

## IceWEB unified storage system performance enhanced by ZFS

*File system provides transactional advantages and delivers tiered storage automatically*

By James E. Bagley

Senior Analyst

Deni Connor

Principal analyst, Storage Strategies NOW

August 2010

Unified Storage provides many advantages, delivering both the block access features of a storage area network for structured data like databases, and file access of network attached storage for unstructured data like documents and media. In order to support multiple external block and file interfaces, IceWEB has chosen the open source ZFS (Zettabyte File System) for the underlying data management in its IceWEB Storage System series. While supporting block, NFS and CIFS access (and others as well), the capabilities of ZFS are integral to the high capacity, performance and cost efficiencies of the storage systems.

The IceWEB implementation has these advantages due to its use of ZFS that we'll discuss in some depth:

- Capacities that will never be exceeded in real world applications;
- Superior data integrity without added operational overhead;
- Rapid snapshot and data cloning capabilities without duplication of data;
- Automatic spindle striping for maximum throughput and self-healing;
- Adaptive caching keeps the hottest data in the fastest storage;
- Built-in compression and deduplication maximizes storage efficiencies;
- Rapid file configuration without wasted space; and,
- The ability to continue to manage and utilize legacy systems without replacing them

### Scalable Capacity

ZFS, (the Z is for Zettabyte, which is  $2^{70}$  bytes) is a 128-bit file system that provides more addressability of data than will ever be needed by even a very large application. Using virtualized storage pools (ZPools), the system can handle 256 ZPools, each ZPool storing up to a zettabyte of capacity. Needless to say, that would be some storage system. In real world applications, the file system provides addressability and capacities that will not be outgrown, ever. IT can start by buying the capacity they need now, and be confident that they can scale the storage infrastructure to any size they will ever need without "forklift" upgrades.

### Data Integrity

A particularly useful feature of ZFS is its method of updating blocks, called 'copy on write.' When a block is updated, a copy is made, and the old data is then marked as invalid so that a subsequent collection operation can determine the disposition of that block – a feature that has advantages in data protection, replication and snapshot functions. Blocks containing active data are never overwritten in place; instead, a new block is allocated, modified data is written to it, then any metadata blocks referencing it is similarly read, reallocated and written. A 32-bit checksum or 256-bit hash code provides error detection and correction, thus ensuring the integrity of the data. Optionally, a second copy of the block is written to another physical unit ensuring no single point of failure within the physical components.

### **Snapshots and Clones**

Copy on write capability provides a benefit for snapshot and clone operations. Since the previous copy of the data already exists due to copy on write, snapshots and clones are created by simply keeping the previous block as part of the snapshot. This also provides special efficiencies since unchanged data is shared between the snapshot and the active data. A writable snapshot or 'clone' can also be created, and a parallel file can be maintained with its own copy on write transactions, yet unchanged blocks remain shared between the clone and the original. This is an extremely efficient method that automatically enforces deduplication at the most elemental level.

### **Maximum Throughput and Self-Healing**

ZFS will automatically stripe across all spindles in the pool. The IceWEB Storage System provides built-in support for RAID-Z (RAID-5), RAID-Z2 (RAID-6) and RAID-Z3 (triple parity). Synchronous mirroring is also a feature that can be implemented, again by using the copy on write transactions. Error correction is automatic and drives can be hot-swapped. Self healing is automatic and data integrity is continuously monitored by block and file error correction codes. The bottom line: the system is completely protected against all component failure and data corruption scenarios.

### **Automated Storage Tiering**

One of the most powerful features of ZFS is its built-in caching functionality. This capability provides two classes of hot data, one based on last cached data, another on data accessed frequently, that is, at least twice. In addition to the two classes, each class has a list of 'ghost' blocks that were recently evicted from the cache. Only metadata is retained on the ghost blocks, but this provides a powerful method of extending the cache to the next tier of storage.

IceWEB customers choose from ultra high-speed, enterprise-class Single Level Cell (SLC) solid state drives with capacities of 200GB, Serial Attached SCSI (SAS) and Enterprise Serial ATA (eSATA) drives to round out the storage tier options. ZFS automatically places the hottest data on the fastest storage, and migrates it out based on the Adaptive Replacement Cache (ARC) algorithm.

In addition, the IceWEB Storage System can be used as a means to manage storage systems from third-party vendors. Rather than facing a fork-lift upgrade, IceWEB products can integrate a customer's existing infrastructure into a new, unified storage environment, preserving the investment in legacy systems.

### **Variable Length Blocks, Compression**

Different workloads can take advantage of variable length blocks of up to 128KB. The administrator can vary this maximum in order to tune for a particular workload. If the customer opts for the compression feature, variable block sizes are used. If a block can be compressed to fit into a smaller block size, the smaller size is used on the physical media to use less storage and improve input-output operations per second (IOPs). An advantage of the IceWEB architecture is that higher performing processors and additional memory can be selected to offset the compression/decompression workload. This additional CPU power is also recommended for using the ZFS deduplication feature.

### **Compression, Thin Provisioning and Deduplication**

With the huge amounts of data being created by virtual servers, virtual desktops, and business applications, storage capacity requirements are spiraling out of control. ZFS has a combination of technologies to maximize the capacity efficiency of the storage infrastructure. Built in compression ensures that highly compressible files are stored efficiently. Thin Provisioning means that blocks are not allocated until written to, meaning that administrators can assign arbitrarily large volumes to servers without paying the penalty of a lot of wasted space. The most significant advantage is inline deduplication. Similar files used by multiple servers or virtual desktops are stored as single images within the storage pool. This can result in efficiencies as high as 95% in environments with many similar servers. This

combination of features is unprecedented and together means that the overall cost per usable terabyte is one of the best in the industry.

### **All-Inclusive and Easy to Manage**

Many of the features discussed in this paper are available on more expensive, traditional storage arrays. However, almost always they are add-on capabilities that are expensive, have a high learning curve, and involve substantial deployment and ongoing management costs. Every feature discussed here is included in the base IceWEB product; no add-ons, no additional GUIs to learn, no future costs, and no decisions about which features you may need. All the features are integrated together under a single UI, and many capabilities (such as thin provisioning) are completely automated – administrators don't need a PhD in storage to enable them.

For customers who are running a virtual server environment, IceWEB offers VM Manager, a unified management console for controlling all storage-related aspects of a virtualized infrastructure for VMware ESX, Citrix XenServer, and Microsoft Hyper-V environments. This adds the ability to easily see which volumes and file systems are being consumed by which virtual machines, and greatly simplifies the task of managing storage for many hundreds of virtual machines.

### **SSG-NOW Assessment**

The growth of storage requirements by organizations of all sizes has created an increased demand for unified storage systems. By supporting block and file access, and providing true heterogeneous system management, the addition of the IceWEB Storage System makes the forklift upgrade a thing of the past. Flexibility in configuring systems with low cost SATA drives, combined with higher performing SAS drives and even SSDs, give customers a level of customization not available in other unified systems. ZFS capacities insure that an organization will not hit an artificial limitation due to architecture and provides storage and performance efficiencies. IceWEB's long experience with the ZFS environment, along with extensive VMware support, is certain to make these systems successful and highly deployed.

Note: The information and recommendations made by Storage Strategies NOW are based upon public information and sources and may also include personal opinions both of Storage Strategies NOW and others, all of which we believe to be accurate and reliable. As market conditions change however, and not within our control, the information and recommendations are made without warranty of any kind. All product names used and mentioned herein are the trademarks of their respective owners. Storage Strategies NOW, Inc. assumes no responsibility or liability for any damages whatsoever (including incidental, consequential or otherwise), caused by your use of, or reliance upon, the information and recommendations presented herein, nor for any inadvertent errors which may appear in this document.