ORACLE OPEN WORLD

Your. Open. World.

Oracle Active Data Guard

Standby on Steroids – DR Included

Joe Meeks, Director, Product Management, Oracle
Shawn Ormond, Database Administrator, Intermap Technologies Inc
Yucheng Liu, Senior Database Administrator, Real Networks
Krishna Kakatur, Senior. Database Administrator, Real Networks

ORACLE

Today's Objectives

Deep Dive

Shine the Light

Program

- Ship Redo / Apply Redo
- Utilize your standby database
  - Active Data Guard
  - Snapshot Standby
- Reduce downtime
- Active Data Guard experiences
  - Real Networks
  - Intermap Inc
- Resources - Q&A

Focus on Oracle Database 11*g* and Redo Apply (physical standby)

Oracle Data Guard

Best Protection at Lowest Cost

Ship Redo

Synchronous Redo Transport (SYNC)– Zero Data Loss

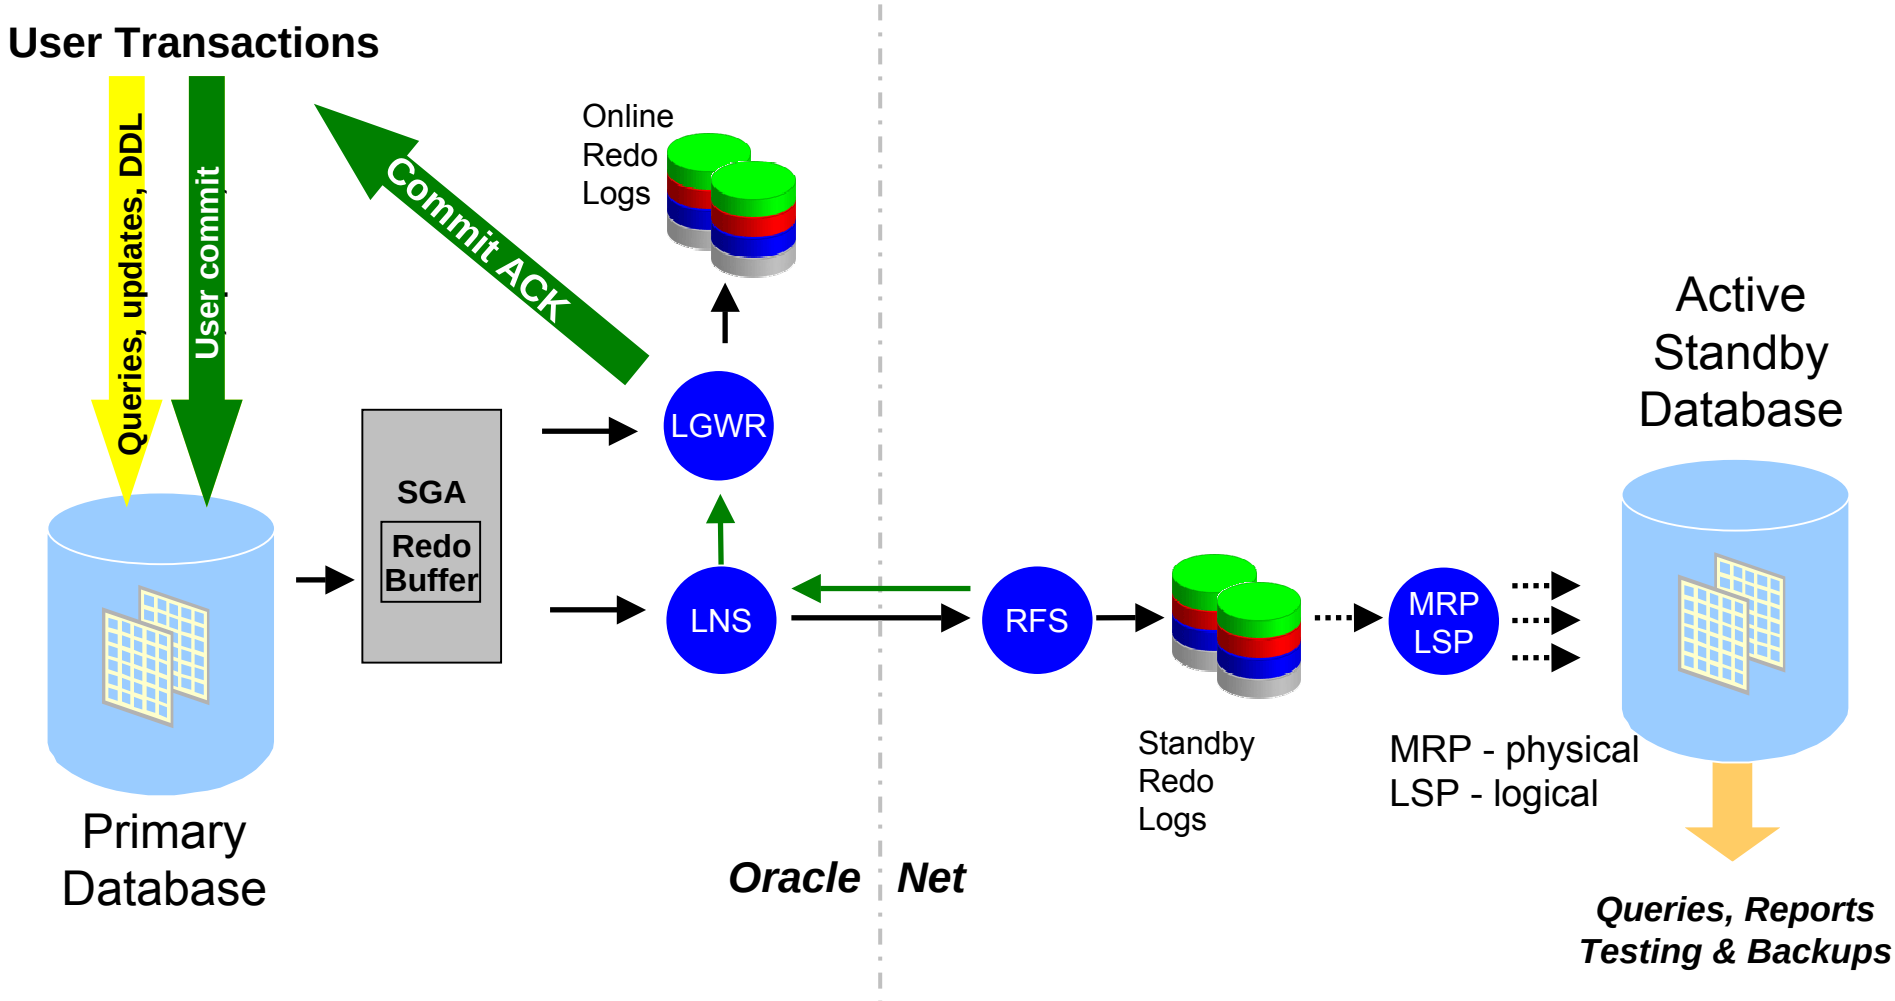

Ship Smart

Just the Redo . . .

- Data Guard ships only redo records
- SCN aware
  - Enables reliable recovery
  - Guarantees commits are applied in order
- Storage remote-mirroring must ship every write
  - 7x greater volume and 27x more network I/Os than Data Guard
  - Round-trip network latency impacts EVERY write to EVERY file

- Data Guard Compared to Storage Remote-Mirroring
  http://www.oracle.com/technology/deploy/availability/htdocs/DataGuardRemoteMirroring.html

Ship Redo

Asynchronous Redo Transport (ASYNC)

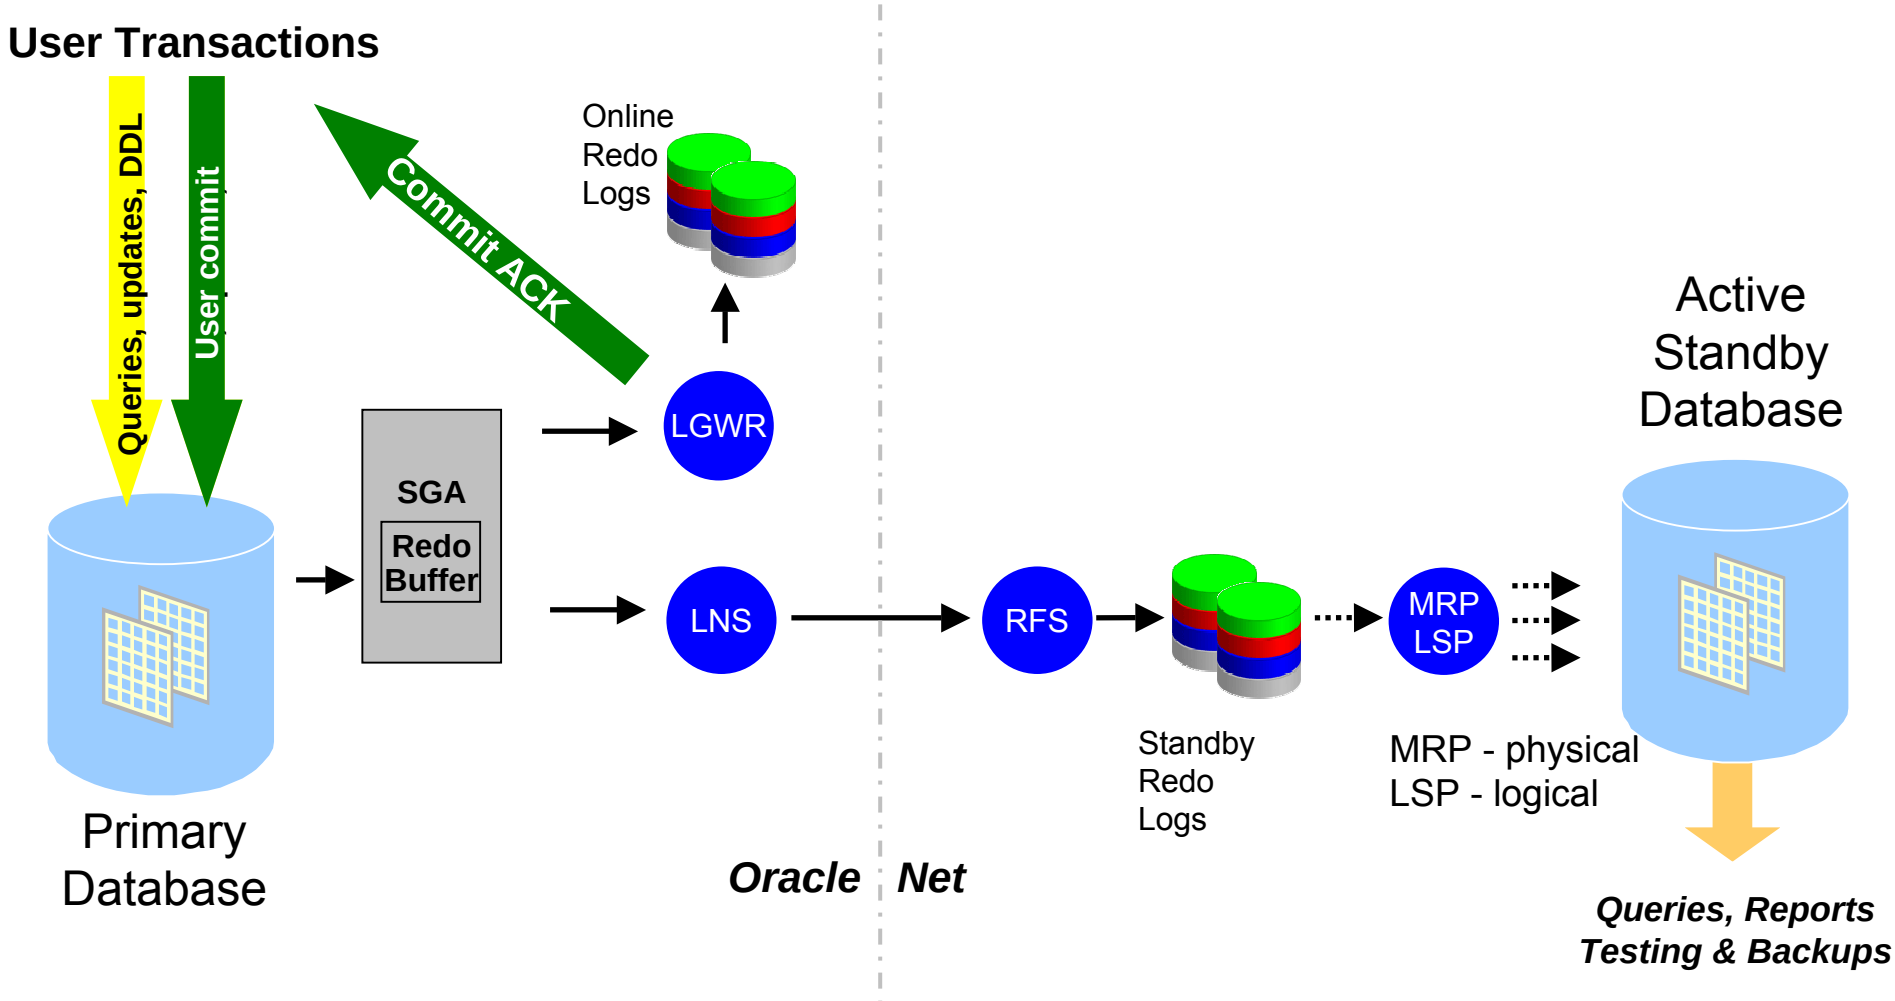

Ship Redo

ASYNC – If Network Can't Keep Pace

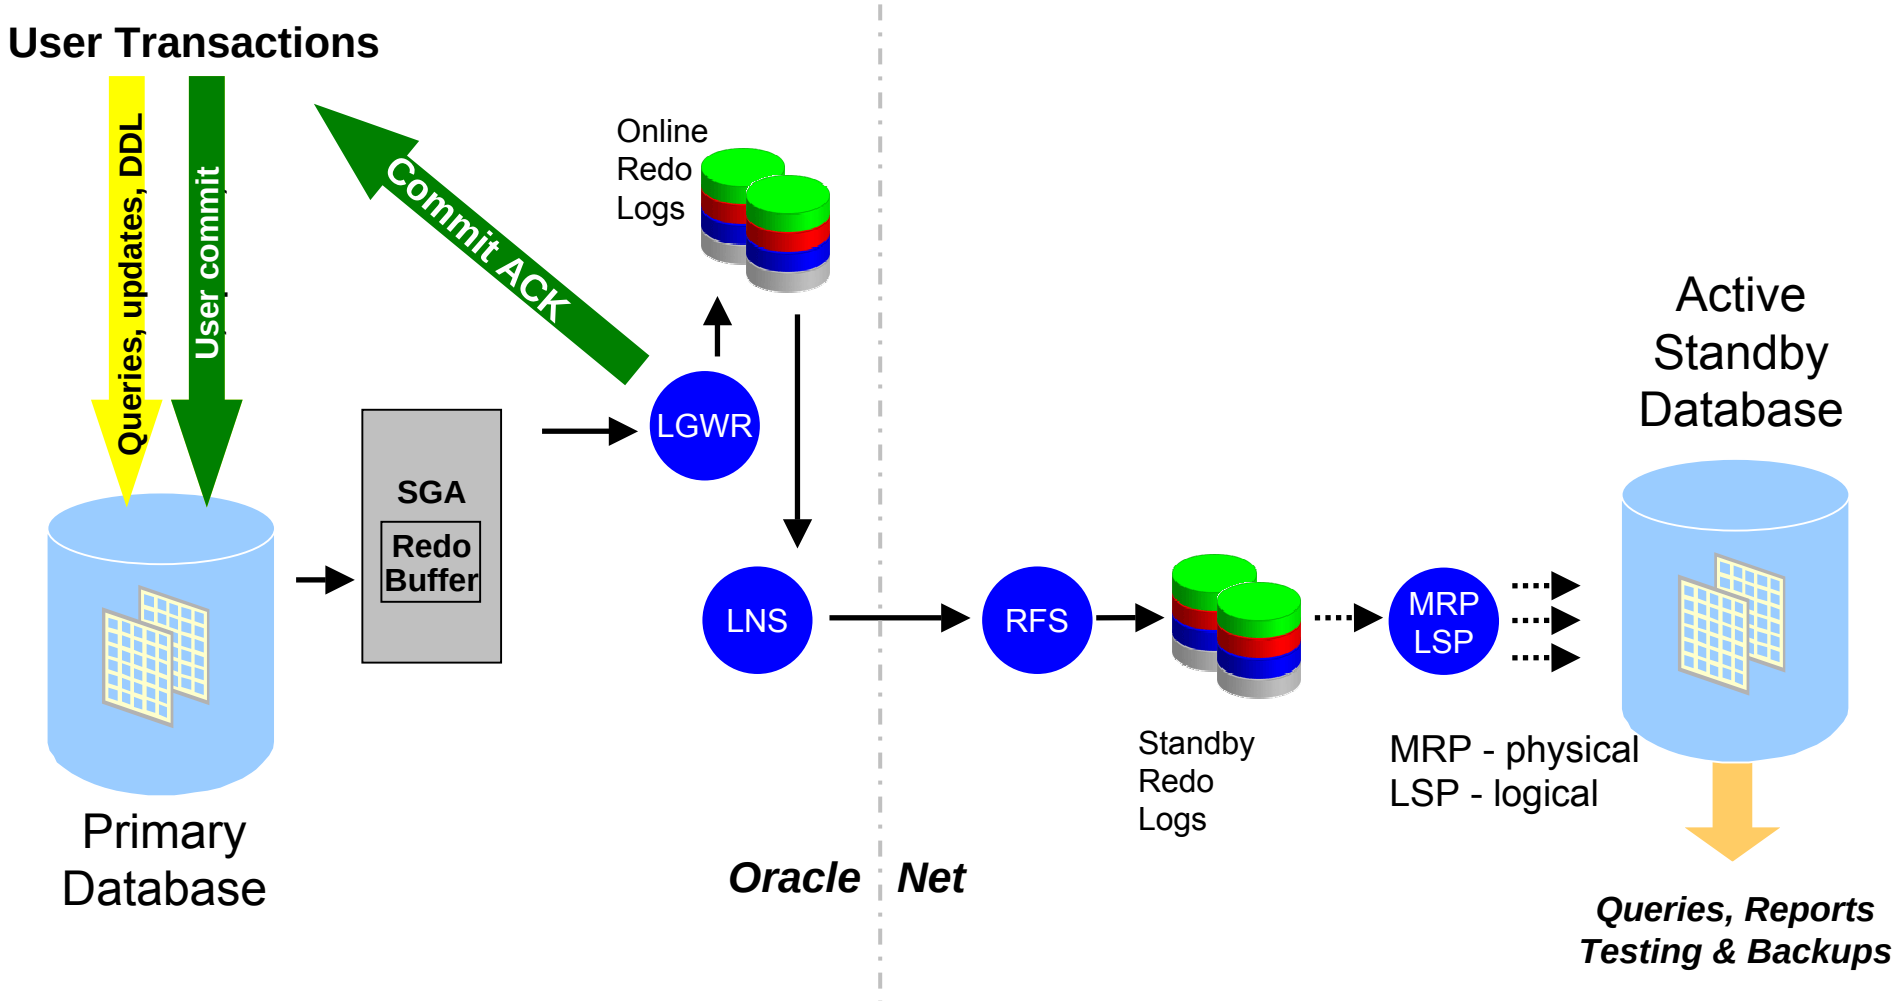

Shipping vs. Protection Mode

Protection Mode Controls Response to Failure Events

Mode	Risk of data loss	Transport	If no acknowledgement from standby:
Maximum Protection	Zero Data Loss Double Failure Protection	SYNC	Stall primary until acknowledgement is received from replica
Maximum Availability	Zero Data Loss Single Failure Protection	SYNC	Stall primary until acknowledgement is received or timeout threshold period expires – then resume processing
Maximum Performance	Potential for Minimal Data Loss	ASYNC	Primary never waits for standby acknowledgement

Shipping vs. Protection Mode

Protection Mode Controls Response to Failure Events

Mode	Risk of data loss	Transport	If no acknowledgement from standby:
Maximum Availability	Zero Data Loss Single Failure Protection	SYNC	Stall primary until acknowledgement is received or timeout threshold period expires – then resume processing

`NET_TIMEOUT parameter of LOG_ARCHIVE_DEST_n`

Data Guard 11*g* default = 30 seconds

Data Guard 10*g* default = 180 seconds

Ship Redo

Automatic Gap Resolution – Primary Pings Standby

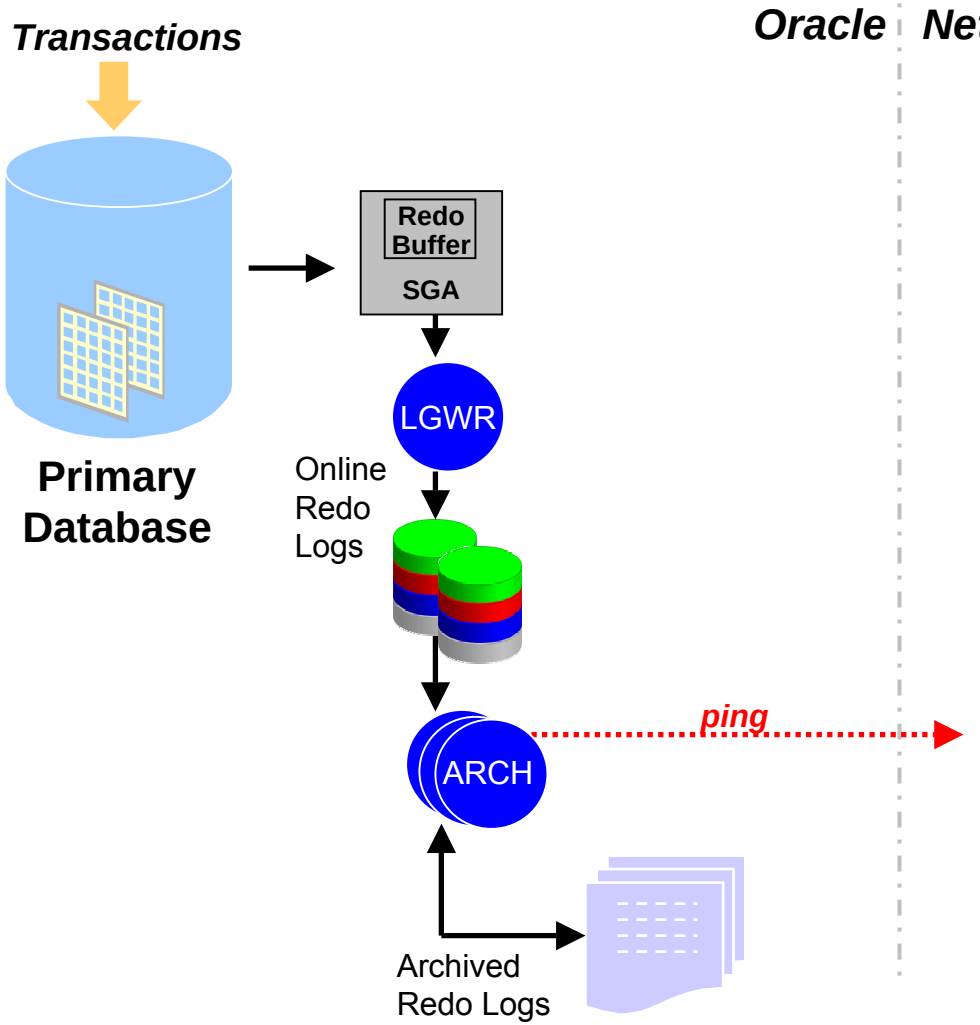

Ship Redo

Automatic Gap Resolution – Connect to Standby

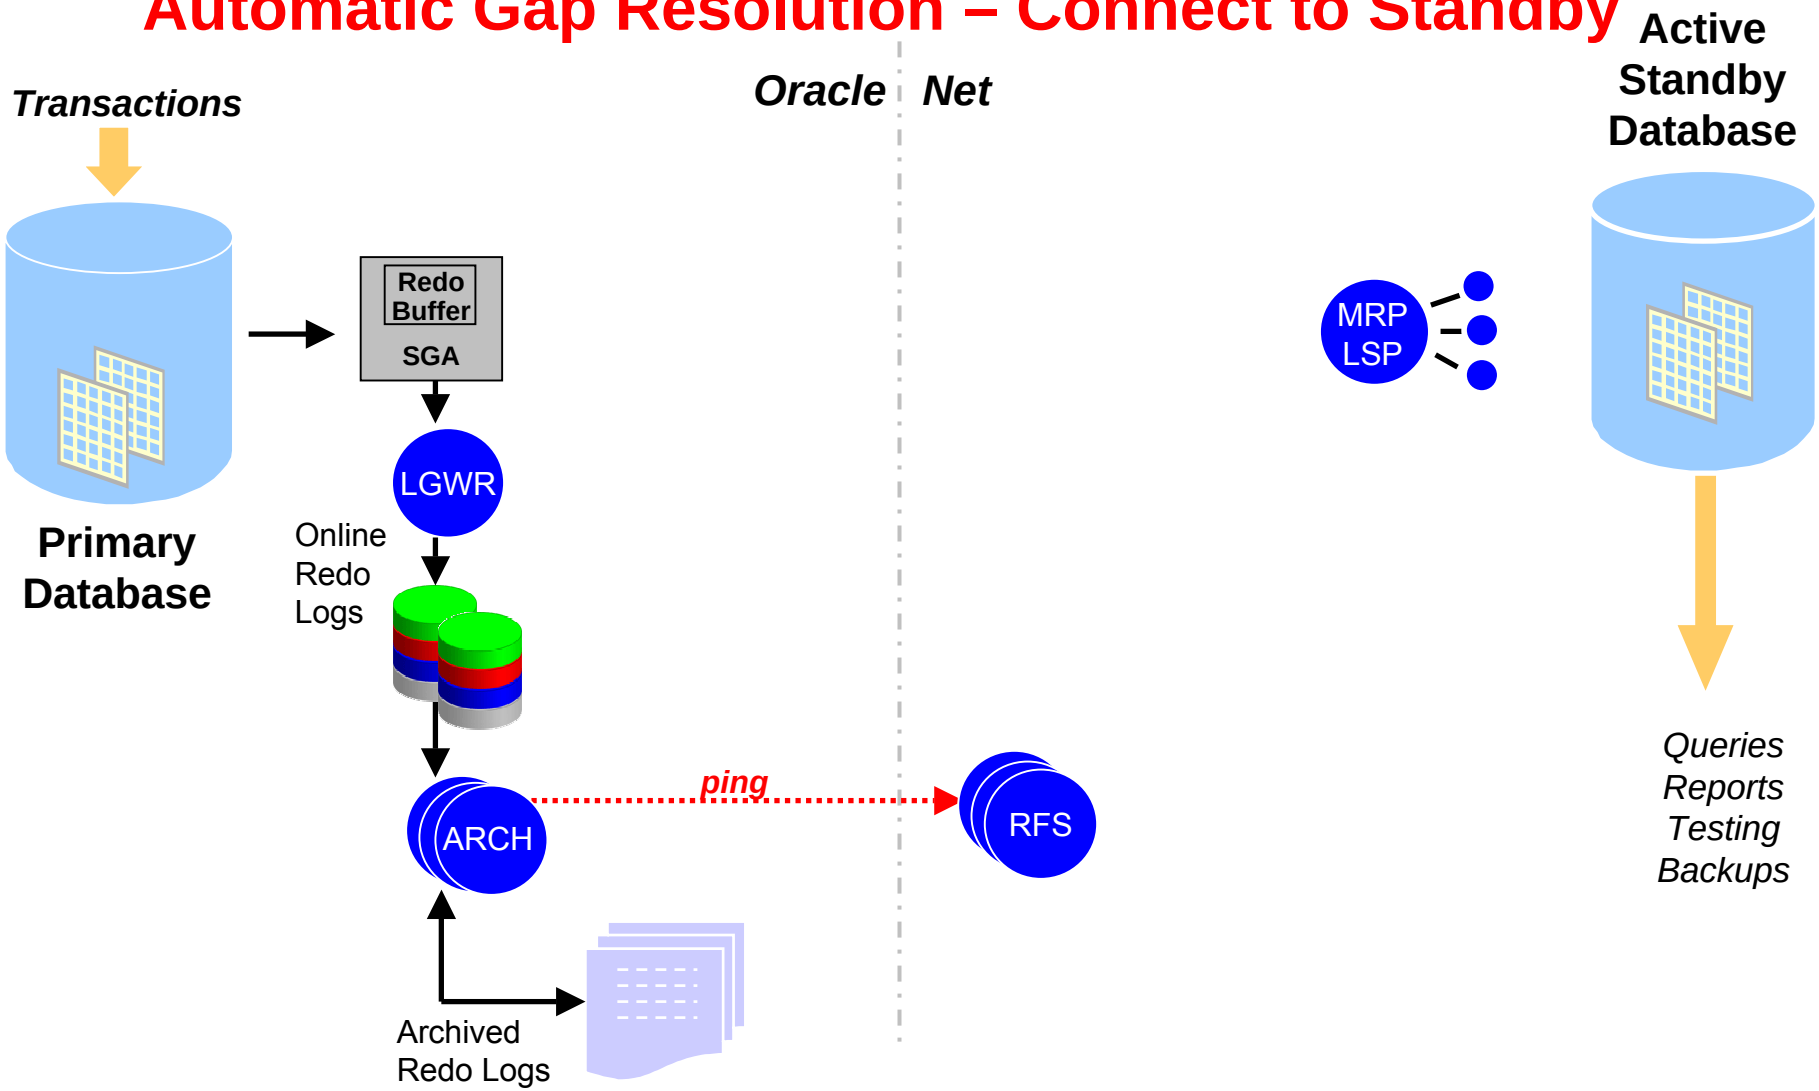

Ship Redo

Automatic Gap Resolution – Ship Log Gap

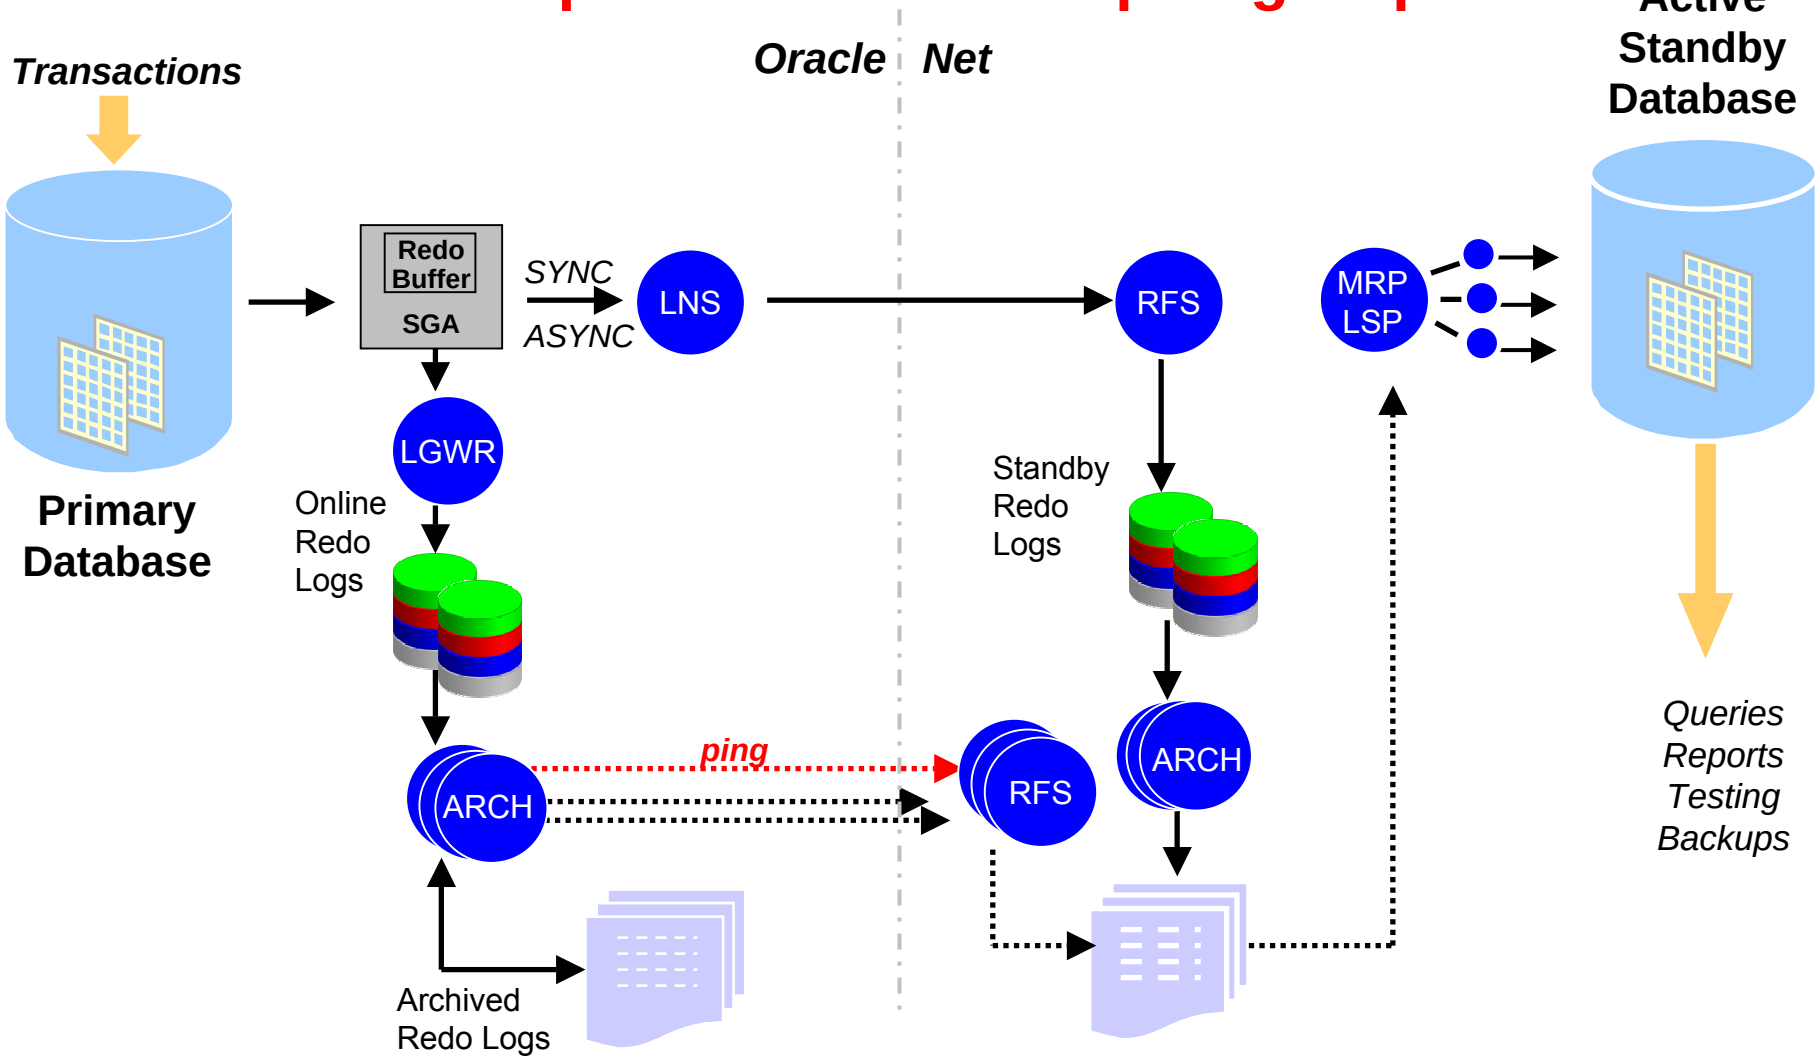

Ship Redo

Automatic Gap Resolution – Gap Resolved

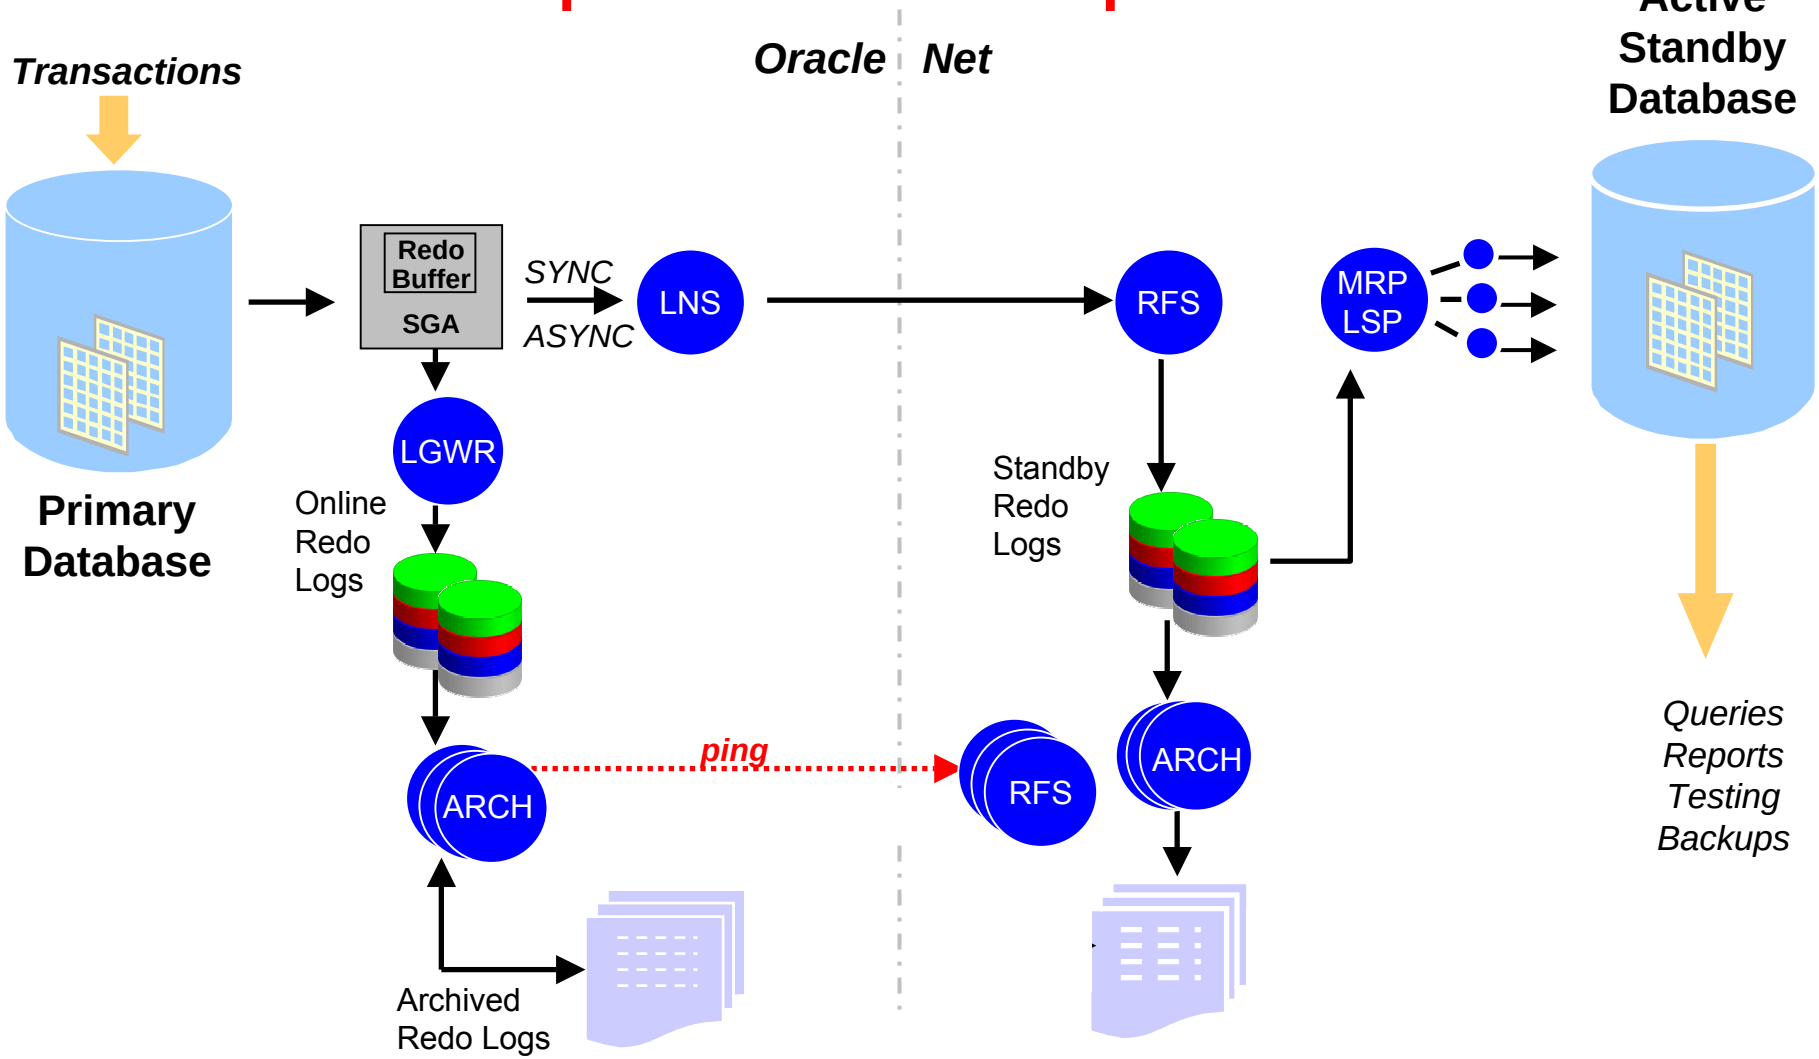

Ship Fast

Network Compression for Gaps

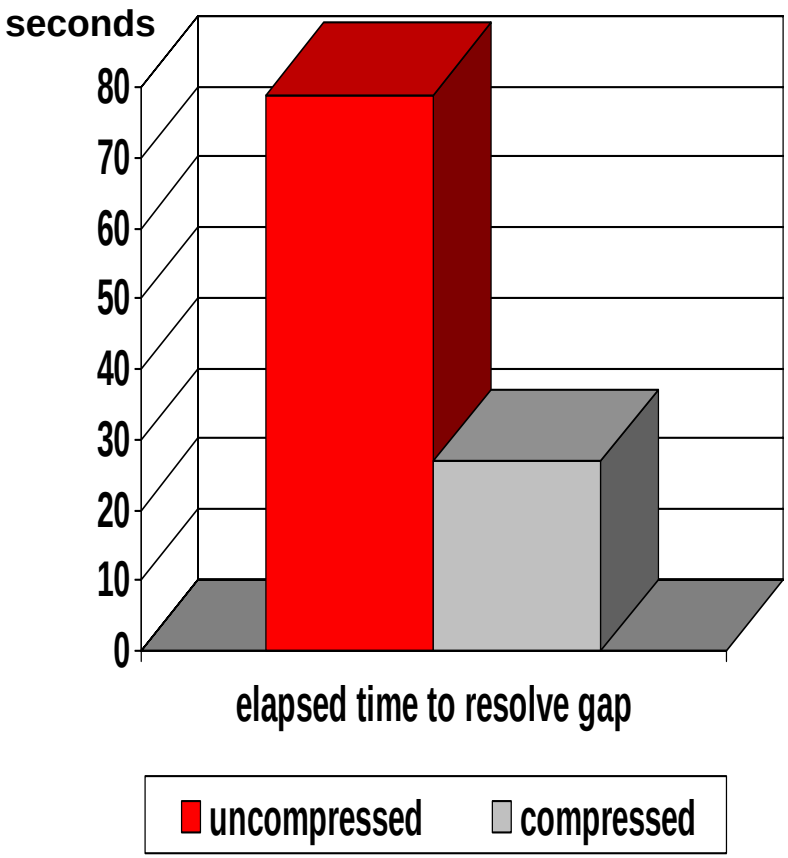

- To enable compression:
  - Set Data Guard broker property, or
  - Set compression attribute of redo transport destination
- Resolves gaps up to 3x faster
  - Better data protection
- Given there is sufficient CPU
  - Negligible impact on response time
  - Negligible impact on throughput
- Requires Oracle Advanced Compression Option 11*g*

Ship Fast

Enable Compression for ASYNC Transport

```
_REDO_TRANSPORT_COMPRESS_ALL=TRUE
```

- Useful when network volume exceeds bandwidth
- Test case:
  - Network Bandwidth: 100Mbps network (12.5MB/s)
  - Redo Rate: 22MB/s
- Results
  - Without compression, transport lag increased linearly over time
  - With compression enabled, transport lag ranged from 4-10 seconds
  - Compression ratio: 60%
- Implementation details - see MetaLink Note 729551.1

Apply Redo

Redo Apply (physical standby) Parallel Media Recovery

- MEDIA RECOVERY COORDINATOR (MRP0)
  - Manages recovery session, merges redo by SCN from multiple instances, parses redo into change mappings partitioned by apply slave
- APPLY SLAVES
  - Read data blocks, assemble redo changes from mappings, apply redo changes to data blocks

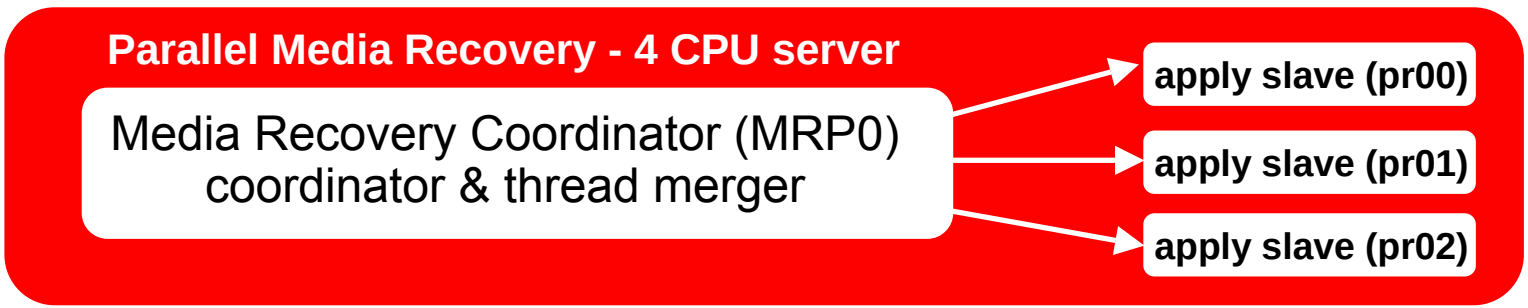

- Automatically configures the # of slaves = # CPUs - 1

Apply Fast

100% Faster than Oracle Database 10*g*

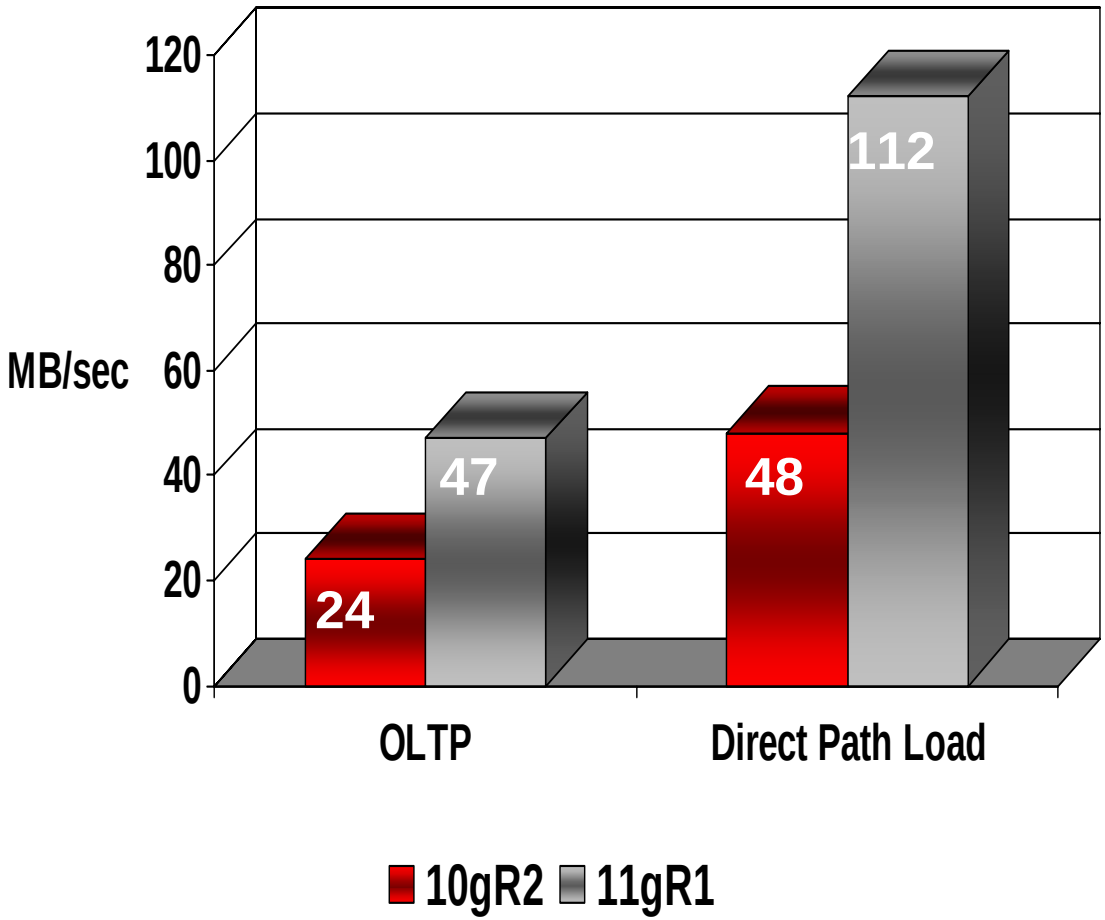

- Increased parallelism
- Less synchronization
- Better utilization of I/O and CPU resources
- Optimizations for direct-path loads
- Self-configuring*

*for ASYNC I/O

Apply Safely

Lost Write Detection

- What is a Lost Write:
  - Storage loses a write that it has acknowledged to Oracle as complete
- Subsequent transactions read stale version of the block and either:
  - Update the same block again
  - Update another block
  - Do something external: print a check, generate an invoice, issue an order
- Primary may continue running for hours or days
  - It may generate an assortment of internal errors, e.g. ORA-00600:[4135], or [4137], or [4152], or [qertbFetchByRowID], depending upon the objects impacted and the writes that are lost
  - Primary may eventually crash
- Any recovery of a block that is victim of a lost write will fail
    - ORA-600 [3020] stuck recovery error

Lost Write Happens

As Reported in SR - Oracle Database 10*g* Release 2

Lengthy outage impacting a multi-terabyte database

- Problems first surface on their standby database

ORA-00600: internal error code, arguments**: [3020]** , [648], [1182463], [2719091455], []
ORA-10567 : Redo is inconsistent with data block (file# 648, block# 1182463)
Recovery interrupted!

- Many hours later – production is down

Noticed odd query results on production
Noticed ORA-600 errors on production this morning for which SGA Heapdump was uploaded.
New info : I was rebuilding an index. After a few minutes, the database took an unexpected crash.
*****please help. it's very urgent, production is down.*****

- Problems traced to lost writes caused by faulty hardware

Not a Problem for Data Guard 11*g*

Capability Unique to Oracle Database

- Detect lost writes using new initialization parameter

```
db_lost_write_protect
```

- Apply compares standby version of block to incoming redo
  - ORA-752 if block SCN from primary is lower than standby
    - 100% certain of a lost write on the primary database
    - Resolve via failover to standby to restore data consistency
  - ORA-600 [3020] if block SCN from primary is higher than standby
    - Possibility of a lost write on the standby database
    - Resolve by re-creating the standby database or affected files

Program

- Ship Redo / Apply Redo
- Utilize your standby database
  - Active Data Guard
  - Snapshot Standby
- Reduce downtime
- Active Data Guard experiences
  - Real Networks
  - Intermap Inc
- Resources - Q&A

Data Guard 11*g*

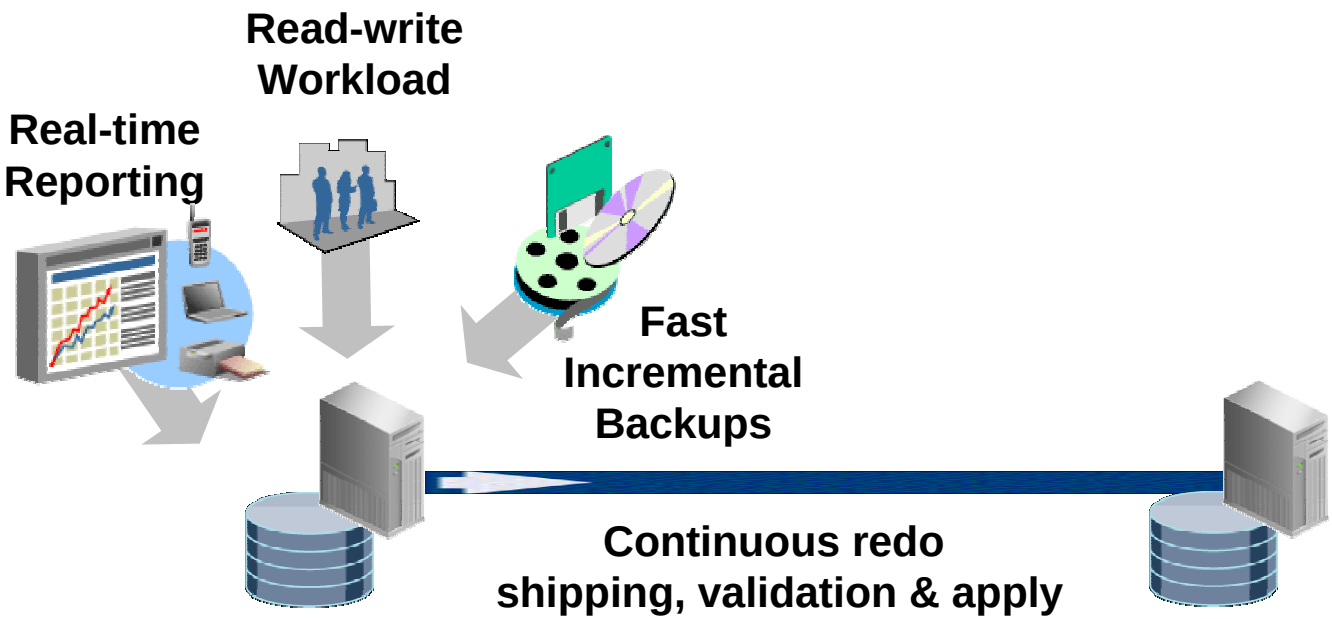

Active Data Guard 11*g*

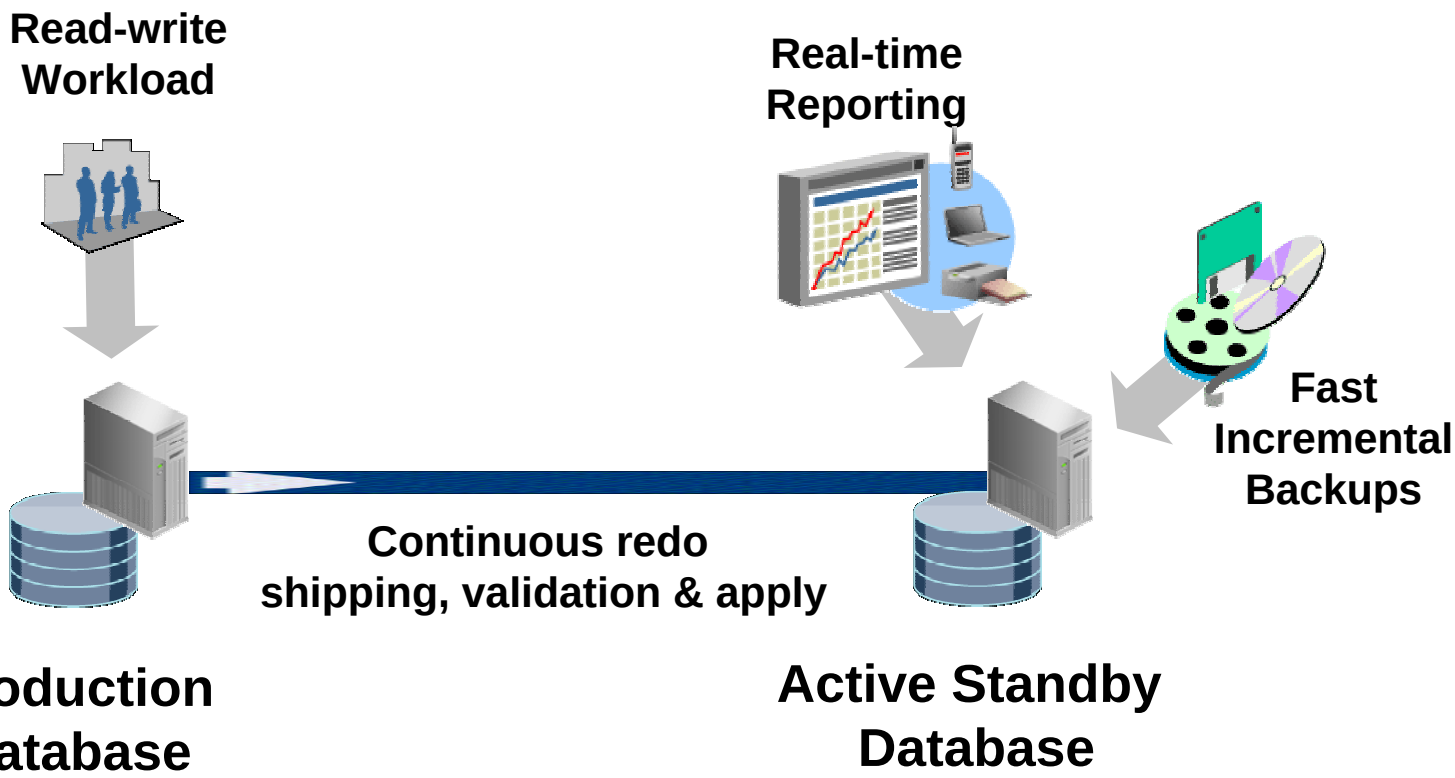

- Offload read-only queries to an up-to-date physical standby
- Use fast incremental backups on a physical standby – up to 20x faster

What's so Different?

Data Guard 11*g*	Active Data Guard Option
• Stop redo apply at 8am • Open read-only for queries	• Redo apply is always on • Always open read only
• By 4pm, data is 8 hours old	• Queries and reports always see latest data
• Any failover will be delayed due to backlog of data that must be applied	• Failover is immediate when needed, standby database always up-to-date

- Active Data Guard MAA Best Practices
  http://www.oracle.com/technology/deploy/availability/pdf/maa_wp_11gr1_activedataguard.pdf

Active Data Guard

Utilize Standby Databases

Active Data Guard

Queries Return Up-to-date Results

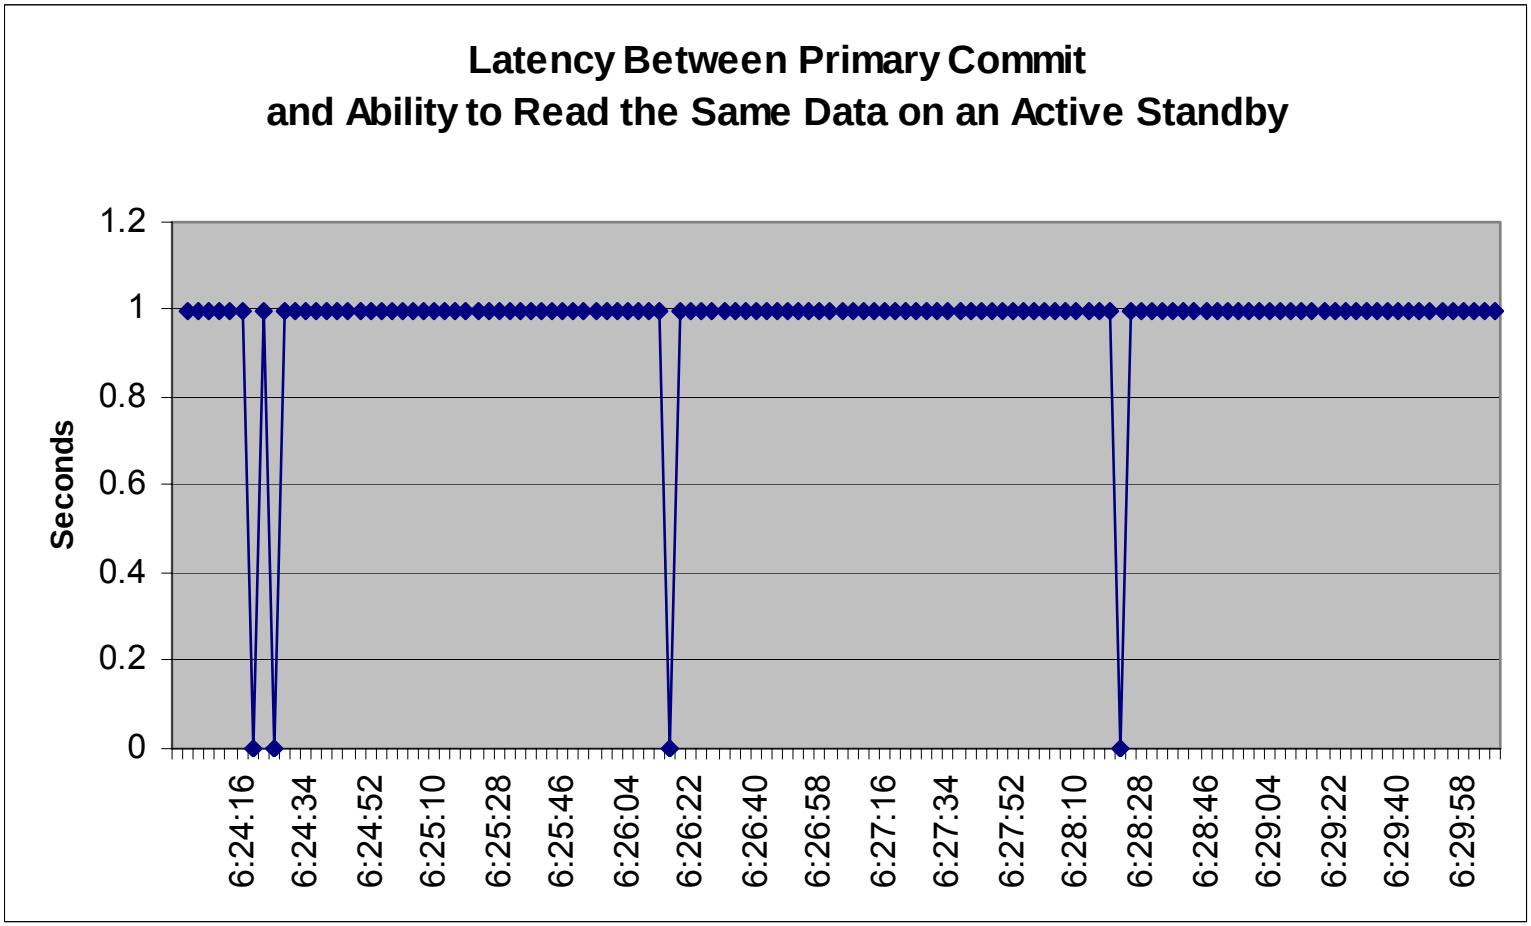

Oracle Business Intelligence Suite

Release 10.1.3.4 Certified for Active Data Guard

- Oracle Business Intelligence Suite EE Plus
  - Suite of BI products offering full range of analysis and reporting
  - Includes Oracle Hyperion reporting products
- Oracle BI server runs on an Active Standby Database
  - Oracle BI server is a read-mostly application
- Configuration highlights
  - Disable BIEE server from creating temp tables on standby
  - Create read-only connection pool
  - Create a write-back connection pool to redirect writes to the primary or a local ‘scratch’ database

- Oracle Business Intelligence and Active Data Guard MAA Best Practices
  http://www.oracle.com/technology/deploy/availability/htdocs/maa.htm

Offload Backups

RMAN Block Change Tracking

- Offload fast incrementals to an Active Data Guard Standby
  - Block change tracking eliminates full scans
  - Incremental backups complete 20x faster (8.3 min vs 2.8 hrs)
  - Minimal overhead on standby database – less than 3%
- If currently using a split mirror to offload backups – consider repurposing that storage
  - Deploy a local standby database instead
  - Realize better HA, better data protection and more reliable backups.
  - Use standby to offload query workload and/or serve as test system

- S298772 - Oracle Recovery Manager (RMAN) Best Practices for Oracle Data Guard and Oracle Streams, Wednesday, 11:30 am – 12:30 pm, Moscone South Room 103

Test System

Oracle MAA Partners Helping with Active Data Guard Testing

EMC CX3-40F UltraScale Storage Systems

- Flare Release 26
- 4 GB RAM per SP
  - Write Cache = 2GB
  - Read Cache = 1GB per SP
- 60 146GB FC drives @ 15K RPM
- All LUNs bound as 1+1 Raid 10
  - Non Vault DATA LUNs 133 GB
  - Vault DATA LUNS 99 GB
  - LUN Prefetch set to Variable with default settings

Dell 6950s

- 4 way Dual-Core AMD Opteron Processor 8212
- 8 GB RAM
- OEL 4.5 x86_64 (2.6.9-55.0.0.0.2.ELsmp)

Snapshot Standby – Data Guard 11*g*

Convert Standby Database to Read-Write Test System

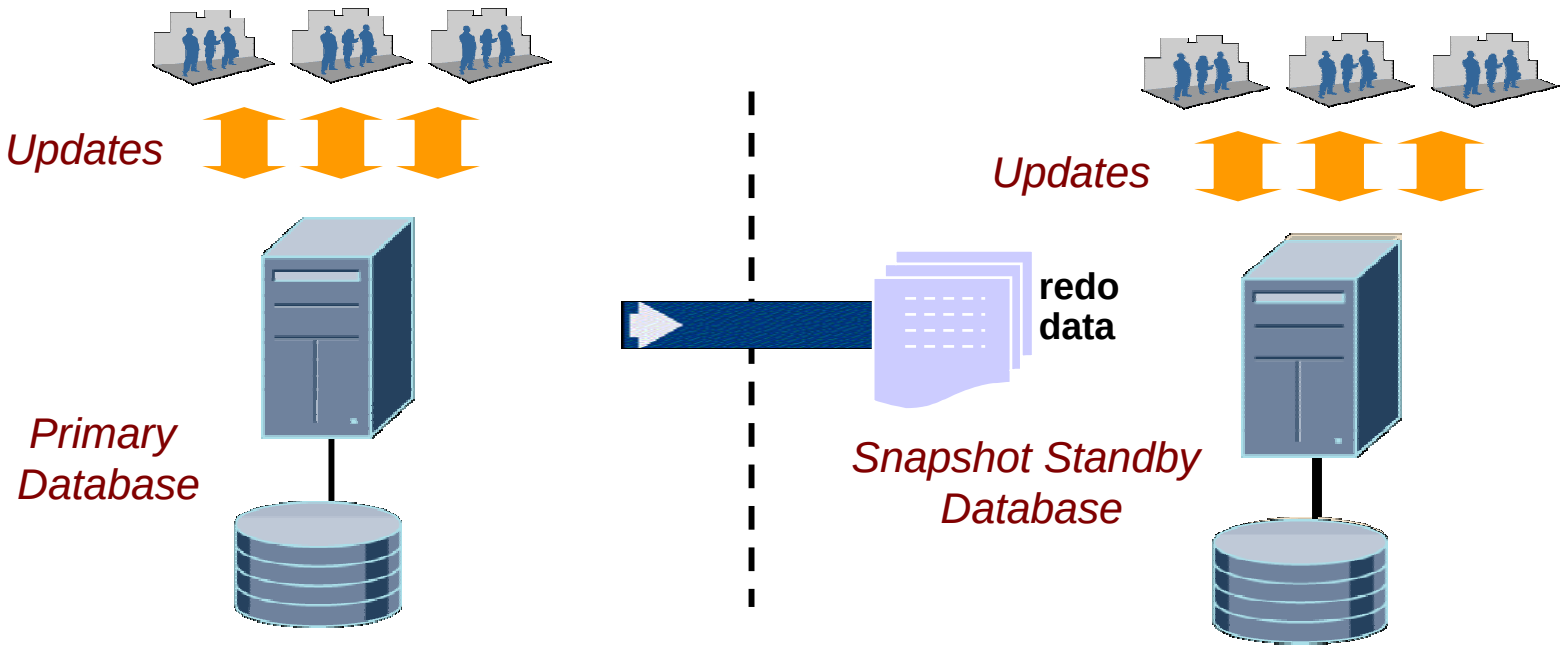


```
DGMGRL> convert database <name> to snapshot standby;
```

Snapshot Standby – Data Guard 11*g*

Convert Back to Synchronized Standby

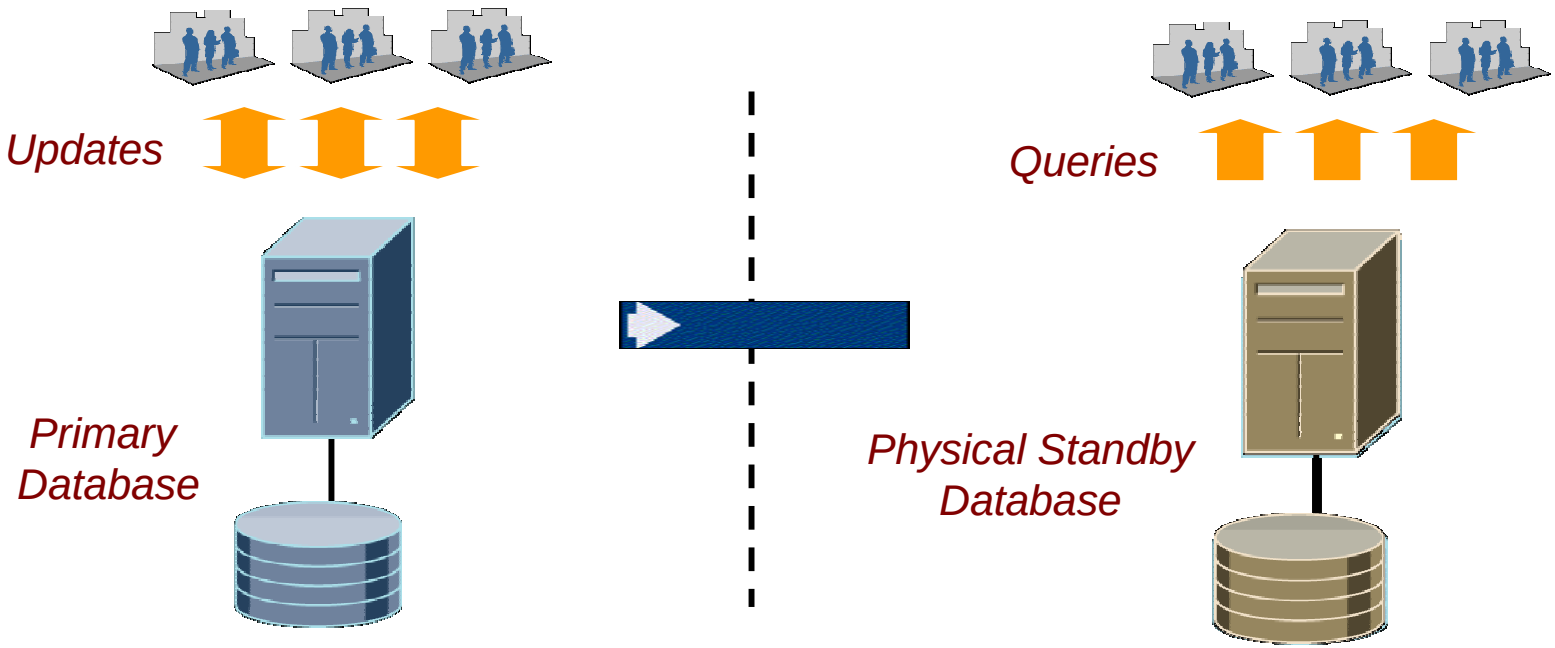


```
DGMGRL> convert database <name> to physical standby;
```

Snapshot Standby 11*g*

Simpler and with better RTO/RPO than Data Guard 10*g*

10.2 – Steps Required

```
Standby
>     alter database recover managed standby database
      cancel;
>     create restore point before_lt guarantee flashback
      database;
Primary
>     alter system archive log current;
>     alter system set log_archive_dest_state_2=defer;
Standby
>     alter database activate standby database;
>     startup mount force;
>     alter database set standby database to maximize
      performance;
>     alter system set log_archive_dest_state_2=defer;
>     alter database open;
```

PERFORM TESTING, ARCHIVE LOGS NOT SHIPPED

```
>     startup mount force;
>     flashback database to restore point before_lt;
>     alter database convert to physical standby;
>     startup mount force;
>     alter database recover managed standby database
      disconnect from session;
Primary
>     Alter system set log_archive_dest_state_2=enable
```

11.1 – Steps Required

```
Standby
>  alter database convert to snapshot standby;
```

PERFORM TESTING, ARCHIVE LOGS CONTINUE TO BE SHIPPED

```
>  alter database convert to physical standby;
```

Program

- Ship Redo / Apply Redo
- Utilize your standby database
  - Active Data Guard
  - Snapshot Standby
- Reduce downtime
- Active Data Guard experiences
  - Real Networks
  - Intermap Inc
- Resources - Q&A

Fast-Start Failover

Reduce Unplanned Downtime

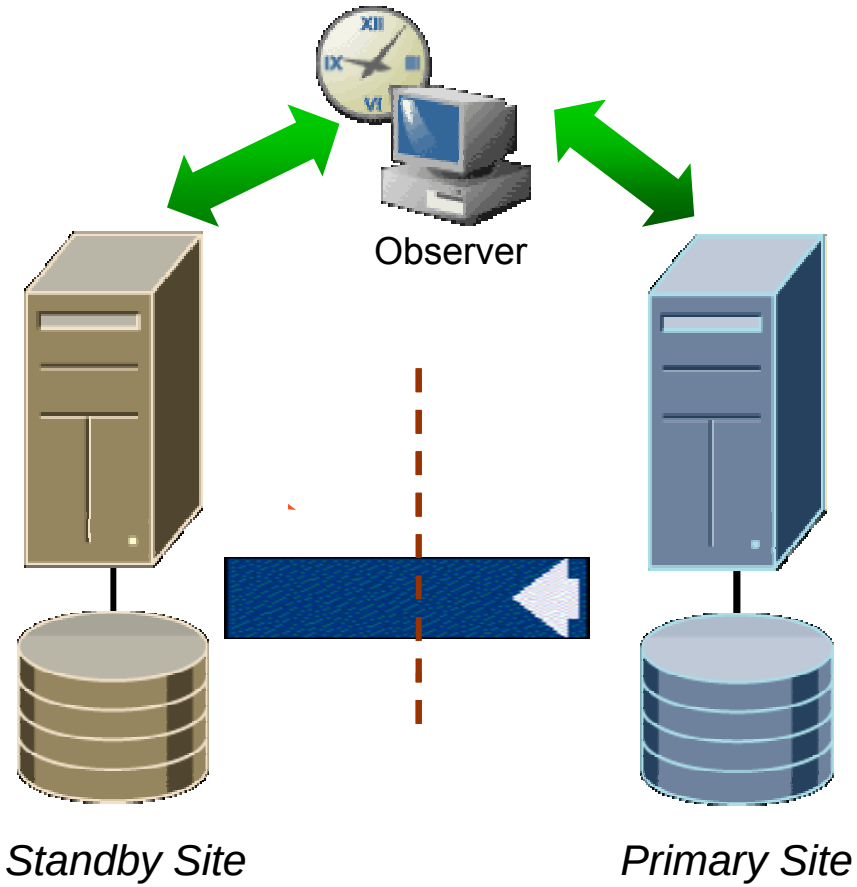

- Maximum Availability
  - SYNC redo transport
  - RPO = zero
- Maximum Performance
  - ASYNC redo transport

`FastStartFailoverLagLimit`

  - Default RPO = 30 seconds
  - minimum threshold = 10 seconds

- Client Failover MAA Best Practices in a Data Guard Configuration
  www.oracle.com/technology/deploy/availability/pdf/MAA_WP_10gR2_ClientFailoverBestPractices.pdf

Fast-Start Failover

Configuration Options for Optimal RPO

- Immediate failover for user-configurable health conditions

```
ENABLE FAST_START FAILOVER [CONDITION <value>];
```

  - Condition examples:
    - Datafile Offline
    - Corrupted Controlfile
    - Corrupted Dictionary
    - Inaccessible Logfile
    - Stuck Archiver
    - Any explicit ORA-xyz error

- Apps can request fast-start failover using API

```
DBMS_DG.INITIATE_FS_FAILOVER
```

Rolling Database Upgrades

Use Physical Standby to Reduce Planned Downtime

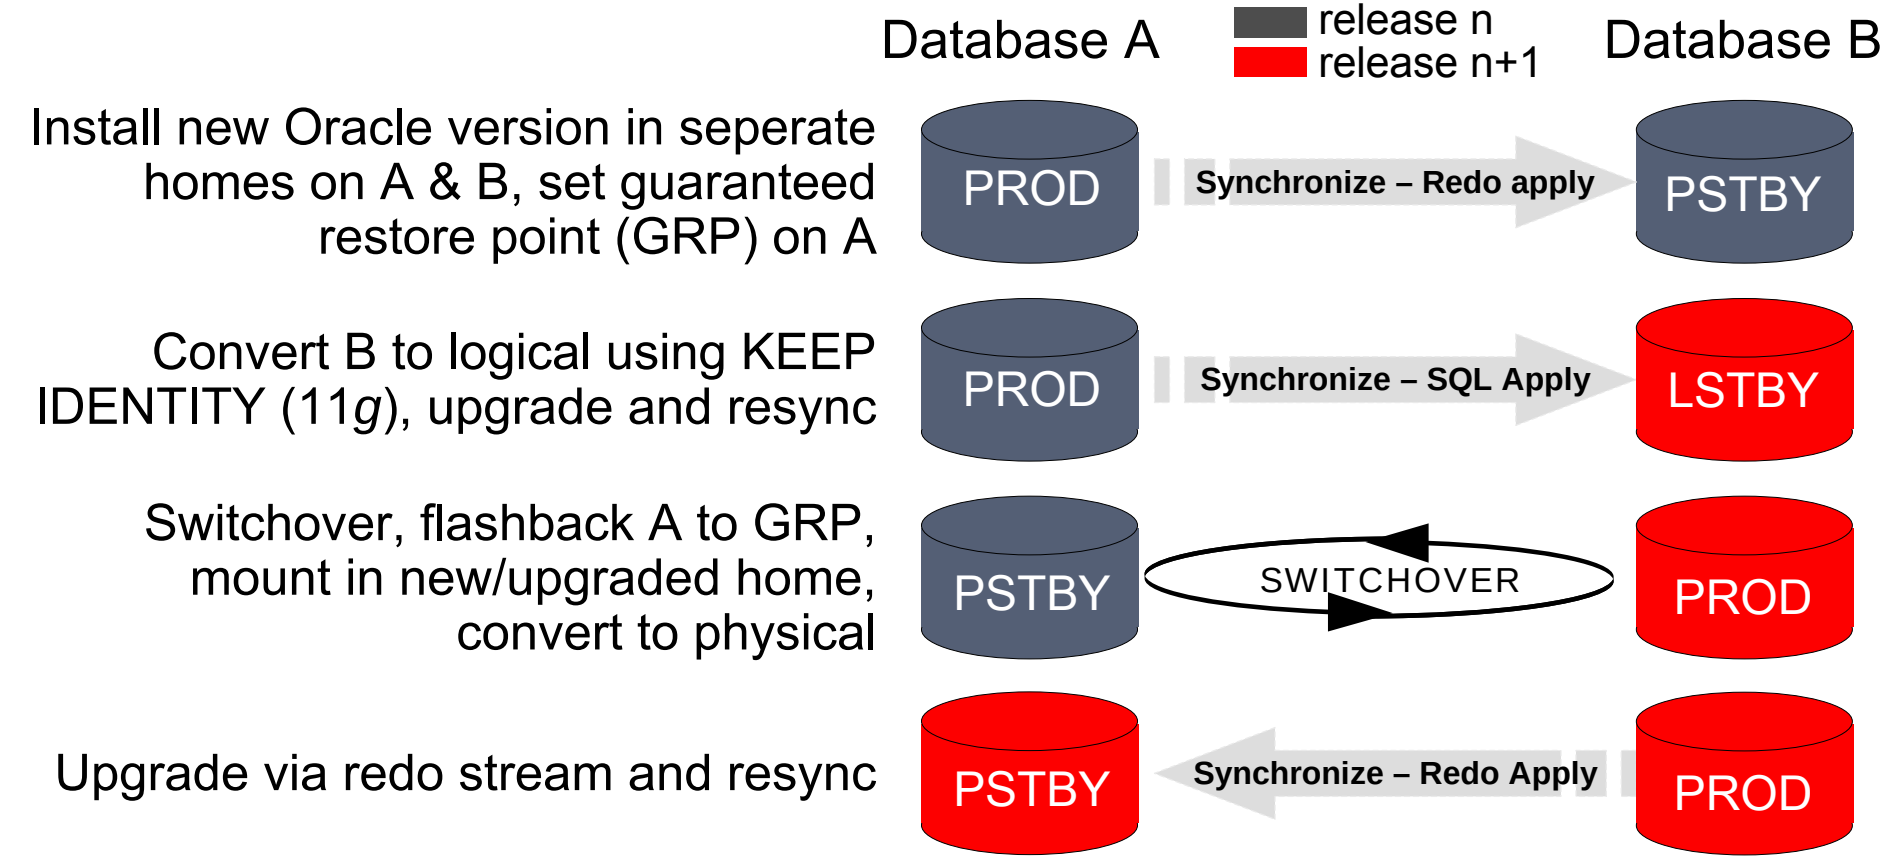

- Rolling Upgrade Best Practices using Transient Logical Standby
  http://www.oracle.com/technology/deploy/availability/pdf/maa_wp_11g_transientlogicalrollingupgrade.pdf

Data Guard Switchover

Reduce downtime for other planned events

- Scheduled power outages and site maintenance
- Data center moves
- Migrations to ASM and/or RAC
- Technology refresh – servers and storage
- Windows/Linux migrations *
- 32bit/64bit migrations*
- HP-UX/PA RISC to HP-UX/IPF migrations*
- Implement major database changes in rolling fashion
  - e.g. ASSM, initrans, blocksize

* see Metalink Note 413484.1

Program

- Ship Redo / Apply Redo
- Utilize your standby database
  - Active Data Guard
  - Snapshot Standby
- Reduce downtime
- Active Data Guard Experiences
  - Real Networks
  - Intermap Inc
- Resources - Q&A

11g Active Data Guard

Yucheng Liu – Sr. Oracle DBA, RealNetworks yliu@real.com
Krishna Kakatur – Sr. Oracle DBA, RealNetworks kkakatur@real.com

What is so Cool About Real Networks

realArcade

realDVD

real RBN
REAL BROADCAST NETWORK

real MTV MUSIC TELEVISION

BIG BROTHER 10
BIG BROTHER LIVE FEED 24/7

film.com

verizon wireless

realPlayer

cingular
raising the bar

Rhapsody

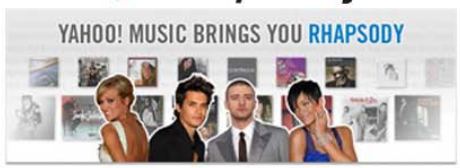

TiVo

TELMEX

RollingStone

Goals

- Consolidate HA and Reporting Data Guard instances
- Off load read-only traffic
- Off load I/O intensive database backups
- Guarantee data and object consistency
- Reduce management of complex mview replication
- Simplify deployment of read-only instances

Current Architecture

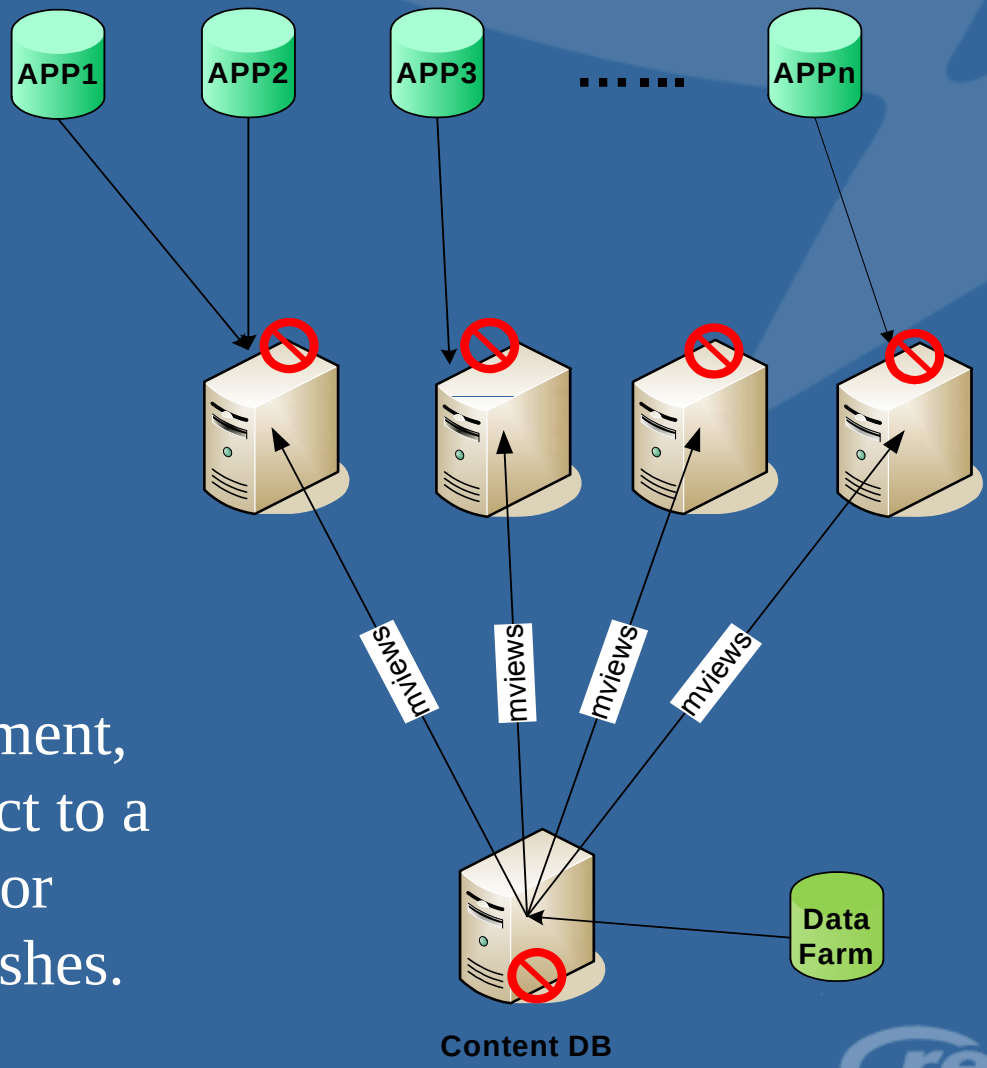

During a code deployment, each database is subject to a lengthy maintenance for DDL and mview refreshes.

Current Architecture Challenges

- Lengthy and complex deployment for DDL changes
  - Each Read-only instance needs ddl scripts + mview refreshes
- Publishing stops if a read only database is down
- No guarantees of data and object consistency
  - Errors can occur, missing indexes, failed mview replication
  - Harder to manage five database images
- Difficult to deploy new nodes for scaling
  - Publishing code must be rewritten for each new node

Future Architecture

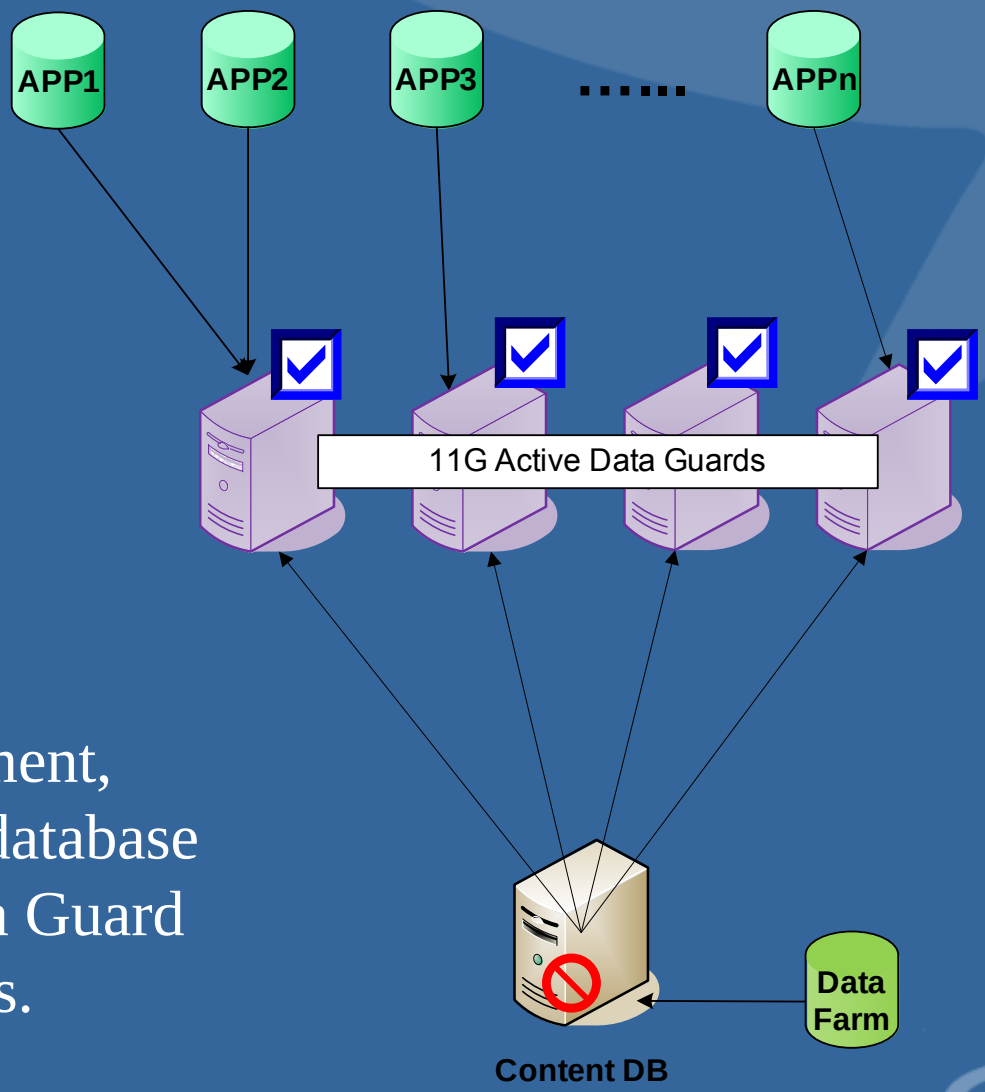

During a code deployment, we can release to one database and rely on active Data Guard to replicate the changes.

Future Architecture Benefits

- Shorter and simpler deployment for DDL changes
  - Deploy changes to Primary instance only
  - No need to rebuild mview logs after a lengthy replication stop
- Guaranteed data and object consistency
- Fewer distinct databases to manage
- New nodes can be deployed without publication outages
- Significant reduction of I/O
  - Backup one instance instead of 4
  - No mview log maintenance
  - No complete or fast refreshes

Active Data Guard Implementation

Proof of Concept Test

- 1 primary + 5 DGs. Primary in MaxAvailability mode. 4 DGs with SYNC redo transport mode, 1 DG with ASYNC mode. 2 DGs with connect-time failover setup.
- Read only queries worked as expected on active DGs.
- Changes populated from primary to active DG as expected.
- Flashback primary + flashback DG tested. No need to recreate DGs.
- Connect-time failover worked correctly on active DGs. (Wow!)

Lessons Learned

Notes

- No downtime on primary. Manual flashback on needed on DG.
- Weapons you need: spfile, service management, Data Guard Broker, Flashback, Grid Control, RMAN
- Switchover bug (#7032374): dbms_service.stop_service does not stop services on standby
- Grid bug (#6379706): Grid does not use RMAN catalog for standby backups
- Export dump (expdp) error: ORA-16000: database open for read-only access

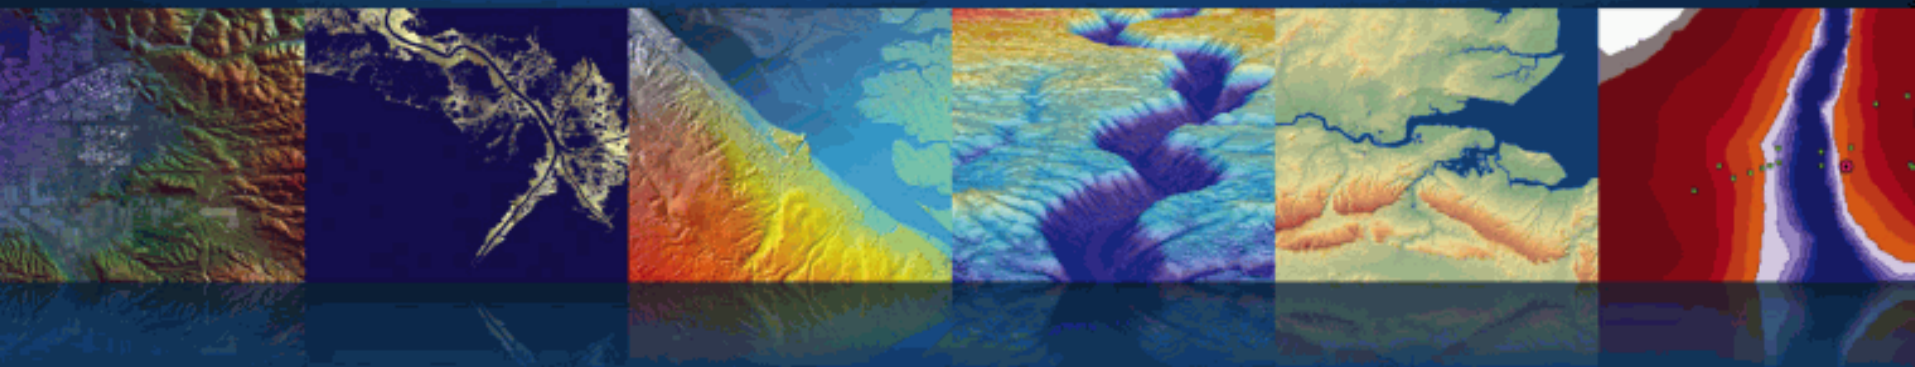

Oracle 11g Active Data Guard

High Availability, Disaster Recovery & Resource Offloading

Presented By: Shawn Ormond, Database Administrator

Who is Intermap Technologies Inc.?

- Intermap is a digital mapping company that is proactively remapping entire countries across the world and building uniform high-resolution 3D digital national data sets which we call NEXTMap®.
- Intermap uses proprietary airborne Interferometric Synthetic Aperture Radar (IFSAR) to collect raw elevation data.
- Intermap produces elevation data models and geometric images of unprecedented accuracy from the IFSAR data.
- These NEXTMap® data sets are used in various commercial and government spatial applications within a many industries:
  - Automotive Safety & Fuel Efficiency
  - Insurance Flood Modeling
  - Global Positioning Systems (GPS)
  - Environmental Planning
  - Wind Power Planning
  - Wireless Communication Planning
  - Other 3D Visualization Applications

INTERMAP™ TECHNOLOGIES

Using Oracle 11g & Active Data Guard

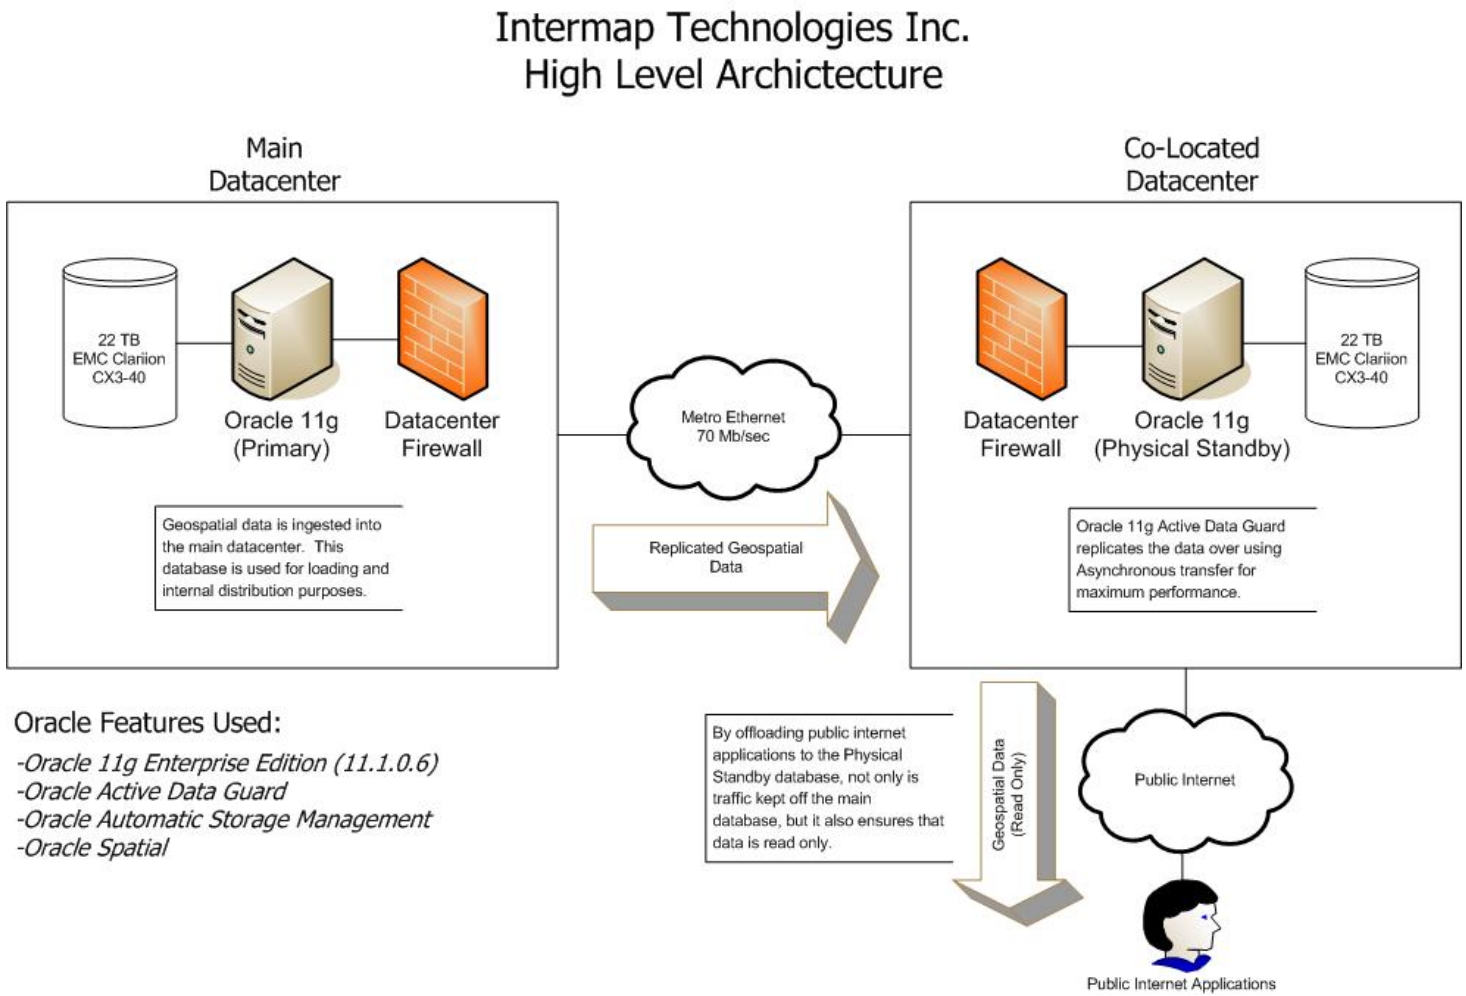

Business Requirements

The Need	The Solution
Disaster recovery site in order to maintain business continuity.	Oracle Active Data Guard (Fast-Start Failover)
Secure data hosting platform for public internet applications.	Oracle Active Data Guard (Read Only Physical Standby)
24x7 availability for customers to retrieve and use Intermap data.	Oracle Active Data Guard (Stays up while applying redo)
Easy storage management to maintain a large database (10 TB and growing).	Oracle Automatic Storage Management (Seamless storage integration for using various storage vendors)
Manage and maintain spatial datasets within Oracle.	Oracle Spatial

INTERMAP™ TECHNOLOGIES

Lessons Learned

- Configuring LUN sizes larger than 2 TB for use in disk groups is not supported by ASMLib in Oracle 11g (11.1.0.6). This has been fixed in 11.1.0.7.
  - https://metalink.oracle.com/metalink/plsql/showdoc?db=Bug&id=6453944
- Some Oracle GeoRaster procedures do not work on a Standby Database as they do on a Primary Database. Specifically SDO_GEOR.mosaic. Oracle provided Intermap Technologies Inc., with a fix to bypass the using of temporary tables on the Physical Standby when assembling tiles of GeoRaster images and instead assembling the tiled GeoRasters into allocated RAM.

INTERMAP™ TECHNOLOGIES

The Oracle Experience

Oracle Active Data Guard –

Prior to Oracle 11g Active Data Guard there was no out-of-the-box solution to meet the business requirements of a co-located geospatial database that was an exact replica of our main production geospatial database and maintain 24x7 availability.

*Active Data Guard was by far the **easiest** component to set up. Since production implementation Intermap Technologies Inc. has experienced **no** problems and has maintained 100% uptime.*

INTERMAP™ TECHNOLOGIES

Conclusion

Standby on Steroids

Yesterday		Today
Disaster protection only	→	HA/DR and performance protection
Complex schemes required to offload query workload	→	Quick win – simple & fast
Systems & storage dedicated to offload backups	→	Flexible use of resources for multiple purposes
Test environments that are a poor match for production	→	Test environments that are a mirror image of production
Low ROI	→	High ROI

Conclusion


9i Standby


Active
Standby
Database
ORACLE DATABASE 11g
CAT
D11T

Resources

- **Oracle Data Guard on OTN**
  http://www.oracle.com/technology/deploy/availability/htdocs/DataGuardOverview.html
- **Oracle HA Portal on OTN**
  http://www.oracle.com/technology/deploy/availability/
- **Maximum Availability Architecture (MAA) white papers and demonstrations**
  http://www.oracle.com/technology/deploy/availability/htdocs/maa.htm
- **Oracle HA Customer Success Stories on OTN:**
  http://www.oracle.com/technology/deploy/availability/htdocs/HA_CaseStudies.html
- **Taneja Group - New Approaches to Data Protection and DR**
  http://www.oracle.com/technology/deploy/availability/htdocs/analysts/tanejagroupdatabasestorage.pdf
- **Enterprise Strategy Group – Data Protection and Disaster Recovery**
  http://www.oracle.com/technology/deploy/availability/htdocs/analysts/enterprisestrategygroupdataguard.pdf

HA Sessions, Labs, Demos From Oracle Development

Mon, Sep 22

- 2:30 pm - Database 11*g*: Next-Gen HA, Moscone South 103

Tue, Sep 23

- 9:00 am - Active-Active Data Centers, Moscone South 103
- 11:30 am - Sharding with Oracle, Moscone South 302
- 11:30 am - HA with Oracle VM, Moscone West 3024
- 1:00 pm - Active Data Guard, Moscone South 104

Wed, Sep 24

- 9:00 am - Fusion Middleware Grid HA, Marriott Nob Hill AB
- 11:30 am - RMAN Best Practices, Moscone South 103
- 1:00 pm - Database in the Cloud, Moscone South 305
- 5:00 pm - Data Guard & Real Application Testing, Moscone 102

Wed, Sep 24 (contd.)

- 5:00 pm - EM in Secure MAA, Moscone West 2001
- 5:00 pm - E-Business Suite HA, Moscone West 2002/04

Thu, Sep 25

- 9:00 am - Oracle Secure Backup, Moscone South 102
- 10:30 am - Streams Replication, Moscone South 102
- 12:00 pm - Rolling Database Upgrades, Moscone South 103
- 1:30 pm - Streams Performance, Moscone South 102
- 3:00 pm - Oracle Grid Computing, Moscone South 303
- 3:00 pm - E-Business Suite R12 MAA, Moscone West 2007
- 3:00 pm - Siebel MAA, Moscone South 308
- 3:00 pm - Fusion SOA HA & Scalability, Marriott Salon 14/15

Hands On Labs - Thu, Sep 25

- 10:30 - 11:30 am, 12:00 - 1:00 pm - Active Data Guard, Marriott Golden Gate A3

DEMOgrounds, Mon-Thu

- Active Data Guard, Streams, Oracle Secure Backup, RMAN/Flashback, MAA