

Oracle Data Guard in Oracle Database 10g

Disaster Recovery for the Enterprise

An Oracle White Paper
December 2003

Oracle Data Guard

Disaster Recovery for the Enterprise

Executive Overview	3
Impact Of Disasters	3
High Availability Challenges	3
Oracle Data Guard	4
Overview Of Oracle Data Guard.....	4
What Is Oracle Data Guard?	4
Oracle Data Guard Functionality.....	5
Benefits Of Oracle Data Guard.....	6
Oracle Data Guard Process Architecture.....	7
Key Technology Components	8
Data Guard Configuration	8
Redo Apply And SQL Apply.....	9
Physical Standby Database – Redo Apply.....	9
Logical Standby Database – SQL Apply	11
Real Time Apply	13
Data Protection Modes	14
Maximum Protection.....	14
Maximum Availability	15
Maximum Performance	15
Failover and Switchover	16
Automatic Resynchronization.....	18
Human Error Protection	18
Rolling Upgrades.....	19
Cascaded Redo Log Destinations.....	19
Enterprise Manager & Data Guard Broker	20
Configuration Options	21
Oracle Data Guard And RAC.....	21
Maximum Availability Architecture	21
Data Guard And Remote Mirroring Solutions	22
Conclusion	24
References	25

Oracle Data Guard

Disaster Recovery for the Enterprise

EXECUTIVE OVERVIEW

Business continuity and disaster recovery are top priorities for the senior management of most global enterprises. Economic fluctuations, rapid changes in market trends, and competitive pressures imply that the global enterprise of today must operate in a 24x7 environment, and must be able to swiftly and efficiently deal with unforeseen business interruptions.

Oracle Data Guard is one of the most effective solutions available today to protect the core asset of any enterprise – its data, and make it available on a 24x7 basis despite disasters and other outages. This paper discusses Data Guard technology in Oracle Database 10g, and demonstrates how it is a key factor in the business continuity infrastructure of any enterprise.

IMPACT OF DISASTERS

With the proliferation of e-business, an enterprise today operates in an extremely complex and a highly networked, global economy, and is more susceptible to interruptions than in the past. The cost of interruptions, or downtime, varies across industries and can be as much as millions of dollars an hour. While that number is staggering, the reasons are quite obvious. The Internet has brought millions of customers directly to the electronic storefronts. Critical and interdependent business matters such as customer relationships, competitive advantages, legal obligations, industry reputation and shareholder confidence are even more critical now because of their increased vulnerability to business disruptions and downtimes.

High Availability Challenges

Downtime that affects a business could be either unplanned or planned. Unplanned downtime may be due to hardware or system failures, data/storage failures, human errors, computer viruses, software glitches, natural disasters and malicious acts. A business may also have to undergo planned downtimes because of scheduled maintenances such as system upgrades.

A company designing its business continuity strategy must create a business continuity plan (BCP) that can effectively deal with these challenges. One of the critical requirements of the BCP is that it must protect business data, because data is one of the most critical company assets – whether it is payroll/employee information, customer records, valuable research, financial records, historical information, etc. If a company loses its data, it is not easily replaced, and rebuilding or regenerating that data will likely be an extremely expensive, if not an impossible task, critically affecting the company's ability to stay in business.

Oracle Data Guard

Oracle Data Guard is designed to address the highly important business continuity need for the enterprise. It provides an extensive set of data protection and disaster recovery (DR) features to help businesses survive disasters, human errors and corruptions that can adversely affect their Oracle databases.

This whitepaper provides an architectural and technological overview of the Oracle Data Guard feature of Oracle Database 10g Release 1. For additional details on Data Guard, please refer to Oracle Data Guard documentation (ref. [1]).

OVERVIEW OF ORACLE DATA GUARD

What Is Oracle Data Guard?

Oracle Data Guard is the management, monitoring, and automation software infrastructure that creates, maintains, and monitors one or more standby databases to protect enterprise data from failures, disasters, errors, and corruptions.

Data Guard maintains standby databases as transactionally consistent copies of the production database. These standby databases can be located at remote disaster recovery sites thousands of miles away from the production data center, or they may be located in the same city, same campus, or even in the same building. If the production database becomes unavailable because of a planned or an unplanned outage, Data Guard can switch any standby database to the production role, thus minimizing the downtime associated with the outage, and preventing any data loss.

Available as a feature of the Enterprise Edition of Oracle Database, Data Guard can be used in combination with other Oracle High Availability (HA) solutions such as Real Application Clusters (RAC) and Recovery Manager (RMAN), to provide a high level of data protection and data availability that is unprecedented in the industry.

The following diagram presents an overview of Oracle Data Guard.




Fig. 1: Overview of Oracle Data Guard Architecture

Oracle Data Guard Functionality

Oracle Data Guard consists of a *production database*, also known as the *primary database*, and one or more *standby database(s)*, which are transactionally consistent copies of the primary database. Data Guard maintains this transactional consistency using redo data. As transactions occur in the primary database, redo data is generated and written to the local redo log files. With Data Guard, this redo data is also transferred to the standby sites and applied to the standby databases, keeping them in synch with the primary database. Data Guard allows the administrator to choose whether this redo data is sent synchronously or asynchronously to a standby site.

The underlying technologies for standby databases are *Data Guard Redo Apply (physical standby database)*, and *Data Guard SQL Apply (logical standby database)*. A physical standby database has on-disk database structures that are identical to the primary database on a block-for-block basis, and is updated using Oracle media recovery. A logical standby database is an independent database that contains the same data as the primary database. It is updated using SQL statements, and has the relative advantage that it can be used concurrently for recovery and for other tasks such as reporting and queries.

Data Guard facilitates switchover and failover operations between the primary database and a chosen standby database, reducing overall downtime during planned outages and unplanned failures.

The primary and standby databases, as well as their various interactions, may be managed by using SQL*Plus. For easier manageability, 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 use either Oracle Enterprise Manager or the Broker's own specialized command-line interface (DGMGRL) to take advantage of the Broker's management capabilities.

The following diagram shows the Oracle Data Guard components.

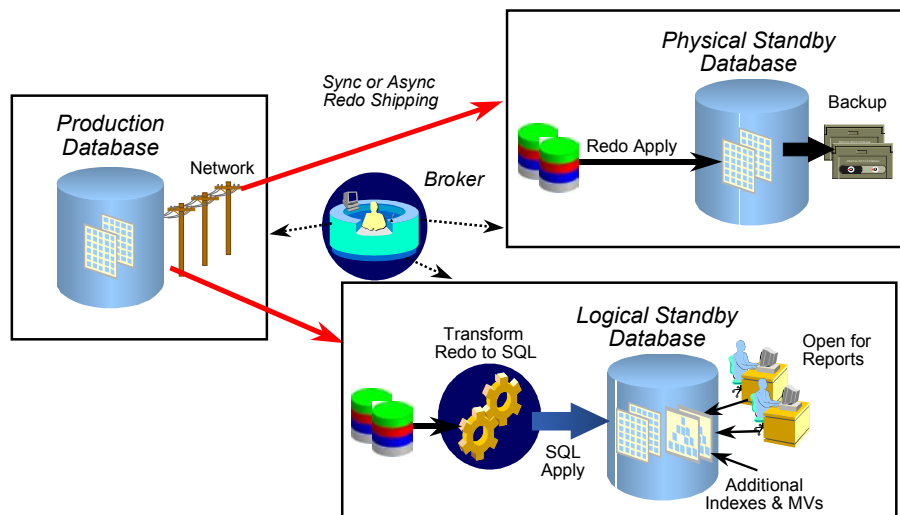


Fig. 2: Oracle Data Guard Architectural Components

Benefits Of Oracle Data Guard

Oracle Data Guard offers the following benefits:

- *Disaster recovery and high availability* – Data Guard provides an efficient and comprehensive disaster recovery and high availability solution. Easy-to-manage switchover and failover capabilities allow role reversals between primary and standby databases, minimizing the downtime of the primary database for planned and unplanned outages.
- *Complete data protection* – With standby databases, Data Guard ensures no data loss, even in the face of unforeseen disasters. A standby database provides a safeguard against data corruption and user errors. Storage level physical corruptions on the primary database do not propagate to the standby database. Similarly, logical corruptions or user errors that cause the primary database to be permanently damaged can be resolved. Finally, the redo data is validated at the time it is applied to the standby database.
- *Efficient utilization of system resources* – The standby database tables that are updated with redo data received from the primary database can be used for other tasks such as backup operations, reporting, summations, and queries, thereby reducing the primary database workload necessary to perform these tasks, saving valuable CPU and I/O cycles. With a logical standby database,

“Data Guard automates disaster-recovery procedures and reduces Fidelity's exposure to data loss by an order of magnitude compared to previous approaches.”

Jonathan Schapiro
Vice President
Data Architecture & Services
Fidelity Investments

users can perform data manipulation operations on tables in schemas that are not updated from the primary database. A logical standby database can remain open while the tables are updated from the primary database and the tables are simultaneously available for read-only access. Finally, additional indexes and materialized views can be created on the maintained tables for better query performance and to suit specific business requirements.

- *Flexibility in data protection to balance availability against performance requirements* – Oracle Data Guard offers the Maximum Protection, Maximum Availability and Maximum Performance modes to help enterprises balance data protection against system performance requirements.
- *Automatic gap detection and resolution* – If connectivity is lost between the primary and one or more standby databases (e.g. because of network problems), redo data being generated on the primary database cannot be sent to those standby databases. Once connectivity is re-established, the missing archive log sequence (or the gap) is automatically detected by Data Guard and the necessary archive logs are automatically transmitted to the standby databases. The standby databases are resynchronized with the primary database, with no manual intervention by the administrator.
- *Centralized and simple management* – The Data Guard Broker automates the management and operational tasks across the multiple databases in a Data Guard configuration. The Broker also monitors all of the systems within a single Data Guard configuration. Administrators may use either Oracle Enterprise Manager or the Broker's own specialized command-line interface (DGMGRL) to take advantage of this integrated management framework.
- *Integrated with Oracle database* – Oracle Data Guard is available as a fully integrated feature of the Oracle Database (Enterprise Edition) at no extra cost.

ORACLE DATA GUARD PROCESS ARCHITECTURE

As shown in the following figure, Oracle Data Guard uses several processes of the Oracle database instance to achieve the automation necessary for disaster recovery and high availability.

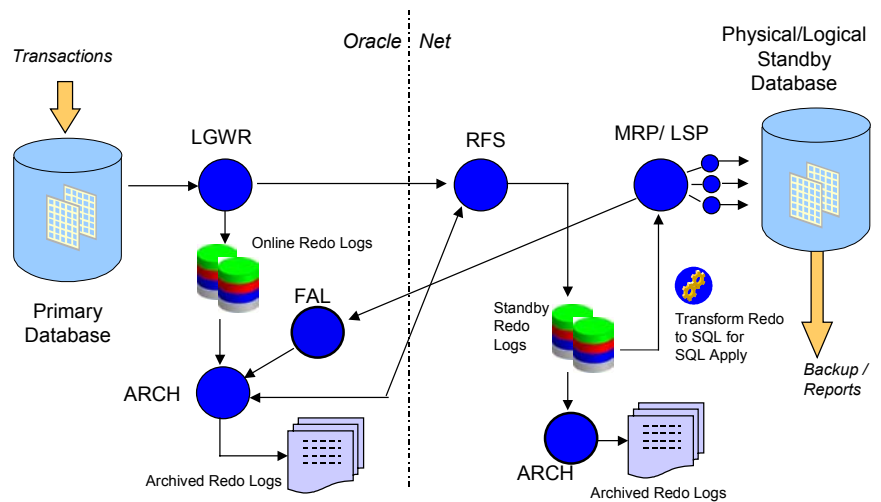


Fig. 3: Oracle Data Guard Process Architecture

On the primary database, Oracle Data Guard uses the *Log Writer Process (LGWR)* or the *Archiver Process (ARCH)* to collect transaction redo data and ship this data to the standby, and the *Fetch Archive Log Process (FAL)* to provide a client-server mechanism for shipping archived logs to the standby following a communication loss between the primary and standby(s), for automatic gap resolution and resynchronization.

On the standby database, Oracle Data Guard uses the *Remote File Server (RFS)* process to receive redo records from the primary database, the *Managed Recovery Process (MRP)* to apply redo information to the physical standby database, and the *Logical Standby Process (LSP)* to apply SQL-translated redo information to the logical standby database.

If the Data Guard Broker is enabled, Oracle Data Guard also uses the *Data Guard Broker Monitor (DMON)* process to manage and monitor the primary and standby databases as a unified configuration.

KEY TECHNOLOGY COMPONENTS

Data Guard Configuration

A Data Guard configuration consists of one primary database and up to nine standby databases. The primary and standby databases can run on a single node or in a Real Application Clusters (RAC) environment. The standby databases are connected to the primary database over standard TCP/IP-based networks (e.g. a Local Area Network (LAN), a Metropolitan Area Network (MAN), a Wide Area Network (WAN)) using Oracle Net Services. There are no

restrictions on where the databases are located, provided that they can communicate with each other. However, for disaster recovery, it is recommended that the standby databases be hosted at sites that are geographically separated from the primary site.

Data Guard requires the operating system architecture on the primary and standby systems to be the same. Thus if the primary database is running the Linux operating system on an Intel architecture, all its standby databases must also be running Linux on Intel – they cannot be Windows systems, for example. In addition, the same release of Oracle Database Enterprise Edition must be installed on the primary database and all standby databases in a Data Guard configuration.

Redo Apply And SQL Apply

A standby database is initially created from a backup copy of the primary database. Once created, Data Guard automatically maintains the standby database as a transactionally consistent copy of the primary database by transmitting primary database redo data to the standby system and then applying the redo data to the standby database.

Data Guard provides two methods to apply this redo data to the standby database and keep it transactionally consistent with the primary, and these methods correspond to the two types of standby databases supported by Data Guard:

- Redo Apply, used for physical standby databases
- SQL Apply, used for logical standby databases

Note that as Fig. 3 indicates, there is no distinction between these two types of standby databases as far as redo data transmission from the primary is concerned. Once this redo data reaches the standby server, it is how the redo data is applied on the standby database that distinguishes these two types of standby databases.

Physical Standby Database – Redo Apply

A physical standby database is kept synchronized with the primary database by applying the redo data received from the primary using Oracle media recovery. It is physically identical to the primary database on a block-for-block basis, and thus the database schemas, including indexes, are the same.

How Redo Apply Works

A log switch on the primary database triggers a log switch on the standby database, causing Archiver processes on the standby database to archive the current standby redo log file to an archive log on the standby database. Thereupon, Data Guard Redo Apply uses a specialized process, called the Managed Recovery Process (MRP), to read from the archive log and apply the redo data to the physical standby database. If the new Oracle Data Guard

feature in Oracle Database 10g, *Real Time Apply*, is enabled, MRP reads redo data directly from the current standby redo log file as it is being filled up by the RFS process.

MRP (and thus application of redo) can be started on the physical standby database by mounting the database and using the following command:

```
ALTER DATABASE RECOVER MANAGED STANDBY DATABASE  
DISCONNECT FROM SESSION;
```

The media recovery processes can be run in parallel for the best performance of Data Guard Redo Apply. In releases prior to Oracle Database 10g, this required the use of a `PARALLEL` clause in the above `RECOVER MANAGED STANDBY DATABASE` command. In Oracle Database 10g, MRP can automatically determine the optimal number of parallel recovery processes at the time it starts (without requiring the `PARALLEL` clause), and this number is based on the number of CPUs available on the standby server.

The physical standby database can be opened read-only, and queries can be run against the physical standby database at that time. The physical standby database cannot run recovery at the same time it is opened read-only. Redo data that are shipped to the standby while it is opened read-only are accumulated at the standby site, and are not applied. However, recovery operations can be resumed on the physical standby at any time, and the accumulated redo data will automatically get applied. This allows the physical standby database to run in a sequence that could involve running in recovery for a while, then being opened read-only to run reports, and then returning to running recovery to apply outstanding redo data.

To open the physical standby read-only, recovery needs to be canceled on the standby using the following command:

```
ALTER DATABASE RECOVER MANAGED STANDBY DATABASE  
CANCEL;
```

and then the database can be opened read-only:

```
ALTER DATABASE OPEN;
```

Benefits of Physical Standby

A physical standby database provides the following benefits:

- *Disaster recovery and high availability* – A physical standby database enables a robust and efficient disaster recovery and high availability solution. Easy-to-manage switchover and failover capabilities allow simple role reversals between primary and physical standby databases, minimizing the downtime of the primary database for planned and unplanned outages.
- *Data protection* – Using a physical standby database, Data Guard can ensure no data loss, even in the face of unforeseen disasters. A physical standby database supports all datatypes, and DDL and DML operations that the

primary can support. It also provides a safeguard against data corruptions and user errors. Storage level physical corruptions on the primary database do not propagate to the standby database. Similarly, logical corruptions or user errors that cause the primary database to be permanently damaged can be resolved. Finally, the redo data is validated when it is applied to the standby database.

- *Reduction in primary database workload* – The physical standby database can be opened read-only for reporting and queries. Besides, using Recovery Manager (RMAN), the physical standby database can be utilized to create backups for the production database, thereby saving valuable CPU and I/O cycles from the production system. RMAN can perform this backup while the physical standby database is performing recovery, or when it is opened read-only.
- *Performance* – The redo apply technology used by the physical standby database applies changes using low-level recovery mechanisms, which bypass all SQL level code layers and therefore is the most efficient mechanism for applying changes. This makes the redo apply technology a highly efficient mechanism to propagate changes among databases.

“We needed to consider the safe-keeping of our data, but we also needed to look at cost. Oracle Data Guard provides everything for a high availability solution at a lower cost than other alternatives.”

Ann Collins
Technical Director
First American Real Estate Solutions

Logical Standby Database – SQL Apply

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. The SQL Apply technology keeps the logical standby database synchronized with the primary database by transforming the redo data received from the primary database into SQL statements and then executing the SQL statements on the standby database. This makes it possible for the logical standby database to be accessed for queries and reporting purposes at the same time the SQL is being applied to it.

Because the logical standby database is updated using SQL statements, it remains open in read-write mode, and the tables that are being updated from the primary database can be used simultaneously for other tasks such as reporting, summations, and queries. These tasks can also be optimized by creating additional indexes and materialized views on the maintained tables. A logical standby database can host multiple database schemas, and users can perform normal data manipulation operations on tables in schemas that are not updated from the primary database.

A logical standby database has some restrictions on datatypes, types of tables, and types of DDL and DML operations. Please refer to [1] for a list of these unsupported datatypes and storage attributes.

How SQL Apply Works

SQL Apply uses a collection of parallel execution servers and background processes that perform the task of applying changes from the primary database to the logical standby database. The following diagram shows the flow of information and the role that each process performs.

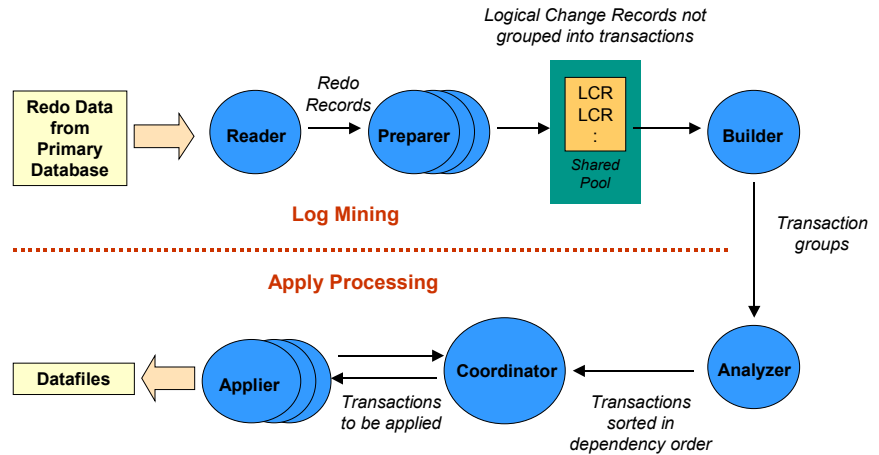


Fig. 4: Data Guard SQL Apply Process Architecture

These various SQL Apply processes can be started by entering this simple command on the logical standby database:

```
ALTER DATABASE START LOGICAL STANDBY APPLY;
```

Considering each of these SQL Apply processes, the *Reader* process reads redo records from the archive logs (or standby redo logs if Real Time Apply is enabled). The *Preparer* processes convert the block changes into table changes, or logical change records (LCRs). At this point, the LCRs do not represent any specific transactions. The *Builder* process assembles completed transactions from the individual LCRs. The *Analyzer* process examines the completed transactions, identifying dependencies between the different transactions. The *Coordinator* process (also known as the Logical Standby Process, or LSP), assigns transactions to the apply processes, monitors dependencies between transactions, and authorizes the commit of changes to the logical standby database. The *Applier* processes apply the LCRs for the assigned transaction to the database and commit the transactions when instructed to do so by the Coordinator.

Data Guard provides helpful views to inspect the state of each process.

Benefits of Logical Standby

A logical standby database provides similar disaster recovery, high availability, and data protection benefits as a physical standby database. It also provides the following specialized benefits:

- *Efficient use of standby hardware resources* – A logical standby database can be used for other business purposes in addition to disaster recovery requirements. It can host additional database schemas beyond the ones that are protected in a Data Guard configuration, and users can perform DDL or DML operations on those schemas any time. Because the logical standby tables that are protected by Data Guard can be stored in a different physical layout than on the primary database, additional indexes and materialized views can be created to improve query performance and suit specific business requirements.
- *Reduction in primary database workload* – A logical standby database can remain open at the same time its tables are updated from the primary database, and those tables are simultaneously available for read access. This enables the logical standby database to be used concurrently for data protection and reporting, thereby off-loading the primary database from those reporting and query tasks, and saving valuable CPU and I/O cycles.

Real Time Apply

With Data Guard's new Real Time Apply feature in Oracle Database 10g, redo data can be applied on the standby database (whether Redo Apply or SQL Apply) as soon as the redo data is written to a Standby Redo Log (SRL). Prior releases of Data Guard require this redo data to be archived at the standby database in the form of archive logs before they can be applied.

To enable real time apply for physical standby databases, recovery should be started on the physical standby database using the following command:

```
ALTER DATABASE RECOVER MANAGED STANDBY DATABASE
USING CURRENT LOGFILE;
```

To enable real time apply for logical standby databases, the apply processes should be started on the logical standby database using the following:

```
ALTER DATABASE START LOGICAL STANDBY APPLY
IMMEDIATE;
```

The Real Time Apply feature allows standby databases to be closely synchronized with the primary database, enabling up-to-date and real-time reporting (especially for Data Guard SQL Apply). This also enables faster switchover and failover times, which in turn reduces planned and unplanned downtime for the business.

The impact of a disaster is often measured in terms of Recovery Point Objective (RPO – i.e. how much data can a business afford to lose in the event of a disaster) and Recovery Time Objective (RTO – i.e. how much time a business can afford to be down in the event of a disaster). With Oracle Data Guard, when the Maximum Protection mode ensuring no data loss even in the event of disasters is used in combination with Real Time Apply, businesses get the benefits of both zero data loss as well as minimal downtime in the event of a disaster and this makes Oracle Data Guard the only solution available today

with the best RPO and RTO benefits for a business.

Data Protection Modes

Oracle Data Guard provides three high-level modes of data protection to balance cost, availability, performance, and transaction protection. These modes can be set easily using any of the available management interfaces. For example, following is a simple SQL statement that can be executed on the primary database for this purpose:

```
ALTER DATABASE SET STANDBY DATABASE TO MAXIMIZE  
{PROTECTION | AVAILABILITY | PERFORMANCE};
```

To determine the appropriate data protection mode, enterprises need to weigh their business requirements for data protection against user demands for system response time. The following table outlines the suitability of each mode from a risk of data loss perspective.

Protection Mode	Risk of Data Loss In the Event of a Disaster	Redo Transport Mechanism
<i>Maximum Protection</i>	Zero data loss; Double failure protection	LGWR SYNC
<i>Maximum Availability</i>	Zero data loss; Single failure protection	LGWR SYNC
<i>Maximum Performance</i>	Minimal data loss – usually 0 to few seconds	LGWR ASYNC or ARCH

The following sections describe all of these aspects in more detail.

Maximum Protection

Maximum Protection mode offers the highest level of data protection for the primary database, ensuring a comprehensive zero-data loss disaster recovery solution. When operating in Maximum Protection mode, redo records are synchronously transmitted by the Log Writer process from the primary database to the standby database, and a transaction is not committed on the primary database until it has been confirmed that the transaction data is available on disk on at least one standby server. This mode should be configured with at least two standby databases, thus offering double failure protection. If the last participating standby database becomes unavailable, processing stops on the primary database. This ensures that no transactions are lost when the primary database loses contact with all of its standby databases.

Because of the synchronous nature of redo transmission, this Maximum Protection mode can potentially impact primary database response time. This impact can be minimized by configuring a low latency network with sufficient bandwidth for the peak transaction load. Stock exchanges, currency exchanges,

and financial institutions are examples of businesses that may require this Maximum Protection mode.

Maximum Availability

Maximum Availability mode has the next highest level of data availability for the primary database, offering zero data loss and protecting against single component failures. As with the Maximum Protection mode, redo data is synchronously transmitted by the Log Writer process from the primary database to the standby database, and the transaction is not complete on the primary database until it has been confirmed that the transaction data is available on disk on the standby server. However, in this mode, unlike the Maximum Protection mode, if the last participating standby database becomes unavailable – e.g. because of network connectivity problems, processing continues on the primary database. The standby database may temporarily fall behind compared to the primary database, but when it is available again, the databases will automatically synchronize with no data loss.

Because of synchronous redo transmission, this protection mode can potentially impact response time and throughput. This impact can be minimized by configuring a low latency network with sufficient bandwidth for peak transaction load.

The Maximum Availability mode is suitable for businesses that want the assurance of zero data loss protection in the event of a severe outage at the production site (assuming that there are no other failures), but do not want the production database to be impacted by network/standby server failures.

Maximum Performance

Maximum Performance mode is the default protection mode. It offers slightly less primary database data protection, but higher performance, than Maximum Availability mode. In this mode, as the primary database processes transactions, redo data is asynchronously shipped to the standby database by the Log Writer process. Alternatively, the Archiver process(es) on the primary database may also be configured to transmit the redo data in this mode. In any case, the commit operation of the primary database does not wait for the standby to acknowledge receipt before completing the write on the primary. If any standby destination becomes unavailable, processing continues on the primary database and there is little or no effect on performance. However, in such cases, error messages are logged to the database alert log and correspondingly, alerts can be set through Enterprise Manager.

In the case of a failure of the primary, there may be some transactions that were committed on the primary that had not completed shipping to the standby. If the network has sufficient throughput to keep up with peaks in redo traffic, the number of lost transactions should be very small or zero.

The Maximum Performance mode should be used when availability and performance on the primary database are more important than the risk of losing a small amount of data. This mode is also suitable for Data Guard deployments over a WAN, where the inherent latencies of the network may limit the suitability of synchronous redo transmission.

Failover and Switchover

Oracle Data Guard offers two easy-to-use methods to handle planned and unplanned outages of the production site. These methods are called switchover and failover, and they can be easily initiated by the administrator using the Oracle Enterprise Manager GUI interface, the Data Guard Broker's command line interface, or directly through SQL.

Failover is the operation of bringing one of the standby databases online as the new primary database when an unplanned catastrophic failure occurs on the primary database, and there is no possibility of recovering the primary database in a timely manner.

The failover operation is initiated on the standby database that will assume the primary role. If the *Flashback Database* feature of Oracle Database 10g was enabled on the original primary database prior to the failover, it may be used after the failover if it is intended to bring back the original primary as a new standby database in the Data Guard configuration. If flashback database was not enabled prior to the failover, and the original primary database needs to be brought back as a standby after the failover, it has to be recreated from a backup copy of the new primary database.

The following figure shows the result of a failover operation from a primary database in San Francisco to a physical standby database in Boston.

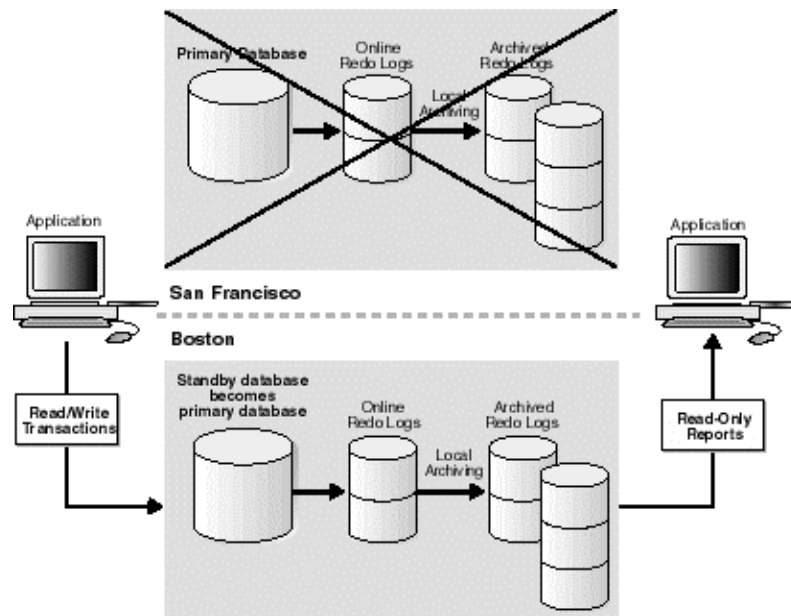


Fig. 5: Failover to a Standby Database

The switchover option, on the other hand, is the planned role reversal of the primary and standby databases, to handle planned maintenance on the primary database. The main difference between a switchover operation and a failover operation is that switchover is performed when the primary database is still available, and it does not require a flashback or re-instantiation of the original primary database. This allows the original primary database to assume the role of a standby database almost immediately. As a result, scheduled maintenance can be performed more easily and frequently. For example, switchover may be used to perform an upgrade on the primary site by switching over all of the database clients to the standby site as hardware is upgraded on the primary site.

At times, the term *switchback* may also be used within the scope of Data Guard role management. A switchback operation is nothing more than a subsequent switchover operation to return the roles to their original state.

The switchover operation always ensures no data is lost during the transition. The failover operation ensures zero data loss if Data Guard was being run in the Maximum Protection mode or Maximum Availability mode at the time of the failover.

It should be stressed that to avoid *false* failovers/switchovers in the event of temporary system or network failures, Data Guard switchover and failover operations are not automatic, but have to be explicitly initiated by the administrator. Once initiated, Data Guard automates the processes involved.

Failover and switchover operations work seamlessly when multiple standby databases are included in the configuration. For example, if multiple standby databases are configured and the primary database goes down, the administrator

has the flexibility to choose one of the available standbys to become the primary. Data Guard fully automates the process of redirecting the other standby databases to use the new primary, including shipping any missing or incomplete redo data.

Automatic Resynchronization

Oracle Data Guard can smoothly handle network connectivity problems that temporarily disconnect the standby (physical or logical) database from the primary database.

When the standby database becomes unavailable (unless this standby database is the last available standby in the Maximum Protection mode in which case the primary database will be shut down), transactions are captured locally at the primary database. When connectivity to the standby is re-established, the accumulated archive logs are automatically shipped and applied to the standby, until the standby has resynchronized with the primary. This process does not require any administrative intervention. Oracle recommends that network capacity be sufficient to handle such resynchronizations if network outages are common in the vicinity of the primary site.

Human Error Protection

When a primary database is open and active, and transactions are in progress, redo data is generated and shipped to standby sites. Considering that human error is the leading cause of system downtime, it may be possible that this redo data contains critical logical user-errors, such as dropping of an important table, and this might have already corrupted the primary database.

Data Guard provides several easy-to-use means to avoid such user errors. The administrator may decide to use the *Flashback Database* feature of Oracle Database 10g on both the primary and standby databases to quickly revert the databases to an earlier point-in-time to back out such user errors. Alternatively, if the administrator decides to failover to a standby database, but those user-errors were already applied to the standby database (say, because Real Time Apply was enabled), the administrator may simply flashback the standby database to a safe point in time (assuming the flashback functionality was already enabled on the standby database). Finally, the administrator has the added option not to use the Real Time Apply feature at one or more standby databases, and instead delay the application of redo data on those standby databases by a configurable amount of time, which provides a window of protection from such user errors or corruptions.

Irrespective of the option chosen, the apply process on the standby database always revalidates the log records to prevent application of physical redo data corruptions on the standby database.

Rolling Upgrades

Oracle Database 10g supports database software upgrades for major release and patchset upgrades (from Oracle Database 10g onwards) in a rolling fashion – with near zero database downtime, by using Data Guard SQL Apply. The steps involve upgrading the logical standby database to the next release, running in a mixed mode to test and validate the upgrade, doing a role reversal by switching over to the upgraded database, and then finally upgrading the old primary database. While running in a mixed mode for testing purpose, the upgrade can be aborted and the software downgraded, without data loss. For additional data protection during these steps, a second standby database may be used.

The following diagram shows the sequence of events in a rolling upgrade process.

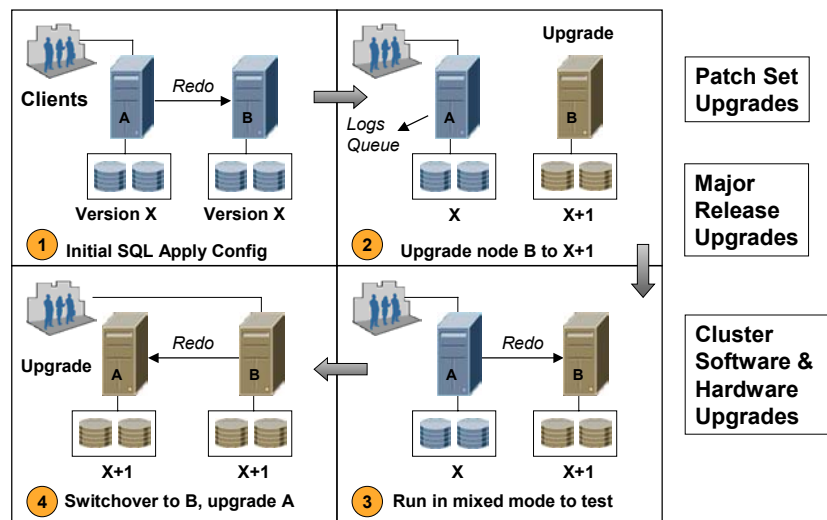


Fig. 6: Rolling Database Upgrades Using SQL Apply

By supporting rolling upgrades with minimal downtimes, Data Guard reduces the large maintenance windows typical of many administrative tasks, and enables the 24x7 operation of the business.

Cascaded Redo Log Destinations

In this case, a standby database receives its redo data from another standby database and not from the original primary database. Since the primary database sends redo data to only a subset of the standby databases, this feature reduces the load on the primary system, and also reduces network traffic and use of valuable network resources around the primary site.

Enterprise Manager & Data Guard Broker

The Oracle Data Guard Broker is a distributed management framework that automates and centralizes the creation, maintenance, and monitoring of Data Guard configurations. All management operations can be performed either through Oracle Enterprise Manager, which uses the Broker, or through the Broker's specialized command-line interface (DGMGRL). The following screenshot shows the Data Guard home page of the Enterprise Manager.



Fig. 7: Oracle Data Guard Configuration through Oracle Enterprise Manager

“The things that impress me the most about Data Guard are its manageability, reliability, and ease of use. It is amazing how easily we could implement a solid Disaster Recovery / High Availability solution with Oracle Data Guard without requiring additional resources to support it.”

Darl Kuhn
Staff Engineer, Database Services
Sun Services Global Engineering
Sun Microsystems

The following list describes some of the operations that the Broker automates and simplifies:

- Creating and enabling Data Guard configurations, which include a primary database and up to nine standby (physical or logical) databases – all or a mix of these databases may be RAC clusters.
- Managing an entire Data Guard configuration from any site in the configuration.
- Implementing switchover or failover operations that involve complex role changes across all systems in the configuration.
- Monitoring log apply rates, capturing diagnostic information, and detecting problems quickly with centralized monitoring, testing and event notification.

The Broker's easy-to-use interfaces and centralized management and monitoring of the Data Guard configuration make Data Guard an enhanced high availability and data protection solution for the enterprise.

CONFIGURATION OPTIONS

Standby databases can be remote (connected over a WAN or MAN) or local (connected on LAN). A Data Guard configuration can contain both local and remote standby databases and can be configured to provide the benefits of both approaches.

Remote standby databases are the best solution for disaster recovery because an event that disables the primary database is unlikely to also disable the standby database. However, performance can be affected by higher latency.

Local standby databases are better suited to address outages related to human errors or data corruptions within a data center. Since LAN provides an inexpensive and reliable network link with low latency, having a dedicated LAN segment just for standby use is an ideal environment for the Maximum Protection mode.

Please refer to [5] for Data Guard network configuration best practice recommendations.

ORACLE DATA GUARD AND RAC

Oracle Data Guard and Oracle Real Application Clusters (RAC) are complementary to each other. RAC addresses system or instance failures. It provides rapid and automatic recovery from failures that do not affect data – such as node failures, or instance crashes. It also provides increased scalability for an application. Data Guard, on the other hand, provides data protection through the use of transactionally consistent primary and standby databases, which neither share disks nor run in lock step. This enables recovery from site disasters or data corruptions.

Data Guard is also natively integrated with RAC – for e.g., some or all of the primary/standby (physical or logical) databases can be RAC databases, and they can be managed using Enterprise Manager or the Broker's command-line interface, or directly using SQL. Customers should use a combination of Oracle Data Guard and Real Application Clusters to get the benefits of both data-level and system-level protection.

MAXIMUM AVAILABILITY ARCHITECTURE

As more system capabilities become available, IT managers, architects and administrators often find it difficult to integrate a suitable set of features to build one unified high availability (HA) solution that fits all of their business requirements. Oracle Maximum Availability Architecture (MAA) is Oracle's best-practices blueprint based on proven Oracle high availability technologies and recommendations, with the goal of removing the complexity in designing the optimal high availability architecture, and maximizing systems availability.

MAA provides the following benefits:

- MAA reduces the implementation costs for a highly available Oracle system by providing detailed configuration guidelines. The results of performance impact studies for different configurations are highlighted to ensure that the chosen highly available architecture can continue to perform and scale according to business needs.
- MAA provides best practices and recovery steps to eliminate or minimize downtime that could occur because of scheduled and unscheduled outages such as human errors, system faults and crashes, maintenance, data failures, corruptions, and disasters.
- MAA gives the ability to control the length of time to recover from an outage and the amount of acceptable data loss under disaster conditions thus allowing mean time to recovery (MTTR) to be tailored to specific business requirements.

Data Guard is an essential component of MAA, and the MAA guidelines include best practice recommendations (ref. [3] – [6]) on various Data Guard configuration aspects, such as a configuration involving both RAC and Data Guard, redo data transport mechanisms, switchover/failover, media recovery, SQL Apply configuration, network configuration, etc. Customers interested in Data Guard implementations are strongly recommended to refer to these MAA best practice guidelines. Besides the MAA whitepapers, standardized Oracle documentation (ref. [2]) on high availability best practices is also available for customers interested in Oracle Database 10g.

DATA GUARD AND REMOTE MIRRORING SOLUTIONS

Remote Mirroring solutions conceptually appear to offer simple and complete data protection. However, Oracle Data Guard is inherently much more efficient, less expensive and better optimized for data protection than remote mirroring solutions. A customer does not need to buy or integrate a remote mirroring solution with Oracle Data Guard. Following are its benefits compared to a remote mirroring solution:

- *Better Network Efficiency*
With Oracle Data Guard, only the redo data need to be sent to the remote site. However, if a remote mirroring solution is used for data protection, then the database files, the online logs, the archive logs and the control file must be mirrored. This means that remote mirroring will send each change at least three times to the remote site. Further, database writes happen a lot more often than log writes because each log write typically contains many changes (known as group commit). This means that the network bandwidth needed for a database redo shipping based solution is considerably less than that of a remote mirroring solution. Even more importantly, this means far fewer network round trips.

Remote mirroring can be very useful for non-database files, but for database data – the combination of better protection and lower cost provide compelling reasons to use Data Guard. An internal analysis of Oracle's corporate e-mail systems, as shown in the following graph, demonstrated that 7 times more data was transmitted over the network and 27 times more I/O operations were performed using a remote mirroring solution, compared to using Data Guard.




Fig. 8: Network Performance: Data Guard vs. Remote Mirroring

- *Better suited for WANs*

Remote-mirroring solutions based on storage systems often have a distance limitation due to the underlying communication technology (Fibre, ESCON) used by the storage systems. This distance can be extended by using specialized devices from third party vendors. These devices convert ESCON/fibre to the appropriate IP, ATM or SONET networks. The problem is that with each such device, latency is introduced in the system, impacting the production database performance, and making such a configuration unsuitable for synchronous transport necessary for the zero data loss capability. This problem may be mitigated by introducing intermediate storage boxes in the communication path, but that only adds to the overall cost. The other solution is to use variations of synchronous transmission – however, depending on the remote mirroring solution, anything other than synchronous transmission of data may not preserve write-ordering across all mirrored volumes that the database resides on. This means such configurations cannot guarantee data consistency at all times, making them unsuitable as a data protection / disaster recovery solution for OLTP data.

Since Data Guard transmits only redo data to the standby sites, using a standard IP network, and preserves transactional consistency across all the protection modes (i.e. whether using synchronous or asynchronous mode of transport), and does not need expensive interim storage boxes, it is a much better DR and data protection solution for a WAN.

- *Better resilience and data protection*

Oracle Data Guard processes are aware of data formats as they read and write information from the primary database. In addition, Oracle Data Guard is integrated with the *Flashback Database* feature, and allows application of changes to be delayed as well. These capabilities prevent many human errors and data corruptions from propagating, and/or affecting the standby database. Remote mirroring does not have this advantage – any inadvertent drop of a critical table will be instantly propagated to, and adversely affect, the remote copy of the database files.

- *Higher ROI*

Using Oracle Data Guard, the standby database can be opened read-only for reporting while changes are still propagating. That is not always the case for a remote mirroring solution. Besides, Oracle Data Guard is an out-of-the-box feature of the core Oracle database. It does not involve extra costs, or extra integration. However, remote mirroring solutions are extra cost purchases and require complex integration with the database. Last, but not the least, many of these remote mirroring solutions are proprietary and can be used with only the storage systems from the same vendor (on both the primary and secondary sites) that manufactures these remote mirroring solutions. Data Guard, on the other hand, does not force any lock-in with a particular storage solution for both the primary and standby sites.

“The Airbus global procurement process requires highly available and very powerful database systems. We’ve been testing Oracle Data Guard 10g since July 2003 and we are very excited with the new features in Data Guard.

We expect to have the highest degrees of data availability and data protection with Data Guard, which will help us achieve our business continuity goals throughout our organization.”

Thomas Brunken
IT Project Manager
Airbus Deutschland GmbH
Infrastructure Design & Projects

CONCLUSION

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

A global enterprise of today cannot provide a mission-critical level of service to its customers and various stakeholders without the kind of technology this paper talks about. It has to be complete, integrated, easy-to-manage, serve multiple purposes and protect all enterprise data. At the same time, such data protection and disaster recovery technology should not be expensive, and should enable businesses to extract value out of their DR investments. Oracle Data Guard is the only solution available today that meets all these needs.

REFERENCES

1. Oracle Data Guard Concepts and Administration Manual, 10g Release 1 (10.1)
2. Oracle High Availability Architecture and Best Practices Manual, 10g Release 1 (10.1)
3. MAA Detailed White Paper,
<http://otn.oracle.com/deploy/availability/htdocs/maa.htm>
4. Oracle9i Media Recovery Best Practices,
<http://otn.oracle.com/deploy/availability/htdocs/maa.htm>
5. Oracle9i Data Guard: Primary Site and Network Configuration Best Practices, <http://otn.oracle.com/deploy/availability/htdocs/maa.htm>
6. Oracle9i Data Guard: SQL Apply Best Practices,
<http://otn.oracle.com/deploy/availability/htdocs/maa.htm>



Oracle Data Guard – Disaster Recovery for the Enterprise
December 2003
Author: 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
www.oracle.com

Oracle is a registered trademark of Oracle Corporation. Various product and service names referenced herein may be trademarks of Oracle Corporation. All other product and service names mentioned may be trademarks of their respective owners.

Copyright © 2003 Oracle Corporation
All rights reserved.