

# Oracle Delivers One-Two Punch with New ZS3 Series

Next Gen Oracle ZFS Storage Designed to Outperform Competitors in both Oracle Environments and Highly Virtualized Data Centers

Deni Connor  
Founding analyst, Storage Strategies NOW  
September 10, 2013

Database and storage administrators face increasing challenges due to the complexities involved with maintaining performance while dealing with the explosion of data. With the amount of information managed increasing by fifty times during the current decade, IT organizations are forced to perform labor-intensive manual tweaking while IT staffing levels are virtually frozen. Oracle's Application Engineered Storage, where the application and storage are communicating continuously and performing tuning functions on an autonomic basis, can reduce manual activities by as much as 65% while providing performance several times faster than competing storage systems.

## ZS3 Next Generation Storage Systems


Oracle ZFS Storage Appliances new ZS3 storage systems provide efficiencies and automation for Oracle Database and applications that cannot easily be duplicated by competitors. Engineered to operate optimally with Oracle's broad suite of database applications, the ZS3 Series has knowledge of workloads via direct communication of performance information between storage and application, which can eliminate bottlenecks even in highly virtualized environments. These software features are built to leverage the advanced performance capabilities that Oracle has engineered into its server and storage systems. These features include four levels of compression as well as deduplication, greatly reducing the amount of storage needed for a given data set.

Oracle's Hybrid Storage Pools (HSP) provide technologically advanced and cost-effective in-system data tiering and cache management functions. Sophisticated algorithms in HSP ensure that hot data is maintained nearest to the computational demand in DRAM or solid state Flash memory, while colder data is maintained on economical rotating media. Systems are configured with tiers of solid state storage including large DRAM pools—up to 2TB in dual-controller configurations—and both read and write cache areas use Flash memory. ZS3 systems are available with single or dual-controller options in two basic models: the ZS3-4, which comes with up to 25TB of cache (2TB DRAM, 23TB Flash) and scales up to 3.5PB of raw uncompressed capacity per cluster; and, the ZS3-2, which comes with up to 15TB of cache (2TB DRAM, 13TB Flash) and as much as 768TB of capacity per cluster.

While competitive systems may have similar features, they have very little DRAM – which acts as the fastest level of cache, and much less than the 25TB of total cache found in the ZS3-4. These competitive systems can also only operate based on storage system historical data; they don't have the intimate communication with the application workload's ensuing storage requirements.

Whereas competitive systems need continual manual tweaking, Oracle Application Engineered Storage performs these tasks automatically and ultimately more efficiently than a manually intervened process. The Oracle Intelligent Storage Protocol (OISP) -- only available on Oracle ZFS Storage systems -- dynamically automates database-to-storage tuning, and enables companies to do 3x more with the same headcount, delivering critical projects significantly faster and more efficiently. For instance, with OISP, it is possible to repurpose six out of ten storage administrators to other critical IT activities.

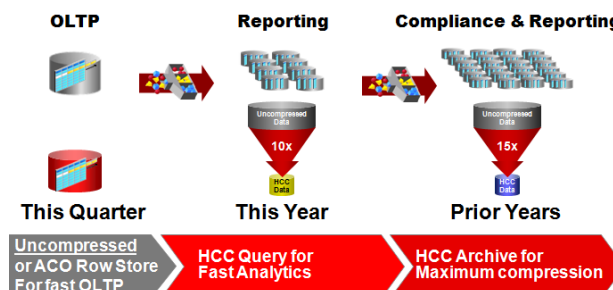
### Oracle Database Intelligent Storage Protocol



### Heat Map

With the ZS3, storage capacity and data access is continually optimized as data usage patterns change. With insight from a Heat Map that spans the entire database, data that changes frequently is advanced row compressed for fast OLTP performance, while data that changes less frequently or has not been recently accessed is compressed at higher ratios and migrated to capacity optimized storage tiers rather than staying on the original performance storage tier. This Heat Map process is performed automatically through a combination of Oracle Database 12c Automatic Data Optimization (ADO) and Hybrid Columnar Compression (HCC). The result is a reduction in storage capacity required, in some cases by as much as fifty times. For customers this means substantial savings, not only in terms of storage outlays, but also data center expenses—floor space, power, cooling, and management.

### Heat Map Process



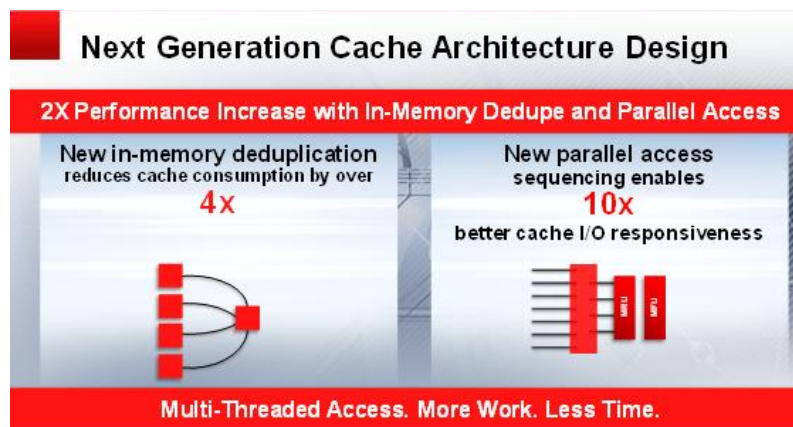
## Next Generation Cache Architecture

Oracle's HSP cache architecture includes the use of data reduction and data placement algorithms, as well as a unique use of massive DRAM pools in combination with Flash memory in both a single level cell (SLC) and multi-level cell (MLC) architecture for IO-specific use cases. DRAM pools are maximized by in-memory deduplication and parallel access sequencing, which combine to, in some cases, double the performance of the system with the same hardware.

SLC Flash is faster for write use cases than MLC, but is more expensive on a capacity basis. Therefore, the ZS3 system uses SLC as a write cache and uses MLC to support a large read cache. This automated real-time migration of data from DRAM to multi-class Flash to multi-class SAS HDD storage is key to maintaining system performance without continual tweaking and is a competitive advantage of the ZS3 Series.

By reference, Oracle pioneered the development and implementation of Hybrid Storage Pools in 2008 and the architecture has been flash-optimized and multi-core processor-enabled since then. The latest iteration features in-memory deduplication and parallel sequencing breakthroughs to deliver up to 2x improvement in overall response time compared to the prior generation. With ZS3's intelligent caching algorithms the vast majority of data is accessed from the highest speed memory—with typically 70%, and in some cases up to 90% of all IO is served from DRAM—up to 2TB per system—which is one thousand times faster than flash drives and 166,000 times faster than hard disk drives.

### ZS3 Series Cache Architecture Design



## Competitive performance advantages

The resulting HSP-based hyper-performance of the ZS3 storage systems is clearly demonstrated in published industry-standard benchmarks.

### SPC-2

In the SPC-2 comparison, which measures sequential I/O performance tests, which include large file processing and large database queries, the ZS3-4 system delivered a new world record with aggregate throughput of 17,244.22 SPC-2 MBPS™ with SPC-2 price performance™ of \$22.53. The nearest competitor spent a whopping \$131 SPC-2 price performance™ and only delivered 15,423.66 SPC-2 MBPS. In the same benchmark, other major competitors were at SPC-2 price performance™ of \$88.34 and \$95.38 and achieved aggregate throughput of 13,147.87 SPC-2 MBPS.

### SPECsfs

ZS3-4 storage system set new world records for dual-node NAS system throughput with 450,702 SPECsfs2008\_nfs.v3 ops/sec and for overall response time with 700 microseconds, and delivered an exceptional a price/performance ratio of \$1.08 per ops/sec. This compares with a competitive NAS supplier's 190,675 operations per second, 1,700 microseconds overall response time, and a price/performance ratio of \$4.24 per ops/sec. and a scale-out NAS supplier that needed a 56-node system to achieve equivalent performance at 456,223 operations per second, a slower overall response time at 3,270 microseconds, and whopping price/performance ratio of \$10.63 per ops/sec.

### 3<sup>rd</sup>-party TCO analysis

When compared to a major NAS supplier, a 400TB storage system with 4 million 50% write IOPs, the competitor's system cost 266% more to own and operate than the ZFS Storage Appliance, primarily due to filer sprawl. This analysis concluded that ZFS Storage Appliances are particularly suited for virtualized environments with high random IO. ZS3 storage appliances feature a highly-threaded symmetric multi-processing (SMP) operating system that leverages the new architecture's massive cache to support thousands of VMs per system, easily handling extreme highly threaded IO, which saturates conventional NAS filers, leading to filer sprawl. As VMware is a symmetrical multiprocessing (SMP) workload, the ZS3's SMP operating system enables it to scale linearly and support hundreds of thousands of threads concurrently, making it ideal for large-scale virtualized data centers and cloud deployments.

## Our Take

With this announcement, Oracle has thrown down the gauntlet on two fronts:

In Oracle Database 12c environments, only Oracle ZFS Storage can take advantage of the company's latest database capabilities to their fullest potential. Automated database-to-storage tuning, automated database-to-storage heat map driven tiering and Hybrid Columnar Compression are all capabilities that separate ZFS Storage from the rest of the competition in Oracle environments. No other storage systems have access to these unique Oracle Database features that help customers buy less storage and rely on less manual database and storage tuning. Running an Oracle Database without these features enabled is like flying a plane without all the engines turned on—you're still flying but you're not getting the maximum output and velocity that's available to you.

In highly virtualized environments running VMware, Hyper-V or Oracle VM, the ZS3 Series delivers the requisite throughput required to support thousands of VMs in a single system without forcing customers to purchase and manage additional systems—the familiar filer sprawl—as they would face with the challenges inherent in the limitations of conventional NAS filers. As we highlighted before, the DTrace Storage Analytics software on ZFS Storage Appliances can prevent a large enterprise from having a single VM user corrupt a 2,000 VM system by enabling customers to trace down an issue, for instance, to user #43 on VM#1,288.



## Snapshot Report

These advantages result in lower capital and operating costs, and most notably reduced manual labor in managing storage systems. According to Oracle, the ZS3 is primarily targeted against NetApp and EMC Isilon and VNX and will be leveraged as Oracle's fastest storage system to gain market share in both Oracle and non-Oracle environments. Fueled by its newly enhanced, hard hitting DRAM-intensive, Flash-centric cache architecture, the ZS3 is well positioned with a unique value proposition for customers across a multitude of industries.

### About Storage Strategies NOW™

Storage Strategies NOW™ (SSG-NOW) is an industry analyst firm focused on storage, server, cloud and virtualization technologies. Our goal is to convey the business value of adopting these technologies to corporate stakeholders in a concise and easy-to-understand manner.

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.

Copyright 2013. Storage Strategies NOW, Inc. All rights reserved.



Storage Strategies NOW | [www.ssg-now.com](http://www.ssg-now.com)  
512.345.3850 | [info@ssg-now.com](mailto:info@ssg-now.com)