RETHINK STORAGE WITH
SUN™ STORAGE 7000 SYSTEMS

White Paper
November 2008

This Page Intentionally Left Blank

Table of Contents

Chapter 1

Introduction

The rapid growth of new digital data demands new storage architectures that offer more flexibility and radically different storage economics. IDC estimates that the total amount of digital information created, captured, and replicated will grow at a rate 58% per year, reaching 1,610 EB by 2011[1]. Web 2.0 applications are growing at a tremendous rate and require highly scalable and affordable storage. High performance computing (HPC) environments also have large datasets where performance needs dictate vast amounts of inexpensive storage and yet HPC systems also have stringent requirements for high I/O bandwidth and low latency.

As storage capacities have grown, traditional means of deploying and managing storage have also become outdated. Most of today's storage systems require a highly trained administrator to effectively manage the environment. This adds time and cost to deployment of new storage systems and increases the ongoing cost of managing the environment.

Perhaps more importantly, with customers implementing or maintaining large storage networks, it has become increasingly difficult to troubleshoot and fix storage performance issues. The tools currently available offer limited visibility because they lack an end-to-end view that encompasses multiple layers from the CPU and application to the storage file system, operating system, and data services.

Customers can no longer tolerate the high costs of proprietary storage or massive licensing fees and are looking for new ways to address their growing storage requirements and their challenges in managing storage environments. Today's IT environments require storage solutions that can offer:

- Greater simplicity and ease-of-use
- Real time diagnostics and tuning
- Massive scalability
- Better storage economics

[1] "The Diverse and Exploding Digital Universe," IDC, March 2008.

Chapter 2

Open Storage

A new market for open storage solutions is developing in response to today's storage requirements. Sun's best estimate is that open storage products and services will represent just under 12% of the total storage market in 2011. With IDC estimating the total storage market (hardware, software, and services) at approximately $90 billion in 2011, Sun expects the open storage portion to be just over $10.6 billion. Sun also predicts that by 2011, open storage will represent more than 20% of the external disk market, or approximately U.S. $5 billion of the U.S. $24 billion external disk market.

What is open storage?

Open storage systems combine industry-standard hardware with open-source software, and are supported by a community of thousands who have a passion to create better storage solutions. This powerful combination helps spur innovation and drives better storage economics. Developers can leverage volume servers and disk drives as well as an open storage software stack to speed storage innovation.

Unlike traditional proprietary storage solutions, open storage solutions offer freedom of choice at every level of the storage system stack (Figure 1).

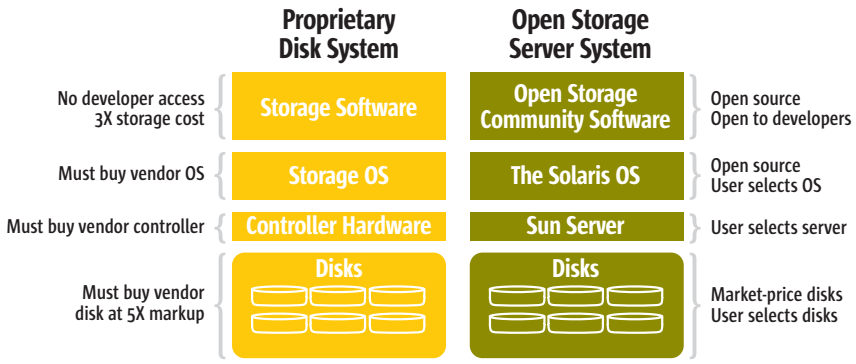

Figure 1. Proprietary versus open storage systems

In proprietary systems, storage software is closed and is typically three times the cost. Customers must generally purchase a storage controller and controller software from the storage vendor and then also pay for individual features, usually with capacity based software licenses. Lastly, customers must buy the disks enclosed within the storage system. In most cases these are commodity disks that are marked up as much as five times the original cost.

"Sun has provided a platform for the democratization of the storage industry. We have found the appropriate level of operating system support we need to run our business through the OpenSolaris storage community, which saves significant time and money. I participate in the community daily and see real business value in the projects that are being created by some of the industry's most important players."

— **Jason Williams, COO and CTO, DigiTar**

By contrast, in open storage architectures, both the operating system and the application software are available as open-source software. The hardware is also industry-standard, so customers can leverage an industry-standard server in place of an expensive, proprietary disk controller. ZFS, which is included in the Solaris™ Operating System (OS) for no additional licensing fee, provides data services such as RAID, error correction, and system management, thus insulating applications from the underlying hardware. Such services have traditionally been tied to specific hardware devices and have been available only when bundled with an expensive controller. Lastly, customers can purchase high performance, affordable disks systems at market prices, enabling much greater price/performance.

Advantages of open-source software

Open-source software has several critical advantages for new storage systems. Freely available source with free distribution rights gives developers a complete platform on which to build, modify, and improve. They can concentrate on unique software value without reinventing an entire storage system. Developers can also work with the open-source community to help solve problems and provide support. This will enable many smaller companies to enter the storage market, bringing innovation and driving progress similar to the advancements on the Web.

The benefits of community development and open-source software have spread widely and can include:

- **Lower costs** – No software licensing fees and low cost community support.
- **Choice** – Users do not need to wait for a single vendor to develop new software functionality when they can find similar functionality already available through the development community or can help develop the functionality they need through the community.
- **Innovation** – Software and application innovation is no longer limited by a vendor's business goals or research and development budget.
- **No lock-in** – Users are not held hostage by a single software supplier.

Sun Open Storage solutions are based on an open storage architecture and the open-source OpenSolaris™ Operating System. Sun leverages the benefits of open storage as listed above to meet today's challenging storage requirements. Additional information about open storage trends and Sun Open Storage offerings is available in the white paper entitled, "Open Storage Adoption," which can be found on the Web at *sun.com/storage/openstorage/OpenStorAdoption.pdf*.

Chapter 3

Sun™ Storage 7000 series systems

The new Sun™ Storage 7000 Unified Storage Systems, provide all the benefits of open storage in an easy to use package so organizations can radically simplify their storage deployment and management. These unified storage systems provide unmatched simplicity and ease-of-use through an intuitive and powerful management interface. Revolutionary business analytics functionality also enables administrators to quickly diagnose and resolve performance issues in production systems.

The Sun Storage 7000 series also changes the economics of storage by utilizing a high-performance hybrid storage pool architecture, industry-standard components, and open storage interfaces. The systems offer up to 576 TB of capacity[2] using a combination of solid-state drives (SSDs), DRAM, and hard disk drives (HDD). Hybrid storage pools utilize SSD for read and write cache, enabling higher performance than traditional storage architectures at up to 75% less cost.

Figure 2. Sun Storage 7000 series family

Breaking through economic and performance barriers

While modern servers are fast approaching processing capabilities in excess of one million I/O operations Per Second (IOPS), mechanical disk drives have failed to keep up with the server performance growth curve. Today's fastest drives are capable of 300 to 400 IOPS. To achieve a level of storage I/O throughput that matches today's server performance and meets the demanding requirements of today's data-intensive applications, most storage vendors recommend a large pool of 15K RPM disk drives. In some cases, a large buffer of expensive DRAM is also deployed so that the entire working set can be stored in memory to reduce latency.

[2] The initial release offers 288 TB in usable capacity and a free software upgrade will be available in the near future to enable the system's full 576 TB of capacity to be utilized.

Not only are these traditional storage architectures complex and costly to acquire, but they are also expensive to operate and manage. High RPM drives consume significantly more power than lower RPM drives. And the software tools that are used to manage these environments typically require a highly trained administrator, adding further costs for ongoing management and making it more challenging to maintain high service levels for storage systems and related applications.

Recent advances in production of flash technology have made solid-state drives (SSDs) much more cost-effective, enabling a new approach to tiered storage. SSDs are interesting because they fall in a cost and performance sweet spot between mechanical drives and DRAM.

Mechanical disk drives are very cost-efficient for storing data, but they are also relatively slow compared to SSDs. SSDs can be used to increase the performance of mechanical drives by providing read and write cache that offers very low latency while also delivering greater power efficiency. While SSDs are significantly more expensive than mechanical drives, they are significantly cheaper and denser than DRAM storage.

Taking advantage of SSD with the ZFS file system

Taking advantage of the performance and cost characteristics of SSD requires an enabling technology that recognizes SSD and can transparently utilize it to drive better application and file system performance.

The ZFS file system can take advantage of SSDs today by transparently caching data on SSDs, overcoming the need to modify applications. Unlike less sophisticated file systems, ZFS recognizes different media types and will optimize how it handles each type to maximize system throughput. For example, ZFS can take advantage of the performance characteristics of high speed SAS drives when they are present. Now ZFS can also leverage SSDs, where available, enabling even more significant performance gains.

"OpenSolaris on an Open Storage platform like the Sun Storage 7000 Unified Storage System is a powerful framework for us, because it liberates us to get the performance and flexibility we need not just the ones we can afford. The FishWorks appliance software expands this freedom to everyone who doesn't have the time or inclination to build an array themselves. Windows, Linux, even OS X—with FishWorks everyone's invited to the Open Storage party."
— **Jason Williams, COO and CTO, DigiTar**

Hybrid storage pools

As shown in Figure 3, ZFS manages DRAM, SSDs, and HHDs in a hybrid storage pool that delivers higher performance at up to 75% less cost than traditional storage architectures. In these hybrid storage pools, SSDs are used as a read cache for actively used data, making the entire storage infrastructure appear as fast as flash storage to applications.

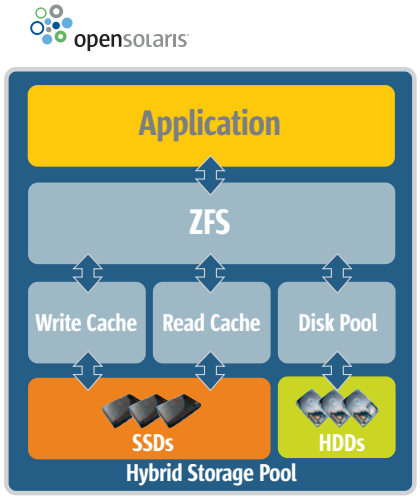

Figure 3. ZFS transparently optimizes data placement across hard disk drives and SSDs to deliver improved throughput.

ZFS transparently executes writes to the pool of low latency SSD media so that writes can be quickly acknowledged, allowing the application to continue processing. The data is then automatically flushed to hard disk drives as a background task performed by ZFS. Another pool of SSD media acts as a read cache and ZFS also manages the process of copying frequently accessed data into this read cache where it can be retrieved with very low latency.

Applications can thus be completely isolated from slower mechanical disk drives, unlocking new levels of performance and higher ROI. This hybrid storage pool approach provides the benefits of high performance SSDs while still saving money with low-cost high-capacity disk drives.

Radically simple storage

The Sun Storage 7000 series also provides unmatched simplicity and ease-of-use through innovative new storage software that includes a browser-based management interface and unprecedented business analytics functionality. A rich set of data services also helps simplify administration and improve service levels.

Easy to deploy

Provisioning and management is dramatically simplified through the easy-to-use management interface that takes the guess work out of system installation, configuration, and tuning (Figure 4). The browser-based interface provides an intuitive environment for administration tasks, visualizing concepts, and analyzing performance data. Systems can literally be set up and configured in minutes using a click and drag installation wizard that does not require a highly trained system administrator, thus saving organizations valuable time.

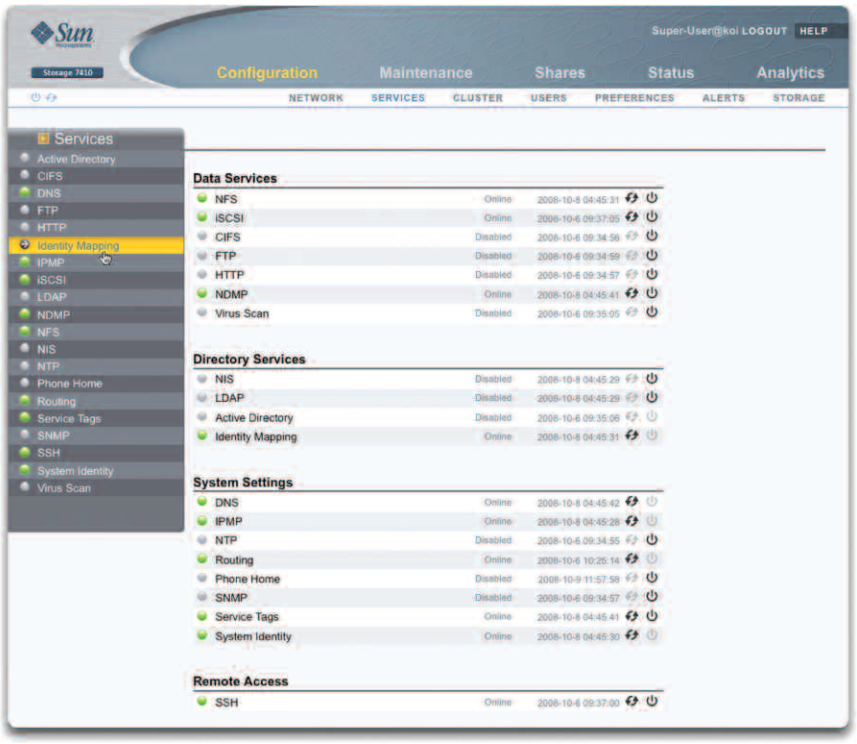

Figure 4. Sun Storage 7000 Storage Manager Software

Easy to scale

Unlike traditional storage architectures, Sun Storage 7000 systems can scale in multiple dimensions to meet specific application needs. As application requirements change, IT managers can choose to:

- Increase computational power by adding more CPUs and cache
- Expand total capacity by adding enterprise class disk drive expansion units
- Increase performance using additional SSD to cache storage reads and/or writes

Increasing system throughput in Sun Storage 7000 systems is as easy as adding SSDs. Since ZFS is designed to dynamically recognize and add new drives, the SSDs can be configured as a cache disk without dismounting the file system that is in use. Once this is done, ZFS automatically optimizes the file system to use the SSDs as high speed disks.

Easy to observe, analyze, and optimize

DTrace Analytics software, which is based on the award winning DTrace facility in the Solaris OS, provides the industry's only comprehensive and intuitive business analytics environment for storage systems. Administrators can drill down for in-depth analysis of key storage subsystems using built-in instrumentation that provides real time visibility throughout the data path. Real time statistical graphs can be used to quickly locate and isolate problems as well as optimize storage performance and capacity utilization—all while systems continue running in production.

Statistics may also be archived to disk so performance data can be reviewed and analyzed throughout the life of the system. Dedicated archive capacity allows the administrator to preserve historical data for up to seven to ten years. Archived data sets can also be suspended or resumed as desired, and can be destroyed if no longer needed.

DTrace Analytics gives administrators all of the tools they need to quickly identify and diagnose system performance issues, perform capacity planning, and debug live storage and networking problems before they become challenging for the entire network.

Most comprehensive self-healing storage system

Sun Storage 7000 systems take advantage of Sun's Fault Management Architecture (FMA) to automatically and silently detect and diagnose underlying problems using an extensible set of agents. When a faulty component is discovered, the self-healing system automatically responds by taking the faulty component offline. Automatic data integrity checking and correction is also provided by ZFS. If its block level checksums identify a corrupt block, ZFS automatically repairs it, providing self-healing capabilities at the block level.

The systems also include easy-to-understand diagnostic messages. The diagnostic messages are linked to articles in Sun's knowledge base to clearly guide administrators through corrective tasks when human intervention is required.

Business value of hybrid storage pools

The following scenarios illustrate how hybrid storage pools can help address some of the most pressing business issues facing today's organizations.

Scenario 1 – Reducing total cost of ownership

A customer requires 46 TB of storage for a high performance file system and is concerned about total cost of ownership as well as data center space. In comparing a popular solution against the Sun Storage 7000 series, the customer finds that the the Sun Storage 7410 system not only offers significantly lower costs, but also provides these key advantages:

- No additional software license fees compared with more than $50,000 in software licenses required with the alternative system.
- Delivers equivalent throughput and better performance without expensive and power-hungry 15K RPM drives
- Provides the same 46 TB capacity in less than one-third the space with its densely packaged 12 RU system

When considering total cost of ownership, the differences were staggering:

- Savings of 75% in total cost of ownership over three years
- Power and cooling costs that are 2.5 times less, resulting in annual savings of more than $6,800
- More than 60% lower annual operating costs for power, space and cooling
- 75% lower initial purchase price for the system and required software licenses

Scenario 2 – Enabling business expansion without more datacenter space

A large online service provider hosts private-branded Web-based reservation systems. Each branded system serves millions of customers and has stringent service level requirements. With traditional storage technologies, meeting the end user response times as outlined in service level agreements requires a storage array with 250 15K SAS drives for each branded system. Unfortunately, the provider has filled their existing data center and is out of space, power, and cooling. Growing the business would therefore require building a brand new and costly datacenter.

Hybrid Storage Pools can provide a way to address these restrictions, allowing the service provider to expand it storage and meet service level agreements with significantly fewer disk drives. Thus the hybrid storage pools can help the service provider grow its business while avoiding the cost of building a brand new datacenter.

Sun Storage 7000 series configurations

To meet a variety of customer needs for capacity, price and performance, the Sun Storage 7000 series family comes in three different configurations plus a cluster configuration that offers maximum availability.

All non-clustered systems come bundled with the same software including data protocols, replication, compression, and DTrace Analytics software for system troubleshooting and performance optimization. Clustered systems include an additional software module for cluster software features.

Sun Storage 7110 system

This entry level easy-to-install storage appliance is well suited for small to medium size businesses and departments as well as remote offices of large corporations. It uses the same software as the high-end configurations, and delivers a 2 TB of raw capacity in a 2U package using 10K RPM SAS drives. With the Sun Storage 7110 system, customers can acquire easy-to-use enterprise data management functionality at entry level costs.

Sun Storage 7210 system

The Sun Storage 7210 system is a mid-range configuration that radically simplifies enterprise storage management and offers extremely high density with up to 44 TB of raw capacity in a 4U system. This system includes write-optimized SSD, so it is ideal for organizations that require high volume throughput for write operations and yet do not require the extensive capacity available in the high-end Sun Storage 7410 system. Such customers can conserve rack space and achieve additional savings in energy consumption by deploying their application using the Sun Storage 7210 system. The system utilizes 7200 RPM hard disk drives to achieve high energy efficiency.

Sun Storage 7410 system

This high-end unified storage system offers radically simplified storage management and up to 576 TB of raw capacity for extreme scalability. The Sun Storage 7410 system also delivers superior performance while reducing energy consumption with its inventive hybrid storage pool architecture.

This system is ideal for organizations with demanding performance requirements and where rapid growth in file-based information is expected. It enables customers to start small (10 TB) and then increase storage capacity, computational power, or read/write storage cache to meet their changing business needs. Like the Sun Storage 7210 system, this configuration also utilizes 7200 RPM hard disk drives for high energy efficiency. With the Sun Storage 7410 system, customers achieve maximum scalability and performance combined with conservation of power, space, and cooling.

For customers that require maximum protection against downtime, the Sun Storage 7410 system also supports a two-node cluster configuration. The Sun Storage 7410 Cluster System features an active-active architecture with no single point of failure, enabling high performance and high availability to maximize business productivity.

Table 1. Configuration information for Sun Storage 7000 series systems

	Key Customer Requirement	Max Storage Capacity	Space (Rack Units)	Write Optimized SSD	Read Optimized SSD	Cluster Option
Sun Storage 7110 system	Low priced entry level system with innovative storage software	2 TB (16 x 2.5" SAS disks)	2U	N	N	N
Sun Storage 7210 system	Mid-range scalability and performance in a compact energy-efficient 4U system	44 TB (48 x 3.5" hard disks)	4U	Up to 36 GB	N	N
Sun Storage 7410 system	Best price/ performance and maximum growth path	576 TB[3] (576 x 3.5" hard disks)	6U for single system, or 8U for cluster configuration[4]	Up to 16 SSDs for cluster configuration	Up to 6 per node (600 GB total)	Y

Get started on the right track

Sun makes it even easier to deploy its Sun Storage 7000 systems with new professional services that can help mitigate the risk of downtime, data loss, and costly delays. Some of Sun's service offerings that are specific to these new systems include:

- **Advanced Installation services for the Sun Storage 7000 family** – Sun's Advanced Installation service for Sun's Unified Storage systems goes beyond basic installation to include configuration into your storage operational environment, configuration with directory services, setup anti-virus scanning service and configuration, and setup of fail-over configuration. Sun provides you with the installation assistance and expertise required to achieve a smooth, successful start-up and to ensure that your systems are running at an optimal level right from the start.
- **Sun Unified Storage Data Migration service** – Minimize disruption and gain full feature functionality of new technology investments when migrating from other platforms. Customers can achieve faster ROI by leveraging Sun's deep technical expertise and proven implementation and migration tools.
- **SunSpectrum℠ Service Plan** – Sun Storage 7000 systems are covered under a single service contract that provides the features of both Sun software and Sun StorageTek™ service plans. The unified service plan can help simplify and reduce the cost of managing an IT infrastructure while delivering greater ROI.
- **Sun Storage 7000 System Administration** – These courses are designed to help administrators gain a deep knowledge of Sun Storage 7000 Systems through knowledge transfer that will enable staff to be more productive.

[3] The initial release offers 288 TB in usable capacity and a free software upgrade will be available in the near future to enable the system's full 576 TB of capacity to be utilized.

[4] Single system includes 2U system plus a 4U expansion array. Cluster system includes two 2U systems plus a 4U expansion array.

Chapter 4

Conclusion

The promise of open storage is freedom from vendor lock-in with a global community sharing a passion to make storage better. Sun's approach to open storage leverages Sun's many years of investment in enterprise reliability and scalability to deliver cost-effective solutions that address some of today's most challenging business issues.

Sun Storage 7000 Unified Storage Systems bring a new level of simplicity and price/performance to the Open Storage market by integrating innovative storage software with hybrid storage pools and an open architecture solution. Customers can benefit by using these new systems to:

- Radically simplify storage deployment and management
- Dramatically reduce the cost of high performance storage
- Increase service levels using DTrace Analytics software to quickly diagnose and troubleshoot storage performance or availability issues

For more information

For additional information on Sun Storage 7000 systems as well as related technologies and services, visit the Web sites listed in Table 2 below or contact a local Sun representative.

Table 2. Web links for additional information

Web Site URL	Description
sun.com/unifiedstorage/	Sun Storage 7000 Unified Storage Systems
sun.com/openstorage	Sun Open Storage products and solutions
opensolaris.com/	OpenSolaris OS
sun.com/storage	Sun Storage solutions
sun.com/service/openstorage/	Sun Services for Open Storage
sun.com/storage/white-papers/	Sun storage white papers

Sun Microsystems, Inc. 4150 Network Circle, Santa Clara, CA 95054 USA **Phone** 1-650-960-1300 or 1-800-555-9SUN (9786) **Web** sun.com

SunWIN #547616 Lit. #SYWP14668-0 11/08