

Fast-Start 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

Fast-Start Failover Best Practices Oracle Data Guard 10g Release 2

Overview	3
Elements of a Fast-Start Failover Configuration.....	3
Configuring Fast-Start Failover	5
Events that Trigger Fast-Start Failover.....	6
Reinstatement after A Fast-Start Failover	6
Oracle Test Results	7
Oracle Best Practices for Fast-Start Failover	7
Production Database Configuration	7
Network Transport.....	8
Standby Configuration	8
Observer.....	9
Fast-Start Failover Threshold	11
Monitoring	12
Flashback Database Configuration	12
Configurations with Multiple Standbys	13
Applications Well Suited to Fast-Start Failover	13
Automatic Client Failover.....	14
DB_ROLE_CHANGE Event.....	14
Conclusion.....	14
References	15

Fast-Start Failover Best Practices

Oracle Data Guard 10g Release 2

OVERVIEW

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 production database, without requiring manual intervention to execute the failover.

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

This paper describes Fast-Start Failover and describes [Maximum Availability Architecture \(MAA\)](#) [2] best practices for its use. For a detailed discussion of best practices for manual failover and planned switchover, please refer to [Oracle Data Guard 10g Release 2, Switchover and Failover Best Practices](#) [3].

ELEMENTS OF A FAST-START FAILOVER CONFIGURATION

Fast-Start Failover is used in a Data Guard configuration under the control of the Data Guard Broker. Data Guard Broker provides centralized management of a Data Guard configuration. Through its command line interface (DGMGRL), Data Guard Broker uses single commands to perform the equivalent work of multiple SQL*Plus statements, greatly simplifying the management of a Data Guard configuration. Data Guard Broker is included with Data Guard and does not require a separate installation.

Fast-Start Failover may be managed using either DGMGRL or Oracle Enterprise Manager 10g Grid Control. Enterprise Manager provides an easy to use graphical user interface for monitoring and control, and enables centralized management of one or more Data Guard configurations.

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

- The production database
- A target standby database
- The Fast-Start Failover Observer

The target standby database will become the new production database following a fast-start failover (note: There can be multiple standby databases in a Data Guard configuration, but only one can be designated as the fast-start failover target. A manual failover would be required in order to execute a failover to any of the additional standby databases within the configuration).




Figure 1 Fast-Start Failover Configuration

The Observer is a separate process incorporated into the DGMGRL client that continuously monitors the production database and the target standby database for possible failure conditions.

The rule is that out of these three participants, (production database, standby database & observer) whichever two can communicate with each other will determine the outcome of fast-start failover. For example, if the production database becomes unavailable, the Observer confirms with the target standby database that the production database is unavailable and that the target standby database is synchronized with the production database, and if so, initiates a fast-start failover to the target standby database. Lacking such agreement, a fast-start

failover cannot occur. This guarantees two important characteristics of Fast-Start Failover:

- Automatic failover will never result in more than one database in a Fast-Start Failover configuration assuming the production 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.
- A fast-start failover will only occur if there is a guarantee that no data will be lost.

CONFIGURING FAST-START FAILOVER

Details for configuring & enabling Fast-Start Failover are discussed in documentation for [Data Guard Broker](#) [4] and [Oracle Data Guard Concepts and Administration](#) [5]. The high level steps are as follows:

- Begin with a Data Guard configuration that includes a production database and at least one standby database configured using Maximum Availability protection mode with LGWR SYNC Redo Transport Services. Flashback Database and Flash Recovery Area must also be enabled on both production and standby databases.
- Designate a third system to be the Observer host. It must have the DGMGRL utility installed and have Oracle Net connectivity to both the production and standby databases. Ideally, the Observer should run on a host that is not located in the same datacenters as the production database, and standby databases. If this is not possible, then locate the Observer at the standby location on a separate system from the standby database, and in a fashion such that it is isolated as much as possible from events that could impact the standby database. Note that the Observer host does not require an Oracle instance be installed.
- Data Guard Real Time Apply is recommended to minimize failover time. This will make the standby database up-to-date with the latest redo received from the production database. This eliminates any delay at failover caused by waiting for the standby database to complete the application of received redo.
- Configure the Fast-Start Failover Target by providing the DB_UNIQUE_NAME of the database that is the failover target.
- Configure the failover threshold. The failover threshold is the number of seconds the Observer will attempt to reconnect to the production database before initiating a failover. Determining the most appropriate value for failover threshold is discussed later in this paper.
- Use Enterprise Manager or DGMGRL to enable Fast-Start Failover and start the Observer.

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 production database, and when the standby database confirms that it is in a “synchronized” state.

REINSTATEMENT AFTER A FAST-START FAILOVER

In many cases it will be possible for administrators to restart the original production database after a fast-start failover, and after the problem that had caused the failover has been resolved. For this reason, following a fast-start failover the Observer periodically attempts to reconnect to the original production database. When the Observer regains network access to the original production database, it initiates a request for the Data Guard Broker to automatically reinstate it as a standby database to the new production database. This quickly restores disaster protection and high availability for the new production database.

Note that there are safeguards in place to control how the old production database opens after a fast-start failover. If an administrator tries to open the old production database with the `STARTUP` command, Data Guard Broker will not allow it to proceed from the *mount* state to the *open* state until at least one other fast-start failover member agrees to that state transition. Because fast-start failover has already occurred and there is a new production database in the configuration, the old production database will not get confirmation from either the Observer or target standby database (now the new production database), thus - preventing a potential split-brain scenario.

If the administrator tries to start the old production database with the `STARTUP MOUNT` command, no error message will be generated, and the Observer will automatically reinstate the old production database as a new standby database in the Data Guard configuration.

If the administrator tries to start the old production database with the `STARTUP NOMOUNT` command, the old production database will not be mounted and the Broker will not pursue reinstatement until further administrative action is taken (e.g. mounting the old production database).

ORACLE TEST RESULTS

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




Figure 2: Fast-Start Failover Test Results

The test databases were each 100GB in size (note: failover timings are not impacted by the size of the database). Each host was connected to the next over a Gigabit Network. The workload on the production database generated 3 MB/second of redo data. 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 (the user configurable amount of time that the observer will wait for a response from the production database before declaring that it has failed and initiating failover) 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.

ORACLE BEST PRACTICES FOR FAST-START FAILOVER

Production Database Configuration

Fast-Start Failover requires that the Data Guard configuration be set at Maximum Availability protection mode using LGWR SYNC AFFIRM redo transport. Maximum Availability insulates the production database from the impact of network or standby server failures (such failures are automatically detected and the production database continues processing). However, due to the synchronous nature of redo shipping, there is potential for the performance of the production

database to be affected by network latency and disk I/O on the standby system. This is because the production database will not return a commit to the database client until it has received acknowledgement that the redo has been written to disk on both production and standby servers. In order to use Fast-Start Failover, network bandwidth and latency must be suitable for synchronous redo shipping.

Network Transport

Tuning operating system parameters that affect Data Guard network throughput can significantly enhance the ability of a network to support a Maximum Availability configuration.

These parameters include the TCP/Receive buffers and settings that regulate the size of the buffer between the kernel network subsystems and the driver for network interface cards. The importance of tuning for applications with significant workload is clear from testing done by Oracle that demonstrated an order of magnitude improvement in throughput by simply tuning TCP Send/Receive Buffers, Device Network Queue Size & SDU (Oracle Net Services session data unit). Figure 3 shows the result of this tuning. For a complete understanding please refer to [Data Guard Primary Site and Network Best Practices](#) [9].




Figure 3 – Impact of Network Tuning

Standby Configuration

To minimize failover time the standby database must be configured with both Standby Redo Logs (SRLs) and Real Time Apply. This enables the Managed Recovery Process on a physical standby, or the SQL Apply Process on a logical

standby, to apply redo to the standby database as it is received, without waiting for a log switch on the production database. This keeps the standby database up-to-date with the production database, minimizing the processing required at failover.

In Data Guard configurations where significant redo volume is generated during peak usage, it may be that default settings are insufficient for apply processes to keep pace with incoming redo. Oracle recommends administrators review Oracle best practices contained in [Oracle Database 10g Release 2 High Availability Best Practices](#) [6]. For further drill down into best practices for tuning the Redo Apply process for physical standby refer to: [Oracle Database 10g Best Practices, Data Guard Redo Apply and Media Recovery](#) [7]. For similar information for SQL Apply processing on a logical standby, refer to: [Data Guard SQL Apply Best Practices in Oracle Database 10g](#) [8].

Observer

Location: For Disaster Recovery requirements it is ideal to install the Observer in a location separate from the production and standby data centers. The observer should be independent from the data centers and when possible, connect to the production and standby database via the same network used by the client application. If this is not possible, then locate the Observer at the standby location on a separate system from the standby database, and in a fashion such that it is isolated as much as possible from events that could impact the standby database.

Installation: Install the Observer by installing 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. If Enterprise Manager will be used, also install the Enterprise Manager Agent on the observer system.

Making the Observer highly available: Enterprise Manager 10.2.0.1 supports automatic restart of the Observer on the same host if it detects that the observer process has failed. Automatic restart is activated when Fast-Start Failover is enabled. This automatically handles Observer outages due to unintended process death or Observer host reboot. It also relieves customers from having to configure a custom procedure to restart the Observer on host boot up, as the Enterprise Manager Agent will automatically start at boot up and will subsequently start the Observer if there is an outstanding unobserved condition.

Enterprise Manager 10.2.0.3 also enables administrators to designate an alternate Observer host to protect against events that could cause the first Observer host to fail. This is easily configured when creating a Fast-Start Failover configuration using the Enterprise Manager setup wizard. Upon detection of an unobserved condition, Enterprise Manager will start the Observer on the alternate host. In cases where the original Observer is still running but is network-isolated from the

primary and/or standby database, a new Observer will be started on the alternate host if it has the required network connectivity.

Enterprise Manager 10.2.0.3 also includes enhancements for self-monitoring an Observer configuration. Enterprise Manager will monitor and alert on conditions that affect its ability to perform an Observer restart, such as alternate Observer host(s) being down or Management Agent unavailability.

Support for multiple fast-start failover Observers on a single host: Customers may wish to use a single host for monitoring multiple production databases, each having its own standby database as a fast-start failover target. Since each observer process monitors a single production/standby pair, multiple Observers must be configured on the single host. Implementation details depend on the management interface chosen. A representative use-case is described below:

Use case:

- Three different (totally unrelated) production databases.
- Three physical standby databases, one for each production database.
- A single server on which will run 3 independent Observers in a single Oracle home, one Observer for each production/standby pair.

Using the Data Guard Broker to configure multiple Observers: The Broker's command line interface, DGMGRL, is used to create fast-start failover Observers for multiple production databases. A dgmgrl log file used to track observer events is specified when you start DGMGRL (using the `-logfile` command line parameter). Simply designate a unique name for each Observer's log file in order to have multiple Observers running in the same Oracle home on a single host,

The Observer also maintains a data file that contains the configuration information for a Fast-Start Failover production/standby pair. By default this file is called FSFO.DAT. To have multiple Observers running in the same home you must specify the location of the Observer data file when you start the observer, using the [FILE= option](#) to the START OBSERVER command [16]. You must ensure that each data file is uniquely named.

For example, following the use-case above, we can use each production database's DB_UNIQUE_NAME in order to uniquely identify the files for each Observer. In the case where the DB_UNIQUE_NAME for each of the production databases are Boston, Chicago, and Washington, you would do the following, (assuming any needed directories have already been created):

```
dgmgrl -logfile $ORACLE_HOME/rdbms/log/Boston.log
DGMGRL> START OBSERVER FILE=$ORACLE_HOME/dbs/Boston.dat';
```

```

dgmgrl -logfile $ORACLE_HOME/rdbms/log/Chicago.log
DGMGRL> START OBSERVER FILE=$ORACLE_HOME/dbs/Chicago.dat';

dgmgrl -logfile $ORACLE_HOME/rdbms/log/Washington.log
DGMGRL> START OBSERVER FILE=$ORACLE_HOME/dbs/Washington.dat';

```

Note that DGMGRL by itself has no context until it connects to a Data Guard configuration. Once DGMGRL is connected to a configuration and a START OBSERVER command is issued, only that configuration will be observed. The broker will disallow multiple observers for the same configuration, hence the requirement for unique file names.

For more information on the DGMGRL commands and options please refer to Chapter 8, Data Guard Command-Line Interface Reference, of the [Data Guard Broker Manual](#) [4].

Using Enterprise Manager to configure multiple Observers: When using Enterprise Manager's Enable Fast-Start Failover wizard, simply specify the host and Oracle Home for the Observer when you configure fast-start failover. Enterprise Manager will automatically ensure that each Observer you configure will have uniquely named Observer data and dgmgrl log files. The Observer data files are located in the \$ORACLE_HOME/dbs directory (and named afoXXXXX.dat) and the dgmgrl log file is located in the \$ORACLE_HOME/rdbms/log directory (and named dgmgrlXXXXX.log). The XXXXX is a unique number generated by Enterprise Manager and is the same for the Observer data file and dgmgrl log file for a given production-standby database configuration. The value for XXXXX will change each time an Observer is restarted.

Fast-Start Failover Threshold

A fast-start failover occurs when the Observer and the Standby Database both lose contact with the production database for a period of time that exceeds the value set for `FastStartFailoverThreshold`, and when both parties agree that the state of the configuration is synchronized. An optimum value for `FastStartFailoverThreshold` weighs the trade-off between the fastest possible failover (thus minimizing downtime), and unnecessarily triggering failover due to fleeting network irregularities or other short-lived events that do not have material impact on availability. The default value set when Fast-Start Failover is enabled is 30 seconds. Recommended settings for `FastStartFailoverThreshold` are:

- Single instance production database, low latency, reliable network = 10-15 seconds
- Single instance production database, high latency network over WAN = 30-45 seconds

- RAC production database = time needed to evict a failed node + 20 seconds. Using this approach to setting a threshold value in a RAC configuration avoids a fast-start failover from occurring when there are surviving nodes that will resume processing once node eviction is complete.

Monitoring

Because a fast-start failover will not occur unless the production and standby databases are in a synchronized status, it is important to respond quickly to any event such as a network outage or standby server crash, so that Data Guard can quickly resolve any resulting redo gap and return the standby configuration to a synchronized status.

Enterprise Manager can be used to continually monitor configuration status and automatically notify administrators of events that require attention. Alternatively the state of the configuration can be monitored via the `FS_FAILOVER_STATUS` column of the `V$DATABASE` view. A `SYNCHRONIZED` status means that the production and standby databases are in sync and a fast-start failover is possible. `UNSYNCHRONIZED` status means the standby has not received all of the redo generated by the production database and a fast-start failover cannot occur.

Because a fast-start failover requires two of the three members in the configuration to agree, Observer status should also be monitored, either through the Enterprise Manager GUI, or by monitoring the Observer via the `FS_FAILOVER_OBSERVER_PRESENT` column of the `V$DATABASE` view.

Flashback Database Configuration

Configure Flashback Database on both the production and the standby databases. Flashback Database is used by the Data Guard Broker to automatically reinstate a failed production database as a standby database to the new production database in the event that a failover occurs. Automatic reinstatement will occur as long as the old production database can be restarted and the failover operation was executed by Fast-Start Failover.

Oracle recommends that when used only for fast instantiation following a Data Guard failover, set `DB_FLASHBACK_RETENTION_TARGET` to a minimum of 60 minutes (default value is 1440). Note that if Flashback Database serves the additional function of providing fast point in time recovery for protection against user error and corruption, then an extended flashback retention period should be set for the amount of time deemed necessary to achieve these goals. For example, Oracle MAA best practices recommend a minimum of a 6 hour retention period for this purpose (such goals are customer specific, trading off increased storage requirements for increased protection).

Also, refer to "[Setup and Maintenance for Oracle Flashback Database](#)" in [Oracle Database Backup and Recovery Basics](#) [10] and "[Setting Up Flash Recovery Areas](#)"

in [Oracle Data Guard Concepts and Administration](#) [4] for information about setting up Flashback Database and flash-recovery areas.

Configurations with Multiple Standbys

It is not unusual for a Data Guard configuration to include more than one standby database for a single production database. Often a “local” standby is utilized to achieve zero data loss protection by utilizing a low latency LAN capable of supporting synchronous data protection without incurring the performance overhead of a high latency WAN. The second standby in such a configuration may be geographically remote; 100’s to 1,000’s of miles away from the production site, utilizing asynchronous data protection to avoid the performance overhead of a high latency WAN. Using Fast-Start Failover in such a configuration, the local standby would be designated as the failover target and would be the system, along with the production database, that the Observer would monitor. At failover, Fast Start Failover would transition the local standby into the production role, and the remote standby would transition seamlessly to the role of standby database to the new production database – affording continuous data and disaster recovery protection while the original production database is repaired.

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 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.

However, there is a class of applications that are extremely sensitive to downtime e.g. 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 system event that when used in conjunction with FAN OCI events, (described below), makes it possible to quickly notify clients that a failover has occurred so that they are automatically redirected to the new production database. This event is described below. For a more detailed discussion of automating client failover, please refer to [Oracle Data Guard 10g Release 2 Best Practices for Client Failover](#) [11].

DB_ROLE_CHANGE Event

Whenever a database transitions from one role to another, a DB_ROLE_CHANGE system event is fired. This is much like the STARTUP system event, except it fires only after a role change. Administrators can develop triggers that execute when this event occurs as a way to manage post role-change tasks. The event fires when the database opens for the first time after the role transition regardless of its new role, (i.e., regardless of whether the role change caused it to open for the first time as a production database or as a logical standby or as a physical standby, in read-only mode).

The DB_ROLE_CHANGE system event may be used to manage/automate post role change tasks. Typical tasks include starting a service / services on the new production database, changing the naming services or connection descriptors so clients will reconnect to the new production database, starting third party applications, adding temporary tablespaces, and so on. DB_ROLE_CHANGE is a flexible mechanism to allow the administrator to automate any actions that can be accomplished via database triggers. For more information on DB_ROLE_CHANGE event refer to [Oracle Database Application Developer's Guide – Fundamentals 10g Release 2](#) [12].

In addition, when a failover operation is coordinated through the Data Guard Broker, a Fast Application Notification (FAN) [13] event is posted on behalf of the failed production database to notify OCI clients of the failure. Further, if the client connection was TAF-enabled ([Transparent Application Failover](#) [14]), the application could automatically fail over to the new production database.

CONCLUSION

Oracle Data Guard has evolved over a number of major Oracle releases. It is the most functional disaster recovery solution 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 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 production database following failover will in most cases eliminate the time and effort required for a "manual rebuild" of the original production database. This makes it easier (and much faster) to execute failovers rather than incur any downtime while administrators troubleshoot failures on the production system.

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 production 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. Oracle Data Guard 10g Release 2, Switchover and Failover Best Practices
http://www.oracle.com/technology/deploy/availability/pdf/MAA_WP_10gR2_SwitchoverFailoverBestPractices.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 Concepts and Administration
http://download-west.oracle.com/docs/cd/B19306_01/server.102/b14239/toc.htm
6. Oracle Database 10g Release 2 High Availability Best Practices
http://download.oracle.com/docs/cd/B19306_01/server.102/b25159/toc.htm
7. Oracle Database 10g Best Practices: Data Guard Redo Apply and Media Recovery
http://www.oracle.com/technology/deploy/availability/pdf/MAA_WP_10gRecoveryBestPractices.pdf
8. Data Guard SQL Apply Best Practices in Oracle Database 10g
http://www.oracle.com/technology/deploy/availability/pdf/MAA_WP_10gR2_SQLApplyBestPractices.pdf

9. Data Guard Primary Site and Network Best Practices

http://www.oracle.com/technology/dep/availability/pdf/MAA_WP_10gR2_DataGuardNetworkBestPractices.pdf

10. Oracle Database Backup and Recovery Basics (Part # B14192-02)

http://download-west.oracle.com/docs/cd/B19306_01/backup.102/b14192/toc.htm

11. Oracle Data Guard 10g Release 2 Best Practices for Client Failover

http://www.oracle.com/technology/dep/availability/pdf/MAA_WP_10gR2_ClientFailoverBestPractices.pdf.

12. DB_ROLE_CHANGE – Oracle Database Application Developers Guide

http://download-west.oracle.com/docs/cd/B19306_01/appdev.102/b14251/toc.htm

13. 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

14. Transparent Application Failover (TAF)

http://download-west.oracle.com/docs/cd/B19306_01/server.102/b14220/high_av.htm - i36759

15. Oracle Database Clusterware and Oracle Real Application Clusters Administration and Deployment Guide 10g Release 2 (10.2).

http://download-west.oracle.com/docs/cd/B19306_01/rac.102/b14197/admcon.htm - sthref29

16. FILE= option

http://download-west.oracle.com/docs/cd/B19306_01/server.102/b14230/dgmgrl.htm#BABJECGA



Oracle Data Guard 10g Release 2 Fast-Start Failover Best Practices

January, 2010

Authors: Joseph Meeks, Michael T. Smith, Ashish Ray, Sadhana Kyathappala, Sean Connolly

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, PeopleSoft, and Retek are registered trademarks of Oracle Corporation and/or its affiliates. Other names may be trademarks of their respective owners.