

PPL

Automatic Failover – Data Guard 10g Release 2 High Availability & Disaster Recovery

“Data Guard Fast-Start Failover provides simple, fast, unattended failover for an outage management system that PPL depends upon to provide critical customer services 24 hours a day and especially during emergencies. While we have used Data Guard for disaster recovery (DR) since Oracle9i, Fast-Start Failover blurs the line between High Availability and DR – enabling us to address both requirements with a single solution. ”

*Chris Carter
Director, Enterprise Technology Services
PPL Services Corp.*



Corporate Profile

- FORTUNE 500® global energy holding company.
- PPL Generation – Owns and controls more than 11,000 megawatts in the United States with a portfolio that includes power plants in Connecticut, Illinois, Maine, Montana, New York and Pennsylvania.
- PPL EnergyPlus - Buys and sells energy in key U.S. competitive wholesale and deregulated retail markets. Firm sales are backed by PPL's generating plants located throughout the United States.
- PPL Electric Utilities - The electric delivery subsidiary serves 1.4 million customers in Pennsylvania
- PPL Global owns and operates distribution businesses in the United Kingdom and Latin America that deliver electricity to 3.7 million customers.
- <http://www.pplweb.com/>

OVERVIEW

Headquartered in Allentown, Pa., PPL controls more than 11,000 megawatts of generating capacity in the United States, sells energy in key U.S. markets, and delivers electricity to more than 5 million customers in Pennsylvania, the United Kingdom and Latin America.



PPL has used [Oracle Data Guard](#) [1] for disaster recovery protection since Oracle9i. As PPL has moved up to Oracle Database 10g Release 2, it has expanded the role it plays in their High Availability & Disaster Recovery Architecture by deploying Data Guard Fast-Start Failover for a mission critical outage management system. One of PPL's most important customer support applications, the outage management system is the front line of delivering customer service during storms and other events that impact electric delivery to over 1.4 million customers in Pennsylvania.

The introduction of [Data Guard Fast-Start Failover](#) [2] with Oracle Database 10g Release 2 fundamentally changed the way PPL utilizes its Data Guard standby

Oracle Technologies used by PPL

- Oracle Database 10g Release 2
- Data Guard Fast-Start Failover
- RMAN
- Flashback Database
- Flash Recovery Area
- Enterprise Manager

databases. Fast-Start Failover enables fast, automatic, zero data loss failover to the standby database should the primary production database fail, delivering both High Availability and Disaster Recovery (HA/DR). Data Guard enables faster failover than a cold cluster failover. It provides a much superior level of high availability and data protection beyond that possible in a cold failover cluster solution. Unlike cold failover clusters, Data Guard maintains a second, completely independent and continually validated synchronized copy of the primary database. This can be accomplished using zero-data loss protection mode even in cases where primary and standby sites are separated by 100's of miles. Thus Data Guard can provide high availability and data protection against a wide range of events - from human error, and hardware-induced data corruption to component failure, complete site failure and widespread natural disasters.

The evolution of Oracle Data Guard 10g Release 2 into an integrated HA/DR solution for Oracle data has occurred without compromising the basic tenants that have made Data Guard so attractive as a disaster recovery solution that include:

- **Low cost:** included as a no-charge feature of Oracle Database Enterprise Edition.
- **Integrated:** Data Guard operates seamlessly with Oracles comprehensive HA feature set - Real Application Clusters, Automatic Storage Management, RMAN, Flashback Database and Enterprise Manager and the Oracle documented [Maximum Availability Architecture](#) [3] best practices blueprint for achieving high availability.
- **Simplicity:** GUI management via Enterprise Manager. Data Guard automates monitoring of a Data Guard configuration and resynchronization following network or standby database outages.

PPL CONFIGURATION DETAILS

Applications

The focus of PPL's initial implementation of Fast-Start Failover is the mission critical database supporting GE Energy's PowerOn® outage management system. The PowerOn product helps PPL reduce service interruptions, improve network reliability and increase customer satisfaction. The PowerOn system provides many features to reduce operational costs and improve network reliability, including:

- Outage prediction
- Dispatching
- Crew management
- Switching
- Outage restoration management

The very nature of the PowerOn application speaks volumes to the requirement for data protection and high availability. A true 24x7x365 mission critical application, it enables PPL to take proactive measures to improve the quality of service, as well

Applications

- [GE PowerOn](#) - Outage Management System from General Electric
- Enterprise Application Integration from WebMethods

Network & Systems Configuration

- Primary and standby datacenters 4 miles apart connected by 2Gbit network with less than 2ms RTT latency
- PowerOn primary & standby servers are HP rp4440s, each with 4 cpus and 12GB memory running HP-UX PA-RISC
- WebMethods primary & standby servers are HP rp3440s, each with 2 cpus and 12GB memory running HP-UX PA-RISC

Data Guard Configuration

- Two primary/standby pairs using Data Guard physical standby and Fast-Start Failover
- Maximum Availability, LGWR SYNC, Real-time Apply – 1.5MB/sec of Redo data at peak call volume of 500 calls/minute
- Fast-Start Failover with user configured Fast-Start Failover Threshold of 4 minutes
- Flashback Database with 1-week retention period for flashback logs
- Flash recovery area used to automate archive log management and to hold RMAN backups on-disk
- Configuration created and managed using Oracle Enterprise Manager
- Both Enterprise Manager repository and RMAN catalog are protected by their own Data Guard standby databases

as to respond to critical customer service requests when events beyond PPL's control impact service delivery.

The PowerOn system also depends upon internal interfaces to WebMethods for coordination with other business processes. Because of this close relationship, PPL has also implemented Data Guard Fast-Start Failover for the Oracle database that supports WebMethods.

Completing the applications related to the outage management system is Oracle Enterprise Manager, used to manage their Data Guard configurations. PPL has protected the Oracle repository used by Enterprise Manager with a Data Guard standby database, as prescribed by Oracle [MAA best practices for high availability for Enterprise Manager](#) [5]. Similarly, PPL uses a Data Guard standby to protect the repository for its RMAN catalog. Refer to the MAA best practice paper [Using Recovery Manager with Oracle Data Guard in Oracle Database 10g](#) for more details [6].

Network & Systems Configuration

Primary and standby databases for each application are located in two different datacenters connected by a 2Gbit network link having RTT network latency less than 2ms. Primary and standby servers for the PowerOn application are HP rp4440s, each with 4 cpus and 12GB memory running HP-UX PA-RISC. Primary and standby servers for WebMethods are HP rp3440s, each with 2 cpus and 12GB memory running HP-UX PA-RISC. Clients connect to the database tier via Citrix servers.

Data Guard Configuration

All Data Guard standby databases are created and managed using Enterprise Manager. Primary and physical standby databases for PowerOn and WebMethods are configured in Data Guard Maximum Availability mode using LGWR SYNC redo transport services using Real-Time Apply.

Flashback Database is configured with a one-week flashback log retention period (note that only a 1 hour retention period is required for Data Guard Fast-Start Failover – the additional retention period used by PPL provides fast point in-time recovery using Flashback database for up to 1 full week.). PPL utilizes the RMAN Flash Recovery Area to hold an on-disk backup and to automate the management of archive log and flashback log files.

Peak call volume of 500 calls/minute generates 1.5MB/second of redo data on the primary database.

Fast-Start Failover Configuration

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

- The primary database
- The target standby database that will become the new production database following a fast-start failover.
- The Fast-Start Failover Observer that continually monitors the state of the configuration and will trigger a failover once required conditions are satisfied.




Figure 1 Fast-Start Failover Configuration

The rule is that of these three participants, (primary database, standby database & 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 is unavailable and that the target standby database was synchronized with the primary when it last made contact. If these conditions are met, the Observer is able to initiate 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 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.
- A fast-start failover will only occur if there is a guarantee that no data will be lost.

Refer to [“Fast-Start Failover Best Practices: Oracle Data Guard 10g Release 2”](#) [2] for detailed discussion of configuration best practices.

Highlights of PPL’s Fast-Start Failover configuration include:

- Flashback Database is used by the Observer to automatically reinstate the original primary database as a new standby database after a failover has occurred. PPL testing demonstrated that automatic reinstatement of the original primary as a standby database for the new primary completed within two minutes of the original primary reaching the mount state. A minimum of one hour of flashback logs must be retained in order for this to occur. A longer flashback retention period can be used if there is a desire to use Flashback for very fast point in time recovery further in the past of the current time. This is the reason why PPL has configured a one-week retention period for Flashback Database.
- Flash Recovery Area Flash Recovery Area has been configured to hold the flashback logs used by Flashback Database, all archive logs, and on-disk backups taken by RMAN. The Flash Recovery Area automates the management of log file space. For example, archive logs are automatically deleted after they are applied to the standby database and have been backed up by RMAN.
- Fast-Start Failover Observer: It is ideal for the Observer to be installed in a location separate from the production and standby data centers. This eliminates the possibility of any single event disabling more than one of the three parties to a Fast-Start Failover configuration. However, as it the case at PPL, it is not always practical to accommodate a third site. PPL installed the Observer at the standby location on a different system from the standby database, isolated as much as possible from failures that can impact the standby systems. The Observer was installed using the Oracle Client Administrator (this results in a small footprint because an Oracle instance is not required). PPL also installed an Enterprise Manager Agent on the Observer system.

PPL went to additional lengths to make 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, automatically handling Observer outages due to unintended process death or Observer host reboot. Enterprise Manager 10.2.0.3 provides the additional ability to identify an alternate observer home on a different host and enable automatic Observer restart on the second host should the first host fail.

Finally, since PPL has two primary/standby pairs (one for PowerON and a second for WebMethods), it requires two separate observer processes. To simplify management, they use a single observer host to install and manage

both observer processes. Details of how to accomplish this are included in the Fast-Start Failover best practice paper referenced above [2].

- 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. By utilizing a user-configurable threshold, administrators can control how quickly an automatic failover will occur. Once this threshold is reached and the failover is triggered, the new database assumes the primary role in a matter of seconds.

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 first enabled is 30 seconds. The optimum value will depend upon your business and technical requirements.

PPL has determined a threshold of four minutes is appropriate for its requirements. They have purposefully taken a conservative approach and allowed for a longer threshold period to provide a greater opportunity for events that impacts the primary system or the ability of the observer and the standby database to communicate with the primary to be resolved without triggering an automatic failover. Once the four-minute threshold has been reached, PPL recovery time objectives require failover to be executed immediately without further delay. Extensive failure testing has demonstrated that once the threshold has been exceeded and a Fast-Start Failover triggered, the target standby database rapidly transitions to the primary role and begins accepting new client connections as quickly as requests are made to do so.

Automating Client Failover

There are many different approaches used to configure client failover in a Data Guard environment. Prior to Fast-Start Failover, however, accommodations had to be made for the time needed for detection and manual intervention required to execute a database failover. This made it difficult to guarantee fast, predictable failover times. Fast-Start Failover addresses this problem by automating fault detection and database failover.

Automating Client Failover

- **Fast-Start Failover** executes an automatic failover when the primary can not be reached for a period of time that exceeds Fast-Start Failover Threshold, and when the standby database confirms the configuration is synchronized and there will be zero data loss
- The Data Guard **DB_ROLE_CHANGE** system event fires a trigger upon database role change to start the production instance of the application service on the new primary
- OracleNet Connect Time Failover redirects clients to the new primary database
- After failover is complete, Fast-Start Failover automatically reinstates the old primary as a new standby database
- The new standby (old primary) is automatically resynchronized with the new primary

Data Guard 10g Release 2 also includes a new **DB_ROLE_CHANGE** system event that fires whenever a database transitions from one role to another to extend this same level of automation to client failover. The new system event enables administrators to develop triggers that can execute post role-change tasks to automatically redirect client connections to the new primary database. For a more detailed discussion of this event and Oracle Database 10g best practices for automating client failover, please refer to [Oracle Data Guard 10g Release 2 Best Practices for Client Failover](#) [7].

PPL has chosen to utilize this event to automate client failover similar to the generic example provided below:

The first step in implementing automatic client failover is to insure that the application service is only available on instances associated with the primary database. This is done by first creating an application service on the Primary database as a unique identifier for the database functioning in the primary role. In this example, the service name is 'sales' and is created as follows:

```
exec DBMS_SERVICE.CREATE_SERVICE('sales','sales');
```

Note that the redo generated by issuing this statement on the primary will be sent to the standby and applied to the standby database, automatically propagating this new application service to the standby (the service will only start on the standby database when a role change occurs).

At switchover or failover, a trigger (written by the user) is fired by the Data Guard **DB_ROLE_CHANGE** event and automatically changes the location of the application service 'sales'. The trigger (see example below) does this by starting the service 'sales' when a database opens as the primary database, and stopping the service on any database that is no longer operating in the primary role.

```
create or replace trigger manage_service
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');
    else
        DBMS_SERVICE.STOP_SERVICE('sales');
    end if;
end;
```

The second part of implementing automatic client failover is redirecting client connections to the new primary system that is offering the application service 'sales'.

Continuing with the example above, a TNS alias 'SALES' is configured using [OracleNet Connect Time Failover](#) [8] to automatically redirect client connections to the new primary database running the production instance. This handles both new connections and the failover of existing connections to the original primary once they have been disconnected.

Note: The client failover best practices paper referenced includes additional provisions for breaking existing connections out of TCP time-outs in order to expedite failover in situations where clients attached to the failed primary do not get an immediate disconnect.

OracleNet Connect Time Failover using a TNS alias 'SALES' is implemented as follows:

```
SALES=
  (DESCRIPTION=
    (address_list=(load_balance=off)(failover=on)
      (ADDRESS=(PROTOCOL=TCP)(Host=BOSTON)(Port=1521))
      (ADDRESS=(PROTOCOL=TCP)(Host=CHICAGO)(Port=1521)))
    (CONNECT_DATA=(SERVICE_NAME=sales)
      (FAILOVER_MODE=
        (TYPE=session)
        (METHOD=BASIC)
        (RETRIES=180)
        (DELAY =5))))
```

Using the alias, the client will failover to the second node in the address_list if the first node in the list does not contain the designated application service 'sales'. If it cannot make immediate connection, it will wait five seconds and retry again up to a total of 180 times (both parameters are user configurable).

Note: The FAILOVER MODE parameters in the above example are part of the 'Transparent Application Failover (TAF)' mechanism and can also be specified for all clients by adding them to the application service 'sales' when it is created. See [Section 13.3 of the Oracle Database Net Services Administrator's Guide](#) for more information on TAF [9].

CONCLUSION

PPL uses Oracle Data Guard 10g Release 2 for both High Availability and Disaster Recovery. PPL simply extended its Oracle9i Data Guard DR architecture to incorporate Data Guard 10g Release 2 Fast-Start Failover. This enables fast, consistent, automatic failover in the event of component or site failure. PPL found Fast-Start Failover to have the combination of features for their HA requirements. It was simple for them to move to Data Guard 10g Release 2 and leverage their existing Data Guard expertise. PPL's extensive fault testing has consistently demonstrated that Data Guard Fast-Start Failover delivers the required level of High Availability and data protection, in the simplest fashion possible.

REFERENCES

1. Oracle Data Guard Overview
<http://www.oracle.com/technology/deploy/availability/htdocs/DataGuardOverview.html>
2. Fast-Start Failover Best Practices
http://www.oracle.com/technology/deploy/availability/pdf/MAA_WP_10gR2_FastStartFailoverBestPractices.pdf
3. Oracle Maximum Availability Architecture -
<http://www.oracle.com/technology/deploy/availability/htdocs/maa.htm>
4. GE PowerON
http://www.gepower.com/prod_serv/products/scada_software/en/poweron.htm
5. MAA best practices for high availability for Enterprise Manager
<http://www.oracle.com/technology/deploy/availability/htdocs/maa.htm - GridControl>
6. Using Recovery Manager (RMAN) with Data Guard 10g
http://www.oracle.com/technology/deploy/availability/pdf/RMAN_DataGuard_10g_wp.pdf
7. Oracle Data Guard 10g Release 2 Best Practices for Client Failover
http://www.oracle.com/technology/deploy/availability/pdf/MAA_WP_10gR2_ClientFailoverBestPractices.pdf
8. OracleNet Connect Time Failover
http://download-west.oracle.com/docs/cd/B19306_01/network.102/b14212/concepts.htm - sthref170
9. Oracle Database Net Services Administrator's Guide and TAF
http://download-west.oracle.com/docs/cd/B19306_01/network.102/b14212/advcfg.htm - sthref1277



PPL Corporation OTN Case Study – Data Guard 10g Release 2

February 2007

Authors: Joseph Meeks, Larry Carpenter

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.