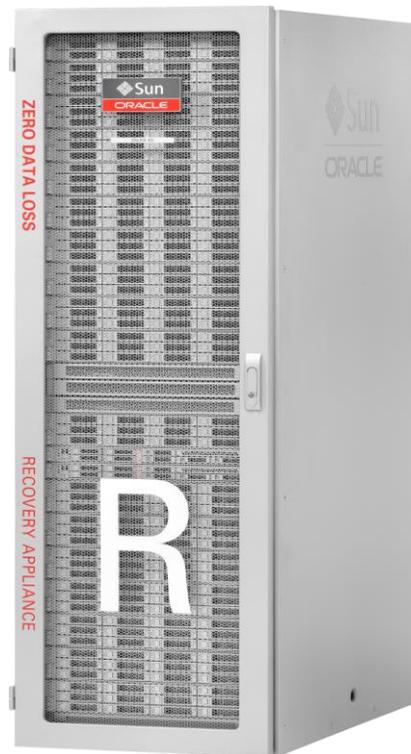


White Paper

Oracle's Zero Data Loss Recovery Appliance – and Comparison with EMC Data Domain



Josh Krischer
December 2015

Contents

Executive Summary.....	3
Oracle’s Zero Data Loss Recovery Appliance.....	4
Oracle’s Zero Data Loss Recovery Appliance – Principles of Operation	5
Oracle’s Maximum Availability Architecture (MAA)	6
Oracle Recovery Manager	6
Oracle’s Zero Data Loss Recovery Appliance – Product Description	7
Recovery Appliance Performance	8
Recovery Appliance Autonomous Tape Archival	8
Reliability, Availability and Data Integrity	8
End-to-End Management of Data Protection	9
Economics.....	10
Case Studies (Source: Oracle OpenWorld 2015).....	11
Deduplication Backup Appliances.....	14
EMC Data Domain Overview	15
Data Domain Architecture	15
Deduplication with Hashing-based Algorithm.....	16
Scalability and Upgrade Path	16
Data Domain Add-on Chargeable Features.....	17
Data Domain Deduplication in Oracle Database Environments	17
Conclusions	18
Appendix 1: Comparison Summary between Oracle’s Zero Data loss Recovery Appliance and EMC Data Domain.....	20

Executive Summary

For many organizations the Oracle Database is crucial for business, providing the highest levels of performance, availability, and data services. Despite the continuous developments in modern data center techniques, there is an “availability gap” between the Always-On Business requirements and what legacy backup solutions can deliver in terms of recovery. The daily backup is no longer sufficient to meet current Recovery Time Objective (RTO) and Recovery Point Objective (RPO) requirements. Backup is important but recovery is critical. Each minute of downtime and each lost datum may cause heavy business and reputation damages to modern organizations.

In 2014 Vanson Bourne, an independent market research organization, conducted an online survey of 760 senior IT decision makers from major world economies that employ more than 1,000 people. Some of the findings:

- 16.74 percent of backups fail to recover,
- The average recovery time of mission-critical applications is 2.86 hours while non-mission-critical applications take an average of 8.45 hours.
- The average downtime lasts 1.33 hours for mission-critical applications and 3.97 hours for non-mission-critical applications.
- The average cost of one hour of downtime for a mission-critical application is \$82,864, and for a non-mission-critical application, \$43,886. This means that an incident of mission-critical application downtime costs, on average, \$110,209; while non-mission-critical downtime costs, on average, \$174,227
- A single incident can cost organizations up to \$341,091 in lost data for mission-critical applications, and up to \$607,551 for non-mission critical.
- In total, a single incident of downtime can cost organizations \$451,300 for mission-critical applications and \$781,778 for non-mission critical, adding the costs of downtime and data loss.
- To answer current business requirements, 78 percent of organizations plan to change their data protection product in 2015 or 2016.

“A new engineered system, Oracle’s Zero Data Loss Recovery Appliance, answers these requirements by offering a new approach to database protection—one that virtually eliminates data loss and backup overhead while delivering cloud-ready scalability”

Backup data requires more and more storage capacity, combined with the need for very quick restores, forcing organizations to look for a solution that is quick and easy to deploy and straightforward to manage. The major requirements from a backup/recovery process are reliability, performance, easy management, flexibility, and costs.

A new engineered system, Oracle’s Zero Data Loss Recovery Appliance (Recovery Appliance), answers these requirements by offering a new approach to database protection—one that virtually eliminates data loss and backup overhead while delivering cloud-ready scalability. The Recovery Appliance can protect thousands of databases and provide

continuous real-time protection from system memory, so data can be recovered up to the last sub-second.

All backup-related processing is offloaded to the Recovery Appliance. This includes time-consuming compression, backup deletion, validation, and maintenance operations. This frees production system resources, even outside the backup window, which increases the performance of the production systems. Transactional changes in the Oracle database are captured by redo records logging. Oracle Database 11g Release 2 and higher databases can now continuously send redo directly from in-memory log buffers to the Recovery Appliance. This provides unique real-time data protection that allows databases to be protected until the last sub-second. All other backup and recovery related processing, including tape backup, is handled by the Recovery Appliance. This is one of the core architectural innovations of the Recovery Appliance, above and beyond today's backup solutions from the likes of EMC and NetApp.

The Recovery Appliance is the only backup/recovery solution built on Oracle's Exadata engineered system architecture. It is "tailored-in" with the Oracle database and exploits the co-engineering with Oracle Database and Oracle Recovery Manager (RMAN) to deliver a level of synergy unavailable to competitive backup systems. As a result, users of the Recovery Appliance 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.

Oracle's Zero Data Loss Recovery Appliance

The Recovery Appliance is tightly integrated with Oracle Database and dramatically reduces data loss exposure for all databases without impacting production environments. It ensures the highest levels of protection, performance, and backup storage efficiencies for all of an enterprise's Oracle databases¹ across all supported database versions, hardware, and OS platforms. The easy scale-out architecture can handle the data protection requirements of thousands of databases. It is the world's first and only engineered system designed specifically for database protection. The Recovery Appliance delivers continuous protection for critical databases while offloading virtually all backup processing overhead from production servers.

¹ Oracle Database 10.2, 11g, 12c

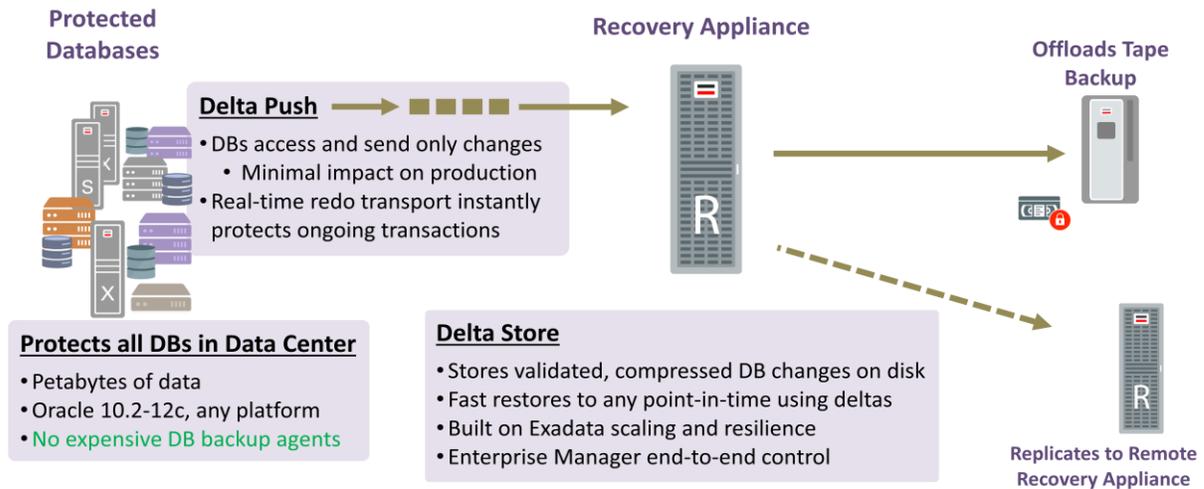


Figure 1: Zero Data Loss Recovery Appliance Overview. Source: Oracle

Oracle's Zero Data Loss Recovery Appliance – Principles of Operation

The principal design goal for the Recovery Appliance is to eliminate the loss of critical database data, and to reduce recovery time (RTO) in comparison to other existing data protection solutions. To accomplish this, Recovery Appliance introduces two industry-first technologies: Delta Push and Delta Store.

Delta Push deals with two types of data backups. First is Real-Time Redo Transport, a methodology based on the proven Data Guard high availability technology. It is important to understand that transactional changes in Oracle databases are done by Redo records logging. Oracle Database 11g Release 2 and higher databases can now continuously send Redo data directly from in-memory buffers to the Recovery Appliance. This provides unique real-time data protection that allows databases to be protected until the last sub-second. Since Redo is sent from database shared memory, the overhead on the production systems is extremely low.

Second, in order to reduce backup-related processing on production database systems to an absolute minimum, the Recovery Appliance uses an Incremental-Forever Backup technique. With this, the user need perform a full backup only once and then incremental backups thereafter. Incremental backups are made up of only changed blocks, which are very efficiently identified within the database itself using RMAN block change tracking. This eliminates the need to read or send unchanged data and is far superior to traditional deduplication methods. Traditional deduplication consumes server and time cycles to segment, compare and create references for duplicate data blocks, whereas this is fully database aware – already knowing where changes occurred without the additional processing to figure that out.

By eliminating backup processing, it allows the production systems to focus on the primary goal - serving business critical workloads. Additionally with this approach, network traffic is greatly reduced allowing low-cost 10 Gig Ethernet to be used for backups instead of Fiber Channel. Another benefit from the reduced network traffic is that it allows the Recovery Appliance to be located further away from the protected databases, even across a WAN in a remote data center.

When backups are received by Recovery Appliance, **Delta Store** validates the incoming changed data blocks and then compresses, indexes and stores them. These changed blocks are used to build Virtual Full Database Backups, which are space-efficient pointer-based representations of physical full backups as of the point-in-time of an incremental backup. This technique can reduce storage requirements by a factor of 10 times or more, depending on the data set and change rate of the protected database. During a restore operation, Delta Store rebuilds a physical full backup based on the latest incremental backup time. This restore operation is done by the powerful hardware (see Oracle's Zero Data Loss Recovery Appliance – Product Description) of the Recovery Appliance. Restoring from a Recovery Appliance is much more effective and faster in comparison to the traditional process of restoring a full backup and then sequentially restoring and applying all relevant incremental backups.

Additional features of the Recovery Appliance include cloud-scale capacity with online expansion to protect thousands of databases and enable data protection-as-a-service (DPaaS) across the enterprise or for external clients, in the case of a service provider. Since many Exadata customers leverage the product for Database-as-a-Service (DBaaS), the Recovery Appliance is highly complementary in that regard. For an extra level of data protection, the Recovery Appliance can provide autonomous tape archival, which enhances tape resource utilization and investment protection. And, the Recovery Appliance provides unified control and automation to simplify end-to-end management of the entire data protection lifecycle – from data creation to disk to tape..

Oracle's Maximum Availability Architecture (MAA)

Maximum Availability Architecture (MAA) is Oracle's best practices blueprint based on Oracle high availability technologies, extensive testing, field experiences, and customer recommendations. The MAA goal is to achieve the optimal high availability architecture at the lowest cost and complexity. It includes (among others) Oracle Database High Availability (e.g. active clustering, replication, continuous Oracle data validation, Oracle aware backup and recovery, online maintenance, etc.) for any platform or application, and integration with Oracle Engineered Systems. It is a unique value-proposition for Oracle Database environments to have this level of well-defined best practices. The Recovery Appliance is a core part of the MAA best practice for Oracle Database.

Oracle Recovery Manager

Oracle Recovery Manager (RMAN) is the Oracle “engine” 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.

To ensure data integrity, RMAN uses block-level corruption detection during backup and restore processes. Oracle RMAN can backup data to disk or tape. Since RMAN is a component of the Oracle Database, there is no need for additional backup servers or extra software licenses to buy, and no third-party technology to purchase and manage. Each block within a backup received by the Recovery Appliance is validated, compressed, and then indexed as part of a new virtual full backup, which appears as a normal full backup in the RMAN catalog. Validated full backups are

available at the point of the incremental, without ever incurring the cost of running full backup operations.

The Recovery Appliance tightly integrates with specific capabilities in the Oracle Database and RMAN to provide unique data protection capabilities and far better backup and restore performance in comparison to the EMC Data Domain deduplication appliances.

Oracle's Zero Data Loss Recovery Appliance – Product Description

The Recovery Appliance's hardware architecture is based on Exadata's proven scale-out technology, enabling it to support hundreds to thousands of protected databases. Storage capacity, compute power, and bandwidth scale together to avoid bottlenecks.

There are two types of servers: Compute servers, which connect to databases, perform replication, and tape backup, and storage servers.

Compute Server hardware (main components):

- 2x Eighteen-Core Intel Xeon E5-2699 v3 Processors (2.3 GHz)
- 256GB DD4 Memory
- 2x QDR (40Gb/s) InfiniBand Ports
- Dual 16 Gb Fibre Channel PCIe Universal HBA, QLogic for tape connectivity (optional)
- Optical and copper 1 and 10 Gb Ethernet ports
- Disk Controller HBA with 1 GB Supercap-backed Write Cache

Storage Server:

- 12 x 8 TB (raw) 7,200 RPM High Capacity disks
- 2 x Eight-Core Intel Xeon Processors
- 2 x QDR (40Gb/s) InfiniBand Ports
- Disk Controller HBA with 1 GB Supercap-backed Write cache

Each of the servers uses two redundant hot-swappable power supplies and redundant hot-swappable fans. The base configuration includes two compute servers and three storage servers internally connected using high-speed InfiniBand. Each base rack contains 2 x 36 port QDR (40

Gb/sec) InfiniBand switches for external connectivity (10GbE recommended connectivity for database backups) and an Ethernet switch for administrative connectivity to servers. The initial configuration is 94 TB of usable capacity. The base rack can be upgraded incrementally by adding storage servers into the rack, up to a maximum of 18 storage servers in a full rack. Each storage server adds 32TB of usable capacity. The total usable capacity of a full rack is 580 TB with an effective capacity of up to 5.8 PB. If additional performance is required, an additional rack can be connected to the first rack via high-speed InfiniBand interconnect. The additional rack includes its own pair

“The base configuration includes 2 compute servers and 3 storage servers internally connected using high speed InfiniBand links”

of compute servers, which add connectivity and processing power to the configuration. The appliance can scale capacity by adding storage servers. Up to 18 storage racks can be connected together, configured and managed as a single appliance and providing 10+ PB of usable capacity.

The 16Gb FC option provides rapid access to Oracle StorageTek tape storage solutions or other tape devices supported by Oracle Secure Backup for customers who want to add an extra layer of data protection or need long-term data archiving to meet regulatory requirements. If using a third-party media management application, copy to tape would go over 10GbE to the backup application's storage directly connected to its media server.

Recovery Appliance Performance

A single-rack Recovery Appliance can achieve a sustained delta Ingest rate of up to 12 TB/hour² and is able to support a restore at the same rate. Because it stores only the change data, this rate is equivalent to full backup, which needs an effective rate of up to 120 TB/hour. As racks are added to the configuration, both performance and capacity increase linearly. An 18-rack Recovery Appliance achieves Virtual Full Backup rates of up to 2 PB/hour, and 216 TB/hour of delta ingest and restore.

Recovery Appliance Autonomous Tape Archival

Several vendors (not having tapes in their tape portfolio) have declared the imminent demise of tape over the years, but in reality tape technology continues to be developed and brings unique benefits to its users. Tape offers the lowest cost solution for long-term data retention and archival, the best environmental values, and the ability to export/import to another site or vault. As opposed to online disks, a tape provides unalterable protection from intentional attacks by hackers or employees and accidental data deletion. However, traditional periodic full backups to tape impose high overhead on production systems.

The Recovery Appliance automates and offloads full, incremental and archived log backups to tape, completely eliminating the impact of tape backup on production database systems. The Recovery Appliance can send data directly to tape libraries using 16Gb Fibre Channel Adapters and highly integrated Oracle Secure Backup media management software. By using the Recovery Appliance, organizations can continue to use their existing tape libraries while eliminating tape backup overhead on production systems. Expensive media manager database backup agents on production servers are not required anymore because all tape activity is offloaded to the Recovery Appliance. All tape hardware products supported by Oracle Secure Backup, including Oracle's StorageTek Tape, are supported by the Recovery Appliance. Alternatively, other vendors' tape backup agents may be deployed on the Recovery Appliance for integration with existing tape backup software, media servers, and processes. Tape archival offload is a significant Recovery Appliance differentiator compared to most currently available backup solutions.

Reliability, Availability and Data Integrity

The Recovery Appliance uses several techniques to ensure availability and data integrity to prevent data loss or data corruption. Each of the compute and storage servers uses dual

² Actual performance metrics will vary depending on the configuration chosen, specifics of the user data, source system throughput, and network connectivity.

processors, two redundant hot-swappable power supplies and redundant hot-swappable fans. Recovery Appliance servers are clustered providing automatic failover, eliminating a single point of failure. The rack itself contains two redundant Power Distributions Units (PDUs). This full redundancy ensures continuous backup and restore operations in the event of a component failure.

To ensure data integrity, the Recovery Appliance uses end-to-end data validation. Database blocks are continuously validated to eliminate data corruption at every stage of transmission or processing. Data is striped and mirrored on Recovery Appliance disks; the Recovery Appliance catalog is triple mirrored; and database backups are stored with double redundancy.

“The synergy between developers of Oracle Database and Recovery Appliance and understanding of internal Oracle Database block formats enables deep levels of data validation.”

The synergy between Oracle Database and Recovery Appliance developers and understanding of internal Oracle Database block formats enables deep levels of data validation. All backup data and redo blocks are automatically validated as they are received by the Recovery Appliance, as they are copied to tape and/or replicated, as well as during restore from either.

The Recovery Appliance is built on the Exadata platform and benefits from Exadata Disk Scrubbing and Exadata Checksum checks capabilities. The backup blocks are periodically validated on disk via background processes to ensure that recovery operations will always restore valid data. If a corruption is discovered during validation, the Recovery Appliance’s underlying storage software automatically reads the good block from a mirrored copy and immediately repairs the corrupted block.

Disaster Recovery schemes can be deployed by using remote replication to another remote Recovery Appliance to protect against total system or site loss on the primary. As disaster protection, the Recovery Appliance can replicate data in real-time to a remote site and regularly archive backups to tape.

End-to-End Management of Data Protection

Many data protection tasks are divided among IT groups such as database administrators, backup administrators, and storage administrators. Such operations may create confusion among the groups. The Recovery Appliance offers Unified Management using Oracle Enterprise Manager Cloud Control, which provides a complete, end-to-end view into the data protection lifecycle through extensive monitoring, alerting, and reporting. It monitors from the time the backup is initiated using RMAN, to the time it is stored on disk, tape, and/or replicated to another Recovery Appliance in a remote data center. All backup locations are tracked by the Recovery Appliance catalog, so that any RMAN restore and recovery operation can retrieve the most appropriate backups, wherever they reside. Recovery Manager’s reports also deliver detailed metrics, such as real-time recovery status, unavailable from any other solution.

In addition to backup/restore operation monitoring, Enterprise Manager calculates the amount of space needed for each database based on its historical backup, current space usage and capacity required to meet the user-defined recovery window goal by database. An automated warning can be generated if space needed is within 15% (or other user-configurable threshold) of total available space.

The Enterprise Manager's user-friendly GUI interface and wizards simplify adding new databases with an appropriate protection policy and establishing the database's credentials. The Database Administrator then uses the Settings page to select the Recovery Appliance as the backup destination and optionally enables Real-time Redo Transport for continuous data protection.

In addition, Recovery Appliance employs policy-based automation to manage SLAs across the enterprise. Utilizing built-in and customizable policies, administrators can easily add databases to established SLA categories and define critical metrics such as retention period on disk, retention period on tape, recovery point objective and more. This not only streamlines adding databases but enables data protection to be delivered on-demand, as a value-added service.

Economics

The Recovery Appliance manages the backup/recovery processes, thus eliminating the need for a backup server hardware, software, and the associated backup applications. This Oracle-on-Oracle solution, which reduces integration costs as well as the complexity and risk that come with managing multi-vendor systems, doesn't require additional licenses for database backup agents.

Storing compressed changes only dramatically reduces the required storage capacity. The ability to use tape allows long-term, cost-effective backup archiving. Simplified management reduces personnel costs through the use of the Oracle Enterprise Manager Cloud Control. The intuitive GUI, intelligent automation, and detailed metrics shrink administration time and reduce the amount of personnel required to execute the tasks.

The ability to recover with minimal, sub-second data loss reduces outage time and data recovery time dramatically, thus reducing lost revenue, increasing productivity, and limiting business damages. These across the board savings significantly reduce the CapEx and the OpEx of the backup/recovery solution based on the Recovery Appliance in comparison to other vendors.

Case Studies (Source: Oracle OpenWorld 2015)



Brazilian Justice Tribunal of Santa Catarina

Covering nearly 7 million residents in 295 cities, the TJSC has 13,000 employees serving internal and public services for courts, judicial processes, appeals, and taxes. Nearly 14 million judicial cases have been recorded since inception of its Oracle-based system, 723,000 just in 2015. All judicial processes are recorded in digital format in a fully paperless administration environment with 24x7 services for legal appeals requirements.

The system uses nearly 40 Oracle production databases, including 15 critical databases with 30TB total volume and 24x7 access requirements. The principal database of 17TB serves more than 8,000 internal users and more than 70,000 web accesses daily. To ensure High Availability, the TJSC designed a redundant infrastructure in storage, network, and backup.

Initial Backup/Restore Infrastructure:

EMC Networker 8.0.2

- Two Node Clustered Media Servers
- 1Gbps communication network

EMC Data Domain

- DD670 configured as VTL
- 50TB usable capacity
- 2 Gbps SAN network
- Shared: Databases and File Servers

Issues with this infrastructure

- Low de-duplication ratio for the Oracle databases
- Slow performance, nearly 3 days to backup the most critical 17TB database (full weekly backup)
- Nearly 3 days to recover it
- Large DBA effort to manage backups (due to space constraints the old backups had to be expired daily)
- Poor space saving by de-duplication
- Inability to meet SLA requirements

Solution Backup/Recovery with Oracle Recovery Appliance

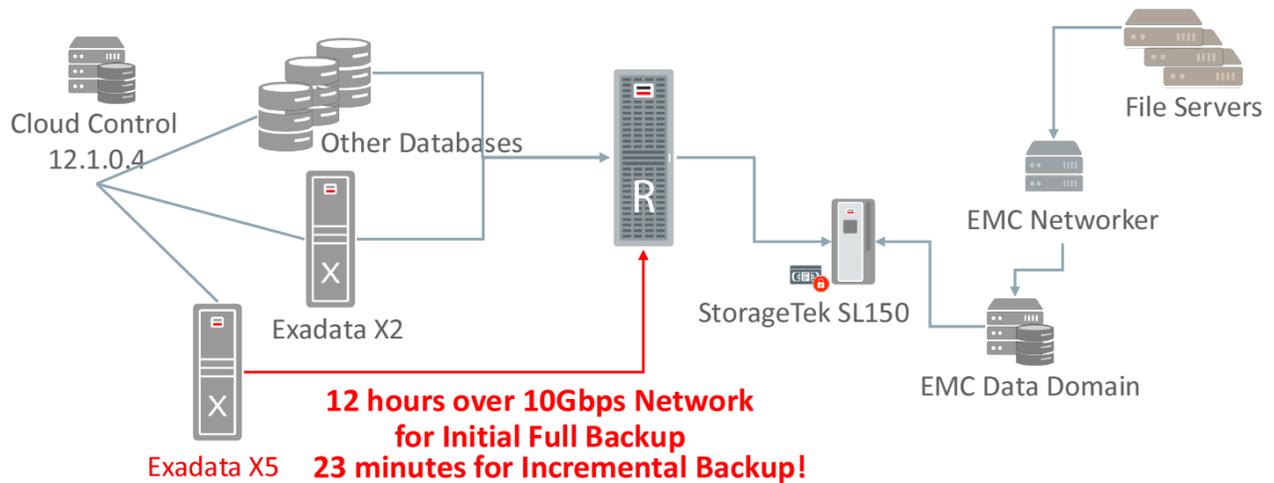


Figure 2: Backup/Recovery with Oracle Recovery Appliance

TJSC Backup/Recovery Benefits with Recovery Appliance:

- Significantly better RPO and RTO
- Higher de-duplication ratio
- No need for the weekly full backup - incremental forever
- **23 minute** incremental backup time (nearly **200X improvement** in backup performance)
- Cost savings - 2 Recovery Appliances less expensive than expanding EMC Data Domain to support same backup requirements
- Less administration efforts to tune and control every VTL of Data Domain
- Wider support: Oracle on Oracle, ONE support organization for the entire environment

Based on the success of the Recovery Appliance, the TJSC plans to deploy full Maximum Availability Architecture (MAA) with remote mirroring to recovery site.

More info :
www.tjsc.jus.br



SK Hynix, headquartered in Korea, is one of the global leaders in producing semiconductors such as DRAM and NAND flash, and System IC including CMOS Image Sensors. As the second largest manufacturer of memory semiconductors, SK Hynix is at the forefront of the IT industry. SK Hynix IT includes 40 Exadata engineered systems in four sites, which support the hi-tech manufacturing main system, Manufacturing Execution System (MES), manufacturing related automation systems and administration modules such as HR and ERP. The databases average 30~50TB each with 10% data daily increments containing business and sensor collected data,

which means structured and unstructured data. In addition to very high transaction processing, the infrastructure controls the manufacturing, stores sensor data, and performs real-time analysis at 7x24 availability requirements.

Before deploying the Oracle Recovery Appliance, Hynix used RMAN with legacy storage with the following issues:

- Backup consumed CPU on database server impacting the real-time transaction and I/O operations on the storage server
- Exadata resource management (DBRM and IORM) capped the backup resource usages under 40%
- Full backup took 13 hours, incremental took 3 hours
- Incremental backup needed periodic validation, and
- Manage RMAN scripts and scheduling and monitoring for 40 systems was complex

The solution was to deploy the Oracle Recovery Appliance as part of MAA infrastructure. To minimize the impact on real-time transactions, the backup is taken from a standby database in the Disaster Recovery site. Multiple databases are backed up by the Recovery Appliance, which relieves server overhead and ensures automatic data validation.

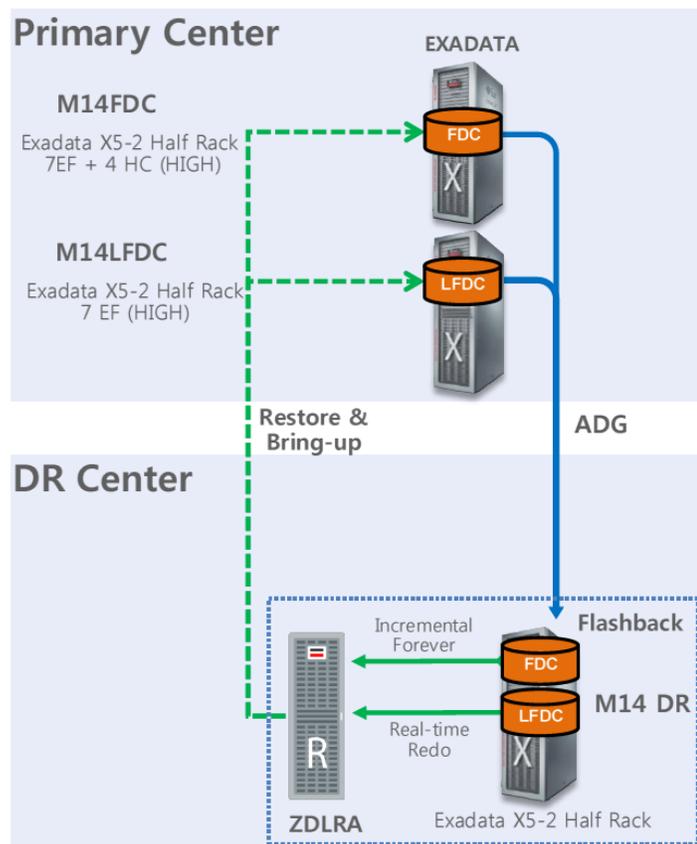


Figure 3: SK Hynix database DR architecture

After the installation of the Recovery Appliances as a part of the MAA structure, Hynix achieved the following benefits:

- 12X reduced backup window
- 3X reduced recovery time
- Up to 6X reduction in space
- Enterprise Manager simplifies backup environment configuration
- Intelligent backup space estimation saves time and capacity, and
- Automated fail-over reduces RTO time.

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 6TB SATA HDD is sold at \$0.045/GB – 24 times more capacity at 93% of the per unit cost. And with this evolution, the business case for stand-alone deduplication appliance silos has greatly diminished. Furthermore, 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 (saving costs) fade away, but the disadvantages such as performance overhead remain.

This is reflected in the market's dynamics. In September 2015, Liz Conner, Research Manager, Storage Systems, IDC, said: "*The worldwide PBBA³ market experienced a year-over-year decline in the second quarter of 2015 as the market continues to evolve. Focus continues to shift away from hardware-centric, on-premise PBBA systems to hybrid/gateway systems.*" The same report shows accelerating year-on-year declines in EMC's traditional storage products, including Data Domain, with -16.9% in 2Q15 in comparison to 2Q14. Another analyst's opinion is from Marc Staimer, October 1, 2015, Dragon Slayer Consulting: "*Deduplication storage appliances are losing their relevancy in the data center. Is This The End For Target Deduplication Storage Appliances?*".

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.

³ Purpose Built Backup Appliance

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 your business at risk just to accommodate a dedupe device.

And then, most importantly, when recovery time is critical and 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.

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 did not change. All the Data Domain models are based on a single controller.. This 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 back up or restore data, and even data loss. Due to the single processor, the microcode upgrades are disruptive as well.

The family includes six models from the entry-level DD2200 to the high-end DD9500. The models differ in performance and capacity. The specifications are shown in table 1.

	DD2200	DD2500	DD4200	DD4500	DD7200	DD9500
Maximum Throughput (Other)	3.8 TB/hr	5.6 TB/hr	10.6 TB/hr	10.6 TB/hr	12.6 TB/hr	27.7 TB/hr
Maximum Throughput (DD Boost)	4.7 TB/hr	13.4 TB/hr	25.6 TB/hr	25.6 TB/hr	28.3 TB/hr	58.7 TB/hr
Logical Capacity ¹	40-860 TB	1.3-6.6 PB	1.8-9.4 PB	2.8-14.2 PB	4.2-21.4 PB	8.6-43.2 PB
Logical Capacity w/ DD Extended Retention			3.7 - 18.9 PB	5.7 - 28.5 PB	8.5 - 42.8 PB	17.2-86.4 PB
Max Usable	Up to 17.2 TB	Up to 133 TB	Up to 189 TB	Up to 285 TB	Up to 428 TB	Up to 864 TB
Max Usable Capacity w/ DD Extended			Up to 378 TB	Up to 570 TB	Up to 856 TB	Up to 1.7 PB
ES30 Shelves	-	2 TB, 3 TB	2 TB, 3 TB	2 TB, 3 TB	2 TB, 3 TB	2 TB, 3 TB
Drive Type	SAS	SAS	SAS, SATA	SAS, SATA	SAS, SATA	SAS, SATA

Table 1: 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 monitor 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 dedupe sprawl, 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 de-duplication may suffer from poor performance under heavy load. To compensate, EMC introduced the Data Domain Boost (DD Boost) software, an agent which runs on and offloads some of the deduplication process from the Data Domain storage system to backup or database production servers, thereby slowing down your production environment. These servers compress and send only unique data segments across the network to the Data Domain storage system speeding up back up 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-processor controllers, 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? This is in marked contrast to the Oracle Recovery Appliance which can accelerate Oracle Database performance on production servers by typically 25% by virtue of its unique Delta Push / backup operations offload capability.

Data Domain Extended Retention

Although the largest usable capacity of the DD990 is 864TB, 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 positioning 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 now available?

Data Domain Performance

As seen in Table 1, EMC claims that the DD9500 top model can backup 27.7TB/hr without the chargeable DD Boost option. This performance is much slower than the Recovery Appliance. A single full rack Recovery Appliance with two Compute Servers and 18 Storage Servers is able to support Virtual Full Backups running at an effective rate of up to 120 TB/hour. Backup is important but restore is vital. EMC has not published restore figures for the new Data Domain models. The Oracle Recovery Appliance restore throughput is similar to its backup throughput. An 18-rack Recovery Appliance can deliver up to 2PB per hour backup and 216 TB per hour restore.

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 de-duplication 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 HCC⁴, 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](#) 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 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 recovery and an expensive solution for customers:

- Single controller failure means no backup, no recovery, and potential data loss
- Slow performance means that backup and recovery extend beyond allotted windows and fail to meet RTO SLAs
- Add-on, separately licensed agents (DD Boost) are required to make up for poor system performance, but slow down application/database servers
- 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 recovery
- No co-engineering with Oracle Database prevents sub-second RPO, incremental-forever, and capacity savings via HCC as well as automated backup and recovery management

Conclusions

A new engineered system, Oracle's Zero Data Loss Recovery Appliance, offers an unparalleled approach to Oracle Database protection—one that virtually eliminates data loss and backup overhead that can protect thousands of databases and provide continuous real-time backup from system memory, so data can be recovered up to the last sub-second.

All backup-related processing is offloaded to the Recovery Appliance freeing the production system resources from backup overhead. This includes time-consuming compression, backup deletion, validation, and maintenance operations.

The ability to restore a single file or even to precisely restore a database table ensures fast recovery and a shortened RTO.

A summary of major differences between the Oracle Recovery Appliance and Data Domain is shown in Appendix 1

⁴ Hybrid Columnar Compression (HCC) is a unique and effective compression option for Oracle Databases that reside on Oracle Storage.

Oracle's Recovery Appliance:

- Delivers better backup and restore performance due to higher processing power and direct connection with high speed, low latency InfiniBand connections
- Provides sub-second recovery capabilities
- Ensures much higher scalability without "forklift" upgrades
- Provides better availability with its dual-controllers clustered configuration, data integrity with checksum end-to-end error detection, and correction of silent data corruption
- Provides additional layers of data protection through replication to another remote Recovery Appliance; copying data to Oracle's StorageTek tape systems or to the Oracle Public Cloud via Oracle Storage Cloud Services

In addition to these technical advantages, the Oracle Recovery Appliance delivers several economical advantages which reduce the CapEx and OpEx:

- No additional backup server hardware and software is required
- A full backup is taken once, followed by incrementals in which only changes are stored, decreasing capacity requirements
- Directly archives backups to low-cost tape storage or the Oracle Public Cloud, offloading this work from production database servers
- Real-time recovery status and detailed recovery window metrics are readily available and displayed through the Enterprise Manager framework, saving administrative time
- Fast performance increases IT productivity and ensures that RPO and RTO SLAs are met

With its modern design and extensive co-engineering with the Oracle Database and RMAN, Oracle's Recovery Appliance provides much-needed data protection assurance for any business, small or large. In contrast EMC Data Domain is an underperforming backup appliance with limited scalability and availability. Which one should any business depend on when disaster strikes?

Josh Krischer is an expert IT advisor with over 46 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: Comparison Summary between Oracle’s Zero Data loss Recovery Appliance and EMC Data Domain

Oracle’s Zero Data Loss Recovery Appliance	EMC Data Domain
Co-engineered for deep integration with Oracle Database and RMAN	Not certified to backup Oracle Database
Ability to restore a single file or table	Not published
Ability to recover to last sub-second of transactions	If using RMAN, recover to time of last backup
High scalability in capacity and performance	No performance scalability (forklift upgrades), much lower maximum capacity
Full redundancy, fail-over cluster in every module and full subsystem	The processor is single-point-of-failure
Restore throughput equivalent to backup throughput	EMC has not published restore performance figures in several years
Practically all backup-related processing is offloaded to the Recovery Appliance, including time-consuming compression, backup deletion, validation, and maintenance operations, freeing production system resources	DD Boost “steals” cycles from the production environment
Background processes which optimize the delta store for performance. Every 7 days a background process runs to store blocks of the most recent backup in a contiguous fashion.	Weekly collection of “trash” data eats capacity, increases fragmentation Weekly trash collection cleaning has significant impact performance
Storing compressed changes only dramatically reduces the required storage capacity	Inferior performance and de-duplication rate leads to storage sprawl as more storage systems are 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 higher CapEx and OpEx