An Oracle Technical White Paper
October 2011

Oracle Data Guard 11g
Data Protection and Availability for Oracle Database

ORACLE

Introduction 1
Oracle Data Guard 11g - Overview..... 2
How Data Guard Works – Technical Details..... 4
Data Guard Transport Service..... 4
Protection Modes..... 6
Data Guard Apply Services 6
Automatic Gap Resolution..... 9
Oracle Data Validation..... 9
Managing a Data Guard Configuration 10
Role Management Services..... 11
Addressing Planned Maintenance 12
Data Guard Compared to Remote-Mirroring..... 14
Data Guard and Oracle Exadata Database Machine 15
Data Guard and Oracle Real Application Clusters..... 15
Data Guard and Oracle GoldenGate 15
Maximum Availability Architecture 16
Data Guard Customers..... 16
Conclusion 17
Appendix: Summary of Data Guard 11g New Features..... 18

Introduction

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

Recovery using off-site backups or storage remote-mirroring, are traditional data protection and disaster recovery (DR) solutions for enterprise data. 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 adequate return on investment due to poor utilization of standby systems that sit idle until a failure occurs.

In contrast, Oracle Data Guard 11g redefines expectations for data protection solutions. Data Guard is the data protection and availability solution for Oracle Database. It 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 can be certain of the ability of a standby database to assume the production role – eliminating business risk at failover time. Data Guard standby databases also 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. Data Guard is unique among Oracle replication solutions in supporting both synchronous (zero data loss) and asynchronous (near-zero data loss) configurations. Administrators can chose either manual or automatic failover of production to a standby system if the primary system fails in order to maintain high availability for mission critical applications. An overview of Data Guard architecture is provided in Figure 1.

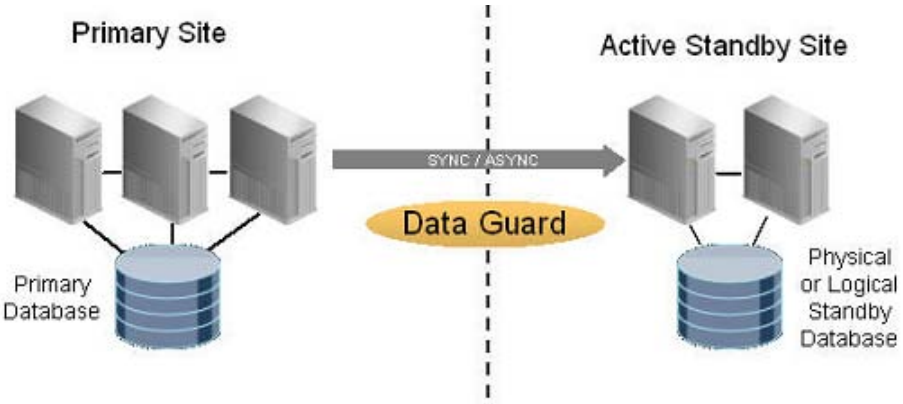

Figure 1 – Data Guard Architecture

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. Physical standby databases provide the best disaster recovery (DR) protection for the Oracle Database.

The second type of standby database uses SQL Apply to maintain a logical replica of the primary database. While a logical standby database contains the same data as the primary database, the physical organization and structure of the data can be different. Beginning with Data Guard 11g

SQL Apply is used in conjunction with a physical standby database to minimize planned downtime when upgrading to new Oracle Database releases or patch-sets (transient logical database rolling upgrade process).

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.

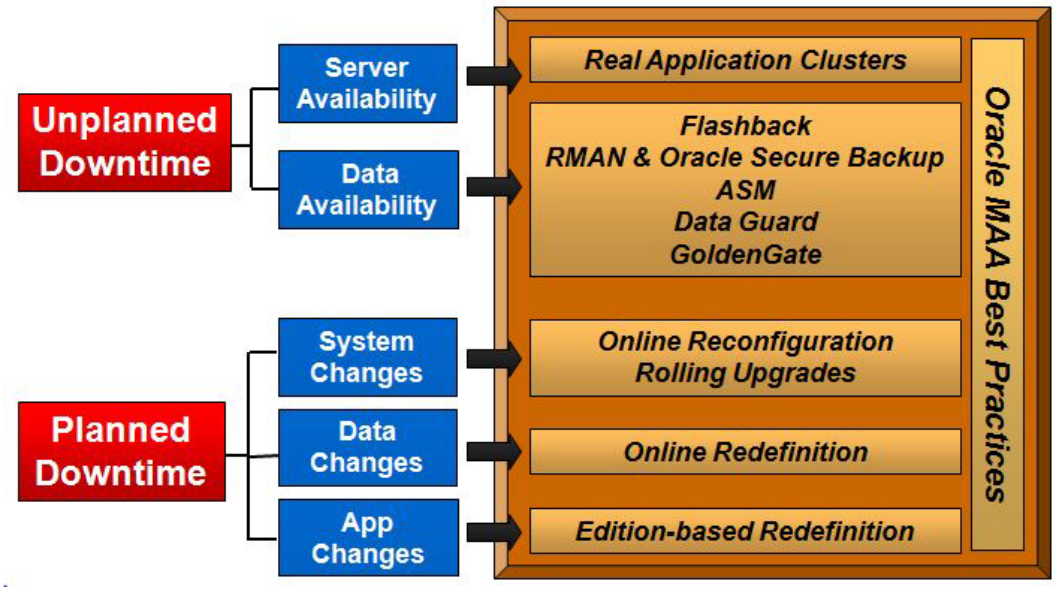

Figure 2 – Oracle Database High Availability Features

In addition to data protection and availability, Data Guard standby databases deliver high return on investment by supporting ad-hoc queries, reporting, backups, or test activity, while in standby role. Specifically:

- The Active Data Guard option (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 return up-to-date results. An Active Data Guard standby database is unique compared to other physical replication methods in its ability to guarantee the same level of read consistency as the primary database while replication is active.
- 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. When testing is complete, the Snapshot Standby is converted back into a synchronized physical standby

"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

database by first discarding the changes made while open read-write, and then applying the redo received from the primary database. Primary data is protected at all times.

- 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. All recovery operations are interchangeable between primary and Data Guard physical standby databases.

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 Service

As users commit transactions at the primary database, Oracle generates and writes redo records to a local online log file. As this occurs, Data Guard transport services transmit the redo directly from the primary database log buffer to the standby database(s) where it is written to a standby redo log file (step one in Figure 3). Data Guard 11g transport is very efficient, it is the only database replication method that does not require any capture processing, database access or disk I/O at the primary database. Data Guard is also the only Oracle database replication product that offers two types of replication methods: synchronous and asynchronous.

Synchronous redo transport (SYNC) requires that the primary database wait for confirmation from a standby database that redo has been received and written to disk (a standby redo log file) before it will acknowledge commit success to the application. This provides a guarantee of zero data loss protection in the event of any single failure, up to and including complete site failure. With the additional configuration of Data Guard Maximum Protection mode (discussed further below), zero data loss protection can even be assured in the case of multiple independent failures.

Data Guard 11g Release 2 reduces the impact that synchronous transport has on primary performance by transmitting redo to the remote standby in parallel with the local online log file write at the primary database (previously Data Guard would wait for the local log file write to complete before transmitting redo from memory to the remote standby). This enhancement effectively eliminates the time required for I/O to complete at the standby database from the total round trip time between primary and standby (assuming local and remote storage have similar I/O performance). The reduction in round-trip time compared to previous Data Guard releases enables even greater geographic separation between primary and standby databases in a synchronous zero data loss configuration. Alternatively, on low latency networks it can reduce the impact of SYNC replication on primary database performance to near zero.

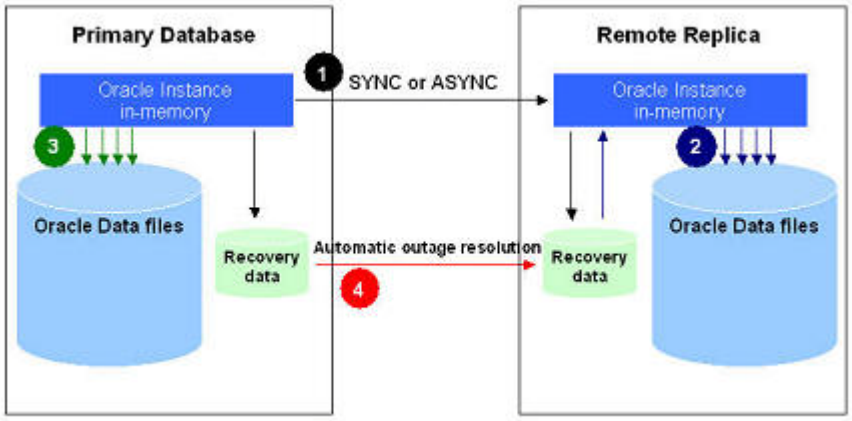

Figure 3 – Data Guard Redo Transport and Apply Services

Asynchronous redo transport (ASYNC) avoids any 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 ASYNC enhancements have virtually eliminated any impact on primary database performance by shipping directly from the primary log buffer (instead of from disk), as well as improving network throughput on high latency wide area networks (WAN). The performance benefit of ASYNC, however, is accompanied by the potential exposure for a small amount of data loss since there can be no guarantee that at any moment in time, all redo for committed transactions has been received by the standby database.

Note that Data Guard is able to support up to 30 standby databases in a single configuration. An increasing number of customers use this flexibility combined with Data Guard 11g redo transport enhancements to deploy both a local Data Guard standby for HA and a remote Data Guard standby for DR. A local Data Guard standby database provides an additional layer of HA

should unexpected events or human error make the production system unavailable even though the primary site is still operational. A local standby database can use synchronous redo transport with minimal, if any, impact on the primary database for zero data loss if a failover is required. The close proximity of the local standby to the primary site application tier also enables fast connection of application clients to the new primary database. Following a failover or switchover to the local standby database, the remote standby database in such a configuration will automatically recognize that a role transition has occurred and begin receiving redo from the new primary database - maintaining DR protection at all times. And finally, the local standby database can also be dual-purposed to offload backups from the primary database, or for use as a test system, or for implementing planned maintenance in rolling fashion (e.g. database rolling upgrades).

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	Signal commit success to the application only after acknowledgement is received from a standby database that redo for that transaction is hardened to disk.
Maximum Availability	Zero data loss Single failure protection	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

Data Guard Apply Services

Apply Services read redo records from a standby redo log file, perform continuous Oracle validation to insure that the redo is not corrupt, and then applies redo changes to the standby database (step two in Figure 3) using either Redo Apply (physical standby) or SQL Apply (logical standby). Data Guard is architected such that apply services are completely independent of redo transport Services. This results in zero data loss protection in synchronous Data Guard

configurations even in cases where apply services have been stopped. Data Guard transport services will continue to transmit redo to the standby where it will be archived for later use when apply is restarted. The isolation between transport and apply also means that the performance of apply services never impacts data protection or primary database performance, regardless of the transport method used.

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. Data Guard physical standby database is the only Oracle database replication method that has no data or storage type restrictions that would prevent the standby from being a complete replica of the production database.

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 on non-Exadata systems (see the Exadata Database Machine section later in this paper for performance data specific to Exadata). Redo Apply is the simplest, fastest, most reliable method (of maintaining a synchronized replica(s) of a primary database.

Redo Apply and Active Data Guard

Active Data Guard is an Oracle Enterprise Edition database option that extends the capabilities of Data Guard Redo Apply and physical standby databases. Unlike other Data Guard features that are included with Oracle Database Enterprise Edition, Active Data Guard requires a separate license purchase. Its capabilities include:

- 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 was previously idle while in standby role (additional active standby databases can be added to the configuration to further scale read-only capacity). Since Active Data Guard is a superset of standard Data Guard functionality, it delivers the same high performance as any other physical standby database. This makes it useful for high throughput applications where it is difficult or impossible for other replication methods to keep pace with the transaction volume generated by the source database.

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

Aris Prassinos, MorphoTrak

- 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) between a change being committed on the primary and the time it can be queried on an active standby database (new with Data Guard 11g Release 2). If the time limit is exceeded the active standby will return an `ORA-3172` error code instead of a stale result. Applications can be coded to respond to this error as if it were 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 transparent to the user and application. 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 repaird using the good version from the primary database.
- Active Data Guard also supports RMAN block-change tracking on a physical standby enabling RMAN fast incremental backups to be offloaded from the primary database.

For more technical information on Active Data Guard see the technical white paper, Active Data Guard Best Practices (including best practices for Oracle 11g Redo Apply).

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 them on a standby database that is open read-write. SQL Apply does have some restrictions - see Oracle documentation for further details.

Use SQL Apply if you meet its prerequisites and:

""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

- 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, then use physical standby and the transient logical database rolling upgrade process. See the section of this paper, *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 (an archive log gap and measured as transport lag). 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 quickly resynchronized using an RMAN fast incremental backup of the primary database taken from the point of the last SCN applied at the standby.

Oracle Data Validation

Data Guard uses Oracle processes to continuously 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 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 isolating 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. 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 is used to update other blocks of the database, thereby spreading corruption. 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 at the standby database to determine if there has been a lost write. This isolates the standby from the impact of lost writes that can occur at the primary database as well as detecting lost writes that may occur independently at the standby. For more details see the technical whitepaper, Preventing, Detecting, and Repairing Block Corruption: Oracle Database 11g.

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 the Broker's command-line interface (DGMGRL) or Oracle Enterprise Manager Grid Control.

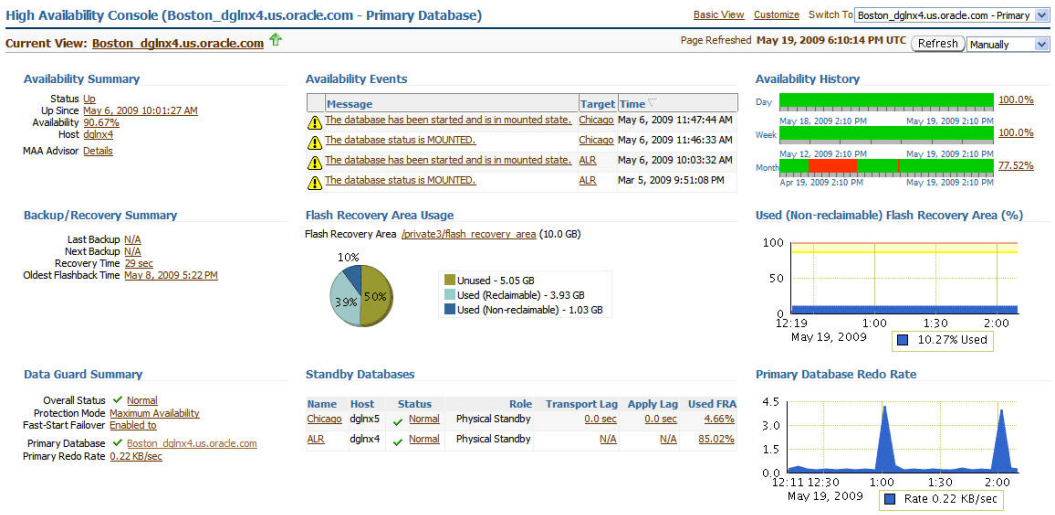

FIGURE 4 – ENTERPRISE MANAGER GRID CONTROL (10.2.0.5) HA CONSOLE

"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

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 the 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.

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 or unnecessary failovers from occurring. Once the failed primary is repaired and mounted, it must establish connection with the Observer process before it can open (preventing split-brain). 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 architecture of fast-start failover makes it ideal 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. For more details see the technical white paper, Client Failover Best Practices for Highly Available Oracle Databases: Oracle Database 11g Release 2.

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, changes to the logical structure of the database, or different hardware architecture, has always been able to use Redo Apply to implement changes in rolling fashion. With each new Oracle release, Redo Apply has gained additional flexibility to address a wider range of planned maintenance requirements. More details are provided in the following sections.

Platform Migration

Since Oracle 10g, there has been increased flexibility in cross-platform support using Redo Apply. In certain Data Guard configurations, primary and standby databases are able to run on systems having different O.S (e.g. Windows/Linux), word size (32bit/64bit) or hardware architectures. Redo Apply is also used to migrate to Automatic Storage Management (ASM), to move from single instance Oracle Databases to Oracle RAC, to perform technology refresh, or to move from one data center to the next. See My Oracle Support Note 413484.1 for details on mixed platform combinations supported in a Data Guard configuration.

Standby-First Patch Apply

Beginning with Oracle Database 11.2, Oracle has introduced Standby-First Patch Apply to enable a physical standby to use Redo Apply to support different software patch levels between a primary database and its physical standby database for the purpose of applying and validating Oracle patches in rolling fashion. Patches eligible for Standby-First patching include:

- Patch Set Update (PSU)
- Critical Patch Update (CPU)
- Patch Set Exception (PSE)
- Oracle Database bundled patch
- Oracle Exadata Database Machine bundled patch
- Oracle Exadata Storage Server Software patch

Standby-First Patch Apply is supported for certified software patches for Oracle Database Enterprise Edition Release 2 (11.2) release 11.2.0.1 and later. Refer to My Oracle Support Note 1265700.1 for more information and the README for each patch to determine if a target patch is certified as being a Standby-First Patch.

Database Rolling Upgrades

Beginning with Oracle Database 10g, 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 and logical standby database. From Oracle Database 11g onward, this can also be accomplished by starting with a Data Guard physical standby database and using the transient logical standby rolling upgrade process. The transient logical process is attractive because it uses an existing physical standby database and only requires a single catalog upgrade to migrate both primary and standby databases to the new Oracle release. SQL Apply is used on a temporary basis during the period when replication is occurring across database versions, but when the

"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

upgrade process is complete, the configuration reverts to its original state of having a primary with a physical standby database, except that both are running at the new version. For more information, see the technical white paper, Database Rolling Upgrades Made Easy.

The only downtime required for maintenance implemented in rolling fashion is the time required to complete a switchover. Switchovers using Redo Apply can complete in less than 60 seconds..

Note that Data Guard 11g Release 2 SQL Apply adds native support for Binary XML, Oracle Advanced Compression (OLTP Table Compression), Oracle Secure Files, and Online Redefinition. SQL Apply also includes the ability to implement extended data type support where native support is not yet available, 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. Combined, these enhancements greatly expand the universe of databases that can be upgraded in rolling fashion with minimal downtime using SQL Apply.

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 the additional I/O volume can be problematic for host or array based remote mirroring solutions - the historical alternative to Data Guard for remote data protection. 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 online redo log files. 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.

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 Oracle Exadata Database Machine

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 Exadata Database Machine. In internal Oracle Database 11g Release 2 tests using Exadata, Redo Apply was able to apply changes to a standby database at a sustained rate greater than 600MB/second. For more details see the technical white paper, Disaster Recovery for Exadata Databases Machine.

Data Guard and Oracle Real Application Clusters

Data Guard, Oracle RAC and Oracle RAC One Node 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).

Data Guard and Oracle GoldenGate

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

Use Data Guard Redo Apply for data protection, high availability, and disaster recovery. Data Guard is the Oracle solution for data protection due to the many advantages it offers over logical replication for such requirements.

GoldenGate is an advanced logical replication product that supports 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 hardware platforms and database management systems. Unlike Data Guard, GoldenGate captures primary database changes by reading redo records from disk, transforming those records into a platform independent trail file format, and

"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

transmitting the trail file to the target database. GoldenGate maintains a logical replica by converting the trail file into SQL and applying SQL to the target database.

Use GoldenGate for data distribution, data integration, for bi-directional and mulitmaster replication, and cross platform migrations or other planned maintenance not supported by Data Guard.

See Active Data Guard and Oracle GoldenGate for a more detailed discussion of the strengths offered by each of these strategic Oracle technologies.

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 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, network configuration and tuning, and use with Exadata Database Machine. Find Data Guard best practices, demonstrations, documentation, user case studies, hands-on labs, technical articles and more, on the Oracle Technology Network.

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 10*g* 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) Support for cascaded standby configurations using Oracle RAC (in addition to single instance)
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, Binary XML 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
October, 2011
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

0109