 meets the needs of a variety of users, as shown in the table below.

<i>Individual User</i>	<i>Buying Needs</i>
Corporate Buyer	Improve business critical application performance for better business response and competitiveness, improve productivity and lower cost
Database Administrator	Improve database performance and SLAs
Storage Administrator	Ease of implementations and management
MIS Manager	Flexible, configurable, scalable and reliable

Audience

Primary Individual Targets: CIOs, CTOs, IT directors/managers/architects, system administrators, storage administrators, database/application administrators, and datacenter operators/managers

Others (Influencers): industry analysts, tech savvy Line-of-Business Managers

Talking with CTOs

Tell them they can reduce overall storage costs and improve application performance by eliminating storage I/O bottlenecks, like getting the performance equivalent of over 3000 high-performance disk drives in 1 RU of space, using less than 300 watts of power.

Talking with IT architects / storage administrators

Point out that they can accelerate application performance, increase service levels, and improve uptime by eliminating storage latencies and I/O bottlenecks — using a very reliable solid state flash array that delivers over 1M IOPS in 1 RU of space. They'll like hearing about the performance equivalent of over 3000 disk drives using less than 300 watts of power and 1 RU of space. Tell them that the Sun Storage F5100 Flash Array reduces the need to purchase unnecessary disk drives to meet performance SLAs. Also the simplified in-band management with Common Array Manager software. Suggest that they take advantage of automated data management with Solaris ZFS or Oracle ASM.

Talking with ISV application providers

Tell the ISV business side that the endorsement of the F5100 Flash Array helps them to become the early majority in tech refresh deals where customers are considering infrastructure consolidation and performance gains in virtualization projects. This means new ISV license revenue in many cases. Talk with the ISV technical side about how endorsement of the F5100 Flash Array opens the door to innovation. ISV designers can review the architecture of the current application to take maximum benefits from a large memory pool, and they can explore metadata architecture and/or service-oriented architecture by looking at the data flow as a natural step in the services flow.

Talking with database administrators (DBAs)

Tell them how the F5100 Flash Array can improve their database application performance, improve their delivery on SLAs, reduce their storage footprint, and even reduce software

licensing costs. Tell them about the enterprise-class reliability that can reduce risk, increase uptime, and increase service levels. The F5100 is the first enterprise-class storage array with no moving parts — only solid state robustness and durability. They can eliminate a lot of their mechanical disks and improve their overall storage efficiency and performance.

Moreover, they'll like the idea of world-record performance using 1/100 of the power and space required by disk for equivalent performance — as well as industry-leading lows in \$/IOPS, W/IOPS, and space/IOPS. And Solaris sites will also be able to use the FMA feature, with which they can automatically and silently detect, diagnose, and resolve problems to further improve their service levels and uptime. These customers can take immediate advantage of the F5100 Flash Array with no change in application code, and turbo-charge their applications instantly.

> **Note:** For benchmarking data on the F5100 Flash array, go to blogs.sun.com/BestPerf/entry/1_4_million_iops_in

Strategy

An appropriate sales strategy for the Sun Storage F5100 Flash Array will be to lead with these four key advantages of the array:

- It instantly accelerates application performance and reduces energy cost for database applications.
- It dramatically improves TCO and storage efficiency by eliminating the need for high-performance disk drives, resulting in reduced power consumption and space requirements.
- It offers greater future scalability so it supports the customer's growth needs w/out investment in lots of new hardware.
- It unleashes application potential by matching the customer's storage tier throughput with the application's speed so it can operate at an optimal level.

Target Markets

Where to Aim

It will be best utilized by datacenter customers to accelerate large databases and database-driven applications. Aim at datacenter customers with I/O-intensive workloads that are bottlenecked by storage performance, who are seeking reliable means not only to improve performance but also to reduce costs for storage power and operations. Since the F5100 is performance optimized for I/O intensive reads and 4K block aligned (multiples of 4K) writes, focus on those applications and workloads that support this. Furthermore, with the F5100 currently not supporting clustering, target single instance databases and not clustering or Oracle RAC

The sweet spot for the Sun Storage F5100 Flash Array is in structured, highly-transactional database application environments such as Oracle and SAP, particularly large datacenters where high speed is required. Propose the F5100 Flash Array to customers with such environments who are challenged with performance bottlenecks. The F5100 Flash Array can accelerate I/O-bound indexes, intent logs, hot tables, and other bottleneck areas.

The F5100 Flash Array will provide an instant performance boost for all database applications, whether they are in new or existing environments. It will allow customers running applications on Windows, Linux, or Solaris to see the immediate benefit of faster performance.

>**Note:** For more in-depth technical information documenting best uses for the F5100 Flash

Array, see the white paper "Lightning Flash Array: Storage Considerations for Accelerating Databases" and the BluePrint "Accelerating Databases with the Lightning Flash Array" at http://www.sun.com/storage/disk_systems/unified_storage/resources.jsp.

Below is an example of the end-user scenarios for which the Sun Storage F5100 Flash Array is suitable, as well as a listing of the key features that apply to each scenario.

End-user Scenario	Key Features
Poor application performance not meeting business needs and SLA	The F5100 delivers unprecedented world record performance exceeding 1.6M read and 1.2M write IOPS to reduce storage latencies and I/O bottleneck that can choke application performance
Excessive power and cooling cost	The F5100 offers the highest power, cooling and space efficiency delivering the performance equivalent to over 3,000 high performance (15K) disk drives with 1/100 th their power and space. The F5100 offers the best IOPS/\$, IOPS/watt and IOPS/space efficiency of any storage product
Performance scalability not meeting business growth and needs	The F5100 not only delivers well over 1.2 million IOPS in a single RU. Multiple F5100 Arrays can be configured in a single rack to satisfy the most demanding business growth needs today and the future.
Escalating storage costs with a shrinking budget	The F5100 allows customer to re-architect their storage hierarchy for greater performance and much lower cost. By using the F5100 for their performance needs, they can then deploy much lower cost capacity disks for their archive storage needs thus allowing them to get the same or greater performance and capacity with up to 80% lower overall storage cost. This hybrid storage deployment is much more cost effective in term of OpEx and CapEx than the traditions all disk storage solution that often requires overpurchase and underutilization of 15K disk drives with "short stroking" in an attempt to increase performance.

Target Applications

The primary target market for the F5100 are database applications, such as Oracle OLTP, large MySQL install and other I/O intensive applications optimized for 4KB (and multiples of 4KB) block writes. Most databases perform I/O on 8KB boundaries.

The F5100 Flash Array offers the best, most efficient way to accelerate Oracle databases - a simple, high performance, eco-efficient solid state storage solution that doesn't require expensive tuning, power or space. With over 1 million IOPS performance and 100X less power and space than traditional disk solutions, the 5100 decreases latencies, improves SLAs and increases transactional scalability at much lower cost than conventional disk drives

Secondary markets include HPC applications optimized for 4KB block operations requiring

highest level IO performance, scalability and low latency with minimum power, cooling and space.

Common needs associated with all F5100 target markets include

- Highest level of I/O performance
- Low latency
- Greater transactional scalability
- Lower power and cooling
- Lower TCO

Applications Not Ideal for Sun Storage F5100 Flash Array at this time:

1. Small block writes. The F5100 is performance optimized for heavy read I/Os and 4K aligned (4K and multiples of 4K) block writes. Most database application take advantage of this more efficient 4K blocks sizing which the whole storage industry is evolving to. Oracle already defaults to a multiple of 4K (8K) making it an ideal target application area.
2. Clustered environments – The F5100 currently does not support clustering nor Oracle Rack due to the fact that SATA Flash Modules lack shared storage support. Clustering and Oracle RAC will be support with be available with SAS versions of the internal Flash Modules which are planned for future release.

Product Availability

Key schedule dates for the Sun Storage F5100 Flash Array are:

<i>Event</i>	<i>Date</i>
Sun Product Introduction (Presto)	September 15, 2009
Public Announcement	October 13, 2009
Revenue Release (RR)	August 28, 2009
General Availability (GA)	September 25, 2009

Enabling Technology for Sun™ Storage F5100 Flash Array

Solid-state Flash Technology

With mechanical disk drives no longer able to keep up with the performance demands of today's servers, enterprise-grade solid state flash technology offers a solution to poor storage performance causing latencies, I/O bottlenecks and chocking applications.

Recent advances in the production of flash technology have made solid state drives (SSDs) and flash arrays products much more cost-effective, enabling a new approach to tiered storage. Solid state flash and SSDs fall in a cost and performance sweet spot between mechanical drives and DRAM. They are non-volatile and significantly cheaper than DRAM. They also offer much higher performance and greater power efficiency than hard disk drives.

Reliability characteristics of enterprise-class flash and SSDs have also improved, yielding MTBF ratings that exceed those of hard disk drives. Like hard disk drives (HDDs), enterprise SSDs also support bad block management, wear leveling, and error correction codes (ECC) to foster the highest level of data integrity and reduce service downtime. The solid state nature of flash also allows enterprise SSDs to withstand significantly higher shock and vibrations than HDDs along with requiring less power and cooling.

These devices supports multiple operating systems and are fully integrates with Solaris ZFS Hybrid Storage Pools management for greatest performance and lowest overall storage costs.

Interface Technologies Overview

SAS interface technology and the Sun Storage F5100 Flash Array.

SAS offers customers more choice, with the ability to deploy a diverse range of system capabilities through a common storage connection interface. Due to its full duplexing capabilities, data can be received and transmitted at the same time, providing increased bandwidth. SAS also breaks free of the 15-drive-per-channel barrier that parallel SCSI technology offered. With its high addressability and connectivity, the SAS interface dynamically increases connectivity to attached nodes.

Additional SAS features include:

- Port and bandwidth aggregation will connect storage devices via x4 wide ports
 - 3 Gb/sec wide-ports combine x4 links providing a cumulative bandwidth of 12 Gb/sec using all lanes to ship frames comprising I/O and data
 - When connecting a server to a F5100 array, wide ports offer higher bandwidth than its parallel SCSI predecessor
- Preserves SCSI middleware
 - Middleware, the software that uses interface commands such as applications,

custom utilities, or scripts, has not changed from parallel SCSI with SAS

The Sun Storage F5100 Flash Array uses SAS direct connect to provide a very high performance host connection delivering over 1.6 million IOPS and 12.8GB/s sequential throughput with less than 0.29ms latency. Because of its blazing performance, placing the F5100 on a FC SAN fabric would only slow it down due to the inherent SAN latency which might be hidden by the slower mechanical disk drives, but not with a high-performance solid-state Flash device, such as the F5100 Flash Array. With an ultra-performance device such as the F5100, you want to locate it as close to the host server as possible – direct SAS connect for best performance. The 5100 takes advantage of SAS direct connect and with its 16 x 4-wide SAP ports (64 lanes) deliver maximum performance and throughput with minimum latency.

Sun StorageTek™ Common Array Manager (CAM) software

Sun StorageTek™ Common Array Manager software, an array management tool with powerful functionality, is used as the standard management software for the Sun Storage F5100 Flash Array as well as the Sun J4xxx, ST25xx and ST6xxx arrays. Common Array Manager software has an easy-to-use, Java-based GUI which provides centralized administration for multiple arrays, including the F5100 and any other Sun storage arrays that the customer may deploy. Common Array Manager software enables online administration, a consistent interface across all operating systems, and the ability to monitor and manage one or all arrays from any location on the network.

For more information on Sun StorageTek Common Array Manager software, please visit: http://www.sun.com/storagetek/management_software/resource_management/cam/

Sun Solaris™ ZFS

Sun Solaris™ and OpenSolaris ZFS (Zettabyte File System) software automates common administrative tasks for arrays, protecting data from corruption and providing virtually unlimited scalability. It uses virtual storage pools to make it easy to expand or contract the F5100 arrays simply by adding or deleting Flash Modules (drives). Solaris ZFS will significantly reduce costs by streamlining storage administration and allowing resources to be shared among file systems. The time required to perform some functions will be reduced by orders of magnitude — from hours to just seconds.

Solaris ZFS Hybrid Storage Pools Management.

Getting the best performance from these and other Flash devices is simplified through the use of Sun's ZFS Hybrid Storage Pools feature included in the Solaris(TM) Operating System (OS). The built-in automated tuning and extra resiliency features make it a popular choice for many customers seeking the highest level of performance, reliability, storage architecture optimization and lower TCO.

The Solaris™ Zettabyte File System (ZFS), available to use with the F5100 Flash Array along with Sun Servers running Open Solaris/Solaris 10 at no added cost, is designed from the ground up to offer customers robust data protection and increasing RAS to meet their demanding data requirements. Solaris ZFS contains many powerful features including end to end data check sums and self-data healing that dramatically improves data integrity and reduces downtime.

For more information about Solaris ZFS, please refer to the following URLs:

Solaris ZFS: http://www.sun.com/software/solaris/zfs_learning_center.jsp

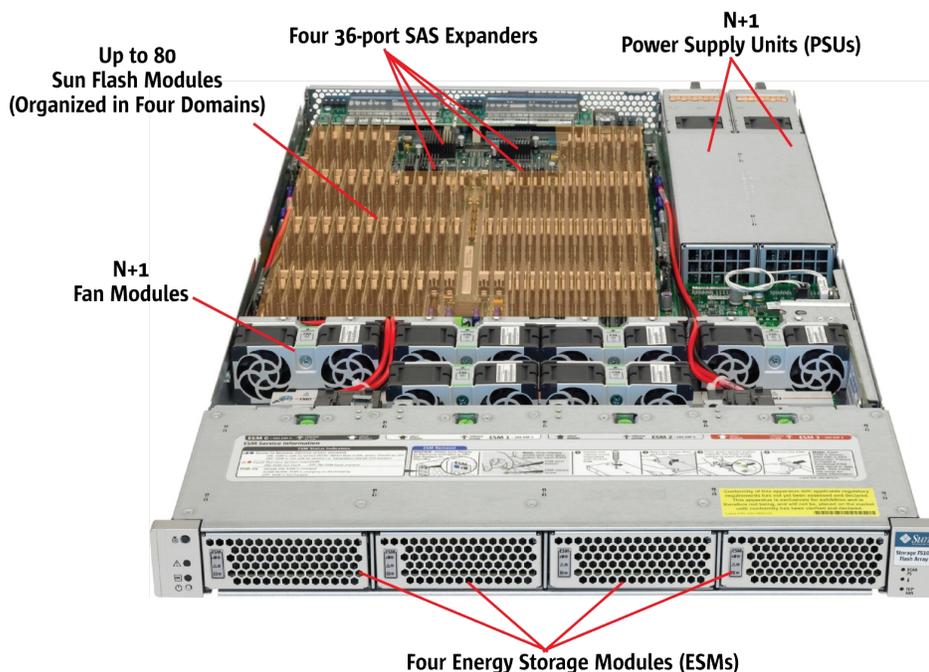
System Architecture: Sun™ Storage F5100 Flash Array

Overview

The Sun Storage F5100 Flash Array is the world's first enterprise flash array and delivers unprecedented performance with well over 1 million IOPS in just 1RU of space with 100X better power and space efficiency than traditional disk. It allows customers to quickly accelerate their databases and I/O intensive applications as well as optimize their overall storage infrastructure for the best performance, eco efficiency, reliability, and cost savings. It offers up to 2TB of solid-state storage with just 300 watts of power and 1RU of space – equaling the I/O performance of over 3,000 enterprise 15K disk drives. It has the best IOPS/\$, IOPS/watt, IOPS/space and IOPS/GB efficiency in the industry.

Architecture

The Sun Storage F5100 Flash Array is a dense, high-capacity, high-performance solidstate storage array that can be attached via a SAS HBA to a server running a variety of operating systems, including OpenSolaris™ 2009.06, the Solaris™ 10 Operating System (Update 8), Microsoft Windows 2003 (SP2), Microsoft Windows 2008 (SP1 or SP2), Red Hat Enterprise Linux 4 (Updates 6 and 7), Red Hat Enterprise Linux 5 (Updates 2 and 3), SuSE Linux 9 (Update 4), and SuSE Linux 10 (SP1 and SP2). The array features a dense, compact design in a 1U chassis. Figure below shows the top cover removed and the locations of several key components.



Sun Storage F5100 Flash Array : Just the Facts
Sun Confidential: Internal and Sun Channel Partners Use Only

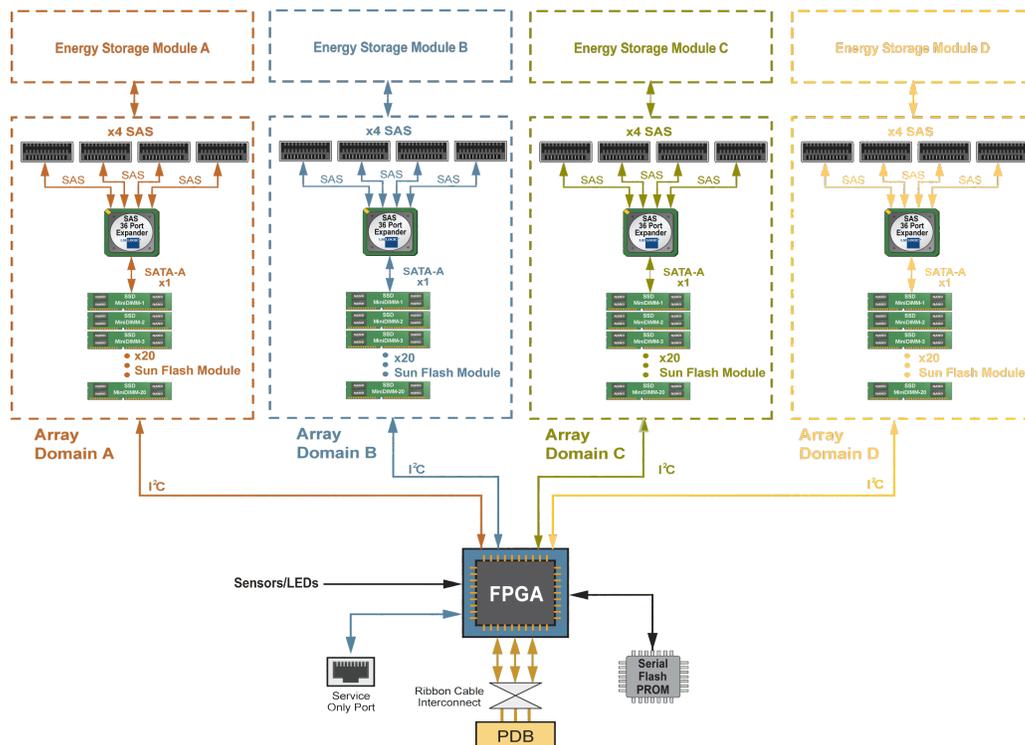
The array is divided into four array domains that function as independent SAS domains for protection, flexibility and optimum server isolation/performance alignment.

Each array domain contains:

- 20 Sun Flash Module slots. When the 20 slots in an array domain are populated, the domain provides up to 480GB of addressable storage capacity.
- Four x4 SAS ports for external host connectivity. Four ports per array domain enables tremendous performance and configuration flexibility (see section on Configurations, Deployments and Best Practices for example configurations).
- An LSI 36-port SAS expander. Each expander interfaces between 20 modules in the array domain and the corresponding four external host SAS connections.
- An energy storage module or ESM (one per array domain). This super capacitor-based unit provides enough energy to flush data from module DRAM to persistent flash storage in each array domain.

Block Diagram

Figure below displays a block diagram for the Sun Storage F5100 Flash Array. Different colors are used to represent components in different array domains, including the ESM associated with each domain.



Sun Storage F5100 Flash Array : Just the Facts
Sun Confidential: Internal and Sun Channel Partners Use Only



An FPGA (Field-Programmable Gate Array) communicates with all four domains via an inter-integrated circuit (I2C) bus and maintains array status information, including the status of Sun Flash Modules and the ESMs, as well as ambient and component temperatures. The FPGA reports status information or device failures to the Sun StorageTek Common Array Management software, which is used for in-band array management.

Sun Flash Modules

Chapter 2 provides specifics on the design of the Sun Flash Modules. When populated, an array domain holds 20 modules. The modules are inserted into motherboard slots that feature color-coded latches corresponding to the array domain.

LSI 36-port expanders

Four 36-port LSI SASx36 expanders reside on the array motherboard, one per array domain. The LSI SASx36 expanders support Serial ATA (SATA) standards and are compliant with ANSI-defined Serial Attached SCSI (SAS) specifications. Each expander provides 20 ports that interface to SATA links from the 20 Sun Flash Modules, along with 16 ports that connect to the four 4-lane SAS external host connectors for the array domain.

Energy storage modules (ESM)

Controlled by the FPGA on the motherboard, the energy storage modules (ESMs) are energy storage devices used to provide backup power to the Sun Flash Modules. On each Sun Flash Module, data is cached on DRAM to enhance performance. In the event that the array loses AC power, the ESMs provide sufficient backup power to flush DRAM contents (both metadata and data) to NAND devices. In this way, the ESMs help to maintain data integrity in the event that a power loss suddenly occurs. Designed by Sun, each ESM is a sealed unit containing 6 super capacitors. The super capacitor-based ESM design provides good reliability, especially compared to other technologies, such as battery-based solutions. Each ESM corresponds to a specific array domain and supplies 5 seconds of backup power to flush module DRAMs when array power is lost.

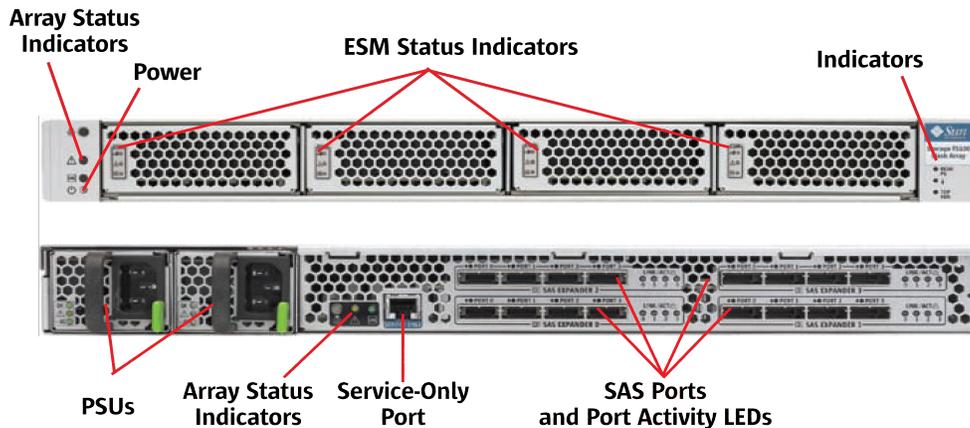
If an ESM experiences a fault condition, it automatically discharges so that the array can be powered down and the unit safely replaced. A Blue “OK to remove” LED on the front of the ESM module indicates that the unit is discharged and safe to remove.

Enclosure features

The Sun Storage F5100 Flash Array enclosure occupies a single rack unit in a standard rack with rail depths of at least 28.125 inches (714 millimeters).

Front and rear perspectives

The following figure shows the front and rear panels of the Sun Storage F5100 Flash Array.



External features and connections

- Front and rear array status indicator lights, reporting “locator” (white), “service required” (amber), and “activity status” (green)
- Front temperature, PSU, and fan status indicator lights
- Front ESM indicators, displaying ESM charging status, an ESM fault condition, and whether the ESM is fully discharged and can be safely removed
- Up to two PSUs (for N+1 redundancy) with integrated fans, each having a single, independent AC plug on the rear panel
- Rear PSU indicator lights, showing the status of each PSU
- Four banks of four x4 SAS ports (SFF 8088), with each bank connecting to a single array domain, along with LEDs representing I/O activity on each SAS connection
- A single service port (password-protected for access only by Sun Service)

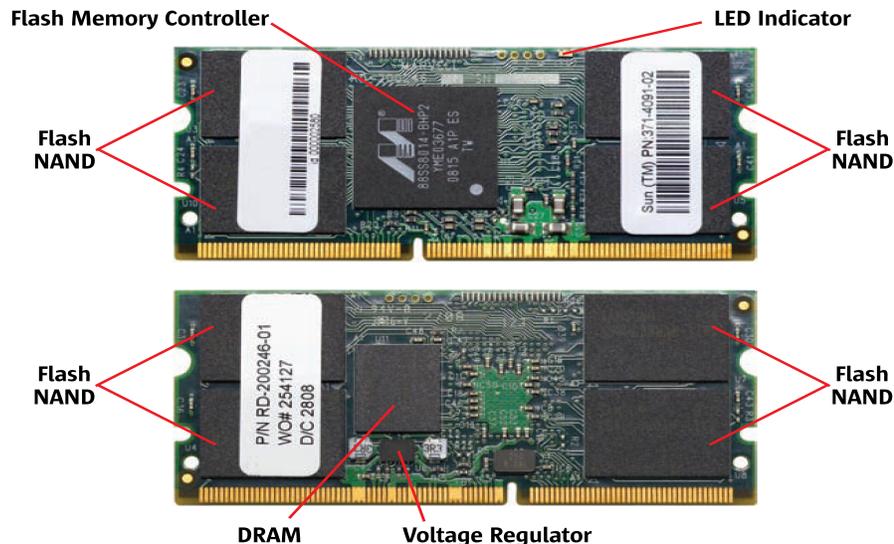
Sun Flash Module

Sun has engineered a novel storage innovation — the Sun Flash Module or “FMod”—that combines SLC NAND flash components and a flash memory controller (FMC) to provide an industry-standard SATA device in a compact, highly efficient form-factor. The Sun Storage F5100 Flash Array takes advantage of Sun Flash Modules to deliver a high capacity and high-performance storage array in an optimal footprint. The array is divided into four separate array domains, each containing 20 Sun Flash Modules for a maximum of 80 modules per array. Peak performance metrics for an individual module are as follows:

- Random write: 15,300 IOPS
- Random read: 20,750 IOPS
- Sequential write: 118 MB/sec
- Sequential read: 265 MB/sec

Flash Module components

Figure below shows key components on the front and rear of each module, which uses a mini-DIMM JEDEC MO-258A form factor.



Flash Module components include:

NAND SLC (Single Level Cell) flash.

Each module contains eight 4GB SLC NAND components (four on the front side and four on the back), for a total of 32GB, of which 24GB is addressable for primary back-end storage. Excess capacity (8GB) is used to optimize performance and longevity. Because of the availability of spare blocks, slower erase cycles can occur independently in the background and faulty blocks can be mapped out so they are not reused.

DRAM.

64 MB of DDR-400 DRAM per each flash module provides a local buffer cache to accelerate flash performance and maintain active data structures. In the event of a loss of power to the array, active data structures located in DRAM are written automatically to the flash devices to help ensure data integrity.

Flash memory controller.

Each Sun Flash Module incorporates a flash memory controller — a SATA-2 controller that

allows each module to communicate using standard SATA protocols. The controller manages the NAND components and buffer cache, and provides a communication interface to systems. To extend the life NAND devices, the controller performs “wear-leveling” and on-the-fly error correction. Wear leveling is a technique that decreases wear by minimizing writes to the same location. The controller is also responsible for tracking and mapping out faulty blocks. Faulty blocks are replaced with spare blocks that are mapped in when needed. In addition, the controller load-balances and interleaves data accesses to back-end NAND devices to accelerate I/O operations.

Voltage regulator

A voltage regulator on each module down-regulates the 3.3V input voltage into 1.2V and 2.5V, which are used by the flash memory controller and DRAM.

Enterprise-quality NAND flash for reliability

To develop enterprise-grade NAND SLC devices, Sun engineers worked closely with NAND manufacturers to make specific reliability enhancements. These design changes allow Sun’s enterprise-quality SLC NAND devices to exhibit greater endurance. In addition, Sun performs extensive quality assurance testing and component screening to optimize NAND device reliability. SLC is usually rated for 100,000 write-erase cycles, whereas MLC Flash is usually rated between 5,000-10,000 write-erase cycles. NAND flash memory on Sun Flash Modules is certified for 2 million hours MTBF.

Low power consumption and fast I/O

The cornerstone of the Sun Storage F5100 Flash Array’s design is the Sun Flash Module, which consumes only 2.1w per module. Because of the module’s low power demands, the array delivers virtually unmatched power and space efficiencies as well as low latency.

F5100 management — Sun StorageTek CAM Software

An internal FPGA maintains array status information, including the status of Sun Flash Modules as well as ambient and component temperatures. The FPGA reports status information or device failures to the Sun StorageTek Common Array Management software, which is used for in-band array management via SCSI Enclosure Services (SES).

When deploying the Sun Storage F5100 Flash Array and other Sun storage products, the StorageTek CAM software provides administrators with a powerful, yet easy-to-use Java™ language-based graphical user interface for management. The CAM software enables online administration, a consistent interface across all operating systems, and the ability to monitor and manage one or all arrays from any location on the network. The software employs a wizard-driven, automated best practices model that takes the complexity out of configurations, saving time and reducing the chance of error. Administrative tasks such as asset discovery, configuration, re-configuration, expansion, and firmware maintenance can all be performed using the StorageTek CAM software.

Starting with version 6.4.1, the StorageTek CAM software supports the concept of zoning. Zoning allows a subset of storage devices to be allocated to specific SAS controllers within a single SAS subsystem. For example, when a single array domain in the Sun Storage F5100 Flash Array is configured using SAS zoning, different zones of flash device-based storage within

a domain can be allocated to a particular HBA. This allows for up to 16 HBAs (4 HBAs per domain) to connect to a single array. The StorageTek CAM software is available at no charge for a wide variety of host platforms, including Linux, Windows, and the Solaris Operating System. It can be downloaded from www.sun.com/storagetek/management_software/resource_management/cam

Performance Sun™ Storage F5100 Flash Array

Performance Data: Four Corners Test

Sun Proprietary/Confidential: Internal Use Only

These results have not been publicly disclosed but may be discussed with customers with a Confidential Disclosure Agreement in place.

Bandwidth & IOPS Four Corners Test

Sun Storage F5100 Flash Array

Achieves 1.6 Million 4K IOPS and 12.8 GB/sec

Significance of Results

The F5100 Flash Array is a high performance high density solid state flash array delivering over 1.6 million IOPS (4K IO) and 12.8GB/sec throughput (1M reads). The Flash Array is designed to accelerate IO-intensive applications, such as databases, at a fraction of the power, space, and cost of traditional hard disk drives. It is based on enterprise-class SLC flash technology, with advanced wear-leveling, integrated backup protection, solid state robustness, over 2M hours MTBF Flash Module reliability.

- The F5100 Array demonstrates breakthrough performance of 1.6 Million 4K random read IOPS
- The F5100 Flash Array can also perform 1.2 Million 4K random write IOPS
- The F5100 Flash Array has unprecedented throughput of 12.8 GB/sec.

Performance Landscape

Results were obtained using four Sun SPARC Enterprise T5240 servers and 16 SAS HBAs

Bandwidth and IOPS Measurements

Test	Flash Modules			
	80	40	20	1
Random 4K Read	1,590,750 IOPS	795,823 IOPS	397,170 IOPS	20,756 IOPS
Maximum Delivered Random 4K Write	1,217,000 IOPS	610,000 IOPS	304,000 IOPS	15,300 IOPS
Maximum Delivered 50%-50% 4K Read/Write	850,000 IOPS	426,000 IOPS	213,000 IOPS	10,600 IOPS
Sequential Read (1M)	12,790 MB/sec	6,395 MB/sec	3,197 MB/sec	265 MB/sec
Maximum Delivered Sequential Write (1M)	9,744 MB/sec	4,833 MB/sec	2,411 MB/sec	118 MB/sec

Sun Storage F5100 Flash Array : Just the Facts
Sun Confidential: Internal and Sun Channel Partners Use Only



Sustained Random 4K Write*	311,000 IOPS	169,000 IOPS	86,000 IOPS	4,400 IOPS
----------------------------	--------------	--------------	-------------	------------

(*) Maximum Delivered values measured over a 1 minute period. Sustained write performance differs from maximum delivered performance. Over time, wear-leveling and erase operations are required that might impact write performance levels.

Latency Measurements

The F5100 Flash Array is tuned for 4 KB or larger IO sizes, the write service for IOs smaller than 4 KB can be 10 times more than shown in the table below. It should also be noted that the service times shown below are both the latency and the time to transfer the data. This becomes the dominant portion the the service time for IOs over 64 KB in size.

Transfer Size	Service Time (ms)	
	Read	Write
4 KB	0.41	0.28
8 KB	0.42	0.35
16 KB	0.45	0.72
32 KB	0.51	0.77
64 KB	0.63	1.52
128 KB	0.87	2.99
256 KB	1.34	6.03
512 KB	2.29	12.14
1024 KB	4.19	23.79

- Latencies are measured application latencies via vdbench tool.
- Please note that the F5100 Flash Array is a 4k sector device. Doing IOs of less than 4k in size, or not aligned on 4k boundaries, can result in a significant performance degradations on write operations.

Results and Configuration Summary

Storage:

- Lightning Flash Array
 - 80 Flash Modules
 - 16 ports
 - 4 domains (20 Flash Modules per)
 - CAM zoning - 5 Flash Modules per port

Servers:

- 4 x Sun SPARC Enterprise T5240
- 4 x 4 HBAs each, firmware version 01.27.03.00-IT

Software:

- OpenSolaris 2009.06 or Solaris 10 10/09 (MPT driver enhancements)
- Vdbench 5.0

Required Flash Array Patches SPARC, ses/sngen patch 138128-01 or later & mpt patch 141736-05

Required Flash Array Patches x86, ses/sngen patch 138129-01 or later & mpt patch 141737-05

Benchmark Description

Sun measured a wide variety of IO performance metrics on the Lightning Flash Array using Vdbench 5.0 measuring 100% Random Read, 100% Random Write, 100% Sequential Read, 100% Sequential Write, and 50-50 read/write. This demonstrates the maximum performance and throughput of the storage system.

Vdbench profile [parmfile.txt here](#)

Vdbench is publicly available for download at: <http://vdbench.org>

Key Points and Best Practices

- Drive each Flash Modules with 32 outstanding IO as shown in the benchmark profile above.
- HBA firmware level needs to be 1.27.3 or newer
- Best to use either single port HBAs or only 1 port per HBA.
- SPARC platforms will align with the 4K boundary size set by the Flash Array. x86/windows platforms don't necessarily have this alignment built in and can show lower performance

Disclosure Statement

Lightning Flash Array delivered 1.6M 4K read IOPS and 12.8 GB/sec sequential read. Vdbench 5.0 (<http://vdbench.org>) was used for the test. Results as of August 31, 2009.

Performance Considerations: 4K Block Alignment

The performance characteristics of the Sun Storage F5100 Flash Array are well-suited to the following I/O access types.

- Access patterns and transfer sizes—Access patterns (LBAs) that are on exact 4 KB boundaries, and read and write transfer sizes that are multiples of 4 KB, can take maximum advantage of the 4,096 byte page size used by Sun Flash Modules.
- Applications—Databases and related applications that perform most I/O in multiples of 4 KB are an ideal fit. Indeed, most databases perform I/O on 8 KB boundaries.
- Block access methods—Most generic device drivers, volume managers, and filesystems assume a 512 byte block size for raw storage. Tools that issue I/O requests to arbitrary 512 byte boundaries or use odd transfer sizes cannot take maximum advantage of the I/O performance advantages of the array. However, a move to larger block sizes is underway. As more systems, interfaces, and best practices optimize for the 4,096 byte block size being implemented in many of today's storage and I/O systems, the greater the performance improvement for applications. In terms of application access patterns and latency sensitivity, the ideal fit and optimal benefit can be found with smaller

datasets, particularly those with very high I/O access densities (greater than 1 IOP/GB), and latency bound applications that do not benefit from caching. Reducing the latencies associated with datasets that fit in the flash array by an order of magnitude or more often can significantly improve system performance. Databases are a prime example. A subset of the database that is 4 KB friendly can achieve a high access density of IOPS/GB. Indeed, use cases and reference architectures have shown double the performance when moving database indexes and hot tablespaces from higher performance traditional RAID systems to the Sun Storage F5100 Flash Array. For more information on these studies, see the technical white papers and resources listed in the reference section at the end of this document.

Aligning Flash Modules for Optimal Performance

Solid State FLASH devices have block alignments typically aligned on 4KByte boundaries, not the 512Byte boundaries of conventional disks. In order to maximize performance, partitions need to be aligned on 4KByte boundaries. In SPARC Solaris systems there are two types of label available, SMI (Sun Microsystems) and EFI (Extensible Firmware Interface). The Format tool selected with defaults ensures that any partition set up using the cylinder values will be correctly aligned.

Solaris Procedures:

The SPARC Solaris format tool using the SMI default, in common with other format and partition tools uses a notation of virtual sectors, tracks and cylinders to describe a disks geometry. This same notation is used for Solid State FLASH devices in common with conventional disk drives.

If EFI labels are used, the partition set up uses sector values. This makes setting and verifying the beginning of a partition easy; the sector size multiplied by the beginning sector needs to be a multiple of 096 bytes. The EFI DEFAULT is NOT 4K aligned.

1. Problem

The storage array uses NAND Flash memory which is aligned on block boundaries that are different from conventional disks that use 512Byte sectors. The F5100, like contemporary SSD's commonly use 4KByte alignment which is the direction that most disk drives are moving toward in the future. Depending on the drive firmware and cache storage on the drive, performance may be adversely affected due to excessive read/modify/write operations when transfers are not aligned.

Partitioning tools in use still use the concept of cylinders, tracks and sectors. This carries over to SSD's as well except that the cylinder, tracks and sectors are now virtual. Some tools maximize the number sectors and tracks, for example in Linux a virtual cylinder is described as 63 Sectors and 255 tracks. A 24GB SSD will be presented as 2987 cylinders, 255 tracks and 63 sectors and a 32GB SSD as 3890 cylinders, 255 tracks and 63 sectors. Setting up a partition based on an arbitrary number of cylinders has a high probability of not being 4K aligned and a subsequent drop in performance.

Bytes per cylinder = 512 bytes/sector * 63 sectors * 255 tracks = 8225280

which is 1962.13740458 4K block. This is clearly NOT 4K aligned.

Unlike Linux and other systems using fdisk for partitioning, Solaris uses Format which assigns different values to sectors per track and tracks per cylinder. Format assigns 128 sectors and 16 heads

Bytes per cylinder = 512 bytes/sector * 128 sectors * 16 tracks = 1048576 which is 256 4K blocks. This means that any partition created on any arbitrary cylinder boundary will be 4K block aligned.

2. Format

The SPARC Solaris tool for creating disk labels and partitioning disk drives is format. If a drive is not labeled then the format command prompts the user with a "Label it now?" message. The format utility includes a verify command which displays the assignment sectors and heads. This allows the bytes per cylinder to be checked and also shows the start in cylinders permitting verification of alignment.

3. SMI vs EFI

The default label created by the format tool is an SMI (Sun Microsystems) label. This uses the a (virtual) cylinder, sector, track notation. An alternative label is the EFI (Extensible Firmware Interface). This label uses simpler sector value to define the beginning of a partition. The value of the sector size is shown by the tool. To set or verify that a partition is on a 4K boundary, the sector size value is multiplied by the start of partition value; the result must be an integer multiple of 4096 to be 4K aligned. Please note there are restrictions on EFI labels, please search the Sun Web for more information on EFI.

There are a number of tools available in Linux to partition drives as well.

Please go to

<http://wikis.sun.com/display/Performance/Aligning+Flash+Modules+for+Optimal+Performance> for more information procedures and examples for different Operating Systems

Performance Considerations: Software

The Sun Storage F5100 Flash Array shares a design similar to four independent SAS JBOD (Just a Bunch of Disks) devices that can each attach to one to four host initiators. As such, the array can operate out of the box with industry-standard SAS storage HBA connections and the latest performance tuned MPT drivers on the Solaris OS. However, there are several key software considerations.

- Because the Sun Storage F5100 Flash Array can deliver extreme performance, it has been qualified with special performance enhancing SAS HBA firmware. This firmware trades off some connectivity to ensure the system can demand high I/O and take advantage of the array without creating bottlenecks.
 - For maximum SAS PCIe HBA performance, a new version of the HBA firmware (firmware level Phase 15 - MPTFW-01.27.03.00-IT) must be downloaded. This version increases maximum performance from ~50K IOPS to over 100K IOPS per HBA. To download this firmware, please go to <http://www.lsi.com/support/sun>.
- Several software tunables can be applied to ensure software aligns I/O requests on 4 KB boundaries and optimize performance. More information on these tunables can be found in the Deployment Considerations section of this JTF.

- Management and monitoring of the Sun Storage F5100 Flash Array is performed inband with the latest version of the Sun StorageTek™ Common Array Management (CAM) software.

Reliability, Availability, and Serviceability (RAS): F5100

RAS features

Corporate data and business information comprise critical business assets. Enterprise computing technologies strive to furnish a high degree of data protection (reliability), to provide virtually continuous access (availability), and to incorporate procedures and components that help to resolve problems with minimal business impact (serviceability). Commonly referred to as RAS, these capabilities are engineered into Sun's mission-critical computing and storage solutions.

The Sun Storage F5100 Flash Array is designed with these RAS features:

- *Backup power to flush DRAM in the event of a power failure.*

Integrated super capacitors in the ESMs provide 5 seconds of backup power, which is sufficient to automatically flush DRAMs on the Sun Flash Modules. This feature helps to protect data integrity — a key factor in maintaining data availability. When ESM failure is detected, FMods associated with that ESM change their behavior. Specifically, the existing contents of each FMod's DRAM buffer, if any, will be written to flash immediately and subsequent writes will not complete until the data has been stored to flash. Depending on the application, this may result in a dramatic reduction in performance. CAM management software monitors the status of the ESMs and notifies administrative and/or service personnel of any failures.

- *Redundant fan modules and power supplies.*

Redundant fan modules and power supply units provide sufficient cooling and power to support continued operation, even if a PSU or fan module fails. When a downtime is scheduled, the faulty unit can be quickly and easily replaced.

- *Accessible components for improved serviceability.*

When a scheduled downtime occurs, ESMs, power supply units, and fans can be replaced without completely removing the array from the rack.

- *Indicator LEDs on the front and back of the chassis enclosure*

Easily visible LEDs allow problems to be identified and isolated easily.

- *Reliability.*

The Sun Flash Modules were designed and built for reliability. They are certified by the manufacturer for 2 million hours MTBF. Reliability benefits of the Sun Flash Module are enhanced by the controller which:

- Uses wear leveling to improve the life expectancy of Sun Flash Modules by minimizing writes to the same location
- Corrects bad data as necessary with ECC
- Takes blocks out of service when their failure rate, detected after a failed write, becomes unacceptable
- Moves data to a known good location (and updates corresponding mapping information)

Solid-state Flash technology

Solid state storage devices based on flash technology do not function like conventional disk drives. Unlike a conventional drive, data is not stored sequentially on a flash-based device. Information that keeps track of the location of the data — the metadata — is also stored inside the flash storage device. The metadata also serves the additional purpose of tracking the number of writes to the individual storage elements. Flash storage also needs to be managed and there are critical operations which can affect access time to the flash module:

- Since an HDD address is by cylinder, track, and sector, the data is laid down sequentially and an LBA (Logical Block Address) is easily translated to disk geometry. A flash memory-based device, on the other hand, can place a block anywhere in the NAND storage element, resulting in an additional level of tracking to manage every block within the flash-based device.
- Wear-leveling requires the movement of data and then updates to the metadata. If a request is made during wear-leveling or other housekeeping operations, the request must be delayed until the operation completes, which can impact latency.
- Maintaining tables of number of writes.
- Managing defragmentation of metadata areas. Along with the data itself, the metadata must be protected to maintain data integrity. If there is insufficient time to write out the metadata to permanent storage, the data becomes corrupted and can not be recovered. The amount of data stored in buffered volatile storage dictates the need for independent energy storage to write out the data in the event of an unexpected power failure. Energy storage is typically implemented using batteries or super capacitors. Both the management tasks and the layout of the storage arrays have impacts on the performance of a solid state device versus a conventional hard drive. In addition, all flash memory has a native block size, and optimum performance is achieved when the size of the read/write data is an integer multiple of the block size and the data transferred is block-aligned. Data transfers that are not block-aligned and do not use sizes that are a multiple of the block size can impact performance, especially for write operations. Sun Flash Modules use a 4 KB block size.

Super Capacitors - ESMS

Flushing the data in the volatile buffer safely to flash storage in the event of sudden power loss requires an energy backup solution, such as batteries or super capacitors. Batteries have a finite and lower functional life than super capacitors. Typical batteries have to be replaced every 2 to 3 years depending on the technology. Batteries also have issues of temperature since both hot and cold affect stored energy. In addition, batteries have higher internal resistance, so if a lot of current is needed for a short duration, batteries cannot provide it without compromising flash storage sizes. With batteries there is also an issue of the availability of an instantaneous charge. Battery chemistry limits immediate availability of energy, whereas a capacitor can instantly supply it. Finally, there is also the problem of detecting when batteries must be replaced. Super capacitors have a much longer life. Although elements of a super capacitor are similar to batteries, they do not suffer from wear out through pure discharge (non rechargeable

batteries) or charge/discharge cycles (rechargeable cells) as severely as batteries. Also super capacitors can provide much higher short-duration current than an equivalent battery and allow for an extended temperature range usage, thus enabling a longer life expectancy. However, super capacitors also have the following limitations:

- Some designs with super capacitors use chemicals with shipping and disposal restrictions. The ESM in the Sun Storage F5100 Flash Array includes a label that describes the proper steps for device disposal.
- All super capacitors have a finite life that is highly temperature-sensitive. The wearout mechanism is a loss of capacitance. If the energy storage falls below what is needed to complete write-back operations, the data becomes at risk for corruption in the event of a power failure. The Sun StorageTek Common Array Manager (CAM) software monitors the status of the ESMs and reports failures as they happen, indicating the need for replacement.

Specifications: Sun™ Storage F5100 Flash Array

Sun Storage F5100 Flash Array Specifications			
	20 Flash Modules	40 Flash Modules	80 Flash Modules
Capacity (#1)			
Capacity	480GB	960GB	1,920GB
Domains	4	4	4
I/O Performance (#2)			
Random 100% Read IOPS (4K)	397K IOPS	795K IOPS	1591K IOPS
Random 100% Write IOPS (4K)	304K IOPS	610K IOPS	1217K IOPS
Random 50% Read-50% Write (4K)	213K IOPS	426K IOPS	850K IOPS
Bandwidth			
Seq. Read (1M transfers) – max delivered	3197 MB/s	6395 MB/s	12790 MB/s
Seq. Write (1M transfers) – max delivered	2411 MB/s	4833 MB/s	9744 MB/s
Latency (#3)			
Read Latency (service time)	0.41 ms	0.41 ms	0.41 ms
Write Latency (service time)	0.28 ms	0.28 ms	0.28 ms
Interfaces			
Ports	16 x 4-wide SAS-1 (64 channels)	16 x 4-wide SAS-1 (64 channels)	16 x 4-wide SAS-1 (64 channels)
SAS Domains	4	4	4
Power			
Input Voltage	100-120 VAC or 200-240 VAC	100-120 VAC or 200-240 VAC	100-120 VAC or 200-240 VAC
Idle Power	127 watts	157 watts	213 watts
Active Power	127 watts	281 watts	386 watts
Peak Turn-on (10 min during ESM charge)	385 watts	445 watts	505 watts
Data Protection			
Cache Power Backup	4 SuperCap Energy Storage Mods	4 SuperCap Energy Storage Mods	4 SuperCap Energy Storage Mods
Mirroring	Host based	Host based	Host based
Reliability			
Endurance - 100% read	Unlimited	Unlimited	Unlimited
Endurance - (50% read/50% write)	> 6 years	> 6 years	> 6 years
Redundant power and fans	Yes	Yes	Yes
Hot swappable Flash Modules	No	No	No
Flash Module MTBF	> 2 million MTBF hours	> 2 million MTBF hours	> 2 million MTBF hours
Environmental			
Temperature	Operating: 5 to 35 deg C; Non-Operating: -40 to 65 deg C		
Relative Humidity	Operating: 1—90%, 27 deg. C max; Non-Operating: 10-93%, 35 deg. C max		
Altitude	Operating: 3,000 meters; Non-Operating: 12,000 meters		
Acoustic	55.4 dB		
Physical Dimensions			
Height	1.746 in (44 mm) – 1RU	1.746 in (44 mm) – 1RU	1.746 in (44 mm) – 1RU
Width	16.750 in (425 mm)	16.750 in (425 mm)	16.750 in (425 mm)
Depth	28.125 in (714 mm)	28.125 in (714 mm)	28.125 in (714 mm)
Weight	35 lb (15.9 kg)	35 lb (15.9 kg)	35 lb (15.9 kg)
Management			
Management Software	StorageTek Common Array Manager		
Host Bus Adapter Firmware	SAS-1 HBA firmware ver. 1.27.03 or greater		

#1 - Not shown is additional 25% of internal capacity used for advanced wear leveling, bad block mapping and performance enhancements

#2 - Performance is system, HBA, application and workload dependent. 4K block optimized. Performance results based on maximum delivered performance tests with 4 Sun SPARC Enterprise T5240 servers and 16 SAS-1 HBAs

#3 - Latencies are measured application latencies via vdbench tool

Sun Storage F5100 Flash Array : Just the Facts
Sun Confidential: Internal and Sun Channel Partners Use Only



System Requirements: Sun™ Storage F5100 Flash Array

The Sun Storage F5100 Product Release notes are posted at the following URL:
<http://dlc.sun.com/pdf/820-6091-10/820-6091-10.pdf>

Supported Host Platforms

The Sun Storage F5100 Flash Array is support the following host platforms:

Current tested and qualified servers,include:#

- Sun Fire X2100
- Sun Fire X2100M2
- Sun Fire X2200
- Sun Fire X2200M2
- Sun Fire X2250
- Sun Fire X4100M2
- Sun Fire X4140
- Sun Fire X4150
- Sun Fire X4170
- Sun Fire X4200M2
- Sun Fire X4240
- Sun Fire X4250
- Sun Fire X4270
- Sun Fire X4275
- Sun Fire X4440
- Sun Fire X4540
- Sun Fire X4600
- Sun Fire X4600M2
- Sun Fire X6220
- Sun Fire X6240
- Sun Fire X6250
- Sun Fire X6440
- Sun Fire X6450
- Sun SPARC Enterprise T5120
- Sun SPARC Enterprise T5140
- Sun SPARC Enterprise T5220

- Sun SPARC Enterprise T5240
- Sun SPARC Enterprise T5440
- Sun SPARC Enterprise T6300
- Sun SPARC Enterprise T6320
- Sun SPARC Enterprise T6340
- Sun SPARC Enterprise M3000
- Sun SPARC Enterprise M4000
- Sun SPARC Enterprise M5000
- Sun SPARC Enterprise M8000
- Sun SPARC Enterprise M9000
- Sun Fire V245
- Sun Fire V445

Supported SAS HBAs

The Sun Storage F5100 Flash Array is support the following Host Bus Adapters (HBAs)

Sun StorageTek external SAS PCIe HBA (*)

Sun PN: SG-XPCIE8SAS-E-Z/SG-PCIE8SAS-E-Z (*)

Sun StorageTek external SAS HBA Express Module for blade servers

Sun PN: SG-XPCI8ESAS-EB-Z/SG-PCIE8SAS-EB-Z

Supported Operating Systems

The Sun Storage F5100 Flash Array has been tested and qualified with the following operating Systems #

- Solaris 10 5/08 SPARC
- Solaris 10 5/08 x86
- Windows 2003 SP2 64-bit
- Windows 2003 R2 SP2 32-bit
- Windows 2008 SP1 64-bit
- Windows 2008 SP2 64-bit
- RHEL4 U5 32/64-bit
- RHEL4 U6 32/64-bit
- RHEL5 64-bit (U-number omitted)
- RHEL5 U2 32/64-bit
- RHEL5 U1 32/64-bit

- SUSE9 SP3 32/64-bit
- SUSE9 SP4 32/64-bit
- SUSE10 SP1 32/64-bit
- SUSE10 SP1 64-bit
- SUSE10 SP2 32/64-bit

Supported Sun Enterprise and Application Software

- Sun StorageTek™ Common Array Manager (CAM 6.5)
- Sun Solaris ZFS

For the latest F5100 Flash Array compatibility information covering HBAs, SAS cables, servers and operating systems, please refer to the following:

Sun Storage Interop Tool posted at:

<https://interopadmin.central.sun.com/interop/interop?cmd=manager>

or

<https://extranet.stortek.com/interop/interop?cmd=report>

Partner site (registration required):

<https://access.sun.net>

Sun Storage F5100 Product Release notes:

<http://dlc.sun.com/pdf/820-6091-10/820-6091-10.pdf>

The data file can also be obtained by internal Sun employees at

<https://onestop.sfbay.sun.com/>

For the latest supported Sun enterprise and application software, and the latest supported third-party software, please refer to the following:

Sun Storage Interop Tool posted at:

<https://extranet.stortek.com/interop/interop?cmd=report>

Management Host System Requirements

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

- Platform: SPARC or x64 server
- Operating Systems: Solaris 8, 9, 10 SPARC, Solaris 10 x64 Operating Environments, Linux RH 4u2, SuSE 9, SuSE 10, Windows XP Professional SP2, Windows 2000 SP4 and Windows 2003 SP1
- Disk Space: 500 MB (full installation), 100 MB (minimum install)
Includes 300 MB in the /opt directory, 200 MB in the /var directory)
- Minimum host memory (2 arrays, 2 users): 512 MB
- Recommended host memory: 1 GB
- Supported Browsers:
 - Netscape Navigator 7.2 or higher
 - Microsoft Internet Explorer 5.5 (7 is not supported)
 - Mozilla / Firefox 1.4 or higher

System Configuration and Deployments: Sun™ Storage F5100 Flash Array

The Sun Storage F5100 Flash Array is designed for maximum performance with minimum space and power. It comes in three simple configurations, consisting of either 20, 40 or 80 Flash Modules. The 20 and 40 Flash Module Configurations can be upgraded up to the full 80 Flash Modules with a purchase of additional Flash Modules, sold in 4-pack of 24GB each.

All configurations come standard with:

- 16 x 4-wide SAS ports (64 channels)
- Redundant, hot swap power and cooling
- 4 Super Capacitor Base Energy Storage Modules (ESM)
- Two 3 meter SAS cables
- 1RU rail kit (snap in type)
- Cable management arm
- CAM software (download)

Flash Module Upgrades and Configurations

Additional Flash Modules can be purchase for upgrading the 20 or 40 Flash Module versions of the F5100 Flash Array up to the full 80 Flash Modules depending on capacity and performance needs. Additional Flash Modules (sold in 4-packs) can be added to any of the 4 internal Domains or across all 4 Domains.

The F5100 Flash Modules can be reconfigured within the systems depending on the need. Flash Modules can be moved from one Domain to another to meet certain configurations, performance and availability needs. A maximum of 20 Flash Modules can be supported on a single SAS-1 HBA. For a fully configured (80 Flash Module) F5100 Flash Array needs a minimum of 4 HBAs with a maximum of 16 HBAs (16 HBAs require Zoning)* for best performance. HBAs and additional SAS cables (up to 3 meter supported lengths) are sold separately.

Min and Maximum Host Connect/HBA s

The F5100 Flash Array uses a direct host connect for maximum performance and lowest latency. It supports up to 4 HBAs without zoning and up to 16 HBAs with Zoning. SAS Zoning is provided via CAM (Common Array Manager) software.

All HBAs can connect to a single server or individual servers. In case a server is limited in the number of available PCIe slots, a Sun IO Expender is recommended.

With CAM zoning, individual servers have access to individual Flash Modules, but the individual

Flash Modules can not be shared between servers (no data sharing) at this time. This is due to SATA Flash Module affiliation restrictions. Data sharing or clustering support is planned for future release with SAS based Flash Modules.

The F5100 Flash Array, with its 4 Domains and 16 SAS ports (4 ports per Domain) can be configured for best performance, best capacity and high availability. Following table gives you an idea of the different configuration options possible with a fully configured (80 Flash Module system) and their tradeoffs. Max performance numbers are just estimates.

	Best Performance	Blend	Best Slot Economy
Number of HBAs*	16	8	4
Modules per HBA/Zone	5	10	20
Zones per Domain	4	2	1
Performance	1.6M IOPS	880K IOPS	440K IOPS

Note:

- All 3 configurations assume a full F5100 of 80 Flash Modules
- For smaller F5100 configurations, the number of HBAs and IOPS will have to be reduced proportionally
- * New high performance HBA firmware (1.27.03) is required and available free at <http://www.lsi.com/support/sun>, This firmware will allow the HBA to more than double its IOPS performance (from some 45K to 105K IOPS), but limit the number devices supported from 128 to some 48.

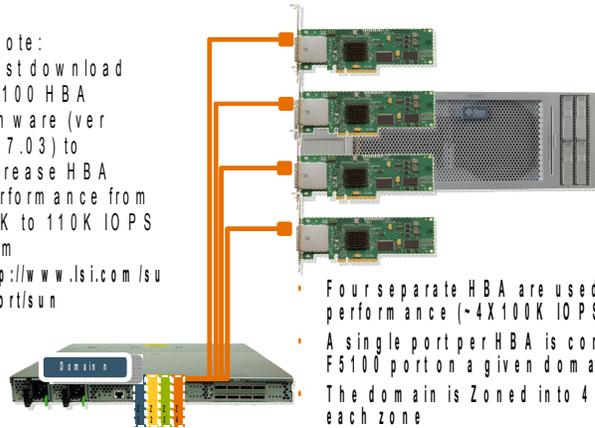
Configuring F5100 for Best Performance

For maximum performance, no more than than 5 Flash Modules should be assigned to a single SAS-1 HBA card. This would imply 4 HBAs per Domain, or 16 per fully configure F5100 with 80 Flash Modules.

Max Performance

- 4 HBA's per domain, 16 per F5100

* Note:
Must download
F5100 HBA
firmware (ver
1.27.03) to
increase HBA
performance from
45K to 110K IOPS
from
<http://www.lsi.com/support/sun>



Sun Storage F5100 Flash Array
Copyright 2009 Sun Microsystems, Inc. All Rights Reserved. Revision A

- Four separate HBA are used for maximum performance (~4X100K IOPS)
- A single port per HBA is connected to a single F5100 port on a given domain
- The domain is Zoned into 4 zones one HBA in each zone
- A single F5100 supports 4X4 zones for a total of 16 HBAs connected for max performance

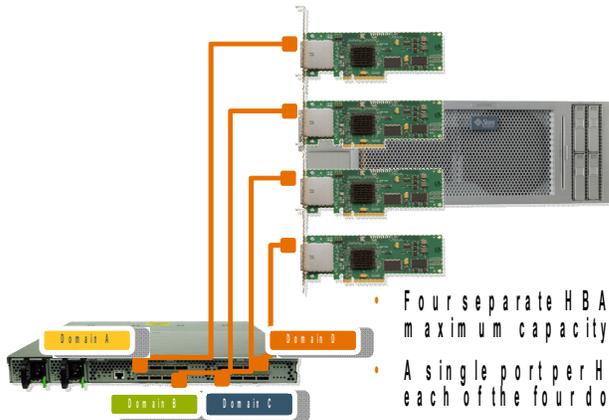


Configuring the F5100 for Highest Capacity

Configuring the F5100 for capacity means reducing the number of HBAs to the minimum, or 1 HBA per Domain (20 Flash Modules) and 4 per fully configure F5100 (80 Flash Modules).

Max Capacity

4 Domains + 4 HBAs
Max Capacity, 1/4 Performance



Sun Storage F5100 Flash Array
Copyright 2009 Sun Microsystems, Inc. All Rights Reserved. Revision A

- Four separate HBA are used for maximum capacity
- A single port per HBA is connected to each of the four domains
- Each HBA controls 20 Fmods

Sun Confidential: Internal Only



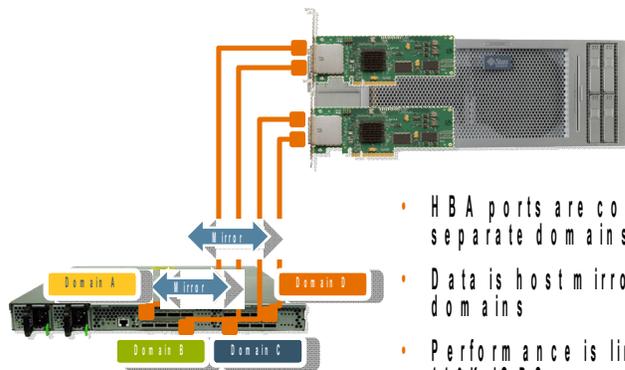
Sun Storage F5100 Flash Array : Just the Facts
Sun Confidential: Internal and Sun Channel Partners Use Only



Configuring the F5100 for High Availability

For enterprise applications, it is highly recommended that the F5100 is mirrored protected. This can be done between individual internal Domains (as shown in the diagram below), or separate F5100 systems. For highest level of protection and redundancy, it's best to use two separate mirrored F5100 systems systems. Host based mirroring is available with the OS such as VSM, Oracle such as ASM, or 3rd party such as Veritas VxVM. The F5100 currently does not support multipathing.

Dual HBAs Driving Four Domains HA is achieved via Host Mirroring



- HBA ports are connected to four separate domains
- Data is host mirrored between domains
- Performance is limited by HBA: 2 X 110K IOPS

Sun Storage F5100 Flash Array
Copyright 2009 Sun Microsystems, Inc. All Rights Reserved. Revision A

Sun Confidential: Internal Only



The F5100 Rack Cabinet Guidelines:

The Sun Storage F5100 is supported in the older Sun racks as well as the new Redwood racks. The rackmount/rail kit offered with the Sun Storage F5100 Flash Array is a snap-in type (similar to that of Sun SPARC Enterprise T5120 servers) and intended for use in a rack with rail to rail depth of at least 28.125 inches (714 mm). The cabinets have been designed to fit within a 1 RU. While the rackmount kit and server have been tested within a limited number of EIA standard racks/cabinets, variances in individual racks may influence the positioning and number of Flash arrays and server chassis that may be installed within any particular rack.

The factors influencing the number of F5100 Flash Arrays and servers that may be installed within a particular rack include available power and cooling, clearances within the rack/cabinet for doors/panels, power distribution and cable management, adjacent equipment and method of mounting, vertical tolerances and positioning of equipment, access requirements, etc.

Sun Storage F5100 Flash Array Physical Dimensions:

Sun Storage F5100 Flash Array : Just the Facts
Sun Confidential: Internal and Sun Channel Partners Use Only



Depth: 28.125 inches (714 mm).
Height: 1RU height or 1.746 inches (44 mm)
Width: 16.75 inches (425 mm)

Note: A 27 inch deep rack can not be used with the F5100 since it cannot accommodate the 28.125 inch long (“Jake”) chassis. A 36 inch standard deep rack will accommodate the Sun Storage F5100 Flash Array.

Note: The Sun Storage F5100 Flash Array ESM (Energy Storage Modules) have a very small amounts of flammable liquid in their Capacitors which makes them Class 9 dangerous good materials that require special air ship considerations. Class 9 is the minimum level of restrictive classification and only involves air ship transports either as a ESM modules alone or if shipped within the F5100 system. Class 9 restrictions do no apply to ground shipments or have any implications to products normal use or configurations.

Cables

SAS interconnect cable lengths which are supported on the F5100 Flash Arrays are .5m, 1m, 2m, and 3m. Six meter cables are not supported.

The F5100 Flash arrays come packaged with two 3 meter SAS cables along with a tool-less snap-in rail kit and cable management arm. For racks that do not support snap-in rail kits, screw in type rails are available as an add-on option.

RAID

RAID protection is provided by host software. There is no hardware RAID support at this time.

Software RAID includes host based mirroring (such as VSM, ASM and VxVM) as well as Solaris ZFS.

Selection of the correct RAID level is equally critical in order for each application to meet the required data availability, performance, and capacity requirements. Mirroring is generally considered to be the best balance for performance and availability on the F5100 (See the Best Practices section for a table that outlines the trade-offs of the various supported RAID configurations).

ZFS software-based RAID functionality

- Sun Solaris 10 ZFS with RAID-Z with support for RAID levels 0, 1, 5, and 6
- Sun OpenSolaris ZFS with RAID-Z with support for RAID levels 0, 1, 5, and 6

When used with the F5100 arrays, ZFS provides data integrity and security, scalability, and ease of management. ZFS spreads writes across all available Flash Modules (additional Flash Modules are added to the ZFS stripe width automatically). ZFS delivers intelligent prefetch: it predicts the next data and pre-caches it. And ZFS offers RAID-Z with support for RAID levels 0, 1, 5, and 6.

Role of CAM in relation to RAID Manager and ZFS software

CAM version 6.5 should be used to perform all the F5100 family component and environmental reporting, service advisor functions, and Auto Service Requests (ASR). CAM v6.5 will also provide the ability to upgrade Flash Module firmware for the F5100 Flash Arrays.

Sun StorageTek CAM version 6.5 and any upgrades can be downloaded from the CAM site.

Multipathing

Multipathing is not supported on the F5100 Flash Array. Redundant and mirrored F5100 Flash Arrays on separate HBAs are recommended for HA configurations

Heterogeneous Hosts

With Domains and heterogeneous hosts features, the F5100 can tailor its behavior to the needs of the host operating system. This provides each individual host with the view of the storage system that it would experience if it had exclusive access to the array. The host types can be completely different operating systems, such as the Solaris OS, Linux, and Windows,

Multi-host support will require SAS zoning capabilities provided through CAM

Deployment and Best Practices

Deployment Considerations

A single Sun Storage F5100 Flash Array is designed to operate similar to four simple JBOD devices, and as such it offers tremendous configuration flexibility. This section includes information about the configuration rules for the Sun Storage F5100 Flash Array and its host bus adapters, as well as several sample configurations that highlight the balancing of functional requirements for availability and performance in database acceleration solutions.

Database acceleration deployment considerations

The Sun Storage F5100 Flash Array brings very low latency random reads to database acceleration environments, making it ideal for index and hot table placement. Storage arrays with NVRAM are still needed to handle logging and data tables. Using these considerations as a checklist can help organizations assess the ability to successfully deploy Sun Storage F5100 Flash Arrays as database accelerators.

Performance improvement assessment

Determining if a database deployment has an I/O bottleneck—and over how much storage—is key to assessing whether the inclusion of flash arrays can help improve system performance.

- Using I/O monitoring utilities that are part of the operating system or environment can help determine if an I/O bottleneck exists, and whether using a flash array can help the situation. In environments running the Solaris OS, the `iostat (1M)` command can help. Sun also provides downloadable utilities that can assist with this analysis, such as the Sun Flash Analyzer available at sun.com/flash. By identifying database index LUNs and storage devices, it is possible to determine if service or wait times are 10 milliseconds or higher. If so, a flash array might help system performance. If these total service times are short, approximately 1 millisecond or less, then the indexes are being cached in the storage subsystem and are likely already optimized.
- Databases often have internal reporting tools that can help identify system bottlenecks. For example, administrators can use the Oracle STATSPACK or Oracle Automated Workload Repository (AWR) reporting tools in existing Oracle Standard Edition (SE) and Oracle Enterprise Edition (EE) database deployments. Go to “Top 5 Wait Events” and check to see if “db file sequential read” is one of the important wait events. The Sun Storage F5100 Flash Array reduces “db file sequential read” events to 1 to 3 ms for affected datasets, which can help dramatically accelerate database performance. If “db file sequential read” is not a widely occurring event, the Sun Storage F5100 Flash Array is not likely to help improve the performance of the deployment.
- In databases that are bottlenecked by storage I/O, a subset of the database typically exhibits the I/O bottleneck, such as deep trees in the indexes. A good rule of thumb is to place 25% of the database (indexes) on the Sun Storage F5100 Flash Array. Each array can support up to 960 GB of protected storage (1.92 TB raw). Use this ratio to determine the number of Sun Flash Modules per array, or number of arrays, needed.

RAS requirements assessment

Sun Storage F5100 Flash Array database acceleration best practices use host-based mirroring

for data protection. Deployments can be designed to avoid single points for failure, except for the compute node itself (single instance database.) Sun Storage F5100 Flash Array database accelerator best practices require some data, such as logs, to be placed on other NVRAM-based storage arrays, making the Sun Storage F5100 Flash Array an ideal add-on component for database acceleration deployments.

Potential RAS requirements include:

- If storage needs to be shared, look for alternate solutions. The SATA Sun Flash Modules cannot be shared by multiple initiators. However, a single flash array can be partitioned using SAS zoning for up to 16 different hosts. • Some deployments use multiple paths to separate, redundant storage devices that are host mirrored. For those that require multipath I/O to database objects, look for alternate solutions. The interface on the SATA Sun Flash Modules do not support multipathing as of this writing.
- Sun Storage F5100 Flash Arrays use Solid State Sun Flash Modules that are inherently more reliable than mechanical rotating hard disk drives. Fans and power supplies are hot swappable. However, all other FRUs require a cold swap, including the Sun Flash Modules. If zero or near zero downtime for servicing is required, then a minimum of two Sun Storage F5100 Flash Arrays are needed, along with Sun Flash Module mirroring across separate Sun Storage F5100 Flash Arrays.
- For Oracle database deployments, use the Oracle Automatic Storage Management (ASM) software to protect database objects, such as indexes, on Sun Storage F5100 Flash Arrays. This is the preferred data protection strategy for Sun Storage F5100 Flash Arrays.
- Other host-based software mirroring, such as that provided by Solaris Volume Manager or Veritas Volume Manager (VxVM) software, can be used to mirror storage—with more complexity and risk, less built-in integration, and potentially higher costs. Documented practices are not provided for the combination of these tools and the Sun Storage F5100 Flash Array.
- Use other storage solutions if hardware data protection, such as RAID and NVRAM, is needed for all database objects.
- If a storage management simplification strategy is in use that dictates a single storage system for holding all parts of the database, deploy an alternate solution.

System deployment assessment

- In addition to performance and RAS considerations, check that the new or existing system deployment supports the addition of flash arrays for use as a database accelerator.
- At least two dual connector external eight lane PCIe SAS HBAs are needed to connect the Sun Storage F5100 Flash Array to servers in deployments. Two HBAs are qualified as of the writing of this document:
 - Sun StorageTek external SAS PCIe HBA (Sun part number SG-XPCIE8SAS-E-Z or SGPCIE8SAS-E-Z)
 - Sun StorageTek external SAS HBA Express Module for blade servers (Sun part number SG-XPCIESAS-EB-Z or SG-PCIE8SAS-EB-Z)
- The SAS HBAs must have firmware level Phase 15 - MPTFW-01.27.03.00-IT. This limits connectively to 20 Sun Flash Modules per HBA, but provides optimum performance results. For example, 40 or fewer Sun Flash Modules require at least two SAS HBAs.

Similarly, 40 to 80 Sun Flash Modules housed in one or two Sun Storage F5100 Flash Arrays require at least four SAS HBAs, and 80 to 160 Sun Flash Modules in two Sun Storage F5100 Flash Arrays require at least eight SAS HBAs. To create a database accelerator solution with a partially populated (40 Sun Flash Modules) array and no single points of failure, four slots are required. Two slots are needed for Fibre Channel HBAs, and two slots are needed for SAS HBAs. If four slots are not available, an alternative solution is required.

- To create a fully mirrored solution using two Sun Storage F5100 Flash Arrays, ten 8-lane PCIe slots are required for the storage alone. For most servers, an I/O expander chassis is likely required. If sufficient PCIe slots are not available on the database server, or if the HBAs are not qualified for the planned database server, look at alternate solutions.
- Use the Solaris 10 OS Update 8 to take advantage of a number of device driver enhancements. Patches are available for earlier Solaris 10 OS updates.
- Each Sun Storage F5100 Flash Array requires only 1 RU of rack space for each Flash array in the deployment. Each array has a nameplate power of 720 watts per power supply to establish its power distribution design. Use the power calculators at sun.com/powercalculators to estimate actual power and cooling consumption for a given deployment. Power consumption for the first five to 10 minutes after power up includes the energy needed to recharge the energy storage modules over and above the steady state chassis power consumption.
- The SAS interconnect cables are limited to three meters. Verify that the rack, rack wiring, and inter-rack wiring if needed, can support interconnects between the server with SAS HBAs and the flash array(s.)

Configuration examples

This section presents three database acceleration configuration examples for the Sun Storage F5100 Flash Array, and discusses how each configuration meets different implementation goals for performance, capacity, and availability.

• Availability considerations.

To meet data availability targets defined in Service Level Agreements (SLAs), array data must be mirrored through the connected host — there is no mirroring in the array itself.

• Performance goals.

Array performance can be influenced by a number of factors beyond the configuration itself, including application-related issues as well as the application-generated I/O mix, and the processing power of the host connected to the array. Configuration decisions that impact performance include the number and type of hosts and the number of independent HBA channels connected to the array.

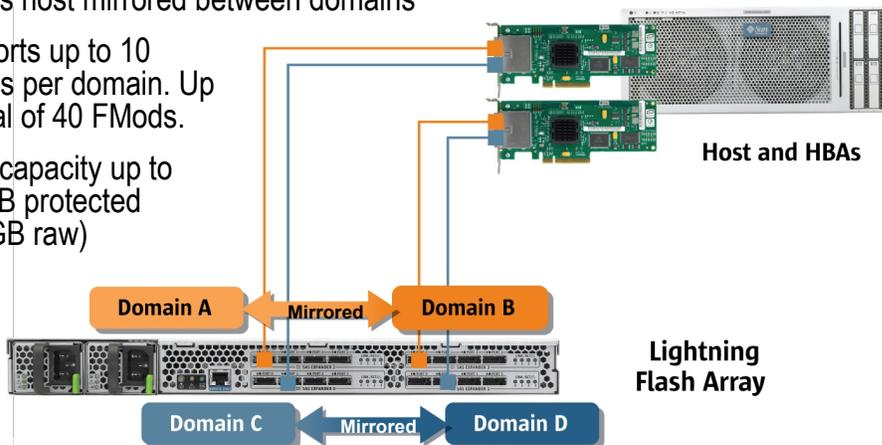
Key factors in maximizing throughput are whether applications produce threaded I/O requests, host compute capabilities, and the number of host HBA cards connected to the array.

The examples below discuss how each given configuration meets design goals for capacity, availability, and performance. Simple availability configurations Figure 6 shows the simplest recommended configuration with two HBA cards installed in the database server. Since each HBA features two SAS channels, this configuration supports connections to all four array domains, and up to 960 GB of raw storage (40 Sun Flash Modules). Host mirroring across HBAs and domains allows data in two array domains to be mirrored to two other domains, reducing the

usable storage capacity to 480 GB. Figure below shows a single host with two HBAs connected to four array domains (two domains per HBA)

Example: 480 GB index for 2TB DB

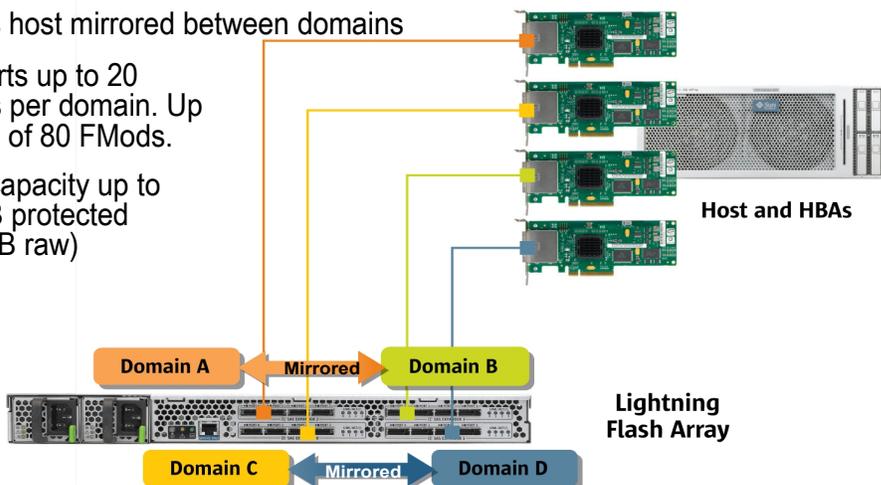
- Requires two 8-lane PCIe slots for SAS HBAs connected via four ports to four separate F5100 domains
- Data is host mirrored between domains
- Supports up to 10 Fmods per domain. Up to total of 40 FMods.
- Total capacity up to 480GB protected (960GB raw)



For maximum capacity with a single Sun Storage F5100 Flash Array, four HBA cards must be used in the database server, as shown in Figure 7, to connect to all four array domains and to access 1.92 TB of storage. As before, host mirroring across SAS domains can be used to increase data availability, which reduces usable capacity to 960 GB. Figure below shows. Single host with four HBAs connected to four array domains (one domain per HBA).

Example: 960 GB index for 4TB DB

- Requires four 8-lane PCIe slots for SAS HBAs connected via four ports to four separate F5100 domains
- Data is host mirrored between domains
- Supports up to 20 Fmods per domain. Up to total of 80 FMods.
- Total capacity up to 960GB protected (1.92TB raw)



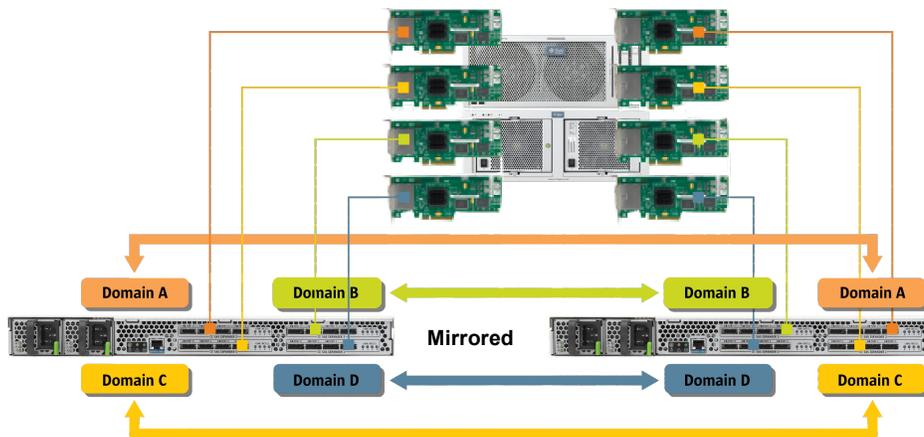
Sun Storage F5100 Flash Array : Just the Facts
Sun Confidential: Internal and Sun Channel Partners Use Only



Figure below depicts a high-performance, high-capacity configuration that reduces downtime when replacing FRUs in one of the mirrored chassis. In this configuration, eight SAS HBA cards and their associated 8 lane PCIe server slots are needed to connect to two fully populated Sun Storage F5100 Flash Arrays and provide 3.84 TB of raw storage. Each domain is mirrored to its counterpart domain in the mirror chassis, reducing usable capacity to 1.92 TB.

Example: 1.92 TB index for 8TB DB

- Requires eight PCIe 8 lane slots for SAS HBAs connected via eight ports to eight separate F5100 domains
- Data is host mirrored between domains on each F5100 chassis. This reduces service downtime for FMod swap, ESM swap, etc



Performance tuning

When first brought online Sun Storage F5100 Flash Arrays are seen as collections of disk drives. Each Sun Flash Module appears as a SAS attached SATA-based solid state drive. Several tuning best practices can help ensure the flash array speeds indexes and delivers maximum performance.

- Be sure the SAS HBA driver is up to date before bringing the storage online.
- Ensure that the firmware in each HBA is qualified for the Sun Storage F5100 Flash Array. The latest firmware can be found at <http://www.lsi.com/support/sun>.
- Adjust block size and alignment parameters that can impact performance. Since Sun Flash Modules use a 4 KB block size, optimal performance occurs using a 4 KB (or a multiple of 4 KB) block size. All partitions on Sun Flash Modules are aligned to start on 4 KB aligned boundaries. This is achieved by using the Solaris OS `format(1M)` command on SPARC® platform, or the `fdisk(1M)` command on x64 platforms, as described in the “Sun Storage F5100 Flash Array Product Notes”. No tuning is required for SPARC

platforms. Simply ensure that standard slices are used to format volumes. On x64 systems running the Solaris OS, block 0 contains the partition table. As a result, formatting starts at block 1, which is not 4 KB aligned.

- Sun Services can help configure and design effective storage solutions using the Sun Storage F5100 Flash Array. Experienced Sun consultants can help to analyze and optimize performance and fine-tune configurations for demanding applications.
- This document recommends using software mirroring to protect raw flash array drives. For Oracle databases, Oracle Automatic Storage Management (ASM) can be used to mirror the systems, yet let the Oracle database manage the raw storage. It is recommended that the Solaris ZFS file system be used to protect key availability portions of the database, such as the archive log and flash recovery areas. It is not recommended to place indexes and hot tablespaces in Solaris ZFS file systems.

Index mirroring, migration, operation, and validation

While a full discussion is beyond the scope of this paper, it is important to describe the final high-level steps for implementing the database acceleration practices defined in this document. Exact determinations and recommendations usually depend more on business needs than the technology. Typically, these steps are handled by the database manager. However, system administrators or storage administrators can be responsible for some of these steps, particularly setting up mirroring and data protection schemes for the array volumes.

- Establish data protection practices—The use of software mirroring is recommended. Third-party volume management tools, such as Symantec's Volume Manager (VxVM), or operating system utilities such as the Solaris Volume Manager (SV), can be used. For Sun Storage F5100 Flash Arrays, it is recommended that the database provide any mirroring needed. For example, using Oracle ASM with Oracle databases simplifies management and delivers optimal performance due to the level of integration inherent with database-managed storage. See *Accelerating Databases with the Sun Storage F5100 Flash Array* for practical examples of balancing the performance and capacity of cost-constrained OLTP and VLDB applications.
- Migrate indexes to the flash array—There are many ways to establish the indexes on the flash array. Working with the various administrative groups that support the database deployment can help determine the best course of action.
 - Add mirrors to existing indexes and resilver to the flash array. Break the mirrors to the old index LUNs. In some environments, these tasks can be completed while the database is online.
 - Perform tablespace migration. Some databases support online tablespace migration (including indexes).
 - Rebuild indexes. Some databases support online index rebuilding to new locations, such as mirrored flash arrays.
 - Use offline techniques if business needs dictate. For example, take the database offline, and backup and restore the indexes to the mirrored flash arrays.
- Operation and validation—Once the database is online, it is important to perform an initial assessment of the service or wait times on the indexes. Service times are expected in the range of 1 to 3 milliseconds, down from the typical 10 to 15 milliseconds. If the index service times remain high, review the installation for discrepancies. If index latencies are reduced, a major database I/O bottleneck has been removed. Large improvements in

total transaction throughput (~2x) and reduced total transaction times are anticipated if the system was principally experiencing I/O read latency bottlenecks. Some systems have other bottlenecks, and easing the I/O bottleneck can help make these issues easier to see, identify, and resolve.

Summary

Businesses face a substantial number of challenges in responding to storage performance requirements for data-hungry applications. Although some high performance storage solutions exist today, economic realities make most of these solutions unattainable. Since many high-performance storage solutions incorporate proprietary controllers and devices, they usually impose high operational costs, burdening the datacenter with high demands for power and cooling, footprint, and administrative complexity. Sun is well-known for its innovative engineering and its ability to deliver value in storage and compute solutions. With the Sun Storage F5100 Flash Array, Sun delivers breakthrough database storage performance at new economic levels. Based on enterprise-grade flash modules designed by Sun, the array raises the bar for I/O performance, as well as performance per watt and performance versus cost. Designed to behave similarly to a JBOD device, the array is a simple building block that integrates easily into existing architectures when higher performance and database acceleration is needed, without the need to re-architect the existing storage infrastructure.

For more information

Following resources are available on sun.com.

Description	URL
Sun Storage F5100 Flash Array	sun.com/storage
Sun Storage F5100 Flash Array Performance Information	blogs.sun.com/BestPerf/entry/1_4_million_iops_in
Flash Performance Tuning	wikis.sun.com/display/Performance/Home#Home-Flash
Sun Power Calculators	sun.com/powercalculators
Sun StorageTek Common Array Management Software	sun.com/storagetek/management_software/resource_management/cam
Solaris ZFS	sun.com/solaris/zfs.jsp
Sun Services	sun.com/services

Other references include:

- Leventhal, Adam. "Flash Storage Memory", Communications of the ACM, July 2008.
- Wright, Jeffrey. *Accelerating Databases with the Sun Storage F5100 Flash Array*

Ordering Information: Sun™ Storage F5100 Flash Array

Ordering information and part numbers for the Sun™ Storage F5100 Flash Array are provided in this section.

Ordering details and marketing part numbers

For technical validation and quality assurance, it is recommended that orders for the Sun Storage F5100 Flash Array be configured through the WebDesk Configuration tool. (Sun UAP for storage is not available through Partner WebDesk.)

There are three URLs for WebDesk, depending on the geography:

- AMER = <http://webdesk.central>
- APAC = <http://webdesk.singapore>
- EMEA = <http://webdesk.holland>

WebDesk Configurator

When the Sun or Sun Partner technical sales representative is configuring the F5100 the Sun Microsystems configuration tool, WebDesk/Partner WebDesk Configurator, will be used as needed. Please note that the Services pricing generated will be a monthly price.

While the three PTO F5100 configurations are very simple, close attention should be paid to assure the right number of HBAs required for optimizing certain configurations .

Configuration Options

The Sun Storage F5100 Flash Array is an external rack-mountable disk array in a 1RU chassis. It is offered in small, medium and a large PTO configurations consisting of either 20/40/80 Flash Modules, 24GB each.

Small	- 0.48TB capacity (20 Flash Modules)
Medium	- 0.96TB capacity (40 Flash modules)
Large	- 1.92TB capacity (80 Flash modules)

The small and medium configurations can be upgraded incrementally with a purchase of additional Flash Module (4 packs) and user installable either in any of the 4 or across all 4 internal domains.

All configurations include the following standard features and components:

- 1 x Lightning Flash Array in 1RU (17x29 in.) consisting of either 20, 40 or 80 internal enterprise-grade 24GB SLC Flash Modules
- 16 x 3Gb x4 lane SAS host connects (4 Domains of 4X4-wide SAS)
- Dual redundant load sharing hot plug power supplies
- Redundant hot plug fans
- 4 ESM (Energy Storage Modules) for power backup – Super Caps based.
- CAM management software
- SAS Zoning
- Rail kit and cable management arm

- 2 x 3m SAS cables

Supported HBAs

Currently, F5100 Flash Array supports the following Sun SAS HBA cards. In the future, other Sun HBA's will be qualified, including SAS-2.

Please note, for maximum SAS PCIe HBA performance a new version of the HBA firmware (rev 1.27.03) must be downloaded which will increase the maximum HBA performance from 45K IOPS to over 105K IOPS. To download this free firmware, please go to <http://www.lsi.com/support/sun>

Marketing Part Number	Description
SG-XPCIE8SAS-E-Z (Pandora)	Sun StorageTek PCIe SAS Host Bus Adapter. 3Gb/s per port. 2X4 external lanes. (Xoption version) http://www.sun.com/storagetek/storage_networking/hba/sas/PCle.pdf
SG-PCIE8SAS-E-Z (Pandora)	Sun StorageTek PCIe SAS Host Bus Adapter. 3Gb/s per port. 2X4 external lanes. (ATO version, for factory configurations) http://www.sun.com/storagetek/storage_networking/hba/sas/PCle.pdf
SG-XPCIESAS-EB-Z (Mimus)	SunStorageTek SAS ExpressModule HBA for Sun Blade Servers. Provides 8-lanes of PCIe 3Gb/s per ports. (Xoption version) http://www.sun.com/storagetek/storage_networking/hba/sas/expressmodule.pdf
SG-PCIE8SAS-EB-Z (Mimus)	SunStorageTek SAS ExpressModule HBA for Sun Blade Servers. Provides 8-lanes of PCIe 3Gb/s per ports. (Factory Configuration version). http://www.sun.com/storagetek/storage_networking/hba/sas/expressmodule.pdf

SAS Cables

The F5100 Flash Array supports standard Sun SAS cables, but only 3 meters in length or less (due to potential performance considerations). Two 3 meter cables are included with each F5100 unit. Additional SAS cables can be purchase as needed.

Marketing Part Number	Description
XTA-0.5M-SAS	Sun™ Storage 0.5m, mini, shielded, SAS cable; For connection

Marketing Part Number	Description
	between array and host; RoHS-6
XTA-1.0M-SAS	Sun™ Storage 1.0m, mini, shielded, SAS cable; For connection between array and host; RoHS-6
XTA-2.0M-SAS	Sun™ Storage 2.0m, mini, shielded, SAS cable; For connection between array and host; RoHS-6
XTA-3.0M-SAS	Sun™ Storage 3.0m, mini, shielded, SAS cable; For connection between array and host; RoHS-6
XTA-6.0M-SAS	Sun™ Storage 6.0m, mini, shielded, SAS cable; For connection between array and host; RoHS-6

IBB Upgrade Advantage Trade-in Program:

The Sun Storage F5100 Flash Array participates in the Sun IBB Upgrade Advantage Trade-in Program. Qualified Sun and non-Sun systems are eligible for trade-in credit when purchasing a new F5100 Flash Array. Trade-in allowances range from 5-15%, depending on the return and the following part numbers. Please refer to the Sun Upgrade Advantage Program IBB website <http://ibb.eng.sun.com/index.html> for more information and trade-in credits for Sun and non-Sun equipment.

IBB Options

Allowance #	Trade-in Allowances
ALW-05-DO-TA5100	UAP 5% allowance for qualified trade-ins of Sun™ systems.
ALW-05-DZ-TA5100	UAP 5% allowance for qualified trade-ins of non-Sun™ systems.
ALW-10-DO-TA5100	UAP 10% allowance for qualified trade-ins of Sun™ systems.
ALW-10-DZ-TA5100	UAP 10% allowance for qualified trade-ins of non-Sun™ systems.
ALW-15-DO-TA5100	UAP 15% allowance for qualified trade-ins of Sun™ systems.
ALW-15-DZ-TA5100	UAP 15% allowance for qualified trade-ins of non-Sun™ systems.

Service and Support: Sun™ Storage F5100 Flash Array

Warranty

Warranty Term : 1 year, Next business day response.

Phone Coverage/Call-Back Response: Monday-Friday/8-5 local business hours.

Customer Installable: Yes

Support

Partners can offer support to their customer by selling a Sun StorageTek Service Plans. The Sun StorageTek Service Plans contain features such as on demand training, online knowledgebase, remote services, and enhanced interoperability support. The SunSpectrum Service Plans provide a simple pricing structure in which a single fee covers support for an entire system, including the Solaris(TM) Operating Environment software.

For more information on the Sun StorageTek Service Plans refer to the following URL:
<http://http://www.sun.com/service/serviceplans/storage.jsp>

- The F5100 support features include Auto Service Request(ASR) via CAM (6.4.1) with "Default-on" enabled.
- Product is customer installable
- Enterprise installation services, Application Readiness Services (implementation) including modules for Flash Storage Optimization for Database Applications and Solaris zfs, iRunbook, parts installation and Data migration services are available for all partners to resell.
- Qualified partners can deliver Sun EIS and DRS services.

Cross-Sell and Up-Sell Opportunities:

Services

Sun Enterprise Installation Services - ensures correct setup and reduces time to deploy. Sun Application Readiness Services - helps achieve optimization of the architecture and take advantage of the advanced availability features of Sun servers and the Solaris Operating System or other applications.

SunSpectrum Support- helps achieve high availability within customer's IT infrastructure, saving time and money.

The recommended service level for this product is 3 years Gold support.

iRunbook - customizable, web-based central repository of streamlined technical "how to" information.

Services targeted to business need

- ARS (Implementation) Flash Storage Optimization for Database Applications - Sun

Sun Storage F5100 Flash Array : Just the Facts
Sun Confidential: Internal and Sun Channel Partners Use Only

Services has extensive expertise with Oracle databases and zfs and can provide implementation customized to customers business needs taking full advantage of the capabilities of zfs.

- Sun host Based Data Migration Services
- Sun Multivendor Support (MVS) or Single Point of Contact Service
- Sun Data Erasure Service
- Sun Managed Services
- Sun Storage Assessment Services
- Sun Parts Installation Service

SunSpectrumSM Service Plans

The SunSpectrum Service Plans are innovative and flexible service offerings that allow 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 Service Plans provide a simple pricing structure in which a single fee covers support for an entire system, including related hardware and peripherals, and the SolarisTM Operating Environment software. Customers should check with their local Sun Enterprise Services representatives for program and feature availability in their areas.

SunSpectrum Service Plans are available both during and after the warranty program. Customers may choose to upgrade the service and support agreement to meet their business needs by purchasing a SunSpectrum Service Plan. There are four levels of SunSpectrum support contracts, ranging from SunSpectrum BronzeSM level to SunSpectrum PlatinumSM level. The recommended Service Plan level for the Sun Storage F5100 Flash Array is gold.

For more information on the SunSpectrum Service Plans refer to the following URL:

<http://www.sun.com/servicessolutions/>

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.

Professional Services

See the professional Services catalog at <https://icexchange.central.sun.com/> for other services and the latest updates on the Sun Services mentioned below.

Consulting Services

Sun Client Services can offer multiple levels of consulting services to ensure a smooth data migration from Sun StorageTek series of arrays to the Sun Storage F5100 array. SunPS 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 quotes for a specific customer environment.

Sun also offers a wide variety of consultative services that will help the customer architect its F5100 array into existing storage infrastructures. These services are generally custom-priced

engagements that can assist with the design and implementation of larger storage architectures. They 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.

Host-Based Data Migration Service

Sun's Host Based Data Migration (HBDM) Service makes relocation or consolidation of data simple and non-disruptive to your business. Whether you are refreshing your Sun servers and disk hardware or adopting Sun systems for the first time, the HBDM service is easy to implement, flexible, and transparent to applications that require access to the data. Our data migration capabilities are vendor-neutral and application-agnostic. We can move data in both mainframe and open systems environments. And we can move your data locally (SAN or direct attach) or globally (via TCP/IP). Whatever your data migration needs, Sun Storage Services provides expert project planning and delivery, so end users no longer have to risk data loss or corruption, or incur expensive downtime in order to move their data.

Backup and Recovery Service

The Sun StorageTek Backup and Recovery Consulting Custom Service includes an analysis of Customer's current backup and recovery environment, with recommendations on how to improve performance, backup window, hardware and software upgrades and migrations; how to implement a backup strategy; or how to completely outsource a customer's data backup needs. This is a custom engagement that must be tailored to the customer's requirements to determine scope and price.

Business Continuity and Disaster Recovery Services

The Sun StorageTek Business Impact Analysis (BIA) service will help a customer gather the necessary business information, and identify the associated operational and financial impacts, in the event of a business disruption. Findings are analyzed and documented in a business impact analysis report that includes recommendations and suggested solutions. This is a 1-2 week fixed-price engagement limited to one customer location and 3 business units at that location.

A custom version of this service is also available.

Glossary

1U or RU	One rack unit as defined by the Electronic Industries Alliances (EIA). A vertical measurement equal to 1.75 inches.
ECC	Error Correcting Code. A type of memory that corrects errors on the fly.
FRU	Field Replaceable Unit.
Hot-pluggable	A feature that allows an administrator to remove a drive without affecting hardware system integrity.
Hot-swappable	A feature that allows an administrator to remove and/or replace a device without affecting software integrity. This means that, while the system does not need to be rebooted, the new component is not automatically recognized by the system.
I/O	Input/output. Transferring data between the CPU and any peripherals.
IOPS	I/O per second. A measure of random access performance
FMod	Flash Module. Similar to a Solid Disk Drive (SSD) but not in a disk form factor. The F5100 can hold up to 80 such Fmods (Flash Modules)
MTBF	Mean Time Between Failures. The average time a component works without failure.
SSD	Solid State Disk
SATA	Serial Attached ATA. The resulting evolution of the ATA (IDE) interface from a parallel to a serial and from a master-slave to a point-to-point architecture with data transfer speeds up to 1.5 Gb/sec.
SCSI	Small Computer Systems Interface. Pronounced “scuzzy.” An ANSI standard hardware interface that allows the connection of up to 15 peripheral devices to a single bus.
SAS	Serial Attached SCSI (SAS). SAS depends on a point-to-point serial protocol that replaces the parallel SCSI bus technology that first appeared in the mid 1980s in data centers and workstations and it uses the standard SCSI command set. Currently SAS-1 is 3Gbit/sec performance, but in 2010 it is expected to double to 6Gbit/sec (SAS-2). SAS-2 is expected to offers backwards-compatibility with SAS-1..
SNMP	Simple Network Management Protocol. A set of protocols for managing complex networks. The first versions of SNMP were developed in the early 80s. SNMP works by sending messages, called protocol data units (PDUs), to different parts of a network. SNMP-compliant devices, called agents, store data about themselves in Management Information Bases (MIBs) and return this data to the SNMP requesters.
X86	Refers to the Intel 8086 family of microprocessor chips as well as compatible microprocessor chips made by AMD and others.

Resources: Sun™ Storage F5100 Flash Array

Following Sun Storage F5100 Flash Array resource materials are available as noted.

Collateral	Description	Purpose	Distribution	Token # or COMAC Order #
Product Literature				
– Sun Storage F5100 Flash Array, Just the Facts (JTF)	Reference guide, FAQs (this document)	Training Sales Tool	SunWIN, Reseller Web	SunWIN #569042
– Sun Storage F5100 Flash Array, Customer Presentation	Customer presentation	Sales Tool	SunWIN, Reseller Web	SunWIN #555647
– Sun Storage F5100 Flash Array, Technical Presentation	Technical presentation	Sales Tool	SunWIN, Reseller Web	SunWIN #569082
– Sun Storage F5100 Flash Array Data Sheet	Data sheet	Sales tool	SunWIN, Reseller Web	SunWIN #555649
– White Paper: A New Approach to Database Storage Optimization: Sun Flash Array	White paper	Sales tool	SunWIN, Reseller Web	SunWIN #567883
– Sun Storage F5100 Flash Array Database Accelerator Deployment Assessment Checklist	Deployment Assessment Checklist	Sales tool	SunWIN, Reseller Web	SunWIN #569382
– Sun Storage F5100 Flash Array Use and Deployment Best Practices	Recorded Webinar	Sales tool	SunWIN, Reseller Web	SunWIN #567730
– Sun Storage F5100 Flash Array Sales Guide	Sales Guide	Sales tool	SunWIN, Reseller Web	SunWIN #555654
External Web Sites				
– Sun Web Site	http://www.sun.com			
– Sun Network Storage Main Page	http://www.sun.com/storage			
– Sun Storage Flash Main Page	www.sun.com/flash			
– Sun Storage F5100 Flash Array Main Page	www.sun.com/f5100			

Collateral	Description	Purpose	Distribution	Token # or COMAC Order #
– Blue Print: Accelerating Databases with F5100 Flash Array	www.sun.com/f5100			
– Sun Storage F5100 Product Tour Video	http://sunbtv.sfbay.sun.com/prod_files/10Q1/10A02201.0/10A02201_00_REV_02.mp4			
– Flash Analyzer Tool	https://cds.sun.com/is-bin/INTERSHOP.enfinity/WFS/CDS-CDS_SMI-Site/en_US/-/USD/ViewProductDetail-Start?ProductRef=SWAT-3.02-G-F@CDS-CDS_SMI			
– Sun Power Calculators	sun.com/powercalculators			
– Sun Storage F5100 Power Calculator	www.sun.com/storage/disk_systems/sss/f5100/calc/index.jsp			
– Sun Solaris ZFS	sun.com/solaris/zfs.jsp			
– CAM (Sun StorageTek Common Array Manager) Software	sun.com/storagetek/management_software/resource_management/cam			
– Sun Services	sun.com/services			
Internal Web Sites				
– Sun Storage F5100 TOI Session on Deployment and Best Practices	https://slx.sun.com/1179275417			
– Training: Sales Essentials Course	http://learning.sun.com/lms/launch.jsp?courseId=WZC-ST5100-301			
– Sun Space/OneStop – Sun Storage F5100 Flash Array	https://sunspace.sfbay.sun.com/display/Onestop/Sun+Storage+F5100+%28aka+Lightning+Flash%29			
– SVP Tools (Sun Internal)	https://valueplatform.sun.com/SVP			
– SVP Tools (for Partners)	https://valueplatform.sun.com/SVP/partner			
– Sun Storage F5100 Flash Array Documentation	http://docs.sfbay.sun.com/coll/f5100flash-array_INT?l=en			
– Sun Storage F5100 Flash Array Performance Information	blogs.sun.com/BestPerf/entry/1_4_million_iops_in			
– Flash Performance Tuning	wikis.sun.com/display/Performance/Home#Home-Flash			

Internal Information: Sun™ Storage F5100 Flash Array

The Sun Storage F5100 Flash is a market leader, delivering well over 1 million IOPS in just 1RU of space with 100x better power and space efficiency than traditional disk. It offers best performance, eco-efficiency, reliability, and cost savings

Partner Information

This section contains information for Sun partners. The information in this section should be shared with Sun partners, Sun sales reps and SEs. This section should not be shared with Sun end user customers.

The Potential

Sun's message about the Sun Storage F5100 Flash Arrays is proving to resonate well with a wide range of customers seeking to improve their applications performance and reduce operating cost. Its world leading performance, density and eco-efficiency provide key value and differentiation in the marketplace. Sun is experiencing an overwhelming response to its Open Storage events, which are exceeding attendance capacity limits and are being extended to more shows in more locations based on demand. The conversion rate from seminar attendees to follow-up leads so far ranges from 20% to over 40% of attendees. The F5100 is an Open Storage product.

As a go-to partner you will be offering the huge potential of this technology in almost limitless new technology areas. You'll also have valuable opportunities for selling into new departments and into new environments within current Sun installed-base customer organizations.

Why Partner with Sun

Sun offers the best products available, with industry-leading IP and strategy.

Sun is the support team for your sales team. As you have seen, our commitment to partners and customers is outstanding. Please review our wide range of sales support programs and resources in the Sales Tools section.

Sun backs up your efforts with OpReg programs that include New Business Incentives, the Competitive Knock-Out Program, OpReg Service Renewals, and more. For full information on Sun's OpReg programs visit the Sun OpReg page.

>**Note:** Access the Sun OpReg page via the Partner Portal at:

<http://partner.sun.com/US/programs-n-policies/partnerfirst/opreg/index.html>

(Must have Sun Partner portal account and password.)

See also *Sun partner OpReg matrix.pdf*

Your Bottom Line

The Sun Storage F5100 Flash Array and Sun's other open storage solutions have been carefully architected with your bottom line in mind. The Sun Storage F5100 has a very strong value proposition that's easy to sell, not only based on improving productivity but also with the purpose of slashing operating costs for a quick ROI.

- **Retained margins:** The pricing and discount categories for Sun Storage F5100 Flash Arrays provide you with the flexibility to manage your business as it works best for you.

Sun Storage F5100 Flash Array : Just the Facts
Sun Confidential: Internal and Sun Channel Partners Use Only



Whether you want to shorten sales cycles with deeper discounts or sell the full value, there are retained margins and dollars that give you room to be very profitable. For these solutions Sun will be applying a list-minus pricing model, with the F5100 Flash Array falling in Category M.

- **Lower costs:** With F5100 Flash Arrays as low as 75% less than the cost of competitors, and with higher performance and scale, space-saving footprints, and energy consumption as little as 1/3 that of leading competitor's solutions, you can now close business that was previously unavailable to you. When the name of the game is performance, the Sun Storage F5100 Flash Array delivers not only record performance, but the best IOPS/\$, best IOPS/watt of power and best IOPS/space of any other competitor or storage device.
- **Closing the deal:** Leverage the Sun Value Tool to help accelerate sales; it provides your prospects with a detailed ROI analysis that lets them see the significant cost savings Sun open storage systems represent.
- **Up-sell and value-add opportunities:** The F5100 Flash Arrays can act as a complement to existing storage solutions, opening the door to quick upgrades. The modular and flexible nature of Open Storage solutions allows you to offer a range of enhancements, such as upgrading processors and memory or adding enterprise-class high-capacity drive expansion units and SSD drives. These upgrades also open a lot of high-margin consulting and integration services opportunities.
- **Alternatives:** You now have a new range of compelling alternatives to many of the over-distributed competitor systems — and a way to resolve any unaddressed customer concerns by providing easy-to-deploy, easy-to-operate storage solutions with unheard-of economy.
- **Lead generation:** With the F5100 Flash Arrays' compelling and highly visible end-customer value proposition, you will be able to generate and convert sales leads with greater efficiency than before.

Open Storage Partner Specialty Program

The Sun Storage F5100 Flash Array will help Sun partners appeal to large new storage opportunities, allowing them to grow their businesses and secure footholds in emerging markets. Supported by the Sun Partner Advantage Program, Sun partners are offered a simple, highly aligned support structure — integrated across systems, software and services — that gives partners the flexibility to resell servers, storage, software products, and services within a single channel-focused program. As part of the Sun Partner Advantage Program, the Open Storage Specialty gives participating partners the skills, tools, and access to Open Storage solutions that they need to expand their revenue. Here are some of the benefits of this program:

- The Open Storage Specialty includes the information, intellectual capital, training, and support partners need to sell and deploy Open Storage Solutions
- Participation in the Open Storage Specialty also enables partners to distinguish themselves to companies through Sun Open Storage Specialty branding and as “go-to” partners with the Sun field sales organization
- Partners have access to Sun's migration tools to deliver Sun's new Unified Storage Data Migration service, enabling them to manage the entire upgrade to F5100 Flash Arrays
- 1,290 partners worldwide are accredited to sell Sun Storage F5100 Flash Arrays through their accreditation on the Storage Entry and Mid platforms; any Sun authorized partner can sell F5100 Flash Arrays

Availability to All Partners

The Sun Storage F5100 Flash Array will be available to ALL partners at GA. (10/25/2009.)

Partner Certification and Training

No certification required. Partner eligibility to sell will be determined by Partner status under the Sun Partner Advantage (SPA) program.

Sun Product Training Portal

Web-based training on Sun Storage solutions is available to Sun Channel Partners worldwide on the Sun iForce Partner Portal:

1. Log into <http://mysun.sun.com>
2. Click on Web Based Training in the Partner Competency section.

Note: If you don't have a My Sun Username, register at <https://pwregister.sun.com>

View new product training and update training modules in the "What's New" section on the main Sun Product Training Portal screen.

Available training courses

The following internal employee accreditations will be available on the at release or shortly after:

- F5100 Flash Array Sales Essentials (available September 12, 2009)
- F5100 Flash Array Technical Essentials (available October, 2009)

The following Sales Essentials and Technical Essentials training courses are part of the SPA program and are scheduled for availability on the following dates:

- F5100 Flash Array Sales Essentials (available September 12, 2009)
- F5100 Flash Array Technical Essentials (available October. 2009)

Cross-Sell and Up-Sell Opportunities

Sun's servers and software

- See Sun.com

Sun storage

- SAS HBAs
- Tape libraries
- Disk drives

Support Services

Where available, customers will be able to upgrade their warranty to the SunSpectrumSM PlatinumSM, GoldSM 7/24, Gold, and SilverSM support levels. Contact local Sun Services sales

representative for details.

SunSpectrumSM Service Plans

With decades of systems and storage expertise, Sun knows how to store your data securely, manage it intelligently, and provide ready access to it all so you can maximize its use to grow your business and gain competitive advantage. Our reliable services and support give you confidence that your information will be available when you need it and your business will be resilient to change.

SunSpectrum Service Plans are available both during and after the warranty period. Customers may choose to uplift the service and support agreement to meet their business needs by purchasing a SunSpectrum Service Plan. The recommended level of support for these products is the Sun StorageTek GOLD Service Plan.

For more information on the SunSpectrum Service Plans refer to the following URL:

<http://www.sun.com/service/serviceplans/index.jsp>

Enterprise Installation Services (EIS)

Installation Services allow the customer to accelerate the deployment of their IT assets by utilizing Sun's proven best practices. Sun will rack, cable and connect the asset to the host.

F5100 Flash Array Sales Program Overviews

Try and Buy Program

- Try and Buy will not be available for the F5100 Flash Array a first release

IBB Upgrade Advantage Trade-in Program

- There is an IBB Upgrade Advantage trade-in program for the Sun Storage F5100 Flash Array. Qualified Sun and non-Sun systems are eligible for trade-in credit when purchasing new F5100 Flash Array. Trade-in allowances range from 5-15%, depending on the return. For more information, please visit one of the following site:
<http://www.sun.com/tradeins/index.jsp>

Competitive Information

F5100 Flash Array Competitive Strengths:

What's Unique About This Product: Key Sun Strengths	
Feature	Competitive Strength
Over 1.6M read IOPS and 1.2M write IOPS in 1 RU of space	<ul style="list-style-type: none"> >First and only solid state storage array to deliver this level of performance in just 1RU >Ideal for accelerating Oracle and SAS databases
Only 300 watts (or 3x 100 watt light bulbs) for 1M IOPS; 200W idle, 300W active	<ul style="list-style-type: none"> >Uses 1/100th the space and power of traditional storage for the same performance level >Most eco-efficient storage product on the market to save power and space
Up to 2 TB in 1 RU of space, scaling up to 80 TB and over 50M IOPS in one storage rack (68M read and 50M write IOPS)	<ul style="list-style-type: none"> >The most space-efficient solid state array, equaling the performance of 120,000 disks in 400 racks > Scalability to meet all performance needs
<ul style="list-style-type: none"> > Greater than 2M MTBF hours Flash Modules, extra write endurance, advanced wear leveling, energy storage protection >Solaris: FMA self-healing 	<ul style="list-style-type: none"> >Most reliable flash array product on the market (2x the life of most disk drives) >Over 6 years write endurance with normal use >Much greater data protection from power outages and no need for battery-backed cache in case of power failure >For Solaris: FMA automatically and silently detects, diagnoses, and resolves problems to yield increased uptime and service levels
<ul style="list-style-type: none"> >Redundancy: hot-swap power, fans, and separate storage domains for mirroring >ZFS provides software RAID protection 	<ul style="list-style-type: none"> >High availability with host based mirroring (Oracle ASM, OS VSM or Veritas VxVM) >With ZFS, the highest level of added error correction and thin provisioning
Industry-leading \$/IOPS, W/IOPS, and IOPS/RU	<ul style="list-style-type: none"> >Best price/performance in its class; best storage performance per dollar, best storage performance per watt of power >Best TCO among comparable products <p>Flash Savings vs. traditional high-performance disk 1million IOPS solution (aka "Short Stroke")</p> <ul style="list-style-type: none"> >Best IOPS/\$, IOPS/Watt, and IOPS/per sq ft <ul style="list-style-type: none"> • Spinning Disk = \$3.00 per IOPS • Sun Flash = \$.08 per IOPS <ul style="list-style-type: none"> >1/100th power and space of HDDs for same performance >Scalable to 80TB of solid state and over 50 Million IOPS in single rack (vs. over 143,000 HDDs in 400 racks)
<ul style="list-style-type: none"> >10x-15x better access service times >Highest access density or throughput per GB (833 IOPS/GB) 	<ul style="list-style-type: none"> >Reduction of operational costs associated with database tuning and deployment; example: less tuning to achieve same client-side response time >Ability to scale to a larger number of clients with less tuning >Capability to run additional value-add processing due to fewer bottlenecks

Key Competitors: Systems and Storage Vendors

There are two levels of competitors to Sun's F5100 Flash Array:

- Systems/storage vendors** include IBM, HP, Dell, and EMC. As with all shifts to newer technology, one-to-one mirror comparisons are challenging. These competitors do utilize various forms of solid state flash/disks, but their offerings are not as mature, nor do they offer as broad a portfolio as Sun. Also Sun's ZFS file system consistently outperforms all others in data management across storage tiers, in independent tests and reviews. Closest in a direct comparison is from NetApp's announcement of an OEM agreement with Texas Memory Systems to resell the company's RamSan 620 and its earlier RamSan 500 products.
- Flash array vendors** include smaller companies such as Violin, Texas Memory Systems, Fusion I/O, etc. These second-tier vendors whose products compete with the F5100 Flash Array do not have the same level of system expertise or credibility to offer the same level of integration, reliability, or performance that Sun delivers. They lack system knowledge, worldwide service, or support. Sales should focus on Sun's systems expertise and broad portfolio of solid state flash products.

<i>Server and Storage System Vendors</i>			
	Flash Arrays	Flash Controllers/PCIe Cards	Flash SSD Drives
IBM	None	Announced resale of Fusion IO, but no shipments (Fusion I/O offers both SLC-based and MLC-based flash PCIe cards ranging in capacity from 80 to 320 GB for SLC and 640 GB for MLC)	In some blades
HP	None	OEM'ing Fusion-io	In some blades
EMC	None	None	In DMX-4 systems (very expensive)
NetApp	Announced resale of Texas Memory Systems RamSan 500 (2/09). Will most likely do RamSan 620 as well.	Announced 256/512GB Performance Acceleration Modules early Sept, 2009	None
<i>What the Competition Offers: Flash Array Vendors</i>			
	Flash Arrays	Flash Controllers PCIe Cards	Flash SSD Drives

Violin Small company, less than 50 employees	Yes , the Violin 1010 offers 2X capacity, but much lower performance	None	None
Texas Memory Systems (TMS) Small company, less than 120 employees, but has recent agreement with NetApp to resell its RamSan 500 Flash Array	Yes (both flash-based and DRAM-based). Introduced a new RamSan 620 with more capacity, but less performance than Sun F5100 Flash Array	Introduce a 450GB RamSan-20 PCIe card mid-2009	None

Competitive Comparisons: Solid State Arrays and Flash Cards

F5100 Competitive Landscape

Competitors → Features ↓	Sun F5100 Flash Array	NetApp/ TMS Ram San-620	NetApp/ TMS Ram San-500	Violin 1010-Flash	TMS Ram San 440
Device Type	Flash Mods. (SLC)	Flash Mods. (SLC)	Flash DIMM (SLC)	Flash Mods. (SLC)	DRAM
Host Interface	3Gb SAS	4Gb FC	4Gb FC	PCIe	4Gb FC/IB
Host Ports	16 x4-wide	2-8 Ports	8	2x4, 1x8	8, 4
Max Capacity	1.92 TB	5 TB	2 TB	4 TB	512 GB
Cache (RAM)	5GB	0	16-64GB	No	n/a
Random Read IOPS (4K)	1.6M	250K	100K	400K	600K
Random Write IOPS (4K)	1.2M	250K	25K	100K	600K
Size (in RU)	1	2	4	2	4
RAID	SW		RAID 8+1	RAID 4+1	RAID
Hot Swap/HA	No	No	Yes	Yes	Yes
Cache Backup	Super Capsx1		Batteries x2	No	Batteries x3
Model	JBOD	Stor. System	Stor. System	Appliance	Stor. System
Power	300 W	300-500 W	300 W	450 W	650 W
Device Management SW	CAM	web mgmt.	web mgmt.	server	web mgmt
Data Management SW	ZFS	none	none	none	none
Ship/Availability	Aug '09	Q3'09	Q2CY08	Q2CY08	Shipping
US List (2 TB)	\$159,995	\$88,000*	\$140,000	\$120,000	> \$150,000
\$/GB (List)	\$83/GB	\$44/GB*	\$70/GB	\$60/GB	
\$/IOPS (4K) – Read	\$0.10/IOP	\$0.35/IOP*	\$1.40/IOP*	\$0.30/IOP	\$0.25/GB
\$/IOPS (4K) – Write	\$0.13/IOP	\$0.35/IOP*	\$5.60/IOP*	\$1.20/IOP	\$0.25/GB

* RamSam-620 pricing on most basic 2TB configuration with 2FC ports

Sun Storage F5100 Flash Array : Just the Facts
Sun Confidential: Internal and Sun Channel Partners Use Only



Sun Flash Accelerator F20 Competitive

Competitors → Features ↓	Sun F20	TMS RamSan-20	FusionIO ioDrive80	FusionIO ioDrive160	FusionIO ioDual 320
Figure					
Usage	Hybrid Storage	Tier-0 Only	Tier-0 Only	Tier-0 Only	Tier-0 Only
Capacity	96GB	450GB	80GB	160GB	320GB
HBA Capability (disk controller)	Yes (2x4-wide SAS-1 ports)	No	No	No	No
HAND Type	SLC	SLC	SLC	SLC	SLC
Random IO (Read)	100K IOPS	120K IOPS	120K IOPS	104 IOPS	185 IOPS
Random IO (Write)	84K IOPS*	50K IOPS	89K IOPS**	93 IOPS**	139 IOPS**
Throughput (Read)	1 GB/s	700MB/s	700MB/s	700MB/s	1.4GB/s
Throughput (Write)	600MB/s	500MB/s	550MB/s	600MB/s	1.0GB/s
RAID Support	SW RAID (ZFS)	Embedded RAID5	No	No	No
OS Support	Solaris, Linux, Windows	Unix, Linux, Window, MacOS	Linux, Windows	Linux, Windows, Solaris	Linux, Windows, Solaris
For Factor	PCIe1, Low profile	PCIe, Full size	PCIe, Low profile	PCIe, Low profile	PCIe2, Full height, ¾ length
Power	18W	15W	9W	9W	PCIe spec 2.0
Custom Drivers	No	Yes	Yes	Yes	Yes
US List	TBD	\$18,000	\$4,000	\$8,000	\$16,000

* Write Specs: Sun F20 = true 100% write (4K blocks), Fusion-io = 75%/25% read/write for "write" spec.

Summary: Sun's main competitors with regard to the Sun Storage F5100 Flash Array are Texas Memory Systems, with its new RamSan 620 product, and Violin, with its Violin 1010 flash product. While both were earlier to market than Sun by 3-9 months, and offer 2 times higher capacity, their performance is far less (¼) than that of the Sun Storage F5100 Flash Array. Both are very small companies who lack many advantages of experience and support that Sun can offer. And as a product, the F5100 Flash Array is superior in several ways:

- World-record performance (Over 4 times the nearest competitor)
- Best Performance Cost and Efficiency
 - Best IOPS/\$
 - Best IOPS/GB
 - Best IOPS/watt of power
 - Best IOPS/space
- Smallest size (1RU versus 2RU)
- Better price/performance, power/performance and space/performance
- Higher reliability (advanced wear-leveling, write endurance, etc.)
- Easier integrations, no special software or drivers required
- Integrated data protection with super capacity-based energy storage modules

Sun Storage F5100 Flash Array : Just the Facts
Sun Confidential: Internal and Sun Channel Partners Use Only



- Higher host connectivity bandwidth
- CAM management software
- Availability of Solaris ZFS
- Worldwide service and support

Frequently Asked Questions: Sun™ Storage F5100 Flash Array

Q, How does the F5100 help accelerate database applications?

A. The Sun Storage F5100 Flash Array can improve database performance by up to 10X by reducing data access times and increasing data access rates by up to 15X over traditional disk drives. For I/O constrained systems, the F5100:

- a) Improves service levels by reducing transaction response times
- b) Increases workloads and transactional scalability without increasing response times

Here are a couple of use case examples and their benefits of:

- 1) Placing just index files on the F5100 and
- 2) Placing both index and data files on the F5100.

Use Case A (small capacity configuration):

Put Index files on F5100 (some 25%-50% of dB size)

=> Improve IO latency performance 10X-15X

Store Data files on disk (some 50%-75% of dB size)

Keep Log files on NVRAM or enterprise disk cache (some 200-400MB)

=>> Overall benefit: Improves dB transaction times 30% or more depending on dB size, IO levels and index depth

Use Case B (larger capacity configuration):

Put Index files and data files on LF

=> Improve IO latency performance 10X-15X

Keep Log files on NVRAM or enterprise disk cache

=>> Overall benefit: Improves dB transaction times over 50% depending on dB size, index and data IO level

Q. How does the F5100 address customers' needs and provides business value?

A. The Sun Storage F5100 Flash Array:

- Accelerates applications, increases business productivity and response
- Reduces latency and eliminates storage IO bottlenecks
- Lowers Power and cooling costs
- Improves space utilization
- Optimizes performance and scaling with less cost

- Improves efficiency, reduces TCO

Q. What is F5100 availability and launch schedule?

A. Presto Announce: September 15, 2009

Revenue Release: August 28, 2009

General Availability: September 24, 2009

Launch Event: Oracle World (OOW, October 12, 2009)

Q. What HBA cards are supported ?

A. The F5100 Flash Array supports the Sun StorageTek external SAS PCIe HBA and the StorageTek SAS Express Module HBAs:

- Sun StorageTek external SAS PCIe HBA *
Sun PN: SG-XPCIE8SAS-E-Z/SG-PCIE8SAS-E-Z *
- Sun StorageTek SAS Express Module HBA for blade servers
Sun PN: SG-XPCIESAS-EB-Z/SG-PCIE8SAS-EB-Z

* Requires new HBA firmware. New HBA firmware will double the HBA performance from approximately 45K IOPS to over 100K IOPS. This new HBA firmware (ver 1.27.03) for the F5100 can be downloaded free by going to the following URL: <http://www.lsi.com/support/sun>

Q. What is needed to maximize the SAS PCIe HBA performance ?

A. New SAS PCIe HBA firmware is available that will more than double the existing HBA IO performance from approximately 45K to 105K IOPS. This firmware is required to take full advantage of F5100 enormous performance capability. This new HBA firmware (ver 1.27.03) for the F5100 can be downloaded free by going to the following URL: <http://www.lsi.com/support/sun>

Q. How many SAS HBAs do I need to maximize performance ?

A. Given F5100's unprecedented performance capability of over 1.6 Million IOPS, the more HBAs the greater the throughput performance. Assuming each HBA can deliver 104K IOPS max (with level 1.27.03 HBA firmware), this would imply 16 HBA cards. A minimum of 1 HBA is required for each 20 Flash Modules ("SSD drives") or 4 HBAs for a fully populated 80 Flash Module (1.92TB) F5100 Flash Array. Up to 16 HBAs can be configured with each F5100 when using zoning (4 without zoning)

Q. What if my server is limited on available PCIe slots ?

A. Please look into the Sun External I/O Expansion Unit for SPARC Enterprise Servers, which provides up to 12 additional PCIe IO expansion slots for Sun SPARC Enterprise M4000, M5000, M8000, M9000, T5120, T5140, T5220, T5240, and T5440 servers.

Q. What are some other necessary considerations for optimizing F5100 performance ?

A. The Sun Storage F5100 Flash Array is performance optimized for heavy IO reads and 4K aligned block writes. 4K block sizing is where the computer industry is moving towards for best performance and storage efficiency. It is not recommended to deploy the F5100 into environments that use < 4K write block sizes. The F5100 has been designed for optimum performance around 4K and 4K aligned write blocks. To better understand how to optimize F5100 performance and set up the OS, file system and the application for 4K block alignment, please refer to the Sun Storage F5100 Flash Array Optimizing Performance and Availability white paper for storage administrators available on Sun.com. This white paper describes the array's innovative architecture and technologies that bring tremendous business benefit to performance-hungry database applications and its use and configuration for optimum performance and availability.

Q. How much performance improvement will a customer see with the F5100 ?

A. It all depends on their use, workloads and configurations. The F5100 Flash Array will eliminate IO bottlenecks and reduce latencies to allow better response time and transactional scalability by cutting the service times for I/O reads and writes by a factor of 10X or more.

In a database environment it also depends on the dB size, what hot files and indexes you put on it along with your IO levels and index depths In the Sun performance lab we were able to demonstrate 1.6 million read and 1.2 million write IOPS (4K)with 1 F5100 unit using 4 x Sun T5240 servers and 16 SAS-1 HBA cards driving it, so the F5100 has the capability of tremendous performance but a what a customer will see depends greatly on their compute environment and the applications and workloads they are running.

In summary, the F5100 can solve most of their storage latency and IO bottleneck problems with room for greater performance scalability, but if they are getting bottlenecked elsewhere in the system the final results will vary. The F5100 has the ability to removes most of their storage performance bottlenecks and unleash the potential for great performance improvements provided that are not constrained elsewhere in the system.

Q. How do I configure the F5100 for High Availability (HA)?

A. For HA configuration, it is recommended to use mirroring. This can be done across the internal F5100 domains or between two different F5100 units for greater redundancy. Host based mirroring software, such as Oracle ASM, can be used to enable this.

Q. Does the F5100 support clustering ?

A. No, currently clustering is not supported with the F5100. While the F5100 supports multiple host connects (up to 16), it does not offer storage sharing to support clustering at this time.

Q. How do I know if my customer can benefit from the F5100 Flash Array?

A. Several tools are available to assess application performance and IO bottlenecks, including Oracle Systems Accounting Tools and Sun Flash Analyzer which is available for a free download from <http://www.sun.com/storage/flash/resources.jsp>.

The Sun Flash Analyzer collects, analyzes, and displays data from Solaris, Windows and Linux systems to assess workloads and IO bottlenecks. It also comes with "sample" collected data so you can experiment with the tool before you collect live data on your application and/or systems.

In addition, please refer to the F5100 Sales Assessment Checklist for target application, environments and best fit for the F5100.

Q. Does the Sun Storage F5100 have certain air ship restrictions ?

A. Yes. Do to a small amount of flammable liquid (electrolyte) contained within the super capacitors of the Energy Storage Modules, the F5100 has been classified as Class 9 dangerous goods item. While Class 9 is the lowest level of dangerous goods designation (which includes such items as rechargeable batteries), it does necessitate certain air ship restrictions and handling considerations. Please refer to the Sun Class 9 Dangerous Goods and F5100 documentation for additional information on air ship restriction. This is not an issue for ground shipping.

Q. Will the F5100 come with a rail kit and SAS cables ?

A. Yes, all three F5100 PTO configurations come standard with a rail kit (snap in type), a cable management arm, and two 3-meter SAS cables.

Q. Will the F5100 upgrade Flash Modules be customer installable?

A. Yes, the Flash Modules (sold in 4-packs as capacity upgrades) will be customer installable and can be added to one or all 4 domains depending on need.

Q. What is the difference between MLC Flash and SLC Flash technology?

A. The MLC (Multi Level Cell) Flash offers more capacity but much lower performance and write endurance. It is mostly use in read oriented applications. The SLC (Single Level Cell) Flash while having less capacity offers over 10X better write endurance, reliability and performance. The F5100 uses only SLC enterprise-class Flash with advanced wear-leveling, write endurance and performance enhancement for the highest level of reliability, durability and performance.

Q. What's the maximum length of SAS cables supported on the F5100?

A. 3 meters and less.

Q. In what environments can I utilize Sun's ARS Flash Storage Optimization for Database Applications module?

A. Non-clustered environments, 4 Flash storage array domains. Only Oracle and MySQL database environments are available at this time.

Q. Is the ARS Flash Storage Optimization for Database Applications module available for non-Solaris systems?

A. No this service is limited to the use of Solaris based systems either Sun or non-Sun hardware. Flash storage arrays are also supported on Windows, but not this service.

Q. What management software will be available for customers to configure and manage the F5100?

A. CAM (Common Array Manager) software provides device management for the F5100 . CAM v6.5 supports the F5100 Flash Array and will be available for a no cost download starting September 15, 2009

The Sun StorageTek Common Array Manager (CAM) is the standard management software for the Sun Storage F5100 Flash Array as well as other Sun storage arrays. With features like storage pools and a built-in performance utility, this software simplifies administration and enables comprehensive management of multiple arrays (e.g., ST2540, ST2530, ST6140 and ST6540). In addition, CAM is included with the array at no additional licensing or maintenance costs.

Solaris ZFS: Automates common administrative tasks, protecting data from corruption and providing virtually unlimited scalability. Solaris ZFS provides software RAID protection and uses virtual storage pools to make it easy to expand or contract the F5100 arrays simply by adding or deleting Flash Module. Solaris ZFS Hybrid Storage Pool provide data management across different memory and storage pools for best performance and lowest overall storage costs.

Q. Where can more comprehensive information about ZFS be obtained?

A. Please access the following URLs for a deeper understanding of Sun's ZFS offering:

Sun.com

http://www.sun.com/software/solaris/data_management.jsp

Learning Center

http://www.sun.com/software/solaris/zfs_learning_center.jsp

How to guides

<http://www.sun.com/software/solaris/howtoguides/zfshowto.jsp>

Data sheet

<http://www.sun.com/software/solaris/ds/zfs.jsp>

Big Admin

<http://www.sun.com/bigadmin/topics/zfs/>

OpenSolaris

<http://opensolaris.org/os/community/zfs/>

Documentation

<http://opensolaris.org/os/community/zfs/docs/>

Best Practices

http://www.solarisinternals.com/wiki/index.php/ZFS_Best_Practices_Guide

Training Class

<http://www.sun.com/training/catalog/courses/SA-229-S10.xml>

Q. Is cascading supported on the F5100?

A. No, cascading multiple F5100 units behind a single HBA will significantly limit its performance. As such, cascading multiple F5100 units behind a single HBA or cascading other storage devices behind a F5100 is not supported.

Q. When will Clustering and Oracle RAC be supported on the F5100?

A. Clustering will be supported with the release of SAS Flash Modules. This will be a future enhancement. Current SATA Flash Modules do not support data sharing between hosts but up to 16 servers can share the F5100 with Zoning and individual Flash Module access (not sharing individual Flash Modules)

Q. Does the F5100 support multipathing?

A. Multipathing is not supported. It is highly recommended that customers buy two units (half populated) and mirror them for high availability and redundancy.

Q. Does the F5100 Flash Arrays come packaged with SAS interconnect cables?

A. Yes. Two 3 meter cables come with every F5100 unit.

Q. Does the F5100 Flash Array come packaged with a rail kit?

A. Yes. A tool-less, snap-in type rail kit comes with each F5100 Array. A cable management arm is also included.

Q. What is the difference between a CRU and a FRU?

A. Field Replaceable Units (FRUs) are typically replaced by Sun authorized and trained personnel. A Customer Replaceable Unit (CRU) has been designed to be easily replaced by the customer and is designated as customer replaceable. FRUs and CRUs are easily identified in the Sun System Handbook, available to customers with a Sun Service Plan.

Q. How will customers receive replacement CRUs in the event of a failure?

A. A customer will open a service call with Sun. A Sun technician will confirm the CRU failure and schedule shipment of a replacement CRU to the customer with a prepaid return label. The customer will remove the failed CRU and return it to Sun in the box in which the new unit was shipped, utilizing the prepaid return label.

Q. How will a customer know how to replace the CRU?

A. Instructions for removal and replacement are available via the Service Advisor which is accessible through StorageTek Common Array Manager. For easy reference, hardcopy removal and replacement instructions will also be provided as part of the CRU shipment.

Q. Are installation services available for this product?

A. Yes. While the system is "customer installable," Sun offers installation services to help accelerate the deployment of this asset.

Q. What are the deliverables of these installation services?

A. Our installation services cover the following:

- Racking and cabling
- Connection to host
- Test for I/O

Q. Which SAS interconnect cable lengths are supported on the F5100?

A. The supported cable lengths are as follows:

- .5m, 1m, 2m and 3m SAS cables. 6 meter cables are not supported.