



An Oracle White Paper  
December 2010

# Achieving Superior Manageability, Efficiency, and Data Protection with Oracle's Sun ZFS Storage Software

<a href="#">Introduction.....</a>	<a href="#">2</a>
<a href="#">Oracle's Sun ZFS Storage Software.....</a>	<a href="#">3</a>
<a href="#">Simplifying Storage Deployment and Management.....</a>	<a href="#">3</a>
<a href="#">Browser User Interface (BUI) .....</a>	<a href="#">3</a>
<a href="#">Built-in Networking and Security.....</a>	<a href="#">4</a>
<a href="#">Transparent Optimization with Hybrid Storage Pools.....</a>	<a href="#">4</a>
<a href="#">Shadow Data Migration.....</a>	<a href="#">5</a>
<a href="#">Third-party Confirmation of Management Efficiency.....</a>	<a href="#">6</a>
<a href="#">Improving Performance with Real-time Storage Profiling.....</a>	<a href="#">7</a>
<a href="#">Increasing Storage Efficiency.....</a>	<a href="#">8</a>
<a href="#">Data Compression.....</a>	<a href="#">8</a>
<a href="#">Data Deduplication.....</a>	<a href="#">9</a>
<a href="#">Thin Provisioning.....</a>	<a href="#">9</a>
<a href="#">Space-efficient Snapshots and Clones.....</a>	<a href="#">10</a>
<a href="#">Reducing Risk with Industry-leading Data Protection.....</a>	<a href="#">10</a>
<a href="#">Self-Healing Fault Management .....</a>	<a href="#">10</a>
<a href="#">Data Integrity.....</a>	<a href="#">10</a>
<a href="#">Triple-parity RAID and Triple Mirroring.....</a>	<a href="#">11</a>
<a href="#">Remote Data Replication .....</a>	<a href="#">11</a>
<a href="#">Conclusion.....</a>	<a href="#">11</a>

## Introduction

Continued growth in storage has made storage efficiency and manageability a top requirement for today's storage systems. Storage administrators often have too little time and lack the proper tools to optimize performance and efficiency for their storage environments. This can impact both enterprise application performance and the overall cost and efficiency of the storage environment. These storage inefficiencies can result in business issues such as:

- Increased storage costs because disk space is used up by duplicate data and because poor housekeeping leaves unused data on primary storage
- Increased risk of service downtime and productivity loss due to unnecessarily long backup and recovery times that result from inefficient disk space utilization
- Difficulty increasing revenue because enterprise application performance bottlenecks are limiting business potential
- Inability to effectively utilize tiered storage to reduce storage costs

Software management features and capabilities in the storage environment provide a key to addressing these business issues. With the proper tools, storage administrators are not only more productive in managing storage environments, but they can also do a better job of improving the performance and efficiency of the storage platforms they manage.

This paper is intended to educate the reader about important criteria to consider when evaluating software features and functionality in unified storage systems. The paper describes how the software features in Oracle's Sun ZFS Storage Appliances can simplify the process of storage management and enable greater storage efficiency.

## Oracle's Sun ZFS Storage Software

Oracle's Sun ZFS Storage software is the key ingredient that enables Oracle's Sun ZFS Storage Appliances to deliver unmatched simplicity, ease-of-use, increased storage efficiency, and industry-leading price/performance. The software includes a rich set of data services, all of which are bundled with Sun ZFS Storage Appliances at no extra cost. Key features such as transparent optimization, data compression, data deduplication, mirroring, snapshots and clones, error correction, and a browser-based system management interface make Sun ZFS Storage software a critical part of Oracle's unified storage offering.

## Simplifying Storage Deployment and Management

In an effort to simplify storage deployment and management as well as reduce costs, many organizations have turned to storage appliances to meet their departmental and enterprise storage needs. Appliances offer a simple and versatile approach to storage and come with storage software preinstalled to help simplify deployment.

Pre-installed software is merely the first step in simplifying storage deployment and management. Oracle's Sun ZFS Storage software brings a new level of simplicity to storage through the software features described in the subsections that follow.

### Browser User Interface (BUI)

Provisioning and management is dramatically simplified through an easy-to-use management interface that takes the guesswork out of system installation, configuration, and tuning. As illustrated in Figure 1, the browser user interface (BUI) provides an intuitive environment for performing administration tasks, visualizing concepts, and analyzing performance data.

In simple configurations, such as when deploying the appliance as a high-speed backup solution, the appliance can literally be set up and configured in minutes. The click-and-drag installation saves valuable time for IT organizations. For more-complex configurations, Oracle offers professional services that can help speed deployment and reduce the risk of downtime or data loss.

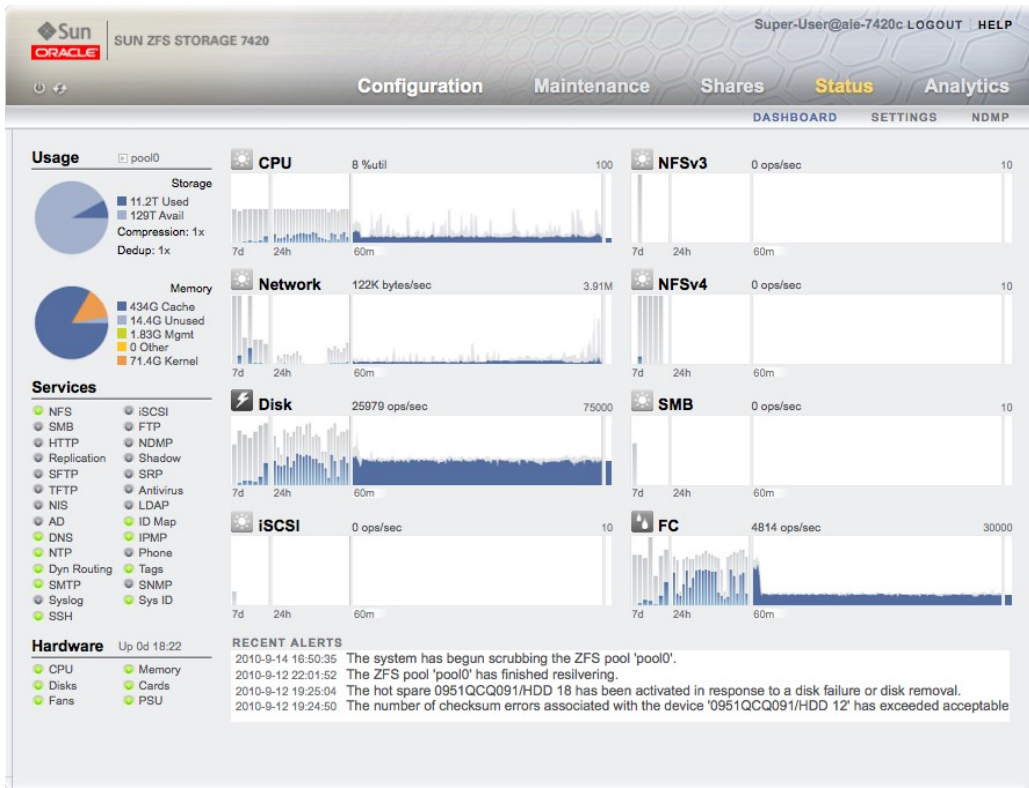


Figure 1. The Sun ZFS Storage Appliance BUI helps simplify deployment and management.

## Built-in Networking and Security

Sun ZFS Storage software supports a range of connection protocols and networking technologies, enabling Sun ZFS Storage Appliances to be easily deployed within an existing heterogeneous environment. The software supports seamless, transparent file sharing and secure access across both UNIX<sup>®</sup> and Microsoft Windows environments. It also includes native support for the Common Internet File System (CIFS) protocol to simplify integration with Windows file sharing. Secure directory services, including role-based access control (RBAC) for both UNIX and Microsoft Windows environments, are also provided via Lightweight Access Directory Protocol (LDAP) and Microsoft Active Directory support.

## Transparent Optimization with Hybrid Storage Pools

Unlike most file systems, Oracle Solaris ZFS recognizes different media types and will optimize how it handles each media type for enhanced file system throughput and application performance. To optimize data placement, Oracle Solaris ZFS makes use of read cache and write cache areas to transparently map data to the right type of media. This automated data placement approach reduces administrative costs by eliminating the need for data classification and the requirement to setup and maintain sophisticated storage management policies for tiered storage.

Sun ZFS Storage Appliances utilize Hybrid Storage Pools (HSPs) to implement the read and write cache areas with read-optimized and write-optimized SSDs (Figure 2). Hard disk drives (HDDs) are used for the high-capacity storage area to enable cost-effective storage of large data sets.




Figure 2. Oracle Solaris ZFS transparently optimizes data placement across different media types.

Data writes are transparently executed to the pool of write-optimized SSDs so that the writes can be quickly acknowledged, allowing applications to continue processing. The data is then automatically flushed to HDDs as a background task performed by Oracle Solaris ZFS. Another pool of SSD media acts as a read cache and Oracle Solaris ZFS manages the process of copying frequently accessed data into this read cache where it can be retrieved quickly.

Oracle Solaris ZFS looks at usage patterns to determine whether and how to use the different storage media. For example, large synchronous writes, such as video streaming, do not benefit from caching, so Oracle Solaris ZFS does not try to copy this type of data to write cache. Similarly, the read cache is populated based on an intelligent algorithm that takes into account not only the most recently used data, but also anticipated read requests and estimated data to be held in DRAM.

Caching frequently accessed data in SSDs can help improve application performance and enable sustained I/O throughput — all without the administrative overhead of ongoing storage optimization efforts.

### Shadow Data Migration

Another factor that affects initial deployment cost is the task of migrating data from a legacy storage environment. Software tools that can automate this process and make the legacy data available more quickly can provide another means of cost reduction in storage system deployment.

Customers have traditionally either paid for professional services to have their data migrated or have relied on manual approaches (e.g. rsync, tar, etc) to migrate data from legacy storage systems. Oracle has developed a new capability called shadow data migration that is included in the Oracle Solaris ZFS file system. Shadow data migration provides a highly efficient means to migrate legacy data and enables the customer to begin leveraging their investment in their new storage system even before the migration is complete.

Shadow data migration involves the creation of a “shadow” file system on the new storage device so that it can transparently pull data from the original source whenever necessary. The native file system on the new storage device is used for reads and writes of files that have already been migrated. When a user requests data that has not yet been migrated, the data is transparently migrated from the legacy file system and returned to the user. This approach makes the legacy data available quickly and also greatly reduces the time and cost of data migration. Rather than spending their time doing manual data migration, administrators can focus on optimizing storage and application performance.

### Third-party Confirmation of Management Efficiency

A recent study by Edison Group confirmed the efficiency of Sun ZFS Storage Appliances using Sun ZFS Storage software. The study showed that a Sun ZFS Storage Appliance can be managed more efficiently than a comparable NetApp system and also offers faster deployment and easier troubleshooting<sup>1</sup>. As illustrated in Figure 3, the study showed the following advantages for the Sun ZFS Storage Appliance:

- 34% less time required to provision and configure the storage and its connectivity
- 31% less time for maintenance and typical configuration changes
- 44% less time spent troubleshooting

---

<sup>1</sup> “Oracle ZFS Storage Appliance Comparative Management Costs Study.” Edison Group, August 2010.

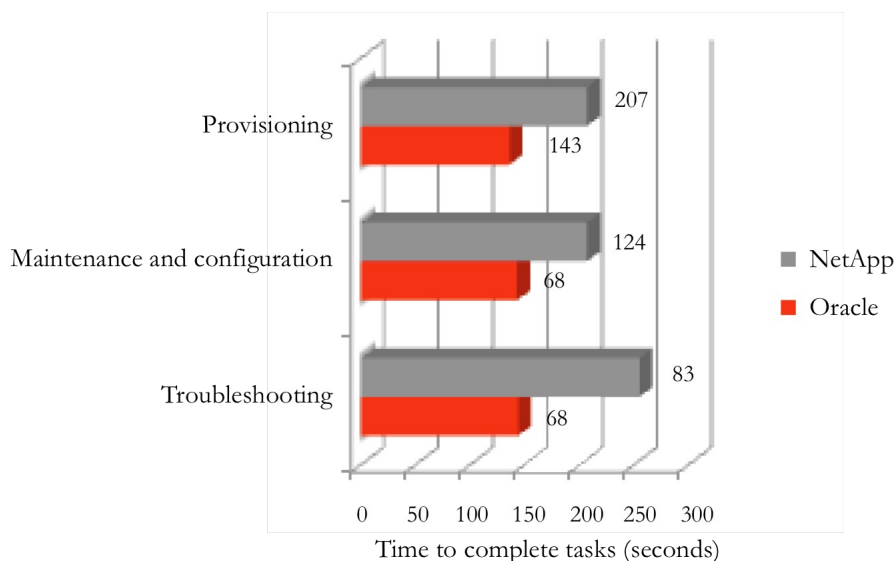


Figure 3. The Edison Group confirms that Sun ZFS Storage Appliances offer dramatic gains in efficiency over NetApp.

## Improving Performance with Real-time Storage Profiling

Without detailed visibility into storage subsystems, optimizing storage performance is a trial and error process that often leads to subpar results. DTrace Analytics software, which is based on the award-winning DTrace facility in Oracle Solaris, provides the industry's first comprehensive and intuitive business analytics environment for storage systems. It 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.

DTrace Analytics software uses built-in instrumentation to give real-time visibility throughout the data path. Its real-time analysis and monitoring functionality enables administrators to drill down for in-depth analysis of key storage subsystems. Graphical displays of performance and utilization statistics can be used to pinpoint bottlenecks and optimize storage performance and capacity usage — all while systems continue running in production.

Figure 4 shows an example of the kind of data that can be visualized with DTrace Analytics. The top part of the screen shows NFS operations per second broken down by client. This kind of detail makes it easy to see peak times and also provides insight into which clients are putting the heaviest load on the storage system. The bottom part of the screen shows network I/O coming in and going out of the storage system.

DTrace Analytics provides many ways to view performance information so that administrators can easily identify CPU or I/O bottlenecks across the storage systems. Administrators can then drill into the applications running on those systems to find out why. This level of observability is an industry first and enables administrators to identify application bottlenecks that might previously have gone undetected.

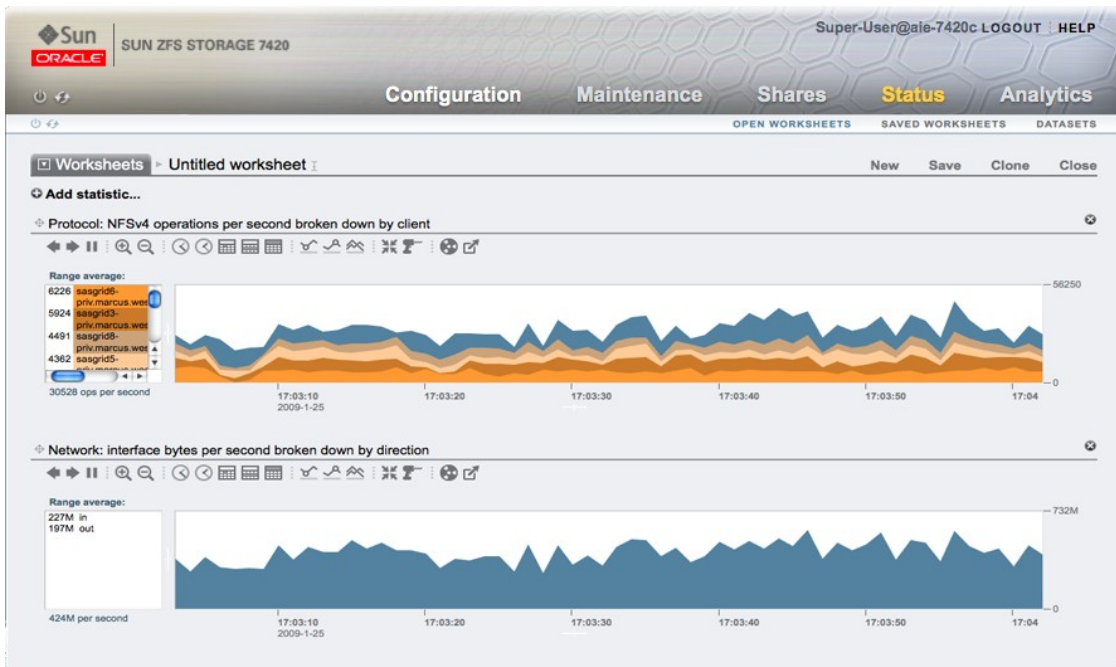


Figure 4.DTrace Analytics screenshot showing real-time storage performance statistics.

Statistics can also be archived to disk so that performance data can be reviewed and analyzed throughout the life of the system. Dedicated archive capacity allows administrators to preserve historical data for up to seven to 10 years. Archive data sets can be suspended or resumed as desired, or they can be destroyed if no longer needed.

## Increasing Storage Efficiency

It's not enough to look for high-density storage systems that can save datacenter floor space and reduce power and cooling costs. Storage software can also provide innovative ways to increase storage efficiency and thus reduce the total storage capacity required.

The sections that follow describe the features of Sun ZFS Storage software that can be used to help save datacenter floor space and reduce storage acquisition costs.

### Data Compression

Oracle Solaris ZFS provides built-in compression that can help reduce the amount of space required to store user data. This in turn increases the effective storage capacity available to applications. Having compression built into the file system not only simplifies management of complex storage architectures, but can also help minimize the impact on application performance by offloading the compression work to the Sun ZFS Storage Appliance. In addition, there are some cases in which compression can improve system performance due to the fact that compression results in fewer bytes of data, and therefore less I/O traffic to and from disks.

Compression typically provides good results for unstructured data in a file sharing environment, often yielding as much as 50% or more savings in storage space. Sun ZFS Storage software offers four different compression algorithms, enabling administrators to optimize compression for a particular application environment. The best way to determine the affect of compression for a given application environment is to run tests with actual data at various levels of compression and observe the results on disk utilization and application performance.

## Data Deduplication

The data deduplication feature in Oracle Solaris ZFS identifies and stores only unique blocks of data during storage writes. This can save significant amounts of storage capacity. For example, without data deduplication, when a 1MB file is emailed to 100 people, the system will store 100 identical copies of the file. By contrast, in a deduplication environment, only one copy of the file is stored.

There are multiple ways to implement deduplication so it's important to understand how the method of implementation might affect storage capacity and application performance. The most common method of deduplication is post-process deduplication. In this approach, the complete set of data is initially written to disk and then a subsequent process is initiated to remove duplicate copies of data. This approach avoids putting an additional load on the storage system or application during the initial writes. However, it incurs a penalty in the sense that it requires enough extra storage capacity to store the duplicate data until the post-processing is complete.

Sun ZFS Storage software uses an alternative approach called in-line data deduplication, which involves performing the deduplication while the data is being written. This approach has the advantage of avoiding an extra processing step and saving storage capacity during the temporary period when multiple copies of data would be stored in a post-process approach.

Sun ZFS Storage software deduplication is performed at the block level rather than at the file level. This means that when saving multiple different versions of large documents such as PowerPoint presentation files, only the blocks that are unique in the different versions would be stored more than once. This can greatly reduce the storage requirements for unstructured data.

For those customers interested in optimal capacity savings, compression and deduplication can be employed on the same data. Oracle Solaris ZFS seamlessly compresses data, which can then be deduplicated, allowing for maximum space savings.

## Thin Provisioning

Thin provisioning in Oracle Solaris ZFS helps optimize utilization of available storage by enabling on-demand allocation rather than requiring up-front allocation of physical storage. A common storage capacity issue is the need to configure databases with room for growth. A database that is expected to grow to 2TB would traditionally have been allocated the full 2TB of capacity from the beginning, even though it may have initially needed only 300GB of space. With thin provisioning, this database would be granted its 2TB size, but the physical media for the 1.7TB of unused capacity would not have to be in place until database has grown enough to make it a requirement.

Upon creation of a thin provisioned volume, the storage array communicates to the operating system the availability of a LUN of a specific size. However, the actual physical storage supporting the LUN can be much less than the reported size because physical storage allocation happens only when blocks are actually written to the storage array. As applications write to the LUN, the Oracle Solaris ZFS file system automatically and transparently grows the LUN in the background until it reaches the maximum size that was defined when the LUN was created.

In practice, some LUNs will never grow to their reservation size, or will only do so at some time well into the future when disk capacity is less expensive. By delaying storage allocation until the storage space is used, thin provisioning enables organizations to use less physical storage and delay storage purchases.

### Space-efficient Snapshots and Clones

Sun ZFS Storage software uses a copy-on-write approach for snapshots and clones so that no storage space is allocated to the snapshot or clone until the data is changed. Snapshots are read-only point-in-time copies of data and clones are writable copies.

With Oracle's copy-on-write approach, snapshots and clones do not occupy any physical storage space at initial creation. Read operations that retrieve data from a snapshot or clone are actually served by the original file system's blocks. When new data is written to the original source data set, the blocks are not modified until they can be copied into the snapshots and clones that have been made from that source. This maintains the point-in-time copy in the snapshot or clone, but defers the usage of disk space until necessary. When writes occur to a clone, this also results in the blocks of the clone being written to physical storage.

## Reducing Risk with Industry-leading Data Protection

Sun ZFS Storage software includes a comprehensive set of data protection features that provide protection from data corruption or hardware failures, as well as support for disaster recovery scenarios.

### Self-Healing Fault Management

Oracle Solaris ZFS takes advantage of the Oracle Solaris Fault Management Architecture to automatically and silently detect and diagnose underlying problems using an extensible set of agents. When a faulty hardware component is discovered, the self-healing capability in the Oracle Solaris Fault Management Architecture automatically responds by taking the faulty component offline.

### Data Integrity

Automatic data integrity checking and correction is provided by Oracle Solaris ZFS. If its block-level checksums identify a corrupt block, Oracle Solaris ZFS automatically repairs it, providing self-healing capabilities at the block level. It provides easy-to-understand diagnostic messages that are linked to

articles in Oracle's knowledgebase where administrators are clearly guided through corrective tasks for situations when human intervention is required.

## Triple-parity RAID and Triple Mirroring

Sun ZFS Storage software is the only software that protects data from as many as three drive failures with triple-parity RAID. With disk drives getting larger and many organizations implementing wider striping across more drives for increased performance, a third-parity bit helps reduce the risk of data loss. Triple mirroring, which maintains two extra copies of a storage volume, is also supported by Sun ZFS Storage software. Triple mirroring provides additional data protection as well as an easy way to increase random read performance for application workloads such as the Oracle database.

## Remote Data Replication

For storage solutions that will support mission-critical applications or a disaster recovery standby site, remote data replication is a critical feature. The remote replication feature in Sun ZFS Storage software enables Sun ZFS Storage Appliances to asynchronously stream data to another remote storage appliance. A project or share at the primary appliance is automatically replicated to a destination appliance, providing a consistent copy that can be used as a backup or as a standby site for disaster recovery. Remote replication is also sometimes used for distribution purposes such as disseminating copies of a new software release that must be installed on remote servers.

The initial replication requires replicating all the data in the source location (an Oracle ZFS project or share) to the destination site. For subsequent replications, only the incremental data that has changed at the source location will be sent to the remote destination. The replication can be configured to replicate either over a dedicated line or over a public network. Data can optionally be encrypted using Secure Sockets Layer (SSL) during transmission.

The following are some highlights and capabilities of the remote replication feature:

- One-to-one, one-to-many, and many-to-one topologies are all supported.
- Replication can be performed on-demand (initiated by an administrator), on a regular schedule, or continuously.
- A dedicated network is not required, although a private link is recommended where possible.
- Replication is supported between different platforms of the Sun ZFS Storage Appliance family, so it is possible to have a larger storage system at the remote site. This is useful when the remote site is also supporting other functions that are not addressed at the source site.

## Conclusion

Storage software is the primary tool that storage administrators have for increasing storage efficiency and reducing the cost of deploying and managing enterprise storage environments. The following list summarizes the most important software features that organizations should look for when considering

unified storage systems. These software features not only help reduce storage costs, but can also help storage systems deliver better performance, protect data integrity, and save disk space.

- Storage analytics that provide detailed visibility throughout the data path
- Transparent optimization by automatically taking advantage of the performance and price characteristics of different media types
- An easy-to-use browser-based storage management interface
- Data protection features such as triple mirroring, triple-parity RAID, and automatic data integrity checking
- Self-healing capabilities in the storage operating system
- Data compression that takes place in the storage system so as not to compromise application performance
- In-line data deduplication
- Space-efficient snapshots and clones
- Thin provisioning
- Easy-to-use remote data replication

For more information about the software features available in Sun ZFS Storage Appliances, visit [oracle.com](http://oracle.com) or call +1.800.ORACLE1 to speak to an Oracle representative.



Achieving Superior Manageability, Efficiency,  
and Data Protection with Oracle's Sun ZFS  
Storage Software  
December 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 © 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, ZFS, and Java are registered trademarks of Oracle and/or its affiliates. Other names may be trademarks of their respective owners.

**SOFTWARE. HARDWARE. COMPLETE.**