ORACLE DATABASE 12c

Oracle Active Data Guard

Real-Time Data Protection and Availability

ORACLE WHITE PAPER | OCTOBER 2015

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 failure can impact the availability of a database. Even a clustered database running on multiple servers is exposed to single points of failure if not adequately protected. While a clustered database can provide excellent server HA, it is ultimately a tightly coupled system running a single database on shared storage.

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 license option) and Data Guard (included in Oracle Database Enterprise Edition) in detail. It is intended for IT managers who are evaluating different alternatives to protect against data loss and downtime, and for technical staff who are seeking a deeper understanding of how Active Data Guard functions.

Oracle Active Data Guard – An Overview

Oracle Active Data Guard[1] capabilities in Oracle Database 12c 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.

Active Data Guard, shown in Figure 1, provides a number of benefits:

- Read-only reporting, ad-hoc queries, and read-mostly applications that write to global temporary tables can be offloaded to standby databases that are also used for disaster recovery. This increases available capacity, it improves response time by isolating competing workloads, and it increases return on investment (ROI) in standby systems while using the simplicity of physical replication.
- Physical block corruptions are repaired automatically wherever they occur, at either the primary or standby, preventing any interruption in service to users and eliminating manual intervention by administrators.
- Zero Data Loss protection can be implemented in configurations where primary and standby databases are separated by thousands of miles, without impacting primary database performance or requiring added complexity or high-cost proprietary storage or network devices. There is no longer a requirement to trade performance for data protection.
- Planned downtime is minimized and the risk of introducing many types of change to a production database environment is reduced using new automation that makes it much simpler and more reliable to perform database rolling upgrades.

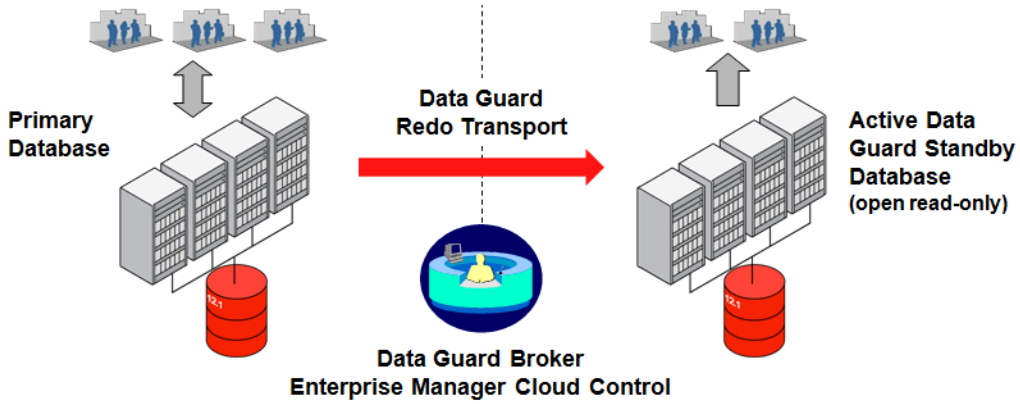

Figure 1: Active Data Guard

Active Data Guard is also a superset of Data Guard capabilities included with Oracle Enterprise Edition. This enables Active Data Guard to provide real-time data protection and availability by eliminating single points of failure. 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 corruptions, human error, and disasters.

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

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.

"Active Data Guard blurs the line between high availability and disaster recovery (DR). VocaLink is able to execute a complete site failover with approximately 8 minutes of total downtime, much faster than our previous disaster recovery architecture. Data Guard Standby-First Patching is also instrumental in reducing downtime for planned maintenance." *Martin McGeough, Database Technical Architect, VocaLink*[2]

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 initially created from a backup of the primary database. Data Guard automatically synchronizes the primary database and all standby databases by transmitting primary database redo - the information used by every Oracle Database to protect transactions - and applying it to the standby database.

Transport Services

Data Guard 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 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 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.

2 http://www.oracle.com/technetwork/database/availability/vocalink-exadata-1940109.pdf

» 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 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 are two new synchronous transport options available in Oracle Database 12c 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. 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 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 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: An increasing number of companies are using Data Guard's ability to support multiple standby databases. An example is a primary database that transmits synchronously to a local standby for HA. The local standby database in turn forwards redo to a second standby database located in a remote location for disaster recovery (DR). When using Real-Time Cascade, an Active Data Guard capability available for Oracle Database 12c, the local standby uses asynchronous transport to send redo to the remote standby database, keeping it closely synchronized with the primary database.

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 and 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. The local standby database can be multi-purposed to offload read-only workload from the primary database using Active Data Guard, to offload fast incremental backups using Active Data Guard, to function as a test system using Data Guard Snapshot Standby, or to perform database rolling upgrades.

Automatic Gap Resolution: In cases where primary and standby databases become disconnected (network failures or standby server failures), and depending upon the protection mode used, the primary database continues to process transactions and accumulate a backlog of redo that cannot be shipped to the standby until a new connection is established (reported as an archive log gap and measured as transport lag). While in this state Data Guard monitors the status of the standby database, detects when connection is re-established, and automatically reconnects and resynchronizes the standby database with the primary.

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 redo changes to the standby database. Redo apply functions independently of redo transport to insure 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 apply is restarted.

Continuous Oracle Data Validation

Data Guard uses Oracle Database processes to continuously 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 memcpy function across the network. Knowledge of Oracle block format is used by Oracle Database to enable corruption-detection checks to occur at a number of 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.

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.

TABLE 1: DATA GUARD PROTECTION MODES

Mode	Risk of data loss	Transport	If no acknowledgement from the standby database, then:
Maximum Protection	Zero data loss Double failure protection	SYNC	Signal commit success to the application only after acknowledgement is received from a standby database that redo for that transaction has been hardened to disk.
Maximum Availability	Zero data loss Single failure protection	SYNC FAST SYNC FAR SYNC	Signal commit success to the application only after acknowledgement is received from a standby database or after NET_TIMEOUT threshold period expires – whichever occurs first
Maximum Performance	Potential for minimal data loss	ASYNC	Primary never waits for standby acknowledgment to signal commit success to the application

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.

Data Guard broker includes a number of enhancements available in Oracle Database 12c, including:

- Configurable thresholds for transport lag and apply lag allow a user to specify tolerance for data loss, also referred to as recovery point objective or RPO. The broker generates a warning status if transport or apply is affected in any way that creates the potential for data loss to exceed RPO.
- A new VALIDATE DATABASE command conducts extensive validation checks to ensure that a Guard configuration is ready for a switchover or failover operations.
- Resumable Switchover: In previous releases a failed switchover would require a broker configuration to be deleted and re-created and all actions to extricate from the failed state would be made from the SQL command line. The Resumable Switchover feature allows you the option to address a failed switchover in any of the following ways:
- Resolve the problem and re-issue broker switchover – broker picks up where it left off
- Use broker to switch back to the original primary while the problem is resolved
- Use broker to switchover to another standby database in a multi-standby configuration

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. 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 2) and on the consolidated HA Console. Enterprise Manager enables automatic notification should any metric exceed pre-configured threshold values.

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 Fast-Start Failover.

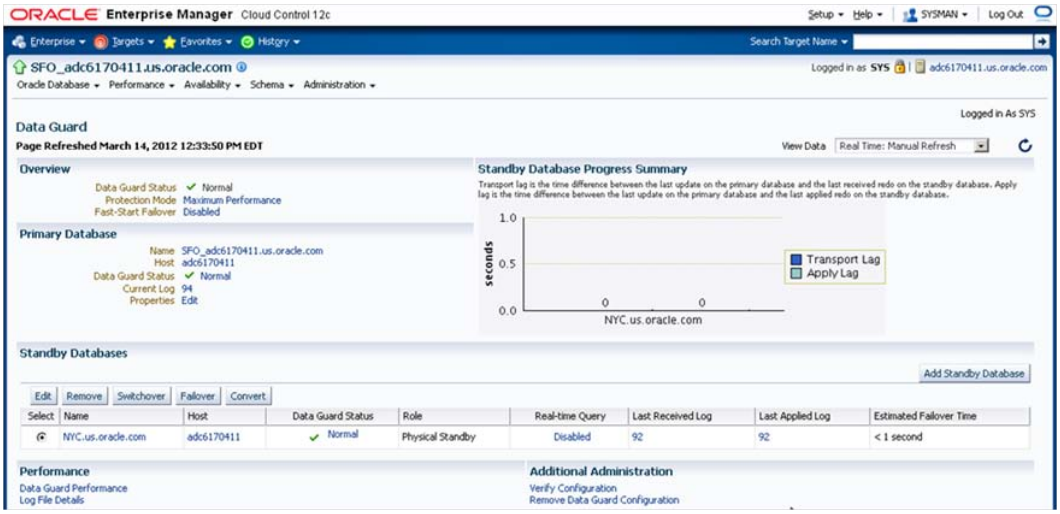

Figure 2: Data Guard Management in Enterprise Manager Cloud Control

Fast-Start Failover

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 that enables transactions that are in-flight when a database failover occurs to complete without needing to be rolled back and resubmitted at the new primary database. Application Continuity is included in Active Data Guard.

Global Data Services (GDS) is a new capability for Oracle Database 12c that extends intelligent load balancing and client failover concepts to globally distributed environments in which there are two or more failover targets that could be used to maintain availability. The multi-standby Data Guard configuration described earlier would be an example of such an environment. GDS is included in 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 database runs unaffected on the primary database while maintenance is being performed at the standby. Downtime is limited to the time required to switch production to the upgraded standby. Specific details of the process used depend upon the type of maintenance being performed.

Platform 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 configuration[3]. Redo Apply can be used to facilitate 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 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[4].

3 https://support.us.oracle.com/oip/faces/secure/km/DocumentDisplay.jspx?id=413484.1

4 https://support.us.oracle.com/oip/faces/secure/km/DocumentDisplay.jspx?id=1265700.1

Transient Logical Database Rolling Upgrade

The transient logical database rolling upgrade process uses 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).

Active Data Guard

Active Data Guard is an Oracle Database Enterprise Edition option. It includes all of the Data Guard functionality described 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.

There are also reporting applications that could use a read-only database except for the requirement that they write to global temporary tables and/or access unique sequences. Active Data Guard includes new capabilities with Oracle Database 12c to allow writes to global temp tables and access to unique sequences at an active standby, further expanding the number of reporting applications that can be offloaded from a production database.

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, 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.

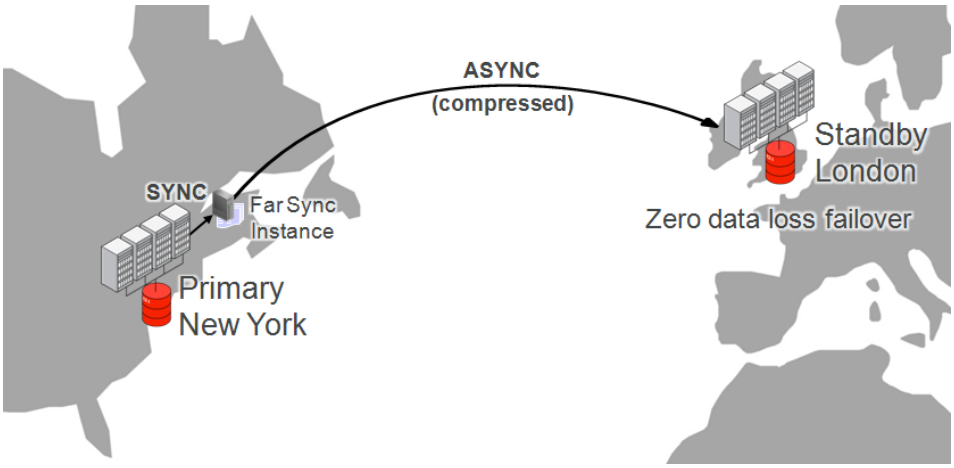

Figure 3: Active Data Guard Far Sync – Zero Data Loss Failover aT Any Distance

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 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.

Database Rolling Upgrades using Active Data Guard, a new capability for Oracle Database 12c, addresses concerns for complexity by replacing forty-plus manual steps required to perform a transient logical rolling upgrade 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. This means that the manual procedure included with Data Guard and described earlier in 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.

You may, however, begin using the new Active Data Guard capability immediately for other maintenance tasks that alter database structure beginning with Oracle Database 12c. 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 is a new application-independent capability for 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) is 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.

Related Technologies

Data Guard and Active Data Guard have close relationships with a number of technologies for high availability and data protection.

Storage Remote-Mirroring

Storage remote-mirroring (for example, EMC SRDF and Hitachi TrueCopy) is a generic approach to maintaining a remote synchronized copy of copy of data on-disk. Storage remote-mirroring plays a complementary role to Active Data Guard for replicating file system data that resides outside of the Oracle Database. Storage remote-miroring is

not best practice for replicating the Oracle Database because it lacks any knowledge of database block and redo structures required to provide the same high levels of protection, availability, functionality, and ROI as an Active Data Guard standby database.

A quick look at the architecture differences between storage remote-mirroring and Active Data Guard shows why this is the case. There are many database processes that generate I/O on an Oracle Database, including writes to data files, control files, flashback log files, online log files, archive log files, and more. While each process is designed for optimal performance and recoverability of a production database, the total I/O volume can be problematic for storage remote mirroring solutions which must mirror every write to every file to maintain real-time synchronization of a remote replica (see Figure 4). Tests show that storage remote-mirroring can transmit up to 7 times the volume, and 27 times more network I/O operations than Data Guard in order to maintain real-time data protection.

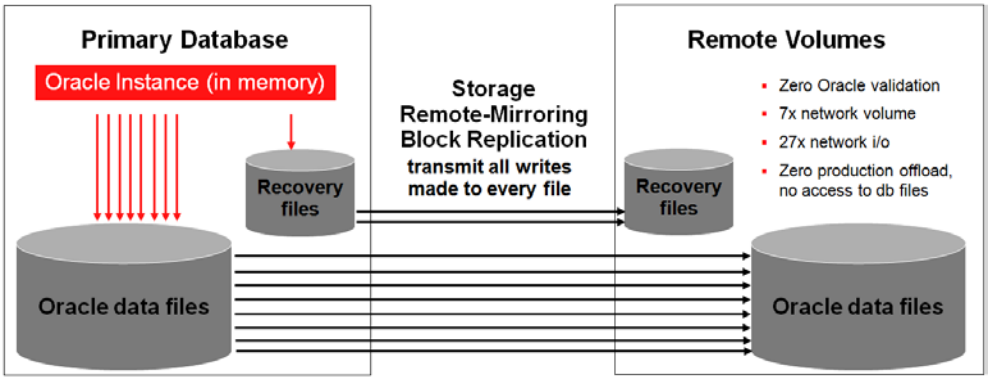

Figure 4: Storage Remote-Mirroring – Mirror Every Write to Every File for Real-Time Protection

In contrast, Data Guard uses a light-weight Oracle-aware replication process to limit bandwidth consumption to the redo volume generated by the primary database – nothing more (a volume equal to the single red stream writing to recovery files drawn in Figure 4). Data Guard transmits this redo directly from memory to eliminate disk I/O impact at the primary database and provide complete isolation from corruptions introduced in the storage layer (see Figure 5).

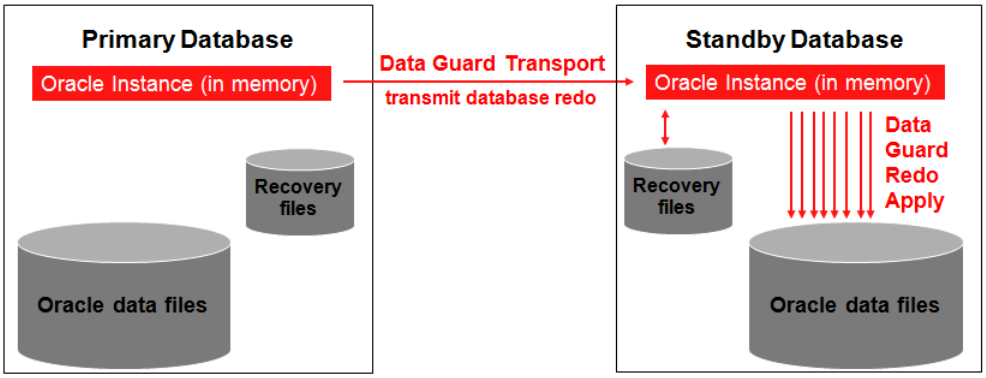

Figure 5: Data Guard – Reduced Network Consumption, Strong Isolation from Corruption

Reduced bandwidth consumption is only one benefit of using Active Data Guard. Other important benefits are:

- Strong physical isolation between primary and standby prevents propagating the impact of administrator error (for example, erroneous deletion of critical database or other files by storage administration).
- Continuous validation using Oracle-aware physical and logical block checks performed before changes are applied to the standby by a running Oracle Database instance. Validation prevents corruptions from propagating from primary to standby, and detects and automatically repairs physical block corruption that can occur independently on-disk at either a primary or standby database using Active Data Guard.
- Production offload of read-only reports, ad-hoc queries and Data Pump exports to an Active Data Guard standby for maximum ROI.
- Risk mitigation and HA. Active Data Guard eliminates the uncertainty over whether a cold start of a database using mirrored volumes will be successful. It also eliminates delay while volumes are mounted. An Active Data Guard standby is always hot and ready for production; it performs continuous Oracle Database validation as well as end-user validation using read-only workload.
- Reduction in planned downtime by performing database maintenance in a rolling fashion.

Oracle GoldenGate

Active Data Guard and Golden Gate are both strategic products within Oracle's software portfolio. Although they both generally fall into the category of replication technologies, each has a very different area of focus.

When to use Active Data Guard: Active Data Guard provides the best data protection and availability for Oracle Database in the simplest most economical manner by maintaining an exact physical replica of the production database at a remote location that is open read-only while replication is active. Use Active Data Guard when the emphasis is on simplicity, best data protection, data availability, and highest performance.

When to use GoldenGate: GoldenGate is an advanced logical replication product that supports bi-directional and multi-master replication, hub and spoke deployment and data transformation, providing customers very flexible options to address the complete range of replication requirements. GoldenGate also supports replication between a broad range of heterogeneous platforms and database management systems.

Unlike Active Data Guard, GoldenGate captures primary database changes by reading redo records from disk, transforming those records into a platform-independent trail file format, and transmitting the trail file to a target database. GoldenGate maintains a logical replica by converting the trail file into SQL and applying SQL to the target database. The target database is open read-write while synchronization occurs.

Use GoldenGate when a replica database must be open read-write while replication is active, or for advanced replication requirements beyond what is addressed by Active Data Guard.

When to use Active Data Guard and GoldenGate Together: It is important to stress that Active Data Guard and GoldenGate are not mutually exclusive, as evident in the following examples of high availability architectures that use both technologies:

- An Active Data Guard standby is used for disaster protection and database rolling upgrades for a mission critical OLTP database. GoldenGate is used to extract data from the Data Guard primary database (or from the standby database using GoldenGate ALO mode) for ETL update of an enterprise data warehouse.
- Active Data Guard is used for disaster recovery but the flexibility of GoldenGate is used to address planned maintenance activities that Data Guard is unable to support directly, such as cross-endian platform migration or application upgrades that modify back-end database objects. GoldenGate is used to perform the migration or upgrade, and once complete, a new Data Guard standby database is deployed to provide disaster protection for the new production environment.

See: Oracle Active Data Guard and Oracle GoldenGate for a more detailed discussion of requirements that are best satisfied by using either or both of these technologies[5].

Zero Data Loss Recovery Appliance

The Zero Data Loss Recovery Appliance is an innovative data protection solution that is completely integrated with RMAN and the Oracle Database. It is an integrated hardware and software appliance that includes substantial technical innovation to both revolutionize and standardize the backup and recovery process for Oracle databases across an enterprise. The recovery appliance enables an incremental forever backup strategy with any point-in-time recovery. It provides multiple levels of backup validation that makes sure recovery is successful, and it does so without any overhead on the source database. It uses Enterprise Manager for end-to-end management control to ensure that administrators can achieve their desired recovery windows.

The recovery appliance eliminates data loss by leveraging Data Guard transport services to transmit redo as soon as it is generated, eliminating the requirement to take archived log backups at a production database. The real-time nature of Data Guard transport enables the recovery appliance to restore even the most recent transaction from a database backup; the equivalent of Data Guard recovery point objectives without requiring a standby database.

The recovery appliance restores availability via a restore of a backup of the database. Data Guard and Active Data Guard, in contrast, provide, fast failover to an already running and synchronized standby database along with additional real-time validation, production offload and other advanced capabilities described in this paper.

Oracle Real Application Clusters (Oracle RAC)

Active Data Guard and Oracle RAC or RAC One Node, are complementary technologies. Oracle RAC provides the ideal HA solution to protect against server failure while also providing unique capabilities for workload management and scalability in a clustered environment. A Data Guard standby is used to protect an Oracle RAC primary database by providing strong isolation and end-to-end redundancy. This protects data and provides an additional level of HA should an event occur that causes an entire cluster to become unavailable, such as data corruption, operator error, or multiple correlated failures that can result in database or site failure. Data Guard also reduces downtime for planned maintenance that cannot be performed online or in rolling fashion across Oracle RAC instances.

Oracle Multitenant

Oracle Multitenant is a new option in Oracle Database 12c that enables a completely new architecture for database consolidation. The multitenant architecture consolidates multiple Oracle Databases (each referred to as a pluggable database, or PDB) to run under a single occurrence of Oracle Database software (referred to as a multitenant container database, or CDB). Architectural separation is enforced between each PDB (user data and metadata) and its CDB (Oracle metadata). PDBs are compatible with traditional Oracle Databases not in a CDB.

Data Guard and Active Data Guard function transparently in a multitenant architecture by providing protection at the container level. For example, a CDB that contains fifty PDBs has a single Data Guard standby and is managed as a

5 http://www.oracle.com/technetwork/database/features/availability/dataguardgoldengate-096557.html

single Data Guard configuration. A single command or mouse-click using Enterprise Manager Cloud Control will failover or switchover all PDBs to a disaster recovery site at one time.

A CDB also provides point-in-time functionality similar to storage consistency groups. In the above example, all fifty PDBs will automatically be at a globally consistent point in time after a Data Guard failover - an important characteristic when there are point-in-time dependencies that span multiple databases. Such databases might previously be placed in the same consistency group if using storage remote mirroring, or would require multiple flashback operations for consistent point-in-time recovery when using Data Guard.

Oracle Engineered Systems and Data Guard

Data Guard is the disaster recovery solution for Oracle Databases running on Oracle Engineered Systems. Data Guard physical standby is the only replication technology able to support the extreme volumes driven by Oracle Exadata Database Machine[6]. Data Guard Redo Apply has proven itself in real customer environments supporting all workloads on Exadata, including large scale data warehouses, web-scale OLTP applications, and database consolidation. Notable examples for each type of workload using Active Data Guard on Exadata include:

- A production data warehouse where Data Guard has applied changes to an Exadata standby database (11.2.0.3) at sustained rates greater than 800MB/second during heavy ETL processing.
- One of the world's most demanding OLTP applications deployed at web-scale and protected by Active Data Guard, as described by PayPal in a presentation delivered at Oracle Open World[7].
- A consolidated database environment deployed by Garmin International on Exadata with Active Data Guard that reduced cost while improving service levels[8].

Customer References

Data Guard functionality was first available with Oracle Version 7 and has continued to evolve and mature with each subsequent Oracle release. Data Guard and Active Data Guard are deployed for mission-critical applications at customer sites worldwide. A number of detailed implementation case studies are available on the Oracle Technology Network[9].

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.

6 http://www.oracle.com/technetwork/database/features/availability/maa-wp-dr-dbm-130065.pdf

7 http://www.oracle.com/technetwork/database/availability/11256-exadata-oltp-paypal-1864630.pdf

8 http://www.oracle.com/technetwork/database/availability/garmin-1667151.pdf

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

Appendix: Summary of New Capabilities with Oracle Database 12c

Area	New Capability Available with Oracle Database 12c
Active Data Guard	DML operations on global temporary tables are supported on an Active Data Guard standby database. Unique global or session sequences created at the primary database can also be accessed by reporting applications running on an Active Data Guard standby database.
	Real-Time Cascade enables a standby database that is receiving redo to forward it asynchronously as soon as it is received to another standby database (a cascaded destination) without waiting for the redo to be archived to a standby redo log file.
	Active Data Guard Far Sync is a new type of remote destination comprised of a light-weight instance having only a control file and log files that receives redo synchronously from a primary database and forwards redo asynchronously to up to 29 other remote destinations. The same failover command used for any Data Guard configuration transparently executes zero data loss failover to any standby serviced by Far Sync. Oracle Advanced Compression may also be used at the Far Sync instance to save network bandwidth while offloading the primary database host of the CPU cycles required by compression.
	Database Rolling Upgrades using Active Data Guard replaces 40+ manual steps with automation provided by three PL/SQL Packages to greatly simplify the transient logical database rolling upgrade process (temporary use of SQL Apply with physical standby to synchronize across database versions during a rolling upgrade).
	Application Continuity is an application-independent feature that attempts to recover incomplete requests from an application perspective and masks many system, communication, hardware failures, and storage outages from the end user. It also ensures that end user transactions are executed no more than once. Active Data Guard includes a license for Application Continuity.
	Global Data Services (GDS) extends the 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 or across multiple datacenters.
Data Guard	The Zero Data Loss Recovery Appliance (ZDLRA) can be configured as an asynchronous Data Guard redo transport destination to enable even the most recent transactions to be restored from a backup
	Fast Sync allows the use of Maximum Availability protection mode with synchronous transport and the NOAFFIRM attribute. This improves primary database performance for synchronous zero data loss configurations by eliminating standby redo log I/O time from total round-trip time.
	A switchover from an Oracle RAC primary database to a physical standby database no longer requires the administrator to shut down all but one primary database instance.
	You can move the location of an online data file from one physical file to another while the database is actively accessing the file. Moves on a primary database do not affect a standby database and vice versa. Active Data Guard is required to move an online data file on a standby while recovery is active.
	Data Guard and Active Data Guard support Oracle Multitenant. Data Guard operates at the level of the multitenant container database (CDB) enabling efficient disaster recovery in consolidated environments.
	SQL Apply support for additional data types: XMLType data for all storage models (if compatibility requirements are met), Oracle Spatial, Oracle Multimedia, Oracle Text, Objects and Collections (including VARRAYs and nested collections), Database File System (DBFS), XDB, Oracle SecureFiles (deduplication), and User-defined types has been added to eliminate previous obstacles to using Data Guard for database rolling upgrades. SQL Apply support for DBMS_SCHEDULER with database role-specific jobs and replication of scheduler jobs in a controlled manner specific to a database rolling upgrade context. Extended Datatype Support (EDS) also provides a mechanism for SQL Apply to support certain data types that lack native redo-based support. For example, a table with a top-level VARRAY column can be replicated using EDS.
	Data Guard now includes a specific administration privilege, SYSDG, that limits privileges to basic administrative tasks

Area	New Capability Available with Oracle Database 12c
Data Guard Broker	The broker now supports a new broker parameter RedoRoutes for simpler configuration and management of more complex multi-standby Data Guard configurations where different transport methods are utilized (SYNC or ASYNC) depending upon database role.
	A new Data Guard broker Validate Database command performs a comprehensive set of database checks prior to a role change. The checks use information available in various Data Guard views as well as the Automatic Diagnostic Repository.
	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.
	Resumable Switchover enables administrator to complete a switchover operation that had encountered an issue which prevented it from completing when the command was first issued. The administrator may either resolve the problem and resume the switchover, or revert to the original state of primary/standby, without requiring the administrator to exit or recreate the broker configuration.
	The Data Guard broker includes support for Data Guard configurations that include: Oracle Multitenant, Cascaded Standby Database, Active Data Guard Far Sync, Database Rolling Upgrade using Active Data Guard, and Global Data Services
Oracle Recovery Manager (RMAN)	Data Guard transport has always supported automatic resynchronization of primary and standby after a standby or network outage using archive log files. In cases of extended outages where needed archive log files may no longer be on disk, an RMAN fast incremental backup can be used for manual resynchronization. The manual resychronization process has been made even simpler in Oracle Database 12c by using a single new RMAN command : RECOVER DATABASE .. FROM SERVICE.

.

Oracle Corporation, World Headquarters
500 Oracle Parkway
Redwood Shores, CA 94065, USA

Worldwide Inquiries
Phone: +1.650.506.7000
Fax: +1.650.506.7200

CONNECT WITH US

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

Hardware and Software, **Engineered to Work Together**

Oracle Active Data Guard, Real-Time Data Protection and Availability
October 2015