

Solaris™ ZFS

Better, safer way to manage your data

Highlights

- Rock-solid data integrity: All data is protected by 256-bit checksums, data consistency is maintained at all times
- No silent data corruption ever: ZFS self healing features 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. 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 data center.

The future-proof file system

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 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, Solaris ZFS is a 128-bit file system. Its theoretical limits are truly mind-boggling — 2128 bytes of storage, and 264 for everything else such as file systems, snapshots, directory entries, devices, and more. And 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 like the Sun Fire 4500.

Cutting-edge data integrity

To ensure that the data on disk is self-consistent at all times, Solaris ZFS combines proven and cutting edge technologies, such as copy-onwrite 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 Solaris 10 OS with Solaris ZFS is the only known OS 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 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, 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, Solaris ZFS stripes data across all available storage devices, balancing I/O and maximizing throughput. And, as disks are added to the storage pool, Solaris ZFS immediately begins to allocate blocks from those devices, increasing effective bandwidth as each device is added. This means 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. Solaris ZFS helps alleviate this problem by automating both common and less frequent administrative tasks.

Solaris 10 allow users to designate ZFS as the default file system and easily boot Solaris from a ZFS root file system. Installing Solaris to a ZFS root file system is also possible. Likewise, users can easily migrate UFS root file systems to 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, as this is handled automatically when they are added to a pool. Solaris ZFS also enables administrators to guarantee a minimum capacity for file systems, or set quotas to limit maximum sizes. Administrators can delegate fine-grained permissions to perform ZFS administration tasks to non-privileged users making it easy to deploy ZFS quickly.

Volume management is a thing of the past

Unlike traditional file systems that require a separate volume manager, 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. 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

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 Solaris 10 OS. 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 Redundant Array of Inexpensive Disks (HW RAID)

Learn More

For additional information, please visit http://www.sun.com/software/solaris/zfs_learning_center.jsp.

controller or Non-Volatile Random Access Memory (NVRAM). 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 Solaris ZFS and take advantage of its industry-leading capabilities. It employs familiar POSIX interfaces, and existing storage infrastructure — device drivers, storage fabric, and devices — work without requiring changes. For applications that prefer to operate directly on block devices, 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, as well as easy and automated manageability, Solaris ZFS simplifies storage and data management for demanding applications today, and well into the future.