

An Oracle White Paper
May 2010

Unleashing Application Performance with Solid-State Drives and Sun Servers

Executive Overview	1
Introduction	1
Computing Cost and Performance Challenges	3
Performance Bottlenecks.....	3
Power and Cooling	4
Increasing Complexity and Decreasing Reliability.....	4
SSDs Address the Need for Speed	5
Flash Economies	5
Fast, Predictable Performance	6
SSDs for Performance and HDDs for Capacity.....	6
Application Economics.....	7
Flashing the Datacenter	7
Better Application Performance	8
Reducing the Footprint	8
Greener IT	8
Less Maintenance and Support	9
The Benefits of SSD-Enabled Computing	9
Oracle Solaris ZFS—Optimized for SSDs.....	10
Conclusion	11

Executive Overview

This white paper, intended for executives and managers, describes how Oracle's Sun servers—equipped with flash-enabled solid-state drives (SSDs)—offer dramatic improvements to high-performance computing (HPC), Web 2.0, and datacenter application performance. In many environments today, high-performance CPUs must wait while data is delivered from the storage subsystem; even the fastest hard disk drives (HDDs) provide only hundredths of the speed of modern servers. Sun servers equipped with SSDs can be configured to use fewer HDDs to deliver the same input/output (I/O) storage performance, resulting in lower capital and operating expenses.

Introduction

Applications running on current multicore, multisocket servers are increasingly held back by storage systems that cannot keep up. Although HDDs provide the capacity needed to handle large amounts of data, their I/O performance capabilities are relatively slow. In fact, storage system I/O performance has increased by only a small fraction of server performance, which seems to be driven largely by Moore's Law (the widely accepted industry prediction made by Intel's cofounder Gordon E. Moore that states that the number of transistors on a chip will double about every two years). The underutilization of server processing power wastes time and decreases return on investment. Currently businesses are dealing with the disk-to-CPU-speed discrepancy in one of three ways:

- Suffering from slow performance and suboptimum productivity
- Overbuying spindles to achieve better I/O performance
- Buying tons of costly RAM

Enterprise SSDs, or flash drives, enhance I/O performance. These devices are breaking down storage bottlenecks and creating substantial improvements in server utilization and application performance. In addition to providing performance gains, SSDs use less power, so their use can lower maintenance-and-support costs. However, in the near term, they cannot provide the capacity needed at the lower price points as HDDs.

The right approach combines the strengths of both HDDs and SSDs. Flash devices can be placed in a new storage hierarchy—called a hybrid storage pool—that will increase storage performance dramatically by holding frequently accessed data. The addition of a flash storage hierarchy results in greater application performance: applications can access the most-used data more than 100 times as fast as when it is stored on a combination of enterprise SSDs and HDDs.

To help customers meet their growing enterprise, Web 2.0, and HPC application performance needs, Oracle offers fully integrated SSD and flash technology in its new Sun Blade, Sun Fire x86, and Sun CoolThreads servers. SSDs provide tremendous benefits for all operating systems available on these servers and are optimized by the Oracle Solaris operating system.

Early software integration for performance optimization, along with the codevelopment of SSD miniaturization for enterprise flash technology with other industry leaders, gave Oracle's Sun products an early edge in the market. The resulting solutions can save time and money for companies seeking to use technology to improve application performance and increase their return on investment. SSDs can help storage professionals who are looking for better ways to increase the efficiency of their datacenters by accelerating I/O-intensive applications without straining already tight computing budgets.

The right approach combines the strengths of both HDDs and SSDs. Flash devices can be placed in a new storage hierarchy—called a hybrid storage pool—which will increase storage performance dramatically by holding frequently accessed data. The addition of a flash storage hierarchy results in greater application performance: applications can access the most-used data more than 100 times as fast as when it is stored on a combination of enterprise SSDs and HDDs.

Computing Cost and Performance Challenges

In the current economy, reducing computing costs—both capital expenses and operating expenses—is a key initiative in every organization. Even so, there are new demands and requirements for applications and services and companies face increasing cost and performance challenges.



Figure 1: Even the fastest HDDs are much slower than today's CPUs and can hold back application performance.

Performance Bottlenecks

Server CPU processing capabilities have steadily increased, roughly doubling in performance every 18 to 24 months. Unfortunately, the ability to get the data to the CPU has not kept pace (see Figures 1 and 2). Disk I/O capabilities have increased only incrementally over the past decade (although storage capacity *has* kept up). The result? The bottleneck for application performance and server utilization is not typically the CPU but more often storage I/O. In many application environments, servers wait while data is delivered from the disk to the CPU.

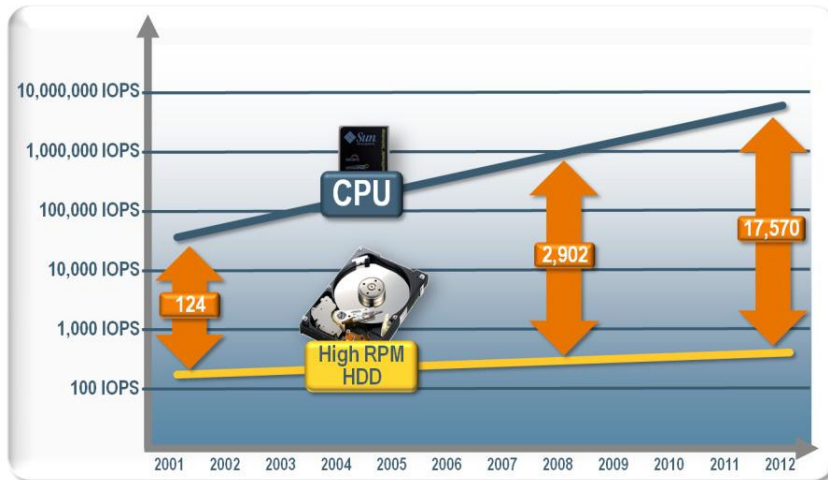


Figure 2: Hard disk drive performance has not kept up with the dramatic increase in CPU performance over the past several years. To keep data flowing to increasingly faster servers, more high-speed HDDs are needed.

Power and Cooling

Adding more servers to roll out new applications can be difficult. Many datacenters and computing operations are constrained by lack of space, power, or cooling. Power and cooling are regarded by 41 percent of Fortune 500 IT executives as a major problem. More than 25 percent of IT budgets are consumed by energy costs (power and cooling), and this percentage continues to grow. As the cost of power and cooling continues to rise, datacenters are using even more energy, because of increased storage requirements and the addition of more system memory.

In an August 2007 report, the U.S. Environmental Protection Agency (EPA) stated that the energy consumption of servers and datacenters had doubled in the previous five years and was expected to almost double again in the following five years, translating into increased operating costs. In many data processing environments, power and cooling are already at capacity; running out of power will severely limit business agility.

Increasing Complexity and Decreasing Reliability

Hard disk drives contain mechanical components that move with great precision but ultimately wear out from wear and tear due to excessive heat and vibration. High-performance HDDs—required to deliver data in high-performance server architectures—spin at 15,000 rpm. This produces more heat and vibration, which contribute to more failures. RAID and failover architectures ensure data safety but also increase deployment and RAS overhead costs.

Payroll is the largest datacenter cost, and power and cooling costs are the second-largest contributor to datacenter total cost of ownership. Installing more servers and storage only drives payroll costs higher.

Source: Patrick Thibodeau, "Servers swamp data centers as chip vendors push ahead," *Computerworld*
www.computerworld.com/hardwaretopics/hardware/server/story/0,10801,108433,00.html.

SSDs Address the Need for Speed

Continuing to roll out new applications in the same old way no longer makes sense. SSDs and flash technology represent breakthroughs in application performance and economics. Deploying new architectures with flash technology can significantly reduce the number of servers and storage devices needed to satisfy the same requirements.

Modern servers are more powerful than ever. Cost-effective four-socket systems (such as the Sun Fire X4450 Server) that use multicore, multithreaded Xeon CPUs are now available and can process data faster than ever. Average servers are fast approaching processing capabilities in excess of one million I/O operations per second (IOPS). However, mechanical disk drives have failed to keep pace: today's fastest drives are capable of only 300 IOPS to 400 IOPS (see Figure 3).

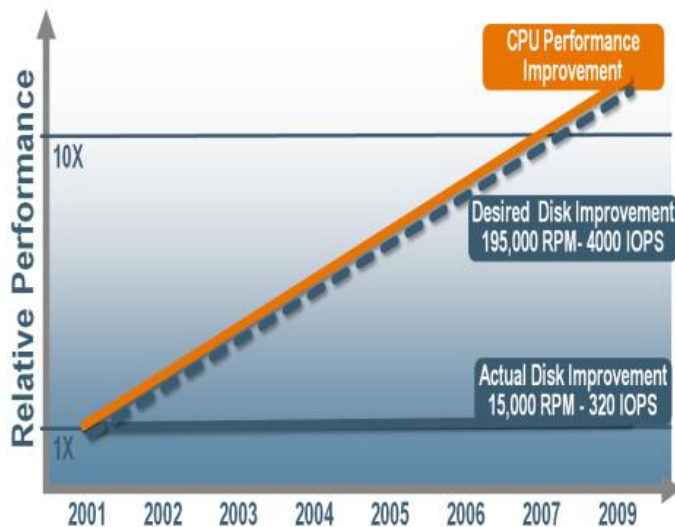


Figure 3: Hard disk drives would need to spin at nearly 200,000 rpm to meet the breakthrough I/O performance capabilities of an SSD.

Flash Economies

Makers of SSDs are benefiting from the high volume of flash-enabled devices—from handheld media players and cell phones to digital cameras and removable USB drives. This high volume has dramatically lowered the cost of flash memory components to the point that it is

economically viable to create enterprise versions as a strong storage alternative that can help rebalance system and storage I/O performance.

Robust data integrity, reliability, and availability—combined with breakthrough performance and power characteristics—have made it possible to create a new class of storage device: flash-based SSDs. With fast read and write performance, compelling price and performance points, and incredibly low power consumption, storage devices based on flash memory are poised to revitalize the storage industry. (See Table 1 for a comparison of data storage using SSDs and other technologies.)

	DRAM	SSD	HDD
Budgetary costs	US\$100/GB	US\$37/GB	US\$5/GB
Reliability (mean time between failures [MTBF], in millions of hours)	NA	2	1.5
Power consumption/units	10.5 W	2.5 W	12 W
Random IOPS/units	1,000,000 IOPS	Up to 35,000 IOPS	350 IOPS
Reaction time (latency)	90 nanoseconds	90 microseconds	10 milliseconds

Fast, Predictable Performance

Keeping systems supplied with data is key to overall performance. Flash technology reads and writes data in microseconds, placing its access time between that of HDDs (which read and write in milliseconds) and random access memory (which reads and writes in nanoseconds). Flash SSDs have highly predictable performance compared to HDDs, which suffer from long seek times and rotational latencies that make data access times extremely variable.

As a result, transferring data to and from solid-state storage media is faster than it is on mechanical disk drives. Enterprise SSDs provide tens of thousands of IOPS, compared to hundreds of IOPS for HDDs.

SSDs for Performance and HDDs for Capacity

The falling price of mechanical disk drives has made these drives extremely cost-effective for storing large amounts of data (about US\$2/GB for bare 15,000 rpm drives). Unfortunately, these mechanical disk drives are relatively slow and expensive to operate—offering only 300 IOPS to 400 IOPS and typically consuming 15 watts to 20 watts of power.

Solid-state drives are fast and consume little power (3,300 write IOPS and 35,000 read IOPS per drive and 2.5 watts) but are expensive, at US\$30 per gigabyte. Although SSD performance and operating costs are appealing, it is clearly not cost-effective in every case to substitute SSDs for mechanical drives in a storage array. Most industry analysts expect HDDs to be a cost-competitive storage component for years to come.

At current prices, using SSDs should be viewed not as a replacement for existing storage but rather as a means to enhance it. Conventional storage systems mix dynamic random access memory (DRAM) and HDDs. Solid-state drives are interesting, because they fall in a cost/performance sweet spot between those two components: SSDs are significantly cheaper and denser than DRAM and significantly faster than HDDs.

Until now, SAS and Fibre Channel 15,000 rpm drives represented the dominant building block for high-performance data storage—the only way to increase performance (more IOPS) was to add more drives—whether or not more capacity was needed. With the Oracle Solaris ZFS Hybrid Storage Pool, storage architects can use the best balance of resources—SSDs, SAS, and/or SATA drives—to achieve the desired capacity and read-and-write performance at the lowest cost.

Application Economics

Many HPC and datacenter applications are licensed on a per-server or per-CPU basis. If these servers are constrained by data I/O—and therefore underutilized—more servers will need to be deployed to serve a given number of users and more application licenses will need to be purchased. When SSD-enabled servers increase application performance, fewer servers—and therefore fewer application licenses—are required, lowering the overall deployment costs.

Flashing the Datacenter

Solid-state drives will significantly change the way storage infrastructure is architected for datacenters. For example, a codevelopment effort among industry leaders produced enterprise flash technology that radically reduced the size of SSDs. Sun servers were the first on the market with software integration that used SSDs in systems and storage solutions. This resulted in

- Significant increases in application performance
- Improved reliability
- Lower power and cooling costs
- Reduced maintenance and administrative costs

Solid-state drives speed up storage and reduce I/O bottlenecks, enabling high-performance servers to increase server utilization and workload capabilities while greatly reducing the requirement for high-performance HDD architectures. (HDD architectures require large quantities of HDDs per server.) Businesses experience an increase in their return on assets—getting more transactions, more users, and more throughput from installed servers while lowering operating costs.

Better Application Performance

The SSDs in Sun servers can instantly turbocharge application performance. For those applications that are I/O-bound or virtualized (such as Web 2.0, media, and HPC), Sun servers configured with SSDs can offer significant performance enhancements. Preliminary benchmark testing shows that databases used to power Web applications can use SSDs to process more than seven times as many transactions per second and can improve response time by a factor of more than 65.

For many application environments (see Table 2) this means that Sun servers equipped with SSDs can support more users and deliver greater performance. And fewer servers mean fewer application licenses, which can further reduce costs.

TABLE 2. APPLICATION ENVIRONMENTS THAT CAN EXPERIENCE SIGNIFICANTLY INCREASED SERVER PERFORMANCE WITH SOLID-STATE DRIVES

APPLICATION TYPE	EXAMPLE APPLICATIONS
Database	MySQL, Oracle Database, IBM DB2, Microsoft SQL Server, Sybase
Web 2.0	Drupal, Alfresco
HPC	Lustre, Nastran, Abaqus, data analytics
Media and entertainment	Audio and video media servers
Enterprise	SugarCRM, TIBCO, Greenplum

SSDs can improve application performance for both new and existing application deployments on all operating system platforms, including Oracle Solaris, Linux, and Microsoft Windows Server.

Reducing the Footprint

Computing infrastructures that deploy SSDs can also reduce floor space requirements in high-performance environments. Some I/O-intensive applications require hundreds of “short-stroked” 15,000 rpm drives to meet service-level requirements and the I/O demands of high-performance servers. Servers with SSDs can now meet performance requirements (IOPS) for less and use much less space—freeing up racks and power for other uses.

Greener IT

Sun servers with SSDs can help conserve datacenter resources. Flash technology SSDs consume approximately 20 percent of the power of both dual inline memory modules (DIMMs) and disk drives. They also have lower cooling requirements, making them “greener” than alternatives. SSDs offer breakthrough capabilities—a dramatic reduction in power consumption and a significant increase in storage performance.

Less Maintenance and Support

In many ways, SSDs are a system component familiar to computer and datacenter operators. They have the same interface and the same plug-and-play form factor as traditional HDDs, and they are also hot-swappable. However, unlike the hard drives that use a motor to spin magnetic media and a read/write head to perform operations, enterprise SSDs contain no moving parts. Data is stored on integrated circuits that can withstand significant shock and vibration. In fact, enterprise SSDs operate in a wider thermal operating range and wider operational vibration range than HDDs and deliver a significantly longer mean time between failures (MTBF)—2.0 million hours for SSDs, compared to 1.5 million hours for HDDs. Simply put, SSDs last longer, and in many applications, fewer are required. Having fewer drives that last longer means that datacenter operators spend less of their time diagnosing and replacing failed devices.

The Benefits of SSD-Enabled Computing

Take a look at a simplistic example of the benefits of SSD-enabled computing that focuses on IOPS. Suppose an application environment needs 30 terabytes of capacity and an aggregated throughput of 30,000 IOPS to meet service-level agreements. A traditional approach would require one hundred 15,000 rpm 300 GB drives (see Figure 4), each of which would consume 12 watts and cost US\$550. The total capacity is 30 TB, the cost is US\$55,000, and the drives consume 1.2 kilowatts of electricity.




Figure 4: Solid-state drives can dramatically reduce capital expenditures and operating expenses. A single SSD has the IOPS performance of one hundred 15,000 rpm drives.

A better alternative for delivering the required 30,000 IOPS is to use a single SSD (approximately US\$1,000) along with a smaller number of high-capacity SATA drives that cost much less than the 15,000 rpm drives. The total capacity is still 30 TB, but the cost of the drives is US\$18,000

and the drives (using 13 watts each) consume a total of 390 watts of electricity. The savings are dramatic:

- The purchase cost of the drives is a 67 percent savings (US\$55,000 versus US\$18,000).
- The amount of electricity consumed provides a 66 percent savings (1.2 kilowatts versus 0.392 kilowatts).

Oracle Solaris ZFS—Optimized for SSDs

Oracle Solaris ZFS is ready to take advantage of SSDs today (see Figure 5). Unlike less sophisticated file systems, Oracle Solaris ZFS can recognize different media types and will seamlessly optimize the file system to use SSDs as high-performance disks to improve read and write throughput.

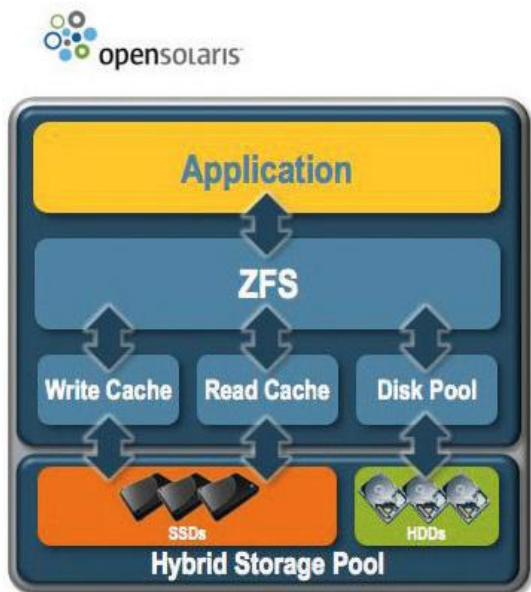


Figure 5: Oracle Solaris ZFS is optimized for SSDs, providing capabilities that instantly increase performance and reliability.

- **Instant performance.** Oracle Solaris ZFS transparently and automatically places the most-frequently-read data into the storage read cache—a combination of system memory (DRAM) and SSD or HDD drives—which instantly improves the performance of most I/O-intensive applications. The SSD can be configured as a cache disk without the need to dismount the file system in use. It automatically optimizes the file system to use the SSDs as high-speed devices that will improve read throughput for frequently accessed data.
- **High-speed write cache.** For certain types of writes, Oracle Solaris ZFS can greatly improve performance with SSDs, especially during use of devices designed for this purpose (such as write-optimized SSDs). The Oracle Solaris ZFS Intent Log (ZIL) enables applications that

demand synchronous writes—such as databases—to use SSDs and benefit from latency reductions by enabling the transaction to move safely to the next operation while data is asynchronously written to HDDs in the background. For example, internal testing showed that Oracle Solaris ZFS wrote the ZIL to the SSD far faster than the milliseconds needed to access an HDD—accelerating the application performance.

- **Solid reliability.** Using the mature, proven Oracle Solaris Fault Management Architecture, the sophisticated data integrity and correction algorithms in Oracle Solaris ZFS automatically detect and repair bad data, providing for self-healing capabilities. One feature—data scrubbing—runs as a low priority in the background to verify the data integrity of every block of data. In addition, an extensible set of agents can automatically take faulty components offline and link diagnostic messages to the support knowledgebase to guide administrators through easy-to-understand corrective tasks when human intervention is required.

Conclusion

Today's SSDs can accelerate application performance and lower operating costs across Oracle's line of Sun servers, including Sun Fire x86 servers, Sun Blade systems, Sun SPARC Enterprise servers, and Sun Netra carrier-grade servers. With flash technology, SSDs are poised to improve application performance and decrease response time dramatically. At the same time, they use less floor space and consume less energy, ultimately costing a fraction of the cost of HDD storage. Better reliability and increased server utilization mean fewer servers—lowering costs for system operators, service, and support—and result in lower capital and operating expenses.

As flash technology continues to develop and gain in popularity, expect significant trends:

- The reliability of enterprise SSDs based on flash technology is anticipated to exceed that of enterprise HDDs by a factor of 2 over the next two years.
- As flash technology makes its way into the storage infrastructure, the primary reason for using enterprise HDDs is likely to shift from performance to providing massive capacity.
- As flash technology is deployed, applications can be automatically accelerated through Hybrid Storage Pools in Oracle Solaris ZFS, thereby minimizing (or even eliminating) the need for tuning—including where to place data for maximum performance. Ultimately, other operating systems will include such advancements and many performance-sensitive applications will be redesigned with flash technology acceleration embedded.
- As more servers incorporate flash technology and provide more capacity in flash memory than in DRAM, the servers will realize greater performance and be more efficient in their use of power. Understanding this, Oracle intends to integrate Oracle Solaris ZFS, flash memory, and new hardware and silicon innovations to deliver high-performance, low-power, general-purpose storage and server appliances that will accelerate software that runs on the Sun SPARC processor-based or Sun x86 platforms.

With an unwavering commitment to getting the best from technological innovation and open storage solutions, Oracle continues to drive Sun server technology forward. With the introduction of flash technology and related hardware and software products, Oracle can help enterprises increase their application performance and lower costs.



Unleashing Application Performance with Solid-State Drives and Sun Servers
January 2010

Oracle Corporation
World Headquarters
500 Oracle Parkway
Redwood Shores, CA 94065
U.S.A.

Worldwide Inquiries:
Phone: +1.650.506.7000
Fax: +1.650.506.7200
oracle.com



Oracle is committed to developing practices and products that help protect the environment

Copyright © 2009, 2010, Oracle and/or its affiliates. All rights reserved.

This document is provided for information purposes only and the contents hereof are subject to change without notice. This document is not warranted to be error-free, nor subject to any other warranties or conditions, whether expressed orally or implied in law, including implied warranties and conditions of merchantability or fitness for a particular purpose. We specifically disclaim any liability with respect to this document and no contractual obligations are formed either directly or indirectly by this document. This document may not be reproduced or transmitted in any form or by any means, electronic or mechanical, for any purpose, without our prior written permission.

Oracle and Java are registered trademarks of Oracle and/or its affiliates. Other names may be trademarks of their respective owners.

AMD, Opteron, the AMD logo, and the AMD Opteron logo are trademarks or registered trademarks of Advanced Micro Devices. Intel and Intel Xeon are trademarks or registered trademarks of Intel Corporation. All SPARC trademarks are used under license and are trademarks or registered trademarks of SPARC International, Inc. UNIX is a registered trademark licensed through X/Open Company, Ltd. 0110

SOFTWARE. HARDWARE. COMPLETE.