

Switchover and Failover Best Practices:
Oracle Data Guard 10g Release 2

*Oracle Maximum Availability Architecture White Paper
July 2010*

Maximum
Availability
Architecture

Oracle Best Practices For High Availability

Switchover and Failover Best Practices: Oracle Database 10g Release 2

Introduction	2
Observations and Best Practices - Overview	2
Failover Best Practices - Overview	4
Switchover Best Practices - Overview	5
Data Guard Role Transitions – Overview.....	6
Failover	7
Minimizing Application Downtime During Failover	7
Manual Failover.....	8
Manual Failover to a Physical Standby Database	8
Manual Failover to a Logical Standby Database.....	9
Fast-Start Failover.....	9
Fast-Start Failover to a Physical or a Logical Standby Database	11
Test Results for Manual Failovers and Fast-Start Failovers.....	12
Single-Instance Databases	12
Multiple-Instance Real Application Clusters	13
Switchover	13
Minimizing Application Downtime During Switchover.....	14
Switchover Steps For Physical Standby Databases.....	15
Using SQL*Plus and Logical Standby Databases	15
Test Results for Switchovers	17
Single-Instance Databases	17
Multiple-Instance Oracle Real Application Clusters	18
Application and Client Failover	18
Conclusion.....	19
References	20

Switchover and Failover Best Practices: Oracle Data Guard 10g Release 2

INTRODUCTION

[Oracle Data Guard](#)[1] is one of the most effective and comprehensive data protection and disaster recovery solutions available today for enterprise data. It offers protection for Oracle Databases, providing customers the ability to protect one of their most important assets: the business' online information. Using Data Guard failover and switchover operations, this information remains available after unplanned outages—such as a network outage or if the production database fails, or during planned outages—such as software upgrades or other regular maintenance.

Data Guard Fast-Start Failover was introduced in Oracle Database 10g Release 2 along with significant improvements to existing failover and switchover features to further reduce the time needed to perform Data Guard *role transitions*.

This white paper provides the [Maximum Availability Architecture \(MAA\)](#)[2] best practices for achieving the fastest Data Guard switchovers and failovers with Oracle Database 10g Release 2. It also provides estimates of switchover and failover timings under various configuration settings. In addition, refer to the complementary white paper: "[Oracle Data Guard 10g Release 2 - Fast-Start Failover Best Practices](#)"[3] for discussions specific to Fast-Start Failover. Taken together, these white papers provide practical advice on best practices for performing role transitions in an Oracle Data Guard environment. Updates to these and other MAA best practice white papers may be found on the [MAA](#) page on the Oracle Technology Network (OTN) [2].

MAA testing with Data Guard 10g Release 2 role transitions included testing switchovers, manual failovers, and fast-start failovers. All testing was performed using both Oracle Enterprise Manager and the Data Guard Broker command-line interface (DGMGRL), as well as SQL*Plus statements.

OBSERVATIONS AND BEST PRACTICES - OVERVIEW

With proper planning and execution, Data Guard role transitions can minimize downtime and ensure that the database environment is restored with minimal effect on the business. Whether using physical standby or logical standby

databases, MAA testing determined that switchover and failover times with Oracle Data Guard 10g Release 2 have been reduced to seconds:

- Automatic failover in seconds for site disasters and database failures
- Automatic failover in seconds for data corruptions
- Manual switchover in seconds to reduce planned downtime for system, hardware, or site changes

Table 1 shows the solutions and recovery times you can attain with Data Guard and related solutions for a variety of outages.

Table 1: Oracle Solutions for Unplanned and Planned Outages

Outage Type	Oracle Solutions	Recovery Time
Computer failure	<ul style="list-style-type: none"> • Fast-Start Failover and Fast-Application Notification (FAN) [8] 	Less than 30 seconds
Storage failure	<ul style="list-style-type: none"> • Fast-Start Failover and Fast-Application Notification • Data Guard and Automatic Storage Management (ASM) [9] 	Less than 30 seconds No downtime ¹
Data Corruption	<ul style="list-style-type: none"> • Automatic validation of redo blocks before they are applied, fast failover to an uncorrupted standby database upon production database corruption 	Less than 30 seconds
Site failure	<ul style="list-style-type: none"> • Fast-Start Failover and Fast-Application Notification (FAN) [8] 	Less than 30 seconds ²
System and cluster upgrades	<ul style="list-style-type: none"> • Switch over to a physical or a logical standby database for system upgrades that cannot be upgraded using an Oracle RAC rolling upgrade (for example, due to system restrictions or cluster firmware upgrades that require downtime) 	Seconds to minutes
All patch set and database upgrades ³	<ul style="list-style-type: none"> • Data Guard SQL Apply and logical standby databases to upgrade an Oracle database. 	Seconds to minutes

¹ Storage failures can be prevented if you use Automatic Storage Management (ASM) with mirroring and its automatic rebalance capability.

² The recovery time indicated applies to database and existing database connection failover. Network connection changes and other site-specific failover activities may lengthen the overall recovery time.

³ Supported only for Oracle Database Release 10.1.0.3 and higher. Also, note that SQL Apply has some data-type restrictions. For more information, see the [Oracle Data Guard Concepts and Administration](#)[5] documentation for a listing.

For information about all of the Oracle high-availability solutions and the benefits of faster Data Guard failover and switchover, refer to the [Oracle Database High Availability Overview](#)[4]. For detailed best practices for the comprehensive use of Oracle Database High Availability features, refer to [Oracle Database High Availability Best Practices 10g Release 2 - Documentation](#) [14].

Failover Best Practices - Overview

To optimize failover processing, use the following best practices:

- Use Fast-Start Failover

MAA tests running Oracle Database 10g Release 2 show that failovers performed using the Data Guard Broker and Fast-Start Failover offer a significant improvement in availability. Oracle recommends users read the comprehensive review of Oracle best practices contained in the white paper: [“Oracle Data Guard 10g Release 2 - Fast-Start Failover Best Practices”](#)[3].

For Disaster Recovery requirements, it is ideal to install the Fast-Start Failover Observer in a location separate from the production and standby data centers. The Observer should be independent from the data centers, and when possible, it should connect to the production and standby database via the same network as any end-user client. If the designated observer fails, Enterprise Manager can detect it and can be configured to automatically restart the observer on the same host. If a third, independent location is not available you should locate the observer in the standby data center, but on a separate host and in a fashion that will isolate it as much as possible from failures impacting the standby database.

- Enable Flashback Database to reinstate failed production databases after a failover operation has completed. Flashback Database provides a second very significant function, enabling fast point in time recovery if needed. Set `DB_FLASHBACK_RETENTION_TARGET` to a minimum value of 60 minutes to enable reinstatement after a failover (default value is 1440 minutes).
- Use Data Guard real-time apply in order to apply redo data to the standby database as soon as it is received.
- For manual failovers that involve Oracle Real Application Clusters, issue the `SHUTDOWN ABORT` statement on all secondary Oracle RAC instances on the standby database prior to performing a failover.
- For logical standby databases, refer to [Oracle 10g SQL Apply MAA Best Practices](#) [6] to obtain an optimal SQL Apply rate.
- For physical standby databases:

- Refer to the [Oracle Database 10g MAA Best Practices: Data Guard Redo Apply and Media Recovery](#) [7] to optimize media recovery for Redo Apply.
- Go directly to the OPEN state from the MOUNTED state instead of restarting the standby database (as required in previous releases).
- When transitioning from read-only mode to Redo Apply (recovery) mode, restart the database.
- Set the LOG_FILE_NAME_CONVERT parameter. As part of a failover, the standby database must clear its standby online logs prior to opening as the new production database. The time needed to complete this I/O can add significantly to the overall failover time. By setting the LOG_FILE_NAME_CONVERT parameter, the standby will pre-clear the standby online redo logs the first time the MRP process is started.

Switchover Best Practices - Overview

- Disconnect all sessions and stop job processing.
- Prior to performing a switchover, cancel any specified apply delay
- Insure that there is no gap in redo on the standby database.
- For logical standby databases:
 1. Refer to the [Oracle 10g SQL Apply Best Practices](#) [6] white paper to obtain an optimal SQL Apply rate.
 2. Prior to performing the actual switchover, perform the LogMiner Multi-versioned Data Dictionary build.
 3. Follow the detailed switchover steps for logical standby databases provided later in this paper.
- For physical standby databases:
 - Refer to the [Oracle 10g Redo Apply and Media Recovery Best Practices](#) [7] white paper to obtain an optimal Redo Apply rate.
 - Follow the detailed switchover steps provided in MetaLink Note 751600.1
- Use Data Guard real-time apply so that redo data is applied to the standby database as soon as it is received, and the standby database should be synchronized with the production database prior to the switchover operation in order to minimize switchover time.
 - Enable Flashback Database so that if a failure occurs during the switchover, the process can be easily reversed. Set DB_FLASHBACK_RETENTION_TARGET to a minimum value of 60

minutes to enable reinstatement after a failover (default value is 1440 minutes).

- Before executing the switchover reduce the number of ARCH processes to the minimum needed for both remote and local archiving. Additional ARCH processes can take additional time to shutdown thereby increasing overall switchover timings. Once the switchover has been completed you can reenable the additional ARCH processes.
- Set the LOG_FILE_NAME_CONVERT parameter. As part of a switchover the standby must clear the standby online logs prior to opening as a new production database. The time needed to complete the I/O can add significantly to the overall switchover time. By setting the LOG_FILE_NAME_CONVERT parameter the standby will pre-clear the standby online redo logs the first time the MRP process is started.

DATA GUARD ROLE TRANSITIONS – OVERVIEW

A Data Guard configuration consists of one database functioning in the production role and one or more databases that function in standby roles. Data Guard maintains these standby databases as synchronized 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.

When unplanned or planned outages occur, Data Guard can change one of the standby databases into the production role quickly, with minimal downtime. Data Guard provides switchover and failover for efficient and rapid recovery from outages, whether you lose a single server or an entire site, to keep your business up and running.

A **switchover** is a planned role reversal between the production database and one of its standby databases to avoid downtime during scheduled maintenance on the production system or to test readiness for future role transitions. A switchover guarantees no data loss. During a switchover, the production database transitions to a standby role, and the standby database transitions to the production role. The transition occurs without having to restart either database. A switchover is performed by an administrator through either Enterprise Manager, Data Guard broker command-line interface, or by issuing SQL*Plus commands.

A **failover** is performed when the production database (all instances of an Oracle RAC production database) fails and one of the standby databases is transitioned to take over the production role, allowing business operations to continue. Once the failover is complete and applications have resumed, the administrative staff can turn its attention to resolving the problems with the failed system. Failover may or may not result in data loss depending on the Data Guard protection mode in effect at the time of the failover.

As of Oracle Database 10g Release 2 there are two distinct types of failover: manual failover and fast-start failover. An administrator initiates manual failover when the production database fails. In contrast, Data Guard automatically initiates a fast-start failover without human intervention after the production database has been unavailable for a set period of time (fast-start failover threshold).

Note: A highly available architecture must achieve not only fast database failover, but it must also execute fast client failover in order for applications to be available to the business. Best practices for client failover in a Data Guard configuration are described in [MAA Best Practices for Client Failover in Data Guard Configurations for Highly Available Oracle Databases](#)[10].

FAILOVER

Performing a failover in a Data Guard configuration converts the standby database into the production database. The following sections describe [manual failover](#) and [fast-start failover](#) in greater detail.

Minimizing Application Downtime During Failover

In order to minimize overall application downtime during failover operations, application connections should seamlessly and automatically reconnect to the database that is being transitioned to the primary database role. Configuring application connections involves:

1. Creating a database service that the application uses for its connection.
2. Configuring the database service used by the application to automatically migrate to the new primary database at failover time.
3. Configure the application Oracle Net alias or JDBC URL to include all host that have the potential host the primary database.
4. Automating application reconnections
 - a. OCI applications can use Transparent Application Failover (TAF) to automate reconnection for existing connections
 - b. JDBC applications can use FAN if the primary and standby are in a Oracle RAC environment. For non Oracle RAC environments JDBC applications can be coded to automatically reconnect upon an exception

For detailed information on configuring your application for client failover please refer to the [MAA Best Practices for Client Failover in Data Guard Configurations for Highly Available Oracle Databases](#)[10].

While the MAA client failover paper discusses client failover for physical standby database, the same steps will work for a logical standby with the exception of FAN OCI support.

Manual Failover

Manual failover is performed by the administrator directly through the Enterprise Manager graphical user interface, or the Data Guard broker command-line interface (DGMGRL), or by issuing SQL*Plus statements. The sections below describe the relevant SQL*Plus commands.

Manual Failover to a Physical Standby Database

Use the following commands to perform a manual failover of a physical standby database:

1. For manual failovers in a Real Application Clusters environment, issue the `SHUTDOWN ABORT` statement on all Oracle RAC instances on secondary standby databases prior to performing the failover.

2. Initiate the failover by issuing the following on the target standby database:

```
SQL> ALTER DATABASE RECOVER MANAGED STANDBY DATABASE FINISH FORCE;
```

Note: Include the `FORCE` keyword to ensure that the RFS processes on the standby database will fail over without waiting for the network connections to time out through normal TCP timeout processing before shutting down.

3. Convert the physical standby database to the production role:

```
ALTER DATABASE COMMIT TO SWITCHOVER TO PRIMARY;
```

4. If the standby database was never opened read-only since the last time it was started, then open the new production database by issuing the following statement:

```
ALTER DATABASE OPEN;
```

If the physical standby database has been opened in read-only mode since the last time it was started, shut down the target standby database and restart it:

```
SQL> SHUTDOWN IMMEDIATE;  
SQL> STARTUP;
```

Note: In rare circumstances, administrators may wish to avoid waiting for the standby database to complete applying redo in the current standby redo log file before performing the failover. (note: use of Data Guard real-time apply will avoid this delay by keeping apply up to date on the standby database). If so desired, administrators may issue the `ALTER DATABASE ACTIVATE STANDBY DATABASE` statement to perform an immediate failover. This statement converts the standby database to the production database, creates a new resetlogs branch, and opens the database. However, because this statement will cause any un-applied redo in the standby redo log to be lost, Oracle recommends you only use the failover procedure described in the above steps to perform a failover.

Refer to these sections in [Oracle Data Guard Concepts and Administration](#) [5]: “[Failovers Involving a Physical Standby Database](#)” for step-by-step failover instructions and “[Recovering Through the OPEN RESETLOGS Statement](#)” to learn how a physical standby database reacts to the new resetlogs branch.

Manual Failover to a Logical Standby Database

Use the following commands to perform a manual failover of a logical standby database:

1. For manual failovers in a Real Application Clusters environment, issue the `SHUTDOWN ABORT` statement on all Oracle RAC instances on all standby databases prior to performing the failover.
2. Initiate the failover by issuing the following on the target standby database:

```
ALTER DATABASE ACTIVATE LOGICAL STANDBY DATABASE FINISH APPLY;
```

This statement stops the RFS process, applies any remaining redo data, stops SQL Apply, and activates the logical standby database in the production role.

Note: To avoid waiting for the redo in the standby redo log file to be applied prior to performing the failover, omit the `FINISH APPLY` clause on the statement. Although omitting the `FINISH APPLY` clause will accelerate failover, omitting the clause will cause the loss of any unapplied redo data in the standby redo log. To gauge the amount of redo that will be lost, query the `V$LOGSTDBY_PROGRESS` view. The `LATEST_SCN` column value indicates the last SCN received from the production database, and the `APPLIED_SCN` column value indicates the last SCN applied to the standby database. All SCNs between these two values will be lost.

Refer to these sections in [Oracle Data Guard Concepts and Administration](#) [5]: “[Failovers Involving a Logical Standby Database](#)” for step-by-step failover instructions and “[Recovering Through the OPEN RESETLOGS Statement](#)” to learn how a logical standby database reacts to the new resetlogs branch.

Fast-Start Failover

Fast-Start Failover is an Oracle Data Guard 10g Release 2 feature that quickly and reliably fails over the target standby database to the production database role, without requiring an administrator to perform manual steps to invoke the failover and with no loss of data.

You must enable the Data Guard configuration and Data Guard Broker for fast-start failover before a failure occurs. Once enabled, an *Observer* process monitors the Data Guard configuration 24x7 and will initiate a failover to the specified

target standby database automatically, whenever the production database becomes unavailable for a specified period of time.

Two of the three members of a Fast-Start Failover configuration - the production database, the target standby database, and the Observer, must agree that all requirements have been met before an automatic failover can occur. This design avoids conditions such as *split-brain* scenarios in which two divergent production databases are able to accept transactions for the same application.

The Observer, (a Broker client), monitors the Data Guard configuration and ensures that it can connect to the production database. If both the Observer and the standby database lose connectivity to the production database, then the Observer will attempt to reconnect for a period of time defined by the administrator. If the Observer or Standby database still cannot contact the production database after this period of time has expired, a failover is initiated.

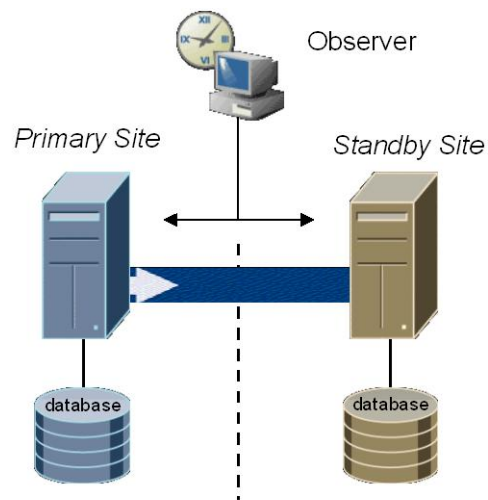


Figure 1 Fast-Start Failover Configuration

Moreover, after failover the Broker automatically *reinstates* the failed production database as a new target standby once that database is restarted (assuming that the database can be restarted and connectivity to the Observer can be established). This enables Data Guard to quickly and automatically resynchronize it with the new production database. The failed (old production) database no longer needs to be restored from a backup of the new production database. This expedites restoring high availability to the Data Guard configuration.

Fast-Start Failover to a Physical or a Logical Standby Database

Fast-Start Failover is used within a Data Guard configuration under the control of the Data Guard Broker. The Data Guard Broker provides centralized management of all resources within a Data Guard configuration through its command line interface - DGMGRL.

The Data Guard Broker uses single commands to efficiently perform the equivalent work of multiple SQL*Plus statements, greatly simplifying the management of a Data Guard configuration. Management can be further simplified by using the Enterprise Manager⁴ graphical user interface that interfaces directly with the Data Guard Broker.

Before enabling fast-start failover, the following prerequisites must be met:

- Enable Flashback Database on both the production database and the target standby database Set `DB_FLASHBACK_RETENTION_TARGET` to a minimum value of 60 minutes to enable reinstatement after a failover (default value is 1440 minutes).
- Configure a Flash Recovery Area on both the production database and the target standby database
- Enable the Data Guard Broker configuration
- Configure Redo Transport Services in LGWR SYNC mode
- Run the Data Guard configuration in Maximum Availability mode
- Ensure the Observer has network connectivity to both the standby and production databases

When you configure Fast-Start Failover with the Broker, it sets up these three essential participants (Figure 1) in the configuration:

- Production database
- Target (physical or logical) standby database

⁴ Enterprise Manager is the preferred interface for fast-start failover for several reasons:

- a. When the Observer is started via Enterprise Manager, it is started as a background process.
- b. Using Enterprise Manager metrics, DBAs can monitor the Observer and be notified when the Observer goes down.
- c. Enterprise Manager automatically restarts the Observer if the host on which it was running is restarted.
- d. If the Observer fails, Enterprise Manager can detect it and can be configured to automatically restart the Observer on the same host.

- Fast-Start Failover Observer

Additional Fast-Start Failover configuration information is available in "[Oracle Data Guard 10g Release 2 - Fast-Start Failover Best Practices](#)"[3] and the [Oracle Data Guard Broker](#) documentation[13]. Also, refer to "[Setup and Maintenance for Oracle Flashback Database](#)" in [Oracle Database Backup and Recovery Basics](#) [12] and "[Setting Up Flash Recovery Areas](#)" in [Oracle Data Guard Concepts and Administration](#) [5] for information about setting up Flashback Database and flash-recovery areas.

TEST RESULTS FOR MANUAL FAILOVERS AND FAST-START FAILOVERS

A number of tests were run to measure failover times using the best practices described in this paper and Oracle Data Guard Release 2. The test databases were each 100GB in size and connected to a Gigabit Network. Although different network latencies were simulated, latency was not a factor in optimizing failover and switchover times. The workload on the production database generated redo at a rate of 3 MB/second. Both single instance databases and Oracle RAC configurations were tested. Tested configurations included failover to a physical standby database (Redo Apply), and a logical standby database (SQL Apply).

To invoke a failover during testing, a failure was simulated by issuing a SHUTDOWN ABORT on the production database. The time for each major section of the failover was measured using the Data Guard Broker and the database alert logs. In all cases, the user-configurable failover threshold (or time to detect the failure) was not included in the failover timing calculation, the test measured only the time required to complete the actual database failover. Total time to complete failover ranged between 10 to 25 seconds, depending on the configuration.

Single-Instance Databases

By using the best practices described in the [Failover Best Practices](#) section of this white paper, the following average failover times were noticed for single-instance databases:

	Manual Failover		Fast-Start Failover
	SQL*Plus Statements	DGMGRL or Enterprise Manager	
Physical Standby	17 seconds	18 seconds	17 seconds
Logical Standby	10 seconds	12 seconds	14 seconds

Multiple-Instance Real Application Clusters

By using the best practices described in the [Failover Best Practices](#) section of this paper, the following average failover times were observed for multiple-instance Oracle RAC configurations:

	Manual Failover		Fast-Start Failover
	SQL*Plus Statements	DGMGRL or Enterprise Manager	
Physical Standby	22 seconds	25 seconds	25 seconds
Logical Standby	14 seconds	17 seconds	16 seconds

Failover results listed for Oracle RAC require Oracle Database version 10.2.0.2.0 or higher. This release includes an optimization to SHUTDOWN ABORT all secondary standby instances, greatly reducing total failover time. To achieve these times with version 10.2.0.1, manually issue a SHUTDOWN ABORT on each secondary standby instance prior to the failover.

Note: During testing, all instances were started so as to simulate the worst-case scenario. However, as a best practice, all secondary standby instances should be closed (using SHUTDOWN ABORT) prior to performing the failover further reducing the overall time needed.

SWITCHOVER

The Data Guard switchover feature is significant for customers who need to reduce system downtime while maintaining a high level of availability. Switchover accomplishes this by providing a method with which the administrator transitions the production database to a standby role, and transitions the standby database to the production role. The role transition occurs with no data loss.

Once the production role has been transitioned, maintenance operations such as operating system or hardware upgrades can occur without affecting application processing. Once maintenance operations are complete, the administrator can simply transition the production role back to the original site. In the same fashion, switchovers can be used to perform rolling upgrades of the Oracle database software as well as testing disaster recovery preparedness.

Switchovers can be performed using Oracle Enterprise Manager, the Data Guard broker command-line interface, or SQL*Plus statements. As part of the switchover, all user sessions are disconnected from the production database. Once all sessions are removed, the production database is converted to the standby role after which the standby database is transitioned to the production role.

If the original production database is still accessible and you desire to execute a role transition, you should always use a Data Guard switchover as opposed to performing a failover. The Data Guard switchover feature proves to be an excellent solution for customers who need to reduce system downtime while maintaining a high level of availability.

Minimizing Application Downtime During Switchover

In order to minimize overall application downtime during switchover operations, application connections should seamlessly and automatically reconnect to the database that is being transitioned to the primary database role. Configuring application connections involves:

1. Creating a database service that the application uses for its connection.
2. Configuring the database service used by the application to automatically migrate to the new primary database after switchover.
3. Configure the application Oracle Net alias or JDBC URL to include all host that have the potential host the primary database.
4. Automating application reconnections
 - a. OCI applications can use Transparent Application Failover (TAF) to automate reconnection for existing connections
 - b. JDBC applications can use FAN if the primary and standby are in an Oracle RAC environment. For non Oracle RAC environments JDBC applications can be coded to automatically reconnect upon an exception

For detailed information on configuring your application for client failover please refer to the [MAA Best Practices for Client Failover in Data Guard Configurations for Highly Available Oracle Databases](#) [10]

While the MAA client failover paper discusses client failover for physical standby database the same steps will work for a logical standby with the exception of FAN OCI support. While FAN OCI is not supported with a logical standby it is not needed for switchover, as application connections are cleanly disconnected.

In addition to the AFTER STARTUP ON DATABASE trigger discussed in the client failover paper, a logical standby requires an additional AFTER DB_ROLE_CHANGE ON DATABASE trigger. Both trigger types are necessary to manage service relocation in a logical standby client failover configuration. Logical standby requires a service management trigger written around the “ON DATABASE STARTUP” system event to start the service if the instance is bounced after the role transition has been completed as well as around the “DB ROLE CHANGE” system event so that the service is started immediately after the role transition. This is due to the fact that a logical standby instance is already

started (open) and does not need to be restarted prior to becoming the primary. The following is an example of the two triggers necessary for a logical standby:

```
CREATE OR REPLACE TRIGGER manage_service1
after startup on database
DECLARE
    role VARCHAR(30);
BEGIN
    SELECT DATABASE_ROLE INTO role FROM V$DATABASE;
    IF role = 'PRIMARY' THEN
        DBMS_SERVICE.START_SERVICE('sales');
    END IF;
END;

CREATE OR REPLACE TRIGGER manage_service2
after db_role_change on database
DECLARE
    role VARCHAR(30);
BEGIN
    SELECT DATABASE_ROLE INTO role FROM V$DATABASE;
    IF role = 'PRIMARY' THEN
        DBMS_SERVICE.START_SERVICE('sales');
    END IF;
END;
```

Switchover Steps For Physical Standby Databases

Detailed switchover steps as well as pre and post switchover operations for physical standby databases, using either SQL*Plus, Data Guard Broker, or Enterprise Manager are described in MetaLink Note 751600.1.

The high level steps include:

1. Disconnecting user sessions and disabling or stopping application processing.
2. If the production and standby databases are in an Oracle RAC configuration, first shutting down all but one of the production instances. Once the primary database has only one instance up then shut down all standby instances except the apply instance (this will leave a single instance running on each cluster).
3. Converting the primary database to a standby database.
4. Converting the original standby database to be the new primary database.
5. If the production and standby databases are configured in an Oracle RAC, then start all instances.
6. Restart the user sessions and application processing.

Using SQL*Plus and Logical Standby Databases

When performing a switchover using SQL*Plus statements, it is possible for the standby database that is to become the new production database to build and transmit the LogMiner dictionary to the current production database (the new

standby database) prior to the actual switchover. This reduces the total time needed to perform the switchover. The following steps describe how to perform this optimized method:

1. Issue the following statement on the production database to enable receipt of redo from the current standby database:

```
ALTER DATABASE PREPARE TO SWITCHOVER TO LOGICAL STANDBY;
```

2. On the current logical standby database, build the LogMiner dictionary and transmit this dictionary to the current production database:

```
ALTER DATABASE PREPARE TO SWITCHOVER TO PRIMARY;
```

Depending on the work to be done and the size of the database, the prepare statement may take some time to complete.

3. Verify the LogMiner Multiversioned Data Dictionary was received by the production database by querying the SWITCHOVER_STATUS column of the V\$DATABASE fixed view on the production database.
4. Initially, the SWITCHOVER_STATUS column shows PREPARING DICTIONARY while the LogMiner Multiversioned Data Dictionary is being recorded in the redo stream. Once this has completed successfully, the column shows PREPARING SWITCHOVER. When the query returns the TO LOGICAL STANDBY value, you can proceed to the next step.

note: for additional information refer to [“Switchovers Involving a Logical Standby Database” in Oracle Data Guard Concepts and Administration \[5\]](#) for step-by-step instructions.

5. If possible, disconnect user sessions and disable or stop application processing.
6. If the production and standby databases are in an Oracle RAC configuration, then cleanly shut down all but one of the production instances. Once the primary database has only one instance available then shut down all standby instances except the apply instance (this will leave a single instance running on each cluster). To optimize the shutdown of the standby instances you can use SHUTDOWN ABORT. Disable threads for all production and standby instances that were shutdown. You can re-enable the threads and start the instances once the switchover operation has completed successfully.
7. When the SWITCHOVER_STATUS column of the V\$DATABASE view returns TO LOGICAL STANDBY, convert the production database to a standby by issuing:

```
ALTER DATABASE COMMIT TO SWITCHOVER TO LOGICAL STANDBY WITH  
SESSION SHUTDOWN;
```

8. Issue the following statement on the old standby database:

```
ALTER DATABASE COMMIT TO SWITCHOVER TO PRIMARY;
```
9. If the production and standby databases are configured in an Oracle RAC, then start all instances.
10. Restart the user sessions and application processing.

TEST RESULTS FOR SWITCHOVERS

Both single instance databases and Oracle RAC configurations were tested. Tested configurations included switchover to a physical standby database (Redo Apply), and a logical standby database (SQL Apply). Total time to complete a switchover via SQL*Plus ranged from 50 to 55 seconds, depending on the configuration.

Single-Instance Databases

Using the best practices described in the [Switchover Best Practices](#) section of this white paper, the testing resulted in switchover times ranging from 50 seconds to 2 minutes and 49 seconds for single-instance databases.

The following table shows the total time needed to perform a switchover with single instance production and logical standby databases.

	Switchover using SQL*Plus	Switchover using DGMGRL or Enterprise Manager
Physical Standby	0:52	2:49
Logical Standby	0:50	1:48

The switchover timings were achieved by performing a switchover using SQL*Plus statements and using the optimal switchover methods described earlier. This method performed the restart of the new standby database old production database at the same time that the old standby (new production) was being converted. In addition, the new production database transitions directly to the OPEN state from the MOUNT state with no need for a database restart.

A switchover performed using Enterprise Manager takes longer because of the sequence in which the instances were restarted during the switchover and because the new production database was restarted. In addition, Data Guard Broker processing time contributed to the overall switchover time.

Multiple-Instance Oracle Real Application Clusters

Using the best practices described in the [Switchover Best Practices](#) section of this white paper, the testing resulted in switchover times for Oracle RAC databases ranging from 53 seconds to 2 minutes and 56 seconds.

	Switchover using SQL*Plus	Switchover using DGMGRL or Enterprise Manager
Physical Standby	0:55	2:56
Logical Standby	0:53	1:54

Testing for Oracle RAC switchovers was performed with all production and standby instances started. The time illustrated in the above table represents the time needed to transition the role from standby to production database, and restart the new standby. It does not represent the time needed to restart the secondary production and standby database instances.

Broker-based logical standby switchover times are higher due to the fact that switchovers with SQL*Plus were fully prepared prior to switchover (using the ALTER DATABASE PREPARE TO SWITCHOVER ... statement) while the broker-managed switchovers do not make use of the functionality.

APPLICATION AND CLIENT FAILOVER

Choosing and implementing the architecture that best fits your availability requirements can be a daunting task. A highly available architecture must achieve not only fast database failover, but it must also address client failover for all types of failures.

New Data Guard 10g Release 2 features provide the added capability to integrate automatic database failover with failover procedures at the middle tier to quickly and automatically redirect clients and applications to the new production database at the standby location, providing an end-to-end solution for achieving business continuity.

Best practices for client failover in a Data Guard configuration are described in [MAA Best Practices for Client Failover in Data Guard Configurations for Highly Available Oracle Databases](#) [10].

CONCLUSION

Data Guard 10g Release 2 enhancements along with the best practices described in this white paper can help you to achieve faster role transitions by overcoming the most common problems:

- Failure detection and reaction can be slow and time consuming. The time required to locate and notify the administrator can be lengthy. Data Guard Fast-Start Failover automatically detects failures and promptly executes failovers when required.
- Assessment of the problem can be time consuming. Determining if the failure warrants a failover adds even more time. Data Guard Fast-Start Failover operates according to established criteria and initiates failover automatically when criteria are met.
- Correct procedure for failover must be followed to incur the least amount of data loss. Data Guard Fast-Start Failover eliminates the chance for human error to impact the failover procedure.
- Rebuilding the old production database following a failover can consume considerable time and resources, and leave the business dangerously exposed to a second failure until the process is completed. Following a fast-start failover, the Observer periodically attempts to contact the old production database. If a reconnection to the old production database is made, the Observer automatically reinstates the old production database so that it can become a standby database to the new production database. These features quickly restore high availability to the Data Guard configuration
- Restarting databases after a switchover or failover can be time consuming. Beginning with Oracle Database 10g Release 2, you can open the new production database from the mount state if the database was previously a physical standby database that was not opened read-only since the last time the database was started.
- Manual database failover can be error prone or problematic, especially because it is an inherently stressful situation to begin with. When Fast-Start Failover is enabled, Data Guard automatically fails over to a previously chosen, synchronized standby database in the event of loss of the production database, without any data loss and without any manual intervention, thus minimizing the probability of any errors that may occur in case of manually administered failover.

REFERENCES

1. Oracle Data Guard
<http://www.oracle.com/technology/deploy/availability/htdocs/DataGuardOverview.html>
2. Oracle Maximum Availability Architecture
<http://www.oracle.com/technology/deploy/availability/htdocs/maa.htm>
3. Oracle Data Guard 10g Release 2 - Fast-Start Failover Best Practices
<http://www.oracle.com/technology/deploy/availability/htdocs/maa.htm>
4. Oracle Database High Availability Overview (Part #B14210-01)
http://download-west.oracle.com/docs/cd/B19306_01/server.102/b14210/toc.htm
5. Oracle Data Guard Concepts and Administration (Part #B14239-01)
http://download-west.oracle.com/docs/cd/B19306_01/server.102/b14239/toc.htm
6. Oracle 10gR2 SQL Apply Best Practices (for logical standby databases)
http://www.oracle.com/technology/deploy/availability/pdf/MAA_WP_10gR2_SQLApplyBestPractices.pdf
7. Oracle 10gR2 Redo Apply and Media Recovery Best Practices (physical)
http://www.oracle.com/technology/deploy/availability/pdf/MAA_WP_10gR2_RecoveryBestPractices.pdf
8. Fast-Application Notification (FAN) references:
 - *Oracle Clusterware and Oracle Real Application Clusters Administration and Deployment Guide*
http://download-west.oracle.com/docs/cd/B19306_01/rac.102/b14197/hafeats.htm#sthref428
 - *Oracle Database High Availability Overview* (Part #B14210-01)
http://download-west.oracle.com/docs/cd/B19306_01/server.102/b14210/hafeatures.htm#sthref54
9. Automatic Storage Management (ASM) references:
 - *Oracle Database Administrator's Guide* (Part #B14231-01)
http://download-west.oracle.com/docs/cd/B19306_01/server.102/b14231/storeman.htm#i1021337

Oracle Database High Availability Overview (Part #B14210-01)
http://download-west.oracle.com/docs/cd/B19306_01/server.102/b14210/hafeatures.htm#sthref43
10. MAA Best Practices for Client Failover in Data Guard Configurations for Highly Available Oracle Databases
http://www.oracle.com/technology/deploy/availability/pdf/MAA_WP_10gR2_ClientFailoverBestPractices.pdf
11. Transparent Application Failover (TAF) – Oracle Clusterware and Oracle Real Application Clusters Administration and Deployment Guide
http://download-west.oracle.com/docs/cd/B19306_01/rac.102/b14197/hafeats.htm#sthref428

Maximum Availability Architecture

12. Oracle Database Backup and Recovery Basics (Part # B14192-02)
http://download-west.oracle.com/docs/cd/B19306_01/backup.102/b14192/toc.htm
13. Oracle Data Guard Broker (Part #B14230-01)
http://download-west.oracle.com/docs/cd/B19306_01/server.102/b14230/toc.htm
14. Oracle Database High Availability Best Practices 10g Release 2
http://download.oracle.com/docs/cd/B19306_01/server.102/b25159/toc.htm



Oracle Data Guard 10g Release 2 Switchover and Failover Best Practices
July 2010

Author: Mike Smith , Lawrence To, and Viv Schupmann

Contributing Authors: Joseph Meeks and 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

Copyright © 2007, Oracle. All rights reserved.

This document is provided for information purposes only and the contents hereof are subject to change without notice.

This document is not warranted to be error-free, nor subject to any other warranties or conditions, whether expressed orally or implied in law, including implied warranties and conditions of merchantability or fitness for a particular purpose. We specifically disclaim any liability with respect to this document and no contractual obligations are formed either directly or indirectly by this document. This document may not be reproduced or transmitted in any form or by any means, electronic or mechanical, for any purpose, without our prior written permission. Oracle, JD Edwards, and PeopleSoft are registered trademarks of Oracle Corporation and/or its affiliates. Other names may be trademarks of their respective owners.