

## ORACLE SOLARIS ZFS

A BETTER, SAFER WAY TO MANAGE  
YOUR DATA

### KEY FEATURES

- **Rock-solid data integrity:** All data is protected by 256-bit checksums, and data consistency is maintained at all times.
- **No silent data corruption ever:** the Oracle Solaris ZFS self-healing feature automatically repairs corrupt data.
- **Mind-boggling scalability:** 128-bit file system, 16 billion billion times the capacity of 32- or 64-bit file systems.
- **Breathtaking speed:** Proven, cutting-edge technologies combine to optimize performance.
- **Near-zero administration:** Complicated storage administration is automatic and simplified, reducing administrative overhead by up to 80 percent.

*Anyone who has ever lost important files, run out of space on a partition, spent weekends adding new storage to servers, tried to grow or shrink a file system, or experienced data corruption knows that there is room for improvement in file systems and volume managers. Oracle Solaris ZFS is designed from the ground up to meet the emerging needs of a general-purpose local file system that spans the desktop to the datacenter.*

### The Future-Proof File System

Oracle Solaris ZFS offers a dramatic advance in data management with an innovative approach to data integrity, near-zero administration, and a welcome integration of file system and volume management capabilities. The centerpiece of this new architecture is the concept of a virtual storage pool, which decouples the file system from physical storage in the same way that virtual memory abstracts the address space from physical memory, allowing for much more-efficient use of storage devices. In Oracle Solaris ZFS, space is shared dynamically between multiple file systems from a single storage pool and is parceled out of the pool as file systems request it. Physical storage can be added to or removed from storage pools dynamically without interrupting services, providing new levels of flexibility, availability, and performance. And in terms of scalability, Oracle Solaris ZFS is a 128-bit file system. Its theoretical limits are truly mind-boggling:  $2^{128}$  bytes of storage, and  $2^{64}$  for everything else such as file systems, snapshots, directory entries, devices, and more. And Oracle Solaris ZFS implements an improvement on RAID-5, RAID-Z, which uses parity, striping, and atomic operations to ensure reconstruction of corrupted data. It is ideally suited for managing industry standard storage servers.

### Cutting-Edge Data Integrity

To ensure that the data on disk is self-consistent at all times, Oracle Solaris ZFS combines proven and cutting-edge technologies, such as copy-on-write and end-to-end checksumming. Data is always written to a new block on disk before changing the pointers to the data and committing the write. And, because the file system is always consistent, time-consuming recovery procedures like fsck are not required if the system is shut down in an unclean manner. Copy-on-write also enables administrators to take consistent backups or roll data back to a known point in time.

The Oracle Solaris 10 operating system with Oracle Solaris ZFS is the only known operating system designed to provide end-to-end checksumming for all data. Oracle Solaris ZFS constantly reads and checks data to help ensure that it is correct, and if it detects an error in a mirrored pool, the technology can automatically repair the corrupt data. This relentless vigilance on behalf of availability protects against costly and time-consuming data loss (even previously undetectable silent data corruption).

### High Performance

This radical new architecture optimizes and simplifies code paths from the application to the hardware, producing sustained throughput at exceptional speeds. New block allocation algorithms accelerate write operations, consolidating what would traditionally be many small random writes into a single, more efficient sequential operation.

Additionally, Oracle Solaris ZFS implements intelligent prefetch, performing read-ahead for sequential data streaming, and can adapt its read behavior on the fly for more-complex access patterns. To eliminate bottlenecks and increase the speed of both reads and writes, Oracle Solaris ZFS stripes data across all available storage devices, balancing I/O and maximizing throughput. And, as disks are added to the storage pool, Oracle Solaris ZFS immediately begins to allocate blocks from those devices, increasing effective bandwidth as each device is added. This means that system administrators no longer need to monitor storage devices to see if they are causing I/O bottlenecks.

### Simplified Administration

Most file system administration tasks are painful, slow operations that are relatively uncommon. Because these types of tasks are so infrequently performed, they are more prone to errors that can destroy a great amount of data very quickly. Oracle Solaris ZFS helps alleviate this problem by automating both common and less frequent administrative tasks.

Oracle Solaris 10 allows users to designate Oracle Solaris ZFS as the default file system and to easily boot Oracle Solaris from an Oracle Solaris ZFS root file system. Installing Oracle Solaris to an Oracle Solaris ZFS root file system is also possible. Likewise, users can easily migrate UFS root file systems to Oracle Solaris ZFS root file systems with the live upgrade feature.

Administering storage is extremely easy, because the design allows administrators to state the intent of their storage policies rather than all of the details needed to implement them. Creating a file system or performing other administrative activities is very fast—less than one second, regardless of size. There is no need to configure (or worse, reconfigure) underlying storage devices or volumes, because this is handled automatically when they are added to a pool. Oracle Solaris ZFS also enables administrators to guarantee a minimum capacity for file systems, or to set quotas to limit maximum sizes. Administrators can delegate fine-grained permissions to perform Oracle Solaris ZFS administration tasks to nonprivileged users, making it easy to deploy Oracle Solaris ZFS quickly.

### Volume Management Is a Thing of the Past

Unlike traditional file systems that require a separate volume manager, Oracle Solaris ZFS introduces the integration of volume management functions. The traditional combination of file system and volume manager maintains one-to-one mapping between the file system and its associated volumes. Oracle Solaris ZFS breaks out of this limitation with the storage pool model. When capacity is no longer required by one file system in the pool, it becomes available to others.

### Reduced Costs

Oracle Solaris ZFS can reduce costs by decreasing the time and complexity of administrative tasks, efficiently using resources, and eliminating volume manager licensing. All administration tasks are performed online, resulting in zero downtime for administration. In addition, the technology does not require a separate support contract because it is part of the Oracle Solaris 10 operating system. This can greatly simplify support issues; there is a single point of contact and only one maintenance contract for all software layers between the application and storage resources. And the RAID-Z feature saves money by maintaining data redundancy without the expense of a hardware RAID controller or nonvolatile RAM. Oracle Solaris ZFS supports new hybrid disk storage devices that include flash technology for greater performance.

### Compatibility

Applications do not need to be changed or modified to use Oracle Solaris ZFS and take advantage of its industry-leading capabilities. It employs familiar POSIX interfaces, and the existing storage infrastructure—device drivers, storage fabric, and devices—works without requiring changes. For applications that prefer to operate directly on block devices, Oracle Solaris ZFS provides the zvol volume emulator, which delivers all of the benefits of transactional integrity and checksums and is compatible with existing block-based, volume manager interfaces.

### Conclusion

By offering data security and integrity, virtually unlimited scalability, and easy and automated manageability, Oracle Solaris ZFS simplifies storage and data management for demanding applications today—and well into the future.

### Contact Us

For more information about Oracle Solaris ZFS, visit [oracle.com/solaris](http://oracle.com/solaris) or call +1.800.ORACLE1 to speak to an Oracle representative.



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. 0310

**SOFTWARE. HARDWARE. COMPLETE.**