

OTN Case Study: Automatic Failover With Oracle Data Guard Fast-Start Failover

“Fast-Start Failover takes the DBA off the critical path. Database failover is automatic. Data Guard can now address recovery time objectives measured in seconds.”

*Ranjit Singh Veen,
Manager Enterprise Systems Management,
Fannie Mae*

OVERVIEW

Fast-Start Failover

- Automatic Failover capability included with Data Guard 10g Release 2
- Unattended database failover executes in seconds,
- Original primary production database is automatically reinstated as a new standby database after failover (assuming the database is in a state where it can be restarted)
 - Zero data loss
 - External, lightweight “Observer” process ensures failover rules are satisfied and eliminates the possibility of “split brain” processing
- Managed through Data Guard Broker & Oracle Enterprise Manager Grid Control

Fast-Start Failover is an [Oracle Data Guard 10g Release 2](#) [1] feature that automatically, quickly, and reliably fails over to a designated, synchronized standby database in the event of loss of the primary database, without requiring manual intervention to execute the failover.

In addition, following a fast-start failover, the old primary database is automatically reconfigured as a new standby database upon its reconnection to the configuration. This enables Data Guard to restore disaster protection in the configuration, returning the database to a protected state as soon as possible.

This paper provides an overview of its capabilities and describes how Data Guard Fast-Start Failover addresses High Availability requirements.

USER EXPERIENCE – AUTOMATIC FAILOVER

Companies universally recognized as leaders in their industries tested Fast-Start Failover prior to its general release. In each case, the business driver generating interest in Fast-Start Failover is a mission-critical application where both high availability (HA) and disaster recovery (DR) are essential to the business function. For this class of applications, any outage is unacceptable, even in the case of an event that makes an entire data-center unavailable. In such instances the goal is a low Recovery Time Objective, or the time that it takes to have applications functioning at a remote standby site, measured in seconds. Fast-Start Failover test participants included Amazon.com, Thomson Legal and Regulatory, Fannie Mae, and Airbus. Quotes referencing their individual experiences are provided in this case study. A more detailed description of the experience of Amazon.com is provided below.

AMAZON.COM AND AUTOMATIC FAILOVER

At \$6.92 Billion (US) in sales and over 7,800 employees, it is not hard to understand why High Availability and Business Continuity are such important topics at Amazon.com. Continuous availability of Amazon’s customer facing

systems is essential. Any disruption in system availability directly impacts Amazon's bottom line.

“The capability of fast, guaranteed zero-data-loss failover with Fast-Start Failover in Oracle Data Guard takes the availability of an Oracle database platform to new levels. Our initial tests running Oracle Database 10g Release 2, show that Fast-Start Failover offers a magnitude of improvement in availability.”

*—Rajesh Sheth,
Manager, Database
Engineering
Amazon.com*

Historically, Amazon has used custom-maintained standby databases to protect business operations and data that would otherwise be affected by system and site-wide failures. Amazon has built-in automation through custom-developed scripts and processes that enable business continuity even in the event of complete system or site failures. Amazon has determined that automation is critical to meeting aggressive service level agreements.

Fast-Start Failover is designed to meet requirements such as those at Amazon. Fast-Start Failover quickly and reliably fails over a target standby database to the primary database role without requiring manual intervention to invoke the failover. Continually seeking to enhance their high availability architecture, Amazon evaluated Fast-Start Failover during early beta testing of Oracle Data Guard 10g Release 2.

Amazon determined that Fast-Start Failover satisfied their High Availability requirements. For the first time, an out-of-the-box solution is available that Amazon can use instead of supporting and maintaining custom scripts and processes. Amazon's findings were presented at Oracle Open World in September, 2005. They are:

- Fast-Start Failover reliably executes fast, automatic failover to a standby site without human intervention.
- Fast-Start Failover reduced failover time from minutes to seconds. Average failover time was 25 seconds based on a Data Guard primary/standby pair with round-trip network latency of 0.5ms.
- Fast-Start Failover achieves zero data loss. Failover will not commence automatically if the failover target is not synchronized with the primary database.
- Because Fast-Start Failover is based on Data Guard Maximum Availability protection mode, the primary production database is not affected by network or standby failures.
- After failover the old primary database can be automatically reinstated as a new target standby (assuming that the database can be restarted), following which, Data Guard quickly and automatically resynchronizes the old primary with the new primary database. It no longer needs to be restored from a backup of the new primary.

The remainder of this paper provides an overview of Fast-Start Failover. A detailed discussion of Fast-Start Failover is provided in the [Maximum Availability Architecture \(MAA\)](#) [2] best practices paper [Fast-Start Failover: Oracle Database 10g Release 2](#) [3] and documentation for [Oracle Data Guard Broker](#) [4], and [Oracle Data Guard Concepts and Administration](#) [5].

OVERVIEW: FAST-START FAILOVER

There are three essential participants (figure 1) in a Fast-Start Failover configuration:

- The primary database
- The target standby database
- The Fast-Start Failover Observer

The target standby database will become the new primary database following a fast-start failover (note there can be multiple standby databases in a Data Guard configuration).




Figure 1 Fast-Start Failover Configuration

Fast-Start Failover is used in 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.

Fast-Start Failover is managed using the Data Guard Broker command line interface, DGMGRL, or Oracle Enterprise Manager 10g Grid Control. Enterprise Manager provides a GUI interface that interacts with the Data Guard Broker. Not only does this provide a GUI interface for monitoring and control, it also enables centralized management of resources in one or more Data Guard configurations.

The Observer is a separate process incorporated into the DGMGRL client that continuously monitors the primary database and the target standby database for possible failure conditions. Installing the Observer is simple – all that is required is

to install the Oracle Client Administrator (choose the Administrator option from Oracle Universal Installer). Installing the Oracle Client Administrator results in a small footprint because an Oracle instance is not included on the observer system.

The underlying rule is that out of these three participants, (primary, standby & observer) whichever two can communicate with each other will determine the outcome of fast-start failover. For example, if the primary database becomes unavailable, the Observer confirms with the target standby database that the primary database is unavailable and that the target standby database is synchronized with the primary database, and if so, initiates a fast-start failover to the target standby database. Lacking such agreement, an automatic failover cannot occur. This insures two important characteristics of Fast-Start Failover:

- There can be no event where more than one database in a Fast-Start Failover configuration can assume the primary role at the same time. This avoids what is commonly referred to as a “split brain” scenario by guaranteeing that only one database in a Fast-Start Failover configuration is able to accept transactions.
- Automatic failover can only occur if there is a guarantee that zero data will be lost.

Following a fast-start failover, the Observer periodically attempts to contact the old primary database. If a reconnection to the old primary database is made, the Observer automatically reinstates the old primary database so that it can become a standby database to the new primary database. This quickly restores high availability to the Data Guard configuration.

EVENTS THAT TRIGGER FAST-START FAILOVER

The following database conditions will trigger a fast-start failover:

- Database instance failure (or last instance failure in a RAC configuration)
- Shutdown abort (or shutdown abort of the last instance in a RAC configuration)
- Datafiles taken offline due to I/O errors

The following network condition will trigger a fast-start failover:

- When both the Observer and the standby database lose their network connection to the primary database, and when the standby database confirms that it is in a “synchronized” state.

ORACLE TEST RESULTS

Oracle tested a Fast-Start Failover configuration comprised of a primary database, standby database, and observer, all running Redhat Linux 3.0. The results of the test are provided in Figure 2 below.

“Airbus testing of Data Guard 10g Fast-Start Failover produced impressive results. Failover executed in less than a minute. This was much faster than a cold failover using third party cluster technology. With Data Guard, Airbus can achieve continuous data protection and high levels of availability using a standard feature of the Oracle Database.”

*—Werner Kawollek
Application Management
Operations
Airbus Deutschland GmbH*

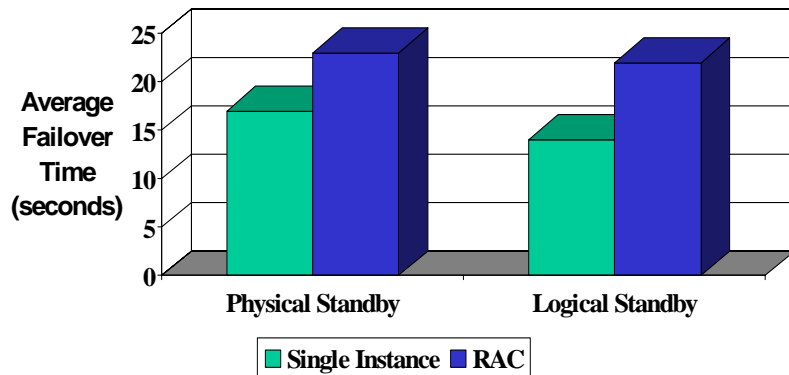


Figure 2 – Fast Start Failover Test Results

The test databases were each 100GB in size. Each host was connected to the next over a Gigabit Network. The workload on the primary database generated 3 MB/second of redo. Both single instance databases and multi-instance RAC configurations were tested. Tested configurations included failover to a physical standby database (Redo Apply), and a logical standby database (SQL Apply). In all cases, the 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 and 25 seconds, depending upon the configuration.

APPLICATIONS WELL SUITED TO FAST-START FAILOVER

Data Guard is typically used to maintain a synchronized standby system in a data center remotely located from the primary production site for the purposes of data protection and disaster recovery. Data Guard has enabled zero data loss protection and fast database failover since Oracle9i, assuming an administrator is immediately available to execute the failover when needed. Enhancements in Oracle Data Guard 10g have further reduced the time it takes to execute a manual failover (for more information please reference the MAA paper, [Oracle Data Guard 10g Release 2 Switchover and Failover Best Practices](#) [6]).

However, there is a class of applications that are extremely sensitive to downtime. Critical manufacturing applications where any downtime translates into lost production, trading systems where downtime results in lost business, online web retailers where downtime directly effects revenue generation and customer satisfaction, to name a few. Businesses with such applications cannot tolerate the additional delay that could result due to a manual process driving the failover. This delay is compounded if an administrator is not immediately available to execute failover when needed.

Fast-Start Failover is very well suited for this class of applications. Fast-Start Failover eliminates the uncertainty of a process that requires manual intervention and automatically executes a zero data loss failover within seconds of an outage being detected.

AUTOMATIC CLIENT FAILOVER

There are several approaches to configuring client failover. In each of these approaches, prior to Fast-Start Failover, accommodations needed to be made for manual intervention required to execute database failover. Fast-Start Failover streamlines the process by making database failover automatic. In addition, Data Guard 10g Release 2 includes a new [DB_ROLE_CHANGE](#) [7] system event that makes it possible to quickly notify clients that a failover has occurred so that they are automatically redirected to the new primary database. For a more detailed discussion of automating client failover, please refer to [Oracle Data Guard 10g Release 2 Best Practices for Client Failover](#) [8].

CONCLUSION

Oracle Data Guard has evolved over a number of major Oracle releases into the most fully functional disaster recovery solution available for the protection of Oracle data and the high availability of applications that require access to that data regardless of the nature or scale of events that impact the primary production system.

Fast-Start Failover further extends Data Guard's ability to address business continuity requirements. Fast-Start Failover monitors the Data Guard configuration 24x7 and executes a failover automatically when specific conditions exist. The automatic nature of Fast-Start Failover avoids delays that can result from human interaction. Automatic failover is also carefully controlled so that any risk of data loss or "split brain" processing of transactions is completely avoided.

Fast-Start Failover's automatic reinstatement of the original primary following failover will in most cases eliminate the time and effort required for a "manual rebuild" of the original primary database. This makes it easier (and much faster) to execute failovers rather than incur any downtime while administrators troubleshoot failures on the primary production system.

New Data Guard 10g Release 2 Role Transition events provide the added capability to integrate database failover with failover procedures at the middle tier to quickly detect Data Guard failovers and automatically redirect clients and applications to the new primary database at the standby location – providing an end-to-end solution for achieving business continuity.

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. Fast-Start Failover: Oracle Data Guard 10g Release 2
http://www.oracle.com/technology/deploy/availability/pdf/MAA_WP_10gR2_FastStartFailoverBestPractices.pdf
 4. Oracle Data Guard Broker (Part #B14230-01)
http://download-west.oracle.com/docs/cd/B19306_01/server.102/b14230/toc.htm
 5. Oracle Data Guard 10g Release 2 Switchover and Failover Best Practices
http://www.oracle.com/technology/deploy/availability/pdf/MAA_WP_10gR2_SwitchoverFailoverBestPractices.pdf
 6. Oracle Data Guard Concepts and Administration
http://download-west.oracle.com/docs/cd/B19306_01/server.102/b14239/toc.htm
 7. DB_ROLE_CHANGE – Oracle Database Application Developers Guide
http://download-west.oracle.com/docs/cd/B19306_01/appdev.102/b14251/toc.htm
 8. Oracle Data Guard 10g Release 2 Best Practices for Client Failover
<http://www.oracle.com/technology/deploy/availability/htdocs/maa.htm>
- A 10gR2 version of this paper will be available soon.



From "Automatic Failover, Data Guard 10g Release 2"
Open World San Francisco, September, 2005
Authors: Joseph Meeks & Mike Smith, Oracle Corporation

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

Copyright © 2005, 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.