



WHITE PAPER

Reducing Storage Complexity with Sun Unified Storage

Sponsored by: Sun Microsystems

Steven Scully

Richard L. Villars

May 2009

EXECUTIVE SUMMARY

IT organizations of all sizes are struggling to deal with the tremendous growth in data and the complexity of managing it all. The march of digital applications and the data they create continues at a rapid pace, even in today's difficult economic times. IDC forecasts that external disk storage capacity will grow 50% through 2010.

Unified storage has emerged as one way that IT managers are using to reduce the complexity of storage implementations while improving the business flexibility and lowering the total cost of storage. The benefits of unified storage include:

- ☒ Unified block- and file-level storage through SAN, NAS and iSCSI interfaces allowing storage administrators to consolidate multiple workloads onto a single storage platform.
- ☒ Advanced software features for provisioning, data protection and system management, improving utilization rates and simplifying management.
- ☒ Fully integrated hardware and software solutions (often in the form of appliances) that simplify the implementation.

Sun Microsystems, a leader in server and storage technologies, has launched the first in a family of open, unified storage appliances. In addition to providing all the benefits of unified storage, Sun's solution includes two unique features:

- ☒ DTrace Analytics – a new way of observing and understanding what the unified storage system and enterprise network clients are doing and how they are behaving, using real-time graphical analysis. IT administrators can use it to diagnose issues and optimize their application/storage environment.
- ☒ Hybrid Storage Pools – a high-performance architecture that integrates flash-based SSDs as a caching tier with capacity-optimized, enterprise class HDDs for all storage.

With all protocols and storage features included in the cost of the system, Sun claims its new unified platform delivers significant performance gains and cost benefits versus proprietary unified storage solutions.

SITUATION OVERVIEW – CHALLENGING TIMES FOR MANAGING DATA

IT organizations of all sizes are struggling to deal with the tremendous growth in data and the complexity of storage management. They are supporting a growing array of data intensive applications (e.g., business analytics) and content rich applications (e.g., digital images or video). They are also improving application availability and disaster recovery through increased use of disk-based data protection/replication. In this information intensive environment, IT managers need to optimize the capacity and performance of storage systems while working to reduce complexity and lower costs. Today's challenging economic times only bring more focus on these requirements.

IT organizations are also adapting to a major shift happening with the use of disk-based backup and recovery solutions, which provide improved data protection with greater availability and continuity than previous solutions (i.e. tape). While they speed up and simplify some operations, they are not without their own management and operational challenges. These solutions improve application and data availability by enhancing system backup performance and reliability, shorten file restore times and are supported by major backup vendors. However, they add to the amount of data being stored on a daily basis and represent another element to manage.

Data Growth

The march of digital data continues at a rapid pace. Even in difficult economic times, applications are generating tremendous amounts of data. Corporate reorganizations and mergers along with increased governmental regulations all contribute. Through the rest of this decade, IDC forecasts that the amount of external disk storage capacity shipped is still growing around 50% a year.

A variety of applications across a broad set of industries is driving this data growth. The use of high performance computing is expanding in a wide range of industries (e.g., financial services, life sciences, energy, media/entertainment). At the same time, companies are leveraging more collaboration and web 2.0 applications to boost productivity and improve customer support. Databases are at the heart of these applications, and the analysis of these databases is spurring greater use of data warehousing and business intelligence to improve efficiencies and market opportunities for organizations.

In many industries ranging from health care to media/entertainment and in specialized markets such as video surveillance and product life cycle management, a shift to digital content is now beyond the point of no return. These digital transformations are already spurring exponential increases in image data and associated content.

The traditional method for dealing with these data increases has been to move to larger disk drives, larger disk systems, and multiple tiers of storage. Many of these technologies and techniques have been effective (up until now) at storing large amounts of data, but they have not necessarily simplified management of the data or the systems that store it.

Economic Pressures

While the need for storage capacity hasn't slowed significantly, IT budgets are under increasing pressure. In more challenging economic times, organizations cannot afford to ignore the upfront or long-term costs of storage. Many IT departments have the same number of people (or less) to manage more and more storage while also supporting an ever more diverse array of data generating applications. IT organizations are looking for every option to make storage more efficient – less costly to acquire and upgrade, more simple to manage, and less costly to power and cool.

One current example is the growing use of server virtualization. Virtualized servers deliver strong economic benefits and make greater use of consolidated, networked storage to achieve the cost savings from server consolidation. However, they also risk adding considerable management complexity to and can disrupt finely tuned storage networks.

The evolution of storage solutions to address the costs of storage and inefficiencies in the use of that storage have been underway for some time but are now receiving increased attention due to the economy. Networked storage, storage tiers, thin provisioning and storage consolidation are all examples of new storage technologies that address cost and efficiency challenges.

Performance Requirements

Digital applications also continue to change customer expectations (both internal customers and external ones) for application and storage performance. Many organizations operate 24x7, support employees and customers with web-based applications, and have database applications that produce thousands of interactions every second. Recent IDC conversations with senior IT architects and administrators reveal that many IT organizations have storage workloads that need higher performance in terms of I/O and throughput. Efficiently integrating these performance requirements into storage solutions while also meeting cost and operational objectives and capacity growth requirements is a critical requirement.

Storage systems providers have responded by developing solutions that use new or advanced technologies to meet customer requirements. Separate (and expensive) systems based on solid state storage have been available for several years for customers with high performance requirements (and the budget to match). Storage systems incorporate high-RPM hard disk drives and practices like short-stroking to provide greater performance. With the introduction of flash-based solid state drives using standard interfaces, systems providers have a new tool to address performance challenges.

REDUCING STORAGE COMPLEXITY

The conditions of dramatic data growth, greater data availability, strong economic pressures, and continued demands for high performance, along with addressing requirements unique to their business or industry, translate into several new challenges for IT organizations looking to add to or refresh their storage architectures. Some of these additional challenges include:

- ☒ Reducing the complexity of storage solutions, while dealing with storage tiers and advanced capabilities
- ☒ Lowering the overall cost of storage, including initial acquisition costs, ongoing upgrades and associated power and cooling.
- ☒ Efficiently troubleshooting all aspects of storage solutions from both the application and storage sides.
- ☒ Improving storage flexibility and application availability, and therefore business flexibility

One of the most important developments in storage (driven by the need to address these challenges) is the reemergence of unified storage platforms. Unified storage solutions like the Sun Storage 7000 provide IT organizations with a flexible platform that can address various types of storage workloads using block and file data with a variety of network interfaces. Implementing unified storage solutions provide IT organizations with numerous benefits including:

- ☒ The ability to leverage industry-standard hardware to lower product costs while also including advanced technologies to boost performance and reliability
- ☒ The flexibility to store file data with NAS protocols as well as block data through SAN protocols using iSCSI and Fibre Channel interfaces, allowing storage administrators to consolidate multiple workloads onto a single storage platform.
- ☒ The inclusion of a number of advanced software features for provisioning, data protection and system management that improve utilization rates and simplify management.
- ☒ The full integration of hardware and software solutions (often in the form of appliances) that simplify the implementation and administration.

Finding the Right Storage Partner

Reducing the complexity of implementing and managing consolidated storage for virtualized servers and tiered storage is an important goal for the storage industry. In today's challenging environment, IT managers need solutions that deliver advanced storage capabilities with lower costs, simple deployment and no limitations on expansion while still accommodating their specific business requirements.

The remainder of this white paper looks at the steps that one leading IT company, Sun Microsystems, is taking to deliver an innovative family of unified storage solutions that address the ongoing needs of IT managers and storage administrators. The paper will also discuss the challenges that Sun must address to further enhance the effectiveness of its offerings.

SUN UNIFIED STORAGE SOLUTIONS

Sun Microsystems is a leading provider of software, server and storage solutions that is committed to open source technologies. Building on this strategy, it developed an advanced unified storage platform, the Sun Storage 7000.

Sun launched the 7000 family in November 2008 as the first in a family of unified, open storage solutions. The 7000 family currently consists of three products, the 7110, 7210 and 7410, which scale from 2TB to 576 TB of capacity. The 7000 series supports both file (including CIFS, NFS, FTPS/SFTP and HTTP/WebDAV protocols) and block data using iSCSI interfaces (with Fibre Channel and Infiniband coming in the near future).

The Sun Storage 7000 is the industry's first unified storage system to integrate flash-based solid state drives as cache rather than as primary storage (as Tier 0 storage). This new architecture provides a unique approach to addressing both the price and performance requirements of storage systems. Sun claims that flash used as cache is faster than the fastest HDDs by a factor of 10 for write and a factor of 100 for reads when measured in IOPs.

Sun has designed the Storage 7000 family to reduce the complexity of storage and address the challenges faced by storage administrators. Some of the key features include:

Advanced troubleshooting and diagnostics. A very important way that Sun is reducing the complexity of managing storage is by including advanced diagnostics and troubleshooting in the 7000 family. DTrace Analytics enables greater system optimization and provides real time graphical analytics (see Figure 1). DTrace Analytics provides performance visibility to both the application/file system side as well as the storage/capacity side of the 7000 environment. It allows storage administrators to isolate, diagnose and resolve problems rapidly, as well as to do capacity planning and identify potential performance bottlenecks.

FIGURE 1

Sun 7000 Advanced Diagnostics



Source: Sun, 2009

Simplified configuration and management. Sun also sees ease of management as a major benefit for its diverse customer base and has developed a very straightforward management approach to the 7000 family. The 7000 products are delivered as appliances with everything installed. Sun claims that customers can go from the box to provisioned storage in under 5 minutes. Installation wizards and simple GUI interfaces manage all aspects of the system which includes role-based management tasks. The dashboard (see Figure 2) provides real-time status for the hardware and software data services with links to the DTrace Analytics.

FIGURE 2

Sun 7000 Dashboard



Source: Sun, 2009

Hybrid Storage Pools. Sun has developed a unique approach with the architecture of the 7000 family called Hybrid Storage Pools (HSPs). HSPs seamlessly integrate flash-based SSD in the storage hierarchy as a caching tier together with DRAM and enterprise-class, capacity-optimized HDDs. Two types of SSDs are used in the HSP, a write-optimized (called Logzilla) and a read-optimized (called Readzilla).

Logzilla is a flash-based SSD that contains a big DRAM cache backed by a supercapacitor so that the cache can effectively be treated as nonvolatile. Logzilla is used as a ZFS intent log device to greatly improve synchronous write performance. Readzilla flash-based SSD is used as a secondary cache tier between DRAM cache (the ZFS ARC) and disk. This second level adaptive replacement cache (ZFS L2ARC) can be as large as 600GB per node in the Storage 7410 system (1.2 TB in the cluster configuration), allowing huge working sets to be serviced with latencies under 100us.

The Sun 7000 file system (ZFS) optimizes the way data is managed across the HSP, treating it transparently as a single hierarchy. HSPs provide for a number of capabilities and benefits in the 7000 family including:

- ☒ Accelerated and optimized system performance with increased performance for write-oriented workloads (synchronous writes) and increased performance for read-oriented workloads (caching frequently accessed data).
- ☒ The ability to configure the system (depending on the model) with the optimum mix of read cache, write cache and disk pool to meet the storage workload requirements.
- ☒ Lowering system power consumption and associated cooling by using both energy efficient SSDs and the capacity-optimized SATA HDDs.

Depending on the configuration, Sun claims that its HSP approach can provide higher performance at better price points and lower power and cooling costs.

Fully Integrated software. Many of the software features in the 7000 family come from open source software initiatives (though not all of the software is open source). By leveraging the open source communities for timely development, Sun can speed the deployment and simplify the release of new capabilities at lower costs. Sun includes these software capabilities as standard, no extra cost features of the 7000 family. Included with the product are CIFS, NFS, iSCSI, Replication, Snapshot/Clones, Anti-Virus, Compression, Thin-Provisioning and Clustering capabilities. On the roadmap are additional software capabilities such as Deduplication and Encryption. By bundling these capabilities with the system, Sun is lowering acquisition costs and simplifying the management of its storage solutions.

CUSTOMER EXAMPLE

IDC recently interviewed Magnus Befwe, Manager, IT Operations and Infrastructure and Johan Hillström, Analyst, IT Operations and Infrastructure from Elanders, a global infomedia and printing group based in Sweden. They are in the process of implementing three Sun Storage 7410 systems in three different locations - two standalone systems and one clustered.

Elanders provides solutions to customers who need to communicate internally or externally, to one or several markets simultaneously. The Group's primary business lines are Infologistics (publishing solutions for e-commerce and electronic document management, database publishing, digital filing and digital print) and User Manuals (printing and composing of user manuals). They recently celebrated 100 years as a company and have a local presence in ten countries on four continents.

In 2008, Elanders went looking for a replacement for several older disk systems used as the primary storage for their organization and for the data they host on behalf of their customers. They investigated systems from many of the leading storage vendors. Key requirements for the replacement systems included:

- ☒ Significant network bandwidth with iSCSI connectivity
- ☒ RAID 6 or something more than RAID 5 (which has caused problems for them in the past)

☒ Replication and Snapshot capabilities along with a good user interface

☒ The ability to scale easily and with enough capacity to meet their growing needs.

After evaluating all the proposals, Elanders selected the Sun Storage 7410 systems as the best fit for their needs, now and into the future. When asked what made them select the Sun 7410, Mr. Befwe said "Price was obviously a factor. But also the licensing model was good – no hidden or extra costs that hit you down the line. The scalability also drew us – very simple to scale up." Mr. Hillström added "What really caught our eyes was the excellent GUI in the Sun system.

Elanders installed the first system in late January 2009 and recently finished installing the third. As of this paper's publication date, they were still migrating data from their current systems to the Sun 7410s and planned to have VMware fully implemented with the systems shortly. They are planning on using the replication capabilities between the systems to augment their traditional backup processes. Having worked with the Sun systems for several months, some of the features that stand out for them include:

☒ Administration tool is better than other systems they investigated. The Analytics feature they feel is also a great tool. They are using it quite frequently now to analyze network bandwidth and determine which clients are generating the traffic.

☒ Hybrid Storage Pools. They are taking advantage of this capability in all three systems, incorporating the write-optimized SSDs to improve performance. They like not having to worry about optimizing the system. They feel the HSP and the system's "powerful file system" will provide good option going forward.

Being a new technology, Elanders reported that the web interface still had the occasional hiccup. A feature they would like to see in the future is direct support for Windows DFS, though a workaround is available.

Overall, Elanders said they were very happy with their Sun Storage 7410 systems. They like having an open software based solution that meets their needs today and will scale and grow to meet their needs in the future.

CHALLENGES FOR SUN

Sun's new Storage 7000 series fits in nicely with Sun's overall strategy around open software initiatives and unified solutions. This strategy has allowed them to build awareness of the products in the market. Still, there are some challenges that IDC feels Sun has to address to make the 7000 product a success. Sun should consider the following:

☒ Focus on educating its sales force and channel partners on how and where to sell its many storage solutions. Sun's Open Storage Partner Specialty Program and several new professional service offerings are designed to help customers take advantage of open storage technologies.

- ☒ Integration with key application vendors will be a key to the success of the 7000 family. Sun needs to continue to expand the alliances it has been building to ensure that major application vendors can take advantage of the advanced features in the 7000 more easily.
- ☒ Extending its unified storage approach to meet the needs of enterprise environments. Sun will have to prove to enterprise customers that this open storage architecture is powerful enough for their needs.
- ☒ Sun must show the value of its architecture by leveraging the open source capabilities to update and enhance system features faster and cheaper than traditional storage systems. Proving this added value will help offset concerns that users might have around open source components.

FINAL THOUGHTS AND GUIDANCE

Sun has taken the first steps in defining a new category of open storage solutions in line with its commitment to be an open technology company. It has shared many of the technologies used in the 7000 family with the open source communities. The company is leveraging its strategic directions in open source software and the development communities that surround them.

Sun's unified storage platform is well positioned to address current market needs by providing a simple to use and scalable storage solution. Industry standard hardware and open source software have the potential to improve the overall economics of storage for customers greatly. Sun's HSP architecture is a unique way to integrated flash-based SSDs into the storage system, providing for higher performance and lower power consumption when compared to traditional approaches. Appliance-based storage with comprehensive analytical tools and numerous bundled software features will help reduce the complexity of managing storage.

When evaluating unified storage solutions, IT managers need to put equal focus on the storage technology and the ability of the storage provider and its business partners to properly deploy and support the solution based on their current and future needs.

Copyright Notice

External Publication of IDC Information and Data — Any IDC information that is to be used in advertising, press releases, or promotional materials requires prior written approval from the appropriate IDC Vice President or Country Manager. A draft of the proposed document should accompany any such request. IDC reserves the right to deny approval of external usage for any reason.

Copyright 2009 IDC. Reproduction without written permission is completely forbidden.