ORACLE

Oracle Active Data Guard versus Storage Remote Mirroring

An Oracle Technical Brief

May 2021, Version 2.1

Executive Overview

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

While a bullet-proof disaster recovery (DR) solution is the ultimate protection for enterprise data, businesses often assign this a lesser priority due to the fact that disaster recovery infrastructure is expensive and rarely used. This leads to under-investment in DR or the deployment of solutions that provide inadequate protection and little confidence that they will actually work when called upon.

Oracle Active Data Guard fundamentally changes what enterprises should expect from a DR solution for the Oracle Database. It provides the best comprehensive data protection, availability and disaster recovery while effectively maximizing the benefits of the standby systems and return on investment (ROI) by offloading production workload and critical activities to standby systems. Data Guard's deep integration with Oracle Database also enables automatic failover for unplanned outages and database rolling upgrades that minimize downtime and risk when performing planned maintenance. This makes 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 database, block and redo change level validation and on top of that auto block repair of physical corruptions, so your DR system is ready for failover when needed.

This brief is intended for I.T. Managers, Database Administrators, 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.

Table of contents

Introduction: Active Data Guard and 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 a product with Oracle 9i and 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 failures, disasters, human error, and data corruptions. Production applications can quickly switch to the standby database if the primary becomes unavailable for any reason. Data Guard adds 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 in standby systems. Active Data Guard is a database option that inherits all Data Guard capabilities 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.

The latest versions, as of 19c make the Standby Database even more active by allowing occasional DML against the standby database.

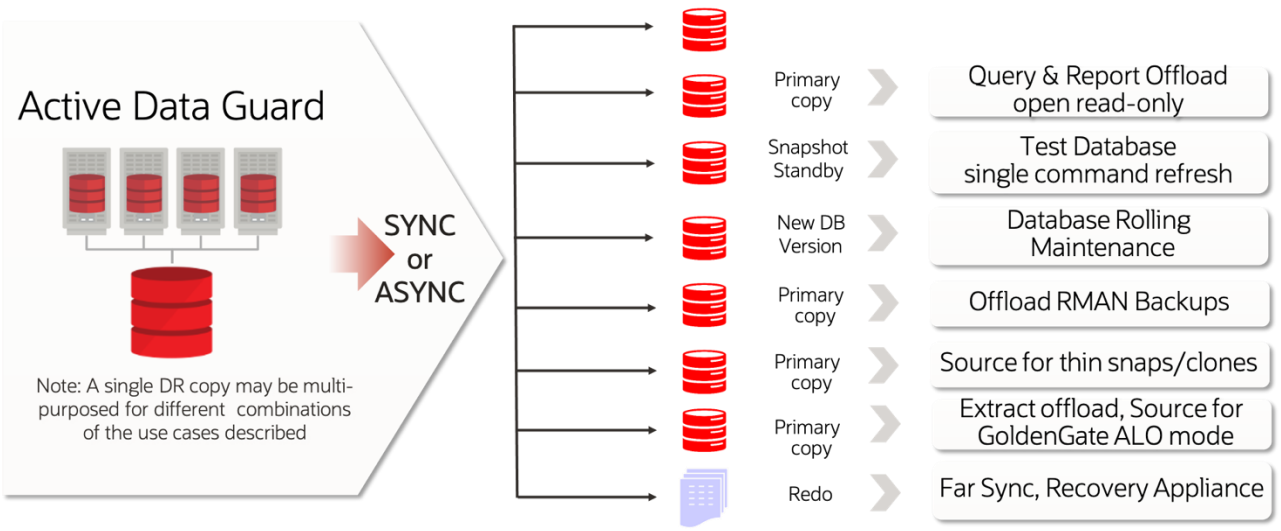

Figure 1: Data Guard and Active Data Guard Use Cases

[1] www.oracle.com/goto/dataguard

Why Data Guard Provides the Best Data Protection

Before Data Guard, storage based remote mirroring (array mirroring) was the most prevalent method of providing real-time protection for the Oracle Database. Storage mirroring is a sophisticated technology promoted as a generic infrastructure solution that makes a simple promise – whatever is written to primary storage will also be written to a mirrored copy of storage at a remote site. Keeping this promise, however, can have disastrous consequences for data protection and availability when the data written to primary storage is corrupt due to failures such as hardware or software defects that cannot be detected by checksum algorithms.

Modern cloud vendors also realize this and have stepped away from storage mirroring in favor of Oracle Data Guard to protect Oracle Relational Databases, which is a clear endorsement of the benefits Data Guard provides. This is also reflected in the service definitions regarding data loss for Oracle Databases versus other databases that lack this kind of functionality.

Data Guard and Active Data Guard are designed to provide greater data protection and availability than is possible using storage technologies alone. Most enterprises have been replacing storage mirroring for Oracle databases 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 storage 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.

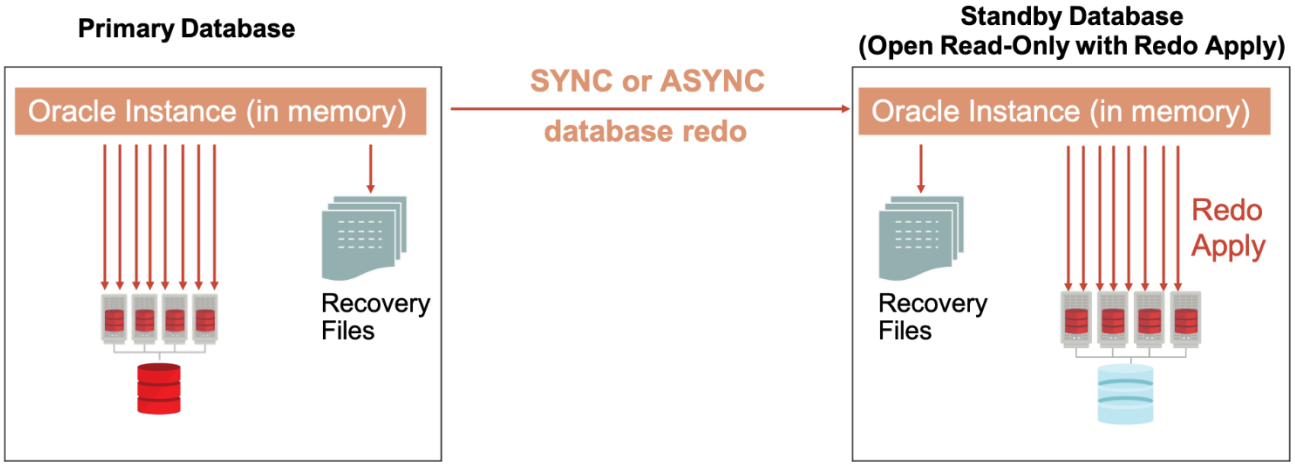

Figure 2: (Active)Data Guard Architecture

Continuous Oracle Data Validation

A Data Guard standby is an independent hardware system and 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. Block media recovery is an action which normally needs to be performed by the Database Administrator, however 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 always maintained.

Lower Cost, High ROI

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

- Data Guard is zero cost from an Oracle Licensing perspective; it is an included feature of Oracle Database Enterprise Edition. Storage-based remote mirroring is often a premium-priced add-on in addition to the cost of storage.
- Data Guard's reduced network volume reduces network bandwidth requirements significantly and can reduce the overall 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 to read-only workloads to standby systems, allowing occasional DML on those standby systems without compromising ACID, enhanced data protection and increased HA.
- Data Guard with Maximum Availability Architecture best practices can achieve very low end-to-end failover timings in seconds (low RTO) thanks to Transparent Application Continuity and with zero or near zero data loss (zero to near-zero Recovery Point Objective (RPO)). Many storage-based failover solutions will take 30 minutes or longer with a potential of data loss.

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

Remote Storage Mirroring – None of the Above

Remote Storage mirroring, in contrast to Oracle Data Guard, has zero awareness of Oracle Database. 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:

- Remote storage mirroring must replicate *every* write made to primary volumes (writes to data files, redo log files, archive log files, flashback log files, control file, TEMP files, etc.). Just to recall a key point from earlier in this paper, 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.
    - Writes to sort or work areas (TEMP) are completely useless to replicate but contribute heavily to overhead using remote storage mirroring.
    - Oracle ASM automatically initiates a Rebalance, after storage configuration changes, such as when you add, drop, or resize disks. When using ASM on top of storage replication can also face issues with rebalance. ASM will rebalance blocks without the blocks being changed, so this can lead to huge impacts if a rebalance occurs. All of those blocks are changed from the storage perspective, even though they are not changed from a database perspective.
- Remote storage 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 it has changed.

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

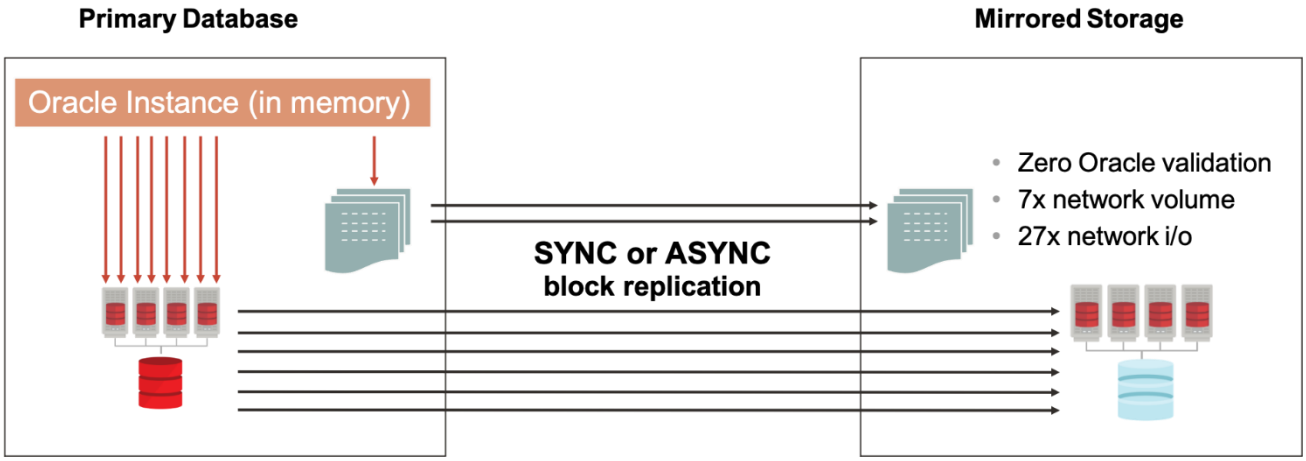

Figure 3: Storage Remote Mirroring Architecture

Storage mirroring also performs zero Oracle validation – it only performs the most rudimentary physical checksum validation providing limited isolation between mirrored storage copies. Limited isolation and zero Oracle validation virtually ensure that physical corruptions written to primary storage 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 storage, making both copies unusable.

Even storage vendors themselves 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 snapshot technology to perform point-in-time recovery when corruptions are mirrored to remote volumes.

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 remote storage mirroring, Data Guard enables problems that emanate from a primary database to be isolated to the primary and not affect the standby. Data Guard involves less downtime and less data loss.

Why Data Guard Provides Better Availability (HA)

Remote storage mirroring cannot provide the same level of high availability as Data Guard and Active Data Guard, at least not for Oracle Databases.

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 remote storage mirroring. There is no Oracle-integrated capability to automate database failover and application redirection to an already running database. Remote storage 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.

Database Rolling Maintenance

Data Guard also provides HA while performing planned maintenance.

- 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 at the new version.
- Total downtime is limited to the time required for switchover.
- The standby can also be used to fully validate the new release before the switchover is performed without compromising data protection. This is not possible using array mirroring.

Performing maintenance in a rolling fashion and seamlessly using a standby database for preproduction testing is not possible using array mirroring.

Better Data Protection means 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 storage array mirroring to maintain a remote copy of a mission critical Oracle RAC database. A storage administrator accidently 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 same file deletion to the remote volumes at the business continuity site making it impossible to quickly resume processing or to recover data from the deleted log file. This could have been avoided if they had been using Data Guard which have ensured the impact of the deletion only impacted the primary site. The Data Guard standby database would have been immediately available to assume the production role and recover the data lost when primary log file was deleted. Not long after this event occurred, array mirroring was replaced by Data Guard.
- **Poor Isolation**: The State of Virginia experienced a publicly reported 5-day outage. They had invested in remote storage mirroring to maintain replicated copies of numerous databases supporting applications providing critical services for state residents. They experienced technical problems with the primary storage that caused all databases to crash. The same problems impacted their secondary storage making it impossible for the standby Oracle Databases to come online. The problems at the second site were not discovered until after the first failure occurred which is not uncommon in these types of situations. Mirroring was disabled, and they attempted to start their Oracle databases. Of course, this failed and resulted in the need to rebuild the databases.
- **Unsuccessful recovery**: American Eagle Outfitters, a popular clothing retailer, experienced a publicly reported 8-day web-site outage. They experienced storage failures followed by a locally mirrored storage failure. 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. One interesting quote from the article, *"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."* Less Risk: You Know it's Working

Active Data Guard further extends Active Data Guard's implicit HA advantage. Read-only workload running on an Active standby database provides continuous user and application-level validation that the standby is ready for failover if needed. This is impossible to do with storage mirroring where the only way to validate if the mirrored storage can support production is to stop the remote mirroring process and open the Oracle Database. This limits storage mirroring users to static, point-in-time validation that also reduces protection for the duration of the test. This explains why storage mirroring users often discover that they have problems at their remote site at the most inopportune time – after a primary outage has occurred.

Ease of Use

Determining whether it's easier to use Active Data Guard or remote storage 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, remote storage mirroring can be perceived as easier to use because storage administration handles configuration of volumes and operation of the mirroring process and recovery on behalf of the database administrator (DBA). The storage administrator uses the management interface provided by the storage vendor or provider to accomplish these tasks. Storage mirroring is often configured based upon storage volume groups – so multiple Oracle Databases can be replicated together in a single group. If storage 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 can be perceived as 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. The DBA can use the same management interface – Oracle Enterprise Manager Cloud Control, that provides integrated monitoring, diagnostics and management of their 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 multiple Oracle Databases configured with Data Guard can be orchestrated using a single command using Oracle Site Guard[3], a component of the Oracle Enterprise Manager Database Life Cycle Management Pack. Data Guard also supports Oracle Multitenant Architecture which was introduced in Oracle Database 12c and has now became the only supported Database architecture as of the Oracle Database 21c release. Using Multitenant architecture, multiple databases can easily be consolidated into a single database – resulting in a single Data Guard configuration to manage "many databases as one".

Weighing each of the above perspectives, one should be influenced heavily by an assessment of which solution is more capable of accomplishing the business objectives for their application from the perspective of data protection and availability to ensure zero data loss and the lowest possible downtime.

What about Storage Consistency Groups?

The storage consistency group feature made remote storage mirroring feasible for remotely mirroring Oracle Databases but has proven to be a less than ideal solution. 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. Remote storage mirroring of Oracle databases would not be possible because remote mirrors would be corrupt (i.e. not consistent). 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 changes are written to the remote volumes in the same order that they were written at the source.

Storage vendors often expand the use-case for consistency groups as a tool for achieving global point in time consistency in cases where there are dependencies that span multiple databases and applications. It is obvious

[3] https://docs.oracle.com/en/enterprise-manager/cloud-control/enterprise-manager-cloud-control/13.4/guard/introduction-oracle-site-guard.html#GUID-8C2FADF5-11F1-420C-B968-05714F276A1F

that storage consistency groups do indeed provide an important value. It is not correct, however, to believe 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 multiple volumes appear to be to a single volume. The problem is that crash consistency does not equal transaction consistency.

When an Oracle Database crashes, incomplete (uncommitted) changes will often be written to disk which applications which should in turn be invisible to the application itself. Normal Oracle Database crash recovery takes care of this. For example, in the process of updating 100 rows a user may not have committed the transaction when the crash occurs. 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 in relation to 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 transactional 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.

The most straightforward solution for 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.

- File system data can be placed into the Oracle Database using DBFS. Once in the database, Data Guard can replicate all content and failover with transactional consistency.
- Oracle Database with Active Data Guard Far Sync enables zero data loss failover at any distance. Concern for point in time consistency becomes a moot point 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 PITR using Flashback Database - for each database participating in the group). There are drawbacks - it doesn't address file system data, nor does it have the simplicity of storage consistency groups, 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 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.[4]

[4] https://blogs.oracle.com/maa/why-is-flashback-often-better-than-backups

Summary

The objective of this paper has been to clearly and objectively illustrate why enterprises gain substantial benefit from Data Guard and Active Data Guard's architectural advantages to provide superior data protection and availability for the Oracle Database in comparison to array mirroring. IT managers are often caught between infrastructure teams that seek generic solutions due to perceived simplicity and application/DBA teams seeking solutions that are optimized for a specific purpose and provide better protection. Data Guard and Active Data Guard provide a common ground to address the complete range of business requirements with a simple-to-use standard infrastructure that's optimized to protect the Oracle Database.

Many other users of Oracle Database have recognized the value of Data Guard and Active Data Guard. The details of many examples of customer deployments are available on the Oracle Technology Network.[5]

When evaluating DR solutions, 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 Database
- Ensuring that service can be resumed within the time frame required.

Remote storage 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 storage-based remote 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.

[5] http://www.oracle.com/technetwork/database/features/availability/ha-casestudies-098033.html

Connect with us

Call +**1.800.ORACLE1** or visit **oracle.com**. Outside North America, find your local office at: **oracle.com/contact**.

blogs.oracle.com facebook.com/oracle twitter.com/oracle