

An Oracle Technical White Paper
September 2009

Oracle Data Guard with Oracle Database 11g Release 2

Introduction	1
Oracle Data Guard 11g - Overview	2
How Data Guard Works – Technical Details	4
Data Guard Transport Services	4
Protection Modes	5
Data Guard Apply Services	6
Automatic Gap Resolution	8
Oracle Data Validation	9
Managing a Data Guard Configuration	9
Role Management Services	10
Addressing Planned Maintenance	12
Data Guard Compared to Remote-Mirroring	13
Data Guard and Exadata	14
Data Guard and Oracle Real Application Clusters	14
Maximum Availability Architecture	14
Data Guard Customers	15
Conclusion	15
Appendix: Summary of Data Guard 11g New Features	16

Introduction

Efficient business operations, high quality customer service, conformance with government regulations, and safeguarding corporate information assets all require the highest possible level of data protection and data availability. Thus it is no surprise that data protection and data availability are among the top priorities for companies of all sizes and industries.

Backup and recovery from tape, storage remote-mirroring, or database log shipping are the traditional data protection and disaster recovery (DR) solutions. Unfortunately, these solutions are unable to reliably deliver aggressive recovery point (RPO - data protection) and recovery time (RTO - data availability) objectives. They also fail to provide an adequate return on investment due to high acquisition costs and poor utilization of standby systems that sit idle until they are called upon to assume a primary role.

In contrast, Oracle Data Guard 11g Release 2 redefines what users should expect from such solutions. Data Guard is included with Oracle Database Enterprise Edition and provides the management, monitoring, and automation software to create and maintain one or more synchronized standby databases that protect data from failures, disasters, errors, and corruptions. It can address both High Availability and Disaster Recovery requirements and is the ideal complement to Oracle Real Application Clusters.

Data Guard has the requisite knowledge of the Oracle database to provide the highest level of protection for Oracle data. Data Guard is straightforward to implement and manage. Administrators are always certain of the ability of a standby database to assume the production role – eliminating business risk at failover time. Finally, in a time when all businesses must reduce expenses, Data Guard standby databases deliver high return on investment when used for queries, reports, backups, testing, or rolling database upgrades and other maintenance, while also providing disaster protection.

"Active Data Guard 11g is a quick win! We easily dual-purposed our ten terabyte standby database for both disaster protection and secure read-only access for our public-facing eCommerce applications. We were happy to discover after much effort evaluating other alternatives, that utilizing our existing Data Guard standby database was the simplest solution to provide customers with continuous access to current information"

Sue Merrigan, Intermap Technologies

Oracle Data Guard 11g - Overview

Oracle Data Guard provides the management, monitoring, and automation software infrastructure to create and maintain one or more standby databases to protect Oracle data from failures, disasters, errors, and data corruptions. There are two types of standby databases. A physical standby uses Redo Apply to maintain a block for block, exact replica of the primary database. A logical standby uses SQL Apply and contains the same logical information as the primary database, although the physical organization and structure of the data can be different.




Figure 1 – Data Guard Overview

Administrators can choose either manual or automatic failover of production to a standby system if the primary fails in order to maintain high availability for mission critical applications. Data Guard architecture is depicted in Figure 1.

Data Guard is one of numerous integrated Oracle Database High Availability (HA) features depicted in Figure 2 that ensure business continuity by minimizing the impact of planned and unplanned downtime.

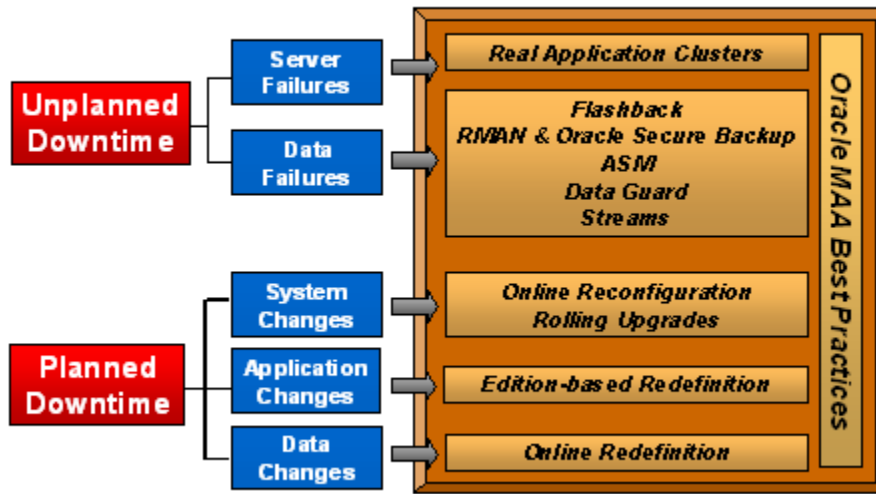


Figure 2 – Oracle Database High Availability Features

Data Guard standby databases provide high return on investment by also supporting ad-hoc queries, reporting, backups, or test activity, while providing disaster protection. Specifically:

- The Active Data Guard option, first available with Oracle Database 11g, enables a physical standby database to be used for read-only applications while simultaneously receiving updates from the primary database. Queries executed on an active standby database receive up-to-date results.
- Snapshot Standby enables a physical standby database to be open read-write for testing or any activity that requires a read-write replica of production data. A Snapshot Standby continues to receive, but not apply, updates generated by the primary. These updates are applied to the standby database automatically when the Snapshot Standby is converted back to a physical standby database. Primary data is protected at all times.
- A logical standby database has the additional flexibility of being open read-write. While data that is being maintained by SQL Apply cannot be modified, additional local tables can be added to the database, and local index structures can be created to optimize reporting, or to utilize the standby database as a data warehouse, or to transform information used to load data marts.
- Standby databases can be used to perform planned maintenance in a rolling fashion. Maintenance is first performed on a standby database. Production is switched over to the standby database when the maintenance tasks are completed. The only downtime is the time needed to effect a switchover operation. This increases availability and reduces risk when performing hardware or O.S. maintenance, site maintenance, or when

"We use Oracle Data Guard instead of direct SAN-to-SAN replication because it helps us control communications costs and ease the load on network hardware"

Craig Gibbons, NRMA Motoring & Services

upgrading to new database patchsets, full database releases, or implementing other significant database changes.

- A physical standby database, because it is an exact replica of the primary database, can also be used to offload the primary database of the overhead of performing backups.

How Data Guard Works – Technical Details

A Data Guard configuration includes a production database, referred to as the primary database, and up to 30 standby databases. Primary and standby databases connect over TCP/IP using Oracle Net Services. There are no restrictions on where the databases are located provided that they can communicate with each other. A standby database is initially created from a backup copy of the primary database. Data Guard automatically synchronizes the primary database and all of its standby databases by transmitting primary database redo - the information used by Oracle to recover transactions - and applying it to the standby database.

Data Guard Transport Services

As users commit transactions at a primary database, Oracle generates redo records and writes them to a local online log file. Data Guard transport services transmit the redo to a standby database either synchronously or asynchronously, where it is written to a standby redo log file (step one in Figure 3). Redo may be transmitted in compressed format to reduce bandwidth requirements by using the Oracle Advanced Compression Option.

Synchronous redo transport (SYNC) causes the primary database to wait for confirmation from the standby database that redo has been hardened to disk before it will acknowledge commit success to the application - providing zero data loss protection. Primary database performance is impacted by the sum of the time required for the standby redo log file I/O to complete and network round-trip time.

Data Guard 11g Release 2 is designed to reduce the impact to primary performance of synchronous transport. Redo is now transmitted to the remote standby in parallel with the local online log file I/O on the primary database, effectively eliminating standby I/O from impacting total round trip time. This enables greater geographic separation between primary and standby databases in a synchronous zero data loss configuration. On low latency networks it can reduce the impact of SYNC replication on primary database performance to near zero, making it

attractive to complement a remote ASYNC standby with a local SYNC standby for zero data loss HA protection against component and database failures (SAN failure for example).

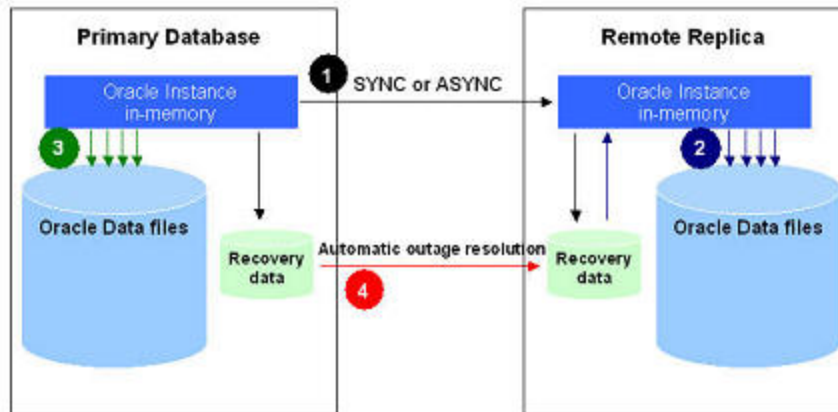


Figure 3 – Data Guard Redo Transport and Apply Services

Asynchronous redo transport (ASYNC) avoids impact to primary database performance by having the primary database acknowledge commit success to the application without waiting for acknowledgment that redo has been received by the standby database. Data Guard 11g enhancements have virtually eliminated any impact on primary database performance by shipping directly from the primary log buffer (instead of from an online redo log file), as well as improving network throughput on high latency wide area networks (WAN). The performance benefit of ASYNC, however, is accompanied by the potential for a small amount of data loss since there is no guarantee that all redo has been received by the standby database.

Protection Modes

Data Guard provides three modes of data protection to balance cost, availability, performance, and data protection. Each mode uses a specific redo transport method, and establishes rules that govern the behavior of the Data Guard configuration should the primary database ever lose contact with its standby. The following table outlines the characteristics of each mode.

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	Stall primary database until acknowledgement is received from the standby database
Maximum Availability	Zero data loss Single failure protection	SYNC	Stall primary database until acknowledgement is received or <code>NET_TIMEOUT</code> threshold period expires – then resume processing
Maximum Performance	Potential for minimal data loss	ASync	Primary never waits for standby acknowledgment

Data Guard Apply Services

Apply Services read redo from a standby redo log file, validates it, and then applies it to the standby database (step two in Figure 3) using either Redo Apply (physical standby) or SQL Apply (logical standby). Note that transport and apply services are completely independent. The status or performance of standby apply has no impact on redo transport or primary database performance. This isolation is very important. Redo transport is the chief determinate of recovery point, the potential exposure to data loss. Anything that negatively impacts transport will increase the potential for data loss. Redo transport in synchronous configurations is also the chief determinate of impact to primary database response time and throughput. Anything that negatively impacts transport in a synchronous configuration can reduce primary database throughput and increase response time. Isolation between transport and apply is designed to optimize database performance, response time, and data protection.

Redo Apply - Physical Standby Database

A physical standby database applies redo received from the primary using the Managed Recovery Process (MRP), a Data Guard aware extension of standard Oracle media recovery used by every Oracle database. A physical standby is identical to the primary database on a block-for-block basis, and thus, the database schemas, including indexes, are the same. The MRP process is highly parallel for maximum performance. Data Guard 11g performance tests conducted by Oracle achieved recovery rates of over 50MB/second for OLTP style workload, and over 100MB/second for direct path loads (see the Exadata section later in this paper for performance data specific to Exadata storage). Redo Apply is the simplest, fastest, most reliable method of maintaining a synchronized replica(s) of a primary database.

"Active Data Guard will enable MorphoTrak to reduce system costs by up to \$100,000 on our larger mission-critical systems. It is simpler to use than disk mirroring or replication. The new features of Active Data Guard 11g Release 2 guarantee that service level agreements for reporting accuracy can be met."

Aris Prassinos, MorphoTrak

Redo Apply and Active Data Guard

The Active Data Guard Option includes a number of features that extend the capabilities of Redo Apply and a physical standby database, including:

- Real-time Query enables read-only access to one or more physical standby database for queries, sorting, reporting, web-based access, etc., while Redo Apply continuously applies changes received from the production database. In cases where read-only workload can be isolated from read-write transactions, Active Data Guard can effectively double production capacity by utilizing an existing physical standby database that previously sat idle while in standby role (additional active standby databases can be added to the configuration to further scale read-only capacity without impacting read-write transactions). Active Data Guard delivers exceptional performance – it can be used for high throughput applications where it is impossible for any other replication method to keep pace with the transaction volume generated by the source database.
- Active Data Guard service level agreements (SLA) can be implemented using the session parameter, `STANDBY_MAX_DATA_DELAY`. The value for this parameter specifies a limit for the amount of time (in seconds) allowed to elapse between when changes are committed on the primary and when they can be queried on an active standby database (new with Data Guard 11g Release 2). The active standby will return an `ORA-3172` error code if the limit is exceeded. Applications can respond to this error similar to a disconnect, and redirect the query to another active standby database or to the primary database to achieve the required SLA.
- Active Data Guard 11g Release 2 enables the automatic repair of corrupt blocks. Block-level data loss usually results from intermittent, random I/O errors, as well as memory corruptions that get written to disk. When Oracle discovers a corruption it marks the block as media corrupt, writes it to disk, and typically returns an `ORA-1578` error to the application. No subsequent read of the block will be successful until the block is recovered manually. However, if the corruption occurs on a primary database that has an Active Data Guard standby, block media recovery is performed automatically, transparent to the application, using a good copy of the block from the standby database. Conversely, bad blocks on the standby database are automatically recovered using the good version from the primary database.

"Data Guard Logical Standby is an important component of a long-term strategic hardware and software platform, dramatically increasing capacity and scalability for our users. After implementing this complete solution, we achieved performance improvements of 50-95% in most batch processing operations."

David Sink, e-Rewards Market Research

SQL Apply - Logical Standby Database

A logical standby database contains the same logical information as the primary database, although the physical organization and structure of the data can be different. SQL Apply keeps a logical standby synchronized by transforming redo received from the primary database into SQL statements and then executing the SQL statements on a standby database that is open read-write. SQL Apply has some restrictions on datatypes, types of tables, and types of DDL and DML operations (see documentation for unsupported datatypes and storage attributes).

Use SQL Apply if you meet its prerequisites and:

- You wish to run reporting applications that require read-write access to the standby database. Note that data maintained by SQL Apply cannot be modified
- You wish to add tables, additional schemas, indexes, and materialized views to your standby database that do not exist on your primary database
- You will perform a database rolling upgrade from a database currently on an Oracle Database 10g release or perform other database maintenance in a rolling fashion to reduce risk and downtime. If your database version is at Oracle Database 11g or later, consider using physical standby and the transient logical standby rolling upgrade process. See the section *Addressing Planned Maintenance* for more information.

Automatic Gap Resolution

In cases where the primary and standby databases become disconnected (network failures or standby server failures), and depending upon the protection mode used, the primary database will continue to process transactions and accumulate a backlog of redo that cannot be shipped to the standby until a new network connection can be established. While in this state, Data Guard continually monitors standby database status, detects when connection is re-established, and automatically resynchronizes the standby database with the primary (step four in Figure 3). No administrative intervention is required as long as the archive logs required to resynchronize the standby database are available on-disk at the primary database. In the case of an extended outage where it is not practical to retain the required archive logs, a physical standby can be resynchronized using an RMAN fast incremental backup of the primary database.

Oracle Data Validation

One of the significant advantages of Data Guard is its ability to use Oracle processes to validate redo before it is applied to the standby database. Data Guard is a loosely coupled architecture where standby databases are kept synchronized by applying redo blocks, completely detached from possible data file corruptions that can occur at the primary database. Redo is also shipped directly from memory (system global area), and thus is completely detached from I/O corruptions on the primary. Corruption-detection checks occur at a number of key interfaces during redo transport and apply. The software code-path executed on standby database is also fundamentally different from that of the primary – effectively secluding the standby database from firmware and software errors that can impact the primary database.

Physical standby also utilizes the parameter: `DB_LOST_WRITE_PROTECT` available with Oracle Database 11g Release 1. A lost write occurs when an I/O subsystem acknowledges the completion of a write, while in fact the write did not 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 corrupting it. When the `DB_LOST_WRITE_PROTECT` initialization parameter is set, the database will record buffer cache block reads in the redo log and this information is used by Redo Apply to determine if there has been a lost write, avoiding downtime and data loss.

Managing a Data Guard Configuration

Primary and standby databases and their various interactions may be managed by using SQL*Plus. 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. Administrators may interact with the Broker using either Enterprise Manager Grid Control or the Broker's command-line interface (DGMGRL).

Enterprise Manager Grid Control 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 included in a new consolidated HA Console (see Figure 4).

Enterprise Manager enables historical trend analysis on the Data Guard metrics that it monitors - for example, how the metric's performance has been in the last 24 hrs, or last 5 days, etc. Also, through Enterprise Manager, it is possible to set up notification-alarms such that administrators may be notified in case the metric crosses the configured threshold value.

“Fast-Start Failover provides simple, fast, unattended failover for our outage management system that PPL depends upon to provide critical customer services 24 hours a day and especially during emergencies. While we have used Data Guard for disaster recovery (DR) since Oracle9i, Fast-Start Failover blurs the line between High Availability and DR – enabling us to address both requirements with a single solution”

Chris Carter, PPL Services Corporation

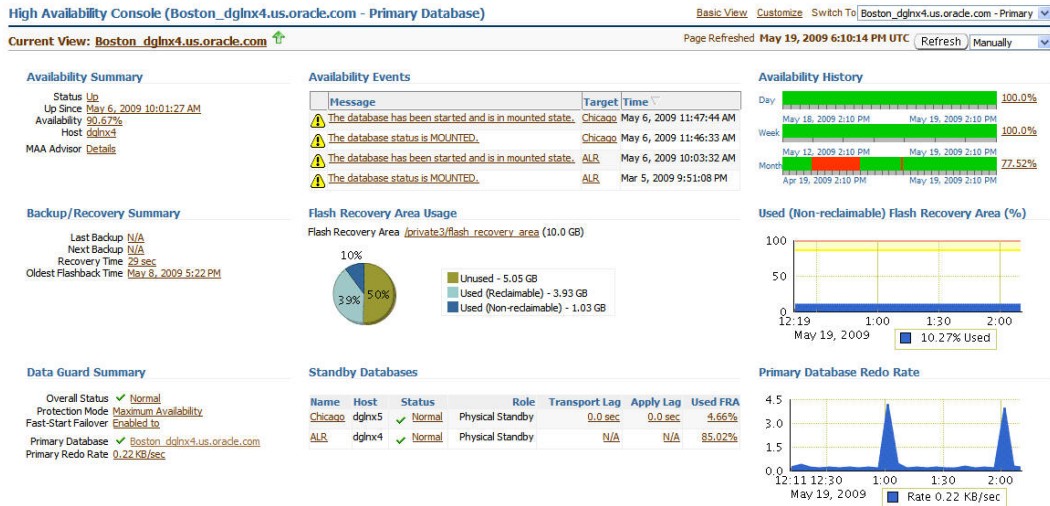


Figure 4 – Enterprise Manager Grid Control (10.2.0.5) HA Console

Role Management Services

Data Guard Role Management Services quickly transition a designated standby database to the primary role. A switchover is a planned operation used to reduce downtime during planned maintenance, such as operating system or hardware upgrades, rolling upgrades of the Oracle database, and other database maintenance. Regardless of the transport service (SYNC or ASYNC) or protection mode utilized, a switchover is always a zero data loss operation.

A failover brings a standby database online as the new primary database during an unplanned outage of the primary database. A failover operation does not require the standby database to be restarted in order to assume the primary role. Also, as long as the database files on the original primary database are intact and the database can be mounted, the original primary can be reinstated and resynchronized as a standby database for the new primary using Flashback Database – it does not have to be restored from a backup.

Manual failover is initiated by the administrator using the Oracle Enterprise Manager GUI interface, the Data Guard Broker's command line interface, or directly through SQL*Plus. Optionally, Data Guard can perform automatic failover in a very controlled manner using Fast-Start Failover.

Fast-Start Failover

Fast-Start Failover allows Data Guard to automatically fail over to a previously chosen, standby database without requiring manual intervention to invoke the failover. A Data Guard Observer process continuously monitors the status of a Fast-Start Failover configuration. If both the Observer and the standby database lose connectivity to the primary database, the Observer attempts to reconnect to the primary database for a configurable amount of time before initiating a fast-start failover. Fast-start failover is designed to ensure that out of the three fast-start failover members - the primary, the standby and the Observer - at least two members agree to major state transitions to prevent split-brain scenarios from occurring. Once the failed primary is repaired and mounted, it must establish connection with the Observer process before it can open. When it does, it will be informed that a failover has already occurred and the original primary is automatically reinstated as a standby of the new primary database. The simple, yet elegant architecture of fast-start failover makes it excellent for use when both high availability and data protection is required.

Automating Client Failover

The ability to quickly perform a database failover is only the first requirement for high availability. Applications must also be able to quickly drop their connections from 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 fast reconnection to the new primary database

In previous Oracle releases, one or more user-written database triggers were required to automate client failover, depending upon configuration. Data Guard 11g Release 2 simplifies configuration significantly by eliminating the need for user-written triggers to automate client failover. Role transitions managed by the Data Guard broker can automatically failover the database, start the appropriate services on the new primary database, disconnect clients from the failed database and redirect them to the new primary database – no manual intervention is required.

"We proved the database rolling upgrade process using transient logical standby works. We reduced application downtime when moving to a new Oracle release to just 4 minutes. Data Guard rolling upgrades achieve our SLA with time to spare "

Kenny Snell, United Parcel Service

Addressing Planned Maintenance

A Data Guard standby database can be used to reduce downtime and risk for many kinds of planned maintenance. The general approach is to implement changes on the standby database, test, and then switchover. Maintenance that does not involve differences in Oracle versions or changes to the logical structure of the database can use Redo Apply. Upgrading to new Oracle Database releases or patchsets or changing the logical structure of a database can be accomplished in rolling fashion using SQL Apply either with a logical standby database or with a physical standby database using transient logical standby.

The only downtime required for such maintenance is the time required to complete a switchover. Switchovers using Redo Apply can complete in less than 60 seconds – see the MAA best practice paper [Data Guard Switchover and Failover Best Practices](#) for more information. Switchovers using SQL Apply are even faster given that a logical standby is already open read-write. SQL Apply has a "GUARD" setting that prohibits any changes to data replicated from the primary while it is in standby role. A SQL Apply switchover formally transitions the standby to the primary role simply by changing the GUARD setting. While timings can vary from one environment to the next, database switchover using SQL Apply can complete in less than 10 seconds – see [Oracle Japan GRID Center Performance Validation: Data Guard SQL Apply on IBM Power Systems](#) for more information.

Details of the various types of planned maintenance that can be accomplished using a Data Guard standby database are described in the following sections.

System Maintenance, Technology Refresh, Select Migrations

Downtime and risk of executing certain platform migrations is minimized by using the flexibility of Redo Apply to support configurations where primary and standby systems may have different CPU architectures, operating systems (e.g. Windows and Linux), operating system binaries (32-bit/64-bit), and Oracle database binaries (32-bit/64-bit) – subject to the restrictions defined in MetaLink Note 413484.1.

Redo Apply is also used to migrate to Automatic Storage Management (ASM), from single instance Oracle Databases to Oracle RAC, from old systems to new systems when performing technology refresh, or to move from one data center to the next.

“We utilize SAN arrays and we've got bandwidth, so we've got the ability to use solutions such as remote-mirroring, but for this critical database system, we went with Data Guard. Data consistency and data integrity were the main drivers.”

David Willen, BarnesandNoble.com

Database Rolling Upgrades

Oracle Database software upgrades for major releases and patchsets (10.1.0.3 onwards) can be performed in a rolling fashion with near zero database downtime using SQL Apply. Alternatively, Data Guard 11g physical standby databases can be temporarily converted to a transient logical standby database and used to upgrade to a new database version in rolling fashion. The transient logical process is attractive because only a single catalog upgrade is required to migrate both primary and standby databases to the new Oracle release. When the upgrade process is complete, the configuration reverts to its original state of having a primary with a physical standby database.

Data Guard 11g Release 2 SQL Apply includes the ability to implement extended data type support, making it possible to support the replication of column objects (with simple or nested user-defined types), Varrays, and Oracle-supplied Spatial type SDO_GEOMETRY – when using SQL Apply for migrations and rolling database upgrades.

Database Maintenance

Data Guard 11g Release 2 SQL Apply added support for Oracle Advanced Compression (OLTP Table Compression), Oracle Secure Files, and Online Redefinition. Logical standby databases can now be used to implement these capabilities or perform other types of database maintenance without any risk of impacting production.

Data Guard Compared to Remote-Mirroring

There are many database processes that generate I/O on an active Oracle Database. The Database Writer Process continually updates data files while control file updates and local archival of online redo log files generate additional I/O. Each process is designed for optimal performance and recoverability, but they can be problematic for host or array based remote mirroring solutions - the historical alternative to Data Guard. Such solutions must replicate every write made to every file, and do so in write-order, in order to maintain real-time synchronization of a remote replica. Data Guard is an Oracle-aware process that only replicates writes made to the online redo log file. Internal tests have shown that array based remote-mirroring can transmit up to 7 times the volume, and 27 times more network I/O operations than needed by Data Guard – see [Data Guard Compared to Remote-Mirroring](#) for more information.

“Collectively, the utilization of Oracle High Availability Features and their implementation utilizing Oracle Maximum Availability Architecture (MAA) best practices has enabled Fidelity National Financial to meet service level agreements at the lowest cost.”

Charles Kim, Fidelity Information Services

Data Guard also provides the advantages of end-to-end Oracle data validation and an open standby database that can quickly assume the primary database role, things that are impossible for remote-mirroring to do given that Oracle cannot be mounted at the standby while array mirroring is active.

Data Guard and Exadata

Data Guard is the only technology that is able to maintain a completely independent physical replica of an Oracle Database residing on Exadata storage to protect against database or site failures. Furthermore, because Data Guard physical standby is the simplest, highest performance solution for maintaining a synchronized independent copy of the Oracle Database, it is the only technology that is able to support the very high volumes driven by an Oracle Database Machine. In internal Oracle Database 11g Release 2 tests on an Oracle Database Machine, Redo Apply was able to apply changes to a standby database at a sustained rate greater than 500MB/second. Please refer to the [MAA home page](#) on the Oracle Technology Network for more information.

Data Guard and Oracle Real Application Clusters

Data Guard and Oracle RAC are complementary technologies providing the highest possible level of scalability, availability, and data protection. Any combination of Oracle RAC and single node databases can participate and assume any role in a Data Guard configuration. Oracle RAC provides the ideal HA solution to protect against server failure simultaneous with providing industry unique capabilities for workload management and scalability. Data Guard provides an additional level of data availability and protection with complete redundancy that minimizes downtime due to storage array failure, operator errors, certain planned maintenance that can not be done in rolling fashion across Oracle RAC nodes, or multiple and correlated failures that can result in database (e.g. SAN array failure) or site failure (e.g. fire, flood, hurricane, or earthquake).

Maximum Availability Architecture

Oracle Maximum Availability Architecture (MAA) is an Oracle tested and customer validated best-practices blueprint for deploying Oracle high availability technologies. The goal of MAA is

"Our recovery strategy has always been based on tape backups. We also set up Oracle Data Guard as a "nice to have" optional extra. Then we had a total SAN failure and a couple of months later a major disk corruption on another SAN, both indirect results of power outages. On both occasions Data Guard enabled us to recover without loss of data. Now I realize it's not "nice to have" – it's essential!"

Rachel Slade, Oxford Brookes University

to remove complexity and accelerate a customer's learning curve for designing and operating the optimal high availability architecture.

[MAA best practices](#) include recommendations on various aspects of a Data Guard configuration, such as a configuration with Oracle RAC, optimizing redo transport, switchover/failover operations, client failover, Redo Apply performance, SQL Apply configuration and tuning, and use with Exadata storage and the Oracle Database Machine.

Data Guard Customers

Data Guard functionality was first available with Oracle Version 7 and has continued to add new functionality and become a more mature technology with each subsequent Oracle release. It is deployed for mission-critical applications at customer sites worldwide. A number of detailed implementation case studies are available on the [Oracle Technology Network](#).

Conclusion

Oracle Data Guard 11g fundamentally changes the traditional disaster recovery paradigm by offering an integrated HA/DR solution with unparalleled data protection and where standby systems simultaneously support production or test functions while they are in standby role.

Data Guard is a comprehensive data protection, data availability, and disaster recovery solution for the Oracle Database. It offers a flexible and easy-to-manage framework that addresses both planned and unplanned outages. Physical and logical standby databases provide high-value data protection while offloading overhead from primary databases. The various data protection modes provide flexibility to adapt to different levels of protection, performance and infrastructure requirements. The Data Guard Broker in combination with Oracle Enterprise Manager provides an easy-to-use configuration and management framework.

Regardless of the length to which high-availability has previously been built into an IT infrastructure using clusters, disk mirroring, and various backup and recovery strategies, it is a fact that data protection, availability, and your return on your IT investment are universally enhanced by including Data Guard in your IT architecture.

Appendix: Summary of Data Guard 11g New Features

DATA GUARD 11G RELEASE 1

AREA	CAPABILITY
Oracle Active Data Guard	Physical standby database open read only while apply is active. Standby queries get up-to-date results RMAN block change tracking for fast incremental backups on an Active Data Guard physical standby
Snapshot Standby	Temporarily open a standby database read-write while still providing disaster protection. Ideal complement to Oracle Real Application Testing
Fast-Start Failover	Asynchronous transport and Maximum Performance – configurable threshold to achieve desired RPO Initiate automatic failover upon previously designated health-check conditions or at application request Fault-tolerant observer for Fast-Start Failover – automatically restart failed observer on a second host
Redo Transport	Asynchronous redo transport enhanced for greater throughput on high latency Wide Area Networks Redo transport compression when resolving archive log gaps
Apply Performance	Redo Apply performance enhancements – double the performance of Data Guard 10g Various SQL Apply performance enhancements, also able to apply parallel DDL in parallel on standby
Planned Downtime	Database rolling upgrades using physical standby databases (transient logical standby) Additional flexibility for mixed primary/standby configurations to facilitate select migrations
Protection	Lost-write corruption protection using physical standby database
Security	SSL authentication can be used in lieu of password file to authenticate redo transmission
Role Transitions	Role specific scheduler jobs on a logical standby database using DBMS_SCHEDULER SQL Apply switchovers no longer require the prior shutdown of all but the first instance in each Oracle RAC cluster, either primary or standby Enterprise Manager jobs and metric thresholds propagated to new primary database upon role transition Data Guard Broker works seamlessly with cold cluster failovers controlled by Oracle Clusterware
SQL Apply Data Types	SQL Apply support for XMLType (when stored as CLOB), Transparent Data Encryption (TDE), DBMS_FGA (Fine Grained Auditing), DBMS_RLS (Virtual Private Database)
Manageability	Standby Statspack for tuning apply performance on an Active Data Guard standby Redo transport response time histogram used to determine appropriate value for NET_TIMEOUT Data Guard SQL Apply parameters set dynamically using DBMS_LOGSTDBY.APPLY_SET Create standby databases direct from the primary database using RMAN without interim storage Convert single instance standby databases to Oracle RAC using Enterprise Manager wizard

DATA GUARD 11G RELEASE 2

AREA	CAPABILITY
Oracle Active Data Guard	Automatically enforce service level objectives for maximum data delay when querying an active standby Automatically repair corrupt blocks online using an active standby
Redo Transport	Synchronous Redo Transport enhancements reduce overhead on primary database Redo transport compression for Synchronous and Asynchronous redo transport Support for up to 30 standby databases for a single primary database (previous limit was 9)
Apply Performance	Redo Apply enhancements that increase the maximum sustained apply rate to over 500MB/sec on the Oracle Database Machine with Exadata storage
Planned Downtime	Transparent support for Oracle Edition-based Redefinition, both Redo and SQL Apply SQL Apply can be used for zero risk, minimal downtime migration when implementing Oracle SecureFiles, Warehouse compression, OLTP table compression, or online redefinition
Protection	Un-sent redo in asynchronous configurations using Maximum Performance may be flushed to a standby before failover to achieve zero data loss (assuming failed primary can be brought to mount state)
Role Transitions	Redo Apply switchovers no longer require any standby instances to be shut down Data Guard Broker uses role-based database services to automate client failover Data Guard Broker managed role transitions work transparently with Oracle Restart
SQL Apply Data Types	Oracle SecureFiles, Warehouse compression, OLTP table compression Enhanced extended data type support for SQL Apply for replication of column objects (with simple or nested user-defined types), Varrays, and Oracle-supplied Spatial type SDO_GEOMETRY
Manageability	Increased performance for very large transactions (greater than 8 million rows) when using SQL Apply A logical standby database can be a source database in an Oracle Streams configuration Triggers can be defined on a logical standby to perform local processing independent of the primary Data Guard Broker has improved status and error reporting Data Recovery Advisor will utilize available standby database for Intelligent Data Repair



Oracle Data Guard
with Oracle Database 11g Release 2
September 2009
Author: Joe Meeks
Contributing Authors:
Larry Carpenter, Ashish Ray

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

Worldwide Inquiries:
Phone: +1.650.506.7000
Fax: +1.650.506.7200
oracle.com



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

Copyright © 2009, Oracle and/or its affiliates. All rights reserved. This document is provided for information purposes only and the contents hereof are subject to change without notice. This document is not warranted to be error-free, nor subject to any other warranties or conditions, whether expressed orally or implied in law, including implied warranties and conditions of merchantability or fitness for a particular purpose. We specifically disclaim any liability with respect to this document and no contractual obligations are formed either directly or indirectly by this document. This document may not be reproduced or transmitted in any form or by any means, electronic or mechanical, for any purpose, without our prior written permission.

Oracle is a registered trademark of Oracle Corporation and/or its affiliates. Other names may be trademarks of their respective owners.