Oracle9*i* Data Guard Switchover/Failover Best Practices

An Oracle White Paper
November 2004

Oracle Data Guard Switchover/Failover Best Practices

Oracle9*i* Data Guard Switchover/Failover Best Practices

EXECUTIVE SUMMARY

Oracle Data Guard is one of the most effective and comprehensive data protection and disaster recovery solutions available today for enterprise data.

This paper provides best practices to configure Oracle9*i* Data Guard switchover and failover operations, and to optimize switchover and failover times. By following these recommendations, you can reduce switchover and failover timings by 30-50% in many cases.

Oracle Data Guard is the management, monitoring, and automation software infrastructure that creates, maintains, manages and monitors one or more standby databases to protect enterprise data from failures, disasters, errors, and corruptions. It maintains these standby databases as transactionally consistent copies of the production database. 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 enabling zero data loss.

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 respectively. They can be easily initiated directly through SQL, or Enterprise Manager, or the Data Guard Broker's command line interface (DGMGRL). This paper focuses on role management using SQL within SQL*Plus.

A failover operation can be invoked when an unplanned catastrophic failure occurs on the primary database, and there is no possibility of recovering the primary database in a timely manner. To bring back this primary database as a standby database in the Data Guard configuration, it has to be re-instantiated from an available backup in Oracle9*i*.

The switchover, in contrast to the failover process, is a planned role reversal of the primary and standby databases to handle planned maintenance on the primary host. A switchover operation does not require re-instantiation of the primary database. This allows the primary database to assume the role of a standby database and the standby database transitions to the primary role. As a result, testing disaster recovery preparedness and 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.

With physical standby database, we have observed switchover and failover times of less than a minute. With logical standby database, we have observed switchover and failover times of less than 30 seconds. A formula and spreadsheet are also provided to help predict your current application's switchover and failover timings. Furthermore, a case study illustrates a typical customer issue and how applying the recommendations reduced the switchover time.

For details on managing application servers and clients during role management see the MAA website [2]. Please refer to this site for the location and future updates of this paper.

BEST PRACTICES

These best practices were derived after testing on Oracle9*i* release 9.2.0.3 and later databases as part of the ongoing studies within the Maximum Availability Architecture (MAA) best practices and recommendations. For more information about MAA, refer to the MAA website [2].

Descriptions of the test environment, test cases (this includes the actual commands), and test results used to identify these best practices are included in subsequent sections of this paper. For further detail on command syntax consult the "Oracle Data Guard Concepts and Administration Release 2 (9.2) " documentation.

This section outlines

- Switchover best practices for physical and logical standby databases
- Failover best practices for physical and logical standby database
- Role Transition timing calculation

Switchover

If the original production database is still accessible, you should always consider a Data Guard switchover first. In Oracle9*i*, a failover requires that the initial production database needs to be reinstantiated as a new standby database, which can be a very expensive operation. In contrast, switchover, which is a planned operation, offers the capability to switch database roles between the production and standby databases without needing to instantiate any of the databases. Switchover can occur whenever a production database is started, the target standby database is available, and all the archived redo logs are available. It is useful in the following situations:

- Scheduled maintenance such as hardware maintenance (e.g. hardware or firmware patches) on the production host
- Resolution of data failures when the production database is still opened
- Testing and validating the secondary resources, as a means to test disaster recovery readiness

Physical Standby Best Practices

- Clear the online redo logs for a "new" standby (following an instantiation or a switchover) 'ALTER DATABASE CLEAR LOGFILE GROUP <n>' command for each online redo log group.
- Use standby redo logs with a physical standby to reduce redo transfer time for unapplied redo
- Run the 'SWITCHOVER TO PRIMARY' command immediately following a successful "SWITCHOVER TO STANDBY" command and in parallel with the shutdown/startup of the new standby database.
- Follow the "Oracle9*i* Media Recovery Best Practices" white paper (available at the MAA website [2]) for recommendations on how to obtain the optimal Redo Apply rate.
- Follow a pre-transition checklist: (details in Section 9.7.2, "*Physical Standby Database Switchover"* of the "MAA Detailed White Paper [1]"
  - Check the status of Log Transport services

- Verify that there are no gaps
- Set the redo apply delay settings to 0, e.g. 'RECOVER MANAGED STANDBY DATABASE NODELAY DISCONNECT'
- Record the current online redo log thread and sequence number(s) on the primary and on the standby
- For a RAC database, ensure that only a single instance is running
- End all jobs and sessions on the remaining active production instance
- Validate that the priomary database SWITCHOVER_STATUS is "TO STANDBY".
- For multiple standby configurations involving a logical standby see Appendix A for the detailed steps

- Understand the factors that affect the physical standby switchover time and test for planned outage timing estimates and use a formula for estimating the switchover time. The factors are:
  - Time taken by the primary for switching over to be a standby
  - Time taken by the standby for switching over to be a primary
  - Redo generation rate at the primary database
  - Redo Apply rate at the standby database
  - Redo apply delay settings
  - Primary and standby database shutdown and startup time
  - Use of LGWR vs. ARCH for redo transport for a physical standby
  - Network round trip time (RTT)
- Follow a post-transition checklist: (details in Section 9.7.2, "*Physical Standby Database: Post- Switchover Steps"* of the "MAA Detailed White Paper [1]"
  - Clear Online Redo log Groups on the Standby Database
  - Check Local and Remote Archive Destinations on the Production database
  - Ensure that the Lag is Set Up Correctly w Archived Redo Logs
  - Ensure that Recovery is Applying New Archived Redo Logs to all standby databases

Logical Standby Best Practices

- Create database links in both directions during the logical standby creation process
- Follow a pre-transition checklist: (details in Section 9.7.3, "*Logical Standyby Database Switchover*" of the "MAA Detailed White Paper [1])
  - Execute a log switch
  - Remove any apply delay
  - Logging off all users and ending all jobs will reduce the time for the "COMMIT TO SWITCHOVER TO LOGICAL STANDBY" command.
- Follow the "Oracle9*i* Data Guard: SQL Apply Best Practices" white paper (available at the MAA website [2].) for recommendations on how to obtain the optimal SQL Apply rate..

Failover

Data Guard failover should be used only when switchover is not possible as a viable role transition option. Typically, a failover may be initiated in the event of an unplanned outage such as:

- Site disaster (e.g. loss of the data center building and it's contents)
- Logical failures
- Data failures

If the original production database is still accessible, you should always consider attempting to fix the problem locally in a timely fashion or using a Data Guard switchover first.

Physical Standby Best Practices

- Clear the online redo logs for a "new" standby (following an instantiation or a switchover) (ALTER DATABASE CLEAR LOGFILE GROUP <n>' command for each online redo log group.
- Use SRL's to reduce data loss
- If the RFS processes are still active, most likely because the primary database can still be communicated with, then either shutdown the primary database or manually kill the RFS process (es). RFS process ID's can be obtained from the V$MANAGED_STANDBY view on the standby database.
- Follow the "Oracle9*i* Media Recovery Best Practices" white paper (available at the MAA website [2].) for recommendations on how to obtain the optimal Redo Apply rate.

Logical Standby Best Practices

- Follow the "Oracle9i Data Guard: SQL Apply Best Practices" white paper (available at the MAA website [2].) for recommendations on how to obtain the optimal SQL Apply rate.

ESTIMATING ROLE TRANSITION TIME

Based on the test results a generic formula can be developed for estimating the switchover and failover timings. The formula differs for physical standby versus logical standby and for a switchover versus a failover. This formula can be used for more accurate planned outage estimates when doing a switchover and for estimating the time until the system is available when a failover is necessary.

In the worst-case, this estimate can vary by as much as the application for a single redo log. These formulas do not take into account the existence of any archive log gaps on the standby. In an optimally monitored system there should be no gaps prior to a switchover.

If there are gaps when a failover is attempted and if the primary is still accessible, the logs can be manually transferred and registered on the standby. If the primary is inaccessible then an incomplete recovery failover would have to be done.

Lastly, these estimates do not account for any application shutdown or startup time, i.e. any they do not account for any timings external to Oracle Data Guard. As with any estimate this should still be tested to validate the estimate as times may also vary depending on the network latency between the primary and standby sites.

A spreadsheet that incorporates the formulas below is available at the MAA website.

Physical Standby

Generic Variables (switchover & failover)

- Redo Generation Rate

  'redo size' 'per second' from statspack snapshot on primary during peak load, unit=K/sec

- Redo Apply Rate

  Set this event on the standby:

  event="10871 trace name context forever, level 1"

Once this event is set then you can use the timestamps from the alert.log to calculate the apply rate as follows:

Log-file-size / AVG ((Log-close-time – log-start-time)), calculate in K/sec

e.g. log file size = 500Mb and the average time to apply is

- New Primary Restart time estimate
- Standby Apply Delay (how much redo still needs to be applied)

- Standby Switchover Command time
  (ALTER DATABASE COMMIT TO SWITCHOVER TO PRIMARY;)
  Estimated time to complete the standby to primary switchover command.

Switchover Variables

- Primary Switchover Command time
  (ALTER DATABASE COMMIT TO SWITCHOVER TO STANDBY;)
  Estimated time to complete the primary to standby switchover command.
- New Standby Restart estimate time
  this time is only necessary in the calculation if the new standby shutdown and startup is not run in the background in parallel to the standby to primary switchover steps.

Failover Variables

Failover Command time
(RECOVER MANAGED STANDBY DATABASE FINISH;)
Estimated time to complete the standby failover command.

Switchover Estimation Formula Using SQL

SO_TIME = Primary Switchover Command time +
Standby Switchover Command time +
New Primary Restart time +
New Standby Restart time +
(((Apply Delay time * 60) * Redo Rate) / Apply Rate)

Note: This formula includes the new standby restart time, but in our case since the best practice is to do this in parallel. If it's deemed that restarting the new standby cannot be done in parallel then the new standby restart time would be non-zero.

e.g. Primary Switchover Command time=10, Standby Switchover Command time=10, New Primary Restart time=35, Apply Delay time=30 mins. (1800 secs.), Redo Rate=750 K/sec, Apply Rate=2000 K/sec, New Standby Restart time=0

SO_TIME=55 + (1800 * 750) / 2000 = 55 + 675 = 730 secs. = 12:10

Whereas using parallel recovery 24 changes the apply rate to 4500 K/sec

SO_TIME=55 + (1800 * 750) /4500 = 55 + 300 = 355 secs. = 5:55

Failover Estimation Formula Using SQL

FO_TIME= Failover Command time +
Standby Switchover Command time +
New Primary Restart time +
(((Apply Delay time * 60) * Redo Rate) / Apply Rate)

e.g. Failover Command time=11, Standby Switchover Command time=3, New Primary Restart time=35, Apply Delay time =30 (mins.), RR Redo Rate=750 K/sec, Apply Rate=2000 K/sec

FO_TIME=49 + (1800 * 750) / 2000 = 49 + 675 = 724 secs. = 12:04

Whereas using parallel recovery 24 changes the apply rate to (AR) to 4500 K/sec

FO_TIME=49 + (1800 * 750) /4500 = 49 + 300 = 349 secs. = 5:49

Logical Standby

Generic Variables (switchover & failover)

- Redo Rate

  'redo size' from statspack snapshot on primary during peak load,

  unit=K/sec
- Logical Apply Read Rate
- Standby Apply Delay (how much redo still needs to be applied)
- Standby Switchover Command
  (ALTER DATABASE COMMIT TO SWITCHOVER TO PRIMARY;)

Switchover Variables

- Primary Switchover Command time
  (ALTER DATABASE COMMIT TO SWITCHOVER TO LOGICAL STANDBY;)
- Logical Apply Start time
  (ALTER DATABASE START LOGICAL STANDBY APPLY;)

Failover Variables

- Logical Apply Stop time
  (`ALTER DATABASE STOP LOGICAL STANDBY APPLY;`)
- Failover Command time
  `(ALTER DATABASE ACTIVATE LOGICAL STANDBY DATABASE;)`

Switchover Estimation Formula Using SQL

SO_TIME= Primary Switchover Command time +
Standby Switchover Command time +
Logical Apply Start time +
(((Apply Delay time * 60) * Redo Rate) / Apply Read Rate)

e.g. Primary Switchover Command time=16, Standby Switchover Command time=16, Apply Delay=30 mins. (1800 secs.), Redo Rate=750 K/sec, Apply Read Rate=1200 K/sec, Logical Apply Start time=11

SO_TIME=43 + (1800 * 750) / 1200 = 43 + 1125 = 1168 secs. = 19:28

Whereas eliminating the apply delay would remove the 1125 seconds and effectively take the logical switxchover about 43 seconds.

Failover Estimation Formula Using SQL

FO_TIME= Logical Apply Stop time +
Failover Command time +
(((Apply Delay time * 60) * Redo Rate) / Apply Read Rate)

e.g. Logical Apply Stop time=8, Failover Command time=9, Apply Delay=30 mins. (1800 secs.), Redo Rate=750 K/sec, Apply Read Rate=1200 K/sec,

FO_TIME=17 + (1800 * 750) / 1200 = 17 + 1125 = 1142 secs. = 19:02

Whereas eliminating the apply delay would remove the 1125 seconds and effectively take the logical failover about 17 seconds.

CASE STUDY

This section is intended to illustrate how the best practices can reduce the time for a planned outage that uses switchover. This scenario is using the 9.2.0.3 RDBMS release.

Physical Standby Switchover

Background

- A Data Guard customer is using physical standby database and achieving 16-20 minutes for a switchover with their physical standby database.
- They would like to reduce the switchover time to be less than 10 minutes.
- Physical standby details:
  - Using maximum performance protection mode.
  - Using the ARCH transport with 500 Mb online redo logs, log switches average every 11 minutes. Averages 4 minutes to transfer a log to the standby.
  - Physical standby database is 380 miles away with a network round trip time (RTT) of 12 ms and a bandwidth of 100 Mbps.
  - They have a redo apply delay of 30 minutes.
  - Peak redo rate is 800 K/sec
  - Redo apply rate is 1600 K/sec

Both systems, primary and standby, are 4 CPU single node (non-RAC) systems employing the stripe and mirror everything (SAME)[3] methodology for the disk layout. Analysis

After reviewing and assessing their environment, the following best practices are recommended:

- Following any switchover or reinstantiation, the first step on the new physical standby is to clear the standby's online redo logs.
- Based on the "Oracle9*i* Media Recovery Best Practices" white paper the following changes were made on the physical standby to increase the redo apply rate to from 1600 K/sec to 3300 K/sec:
  - Set parallel recovery=2 x CPUs to optimize the redo apply rate.
  - Disabled data block checking for faster redo apply rates, DB_BLOCK_CHECKING=FALSE.
- Execute 'commit to primary' immediately after a successful 'commit to standby' command, running the new standby restart in parallel (in the background).
- They have switched the transport from ARCH to LGWR ASYNC with standby redo logs, thereby eliminating the potential for having to transfer a

complete archive log which could be an additional 4 minutes to the switchover. This change also reduces the data loss potential.

- Reducing the apply delay from 30 minutes to 20 minutes was investigated and tested and this would further reduce the switchover time by about 2 minutes and 30 seconds. However, the 30-minute delay time was left intact since the other best practices made a significant enough impact to the switchover time.
- Log out users prior to switchover.

The above changes improved their switchover time from 16 minutes and 41 seconds to 8 minutes and 11 seconds, a 51% reduction. They could further reduce switchover times to 5 minutes and 45, seconds, a 65% reduction, by reducing the apply delay to 20 minutes but chose to keep the 30 minute delay in based on their ability to react to any primary database corruptions.

This scenario illustrates how following the best practices identified in this paper can reduce switchover time by over 50%.

TEST DESCRIPTION

The following experiments are documented:

- Performance of switchover for a physical standby in a LAN and simulated WAN network environment.
- Performance of failover for a physical standby.
- Performance of switchover for a logical standby in a LAN
- Performance of failover for a logical standby in a LAN

Tests were run for physical and logical standby databases. These tests were run in separate environments as described in the Test Environment section. For each of the standby database types, physical and logical, switchover and failover tests were run.

Physical Standby Tests

Switchover

The scripts in Appendix B were used for these tests. To summarize, the following tests were run:

- **Serial Test** (the background submit for the new standby start was removed) Switchover with all steps run serially on a LAN.

  Serial Steps

  On Primary

  1. ALTER DATABASE COMMIT TO SWITCHOVER TO STANDBY with session shutdown;
  2. shutdown immediate
  3. startup nomount
  4. alter database mount standby database ;
  5. alter system set log_archive_dest_state_2=defer;
  6. recover managed standby database disconnect;

  On Standby

  7. ALTER DATABASE COMMIT TO SWITCHOVER TO PRIMARY with session shutdown;
  8. alter system set log_archive_dest_state_2=enable;
  9. shutdown immediate
  10. startup
- **Parallel Test**
  Switchover with the new standby shutdown/startup, steps 2-6 above, run in the background, i.e. in parallel to the 'ALTER DATABASE COMMIT TO SWITCHOVER TO PRIMARY' command.
- **Parallel Test with Pre-clear of standby's Online Redo Logs**
  - Switchover with the new standby shutdown/startup, steps 2-6 above, run in the background, i.e. in parallel to the 'ALTER DATABASE COMMIT TO SWITCHOVER TO PRIMARY' command.

- Additionally, the standby database's online redo logs were cleared prior to the test using the 'ALTER DATABASE CLEAR LOGFILE GROUP <n>' command for each online redo log group.

Failover

The following failover steps were run via SQL*Plus for the Maximum Performance and Maximum Protection protection modes:

1. RECOVER MANAGED STANDBY DATABASE FINISH
2. ALTER DATABASE COMMIT TO SWITCHOVER TO PRIMARY;
3. shutdown immediate
4. startup
5. exit

The following tests were run with the above failover steps:

- Failover without pre-clearing the standby's online redo logs
- Failover with Pre-clear of standby's Online Redo Logs
  - Additionally, the standby database's online redo logs were cleared prior to the test using the 'ALTER DATABASE CLEAR LOGFILE GROUP <n>' command for each online redo log group.

Logical Standby Tests

Switchover

There is no need to shut down and restart any logical standby databases that are in the Data Guard configuration.

On Primary

1. ALTER DATABASE COMMIT TO SWITCHOVER TO LOGICAL STANDBY;
2. ALTER SYSTEM SET log_archive_dest_state_2='DEFER' SCOPE=BOTH;
3. ALTER DATABASE START LOGICAL STANDBY APPLY NEW PRIMARY location1;

On the original logical standby database

4. ALTER DATABASE COMMIT TO SWITCHOVER TO PRIMARY;
5. ALTER SYSTEM SET log_archive_dest_state_2='ENABLE' SCOPE=BOTH;

Failover

```
1. Ensure that all redo logs were mined and applied:
   SELECT APPLIED_SCN, NEWEST_SCN FROM
   DBA_LOGSTDBY_PROGRESS;

2. ALTER DATABASE STOP LOGICAL STANDBY APPLY;
3. ALTER DATABASE ACTIVATE LOGICAL STANDBY DATABASE;
```

Performance Metrics

The test timings were captured using the alert log messages from each database, the primary and the standby. Samples of the alert logs are contained in the "Test

Results Data" Appendix. The general database performance and the operating system performance metrics were monitored as well to make sure there were no bottlenecks in these areas. Tuning for these areas is beyond the scope of this paper but should not be overlooked as part of standard operating procedures. For further details consult the following resources:

MAA Papers
http://otn.oracle.com/deploy/availability/htdocs/maa.htm

Oracle 9i Performance Tuning Guide and Reference
http://download-west.oracle.com/docs/cd/b10501_01/server.920/a96533/toc.htm

TEST RESULTS

Physical Standby

As can be seen by the chart below, following the best practices;

1. shutdown/startup the new standby in parallel to the standby 'switchover to primary' command, and

2. pre-clear the standby database online redo logs,

can improve the switchover time by over a minute for a physical standby.

The startup and shutdown time for the new primary database for switchover or failover will depend on the database characteristics: e.g. database size, SGA size, and number of datafiles.

Switchover

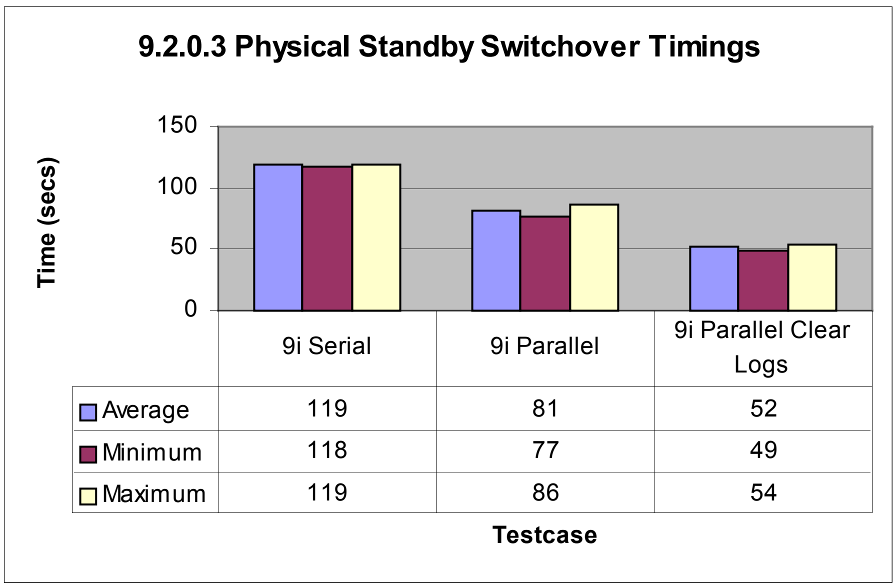

	9i Serial	9i Parallel	9i Parallel Clear Logs
Average	119	81	52
Minimum	118	77	49
Maximum	119	86	54

Figure 1 Switchover Test Summary

Serial Testcase

As it says, the switchover steps were run serially as described under the "Physical Standby Tests". This is virtually the same steps that the Data Guard Manager uses to do the switchover.

Parallel Testcase

This method eliminates the time required to shutdown and startup the new standby database since following the completion of the primary being switched to a standby the switchover to primary of the current standby begins rather than waiting for the new standby to recycle as in the serial testcase. The new standby shutdown/startup is done in parallel by running submitting it as a background job.

Parallel with the Standby's Online Logs Pre-cleared Testcase

Pre-clearing the standby database online redo logs saves the switchover operation from having to do it. Clearing the online redo logs prior to the switchover saved from 10-20 seconds during the switchover. To clear the standby database online redo logs requires managed recovery to be stopped.

Failover

Summary

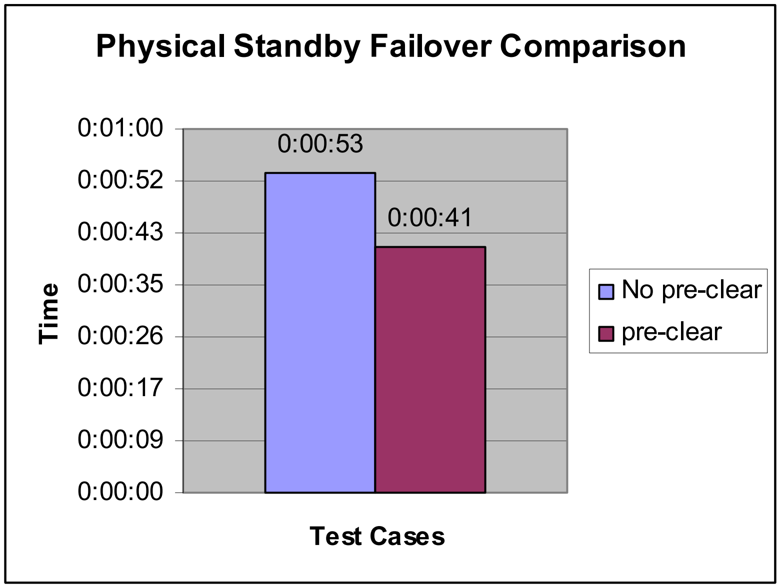

Figure 2 Physical Standby Failover

Logical Standby

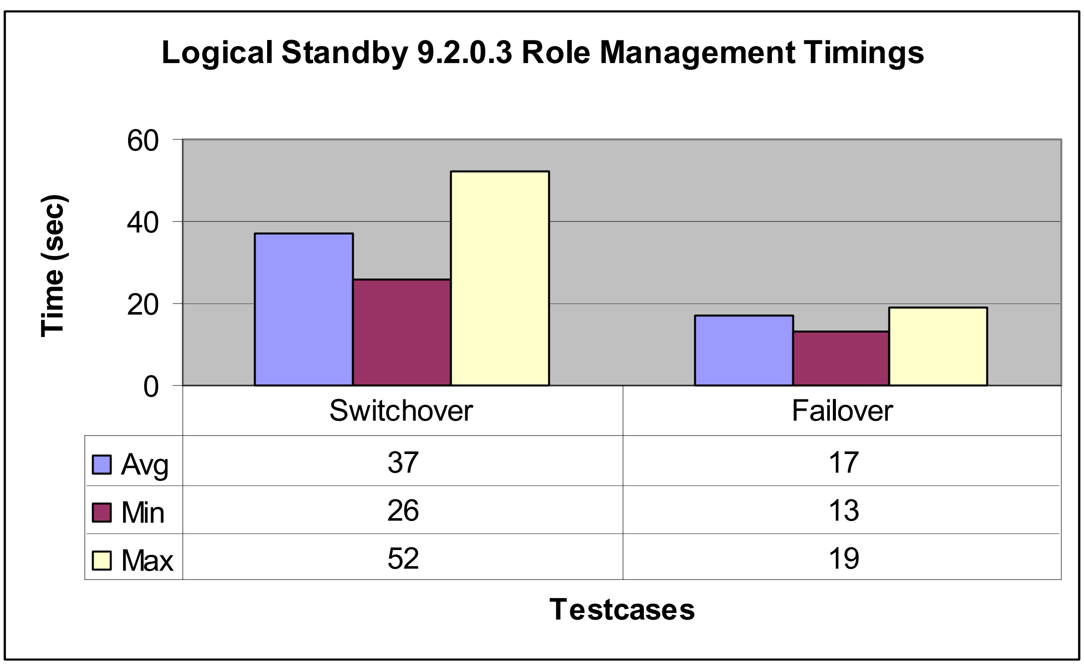

CONCLUSION

Understanding, testing and using Data Guard switchover and failover are key to meeting service availability. Using the identified best practices in testing and implementing Data Guard will give optimal and consistent timings for switchover and failover during planned and unplanned outages. This knowledge and testing experience will assist in meeting application service level agreement (SLA) requirements consistently.

APPENDIX

A. Test Environment

The environment was setup to follow the recommendations from the Maximum Availability Architecture (MAA) paper. See the Oracle Technology Network at http://otn.oracle.com/deploy/availability/htdocs/maa.htm for details on those recommendations.

Physical Standby Environment

Database

1. Two separate databases were used, both with a 2-node RAC primary and a 2-node RAC standby using a single node from the cluster:

2. A TPCC type environment with a simplified OLTP transaction profile:

- Storage: approximately 200 GB
- full archive files being 256 MB in size.
- Total System Global Area 1125996256 bytes (1.073 GB)

```
Fixed Size                      456416 bytes
Variable Size                335544320 bytes
Database Buffers             788529152 bytes
Redo Buffers                   1466368 bytes
```

3. A small functional database with no load being generated against it::

- Storage: approximately 2 GB
- full archive files being 50 MB in size.
- Total System Global Area 287134452 bytes (275 MB)

```
Fixed Size                      455412 bytes
Variable Size                117440512 bytes
Database Buffers             167772160 bytes
Redo Buffers                   1466368 bytes
```

Storage for both databases was configured using the Stripe And Mirror Everything[1] (SAME) methodology with a stripe size of 1MB.

The standby system is configured identically to the production environment. Recovery uses half of the CPU, memory and I/O channel resources on the standby host since recovery is limited to one node of the RAC cluster.

Hardware

2-node Primary RAC, 2-node Standby RAC. Each node in the RAC Primary and the RAC Standby has the following configuration:

- 8 400 Mhz CPU's per node, 8Gb memory
- Sun Solaris 2.8 64-bit (SunOS 5.8 Generic_108528-12 sun4u SPARC SUNW, Ultra-Enterprise)

---

[1] For more information about SAME, refer to http://otn.oracle.com/deploy/availability/pdf/oow2000_same.pdf

- EMC SYMMETRIX-SUNAPE Shared disk configured following the SAME[1] methodology, using a 1MB stripe
- Archive destinations on a clustered file system using the SAME methodology, 1 MB stripe size

Software

- Sun Cluster 3.0
- Oracle Enterprise Edition Release 9.2.0.3 – Production with the Partitioning and Real Application Clusters option

Network

- 100 MB/s dedicated private network

Logical Standby Environment

Database

The production database consists of a 2-node Oracle Real Application Clusters (RAC). A TPCC type environment with a simplified OLTP transaction profile.

The standby system was configured identically to the primary system; the logical standby process (LSP0) was run on a single node of the standby system in the Real Application Clusters system.

Storage for both databases was configured using the Stripe And Mirror Everything[2] (SAME) methodology with a stripe size of 1MB.

Hardware

2-node Primary RAC cluster, 2-node Standby RAC cluster. Each node in the RAC Primary and the RAC standby cluster has the following configuration:

- 8 x 440Mhz CPU's per node
- 16GB memory per node
- HP StorageWorks Virtual Array va7100 for file systems and archive destinations
- HP StorageWorks Virtual Array va7400 for the database files using RAID 1+0
- HP HyperFabric cluster interconnect

Software

- HP-UX v11.11 64-bit.
- HP ServiceGuard eRAC edition v11.13
- Oracle Enterprise Edition Release 9.2.0.4.0 – Production with the Partitioning and Real Application Clusters option

Network

- 1 GB/s dedicated private network

---

[2] For more information about SAME, refer to http://otn.oracle.com/deploy/availability/pdf/oow2000_same.pdf

B. References

[1] Maximum Availability Architecture (MAA) paper,

See the Oracle Technology Network at
http://otn.oracle.com/deploy/availability/htdocs/maa.htm

[2] MAA OTN Website

http://otn.oracle.com/deploy/availability/htdocs/maa.htm

[3] Optimal Storage Configuration Made Easy - Stripe and Mirror Everything (SAME) paper

http://otn.oracle.com/deploy/availability/pdf/oow2000_same.pdf

Oracle Data Guard Switchover/Failover Best Practices
November 2004
Authors: Ray Dutcher, High Availability Systems Team
Contributing Authors: Lawrence To, Andrew Babb, Rick Anderson, HA Systems Team, 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