

An Oracle White Paper
January 2015

Oracle Active Data Guard vs Storage Remote Mirroring

Executive Overview

A critical objective for any enterprise is the protection of corporate assets, including data. Enterprises are equally concerned about the impact of application downtime when databases become unavailable due to outages caused by data corruptions, component, system and site failure, or human error. Increased cost, lost revenue, negative publicity, and regulatory non-compliance are just a few of the many negative consequences resulting from data loss and downtime.

Even though a bullet-proof disaster recovery (DR) solution offers the ultimate protection for enterprise data, businesses often don't place a high priority on DR due to the fact that disaster recovery infrastructure is expensive and rarely used. This leads to under-investment in DR and/or the deployment of solutions that provide inadequate protection and little assurance that they will actually work when called upon.

Oracle Data Guard and Active Data Guard fundamentally changes what enterprises should expect from a DR solution for the Oracle Database. It provides the best possible data protection and application availability and simultaneously reduces the cost of DR by utilizing standby systems to offload production workload while in a standby role. Data Guard's deep integration with Oracle Database also enables automatic failover for unplanned outages and database rolling upgrades to minimize downtime and risk when performing planned maintenance. These capabilities make Data Guard a comprehensive solution for HA and DR.

Active Data Guard also eliminates the risk inherent in generic data protection offered by storage-centric solutions such as array-based remote mirroring. Only Active Data Guard provides continuous real-time application level validation that the DR system is ready for failover if needed.

This brief is intended for I.T. Managers and Architects who are evaluating disaster recovery solutions for the Oracle Database and describes why Data Guard and Active Data Guard are preferred to traditional DR solutions based on storage technologies.

Introduction: Data Guard and Active Data Guard

Managed Standby, the precursor to Data Guard, first appeared in Oracle 7. It offered very basic archive log shipping capabilities that required complementary scripts to maintain a synchronized replica of a production database at a remote destination for DR. Data Guard was introduced as an included feature of the Oracle Database with Oracle 9i. It represented a major evolution in technology, eliminating the need for external scripts and providing complete management, monitoring, and automation software to create and maintain one or more replicas (standby databases) of the production database (primary).

Standby databases protect Oracle data from data corruptions, system and software failures, human error and disasters. Production applications can quickly switch to the standby database if the primary becomes unavailable for any reason. Data Guard with Oracle Database 10g added significant high availability (HA) features, making it a comprehensive solution for HA/DR optimized for the Oracle Database.

Active Data Guard was introduced with Oracle Database 11g Release 1 to provide important extensions to basic Data Guard functionality that further enhance data protection, availability, and return on investment (ROI) in standby systems. Active Data Guard is a separately priced database option that inherits all Data Guard capabilities while adding numerous advanced features.¹ The introduction of Oracle Multitenant Architecture in Oracle Database 12c extends all of Active Data Guard's benefits to consolidated database environments whether on premises or in the Cloud. Active Data Guard's many use-cases are described in Figure 1.

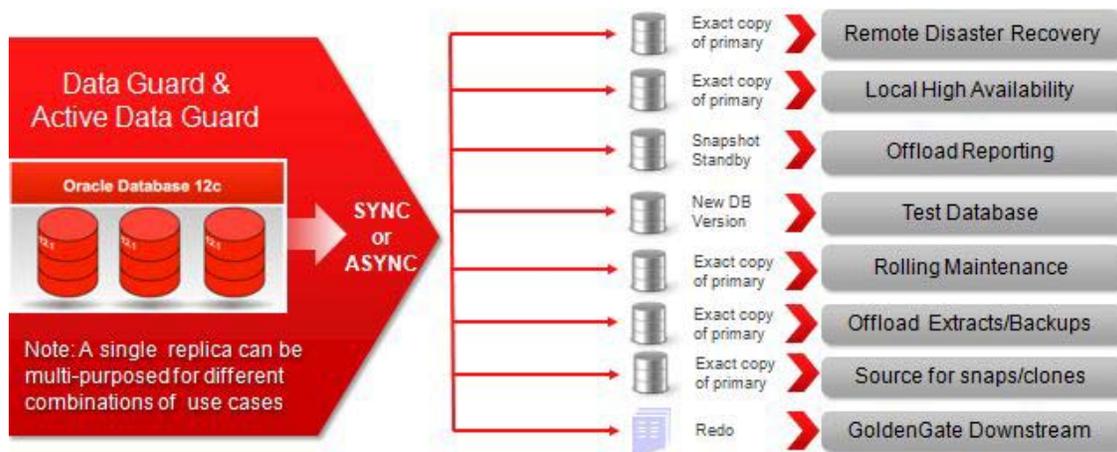


Figure 1: Data Guard and Active Data Guard Use Cases

¹ www.oracle.com/goto/dataguard

Why Data Guard Provides the Best Data Protection

Before Data Guard was introduced, storage based remote mirroring (array mirroring) was the most prevalent method of providing real-time protection for the Oracle Database. Array mirroring is a sophisticated technology promoted as a generic infrastructure solution that makes a simple promise – whatever is written to a primary volume will also be written to a mirrored volume at a remote site. Keeping this promise, however, can have disastrous consequences for data protection and availability when the data written to primary volumes is corrupt.

Data Guard and Active Data Guard are designed to provide greater data protection and availability than is possible using storage technologies alone. Many enterprises have been replacing array mirroring with Active Data Guard for their business critical databases for the following reasons:

Superior Isolation, Bandwidth Efficient

Simply stated, it is architecturally impossible for generic infrastructure solutions based upon array mirroring to provide the same level of data protection as Data Guard. Data Guard is a light-weight Oracle-aware solution that provides superior isolation between the production database (the primary) and its standby database(s). Isolation from faults that can impact the primary copy is the most critical aspect of data protection. A high level description of Data Guard architecture is provided in Figure 2.

Data Guard replicates only the information needed to recover an Oracle transaction (redo) which represents a small percentage of the total write volume generated by an Oracle database. Data Guard replicates data directly from the memory of the primary database ensuring that the standby is isolated from corruptions that can be introduced by the I/O stack.

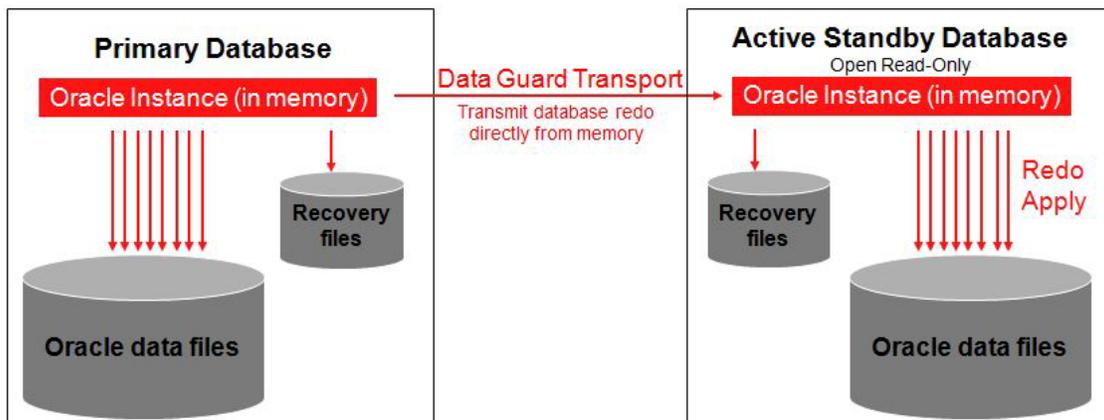


Figure 2: Data Guard/Active Data Guard Architecture

Continuous Oracle Data Validation

A Data Guard standby is an independent Oracle Database that uses media recovery to apply the changes to the standby database to maintain a synchronized physical replica of the primary. Oracle Database recovery processes perform continuous validation as changes are applied to the standby. This validation uses knowledge of Oracle redo and data block structures to check for physical data corruption, logical intra-block corruption and lost write corruption to insure the highest level of isolation between primary and standby.²

Automatic Repair

Active Data Guard makes block level corruption invisible to users with no changes to application code. Block level corruption is caused by intermittent random I/O errors that can occur independently at either primary or standby databases. Under normal operation when an Oracle Database reads a block and detects corruption it marks the block as corrupt and reports the error to the application. No subsequent read of the block is successful until the block is recovered manually - unless you are using Active Data Guard. Active Data Guard automatically repairs physical block corruption at a primary database by retrieving a good version of the block(s) from the active standby. Conversely, corrupt blocks detected by either the apply process or by read-only users on the standby database are automatically repaired using the good version from the primary database. Both HA and data protection are maintained at all times.

Lower Cost, High ROI

Data Guard and Active Data Guard significantly reduce cost and increase return on investment compared to Array Mirroring along several dimensions:

- Data Guard is zero cost from an Oracle licensing perspective; it is an included feature of Oracle Database EE. Array-based remote mirroring is a premium-priced add-on in addition to the cost of storage.
- Data Guard's reduced network volume reduces network cost
- Data Guard's integration with Oracle Database and management tools reduces administrative cost.
- Advanced features included with the Active Data Guard option provide high ROI by enabling the offload of read-only workloads to standby systems, enhanced data protection and increased HA.

² <https://support.us.oracle.com/oip/faces/secure/km/DocumentDisplay.jspx?id=1302539.1>

Array Mirroring – None of the Above

Array mirroring, in contrast, has zero Oracle awareness. The very fact that it is a generic tool for block-level replication requires it to replicate a much higher volume of data than Data Guard in order to maintain real-time protection. This is due to two characteristics inherent to storage remote mirroring:

- Array mirroring must replicate every write made to all primary volumes (writes to data files, undo and temp files, online redo log files, archive log files, flashback log files, control file, etc). Recall from earlier discussion that Data Guard only replicates a volume equal to the writes made to online redo log files – a small portion of the total write volume of a database.
- Array mirroring further increases the volume of replicated data because it must replicate the entire block (or potentially a 1 to 4 megabyte sector size), even if only a small portion of the data within the block has changed.

A high level architecture for storage remote mirroring is provided in Figure 3.

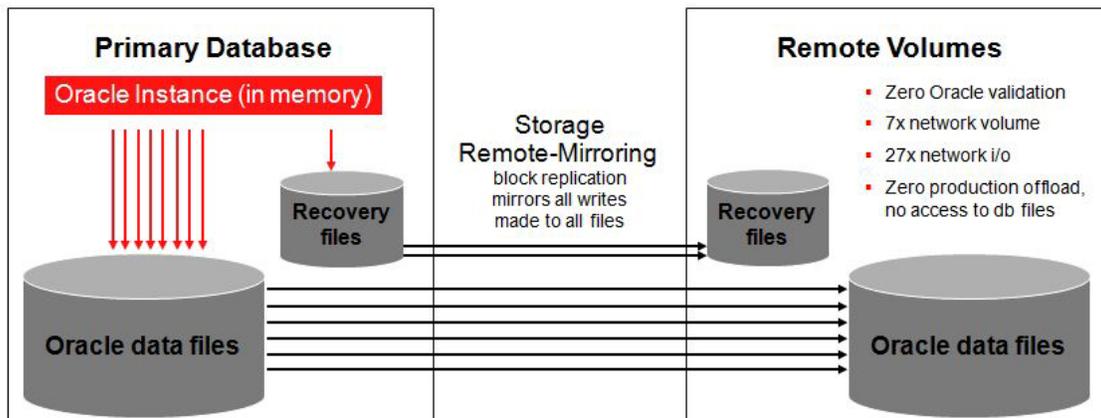


Figure 3: Storage Remote Mirroring Architecture

Array mirroring also offers zero Oracle validation – it only performs basic storage-level checksum validation – providing limited isolation between mirrored volumes. Limited isolation and zero Oracle validation virtually ensures that physical corruptions written to a primary volume as well as administrative errors that occur out-of-band of the Oracle database (e.g. accidental deletion of data files or log files) are faithfully replicated to remote volumes, making both copies unusable.

Even storage vendors acknowledge these limitations as shown in a quote from an executive at a leading vendor, "The SRDF model is a remote mirroring mode - which means that any potential data corruption would be copied faithfully and expeditiously to the other side." For this reason the vendor encourages complementary

use of additional storage snap-shot technology to perform point-in-time recovery when corruptions are mirrored to remote volumes.³

Oracle believes differently. When corruption occurs it should be isolated to the primary database. Point-in-time recovery and the accompanying downtime and data loss should be avoided in the first place. Data Guard understands native Oracle block and redo structure and can use that knowledge to perform continuous validation before changes are applied to a standby database. Unlike array-mirroring, Data Guard enables problems that emanate from a primary database to be isolated to the primary and not affect the standby. Data Guard = less downtime and less data loss.

Why Data Guard Provides Higher Availability

Array mirroring also cannot provide the same level of high availability for the Oracle Database as Data Guard and Active Data Guard.

Fast, Automatic Failover

Data Guard includes an automatic failover capability that uses the same client failover infrastructure as Oracle RAC to automatically fail application connections over to a new primary database should an outage occur. Failover is fast because Oracle is already running at the standby database. Data Guard automatic database failover is carefully constructed to ensure that recovery point objectives (zero or a maximum allowable data loss threshold) are met. It provides safeguards to protect against a split brain condition (the case where there are two independent databases that each function as primary). Data Guard includes intelligent automation to reinstate a failed primary database as a synchronized standby, quickly returning the configuration to a protected state.

Contrast the aforementioned Data Guard HA capabilities with array mirroring. There is no Oracle-integrated capability to automate database failover and application redirection to an already running database. Array mirroring also requires time-consuming reconfiguration and start-up procedures just to arrive at a state comparable to that of a Data Guard standby prior to the failure occurring. For example, when the primary database fails, volumes must be mounted before the new Oracle Database and relevant database services can be started. These additional steps increase downtime and the risk of something else going wrong which can lengthen the outage period.

³ http://chucksblog.emc.com/chucks_blog/2013/04/are-snaps-dead.html?utm_source=feedburner&utm_medium=feed&utm_campaign=Feed:+emc/Ykrh+%28Chuck%27s+Blog%29

Database Rolling Maintenance

Data Guard also provides HA during planned maintenance in the following manner:

- Upgrades and many other types of maintenance are first performed at the standby database.
- Once implemented and fully tested, production is quickly switched to the standby database already running at the new version.
- Total database downtime is limited to the time required for switchover (less than 60 seconds for an individual Oracle Database when using Oracle best practices).
- The standby can also be used to fully validate the new release before the switchover is performed without compromising data protection.

Performing maintenance in rolling fashion and seamlessly using a standby database for preproduction testing is not possible using array mirroring. See the box below for the experiences of a Data Guard user during a planned DR test for their SAP landscape.

Data Guard: HA/DR for mission critical databases

“We recently ran a test of our DR process using Data Guard for our SAP landscape. The database failover process is allocated 10 minutes of the 1-hour RTO. We successfully failed over 26 databases in our SAP landscape in 4 minutes and 36 seconds. We used another 3 minutes and 46 seconds for a validation process that is part of the design. Total time for the database failover was 8 minutes 22 seconds.”

Director of I.T. of a Multinational Agribusiness

Better Data Protection = Better HA

The same attributes of Data Guard that result in better data protection - strong isolation and continuous Oracle data validation – also result in better HA. Protecting a standby database from events that impact the primary makes it possible to quickly restore service following an outage. Data Guard avoids the negative consequences of inadequate data protection shown in several representative examples where reliance upon storage technologies led to extended outages:

- **Propagating error:** A national grocery retailer was using asynchronous array mirroring to maintain a remote copy of a mission critical Oracle RAC database. A storage administrator accidentally deleted online log files for one of the primary RAC database instances causing the database to crash and resulting in data loss. Array mirroring faithfully replicated the error by deleting the same files on the remote volumes at the business continuity site, making it impossible to quickly resume processing or to recover data in the deleted log files. This problem could have been easily avoided if they were using Data Guard; it would have isolated the error and contained the impact of the deletion to the primary site. The Data Guard standby database would have been immediately available to assume the production role and recover the data that was lost when primary log file was deleted. Shortly after this event occurred, the retailer replaced array mirroring with Data Guard.

- **Poor isolation:** The State of Virginia experienced a publicly reported 5-day outage. The state had invested in array mirroring to maintain replicated copies of numerous databases that support applications providing critical services for state residents. They experienced technical problems with the primary array that caused all databases to crash. The same problems impacted the secondary array making it impossible for the Oracle Database to start.⁴
- **Unsuccessful recovery:** American Eagle Outfitters, a popular clothing retailer, experienced a publicly reported 8-day web-site outage. They experienced disk failures followed by locally mirrored disk failures. Attempts to restore from a local backup then failed – likely due to corruptions that originated from the same issues that led to the disk failures. They finally attempted to restore using a copy at their DR site, but that copy failed as well. An interesting quote from the Computer World that article that refers to a project that had not yet been implemented highlights their perspective on Data Guard: *“I know they were supposed to have completed it with Oracle Data Guard, but apparently it must have fallen off the priority list in the past few months.”*⁵

Less Risk: You Know it’s Working

Active Data Guard enables a standby database to service read-only workloads while it is being synchronized with the primary database. Workloads running on an active standby database provide continuous user and application level validation that the standby is ready for failover if needed. This online validation is impossible to accomplish with array mirroring. The only way to validate that mirrored volumes are able to support production is to stop mirroring, mount volumes, and open the Oracle Database. This explains why users of array-mirroring often discover that they have problems at their remote site at the most inopportune time – after a primary outage has occurred.

The benefit of continuous Active Data Guard validation was observed by Paychex, a provider of payroll and human services applications to 500,000 small businesses, when read-only workloads encountered silent corruptions caused by a file system bug at the standby database.⁶

Ease of Use

Determining whether it’s easier to use Active Data Guard or array mirroring is truly a matter of perspective. When comparing the different approaches it’s critical to consider which solution is more capable of accomplishing business objectives for data protection and availability.

From one perspective, array mirroring may appear to be easier to use because storage administration staff handle the configuration of volumes, operation of the mirroring process and recovery on behalf

⁴ http://www.computerworld.com/s/article/9182719/Update_Virginia_s_IT_outage_continues_3_agencies_still_affected

⁵ http://www.computerworld.com/s/article/9182159/American_Eagle_Outfitters_learns_a_painful_service_provider_lesson

⁶ <http://www.oracle.com/technetwork/database/availability/paychex-513965.pdf>

of the database administrator (DBA). The storage administrator uses the management interface provided by the storage vendor to accomplish these tasks. Array mirroring is configured based upon volume groups – so multiple Oracle Databases can be replicated together in a single group. If array mirroring is already in place, it can be simpler to continue using established processes and practices rather than introducing something new. It also provides a single mechanism for replicating any and all data between sites – whether it is casual email, critical financial transactions or sensitive personal data.

From a different perspective, Data Guard appears easier to use because the database administrator is in complete control of a replication and recovery process that is tightly integrated with other Oracle HA capabilities (e.g. Oracle RAC, ASM, RMAN, Flashback). This control makes a standby database immediately available for the DBA to utilize for different types of recovery tasks, there is no need to coordinate activities with storage administrators. The DBA is able to use the same management interface – Oracle Enterprise Manager Cloud Control, for centralized, integrated management of the entire Oracle environment. A DBA can easily monitor critical Data Guard functions and execute database failover in seconds with a single mouse-click.

While a single Data Guard configuration consists of a primary database and one or more standby databases, multiple Data Guard configurations can be managed in concert using features provided by Cloud Control. For example, complete site failover for many Data Guard configurations can be executed with a single command using Oracle Site Guard⁷, a component of the Enterprise Manager Life Cycle Management Pack. Data Guard also supports Oracle Multitenant Architecture introduced with Oracle Database 12c. Many different databases can be easily consolidated into a single database – resulting in a single Data Guard configuration to manage.

Additional information comparing Active Data Guard to array mirroring is available on the Oracle Technology Network (OTN).⁸

What about Storage Consistency Groups?

This topic merits discussion because it is often used by storage vendors to diminish the value of Data Guard. A storage consistency group is a composite group of storage devices created with special properties to maintain dependent write consistency across all devices, and across one or more storage arrays. In an Oracle Database context, a consistency group ensures crash-consistency for Oracle Database files that span multiple volumes. Consistency groups compensate for the fact that a storage array has zero intrinsic knowledge of the application data it is attempting to protect. Consistency groups are also required in order to provide a crash consistent copy when array mirroring is used - by ensuring that writes are mirrored to the remote volumes in the same order that they were written at the source.

⁷ http://docs.oracle.com/cd/E24628_01/server.121/e52894/toc.htm

⁸ <http://www.oracle.com/technetwork/database/features/availability/dataguardremotemirroring-086151.html>

Storage vendors often expand the use-case for consistency groups by positioning them as a tool for achieving global point in time consistency, for example in cases where dependencies span multiple databases and applications. It is evident that storage consistency groups provide an important benefit by enabling data consistency at the I/O level. It is incorrect, however, to conclude that consistency groups by themselves provide application or transactional level consistency, for either a single database or a set of databases.

I/O Consistent Crash Point versus Transactional Consistency

Consistency groups provide storage level crash consistency. They guarantee that writes to a number of volumes are applied as if there were a single volume. The problem is that crash consistency does not equal transaction consistency.

When an Oracle Database crashes there will be uncommitted transactions that were written to disk which applications should not see. For example, in the process of updating 100 rows a user may not have committed the transaction when the crash occurred. Imagine a similar scenario for workload on each database participating in a given consistency group. When an outage occurs the storage brings all participating databases up at the same [consistent] crash point. However, each Oracle Database will then perform additional recovery that rolls it back to its own [different] transaction boundary than the crash point. Achieving any level of application consistency would be highly unlikely since there is no application aware mechanism, such as a transaction monitor, that coordinates database crash recovery across multiple databases in the consistency group. Additional reconciliation at a transaction level is still required across the different databases in the consistency group.

This explains why storage provides I/O consistent crash points instead of transactionally consistent crash points as required by the application. A similar outcome occurs when using consistency groups to recover a mix of databases and non-database files. Each database will be recovered to a point in time that is different from the crash consistent point that the storage system works so hard to present. This perspective is not Oracle's alone. IBM also agrees as shown in the box below:

IBM confirms that i/o consistency at a physical level is not enough:

“With Data Replication Consistency Groups, cross volume data integrity / data consistency dependent write I/O consistency is maintained on the physical level. However, data consistency at this level is very different from data consistency at a database, file system, or application level... Consistency Groups enable databases to perform a database RESTART function, which examines the DB logs and tables, backing out “in flight and In doubt ‘transactions,’ making the DB “transaction” consistent. Depending on the application, this may or may not be enough. A “business transaction” may involve a single or multiple transactions across a single or multiple applications. So, the recovery process at a “business transaction” level may require multiple levels of application(s) restarts in which each may involve a single or multiple database and/ or file system restart(s)”.

⁹ http://www-03.ibm.com/systems/resources/systems_z_pdf_Data_Replication_Consistency_Groups_-_April_2007.pdf

The most straightforward solution for maintaining transactional consistency is two-phase commit / distributed transactions. Interestingly, these protocols function without consistency groups - demonstrating that consistency groups are not useful if global point in time consistency from an application perspective is the desired objective.

How to Achieve Globally Consistent Point in Time using Oracle Technologies

Oracle Database offers several options that can be used to achieve transactional consistency after disaster recovery. Each addresses the inherent shortcomings of consistency groups and meets the true requirement, which is to achieve transactional consistency after a failure has occurred.

- Data Guard using Oracle Multitenant with Oracle Database 12c transparently addresses global point in time consistency across all pluggable databases (PDBs) that reside within a container database (CDB). A single Data Guard command or mouse-click using Enterprise Manager executes the failover of all PDBs to the standby container at a globally consistent point in time. It is important to note that each PDB is still logically independent, similar to a storage consistency group. Therefore it is not possible to guarantee global transaction consistency unless the application is designed to update multiple PDBs under a single atomic transaction. Oracle Multitenant, however, offers a very efficient mechanism for achieving global point in time consistency in an environment where all participating databases have been consolidated in the same container database.
- File system data can be placed into the Oracle Database using DBFS. Once in the database, Data Guard is able to replicate all content and failover with transactional consistency.
- Oracle Database 12c with Active Data Guard Far Sync enables zero data loss failover at any distance. Concern over point in time consistency becomes unnecessary when failover is a zero data loss event.
- Oracle has published a support note: [Recovery for Global Consistency in an Oracle Distributed Database Environment \(Doc ID 1096993.1\)](#). This support note describes how to achieve global point in time consistency across multiple independent databases (a combination of heartbeat transaction and point in time recovery [PITR] using Flashback Database - for each database participating in the group). Flashback Database does not address file system data or offer the simplicity of storage consistency group; but unlike storage consistency groups it does achieve transactional consistency when multiple databases are involved.
- Oracle Flashback Database can also be used to sync databases with file system data to the same point in time. After failover to the DR site the administrator first determines the point in time of the file system data, then uses Flashback Database to rewind the participating databases to that point in time.

Proof Points

The objective of this paper has been to clearly and objectively illustrate why enterprises gain substantial benefit from the architectural advantages of Data Guard and Active Data Guard compared to storage array mirroring when providing data protection and availability for the Oracle Database. IT managers

are often caught between infrastructure teams that seek generic solutions for perceived simplicity and application/DBA teams seeking solutions that are optimized for a specific purpose. Data Guard and Active Data Guard meets the needs for both groups by addressing the complete range of business requirements with a simple-to-use standard infrastructure that's optimized to protect the Oracle Database.

Many Oracle Database users have recognized the value of Data Guard and Active Data Guard. Detailed examples of numerous customer deployments are available on the Oracle Technology Network.¹⁰ The following are representative examples taken from presentations in the public domain:

- **JPMC:** Active Data Guard replaced storage replication. Data Guard is also used for backup, testing, reporting and upgrades.¹¹
- **Fidelity Investments:** Data Guard is one part of an integrated Oracle HA architecture for data protection and availability used to protect against hardware, site, and region failures.¹²
- **Enterprise Rent-A-Car:** Deployed a DR solution with automatic database failover and read-only offload across WAN.¹³
- **Intel:** Enabled automatic database failover with zero data loss.¹⁴
- **CERN:** Protect petabytes of costly research data, easily scale read performance, and use standby database to provide HA during planned maintenance.¹⁵
- **Metlife:** Eliminated a Shareplex reporting instance and instead directs their Peoplesoft OBIEE and Hyperion reporting to an Active Data Guard standby that is also used for DR.¹⁶
- **AmerisourceBergen:** Active Data Guard replaced storage remote-mirroring for high-volume SAP system.¹⁷

¹⁰ <http://www.oracle.com/technetwork/database/features/availability/ha-casestudies-098033.html>

¹¹ <http://www.oracle.com/technetwork/database/availability/8395-jpmc-oow2012-1967243.pptx>

¹² <http://www.oracle.com/technetwork/database/availability/fidelity-2030450.pdf>

¹³ <http://www.oracle.com/technetwork/database/features/availability/13447-enterprise-515168.pdf>

¹⁴ <http://www.oracle.com/technetwork/database/features/availability/13441-intel-515449.pdf>

¹⁵ <http://www.oracle.com/technetwork/database/availability/active-data-guard-at-cern-1914404.ppsx>

¹⁶ <http://www.oracle.com/technetwork/database/availability/metlife-513964.pdf>

¹⁷ <http://www.oracle.com/us/corporate/customers/customersearch/amerisourcebergen-1-exadata-ss-2254840.html?ssSourceSiteId=otnen>

Summary

When evaluating DR solutions for the Oracle Database it is important to focus on the main objectives:

- Having the highest degree of confidence that data is safe from problems that can impact the primary.
- Ensuring that service can be resumed within the time frame required.

Array mirroring is architecturally limited by the degree of fault isolation it can provide and application knowledge it can apply to data protection and HA. Storage-centric solution proponents often promote consistency groups as the reason to dismiss these shortcomings of array mirroring. The facts show that doing so places data and HA at risk, reduces ROI, and fails to achieve transactional consistency across multiple databases.

Data Guard, both in architecture and in practice, far exceeds the level of data protection and availability that can be offered by array mirroring. Active Data Guard provides an additional level of return on investment and HA for substantial business benefit and the additional confidence that the standby will work when needed.



Oracle Active Data Guard vs
Storage Remote Mirroring

January 2015
Oracle High Availability Product Management

Oracle Corporation
World Headquarters
500 Oracle Parkway
Redwood Shores, CA 94065
U.S.A.

Worldwide Inquiries:
Phone: +1.650.506.7000
Fax: +1.650.506.7200

oracle.com



Oracle is committed to developing practices and products that help protect the environment

Copyright © 2015, Oracle and/or its affiliates. All rights reserved. This document is provided for information purposes only and the contents hereof are subject to change without notice. This document is not warranted to be error-free, nor subject to any other warranties or conditions, whether expressed orally or implied in law, including implied warranties and conditions of merchantability or fitness for a particular purpose. We specifically disclaim any liability with respect to this document and no contractual obligations are formed either directly or indirectly by this document. This document may not be reproduced or transmitted in any form or by any means, electronic or mechanical, for any purpose, without our prior written permission.

Oracle and Java are registered trademarks of Oracle and/or its affiliates. Other names may be trademarks of their respective owners.

AMD, Opteron, the AMD logo, and the AMD Opteron logo are trademarks or registered trademarks of Advanced Micro Devices. Intel and Intel Xeon are trademarks or registered trademarks of Intel Corporation. All SPARC trademarks are used under license and are trademarks or registered trademarks of SPARC International, Inc. UNIX is a registered trademark licensed through X/Open Company, Ltd. 1010

Hardware and Software, Engineered to Work Together