# Fannie Mae & Data Guard 10g Release 2
## Zero Data Loss Protection at Very High Throughput

*"Data Guard is our preferred Disaster Recovery solution for Oracle data. We have very high performance applications. No other 3rd party product has approached the throughput that Data Guard can attain in zero data loss configurations."*

*Manohar Malayanur*
*Manager and Infrastructure Architect*

**Corporate Profile**

- **A private, shareholder-owned company.**
- **Second largest US Corporation measured in value of assets.**
- **4,700 employees**
- **www.fanniemae.com**

**Disaster Recovery**

- **Oracle Data Guard**
- **Redo Apply - Physical Standby**
- **Maximum Availability - Zero Data Loss**
- **Oracle Database 10g Release 2**
- **Zero data loss for redo volumes as high as 21.4MB/sec**
- **Asynchronous protection with minimal data loss exposure for volumes as high as 29MB/sec**
- **Following Maximum Availability Architecture Best Practices yields performance gains as high as 80%**

## OVERVIEW

Fannie Mae has deployed Oracle Data Guard [1] extensively for more than 24 production databases (Oracle9i & Oracle Database 10g). In preparation for upgrading to Oracle Database 10g Release 2, Fannie Mae tested the performance of synchronous (zero data loss) and asynchronous (minimal data loss) Data Guard redo transport modes. Their results are provided in the table below.

| Feature tested | Max redo rate achieved |
|---|---|
| Synchronous transport: zero data loss | 21.4 MB/sec |
| Asynchronous transport: near zero data loss | 29 MB/sec |
| Redo apply performance at standby database | 32 MB/sec |

Tests proved that zero data loss protection was practical for application workloads generating redo at rates up to 21.4MB/sec. Testing also demonstrated how Data Guard's flexible configuration options can sustain database throughput as high as 29MB/sec using asynchronous transport mode, for applications where the potential for a small amount of data loss is an acceptable trade-off in return for maximum database performance. Finally, redo apply performance tests confirmed that redo data could be applied to the standby database as quickly as it is received, minimizing the time required to execute switchover/failover operations.

Fannie Mae also observed how by following Oracle Maximum Availability Architecture (MAA) [2] best practices for optimal tuning recommendations, it was possible to achieve significant gains in network throughput compared to using default parameters.

## INTRODUCTION

Fannie Mae is a private, shareholder-owned company whose mission is to insure that mortgage money is available for people in communities across America. Fannie Mae does not lend money directly to homebuyers. Instead, it works with lenders to make sure mortgage funds are always available, helping more people to achieve the dream of home ownership. Fannie Mae has helped make possible over $6.3 trillion in home financing for more than 63 million American families.

**System Configuration**

- **Primary & Standby servers: SUN E25K, each with196GB memory and 48 dual core CPUs**
- **(2) EMC DMX3000 SAN with 16FA ports**
- **1Gbit network bandwidth**
- **4ms RTT (round trip network latency)**
- **Oracle Database 10g Release 2**
- **2GB online logs**

# TEST OBJECTIVES:

Fannie Mae utilizes Data Guard for databases where high performance and superior data protection are critical requirements. In preparation for upgrading to Oracle Database 10g Release 2 performance tests were conducted to assess Data Guard configurations using synchronous and asynchronous redo transports. Fannie Mae had the following test objectives:

- Determine the maximum production database throughput that can be sustained in a synchronous configuration with zero data loss protection
    - Data Guard Maximum Availability protection mode and LGWR SYNC redo transport services.
- Determine the maximum production database throughput that can be sustained in an asynchronous configuration with minimal data loss exposure.
    - Data Guard Maximum Performance protection mode and LGWR ASYNC redo transport services.
- Determine the maximum apply throughput possible on the standby database using Data Guard Redo Apply (physical standby). The obvious goal is to apply redo as fast as it is received, such that the standby database is always up to date and there is no backlog of redo processing that would delay failover or switchover during role management operations.
- Quantify the benefits of MAA recommendations for tuning network parameters to assess the potential value of engaging in such efforts across Fannie Mae's many Data Guard configurations.

# TEST CONFIGURATION & PROCESS

The primary and standby database servers in Fannie Mae's Data Guard configuration are SUN E25K servers located in separate facilities that are on a metropolitan area network (MAN) several kilometers apart and connected by a 1 Gbit network link. Average RTT network latency (round trip) between primary and standby is 4ms. The workload generated by the test application grows until it consumes 100% of available capacity or until some other resource limit places a cap on throughput.

Fannie Mae's test scenarios measured the following performance aspects of a Data Guard configuration:

- **Data Guard Redo Transport Services:** These services control the automated transfer of redo data from a production or primary database to one or more remote (standby) database destinations. The transport layer also manages the process of resolving any gaps in redo log files that can occur due to a network failure or standby server outage. From a performance perspective, Data Guard redo transport services must be able to ship redo data from the primary to the standby server quickly enough to achieve Recovery

Point Objectives. If the rate that the redo is shipped falls behind the rate that it is generated on the primary server, a backlog of unshipped redo accumulates on the primary server increasing the exposure to data loss should the primary server fail. The two modes of transport tested by Fannie Mae and discussed in this paper are synchronous (LGWR SYNC) and asynchronous (LGWR ASYNC) transport modes. An overview of Data Guard redo transport services is provided in Figure 1.

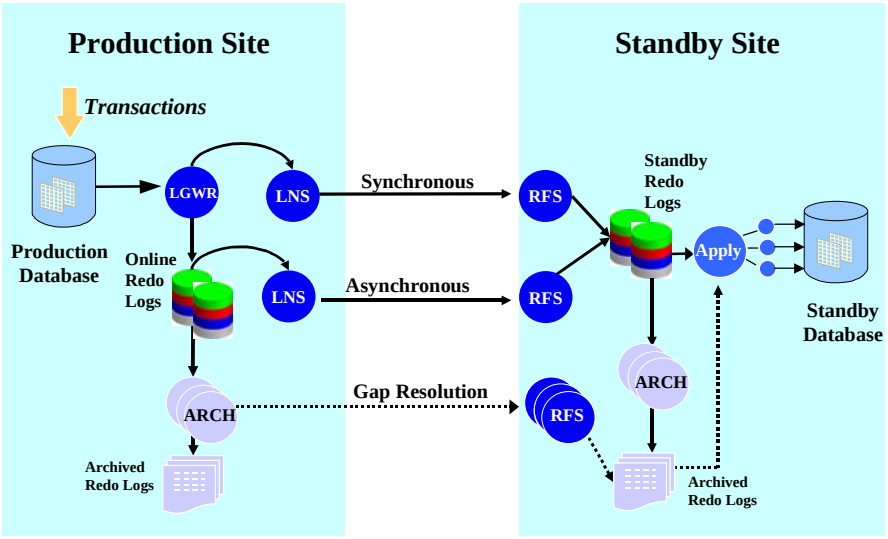


Figure 1: Data Guard Transport Services

- **Data Guard Redo Apply Services** automatically apply redo data on the standby database to maintain consistency with the primary database. Fannie Mae's Data Guard configuration uses Redo Apply technology, which maintains a physical replica of the primary by applying redo data to the standby database. Redo Apply is a very efficient process that is able to scale with the highest level of workload. Because Redo Apply can keep the standby database in sync with the primary, database switchover/failover operations can be completed in seconds, enabling the most aggressive Recovery Time Objectives.

# TEST RESULTS

## Synchronous Redo Shipping – Zero Data Loss

Synchronous redo shipping is enabled by configuring a remote destination to use the LGWR SYNC redo transport. On the primary database, once the redo buffers for committed transactions are written to an online redo log, a Data Guard process (LNS) ships the redo directly from the in-memory buffer to a remote standby database destination. A Data Guard process (RFS) at the remote standby database receives the redo, and writes it to disk into a standby redo log file (SRL). When this write is complete, an acknowledgment is returned to the primary database. The primary database then notifies the client that the commit is complete, enabling the

**Test 1 – Zero Data Loss**

- **LGWR SYNC**
- **21.4 MB/sec maximum primary database throughput achieved using synchronous redo shipping**

LGWR process on the primary database to proceed with the next set of committed transactions.

Fannie Mae testing measured the maximum network throughput that could be achieved using this synchronous mechanism in terms of megabytes of redo data generated/second (MB/sec).

Sufficient network bandwidth was made available such that it did not impact the test results. Network latency was minimal, simulating a Data Guard configuration where primary and standby would be located within the same metropolitan area.

The maximum database and network throughput achieved using synchronous transport services was 21.4 MB/sec of redo data.

This strong performance was influenced by the fact that Fannie Mae's application is characterized by batch processing of large transactions. Data Guard redo transport service using LGWR SYNC is very efficient with this transaction profile when combined with a high-bandwidth, low latency network. While the efficiency of Data Guard synchronous transport services will benefit any transaction profile compared to traditional replication methods such as remote-mirroring [3], Oracle always recommends a high priority be placed on performance testing when using synchronous redo transport since different network latencies and transaction profiles can yield different results

## Asynchronous Redo Shipping – Minimal Data Loss

**Test 2 - Minimize Potential Data Loss while Maximizing Performance**

- **LGWR ASYNC**
- **29MB/sec maximum network throughput achieved by asynchronous redo transport services**

Asynchronous redo shipping is enabled by configuring a remote destination to use the LGWR ASYNC redo transport. In contrast to synchronous transport, once the LGWR writes redo data to the local online redo log, it does not wait for acknowledgement by the standby database. Asynchronously, a Data Guard network process (LNS) reads from the local online log and transmits the redo to the remote database destination. The ASYNC attribute of this transport mode causes all network I/O to be performed asynchronously and control is returned to the executing application immediately. The LGWR process continues processing the next request without waiting for the network I/O to complete.

Asynchronous redo shipping achieves the goal of minimal or no performance impact on the primary production system.

The maximum database and network throughput Fannie Mae was able to achieve using LGWR ASYNC was 29MB/second of redo data.

**Test 3 – Standby Database Apply Throughput**

- **Redo Apply – Physical Standby Database**
- **32 MB/sec maximum redo apply throughput on standby database**

## Redo Apply Performance on Standby

Testing proceeded to the final element of Data Guard performance – the apply rate – or the rate at which redo data ca n be applied to the standby database. The maximum apply rate achieved in Fannie Mae testing was 32 MB/second.

## Optimizing Network Throughput

Tuning of network parameters was required in order to achieve the throughput reported in this paper. Fannie Mae optimized network settings by referring to the Maximum Availability Architecture white paper "Data Guard Redo Transport and Network Best Practices" [4].

The most significant gain was realized by increasing the TCP socket buffer size. From the Redo Transport and Network Best Practices white paper:

- TCP socket buffer settings will control how much network bandwidth can be used. Socket buffer sizes need to be increased from their default values in order to improve utilization of available bandwidth. When network latency is high, larger socket buffer sizes are needed to fully utilize network bandwidth.

Fannie Mae determined that a 3MB buffer size was optimum for their configuration. This produced gains of over 30% for synchronous redo transport, and up to 80% for asynchronous redo transport. Details for determining optimum value for TCP socket buffers and for several other parameters key to optimizing Data Guard network utilization are provided in the MAA best practice paper referenced above.

# CONCLUSION

Fannie Mae requirements are simple – protect mission critical data and keep it available. Achieving these simple goals becomes very challenging due to the very high volume of data generated by Fannie Mae applications. Fannie Mae has deployed Data Guard on many systems with high levels of confidence that come from extensive pre-production test experience. Fannie Mae knows it can achieve zero data loss protection for very demanding applications. For applications where the performance impact of synchronous data protection is beyond acceptable levels, Fannie Mae takes advantage of Data Guard's flexibility by utilizing asynchronous transport services to maximize performance while keeping potential data loss exposure to a minimum.

# 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. Data Guard & Remote Mirroring http://www.oracle.com/technology/deploy/availability/htdocs/DataGuardRemoteMirroring.html
4. Oracle Database 10g Release 2, Data Guard Redo Transport and Network Best Practices - http://www.oracle.com/technology/deploy/availability/pdf/MAA_WP_10gR2_DataGuardNetworkBestPractices.pdf

**Fannie Mae OTN Case Study – Data Guard 10g Release 2**
**August 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**