

White Paper

Exadata Backup – Synergy with Oracle ZFS Backup Appliance ZS5 Series and Comparison with EMC Data Domain



Josh Krischer
December 2016

Contents

Executive Summary.....	3
Oracle Exadata Database Machine	4
Oracle Recovery Manager	5
Oracle ZFS Backup Appliance ZS5 Series	5
Oracle ZFS Backup Appliance ZS5 Series – Product Description.....	6
Oracle Integration and Synergy	7
Oracle Intelligent Storage Protocol	7
Oracle Hybrid Columnar Compression (HCC)	7
Oracle Enterprise Manager Plug-in for Oracle ZFS Storage	8
Oracle Snap Management Utility for Oracle Database	8
ZBA ZS5 Series Cloud Integration.....	8
Reliability, Availability and Data Integrity	9
Performance.....	10
Economics.....	10
EMC Data Domain Overview	11
Data Domain Architecture	11
Deduplication with Hashing-based Algorithm.....	11
Scalability and Upgrade Path	12
Data Domain Add-on Chargeable Features.....	12
Data Domain Deduplication in Oracle Database Environments	14
Deduplication Backup Appliances.....	14
Conclusions and Recommendations.....	15
Appendix 1: Oracle ZBA ZS5 Series Software.....	18
Appendix 2: Comparison Summary between Oracle ZBA ZS5 Series and EMC Data Domain. ...	19

Executive Summary

The exponential growth in the amount of data stored for analysis creates new challenges for system and database administrators for effectively managing the backup and restore of this data in the 24/7 economy. Backup data requires more and more storage capacity, combined with the need for very quick restores, forces organizations to look for a solution that is quick and easy to deploy and straightforward to manage. The current requirements from a backup/restore process are reliability, performance, easy management, flexibility and costs.

For Oracle Exadata Database Machine, the Oracle ZFS Backup Appliance (ZBA) ZS5 Series

“the “tailored-in” Oracle ZS5 Series exploits the co-engineering with Oracle Database Oracle Exadata and Oracle Public Cloud to deliver a level of synergy unavailable to competitive backup systems”

provides a high performance backup solution that dramatically reduces Oracle Database backup and restore times at significantly lower costs than competitive products. As one of Oracle’s engineered storage solutions, the “tailored-in” ZBA ZS5 Series exploits co-engineering with Oracle Database, Oracle Exadata and Oracle Public Cloud to deliver a level of synergy unavailable to competitive backup systems. As a result, users of Oracle Exadata with Oracle ZBA ZS5 Series storage do not need to ensure the compatibility or interoperability of the server, operating system, firmware, storage and networking, dramatically reducing integration and deployment time and greatly simplifying systems operation. This results in significant

reductions in CapEx and OpEx and solidifies the Oracle ZBA ZS5 Series’ position as a superior backup and restore solution for Oracle Exadata.

A comparison of the Oracle ZBA ZS5 Series with the EMC Data Domain deduplication storage systems in backing up Oracle Database shows that the Oracle ZBA ZS5 Series provides a much better solution. The Oracle ZBA ZS5 Series:

- Delivers better backup and restore performance due to higher processing power and direct connection to Exadata with high speed, low latency InfiniBand
- Supports all levels of Oracle’s Hybrid Columnar Compression (only available with Oracle storage) for highly effective Oracle Database data reduction
- Ensures much higher scalability (4.5x more usable capacity than Data Domain) without “forklift” upgrades
- Provides better availability with its dual-controller clustered configuration, data integrity with checksum end-to-end error detection, and correction of silent data corruption

In addition to these technical advantages, the ZBA ZS5 Series delivers several economical advantages which reduce the CapEx and OpEx of installations:

- No additional backup server hardware and software are required
- Oracle Hybrid Columnar Compression (only available with Oracle storage) compresses data up to 50x, reducing the amount of storage capacity required by 3x-5x or more
- Compressed data on the ZBA ZS5 Series can be immediately leveraged for secondary uses such as application development, test, QA, and analytics without rehydration

- Efficient management and troubleshooting through a user-friendly GUI and sophisticated storage analytics software reduces administration time
- Fast performance increases IT productivity and ensures that RPO and RTO SLAs are met

Oracle Exadata Database Machine

The Oracle Exadata Database Machine is purpose-built to run the Oracle Database. It leverages industry-standard hardware and unique software algorithms to deliver higher performance for Online Transaction Processing (OLTP), Data Warehousing (DW), On-line Analytics Processing (OLAP) and consolidation of these mixed workloads than competing systems at a lower cost. Exadata Database Machine is a 'turn-key' solution that includes all the hardware needed to run the Oracle Database, including the database servers, storage servers and InfiniBand networking— all pre-configured, pre-tuned, and pre-tested by Oracle.

The **Exadata Storage Server** (Exadata storage or Exadata storage cells) is used as the storage for the Oracle Database in the Database Machine. It runs the Exadata Storage Server Software that provides the unique and powerful Exadata technology including features such as Smart Scan, Smart Flash Cache, Smart Flash Logging, IO Resource Manager, Storage Indexes and Hybrid Columnar Compression. Exadata Storage Expansion Racks can be used to add capacity and bandwidth to the system.

The Exadata Database Machine offloads data intensive SQL operations into the Exadata Storage Servers. By doing that, data filtering and processing occur immediately and in parallel across all storage servers as data is read from disk. This storage offload reduces database server CPU consumption as well as the amount of data moved between storage and database servers. Some other interesting features include:

Exadata Smart Flash (each Exadata Storage Server includes 4 PCI flash cards) accelerates Oracle Database processing by speeding I/O operations. The Flash provides intelligent caching of database objects avoiding physical I/O operations. The Exadata Storage Server Software also provides the Exadata Smart Flash Logging feature to speed database log I/O.

Hybrid Columnar Compression (HCC) - Exadata storage provides an advanced compression technology, that typically provides 10x and higher (up to 50x) levels of data compression and significantly improves the effective data transfer. The HCC technique utilizes a combination of both row and columnar methods for storing data and achieving the compression benefits of columnar storage, while avoiding the performance shortfalls of a pure columnar format. A logical construct called the compression unit stores a set of Hybrid Columnar-Compressed rows as well. Queries run directly on Hybrid Columnar Compressed data and do not require the data to be decompressed. HCC is available for all Oracle engineered systems and Oracle Engineered Storage including the ZFS Storage Appliance ZS5 Series and FS Storage Systems when attached to Oracle Database.

Oracle Recovery Manager

The Oracle Recovery Manager (RMAN) utility is a component of the Oracle Database which is used for backing up and recovering the Oracle Database. It saves storage space and data transfer times by using file multiplexing and compression features. Oracle RMAN uses the incremental backup technique by backing up only the RMAN database blocks that have changed since the last backup. Oracle RMAN merges these changed blocks into the original image backup to create a new image of the Oracle data files which enables a full restore without the need to merge incremental backups into a full backup as part of the restore operation. This technique in addition to reducing the storage requirements also reduces significantly backup and in particular the restore times.

To ensure data integrity RMAN uses block-level corruption detection during backup and restore processes. Oracle RMAN offers data encryption capabilities as well as three levels of compression which optimize CPU utilization and the compression ratio to reduce storage capacity and network bandwidth requirements. Oracle RMAN can backup data to disk or tape. By leveraging Oracle RMAN, customers have no need to buy additional backup servers, extra software licenses, nor buy, integrate, and manage third-party technology.

Oracle ZFS Backup Appliance ZS5 Series

The Oracle ZFS Backup Appliance ZS5 Series is based on a state-of-the-art hardware and software architecture, including powerful multi-core processors (in a fully redundant dual-controller active-active configuration) and an SMP multithreading OS which allow running multiple workloads and advanced data services without performance limitations. Oracle's ZBA ZS5 Series complements the extreme performance of Oracle Engineered Systems, including Oracle Exadata, with throughput of up to 50 TB/hr for full backups and 60 TB/hr for full restores. As result, the ZBA ZS5 Series significantly reduces backup and restore times for Oracle Exadata, ensuring that backup windows and Recovery Time Objectives (RTOs) are met for high demand SLAs. The "tailored-in" ZBA ZS5 Series is co-engineered with Oracle Database and Oracle Exadata exploiting the synergy among the three products to deliver a solution that is superior to competitive backup systems.

The Oracle ZBA ZS5 Series is accessed via the NFS protocol. A faster protocol, Oracle Direct NFS (dNFS) significantly accelerates backup and restore performance by performing concurrent direct I/O, which bypasses any operating system level caches and eliminates any operating system write-ordering locks. In addition, dNFS Client performs asynchronous I/O, which allows processing to continue while the I/O request is submitted and processed.

The ZBA ZS5 Series also supports NDMP enabling it to be equally effective as a backup target for Oracle Database installations running on Oracle SPARC systems, non-Oracle servers, as well as providing backup to tape storage.

Oracle ZFS Backup Appliance ZS5 Series – Product Description

The ZBA ZS5 Series is available in two models: the high-end ZS5-4 and the mid-range ZS5-2. Both are available in single or dual-clustered controller configurations, and include a rich set of data services, the Hybrid Storage Pool intelligent cache architecture, multi-threaded SMP operating system, as well as a DRAM-centric system design and Oracle dNFS to power its superior backup and restore performance. The file system used is the advanced Oracle Solaris ZFS with 128-bit addressability.

The ZBA ZS5 Series contains one to forty-eight disk shelves. Each controller includes six 2.5" disk slots that accommodate two system disks.

	Oracle ZFS Storage ZS5-2	Oracle ZFS Storage ZS5-4
Architecture	Single-controller or dual-controller HA cluster with external disk shelf storage (stated specs assume active-active cluster)	Single-controller or dual-controller HA cluster with external disk shelf storage (stated specs assume active-active cluster)
Processor	4x 18-core 2.3 GHz Intel® Xeon® processors	8x 18-core 2.6 GHz Intel® Xeon® processors
DRAM cache	768 GB or 1.5 TB	3 TB
Read flash cache	0–102.4 TB	0–230.4 TB (0-307 TB with DE3*)
	<ul style="list-style-type: none"> 12 TB to 3.0 PB scalability (DE3*) Choose 20 or 24 HDDs per disk shelf If 20 HDDs, choose 0–4 SSD read and/or write accelerators per disk shelf Attach 1–16 DE2 or DE3* disk shelves for storage <p>Oracle ZFS Storage ZS5-2 default (one SAS HBA) can support up to 12 disk shelves. One additional HBA is required to support the maximum configuration. Two HBAs are the maximum number supported.</p>	<ul style="list-style-type: none"> 12 TB to 9.0 PB scalability (DE3*) Choose 20 or 24 HDDs per disk shelf If 20 HDDs, choose 0–4 SSD read and write accelerators per disk shelf Attach 1–36 DE2 or 1–48 DE3* disk shelves for storage <p>Oracle ZFS Storage ZS5-4 default (two SAS HBAs) can support up to 24 disk shelves. One additional HBA is required to support the maximum configuration. Four HBAs are the maximum number supported.</p>
Disk shelf/HDD options	<ul style="list-style-type: none"> Oracle Storage Drive Enclosure DE2-24C: 8 TB SAS-3 3.5" 7,200 RPM HDDs Oracle Storage Drive Enclosure DE2-24P: 600 GB / 1.2TB SAS-3 2.5" 10,000 RPM HDDs Oracle Storage Drive Enclosure DE3-24C: 8 TB SAS-3 3.5" 7,200 RPM HDDs* Oracle Storage Drive Enclosure DE3-24P: 600 GB / 1.2TB SAS-3 2.5" 10,000 RPM HDDs* 	
Standard and Optional Interfaces		
Integrated network	8x 10 Gb Base-T Ethernet ports	8x 10 Gb Base-T Ethernet ports
Optional network connectivity	10 Gigabit Ethernet, QDR InfiniBand HCA, 8 Gb FC HBA, 16 Gb FC HBA	
Optional tape backup HBA	Dual-channel 16 Gb FC HBA	
Maximum Ports per System		
10 GbE Base-T/10 GbE Optical/IB/ 16 Gb FC	24/16/16/16	32/24/16/16

Table 1: Oracle ZFS5 processors. Source Oracle Corp.

The ZBA ZS5 Series models support performance, capacity HDDs, and flash SSD in the storage shelves. The ZBA ZS5 Series dynamic caching supports two levels of caching, on flash drives and DRAM. The data movement between DRAM, flash cache, and disk drives is optimized by sophisticated algorithms ensuring fastest access to the required data.

The ZBA ZS5 Series offers both disk-to-disk (D2D) and disk-to-disk-to-tape (D2D2T) backup connectivity options. It is optimized as a target for Oracle RMAN backups, exploiting the advantages of Oracle Database and Oracle RMAN features to speed up data backup and restore processes.

Host connectivity: Both ZBA ZS5 Series storage systems support QDR InfiniBand, 10Gb Ethernet and 8/16Gb Fibre Channel connectivity options. InfiniBand connectivity, available in the ZS5 Series and used for high-speed direct connection to Oracle Exadata, is a unique feature among all backup appliances in the market. The 10GbE and 16Gb FC options provide rapid access to Oracle StorageTek tape storage solutions for customers who want to add an extra layer of data protection or need long-term data archiving to meet regulatory requirements.

Oracle Integration and Synergy

The Oracle ZFS Backup Appliance exploits the synergies and tight integration with the Oracle Database, increasing efficiencies, lowering TCO and ensuring high availability. With the advantage of “one company” engineering, developing, testing, the ZBA ZS5 Series delivers unique features to ensure Oracle software runs fastest and most efficiently on Oracle storage. The full table of software functionality is shown in Appendix 1. However, it is important to mention some unique functions here.

Oracle Intelligent Storage Protocol

The Oracle Intelligent Storage Protocol feature is exclusive to Oracle Database 12c and the Oracle ZFS Backup Appliance. This feature enables the storage to receive information from Oracle Database on incoming I/Os, providing storage visibility into the database on an unprecedented level. Oracle Database sends ZBA ZS5 Series metadata about each operation, enabling automated tuning and the most effective I/O processing to achieve optimal performance. With OS 8.6 and above, Oracle Intelligent Storage Protocol 1.1 includes integrated analytics. This capability provides advanced statistics on individual pluggable databases when ZBA ZS5 Series is used with Oracle Database.

Oracle Hybrid Columnar Compression (HCC)

Oracle Hybrid Columnar Compression (HCC) is one of the best examples of the synergy among Oracle Database and the ZS5 Series. Hybrid columnar compression (HCC) changes the way row data is stored to obtain dramatically better compression ratios, often as much as 50x. Rather than storing a series of rows inside a data block, HCC operates on an aggregate group of data blocks called a *compression unit*. Within the compression unit, HCC-compressed data contains the data for each column. And because most data tables tend to repeat the same column values, HCC saves space by storing each value once. By reducing storage space, HCC also reduces disk I/O requirements and uses less cache memory, too. This unique functionality, not supported by any other storage platform, is a standard feature at no additional costs for users.

With HCC, compression ratios up to 50x can be achieved, leading to storage capacity savings of 3x-5x over competitive systems, while query performance improves by factor of 3x-8x. In addition

to HCC, the Oracle ZS5 Series supports four levels of data compression¹ as well as in-line, block-level deduplication.

Oracle Enterprise Manager Plug-in for Oracle ZFS Storage

With Oracle Enterprise Manager Plug-in for Oracle ZFS Storage Appliance, user can monitor, manage, and provision distributed ZFS storage systems from a single console. It provides graphical and customizable single-pane-of-glass monitoring and storage provisioning for Oracle ZBA ZS5 Series across the enterprise. With this plug-in DBAs and storage administrators have access to an easy-to-read graphical dashboard in the Oracle Enterprise Manager console that displays key metrics and analytics, enabling deep insights into performance, capacity utilization, health, and availability.

Oracle Snap Management Utility for Oracle Database

The Oracle Snap Management Utility for Oracle Database is a standalone management tool for Oracle ZFS Backup Appliance. It provides a fast, efficient, and automatic way to snapshot, clone, and provision instances of Oracle Database from the copies stored on Oracle ZFS Backup Appliance. This feature enables administrators to reduce application development cycles, and shorten deployment times while increasing productivity and reducing costs.

In addition to Snap Management Utility, a rich portfolio of enterprise licensed software is available with features such as remote replication, encryption, cloning, etc.

ZBA ZS5 Series Cloud Integration

Over 400PB of Oracle ZFS Storage is installed in the Oracle Public Cloud, delivering IaaS, PaaS and SaaS services. They sustain high-bandwidth, high-throughput, transaction-intensive cloud workloads in addition to scalability and availability. The ZS5 Series is designed for cloud access by supporting OpenStack Cinder, Manila, Swift, and API calls to manage all functionalities. With the ZS5 Series, users can leverage Oracle Public Cloud services for object storage, elastic block storage, and archive storage while providing visibility and diagnostics from on-premise environments to the Oracle Public Cloud. Seldom used “cold data” can be migrated to Oracle Storage Cloud—Archive Storage at a very attractive price, dramatically decreasing the organization’s storage expenditures.

¹ The 4 levels of compression differ by compression ratio and the required server overhead. The user can select the optimal level suiting his requirements.

Reliability, Availability and Data Integrity

The ZBA ZS5 Series uses several techniques to ensure availability, data integrity and to prevent data loss or silent data corruption. The dual-clustered controllers provide redundancy during maintenance operations or in the event of a controller outage, backup and restore operations will

continue unabated. The ZBA ZS5 Series supports end-to-end *check summing* which reads and compares data to ensure that it's correct. Several predictive and self-healing capabilities ensure system availability by automatically diagnosing, fencing, and recovering from faults. The system also provides detailed diagnostic messages that link to Oracle's Knowledgebase, guiding administrators through corrective tasks when human intervention is required. In addition, the enterprise quality disk drives have better Mean Time Between Failures (MTBF) while industry-leading triple-parity RAID further reduces the risk of data loss.

Disaster Recovery schemes can be deployed by using remote replication to another ZBA ZS5 Series system to protect against total system or site loss on the primary. Additional data protection can be provided by the D2D2T option by backing up to a tape library system, such as Oracle's StorageTek SL8500 modular library system, shown in Figure 1. The ZBA ZS5 Series maintains a complete backup on disk for fast restore times, and additional copies can be archived to tape. To reduce eventual performance degradation, the tape backups are initiated from Oracle RMAN and utilize the most recent backup file stored on the ZS5 Series to create a copy for the tape library.




Figure 1. The Oracle Exadata backup and recovery solution utilizing the Sun ZFS Backup Appliance provides both disk-to-disk (D2D) and disk-to-disk-to-tape (D2D2T) options.

Performance

The Oracle ZBA ZS5 Series storage systems feature an innovative architecture that maximizes I/O throughput. This architecture includes an intelligent cache design based on a multi-threaded SMP operating system that takes full advantage of the large number of CPU cores across two controllers, Hybrid Storage Pools (HSPs) that provide very large DRAM and flash cache capacity, enterprise-class disks. This architecture was the foundation for the previous generation ZS3 and ZS4 Series systems record setting performance on the SPC-2 benchmark with a world record (at the time) of 17,244.22 SPC-2 MBPS™ and the second best overall price-performance with a result of \$22.53 SPC-2 price-performance™.² SPC-2 figures for the ZS5 Series are not available yet, but initial tests show very promising results.

The SPC-2 benchmark provides a source of comparative storage performance for streaming data workloads and is a good predictor of a system's backup and restore performance. Indeed, these robust capabilities coupled with the synergy with Oracle Exadata and Oracle Database enable the ZS5-4 system to achieve throughput of up to 50TB/hr for full backups and 60TB/hr for full restores—a 92 percent increase for backups and 353 percent increase for restores over the ZS3 generation.

“The Oracle ZBA ZS5 Series storage systems connect directly to Oracle Exadata's internally managed InfiniBand network, thus eliminating the need for a backup server hardware, software and the associated backup applications.”

Economics

Since Oracle RMAN and Hybrid Columnar Compression are included with the Oracle Database and the Oracle ZBA ZS5 Series can be connected directly to an Oracle Exadata Database Machine via an internally managed InfiniBand network, there is no need for a backup server, backup software and other associated infrastructure. This Oracle-on-Oracle solution for protecting Oracle Database(s) reduces integration costs as well as the complexity and risk that come with managing multi-vendor systems.

Simplified management reduces personnel costs through the use of an intuitive GUI that shrinks administration time by more than 30 percent by taking the guesswork out of system configuration, provisioning and tuning. In addition, Oracle's DTrace Analytics provides deep visibility for administrators to monitor crucial system parameters that can affect backup and restore of Oracle Exadata environments, speeding the resolution of performance bottlenecks and other issues. DTrace monitors, in real-time, the Oracle ZBA ZS5 Series' processor utilization, cache usage, data transfers and other system-related data. Administrators can drill down to areas of concern to get more precise information which helps in problem resolution. These across the board savings significantly reduce the CapEx and the OpEx of the backup/restore solution based on Oracle Exadata and ZBA ZS5 Series storage systems in comparison to those of other vendors.

² Results as of Sept. 10, 2013. Full results at http://www.storageperformance.org/results/benchmark_results_spc2#b00067.

Additional savings can be achieved using the Oracle Public Cloud. Infrequently accessed data (so called “cold data”) can be migrated to Oracle Archive Cloud at a very attractive price, decreasing the organization’s storage expenditures.

EMC Data Domain Overview

EMC Data Domain is a family of general purpose storage systems whose claim to fame is in-line deduplication for backup and archiving. Data Domain supports all major backup applications, including Oracle RMAN. However, it cannot leverage the same synergy that the Oracle hardware and software can. EMC discloses very few technical details about the Data Domain models. Internet search results provide mainly marketing information. There is no public information on which processors are used, cache size, flash, RAID levels, etc. EMC does not disclose restore times.

Data Domain Architecture

Data Domain launched in 2003. It was acquired by EMC in 2009, but over all these years the basic architectural design has changed very little. The Data Domain models are based on a single controller. The single controller represents a Single Point of Failure (SPOF). A component failure on the controller may cause a total system outage with potentially dire consequences, such as the inability to backup or restore data, and even data loss. Due to the single controller, microcode upgrades are disruptive as well. To provide higher availability, the three top models require investment in a second, cross-coupled, passive controller in standby mode as opposed to the Oracle ZBA ZS5’s dual active/active configuration.

The new EMC Data Domain family includes the entry-level DD2200 and four new models. The models differ in performance and capacity. The specifications are shown in table 2.

	DD2200	DD6300	DD6800	DD9300	DD9800
MAXIMUM THROUGHPUT	3.8 TB/hr	8.5 TB/hr	14 TB/hr	20 TB/hr	31 TB/hr
MAXIMUM THROUGHPUT (DD BOOST)	4.7 TB/hr	24 TB/hr	32 TB/hr	41 TB/hr	68 TB/hr
MAX USABLE CAPACITY	Up to 17.2 TB	Up to 178 TB	Up to 288 TB	Up to 720 TB	Up to 1 PB
W/ DD EXTENDED RETENTION²	-	-	Up to 576 TB	Up to 1.44 PB	Up to 2 PB
W/ DD CLOUD TIER²	-	-	Up to 864 TB	Up to 2.16 PB	Up to 3 PB
ES30 SHELVES	-	2 TB, 3 TB, 4 TB ³	2 TB, 3 TB, 4 TB ³	2 TB, 3 TB, 4 TB ³	2 TB, 3 TB, 4 TB ³
DRIVE TYPE	SAS	SAS, SATA	SAS, SATA	SAS, SATA	SAS, SATA
DS60 SHELF	N/A	3 TB, 4 TB	3 TB, 4 TB	3 TB, 4 TB	3 TB, 4 TB
DRIVE TYPE	N/A	SAS	SAS	SAS	SAS

1. Mix of typical enterprise backup data (file systems, databases, email, developer files). The low end of capacity range represents a full backup weekly or monthly, incremental backup daily or weekly, to system capacity. The top end of the range represents full backup daily, to system capacity. All capacity values are calculated using Base10 (i.e., 1TB = 1,000,000,000,000 bytes).

2. DD Cloud Tier and DD Extended Retention are mutually exclusive long-term retention solutions

3. No SATA Drive for 4TB drives

4. The following systems support High Availability Active/Standby configuration: DD9800, DD9300 & DD6800

Table 2: Data Domain family (source EMC)

As stated above, the Data Domain models are typical Purpose Built Backup Appliances with integrated deduplication that are not specially designed for nor have any unique integration points with Oracle Database. Data Domain Management Center is a dashboard-based virtual appliance which manages and monitors up to 75 Data Domain subsystems through a single interface. Data Domain has no scale-out architecture. Each box is a silo. This causes data fragmentation and leads to sprawl to achieve performance and capacity goals, much like NetApp leads to filer sprawl.

Deduplication with Hashing-based Algorithm

Hashing is CPU- intensive and the hash tables must be kept in memory to maximize performance. A major problem with the hash-based algorithm is the very large index that it requires. If the repository grows to the extent that the hashing tables cannot be contained in memory, performance will drop dramatically. This can be seen in particular with low-end models of the Data Domain family with less processing power and less memory for the hash table.

Scalability and Upgrade Path

Data Domain upgrades are not smooth and require “forklift” upgrades; that is, physical replacement. It supports data-in-place upgrades via swapping controllers but cannot scale out capacity and performance separately.

Physical upgrades may require data migrations, cause operation interruptions and may interfere with amortization time. For government agencies, forklift upgrades may require issuing a new RFP.

Data Domain Add-on Chargeable Features

Data Domain Boost (DD Boost)

In-line deduplication may suffer from poor performance under heavy load. To compensate, EMC introduced the Data Domain Boost (DD Boost) software, an agent which runs on database production servers and offloads some of the deduplication process from the Data Domain storage system to backup or, thereby slowing down the production environment. These servers compress and send only unique data segments across the network to the Data Domain storage system speeding up backup and reducing networking bandwidth requirements, according to EMC. While the Data Domain Boost software compensates for the relative low processing power of the Data Domain single controller, why should users have to pay for additional licenses and experience slower database and application performance due to the CPU load caused by running deduplication on their servers to correct Data Domain’s lack of performance due to a design flaw?

Data Domain Extended Retention

Although the largest usable capacity of the DD9800 is 1PB, it can be extended by another chargeable feature called Data Domain Extended Retention. This feature creates two tiers of storage on the Data Domain storage system; tier 2 is positioned for archiving or long-term data

retention. Again, users are forced to pay to make up for another of Data Domain’s design flaws: lack of capacity. Why would a company want to archive data on a single controller dedupe box when much more economical options such as deep cloud archival and on-premise tape archives are available?

Data Domain Performance

As seen in Table 2, EMC claims that the DD9800 top model can backup 31TB/hr without the chargeable DD Boost option. This performance is much slower, in fact 61% slower, than the ZBA ZS5 Series with 50TB/hr. The DD Boost option may increase the potential throughput to 68TB/hr with the production server penalty. Backup is important but restore is vital. As in the past, EMC has not published restore figures for the Data Domain models. There are several anecdotal reports that Data Domain restore may be a very lengthy process. A user tried to restore a database of ~6TB in size. He wrote in TechTarget: *“Running the restore off of the EMC Data Domain, was painfully slow. I canceled it after about 24 hours. It was at ~2% complete. Doing a little bit of math that database restore was going to take 25 days. While the restore was running we tried calling EMC support to see if there was a way to get the EMC Data Domain to allow the restores to run faster, and their answer was no, that’s as fast as it’ll run.”*. The whole story can be read: <http://itknowledgeexchange.techtarget.com/sql-server/if-you-thought-database-restores-were-slow-try-restoring-from-an-emc-data-domain/>.

Another case was Brazilian Justice Tribunal of Santa Catarina, which experienced database restore of 3 days. The cases happened with older Data Domain models. However, the design concepts of the new models have not changed much.

In summary, ZFS Backup Appliance is 6X the performance of Data Domain at approximately the price, with significant advantages both on and off Oracle Database environments. The following table is a performance examples comparison as claimed by Oracle:

ZFS Storage—Real World Backup & Restore Results
Never Lost a POC to Data Domain

	Data Domain 990	Oracle ZFS Storage	Oracle Advantage	
\$30 Billion Cloud Provider	Backup	8 hours	87 minutes	5.5X
	Restore	40 hours	127 minutes	19.0X
Fortune 500 Insurance Company	Backup	48 hours	44 minutes	65X
\$200 Billion Financial Services Leader	Backup	2 TB/hr	13 TB/hr	6.5X
	Restore	1 TB/hr	7 TB/hr	7.0X
\$100 Billion Financial Services Leader	Backup	555MB/s	3.5GB/s	6.2X

ORACLE

Table 3: Performance comparisons. Source Oracle

Data Domain Deduplication in Oracle Database Environments

EMC claims that Data Domain deduplication can reduce backup and archive storage capacity requirements by an average of 10-30 times. This figure may be achievable when backing up files such as documents, images, e-mails or Microsoft Share Point where duplicate files are prevalent. In contrast, relational databases, such as the Oracle Database, usually store data only once therefore the deduplication factor is much lower.

Further, Oracle RMAN uses an advanced, Oracle-specific format which makes the backup stream largely opaque to third-party backup applications. This opaqueness combined with RMAN's own compression or HCC³, leaves little duplicated data left for Data Domain to act on. In addition, deduplication is completely ineffective when data is encrypted by RMAN. In fact, EMC published a white paper (before HCC was available) titled EMC Backup and Recovery for Oracle 11g OLTP (no longer available) which shows a deduplication factor of only **6.3:1** in backing up Oracle Database to Data Domain. For maximum deduplication in backing up an Oracle Database, EMC recommends turning off HCC and performing daily full backups. In addition, data should not be encrypted and archived log files should not be included.

In summary, EMC Data Domain has several design flaws that make it a poor choice for Oracle Database backup and restore and an expensive solution for customers:

- Single controller failure means no backup, no recovery, and potential data loss
- Adding “stand-by” controller for availability increases the cost
- Slow performance means that backup and restore extend beyond allotted windows and fail to meet RPO and RTO SLAs
- Add-on, separately licensed agents (DD Boost) are required to make up for poor system performance, but slow down application/database servers, and add to the overall cost
- Deduplication ratios can only be met if HCC and encryption are turned off, violating Oracle best practices for database security
- Inability to scale leads to sprawl and complex backup and restore processes

Deduplication Backup Appliances

The development pace of storage technology is the fastest in history. Back in August 2005, a typical capacity HDD had 250GB capacity and cost \$160.00 (or \$0.64/GB). To contend with data growth with these small capacities and high costs, deduplication appliances were introduced in order to reduce backup and archiving storage capacities, lower costs and reduce bandwidth requirements for replication.

Meanwhile today, a 8TB SATA HDD is sold at \$0.028/GB – 32 times more capacity at 4.4% of the per GB cost which means much lower absolute savings. And with this evolution, the business

³ Hybrid Columnar Compression (HCC) is a unique and effective compression option for Oracle Databases only available with Oracle Storage.

case for stand-alone deduplication appliance silos has greatly diminished. Furthermore, Data Domain deduplication backup appliances carry disadvantages such as performance overhead, single-point-of-failure, and limited 24/7 operation due to internal reorganization (cleaning time). The main advantages—cost saving—fade away, but the disadvantages such as performance overhead remain.

A particular situation occurs with deduplication of Oracle databases. The RMAN backup block format is largely opaque to third-party deduplication products, which causes general purpose dedupe appliances to achieve very low deduplication ratios for Oracle Database RMAN backups and log files.

The same applies for encryption. If data is encrypted high up in the stack, at the application or database level, it provides greater end-to-end protection and reduces risk more broadly. However, deduplication is rendered largely ineffective because encryption, by its very nature, makes data blocks unique and illegible. Therefore, deduplication appliances require encryption to be turned off in order to realize higher deduplication ratios. Turning off encryption on Oracle Database voids Oracle Best Practices and puts businesses at risk just to accommodate a dedupe device. In addition, Data Domain cannot turn off deduplication to support a non-deduplicated disk shelf, as requested by some customers.

“High-availability features such as active controller clustering for failover, a self-healing file system architecture that ensures end-to-end data integrity, and a rich set of enterprise-class data services make Oracle ZBA ZS5 Series an ideal choice for enterprise storage and backup”

Some Data Domain users experienced much lower deduplication rates than expected. One of them was the already mentioned Brazilian Justice Tribunal of Santa Catarina. Another one was American Northern Trust which acquired four DD990s in 2012 mainly due to the deduplication ability. Due to poor performance and low deduplication factors, they had to add four more in 2014 and finally decided to pursue other solutions in 2015.

And then, most importantly, when restoring the deduplicated database is most urgent, the data has to be rehydrated before it can be restored, much less recover the database. This can dramatically hinder performance and extend recovery time beyond that specified by the enterprise’s SLAs.

Conclusions and Recommendations

Users demand simplicity, scalability, performance, manageability, and ease of use—all can be delivered by the Oracle ZFS Backup Appliance ZS5 Series. Through its acquisition of Sun Microsystems, Oracle inherited Sun’s server experience and Sun’s StorageTek division, which can look back at 45 years of experience in storage technologies. The co-engineering among these groups has resulted in the synergy from which the Oracle Exadata, Oracle Database, Oracle RMAN and Oracle ZBA ZS5 Series emerged as an enterprise-grade backup/restore solution to protect the enterprises’ mission-critical data. This combination ensures high

performance, lower capital and operational costs and can be deployed faster than the integration of a third-party backup platform. Customers further benefit from single-vendor support for all software and hardware components and avoid “finger pointing” in problem determination and resolution.

Users can significantly decrease storage expenditures by changing the game and moving some data to cloud. Oracle’s ZFS Backup Appliance has been proven in a variety of private and public cloud implementations worldwide and is the NAS and data protection backbone of Oracle’s own cloud.

A comparison of the ZBA ZS5 Series with the EMC Data Domain deduplication storage systems in backing up Oracle Database shows that the Oracle ZBA ZS5 Series provides a superior solution (a comparison summary is provided in Appendix 3). The Oracle ZBA ZS5 Series:

- Delivers better backup and restore performance due to higher processing power and direct connection with high speed, low latency InfiniBand (8x faster than Data Domain),
- Promises lower price per GB with 8TB SAS drives than Data Domain’s 4TB SATA HDDs
- Provides fast enterprise quality disks with better MTBF than Data Domain’s SATA drives
- Supports Oracle Hybrid Columnar Compression—only available with Oracle storage—for highly effective Oracle Database data reduction
- Ensures much higher scalability (4.5x more capacity than Data Domain) without “forklift” upgrades
- Provides better availability with its dual-controller clustered configuration, data integrity with checksum end-to-end error detection technique, and correction of silent data corruption
- Provides cloud access by supporting OpenStack Cinder, Manila, Swift, and API calls available to manage all functionalities
- ZFS Backup Appliance is 6X the performance of Data Domain at approximately the same price, with significant advantages both on and off Oracle Database environments.

In addition to these technical advantages, the Oracle ZBA ZS5 Series delivers several economical advantages:

- In Oracle Engineered Systems environments no additional backup server hardware and software are required
- Oracle Hybrid Columnar Compression—only available with Oracle storage—compresses data up to 50x, reducing the amount of storage capacity required by 3x-5x
- Compressed data on the ZBA ZS5 Series can be immediately leveraged for secondary uses such as application development, test, QA, and analytics without rehydration
- Efficient management and troubleshooting through a user-friendly GUI and sophisticated storage analytics software reduces administration time
- Fast performance increases IT productivity and ensures that RPO and RTO SLAs are met

High-availability features, such as active controller clustering for failover, a self-healing file system architecture that ensures end-to-end data integrity, and a rich set of enterprise-class data services, make the Oracle ZFS Backup Appliance an ideal choice for enterprise storage and backup.

Josh Krischer is an expert IT advisor with over 47 years of experience in high-end computing, storage, disaster recovery, and data center consolidation. Currently working as an independent analyst at Josh Krischer & Associates GmbH, he was formerly a Research Vice President at Gartner, covering mainframes, enterprise servers and storage from 1998 until 2007. During his career at Gartner he was responsible for high-end storage-subsystems and disaster recovery techniques. He spoke on these topics and others at a multitude of worldwide IT events, including Gartner conferences and symposia, industry and educational conferences, as well as major vendor events.

Find more on: www.joshkrischer.com

Appendix 1: Oracle ZBA ZS5 Series Software.

Oracle ZFS Storage Appliance Software	
Included Features	Details
Oracle Intelligent Storage Protocol	Oracle Database 12c sends metadata to the ZFS Storage Appliance about each I/O, enabling storage to dynamically tune itself for optimal performance
File system	Oracle Solaris ZFS (128-bit addressability)
File-level protocol	NFS v2/v3/v4, CIFS, HTTP, WebDAV, FTP/SFTP/FTPS
Block-level protocol	ISCSI, Fibre Channel, iSER, SRP, IP over InfiniBand, RDMA over InfiniBand
Data compression	Four distinct compression options to balance data reduction with performance for specific workloads
Hybrid Columnar Compression	3x to 5x reduction in storage footprint with existing Oracle Databases for OLTP , data warehousing, or mixed workloads
Data deduplication	Inline, block-level deduplication
Monitoring	DTrace Analytics (for system tuning and debugging); dashboard monitoring for key system performance metrics; plug-in available for Oracle Enterprise Manager
Automated serviceability	“Phone Home” capability with automatic case creation, configurable alerts
RAID	Striping, mirroring, triple-mirroring, single-parity RAID, double-parity RAID, triple-parity RAID, wide stripes
Remote management	HTTPS, SSH, SNMP v1/v2c, IPMI
Snapshots	Read only, restore, Microsoft Volume Shadow Copy Support (VSS)
Directory services	NIS, AD, LDAP
Data security	Checksum data and metadata, antivirus quarantine
Network services	NTP, DHCP, SMTP
Backup	NDMP v3/v4, ZFS NDMP
Local replication	Replication within same Oracle ZFS Storage Appliance configuration (single or cluster)

Separately Licensed Features	Details
Clones	Writable snapshots
Remote replication	Replication from one Oracle ZFS Storage Appliance product to another. 1:N, N:1, manual, scheduled, or continuous
Oracle Snap Management Utility for Oracle Database	Fast, efficient, and automatic way to back up, restore, clone, and provision Oracle Database when one or more databases are stored on Oracle ZFS Storage Appliance

Appendix 2: Comparison Summary between Oracle ZBA ZS5 Series and EMC Data Domain.

Oracle ZBA ZS5 Series	EMC Data Domain
Co-engineered for deep integration with Oracle Database, Oracle Exadata Database Machine and other Oracle Engineered Systems.	Not certified for backup with Oracle Engineered Systems.
Industry-leading performance with Oracle Engineered Systems, with backup and restore throughput rates of 50TB/hr and 60TB/hr respectively.	EMC Data Domain 9800 delivers 31TB/hr backup; no numbers for restore.
Native support for Hybrid Columnar Compression (HCC) for 10x-50x compression ratios and up to 8x faster query performance.	HCC is not supported on EMC Data Domain or any other storage system.
Short backup and restore windows with high-throughput architecture, high-speed InfiniBand connectivity and optimized Direct NFS.	Long backup and restore times.
Two controllers and a clustered active/active configuration mean even a planned software upgrade won't take backup and restore systems offline.	A single controller means all backups and restores will be unavailable in the event of a hardware failure or controller software upgrade. Only the 3 top models support second controller in active/passive standby configuration.
Backups can be used for development, test, QA, or analytics with immediate and full access to HCC data from Oracle Recovery Manager (RMAN) images without the need for decompression.	No ability to access or deduplicate HCC data in RMAN image backups.
HCC 10x-50x compression and higher throughput performance reduces backup and secondary processing footprints.	Inferior performance and compression leads to storage sprawl as more equipment is needed to meet capacity demand.
Superior performance and efficiency mean fewer systems are required—lowering capital and operational costs.	Storage sprawl and numerous integration points mean more complexity and add up to higher CapEx and OpEx.
Seamless access to Oracle Public Cloud storage services.	Dell/EMC has no public cloud; provides customers one-way ticket to AWS.