

# WHAT IS OPEN STORAGE?

**Sun Systems: Storage**  
White Paper

**Contributing Authors:**

Graham Lovell, Taylor Allis, Bruce Norikane and Chris Ilg  
Sun Microsystems  
May 2008

**Abstract**

Using industry-leading server and operating system technology, Sun has opened the closed world of disk systems to innovation and market choice. By delivering industry-standard hardware and releasing critical storage software to open source via the OpenSolaris™ project, Sun has broken the final barriers to building state-of-the-art storage systems with truly open architectures.

## Table of Contents

<b>Executive Summary</b> .....	<b>3</b>
<b>The Open Storage Market</b> .....	<b>4</b>
<b>Open Storage Concepts and Definitions</b> .....	<b>5</b>
Open Versus Closed Systems .....	5
Open Source .....	5
Open Storage .....	5
Sun Open Storage Platform .....	6
Sun Open Storage Software .....	6
<b>The Evolution of Storage Systems</b> .....	<b>7</b>
Single Large Expensive Disks (SLEDs) .....	7
Enterprise RAID .....	7
Modular Storage .....	8
Redundant Array of Independent Nodes (RAIN) and Highly Scalable Storage .....	8
Open Storage Systems .....	9
Sun Open Storage Versus Traditional Storage .....	9
<b>Sun Open Storage Value Proposition</b> .....	<b>11</b>
Enterprise-Class Quality .....	11
Open Storage Software Community .....	12
Breakthrough Economics .....	13
Dynamic Scalability .....	16
<b>Sun Open Storage Offerings</b> .....	<b>17</b>
Sun Open Storage Servers .....	17
Sun Open Storage Archive .....	18
Sun Open Storage HPC .....	19
Sun Open Storage Software .....	20
<b>Conclusion</b> .....	<b>22</b>

## Chapter 1

# Executive Summary

Today's IT industry consists largely of closed, proprietary, and expensive storage systems. Enterprise disk arrays are generally closed systems with components available only from a single vendor. Customers cannot select the best disks, software, or controller hardware from the open market to use within their systems — they are locked into buying the vendor's limited and expensive components. Even as storage vendors start to utilize commodity-priced components such as disk drives, software, processors, and servers within the design of their enterprise storage systems, they still choose to limit customer choice and to charge premium prices for the entire system — even for minor upgrades.

Years ago, servers were large, proprietary, and expensive. When smaller, industry-standard servers became more powerful and were able to scale, many of the proprietary computer architectures were unable to achieve the same economics and were quickly replaced for scale-out applications. Today's server vendors, in fierce competition, drive rapid technical innovation at market-driven prices. Customers reap the benefits with the best technology at the best economics. This same closed-to-open market shift is now happening in storage.

Sun is leading the open storage revolution by combining open-source software with industry-standard system components to reduce storage costs by up to 90 percent.

Using industry-leading server and operating system technology, Sun has opened the closed world of disk systems to innovation and market-driven pricing. By delivering industry-standard hardware and releasing critical storage software to open source via the OpenSolaris project, Sun has broken the final barriers to building state-of-the-art storage systems with truly open architectures. This enables customers to develop more innovative enterprise systems faster and at a lower cost — without an army of engineers.

## Chapter 2

# The Open Storage Market

Sun believes that open storage products and services will represent just under 12 percent of the total storage market in 2011.<sup>1</sup> With IDC estimating that the total storage market (hardware, software, and services) will generate approximately \$90 billion in 2011, the open storage portion is just over \$10.6 billion. This includes storage hardware, storage software, and storage services. It is also important to note that in open storage, unit volume increases, which means that \$10.6 billion in revenue buys much more storage for the money.

Sun believes that the earliest adoption and fastest growing hardware segments for open storage will appear in several different markets, including NAS, unified storage (file, iSCSI, and FC storage), and object-based storage.

What's driving open storage adoption? The rapid growth of new digital data demands new storage architectures that offer more flexibility and better storage economics. Web 2.0 applications are growing at a tremendous rate and require more scalable and affordable storage. Customers can no longer tolerate the higher cost per GB of proprietary storage in areas more suited to open storage technologies. Additionally, there are specific market segments or trends that benefit from new, open storage architectures:

- Eco-responsible IT requirements can leverage the economic and consolidation advantages of open storage
- High-performance computing (HPC) environments require open storage architectures to efficiently manage storage I/O bandwidth and latency needs
- Virtualized server environments can also leverage the flexibility and consolidation advantages of open storage

For more information on market growth, adoption, and drivers, see the Sun white paper *Open Storage Adoption*.

1. Sun Storage Intelligence Model, 2008.

## Chapter 3

# Open Storage Concepts and Definitions

Open storage is an emerging market, so it is important to define what open storage is and what it is not. The following sections describe how Sun defines open storage as well as Sun Open Storage offerings.

## Open Versus Closed Systems

When applied to computer systems, the word “open” is used in two related terms: open systems and open source. “Open systems” generally refers to computer systems built with industry-standard hardware and open operating systems. “Open source” is described in more detail below.

The term “closed” applies to systems in which the components of a computer or storage system are proprietary. The vendor might have policies that prevent customers from selecting accessories or options from third parties. This vendor lock-in enables the vendor to charge premium prices for products.

## Open Source

Open source refers to software that has freely available source code and is free from license fees. Open-source software has been part of computing from the earliest days but became a market force in the 1990s, roughly paralleling the rise of the Web.

Open source has many advantages for customers and for developers. Developers have a solid base upon which to build new systems, a community to help solve problems, and a license to use and distribute the open-source platform. Market segments with a strong open-source community burst with innovation. The incredible growth of the Web market is built on the open-source LAMP (Linux, Apache, MySQL, Perl/Python/PHP) and SAMP (The Solaris™ OS, Apache, MySQL, Perl/Python/PHP) platforms.

## Open Storage

Open storage refers to storage systems built with an open architecture using industry-standard hardware and open-source software. In an open architecture, customers can select the best hardware and software components for their individual needs.

In contrast, almost all modern disk arrays and NAS storage offerings are closed systems. Customers are locked into using the vendor’s disk drives, controllers, and proprietary software features.

## Sun Open Storage Platform

This refers to Sun's portfolio of industry-standard hardware and OpenSolaris technology. Customers can combine open storage applications with standard server and storage hardware to greatly increase system flexibility and scalability while reducing the cost of storage. (See Sun Open Storage Offerings section.)

## Sun Open Storage Software

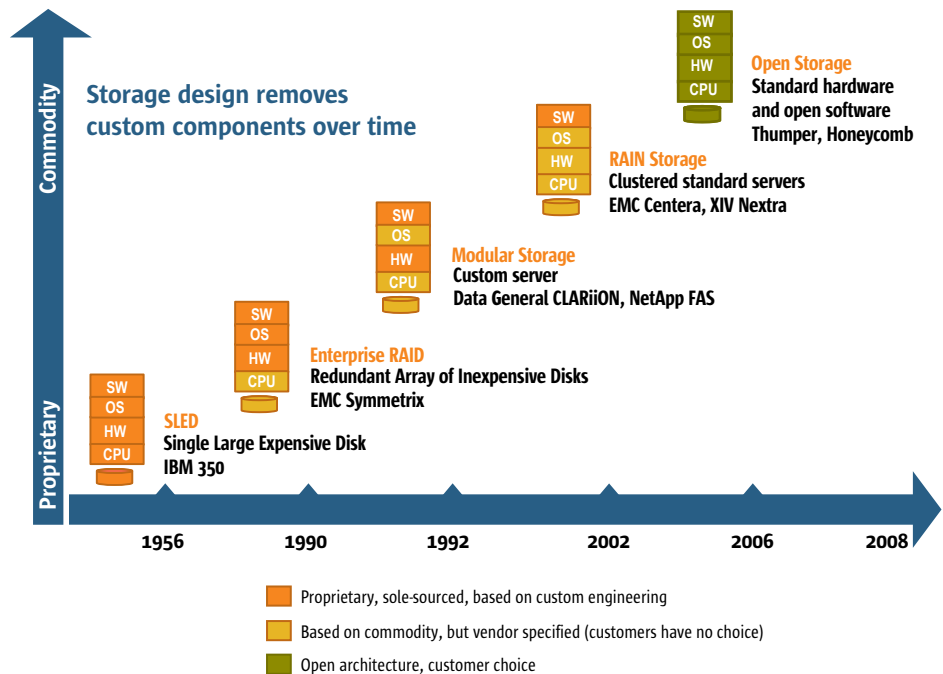
This refers to Sun's open-source software developed through the OpenSolaris storage community initiative, including the Solaris ZFS file system, Sun StorageTek™ Archive Manager, and Sun StorageTek™ Availability Suite software. (See Sun Open Storage Offerings section.)

## Chapter 4

## The Evolution of Storage Systems

Early disk systems were custom engineering projects with few standard components. Over the years, storage vendors have incorporated many standard components into their proprietary architectures. However, every part must be purchased from the system vendor. Almost all storage systems have been completely closed — until now. The following chart shows the evolution of disk systems from closed, proprietary systems starting 52 years ago to today's first true open storage offerings:

Figure 1. Evolution of Disk Architecture



Source: Wikipedia.

## Single Large Expensive Disks (SLEDs)

The first disk system, or SLED, was developed by IBM in 1956. It was entirely proprietary technology including custom controllers, custom magnetic heads, and even custom magnetic coating on the platters.

## Enterprise RAID

During the 1980s, a high-volume disk market had developed for personal computers and open-systems servers. These drives were smaller and much less expensive than the proprietary SLEDs and were built with industry-standard interfaces such as ATA and SCSI. Optimized for the desktop market, these drives had lower capacity, reliability, and performance than SLEDs.

In the early 1990s, several companies released enterprise disk systems called disk arrays, or RAID, that incorporated a large number of small, desktop drives. Enterprise RAID systems consisted of a monolithic enclosure that held a custom disk controller and hundreds of industry-standard drives.

Exploiting the volume market for desktop drives gave RAID vendors several advantages. The most dramatic were lower cost and better storage economics. Almost as important was the high-volume marketplace. The RAID vendors could ride the technology investments of the drive market, which yielded 60 percent annual improvements in disk capacity at lower costs.

## Modular Storage

The original monolithic RAID systems were designed for the mainframe market, which required reliability and performance but could tolerate higher costs. Although the monolithic RAID systems were extended to accommodate the open systems market, their basic design was too big and expensive for open system storage.

In the mid-1990s, midrange computer companies developed smaller, modular RAID systems. Unlike the monolithic RAID systems, modular RAID was composed of a separate controller and disk enclosures that fit into a standard computer rack. Modular RAID systems typically started at 10 percent of the lowest price of a monolithic RAID.

## Redundant Array of Independent Nodes (RAIN) and Highly Scalable Storage

About a decade after the first dual-controller modular systems, distributed storage architectures based on server technologies started to appear. The distributed storage field is still emerging, and there are many variants, but most are based on standard servers with disks, networks, and operating systems. Vendors are adding a software layer to connect multiple servers into a seamless storage system. Distributed storage systems enable customers to scale horizontally by increasing storage performance and capacity by adding nodes.

One distributed architecture in the storage market is called RAIN, and RAIN systems scale by adding a node that includes a server with embedded drives, so capacity and performance scale simultaneously. RAIN systems include clustering software that creates a unified pool of storage by linking many nodes.

RAIN and similar distributed storage architectures exploit the rapid advances in server technology. By concentrating development on the software layers, RAIN systems gain performance and reduce cost at the rate of volume servers. This creates a dramatic advantage over closed storage architectures.

However, although RAIN storage systems are based on standard servers and operating systems, they tend to be closed systems with locked-in components from traditional vendors. Customers typically cannot select disks on the open market, and vendors block any addition or alteration of software. All the disks and software for Centera and NEXTRA must come from EMC and IBM respectively, and system costs reflect the lack of competition.

## Open Storage Systems

Open storage systems combine industry-standard hardware with open-source software. This powerful combination helps spur innovation and drives better storage economics. Developers can leverage both volume servers and an open-source storage stack.

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.

## Sun Open Storage Versus Traditional Storage

The following is a selection of products in the enterprise storage market that compares vendor products, first ship dates, and storage system components. Red items are system components that are closed, proprietary, and available through a single vendor. Yellow items are built from components that are market-priced yet sold as part of a closed, proprietary system. Green items are truly open components. They are industry-standard or open-source software, giving customers choice over how their systems are built.

Storage software is the last major storage infrastructure software to become open source, and Sun's OpenSolaris community is leading the change. Of most importance, since 1992, storage systems have started to utilize open-source, commodity, or industry-standard technologies. OpenSolaris, FreeBSD, and Linux are used in several systems from vendors such as Nexenta, EMC, NetApp, and IBM. However, these open technologies are deployed in a closed-system approach at premium prices — keeping reduced costs from becoming savings for storage users. Industry-standard servers based on x86 processors and commodity disk drives are also largely included in closed, proprietary storage systems today at a premium price, again keeping reduced costs from becoming savings for customers. This in a time in which a new storage economic model is desperately needed.

Figure 2. Storage Architecture Comparisons

Product	Vendor	First Ship	Storage Software	OS	Controller Hardware	Processor	Disk Tray	Disk
<b>Open Storage</b>	<b>Sun</b>	<b>2008</b>	<b>Open Storage</b>	<b>Open Solaris</b>	<b>Server</b>	<b>Intel® Xeon®/AMD® Opteron®/SPARC® processor</b>	<b>JBOD</b>	<b>SAS/SATA</b>
StorageTek™ 6000 array	Sun	1995	Santricity	LSI	LSI, x86	Intel Xeon Processor	Sun	FC/SATA
Symmetrix	EMC	1991	Enginuity	Enginuity	DMX	DMX	DMX	FC/SATA
CLARiiON	EMC	1992	FLARE	Windows	CX3/x86	Intel Xeon Processor	CX3	FC/SATA
Centera	EMC	2000	CentraStar	Linux	x86	Intel Xeon Processor	Internal	SATA
CDL	EMC	2001	FalconStor	Linux	x86	Intel Xeon Processor	CX3	SATA
FAS	NetApp	1994	OnTap	BSD	x86+	Intel Xeon Processor	FAS	FC/SATA
DS8000	IBM	1996	IBM	AIX	P-Series+	PowerPC	IBM	FC/SATA
SVC	IBM	2003	IBM	Linux	x86	Intel Xeon Processor	None	None
USP	Hitachi	2004	BOS	BOS	USP	USP	USP	FC/SATA
Nextra	IBM	2006	XIV	Linux	x86	Intel Xeon Processor	Internal	SATA

■ Proprietary, sole-sourced, based on custom engineering  
■ Based on commodity, but vendor specified (customers have no choice)  
■ Open architecture, customer choice

As shown in the previous figure, Sun is leading the storage industry in delivering truly open storage systems and architectures. But what is the true value to storage customers and the storage community? Open storage's value proposition to enterprise storage buyers and storage developers is explored further.

## Chapter 5

# Sun Open Storage Value Proposition

The market shift to open storage systems and Sun's open storage approach offer four distinct advantages:

1. Freedom from vendor lock-in with a global community sharing a passion to make storage better
2. Breakthrough economics that can save storage users up to 90 percent over closed, proprietary storage
3. Enterprise-class quality storage systems that are reliable and offer some of the highest data integrity in the industry
4. The ability to easily and dynamically scale IT infrastructures to meet changing business needs

## Enterprise-Class Quality

One of the most significant advancements in the storage industry is that enterprise-quality storage systems can now be developed from industry-standard hardware and open-source software.

The Sun StorageTek™ 5800 system is designed for fixed-content archives and built from industry-standard servers, storage, and open-source software. In a recent *InfoWorld* product review, the Sun StorageTek 5800 system scored a 9.3 out of 10, with perfect 10s in reliability and scalability.<sup>2</sup> According to *InfoWorld's* senior analyst, Mario Apicella:

Sun has taken a different approach to companion software than vendors such as EMC, Hitachi, and HP, which have married their fixed-content archiving solutions to compliance applications.... Sun has not wedded Honeycomb [the Sun StorageTek 5800 system] to any specific application, leaving that task to partners and customers. The upside of Honeycomb's [the Sun StorageTek 5800 system's] openness is that the possibilities are endless. In fact, Honeycomb's [the Sun StorageTek 5800 system's] powerful, built-in administrative software is complemented by an SDK that allows Java™ or C developers to define their own metadata schemas consistent with the specifics of their application.

Conventional NAS simply isn't designed for long-term archiving. The typical NAS would choke under the load of storing multiple large objects at the same time, and it would die with its third consecutive drive failure. Honeycomb [the Sun StorageTek 5800 system] addresses the performance and resilience requirements of content archiving with a new architecture. Unlike plain NAS solutions — and fixed-content archiving solutions built on conventional storage systems (think EMC Centera) — it's made for the job.

2. "Product review: Sun's StorageTek™ Honeycomb is sticky and sweet," Mario Apicella, *InfoWorld*, 3/24/08.

Storage advancements such as the Sun StorageTek 5800 system demonstrate that industry-standard hardware coupled with open-source software can be used to build enterprise systems that offer the same level, if not a higher level, of quality compared to closed systems.

Commercial distributions of Sun open-source software have gone through rigorous testing and certification as well. For example, Solaris ZFS offers “19 nines” of data integrity, while Predictive Self-Healing offers automatic system diagnosis. The Solaris OS is the only known operating system designed to provide end-to-end checksumming for all data. Solaris ZFS constantly reads and checks data to help ensure it is correct, and if Solaris ZFS detects an error in a mirrored pool, the technology can automatically repair the corrupt data.

Unlike HP-UX, Red Hat Linux, and Microsoft Windows, the Solaris OS has the capability to monitor historical failure data to predict faults and take proactive, preventive action to avoid future failures. The Solaris OS has the ability to track historical errors, perform root-cause analysis, and shut down any components dependent on the failed service or component. This capability is unmatched by HP-UX, Red Hat Linux, and Microsoft Windows.

Open-source software with a vibrant community has typically matched or surpassed its closed counterparts in quality and reliability.

## Open Storage Software Community

More storage companies are hiring developers or leveraging development resources, in order to include the functionality, innovation, customization, and time-to-market advantages that the open-source community provides.

Sun’s OpenSolaris project has more than 96,000 registered members.<sup>3</sup> On April 10, 2007, Sun announced the donation of storage technologies for storage developers within the OpenSolaris project.<sup>4</sup> More than 60,000 people have viewed or participated in the OpenSolaris storage community initiative, which started with a couple of open-source storage projects and now tracks more than 30 projects. Community members share their expertise to find answers more quickly and then take an active role in the evolution of features and source code.

Through the OpenSolaris storage community, developers have access to a complete storage stack — from client-side file, object, and block access technologies to complete server-side file and block services, data services, storage management, and extensive storage connectivity components. The community also fuels innovation: Nexenta Systems and Prominic.net have created new industry NAS appliances through the OpenSolaris project. Open-source software enables developers to quickly customize storage applications as needed, with continual support from the community, Sun, and other developer services.

3. <http://www.opensolaris.org/os/community/marketing/metrics/latest/>.

4. <http://www.sun.com/aboutsun/pr/2007-04/sunflash.20070410.1.xml>.

Unlike IBM, Sun offers indemnification against patent infringement claims, to protect customers and developers as well as continuing to invest in community code development after code contribution.


Open source also ensures that OpenSolaris will continue to support multiple systems and storage platforms. The Solaris OS is currently compatible with more than 1,000 systems (917 x86/x64 systems and 83 SPARC technology-based systems) and 1,612 various components.

Not only is the Solaris OS open, but so are critical tools such as Sun Java™ Studio and NetBeans™ software. Unlike with other operating systems, the Solaris OS' developer tools are integrated into the OS distribution, resulting in lower development costs and less complexity.

## Breakthrough Economics

Open storage changes the economics of IT as well. Consider the following general diagram of a closed storage architecture:

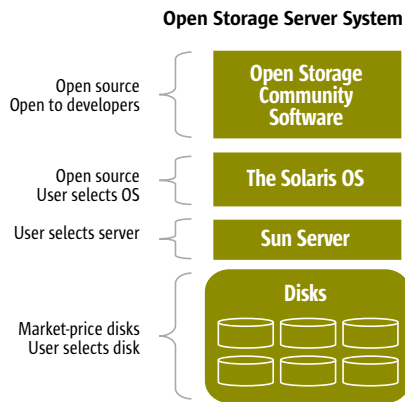
Figure 3. A Closed Storage Architecture



Storage software is closed and is typically three times the cost. Even if the operating system is based on open source code, customers are typically paying for the operating system from the vendor as part of the system. Customers must also purchase a storage controller and controller software from the vendor. Lastly, customers must buy the disk embedded in the storage system, which in most cases is a commodity disk marked up to five times the original cost by the proprietary storage vendor.

Consider an open storage architecture:

Figure 4. An Open Storage Architecture

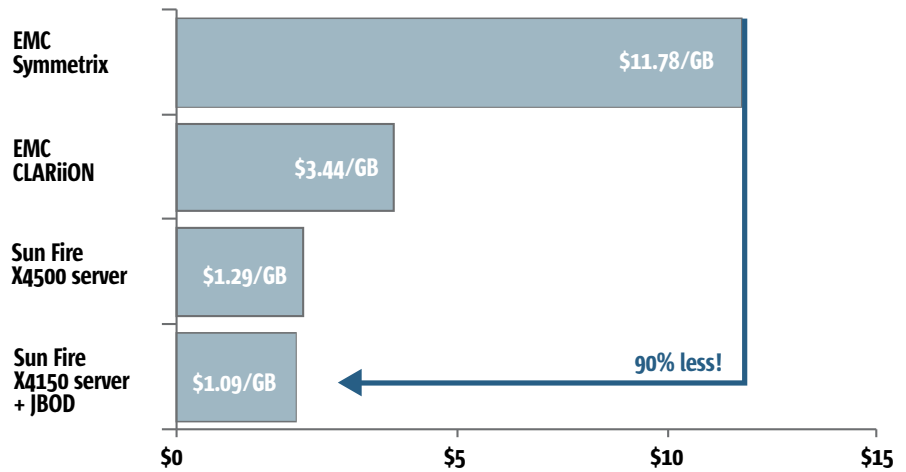


In an open storage architecture, the operating system is available as open-source, the application software is open-source, and the hardware is industry-standard. Being able to freely select the storage software is useful to enterprises that have IT staff more familiar with a specific operating system. In place of an expensive, proprietary disk controller, users can now leverage an industry-standard server. This is possible through innovations such as Solaris ZFS, which freely provides data services such as RAID, error correction, and system management — traditionally limited to expensive controllers. Lastly, market price disks can be used without the markup. Users can purchase high-performance, affordable disks systems — even JBOD — to increase performance and save cost by leveraging market-price disks with Solaris ZFS.

The economic difference between an open storage architecture and a closed storage architecture is staggering. The following diagram compares EMC's closed storage to two Sun Open Storage systems: the Sun Fire™ X4500 server and a Sun Fire™ X4150 server with JBOD disk attachment.

This model uses 2008 industry U.S. list pricing from independent pricing analyst Ideas International:

Figure 5. Closed Versus Open Storage \$/GB Comparison

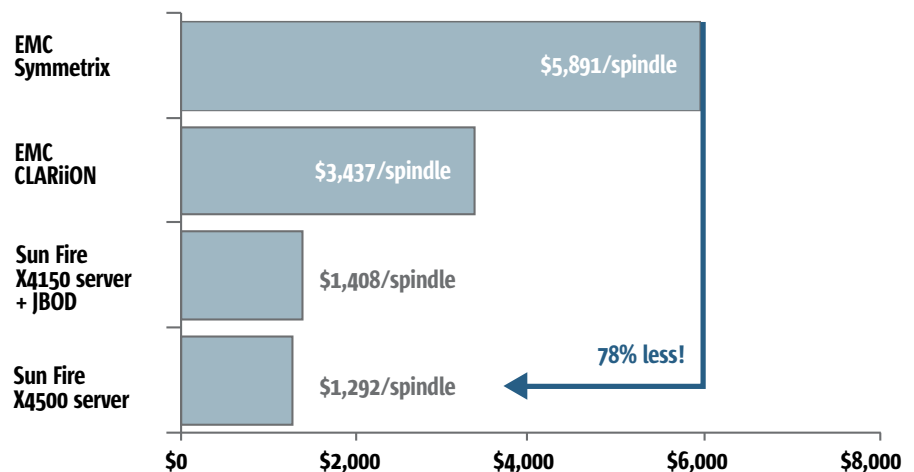


Sun’s open storage architecture offers 90 percent less \$/GB storage compared to a closed storage architecture.

Additionally, according to Ideas International, EMC CLARiiON CX3-20 replication software costs \$86,100 list. OpenSolaris technology is free to use without any licensing fees.

Another economic metric to evaluate in storage systems is the price per disk drive or spindle in any given system. Again, using Ideas International U.S. list pricing, closed and open storage systems are compared:

Figure 6. Closed Versus Open Storage \$/Disk Spindle Comparison



Again, open storage offers economic benefits close to 80 percent better than closed systems, when evaluating how much a customer is paying for disk drives or spindles in an enterprise storage system.

Industry-standard hardware and open-source software offer radically lower price points and breakthrough economics compared to closed storage. Additionally, vendor lock-in is avoided with open storage systems.

It is important to note that there are certain services and features that are available to run on proprietary higher-end disks today. But at \$1.09/GB and \$1,292/spindle, customers will be compelled to design areas of their storage infrastructure to suit open storage architectures where they can take advantage of the breakthrough economics.

## Dynamic Scalability

Web-scale applications, HPC storage, and traditional storage applications all demand dynamic, massively scalable storage. Open storage offers massive, dynamic scalability at significantly lower costs.

Sun offers the most scalable storage systems in the industry, ranging from less than 10 TB to more than 100 PB. The Solaris OS is the only operating system to deliver both vertical and horizontal multiplatform scalability with the same code base. With Solaris ZFS, users can also quickly add any device to their storage pools. When storage users reach the limit of their current infrastructure, they don't always have to migrate to new storage devices with new management tools. They can dynamically change or upgrade components as needed.

Unlike traditional storage deployments, Sun Open Storage environments enable customers to freely mix, match, and maximize the components of their computing infrastructures as business needs change. Customers can also repurpose and reuse hardware simply by adding new software to an open platform.

Lastly, Solaris ZFS is a 128-bit file system, and it can provide 16 billion times the capacity of 32-bit or even 64-bit file systems. This means that Solaris ZFS can support more file systems, snapshots, and files in a file system than can possibly be created in the foreseeable future.

## Chapter 6

# Sun Open Storage Offerings

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 offers enterprise reliability and scalability at one-tenth the cost of closed, proprietary storage. Sun offers the open-source Solaris ZFS, which delivers incredible data integrity and can significantly reduce downtime. Additionally, Sun's innovative use of industry-standard hardware in systems such as the Sun Fire X4500 server delivers two to three times the density, using 50 percent less in power and cooling than competing closed storage.

Sun Open Storage also empowers developers to create services quickly for multiple platforms. Sun's open-source software is open, secure, and freely available. With the Solaris OS, developers, startups, and Web 2.0 companies can quickly develop highly scalable and secure storage services today — reliably, cost-effectively, and on the broadest set of platforms (UltraSPARC® and x64/x86) of any operating system. Unlike the competition, Sun remains active in the community after contributing code, so users can rely on Sun for expert support and service, even on certified third-party hardware.

Sun currently offers the following industry-standard hardware, software, and systems built from standard components and open-source software:

## Sun Open Storage Servers

**Sun Fire X4500 server:** The Sun Fire X4500 server leverages industry-standard hardware and software in a unique package. It is a Dual-Core AMD Opteron™ processor-based server with 48 hot-swappable SATA drives in a single 4 U chassis that can achieve 48 TB of raw capacity with 1 TB SATA drives. It includes 4 GB NICs and 4 GB of RAM, and ships with the Solaris 10 OS. It can run multiple operating systems including OpenSolaris. The Sun Fire X4500 server recently won *InfoWorld's* 2008 Technology of the Year award for best storage server.

**Sun industry-standard servers:** In the x64 market, Sun offers a complete range of servers from one CPU to eight CPU and with a range of storage capacities. For example, an extensive range of modular blade systems: the Sun Blade™ 6000 and 8000 series, and the Sun Fire X4150 server with two CPUs, eight cores, and eight disks in a compact 1 U form factor, and the Sun Fire™ X4600 M2 server with up to eight CPUs in 4 U.

## Sun Open Storage Archive

Sun StorageTek 5800 system: The Sun StorageTek 5800 system also leverages industry-standard components. The Sun StorageTek 5800 system is the first integrated, fixed-content archiving system built using open-source software.

In a recent *InfoWorld* product review, the Sun StorageTek 5800 system scored a 9.3 out of 10, with perfect 10s in reliability and scalability.<sup>5</sup> According to Mario Apicella, the Sun StorageTek 5800 system provides:

Impressive resilience together with excellent performance and powerful administrative tools make “Honeycomb” [the Sun StorageTek 5800 system] one of the most interesting solutions in the emerging fixed-content archiving space. With a foot in the open source community, Honeycomb [the Sun StorageTek 5800 system] promises to deliver more software features faster than competing proprietary solutions, and customers that can’t wait have an easy and free alternative with a flexible SDK.

Sun StorageTek 5800 Open Edition: This freely downloadable binary of the Sun StorageTek 5800 system software fully implements a digital archive with fast, searchable content and rich metadata that runs on virtually any x86 hardware device. Like the Sun StorageTek 5800 system, it can store and manage large amounts of fixed content (videos, x-rays, digital books). The OpenSolaris project is focused on client and server implementations of this object-oriented storage system, with traditional Java technology and C interfaces to be later expanded by a “StorageBean” Java interface. Large data repository applications access the fixed content through these interfaces, which are designed to manage data collections that can total up to 100 million objects or petabytes of storage.


Another goal of the project is to add the Storage Networking Industry Association’s industry-standard eXtensible Access Method (XAM) specification to both the Sun StorageTek 5800 system and the Solaris OS.

Sun StorageTek Archive Manager and Sun StorageTek QFS software: Sun StorageTek QFS software used with Sun StorageTek Archive Manager provides a shared file system and storage archive management solution for tiered storage solutions in the HPC, data protection, and archive markets.

5. “Product review: Sun’s StorageTek Honeycomb is sticky and sweet,” Mario Apicella, *InfoWorld*, 3/24/08.

## Sun Open Storage HPC

The Sun Constellation System builds on cost-effective, off-the-shelf components and state-of-the-art technologies to deliver an open, petascale architecture. Using a holistic approach that includes servers, software, storage, and services, Sun has created one of the most powerful HPC platforms in the world.



The Sun Constellation System requires less energy to operate than competitive solutions because of its power and cooling efficiencies. Applications can be created quickly, using open tools and interfaces in small environments, and then rapidly deployed to environments capable of providing up to 1.7 petaFLOPS of computing power.

The Sun Constellation System leverages the following open storage components:

- OpenSolaris technology: Offering key HPC functionality, including performance enhancements, system analysis tools, and high-performance file systems such as Solaris ZFS
- Open-source Lustre file system for unmatched scalability
- Open-source Sun StorageTek QFS software for maximum scalability, data management, and throughput
- Sun Open Storage Sun Fire X4500 servers: Delivering almost .5 PB of storage in a single rack, all accessible from the same IB network

The Sun Constellation System also supports other industry-standard and open-software components and interfaces including:

- Linux
- Intel Xeon and/or AMD Opteron processors
- Sun HPC ClusterTools™ software based on Open MPI
- CLI, IPMI, and SNMP protocols
- Fortress programming language

The Sun Constellation System leverages other industry-leading products from Sun as well, including:

- An ultradense 3,456-port InfiniBand switch — with 3,456 ports, the Sun InfiniBand switch helps reduce cost and overcomes the complexity typically associated with large-scale HPC clusters. (A single InfiniBand switch replaces hundreds of discrete switching elements.)
- UltraSPARC T1 processors with CoolThreads™ technology
- Sun StorageTek™ modular library systems
- Sun StorageTek™ tape drives and libraries for data backup and archive
- Factory integration, services, and support

The first customer install of the Sun Constellation System is the Texas Advanced Commuting Center (TACC) Ranger cluster. With more than 500 teraFLOPS of CPU power, the Ranger cluster is one of the most powerful general-purpose computing platforms in the world. (For more information on the TACC implementation, see the Sun white paper *Open Storage Adoption*.)

## Sun Open Storage Software

OpenSolaris technology: OpenSolaris technology is the cornerstone of Sun Open Storage offerings and provides a solid foundation as an open storage platform. The origin of OpenSolaris technology, the Solaris OS, has been in continuous production since September 4, 1991. OpenSolaris technology offers the most complete open-source storage software stack in the industry. Below is a list of current and planned offerings:

- At the storage protocol layer, OpenSolaris technology provides SCSI, iSCSI, iSNS, FC, FCoE, InfiniBand software, RDMA, OSD, SES, and SAS
- At the storage presentation layer, OpenSolaris technology offers Solaris ZFS, UFS, SVM, NFS, Parallel NFS, CIFS, MPxIO, Shared QFS, FUSE, and the Sun StorageTek 5800 system
- At the storage application layer, OpenSolaris technology offers MySQL™ software, Postgres, BerkeleyDB, AVS, SAM-FS, Amanda, and Filebench

OpenSolaris technology provides an end-to-end storage platform and includes these essential features:

### *Solaris ZFS:*

Another cornerstone of Sun's open storage platform is the Solaris ZFS file system. Solaris ZFS can address 256 quadrillion zettabytes of storage and handle a maximum file size of 16 exabytes. Solaris ZFS deploys several storage services including snapshots, point-in-time-copy, volume management, administration, and data integrity features such as copy-on-write and RAID.

Vendors of closed storage appliances typically charge customers extra software licensing fees for data management services such as administration, replication, and

volume management. The Solaris OS with Solaris ZFS moves this functionality to the operating system, simplifying storage management and eliminating layers in the storage stack. In doing this, Solaris ZFS changes the economics of storage. A closed and expensive storage system can now be replaced by a storage server running Solaris ZFS, or a server running Solaris ZFS attached to JBOD.

Solaris ZFS recently won *InfoWorld's* 2008 Technology of the Year award for best file system. In the *InfoWorld* evaluation, the reviewer stated, "Soon after I started working with [Solaris] ZFS (Zettabyte File System), one thing became clear: The file system of the next 10 years will either be [Solaris] ZFS or something extremely similar."<sup>6</sup>

#### *Solaris DTrace:*

Solaris DTrace provides an advanced tracing framework and language that enables users to ask arbitrary diagnostic questions of the storage subsystem, such as "Which user is generating which I/O load?" and "Is the storage subsystem data block size optimized for the application that is using it?" These queries place minimal load on the system and can be used to resolve support issues and increase system efficiency with very little analytical effort.

#### *Solaris Fault Management Architecture:*

Solaris Fault Management Architecture provides automatic monitoring and diagnosis of I/O subsystems and hardware faults and facilitates a simpler and more effective end-to-end experience for system administrators, reducing cost of ownership. This is achieved by isolating and disabling faulty components and then continuing the provision of service through reconfiguration of redundant paths to data, even before an administrator knows there is a problem. The Solaris OS' reconfiguration agents are integrated with other Solaris OS features such as Solaris Zones and Solaris Resource Manager, which provide a consistent administrative experience and are transparent to applications.

#### *Sun StorageTek Availability Suite:*

Sun StorageTek Availability Suite software delivers open-source remote-mirror-copy and point-in-time-copy applications as well as a collection of supporting software and utilities. The remote-mirror-copy and point-in-time-copy software enable volumes and/or their snapshots to be replicated between physically separated servers. Replicated volumes can be used for tape and disk backup, off-host data processing, disaster recovery solutions, content distribution, and other volume-based processing tasks.

#### *Lustre™ file system:*

Lustre™ is Sun's open-source shared disk file system that is generally used for large-scale cluster computing. The Lustre file system is currently used in 15 percent of the top 500 supercomputers in the world, and six of the top 10 supercomputers. Lustre currently supports tens of thousands of nodes, petabytes of data, and billions of files. Development is underway to support one million nodes and trillions of files.

6. "Sun ZFS breaks all the rules," Paul Venezia, *InfoWorld*, 6/7/07.

## Chapter 7

# Conclusion

Open storage is a new concept for the storage industry. Today's storage products consist of closed, proprietary, and expensive storage systems. Sun has evidence that storage users have started to adopt open storage architectures — storage systems that leverage industry-standard hardware and open-source software. The high quality of open storage offerings such as Sun Open Storage, coupled with the innovative power of the open-source community, can drive economic and scalability benefits to those who move to open storage.

Sun is leading the open storage revolution by combining open-source software with industry-standard system components, which can reduce storage costs by up to 90 percent. Sun predicts that the economic and innovation advantages of open storage can rapidly drive customer adoption, representing approximately 12 percent of the total storage market by the year 2011.

