Oracle Maximum Availability Architecture

Oracle (Active) Data Guard 19c

Real-Time Data Protection and Availability

WHITE PAPER / MARCH 7, 2019

ORACLE

TABLE OF CONTENTS

INTRODUCTION

Successful high availability (HA) architectures prevent downtime and data loss by using redundant systems and software to eliminate single points of failure. The same principle applies to mission critical databases. Administrator error, data corruption caused by system or software faults, or complete site failures can affect the availability of a database. Even a clustered database running on multiple servers using shared storage can be exposed to single points of failure if not adequately protected.

The only way to prevent being impacted by single points of failure is to have a completely independent copy of a production database already running on a different system and ideally deployed at a second location, which can be quickly accessed if the production database becomes unavailable for any reason.

Oracle Active Data Guard is the most comprehensive solution available to eliminate single points of failure for mission critical Oracle Databases. It prevents data loss and downtime in the simplest and most economical manner by maintaining a synchronized physical replica of a production database at a remote location. If the production database is unavailable for any reason, client connections can quickly, and in some configurations transparently, failover to the synchronized replica to restore service. Active Data Guard eliminates the high cost of idle redundancy by allowing reporting applications, ad-hoc queries, and data extracts to be offloaded to read-only copies of the production database. Active Data Guard's deep integration with Oracle Database and complete focus on real-time data protection and availability avoids compromises found in storage remote mirroring or other host-based replication solutions.

This paper describes both Active Data Guard (a licensed option) and Data Guard (included in Oracle Database Enterprise Edition) in detail and is tailored to IT managers, Database Administrators and technical staff, who are evaluating different alternatives to protect against data loss and database downtime.

ORACLE ACTIVE DATA GUARD – AN OVERVIEW

Oracle (Active) Data Guard capabilities in Oracle Database 19c further enhance its strategic objective of preventing data loss, providing high availability, eliminating risk, and increasing return on investment by enabling highly functional active disaster recovery systems that are simple to deploy and manage. It accomplishes this by providing the management, monitoring, and automation software infrastructure to create and maintain one or more synchronized standby databases that protect Oracle data from failures, data corruption, human error, and disasters.

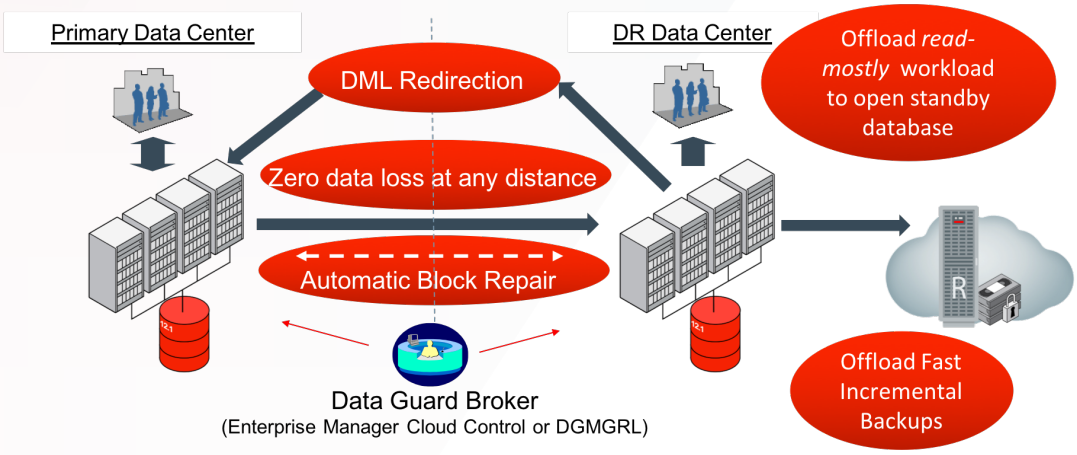

Figure 1: Oracle Active Data Guard Architecture Overview

Active Data Guard uses the simplicity of physical replication, but its deep integration with Oracle Database provides unique isolation between primary and standby databases to deliver the highest level of protection against data loss. Active Data Guard supports both synchronous (guaranteed zero data loss) and asynchronous (near-zero data loss) protection. To maintain high availability for mission critical applications, database administrators can choose either manual or automatic failover to a standby should the primary system become unavailable for any reason.

Active Data Guard is a licensed option for Oracle Database Enterprise Edition. All capabilities described in the following sections that are explicitly referred to as being 'Active Data Guard' require an Active Data Guard license. All capabilities that are explicitly referred to as 'Data Guard' are included with Oracle Enterprise Edition; no option license is required. Active Data Guard is a superset of Data Guard thus inherits all Data Guard capabilities.

One of the big advantages of Active Data Guard 19c is the better capability to offline read intensive applications to the standby. It is now possible to also issue occasional DML against the standby database, so this is now a fully functional reporting database. This leverages the return on investment as the primary database is used in a more optimal way and the resources of the DR system are used in an optimal way.

DATA GUARD 19C NEW FEATURES

ACTIVE DATA GUARD DML REDIRECT

This is an Active Data Guard Only Feature, which enables DML operations on the standby database to be redirected to the primary database to allow and accommodate for reporting applications that make infrequent writes to actively run on the Active Data Guard standby database.

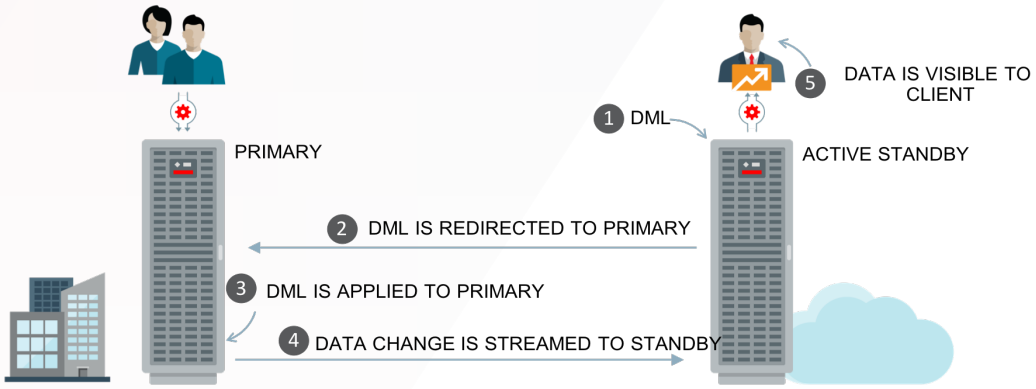

Figure 2: DML Redirect

Applying DML on the standby database can be achieved in 5 simple steps:

1. The user issues DML against the open standby database
2. This DML is redirected to the primary database
3. The DML is then applied in the primary database
4. The redo information generated the change is streamed back to the standby database
5. The application of the change-based redo information completes the DML redirect

DML Redirect can be configured for all sessions connection to the standby database, by setting the system initialization parameter “ADG_REDIRECT_DML” to TRUE. Alternatively, and to override the system parameter, “ADG_REDIRECT_DML” can be used in an alter session command to enable the DML redirect for the current session only:

ALTER SESSION ENABLE ADG_REDIRECT_DML;

In either case, DML Redirect should mainly be used for read-mostly, occasional updates applications.

FAST-START FAILOVER

Fast-Start Failover (FSFO) is a feature of the Oracle Data Guard Broker that enables an automatic failover to the standby database upon failure of the primary. FSFO can be configured in either an active or an Observer only mode. The benefit of the observer only mode is that it allows for tracking the behavior of the Data Guard Broker and see the interaction that would have occurred during normal production processing.

This allows the user to tune the FSFO properties more precisely and to discover under what circumstances an automatic failover would have occurred in their environment. This makes it easier to justify using automatic failovers in order to reduce the recovery time.

The fast-start failover target (active or observer only) can be changed dynamically and without disabling fast-start failover as well as without impacting the current environment, which allows users to test how fast-start failover will work by using the observe-only as needed.

NEW PARAMETERS FOR TUNING AUTOMATIC OUTAGE RESOLUTION WITH DATA GUARD

Oracle Data Guard has several processes on the Primary and Standby databases that handle redo transport and archiving which communicate with each other over the network. In certain failure situations, network hangs, disconnects, and disk I/O issues, these processes can hang potentially causing delays in redo transport and gap resolution. Data Guard has an internal mechanism to detect these hung processes and terminate them allowing the normal outage resolution to occur.

The following parameters allow the waits times to be tuned for a specific Data Guard configuration based on the user network and Disk I/O behavior:

- DATA_GUARD_MAX_IO_TIME
    - This parameter sets the maximum number of seconds that can elapse before a process is considered hung while performing a **regular I/O** operation in an Oracle Data Guard environment. Regular I/O operations include read, write, and status operations.
- DATA_GUARD_MAX_LONGIO_TIME.
    - This parameter sets the maximum number of seconds that can elapse before a process is considered hung while performing **a long I/O** operation in an Oracle Data Guard environment. Long I/O operations include open and close operations.

SIMPLIFIED DATABASE PARAMETER MANAGEMENT IN A BROKER CONFIGURATION

Users can now manage all Data Guard related parameter settings using the SQL*Plus ALTER SYSTEM commands or in DGMGRL with the new EDIT DATABASE ... SET PARAMETER command. Parameter changes made in the DGMGRL interface are immediately executed on the target database.

In addition, this new capability allows the user to modify a parameter on all databases in a Data Guard configuration using the ALL qualifier, eliminating the requirement to attach to each database and execute an ALTER SYSTEM command or set a Broker property for each database with multiple EDIT PROPERTY commands.

The SHOW command has also been updated to show the current setting of a parameter in the target database.

SUPPORT FOR MULTI-SHARD QUERY COORDINATORS ON SHARD CATALOG STANDBY DATABASES

Before Oracle Database 19c, only the primary shard catalog database could be used as the multi-shard query coordinator. In Oracle Database 19c, you can also enable the multi-shard query coordinator on the shard catalog's Oracle Active Data Guard standby databases.

RESTORE POINT REPLICATION

The process of flashing back a physical standby to a point in time that was captured on the primary is simplified by automatically replicating restore points from primary to the standby. These restore points are called replicated restore points. Irrespective of whether a restore point on the primary database is a guaranteed restore point or a normal restore point, the corresponding replicated restore point is always a normal restore point.

The replication of restore points depends on 2 conditions:

1. The COMPATIBLE initialization parameter for both the primary database and the standby database is set to 19.0.0 or higher
2. The primary database is open. A restore point that is created on a primary database when the primary is in mount mode is not replicated. This restriction is because the restore point information is replicated though the redo.

These restore points can be Identified by "_PRIMARY" at the end of the original name and are displayed in V$RESTORE_POINT. This view has been updated and has new column 'REPLICATED'.

When you delete a restore point on the primary, the corresponding replicated restore point on the standby is also deleted.

The managed redo process (MRP) manages the creation and maintenance of replicated restore points. If restore points are created on the primary database when MRP is not running, then these restore points are replicated to the standby database after MRP is started.

PHYSICAL STANDBY RECOVERY

When flashback or point-in-time recovery is performed on the primary database, a standby that is in mounted mode can automatically follow the same recovery procedure performed on the primary.

This means that when the standby database was in mount mode on time of the recovery operation of the Primary database, that no user intervention is needed.

When the Standby database was open. It is necessary to restart the standby database in mount mode and restart the recovery. This recovery will automatically flashback the standby database when necessary, restart itself and follow the Primary database. For this to succeed you will need to set the parameter DB_FLASHBACK_RETENTION_TARGET to a sufficiently high value so the standby database can perform these operations.

HOW DATA GUARD SYNCHRONIZES STANDBY DATABASE(S)

A Data Guard configuration includes a production database referred to as the primary database, and up to 30 directly connected replicas referred to as standby databases. Primary and standby databases connect over TCP/IP using Oracle Net Services. There are no restrictions on where the databases are physically located provided they can communicate with each other. A standby database is created from a backup of the primary database without requiring any downtime of the Production application or database. Once a standby database has been created and configured, Data Guard automatically synchronizes the primary database and the standby database by transmitting the primary database redo - the change vector information used by every Oracle Database to protect transactions – as it is generated at the Primary database and applying it to the standby database.

REDO TRANSPORT SERVICE

Data Guard redo transport services handle all aspects of transmitting redo from a primary to a standby databases(s). As users commit transactions at a primary database, redo records are generated and written to a local online log file. Data Guard transport services simultaneously transmit the same redo directly from the primary database log buffer (memory allocated within system global area) to the standby database(s) where it is written to a standby redo log file. Data Guard redo transport is very efficient for the following reasons:

- Data Guard's direct transmission from memory avoids disk I/O overhead on a primary database. This is different from how other host-based replication solutions increase I/O on a primary database by reading data from disk and writing captured data back to disk in special-purpose files utilized by their replication processes.
- Data Guard transmits only database redo. This is in stark contrast to storage remote-mirroring which must transmit every changed block of every file in order to maintain real-time synchronization. Oracle tests have shown that storage remote-mirroring transmits up to 7 times more network volume, and 27 times more network I/O operations than Data Guard.
- Data Guard physical standby also avoids the I/O overhead of supplemental logging at the primary database required by logical replication solutions. The advantages of physical replication in minimizing I/O impact also extend to the standby database where, unlike logical replication, the Data Guard apply process does not generate local redo that must be written and archived to disk

Data Guard offers two choices of transport services: synchronous and asynchronous

SYNCHRONOUS REDO TRANSPORT

Synchronous redo transport requires a primary database to wait for confirmation from the standby that redo has been received and written to disk (a standby redo log file) before commit success is signaled to the application. Synchronous transport combined with the deep understanding of transaction semantics by Data Guard apply services provides a guarantee of zero data loss if the primary database suddenly fails.

Although there is no physical limit to the distance between primary and standby sites, there is a practical limit to the distance that can be supported. As distance increases, the amount of time that the primary must wait to receive standby acknowledgement also increases, directly impacting application response time and throughput. There were two new synchronous transport options implemented in Oracle Database 12c Release 1 designed to address this performance concern:

- **Fast Sync** provides an easy way of improving performance in synchronous zero data loss configurations. Fast Sync allows a standby to acknowledge the primary database as soon as it receives redo in memory, without waiting for disk I/O to a standby redo log file (SYNC NOAFFIRM). This reduces the impact of synchronous transport on primary database performance by shortening the total round-trip time between primary and standby. Fast Sync can introduce a very small exposure to data loss should simultaneous failures impact both primary and standby databases before the standby I/O completes. The time interval, however, is so brief (both failures must occur within milliseconds of each other) and the circumstances so unique that there is a very low likelihood that this would occur. Fast Sync is included with Data Guard
- **Far Sync** enables a zero data loss failover to a remote standby database even if it is located thousands of miles away, without affecting primary database performance or materially increasing cost or complexity. Far Sync is included with Active Data Guard (see the Active Data Guard section of this paper for more details).

ASYNCHRONOUS REDO TRANSPORT

Asynchronous redo transport avoids any impact to primary database performance by acknowledging commit success to the application as soon as the local log-file write is complete; it never waits for the standby database to acknowledge receipt. This performance benefit comes with the potential for a small amount of data loss because there can be no guarantee that at any moment in time all redo for committed transactions has been received by the standby.

DATA GUARD TRANSPORT AND MULTI-STANDBY CONFIGURATIONS

Data Guard transport and multi-standby configurations avoids any impact to primary database performance by acknowledging commit success to the application as soon as the local log-file write is complete; it never waits for the standby database to acknowledge receipt. This performance benefit comes with the potential for a small amount of data loss because there can be no guarantee that at any moment in time all redo for committed transactions has been received by the standby.

A multi-standby configuration having both a local and remote standby databases provides the following benefits:

- Best data protection. The close proximity of the local Data Guard standby enables zero data loss failover with minimal impact to database performance. Data Guard Fast-Start Failover can also be used to automatically failover to the local standby without manual intervention.
- Highest availability. Client database connections can rapidly and transparently failover to the local standby using Transparent Application Failover and Fast Connection Failover. In-flight transactions also failover transparently using Application Continuity, new with Oracle Database 12c Release 1and included with Active Data Guard or Oracle RAC.

- Simple operation with continuous data protection. Following a failover to the local standby, the remote standby database automatically recognizes that failover has occurred and begins receiving redo from the new primary database - maintaining DR protection at all times.
- Cost effective and flexible. While always ready to serve as the Production database in case of a failure, the standby databases can be multi-purposed to function as a test system using Data Guard Snapshot Standby. In addition, they can be used offload read-only workloads from the primary database, offload fast incremental backups, or to perform database rolling upgrades using Active Data Guard.

AUTOMATIC GAP RESOLUTION

In cases where primary and standby databases become disconnected (network failures or standby server failures) redo stops being shipped to that standby database. The primary database continues to process transactions and accumulate a backlog of redo until a new connection to the standby database has been established. This disconnected period is reported as an archive log gap and measured as transport lag. While in this state, Data Guard monitors the status of the disconnected standby database, detects when the connection is re-established, and automatically reconnects and resynchronizes the standby database with the primary by sending the archive log files generated during the disconnected period. Note that in Maximum Protection mode if the disconnected standby database is the last remaining synchronous redo destination then there cannot be a redo gap, as the Primary database will abort itself to guarantee zero data loss. For more detail, see Protection Modes later in this paper.

REDO APPLY SERVICES

Redo Apply services run on a physical standby database. Redo Apply reads redo records from a standby redo log file, performs Oracle validation to ensure that redo is not corrupt, and then applies those redo changes to the standby database. Redo apply functions independently of redo transport to ensure that the primary database performance and data protection (Recovery Point Objective - RPO) is not affected by apply performance at the standby database. Even in the extreme case where apply services have been stopped, Data Guard transport continues to protect primary data by transmitting redo to the standby where it is archived for later use when the apply process is restarted. The Redo Apply processes run on one node of the Physical Standby system even if there are multiple nodes in the standby cluster. Starting with Oracle Database 12c Release 2 the Redo Apply services can be spread across multiple nodes, referred to as “Multi-Instance Redo Apply”, increasing the apply rate at an almost linear rate.

CONTINUOUS ORACLE DATA VALIDATION

Data Guard uses Oracle Database processes constantly validate redo before it is applied to the standby database. Redo is completely isolated from I/O corruptions on the primary because it is shipped directly from the primary log buffer – the equivalent of a memory copy (memcpy) function across the network. Knowledge of the Oracle block format is used by the Oracle Database to enable corruption-detection checks to occur at several key interfaces during redo transport and apply to ensure both physical and logical intra-block consistency. The software code-path executed on a standby database is also fundamentally different from that of the primary - effectively isolating the standby database from firmware and software errors that can affect a primary database.

Data Guard also detects silent corruption caused by lost-writes. A lost-write occurs when an I/O subsystem acknowledges the completion of a write that did not actually occur in the persistent storage. On a subsequent block read the I/O subsystem returns the stale version of the data block which can be used to update other blocks of the database, thereby spreading corruption. Data Guard prevents this by performing lost-write validation at the standby database (offload the primary database of this overhead). Data Guard detects lost-write corruption whether it occurs at the primary or at the standby. As of Oracle Database 19c, Lost Write detection (called Shadow lost write protection) can be implemented on a Production database to detect lost writes even if a standby is not configured or is not applying redo at that time. Shadow lost write protection detects a lost write before it can result in a major data corruption. You can enable shadow lost write protection for a database, a tablespace, or a data file without requiring an Oracle Data Guard standby database. Shadow lost write protection provides fast detection and immediate response to a lost write, thus minimizing the data loss that can occur in a database due to data corruption.

PROTECTION MODES

Data Guard provides three different modes to balance cost, availability, performance, and data protection shown in Table 1. Each mode uses a specific redo transport method and defines the behavior of the Data Guard configuration if a primary database loses contact with its standby

Maximum Availability	Maximum Performance	Maximum Protection
AFFIRM	NOAFFIRM	AFFIRM
SYNC	ASYNC	SYNC

Table 1

These values are configured in the SERVICE descriptor of the LOG_ARCHIVE_DEST_N parameter for redo transport.

MANAGING A DATA GUARD CONFIGURATION

You can use SQL*Plus to manage primary and standby databases and their various interactions. Data Guard also offers a distributed management framework called the Data Guard broker, which automates and centralizes the creation, maintenance, and monitoring of a Data Guard configuration. Note that the actual creation of the standby database is performed outside the broker using one of the prescribed methods, Enterprise Manager Cloud Control, RMAN duplicate commands or by using the Database Creation Assistant (DBCA), which is new in Oracle Database 19c.

Database Administrators (DBAs) interact with the broker using either the broker's command-line interface or Oracle Enterprise Manager Cloud Control. Enterprise Manager includes wizards that further simplify the creation of a Data Guard configuration and its standby databases. Key Data Guard metrics such as apply lag, transport lag, redo rate and configuration status are displayed on both the Data Guard management page (see Figure 3) and on the consolidated HA Console. Enterprise Manager also enables automatic notification should any metric exceed pre-configured threshold values.

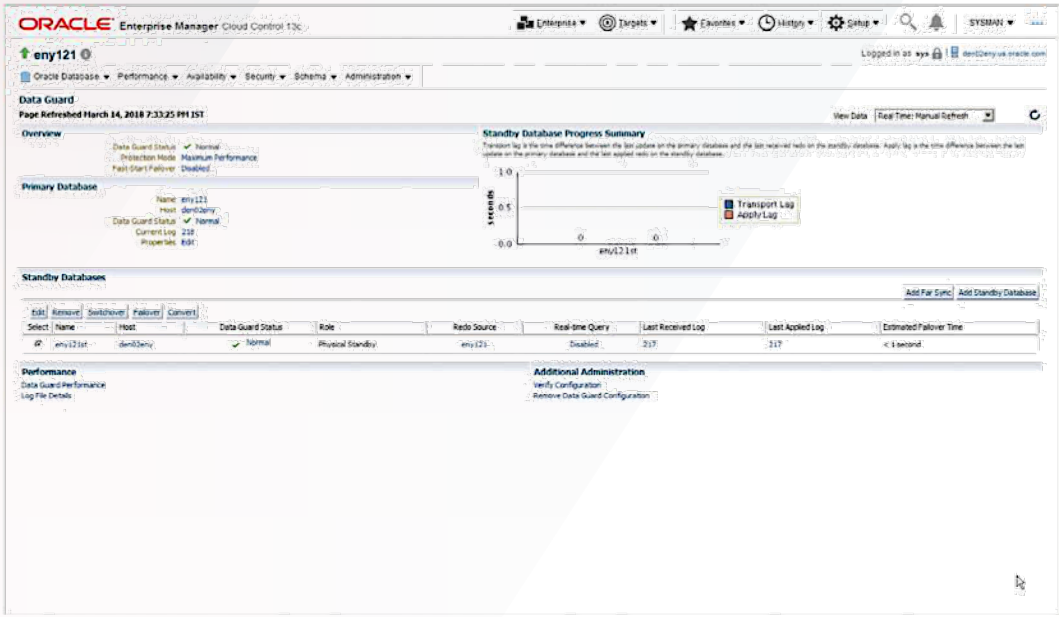

Figure 3: Data Guard Management in Enterprise Manager Cloud Control

ROLE MANAGEMENT SERVICES – SWITCHOVER AND FAILOVER

Data Guard role management services quickly transition a designated standby database to the primary role. A switchover is a planned event used to reduce downtime during planned maintenance, such as operating system or hardware upgrades, rolling upgrades of Oracle Database, and other database maintenance. Maintenance is first performed at a standby database and a switchover moves production from the primary to the standby operating at the new version. A switchover is always a zero data loss operation regardless of the transport method or protection mode used.

A failover brings a standby online as the new primary during an unplanned outage of the original primary database. A failover does not require the standby database to be restarted in order to assume the primary role. Also, as long as the original primary database can be mounted and its files are intact, it can be quickly reinstated and resynchronized as a standby database using Flashback Database; there is no need to restore from a backup.

Manual failover is initiated by the DBA using the Oracle Enterprise Manager GUI interface, the Data Guard broker's command line interface, or SQL*Plus. Optionally, Data Guard can perform automatic failover using the broker's Fast-Start Failover (FSFO).

FAST-START FAILOVER

The Data Guard Broker's Fast-Start Failover allows Data Guard to automatically failover to a previously chosen standby database without requiring manual intervention to invoke the failover. Data Guard continuously monitors the status of the configuration and initiates a failover if needed. Fast-Start Failover has built-in controls to prevent split-brain (a condition where more than a one database believes it is the primary at the same time). This simple yet tightly controlled architecture makes fast-start failover ideal when both HA and DR are required.

AUTOMATING CLIENT FAILOVER

The ability to quickly perform a database failover is only the first requirement for HA. Applications must also be able to quickly drop their connections to a failed primary database and quickly reconnect to the new primary database.

Effective client failover in a Data Guard context has three components:

- Fast database failover
- Fast start of database services on the new primary database
- Fast notification of clients and reconnection to the new primary database

Role transitions managed by the Data Guard broker can automatically transition a standby database to the primary role, start database services appropriate for the primary role, notify application clients to disconnect from the failed primary (breaking them out of TCP time-out), and direct them to the new primary database, all without manual intervention. Data Guard role change events can also be used to automate cases where a global load balancer and DNS failover are used to redirect user connections to a new middle-tier.

Application Continuity is a new capability for Oracle Database 12c Release 1 and beyond that enables transactions that are in-flight when a database failover occurs to complete without needing a rollback of the transaction and resubmitting it at the new primary database. Application Continuity is included with Active Data Guard.

Global Data Services (GDS) is a new capability for Oracle Database 12c Release 1 and beyond that extends intelligent load balancing and client failover concepts to globally distributed environments in which there are two or more failover targets that can be used to maintain availability. The multi-standby Data Guard configuration described earlier would be an example of such an environment. GDS is included with Active Data Guard.

USING DATA GUARD TO REDUCE PLANNED DOWNTIME

Data Guard can be used to reduce downtime and risk for many kinds of planned maintenance. The general approach is to first implement changes on a standby database, test, and then switchover. The production applications run unaffected on the primary database while maintenance is being performed at the standby database. Downtime is limited to the time required to switch production processing to the upgraded standby database. Specific details of the process used depend upon the type of maintenance being performed.

PLATFORM/CLOUD MIGRATION, HARDWARE AND O.S. MAINTENANCE, DATA CENTER MOVES

Data Guard Redo Apply offers some flexibility for primary and standby databases to run on systems with different operating systems or hardware architectures. See My Oracle Support Note 413484.1 for details on mixed platform combinations supported in a Data Guard configuration3. Redo Apply can be used to facilitate migration of On-Premise Production databases to Oracle's Cloud, perform technology refresh and some platform migrations with minimal downtime. Redo Apply can also be used to migrate to Automatic Storage Management and/or to move from single instance Oracle Databases to Oracle RAC, and for data center moves.

PATCH ASSURANCE USING STANDBY-FIRST PATCHING

Standby-First Patch Apply (Oracle Database 11.2.0.1 onward) enables physical standby databases with Redo Apply to support different software patch levels between a primary and standby database for the purpose of applying and validating Oracle patches in rolling fashion. Eligible patches include:

- Patch Set Update, Critical Patch Update, Patch Set Exception, and Oracle Database bundled patch
- Oracle Exadata Database Machine bundled patch, Exadata Storage Server Software patch

Refer to My Oracle Support Note 1265700.1 for more information

ACTIVE DATA GUARD

Active Data Guard is an Oracle Database Enterprise Edition option. It includes all of the Data Guard functionality described up to this point, as well as capabilities described in the following sections

REAL-TIME QUERY - PERFORMANCE AND ROI

Active Data Guard enables the offloading of read-only reporting applications, ad-hoc queries, data extracts, and so on, to an up-to-date physical standby database while also providing disaster protection. Active Data Guard is unique in having a highly parallelized apply process for best performance while also enforcing the same read consistency model at the standby as is enforced at the primary database. No other physical or logical replication solution does this.

Offloading work to an Active Data Guard standby database yields two significant benefits.

- Increased ROI in standby systems by productively using them at all times, putting an end to expensive assets that sit idle until an outage occurs.
- Eliminating risk of the unknown through continuous user-validation that an active standby is ready for failover if needed. An active standby is already working, all the time.

AUTOMATIC BLOCK REPAIR – HIGH AVAILABILITY

Block-level data loss usually results from intermittent random I/O errors, as well as memory corruptions that get written to disk. When 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 will be successful until the block is recovered manually unless you are using Active Data Guard.

Active Data Guard automatically performs block media recovery that is transparent to the application. Active Data Guard repairs physical corruption on a primary database using a good version of the block retrieved from the standby. Conversely, corrupt blocks detected on the standby database are automatically repaired using the good version from the primary database.

Physical corruption on an active standby database is also detected and automatically repaired even in cases where a block has never been changed at the primary database or read by applications running at the standby. This is done by enabling Data Guard lost-write protection at both primary and standby databases; a standard best practice for detecting silent corruption resulting from transactions that use stale data. Lost-write protection has a secondary benefit of dramatically increasing the overall level of validation for physical corruption performed at a standby database. Lost-write validation occurs at the standby database for every block that is read at the primary, whether or not the data is changed. Reading the standby version of the block in this manner triggers additional checks for physical block corruption to detect faults that occur only at the standby database and not at the primary.

FAR SYNC – ZERO DATA LOSS PROTECTION AT ANY DISTANCE

The impact that synchronous zero data loss protection has on database performance can lead to undesirable compromises. Customers with large distance between sites must compromise on protection and use asynchronous transport, accepting data loss in return for acceptable performance. Customers who absolutely require zero data loss must compromise on geo-protection and locate all sites within the same metropolitan area. Before Oracle Database 12c, the only viable option to achieve zero data loss across long distances is a 3-site architecture characterized by one or more of: expensive proprietary storage arrays, special purpose network devices, multiple Data Guard standby databases (local and remote), and complex administrative procedures.

Active Data Guard Far Sync, a new capability for Oracle Database 12c Release 1, eliminates compromise by extending zero data loss protection to any standby database located at any distance from a primary database, and doing so at minimal expense and without additional complexity.

Far Sync is a new type of Active Data Guard transport destination, referred to as a far sync instance, that receives redo synchronously from a primary database and forwards that redo asynchronously to as many as 29 remote destinations. A far sync instance is a light-weight entity that manages only a control file and log files. It requires a fraction of the CPU, memory, and I/O of a standby database. It does not have user data files, nor does it run Redo Apply. Its only purpose is to transparently offload the primary database of the overhead of transmitting redo to remote destinations. Far Sync can also save network bandwidth by offloading the primary database of overhead from redo transport compression incurred when using Oracle Advanced Compression.

Take for example an existing Data Guard configuration that uses asynchronous transport between a primary in New York, and a standby in London. Upgrade to Active Data Guard and implement zero data loss by simply deploying a far sync instance at a third location within synchronous replication distance (estimated at 30-150 miles) of New York, (see figure 3). Any server that is compatible with the primary will suffice. No proprietary storage, no special network devices, no additional licensing, and no complex management are required. If the primary fails, the same failover command used in any Data Guard configuration or automatic failover using Fast-Start Failover will quickly transition the database in London to the primary role, with zero data loss.

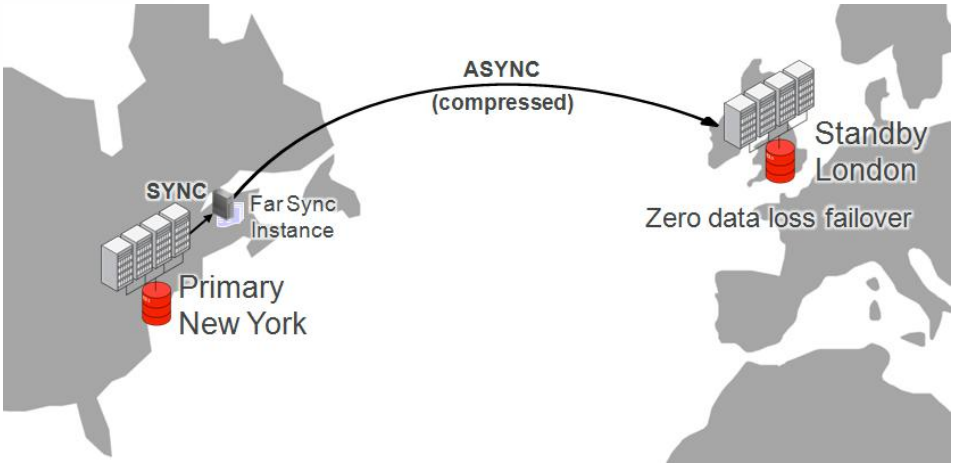

Figure 4: Active Data Guard Far Sync – Zero Data Loss Failover at Any Distance

DATABASE ROLLING UPGRADES USING ACTIVE DATA GUARD

Companies are placing increasing priority on reducing planned downtime and risk when introducing change to a mission critical production environment. Database rolling upgrades provide two advantages:

- Minimizing downtime: Database upgrades and many other types of planned maintenance that alter the physical structure of a database (other than changing the actual structure of a user table), can be implemented at the standby while production continues to run at the primary database. Once all changes have been validated, a switchover moves the production applications to the standby database, enabling the original primary to be upgraded while users run on the new version. Total planned downtime is limited to the brief time required to switch production to the standby.

- Minimizing risk: All changes are implemented and thoroughly tested at the standby database with zero risk for users running on the production version. Oracle Real Application Testing enables real application workload to be captured on the production system and replayed on the standby for the most accurate possible test result – real production workload running on a complete copy of the production database in a tightly controlled environment where it is impossible to impact production service levels. Even in cases where maintenance could otherwise be performed online at a production database, database rolling upgrades can be used by those who prefer to perform maintenance on a separate copy completely isolated from production.

Database Rolling Upgrades using Active Data Guard, a new capability for Oracle Database 12c Release 1, addresses concerns for complexity by replacing forty-plus manual steps required to perform a transient logical rolling upgrade (See Appendix B) with three PL/SQL packages that automate much of the process.

Database Rolling Upgrades using Active Data Guard can be used for version upgrades starting with the first patchset of Oracle Database 12c Release 1. This means that the manual procedure included with Data Guard and described in Appendix B of this paper must still be used for rolling upgrades from Oracle Database 11g to Oracle Database 12c, or when upgrading from the initial Oracle Database 12c release to the first patchset of Oracle Database 12c Release 1 or beyond.

This new Active Data Guard rolling upgrade capability can be used for other maintenance tasks that alter database structure. Such tasks include:

- Adding partitioning to non-partitioned tables
- Changing BasicFiles LOBs to SecureFiles LOBs
- Changing XMLType stored as CLOB to XMLtype stored as binary XML
- Compressing tables

APPLICATION CONTINUITY

Fast Application Notification (FAN) is a capability of Oracle Database that quickly delivers exception conditions to an application, but it does not report the outcome of the last transaction nor recover an in-progress request from an application perspective. As a result, outages can become visible leading to inconvenience for users and lost revenue. Users could also unintentionally make duplicate purchases and submit multiple payments for the same invoice. Developers would have no alternative other than to write and maintain custom application code to address these shortcomings, complicating support and ongoing development.

Application Continuity was a new application-independent capability in Oracle Database 12c that recovers incomplete requests from an application perspective and masks many system, communication, and hardware failures, and storage outages from the end-user. It also ensures that end-user transactions are executed no more than once. Application Continuity is included with Active Data Guard.

GLOBAL DATA SERVICES

Oracle Global Data Services (GDS) was a new capability for Oracle Database 12c that extends familiar RAC-style connect-time and run-time load balancing, service failover and workload management capabilities to a collection of replicated databases, be it within a single datacenter or across multiple datacenters. GDS is included with the Active Data Guard

CONCLUSION

Active Data Guard provides the best data protection and availability for Oracle data in the simplest most economical manner by maintaining an exact physical replica of a production database at a remote location. Although other technologies are also capable of maintaining a synchronized copy of a production database, such as storage remote-mirroring or logical replication, each makes significant compromises in one or more of the following areas when used to protect Oracle data: cost, complexity, corruption detection, automatic repair, availability, and return on investment. Active Data Guard eliminates compromise though deep integration with Oracle Database and through the simplicity achieved by complete focus on providing real-time data protection and availability for Oracle data

APPENDIX A: SUMMARY OF (ACTIVE) DATA GUARD FEATURES BY VERSION

Area	New Capability Available with Oracle Database 19c
Active Data Guard	DML operations on the Read Only Standby are redirected to the Primary database to allow for some reporting applications that make infrequent writes to run on the ADG Standby.
	You can enable the Oracle Database In-Memory Column store and Data Guard Multi-instance Redo apply at the same time on an Active Data Guard standby database.
Data Guard	You can dynamically change the fast-start failover target without disabling fast-start failover
	Without impacting your current environment, you can test how fast-start failover will work by using the observe-only mode of fast-start failover.
	The process of flashing back a physical standby to a point in time that was captured on the primary is simplified by automatically replicating restore points from primary to the standby
	When flashback or point-in-time recovery is performed on the primary database, a standby that is in mounted mode can automatically follow the same recovery procedure performed on the primary

Area	New Capability Available with Oracle Database 18c
Active Data Guard	Users on the standby will be able to continue exactly where they left off after a role change (switchover of failover) with the same performance as the Buffer Cache is now maintained over the role change operation.
	Global Temporary tables and Sequences can be created dynamically while connected to the standby removing the requirement that they be pre-created at the Primary before they can be used.
	Log in security is enhanced to allow user accounts that exceed their login failure count to be locked across the entire Data Guard environment.
	Users can now be moved during a Data Guard Role transition using the drain service capability and any users attached to the standby in read mode no longer get disconnected but maintain their state through the role change operation.

Data Guard	Lost writes can now be detected on the Primary database even if a standby is not available using shadow tablespaces.
	Primary database nologging operations can be automatically repaired on the standby on Oracle's Engineered Systems and in Oracle's Cloud.

Area	New Capability Available with Oracle Database 12c Release 2
Active Data Guard	Multi-Instance Redo Apply for Physical Standby databases. When the standby is a Real Application Cluster all the nodes in the cluster are used to apply redo thereby increasing the rate at which the standby can keep up with high workload production databases.
	The Oracle Database In-Memory feature can be used on Active Data Guard Standby databases when those standbys are on Oracle Engineered Systems or running in Oracle's Cloud.
	Tuning queries and redo apply performance can now be done completely on the standby database using Oracle's Automatic Workload Repository (AWR) and the SQL Tuning Advisor.
	Users can now be moved during a Data Guard Role transition using the drain service capability and any users attached to the standby in read mode no longer get disconnected but maintain their state through the role change operation.
	The Automatic Block Repair feature can now detect and repair almost all potential block corruptions at the Primary or Standby databases.
	Full Data Guard Broker support for DBMS_ROLLING upgrades.
	RMAN and Enterprise Manager can be used to create Far Sync instances
Data Guard	The Database Creation Assistant (DBCA) and the Enterprise Manager Command Line Interface (EMCLI) can now be used to create standby databases
	Failover an individual PDB from a standby container database to another Primary container using the Broker MIGRATE command in Multitenant databases

	Password file changes done on the primary database are now automatically propagated to standby databases. The only exception to this is far sync instances. Updated password files must still be manually copied to far sync instances because far sync instances receive redo, but do not apply it.
	Nologging operations performed at the Primary can be easily repaired on the standby using the RMAN RECOVER NONLOGGGED BLOCK command.
	Zero Data Loss failovers to Maximum Performance standby databases can now be done if there is a Production database storage failure.
	Multiple Synchronous standby destinations can be configured to only impact Production performance for the amount of time required to receive acknowledgement from the fastest standby database connection.
	Multitenant Standby databases can be a subset of the Production database on a PDB level.
	Convert Primary and Standbys databases to TDE with little Production downtime by encrypting the standby database first and using switchover to move users to the encrypted database.
Data Guard Broker	The Broker now has a block comparison tool that allows the user to compare blocks between the Production database and any or all standby databases.
	Fast_Start Failover can now support multiple failover targets and up to 3 Observers. Fast_start Failover also supports the Maximum Protection mode.
	Data Guard broker enables user configurable thresholds for apply lag and transport lag automatically signal if the potential for data loss exceeds the desired recovery point objective.
	A new Broker Property, RedoRoutes, allows for more complex redo transport configurations supporting Far Sync instances, Real-Time Cascading using the improved Alternate destination capability of Data Guard.
	Scripting with the Broker DGMGRL command line interface

APPENDIX B: TRANSIENT LOGICAL DATABASE ROLLING UPGRADE

Database rolling upgrades require the use of Data Guard SQL Apply. Oracle Database 11g introduced the Data Guard transient logical rolling upgrade process (accompanied by a set of complex manual procedures) to enable physical standby users to perform database rolling upgrades. Because complexity always increases risk, many physical standby users have understandably favored the relative simplicity of traditional upgrades. Traditional upgrades, however, result in the two things that companies wish to avoid: longer downtime and a different element of risk as described above. See Appendix B for more information.

When upgrading from an Oracle Database Release below 12.1.0.2, the transient logical database rolling upgrade process can be used with a Data Guard physical standby database to install a complete Oracle Database patch set (e.g. Oracle 11.2.0.1 to 11.2.0.3), or major release (e.g. Oracle 11.2 to 12.1) and upgrade the database with minimal downtime. The same process is also useful for customers who prefer to use an offline copy of the production database to perform various types of planned maintenance that change the logical structure of the database, validate, and then switch production to the changed version.

The transient logical process begins with a primary and physical standby database. The standby is upgraded first just as with standby-first patching, but in this case Data Guard logical replication (SQL Apply) is used on a temporary basis to replicate from the primary operating at the old version to the standby operating at the new version. Unlike Redo Apply, logical replication uses SQL to synchronize across versions and is unaffected by differences in physical redo structure that can exist between different Oracle releases.

A switchover (the only downtime) moves production to the new version running at the standby after the upgrade is complete. The original primary is then flashed back to the time when the upgrade process began and is converted to a physical standby of the new primary. The physical standby is mounted in a new Oracle home and then is upgraded and resynchronized using redo it receives from the new primary (a second catalog upgrade is not required).

ORACLE CORPORATION

Worldwide Headquarters
500 Oracle Parkway, Redwood Shores, CA 94065 USA

Worldwide Inquiries
TELE + 1.650.506.7000 + 1.800.ORACLE1
FAX + 1.650.506.7200
oracle.com

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

Integrated Cloud Applications & Platform Services

White Paper **Oracle (Active) Data Guard 19c**
March 2019
Author: [OPTIONAL]
Contributing Authors: [OPTIONAL]

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

ORACLE