

# Synchronous Data Protection Across 300 Miles

## NeuStar & Oracle Data Guard

### OVERVIEW

#### NeuStar Corporate Profile

- 2007 Revenue: \$420 million
- Leading provider of Operations Support Systems (OSS) clearing house and Ultra NGM services
- Manages the North American registry of telephone numbers
- Operates the Internet Registry supporting .biz, .cn, .us, & other domain addresses
- <http://www.neustar.com/>

#### High Availability and Disaster Recovery

- Oracle Database 10g
- Oracle Data Guard
- Synchronous data protection: Zero data loss
- Production data center & remote disaster recovery site on Wide Area Network separated by 300 miles

Founded to meet new technical and operational challenges of the communications industry when the U.S. government mandated local number portability in 1996, NeuStar has evolved to create customer-responsive technologies that deliver a broad range of essential clearinghouse capabilities for service providers. NeuStar's clearinghouse services enable customers to manage critical activities such as record exchanges, subscriber growth, network optimization, content management, and inter-network call origination and termination. NeuStar has also designed its automated solutions to evolve in response to the rigors of tomorrow's communications services - including VoIP, wireless data, and law enforcement compliance.

NeuStar has deployed [Oracle Data Guard](#) [1] to protect more than 60 mission critical databases essential to its operations. Most Data Guard configurations use Maximum Availability protection mode in which data is transmitted synchronously to remote disaster recovery sites, assuring no loss of data should the production database fail. NeuStar has successfully implemented such Data Guard configurations across sites separated by up to 300 miles, providing optimal protection from widespread disasters.

This case study describes how NeuStar has deployed Oracle Data Guard as a major component of its Business Continuity strategy.

### NEUSTAR REQUIREMENTS

NeuStar's Business Continuity requirements for mission critical databases are:

- Recovery Point Objective (RPO) is zero. There must be no data loss should the production database fail due to corruptions, system failures, storage failures, site failures, or any other unforeseen event.
- Recovery Time Objective (RTO) is the total time required for database recovery after a failure. Service Level Agreements require RTO's no greater than 5-10 minutes.
- Zero data loss RPO must be achievable for WAN deployments where the remote standby location is up to 300 miles away from the production location.
- Minimize costs by utilizing industry standard enterprise server, SAN, and networking components.
- Protect against corruptions by validating data before updates are applied to the standby database.
- If the production database connection to the standby is lost, the production database must be able to continue processing without interruption. When the production database regains connection to the standby, it must automatically

resynchronize the standby database. Likewise, if the standby database detects a corruption in data transmitted by the production database, it must guarantee that the corruption is not applied to the standby database, and it must automatically request a new copy of the data in question, and automatically resynchronize with the production database should the new copy prove valid.

- Role management services must allow for the reliable transition of roles from production to standby and back to production, with the guarantee of zero data loss.
- The ability to automate switching of both the application and database to the standby site using a single procedure.

## NEUSTAR'S CHOSEN ARCHITECTURE

### Oracle Data Guard

NeuStar determined that Data Guard is the optimum solution to address the full range of their requirements. An overview of NeuStar's Business Continuity architecture is described in Figure 1.

#### System & Network Configuration

- OS: AIX5L (64-bit)
- Network: T3 network
- RTT latency: 20ms
- WAN – 300 miles

#### Data Guard Configuration

- Oracle Database 10g
- Redo Apply (physical standby)
- Zero Data Loss: Maximum Availability - LGWR SYNC AFFIRM
- 10 MB/sec peak redo generation
- 60 different Data Guard configurations company-wide




Figure 1: NeuStar Business Continuity Architecture

NeuStar utilizes Data Guard Redo Apply (physical standby) to maintain a synchronized copy of the production database at a remote data center located 300 miles from the primary site. Data Guard transmits Oracle recovery data (redo data) from the primary site to the standby location to maintain complete synchronization between primary and standby databases. Peak redo generation at primary databases reach 10 MB/sec. Data Guard is configured in Maximum Availability Mode using LGWR SYNC Redo Transport Services. NeuStar utilized [Oracle's Maximum](#)

[Availability Architecture](#) [2] best practices blueprint, as well as the technical white paper “[Oracle Data Guard: Primary Site and Network Configuration Best Practices](#)” [3] for technical guidance during deployment.

## Maximum Availability Protection Mode

Maximum Availability provides the highest level of data protection possible without compromising the availability of the primary database. Figure 2 provides an overview of the Maximum Availability process architecture. It begins with a transaction on the primary database. LGWR reads the resulting redo records from the redo buffer in SGA and writes to an online redo log file. A Data Guard process, the Log Network Server (LNS), also reads the redo records from the redo buffer and synchronously transmits it to the standby server. A second Data Guard process running on the standby, the Remote File Server (RFS), receives the redo data and writes it to a standby redo log file (SRL). The RFS sends an acknowledgement back to the primary database that the redo has been received and written to disk. When the primary database receives this acknowledgement, the LGWR is




Figure 2 – Maximum Availability Process Architecture

acknowledges the commit to the client application and is then able to process the next transaction. Note that an SRL is similar in all ways to an online redo log except that SRLs are only used by standby databases, and only when databases are in a standby role.


Beginning with Oracle Database 10g, Data Guard Real-Time Apply enables the [Managed Recovery Process \(MRP\)](#) [4] to apply redo to the standby database as it is written to the SRL on the standby, without waiting for a log switch on the primary database.

In Maximum Availability protection mode the primary database does not shut down if a fault (e.g. a network disconnect) prevents it from receiving an acknowledgement from the standby database. If such a fault occurs, the primary database momentarily pauses until the Data Guard NET\_TIMEOUT threshold

(configurable and specified in seconds) has been reached, after which the primary database will continue to process new transactions. When the fault is corrected, Data Guard automatically resolves the resulting log file gap on the standby server and resumes operating in Maximum Availability mode (details below).

### Automatic Resynchronization

Data Guard automatically detects and resolves occurrences where the standby has not received all of the redo generated by the primary database. The resynchronization process is described in Figure 3.



**Figure 3 – Automatic Resynchronization**

All of the solid lines in Figure 3 represent the primary path used to transmit the current redo stream in Maximum Availability protection mode. The dotted lines in Figure 3 reflect the path taken by redo needed to resolve a gap between the primary and standby databases (gaps occur due to network outage, standby failure, or physical corruptions detected by the Data Guard redo apply process). For example, assume a failure of the network link between primary and standby locations. Data Guard detects the fault, but processing continues on the primary database generating and archiving the redo locally to archived redo logs. At the same time, an ARCH process on the primary database continually pings the standby server. When it can reconnect to the standby server, the resynchronization process begins. Data Guard determines which archive logs are missing or incomplete at the standby server. One or more ARCH process automatically resynchronize the standby database by sending the required archive logs to the standby server (note that administrators may configure a maximum of 10 ARCH processes in Oracle Database 10g Release 1, and a maximum of 30 ARCH processes using Oracle Database 10g Release 2.). An RFS process receives the redo, writes it to an archive log on the standby server, and registers the completed log file in the standby control file. The Managed Recovery Process (MRP) applies

the redo to the standby database. This is the process that follows the dotted lines in Figure 3.

Meanwhile, at every log switch, Data Guard also attempts to reestablish a SYNC LNS connection between the primary and standby databases. When it succeeds, the LNS resumes transmitting the current redo stream. This prevents the standby database from falling any further behind. Once the redo log file gap has been resolved, the MRP process automatically transitions to applying current redo records directly from the SRL.

## **NeuStar Role Transition Procedures**

NeuStar has configured both database and application servers at the remote DR site.

In addition to following [Oracle Best Practices for Switchover and Failover](#) [5], NeuStar has developed a custom GUI used by their administrators to facilitate failover for several of their Data Guard systems. The GUI also provides administrators the ability to do planned switchovers of the database and start the application at the remote site. There is an SLA that allows a maximum of 10 minutes to complete this process. Role transitions usually complete in 5 minutes.

For other Data Guard configurations, system administration staff and DBAs coordinate database and application switchover/failover. Application administrators shutdown the applications at the primary site and a DBA initiates a database switchover script. When switchover is complete, the Application team starts the application on servers at the remote site, which is pre-configured to connect to the database on that site. This manual process completes within 10 minutes.

NeuStar has also implemented Oracle Enterprise Manager Grid Control and is evaluating Data Guard Fast-start Failover (described below) to completely automate database failover and reinstatement of the original primary database.

## **DATA GUARD 10g ENHANCEMENTS**

There are many new enhancements in Data Guard 10g. Several significant enhancements directly relevant to NeuStar's configuration are highlighted below.

### **Real Time Apply**

Real-time apply enables the standby database to apply redo data as it is received, without waiting for the current standby redo log file to be archived. The Managed Recovery Process used by Redo Apply on the standby database reads directly from standby redo logs. This results in faster switchover and failover times because all of the redo received by the standby database has already been applied when failover or switchover begins.

## Flashback Database Support

Data Guard 10g supports the Flashback Database feature that allows a standby database to be quickly and easily flashed back to an arbitrary point in time. This feature provides the following benefits when used with Data Guard:

- Flashback Database removes the need to re-create the original primary database after a failover. Following a failover, the original primary database can be flashed back to a point in time before the failover occurred and converted into a standby database. At this point, Data Guard can automatically resynchronize the original primary database (now a standby) with the new primary database. Once all redo backlog has been applied a switchover can be executed that will return all databases back to their original roles.
- Flashback Database provides an alternative to the method of delaying the application of redo data to the standby database to enable fast recovery from user errors or logical corruptions. Flashback Database, a “rewind button” for Oracle databases, makes it quicker and easier to recover from such events while allowing a standby database to always be completely synchronized with the primary database, reducing failover and switchover time.

## Fast-Start Failover

Fast-Start Failover is a Data Guard 10g Release 2 feature. It enables an automatic failover to a previously chosen, synchronized standby database in the event of loss of the primary database, without requiring any manual steps to invoke the failover. Fast-start Failover is used in a Data Guard Broker configuration and is configured through DGMGRL or Enterprise Manager. Fast-Start Failover also requires that the Data Guard configuration use the Maximum Availability Protection mode.

A fast-start failover configuration is monitored by a separate Data Guard “Observer” process. The Observer is a lightweight process integrated in the DGMGRL (Data Guard Broker) client-side component. It runs on a different computer from the primary or standby databases. It continuously monitors the fast-start failover configuration to ensure the primary database is available. If both the Observer and the standby database lose connectivity to the primary database, the Observer waits for a configurable amount of time and then initiates a fast-start failover without any human intervention.

## CONCLUSION

NeuStar has made extensive use of Data Guard capabilities to implement zero data loss protection across a Wide Area Network as a central component of its Business Continuity strategy. NeuStar utilizes the inherent advantages of Data Guard’s synchronous redo transport services to achieve geographic separation between primary and DR sites to an extent that cannot be achieved using traditional remote disaster recovery solutions [6].

## **ADDITIONAL REFERENCES**

1. Oracle Data Guard Overview -

<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. Oracle10g Data Guard: Primary Site and Network Best Practices -

[http://www.oracle.com/technology/deploy/availability/pdf/MAA\\_WP\\_10gR2\\_DataGuardNetworkBestPractices.pdf](http://www.oracle.com/technology/deploy/availability/pdf/MAA_WP_10gR2_DataGuardNetworkBestPractices.pdf)

4. Media Recovery Best Practices for Redo Apply (Physical Standby)

[http://www.oracle.com/technology/deploy/availability/pdf/MAA\\_WP\\_10gR2\\_RecoveryBestPractices.pdf](http://www.oracle.com/technology/deploy/availability/pdf/MAA_WP_10gR2_RecoveryBestPractices.pdf)

5. Switchover/Failover Best Practices

[http://www.oracle.com/technology/deploy/availability/pdf/MAA\\_WP\\_10gR2\\_SwitchoverFailoverBestPractices.pdf](http://www.oracle.com/technology/deploy/availability/pdf/MAA_WP_10gR2_SwitchoverFailoverBestPractices.pdf)

6. Oracle Data Guard and Remote Mirroring Solutions -

<http://www.oracle.com/technology/deploy/availability/htdocs/DataGuardRemoteMirroring.html>



**Data Guard OTN Case Study**

**Authors: Larry Carpenter, Manoj Gupta, Joseph Meeks & Ashish Ray  
December 2008**

**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 © 2008, 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.**