
ORACLE®

# Oracle Data Guard 11g Release 2: High Availability to Protect Your Business

Joseph Meeks
Director,
Product Management
Oracle USA

Aris Prassinos
Distinguished Member
of Technical Staff
MorphoTrak, SAFRAN Group

Michael T. Smith
Principal Member of
Technical Staff
Oracle USA

# Program

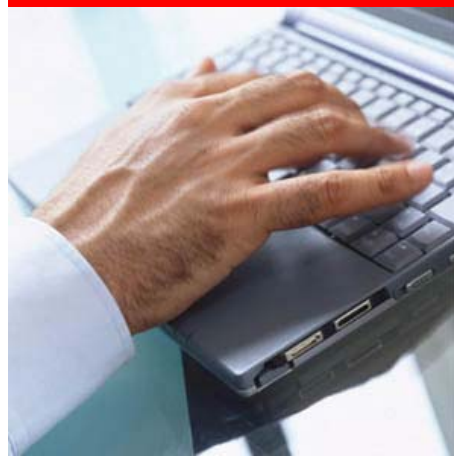

- Traditional approach to HA
- The ultimate HA solution
- Active Data Guard 11.2
- Implementation
- Resources

# Buy Components That Never Fail

# Deploy HA Clusters That Never Fail

## (to compensate for components that fail)

# Hire People That Never Make Mistakes

## (to manage HA clusters that never fail)

# Three Production Examples

## (that never said never)

# Oracle - 90,000 Users

## Beehive Office Applications

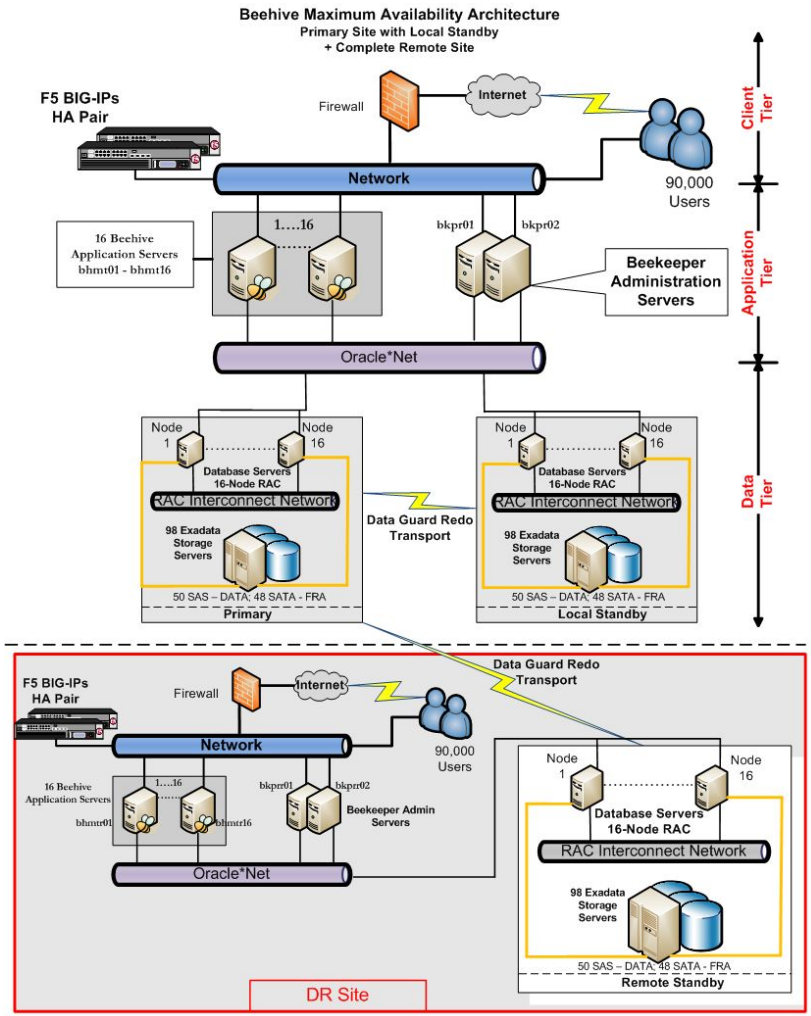


- Beehive – Oracle's unified collaboration solution
  - Email, instant messaging, conferencing, collaboration, calendar…
  - Oracle Database 11.1.0.7
  - 16 node RAC clusters
  - 98 Exadata storage cells / site
  - Data Guard
    - Local standby for HA
      - Offload read-only workload
      - Offload backups
    - Remote standby for DR
      - Dual purpose as test system

# Major Credit Card Issuer

## Website Authentication and Authorization

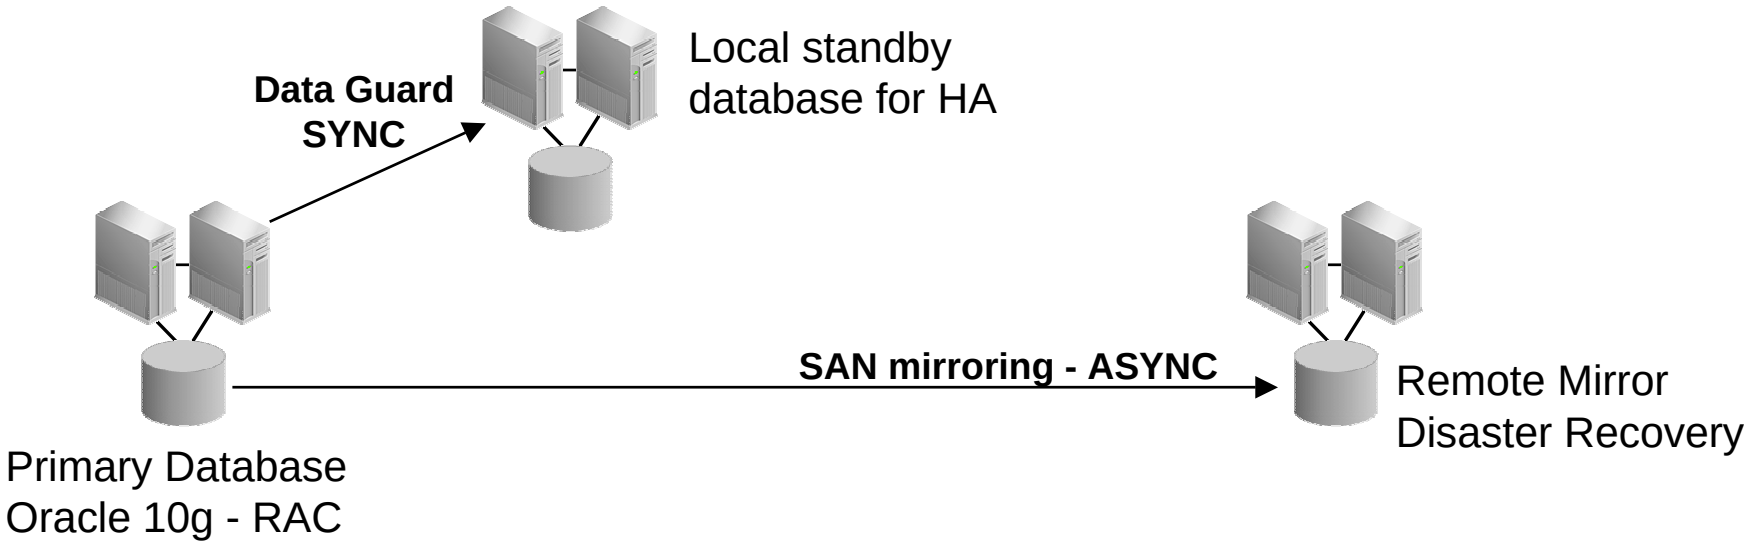


- Single-Sign-On Application
  - Internal and external website authentication and authorization, including web access to personal accounts

# MorphoTrak

## Aris Prassinos - Distinguished Member of Technical Staff

- US subsidiary of Sagem Sécurité, SAFRAN Group
- Innovators in multi-modal Biometric Identification and Verification
  - Fingerprint, palmprint, iris, facial
  - Printrak Biometrics Identification Solution
- Government and Commercial customers
  - Law enforcement, border management, civil identification
  - Secure travel documents, e-passports, drivers' licenses, smart cards
  - Facility / IT access control
- Recently chosen by the FBI as Biometric Provider for their Next Generation Identification Program

  http://www.sagem-securite.com/eng/site.php?spage=04010847

# MorphoTrak

## Printrak Biometrics Identification Solution

- Goal – high availability and disaster recovery at minimal cost

Read-write transactions

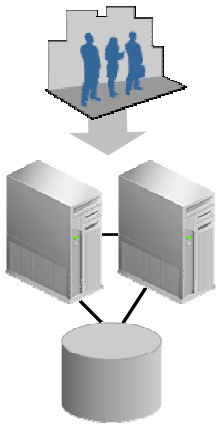

Data Guard Maximum Availability - SYNC

continuous redo shipping, validation and apply
(up to 10ms network latency - approx 60 miles)

Read-only transactions

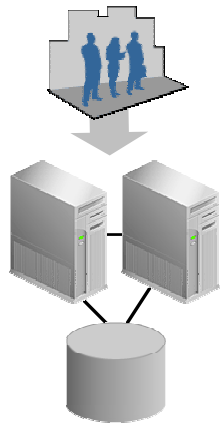

Active Data Guard

- Oracle 11.1.0.7
- Oracle RAC, XML DB, SecureFiles, ASM
- 15TB, 2MB/sec redo rate
- Mixed OLTP – read intensive
- At 10ms network latency, SYNC has 5% - 10% impact on primary throughput

- Automatic database failover (Fast-Start Failover)
  - Complements RAC HA
  - Remote location provides DR
- Off-load read-only transactions to active standby
  - Full utilization reduces acquisition cost
  - Simpler deployment reduces admin cost

# Program

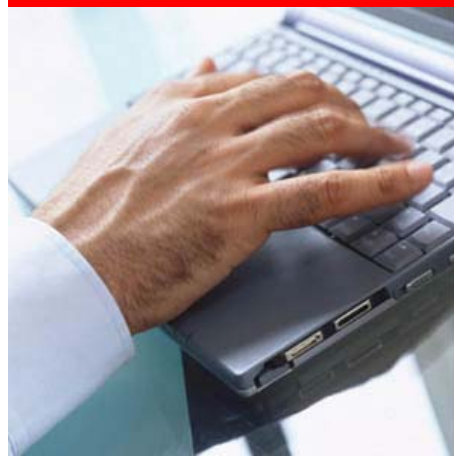

- Traditional approach to HA
- The ultimate HA solution
- Active Data Guard 11.2
- Implementation
- Resources

# High Availability Attributes

| Attribute | Why Important |
|---|---|
| 1. Redundancy with isolation | No single point of failure, failures stay put |
| 2. Zero data loss | Complete protection, no recovery concerns |
| 3. Extreme performance | Deploy for any application |
| 4. Automatic failover | Fast, predictable |
| 5. Full systems utilization | Fast recovery, high return on investment |
| 6. Management simplicity | Reliable, reduced administrative costs |

# Cluster

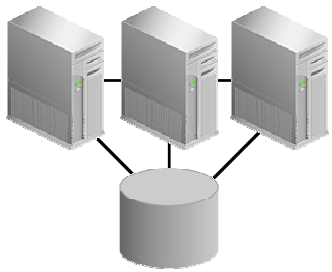

Production
Database

# Cluster

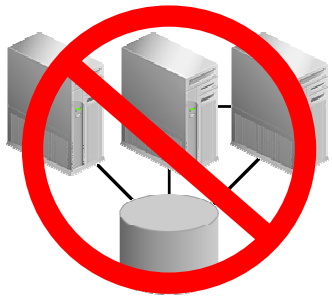

Production
Database

| | |
|---|---|
| ✗ Redundancy with isolation | ✓ Automatic failover |
| ✗ Zero data loss | ✓ Full systems utilization |
| ✓ Extreme performance | ✓ Management simplicity |

# Cluster with Remote DR Site

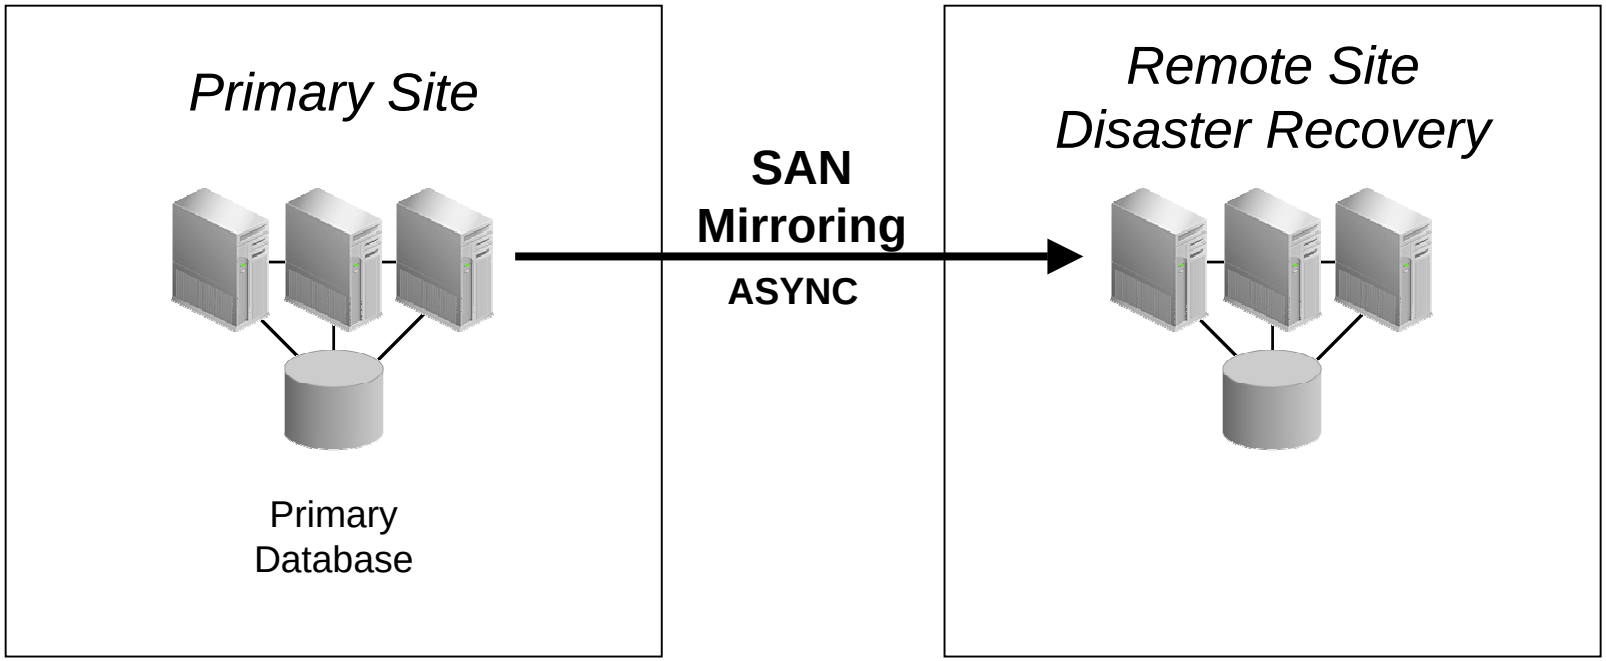

# Cluster with Remote DR Site

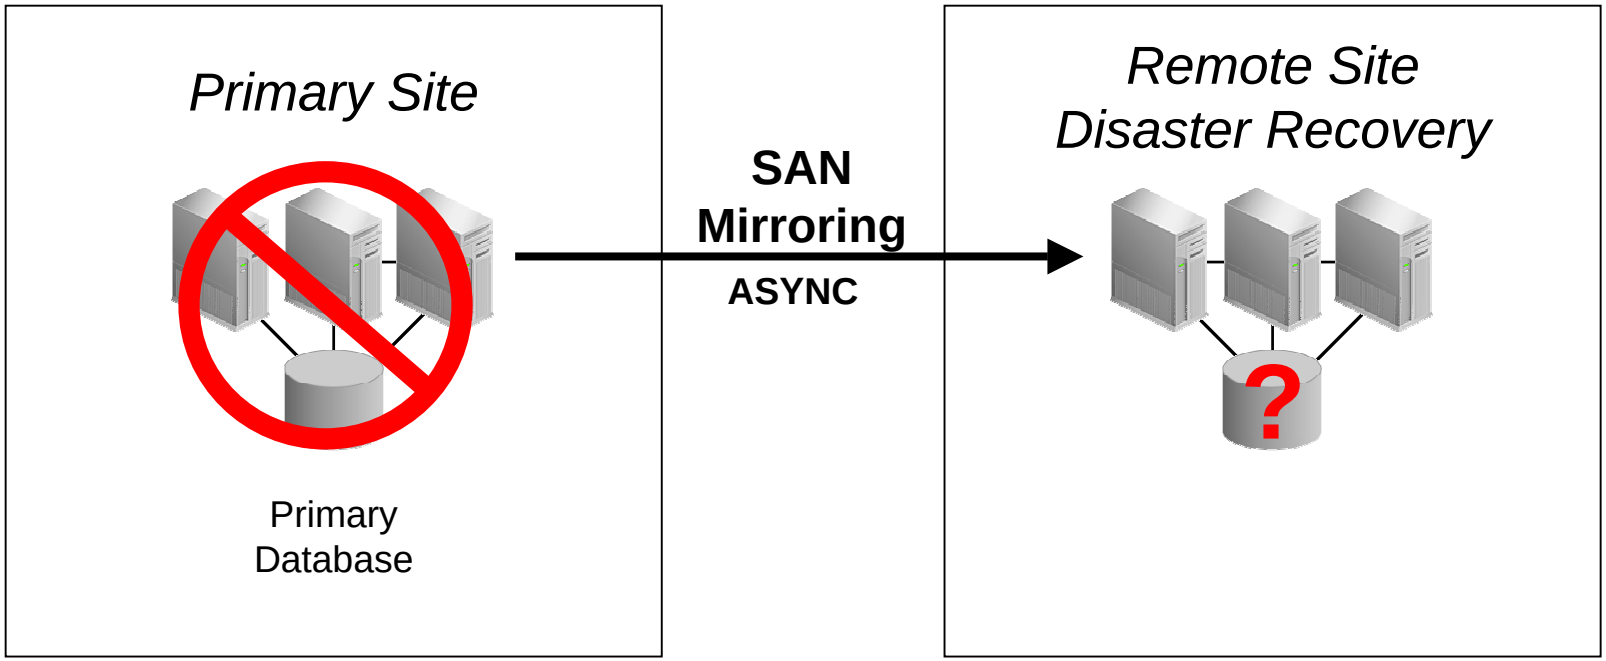


| | |
|---|---|
| ✗ Redundancy with isolation | ✗ Automatic failover |
| ✗ Zero data loss | ✗ Full systems utilization |
| ✓ Extreme performance | ✗ Management simplicity |

# Cluster with Remote DR Site

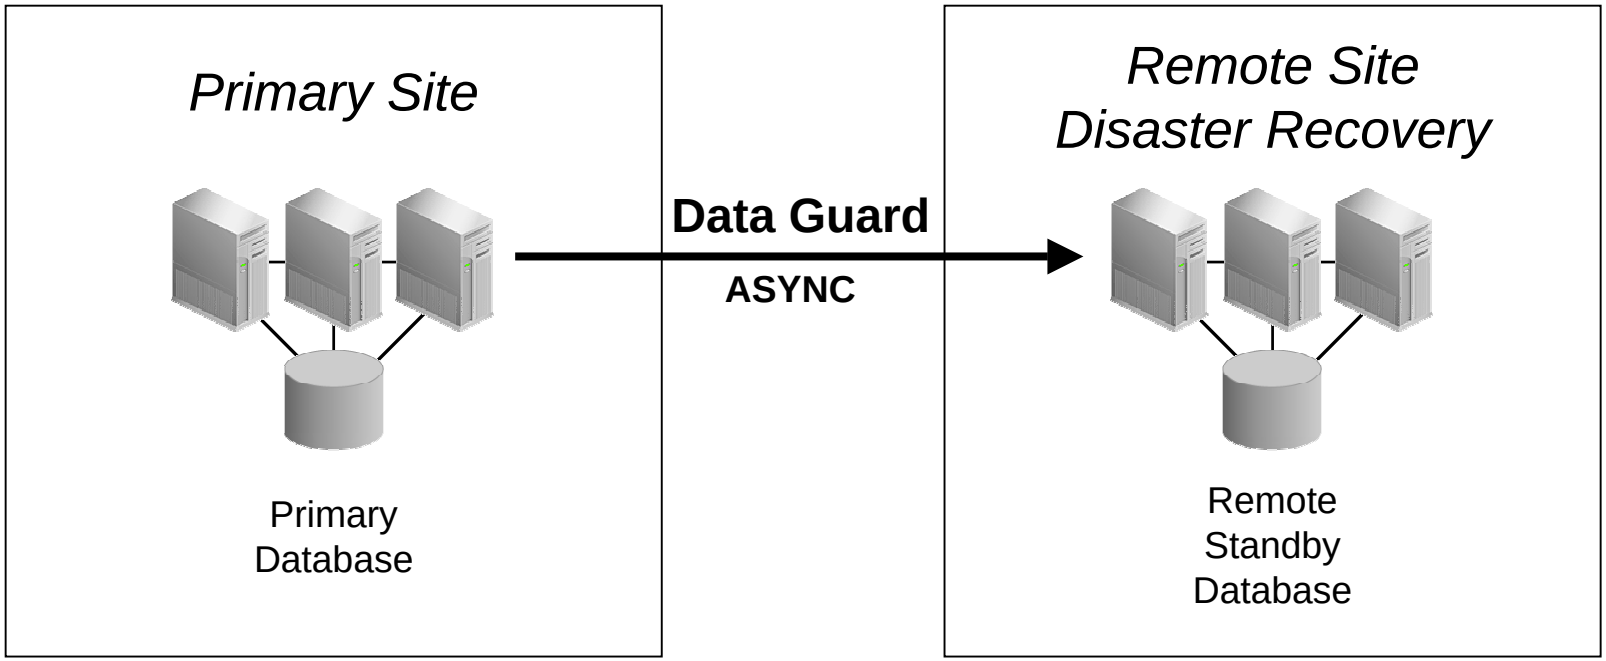

# Cluster with Remote DR Site

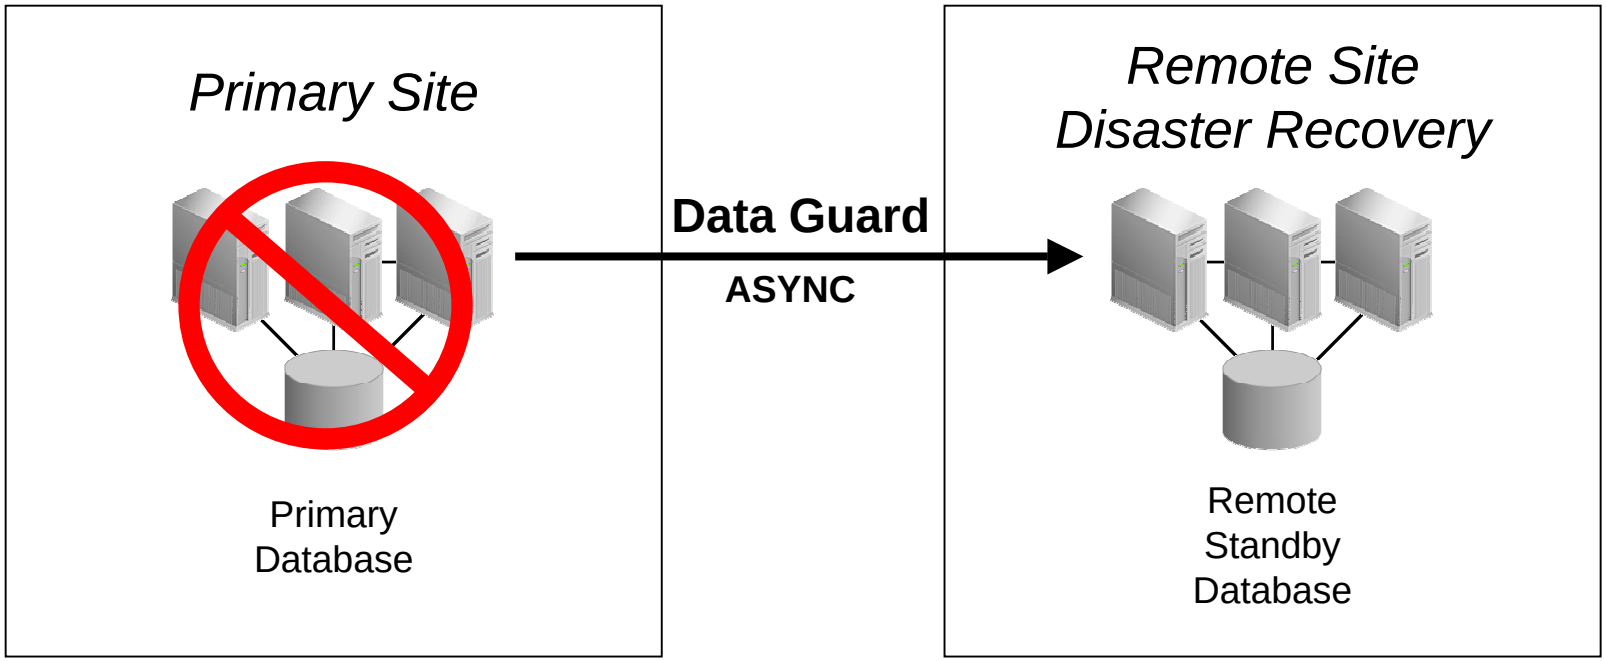


| | |
|---|---|
| ✓ Redundancy with isolation | ✓ Automatic failover |
| ✗ Zero data loss | ✓ Full systems utilization |
| ✓ Extreme performance | ✓ Management simplicity |

# Cluster with Data Guard Local and Remote Standby

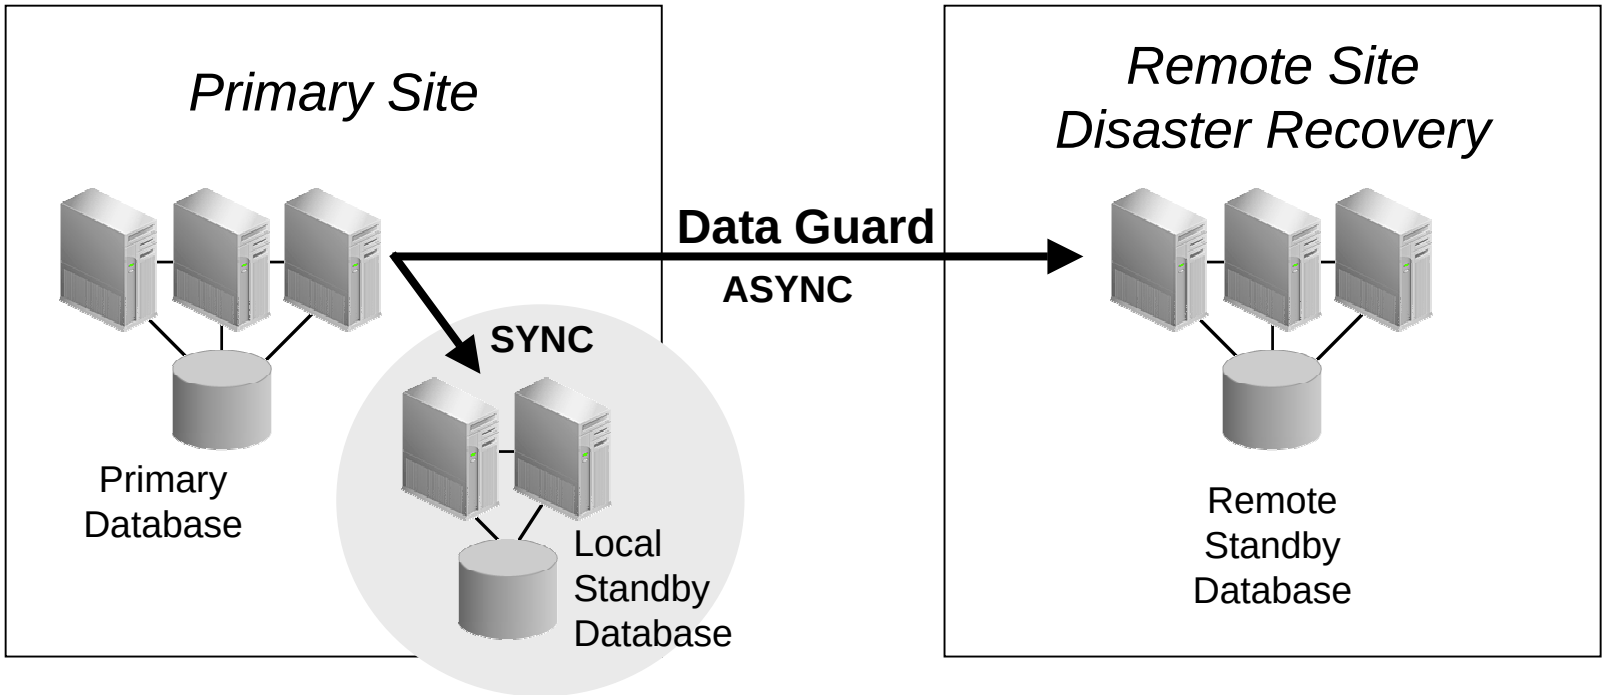

# Cluster with Data Guard Local and Remote Standby

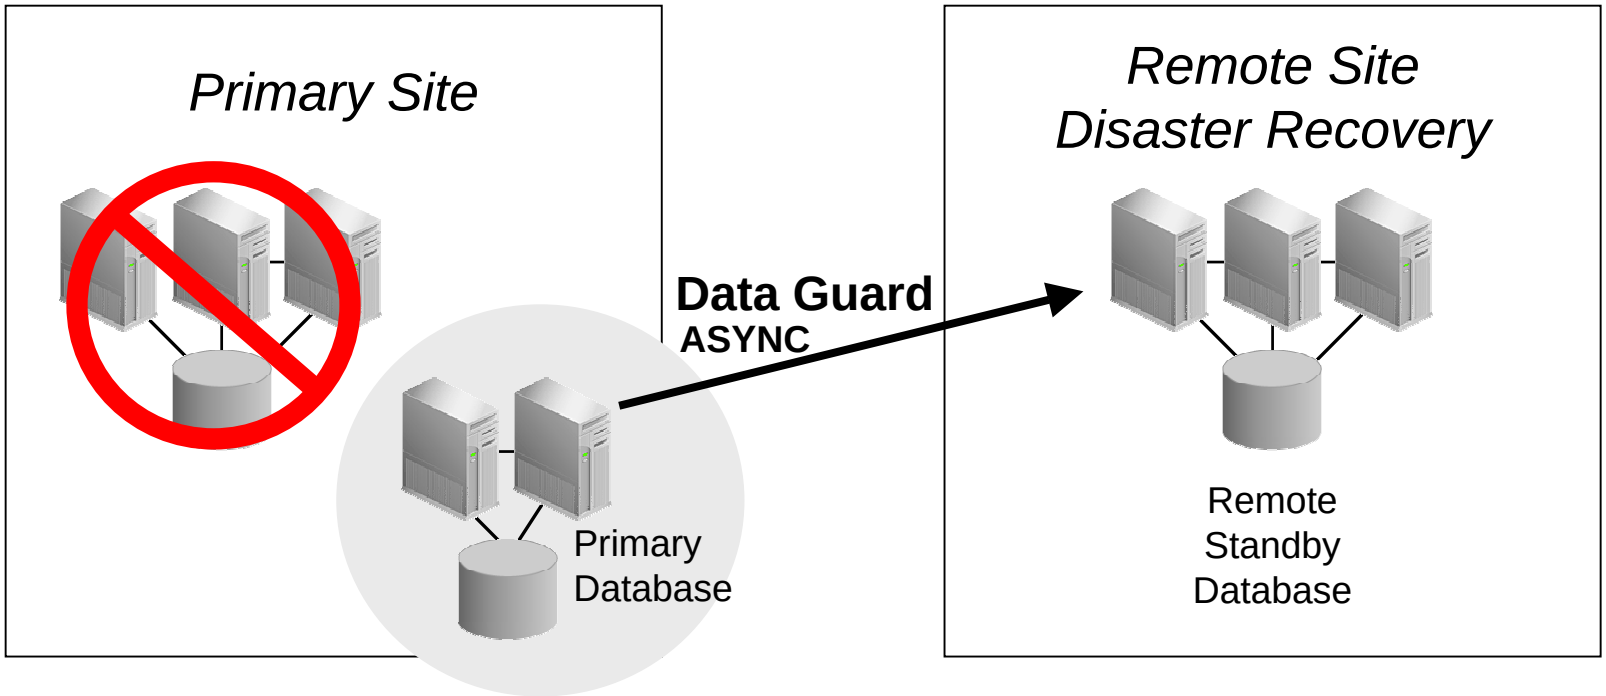


| | |
|---|---|
| ✓ Redundancy with isolation | ✓ Automatic failover |
| ✓ Zero data loss | ✓ Full systems utilization |
| ✓ Extreme performance | ✓ Management simplicity |

# Program

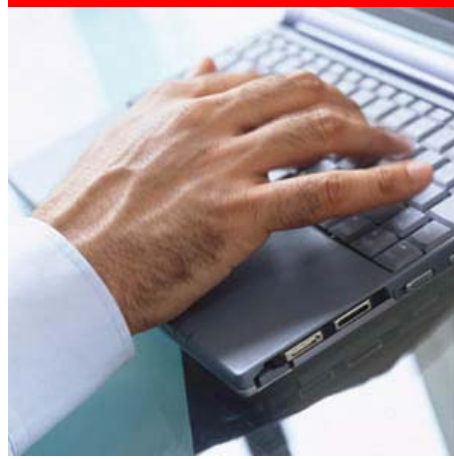

- Traditional approach to HA
- The ultimate HA solution
- Active Data Guard 11.2
- Implementation
- Resources

# What is Active Data Guard?

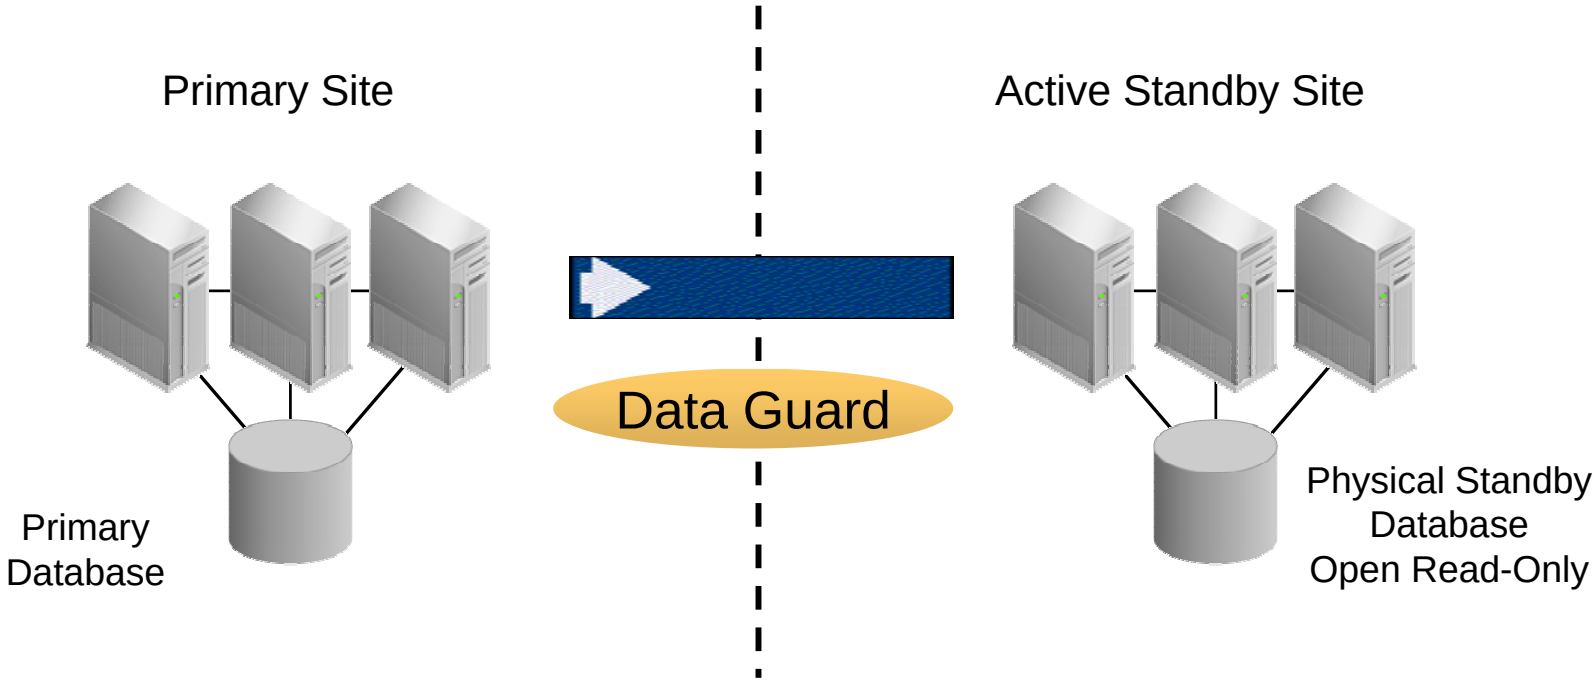


- Data availability and data protection for the Oracle Database
- Up to thirty standby databases in a single configuration
- Physical standby used for queries, reports, test, or backups

# High Availability Attributes

## How Does Active Data Guard Stack Up?

| Attribute | Why Important |
|---|---|
| 1. Redundancy with isolation | No single point of failure, failures stay put |
| 2. Zero data loss | Complete protection, no recovery concerns |
| 3. Extreme performance | Deploy for any application |
| 4. Automatic failover | Fast, predictable |
| 5. Full systems utilization | Fast recovery, high return on investment |
| 6. Management simplicity | Reliable, reduced administrative costs |

# HA Attribute: Redundancy with Isolation

## Data Guard Transport and Apply

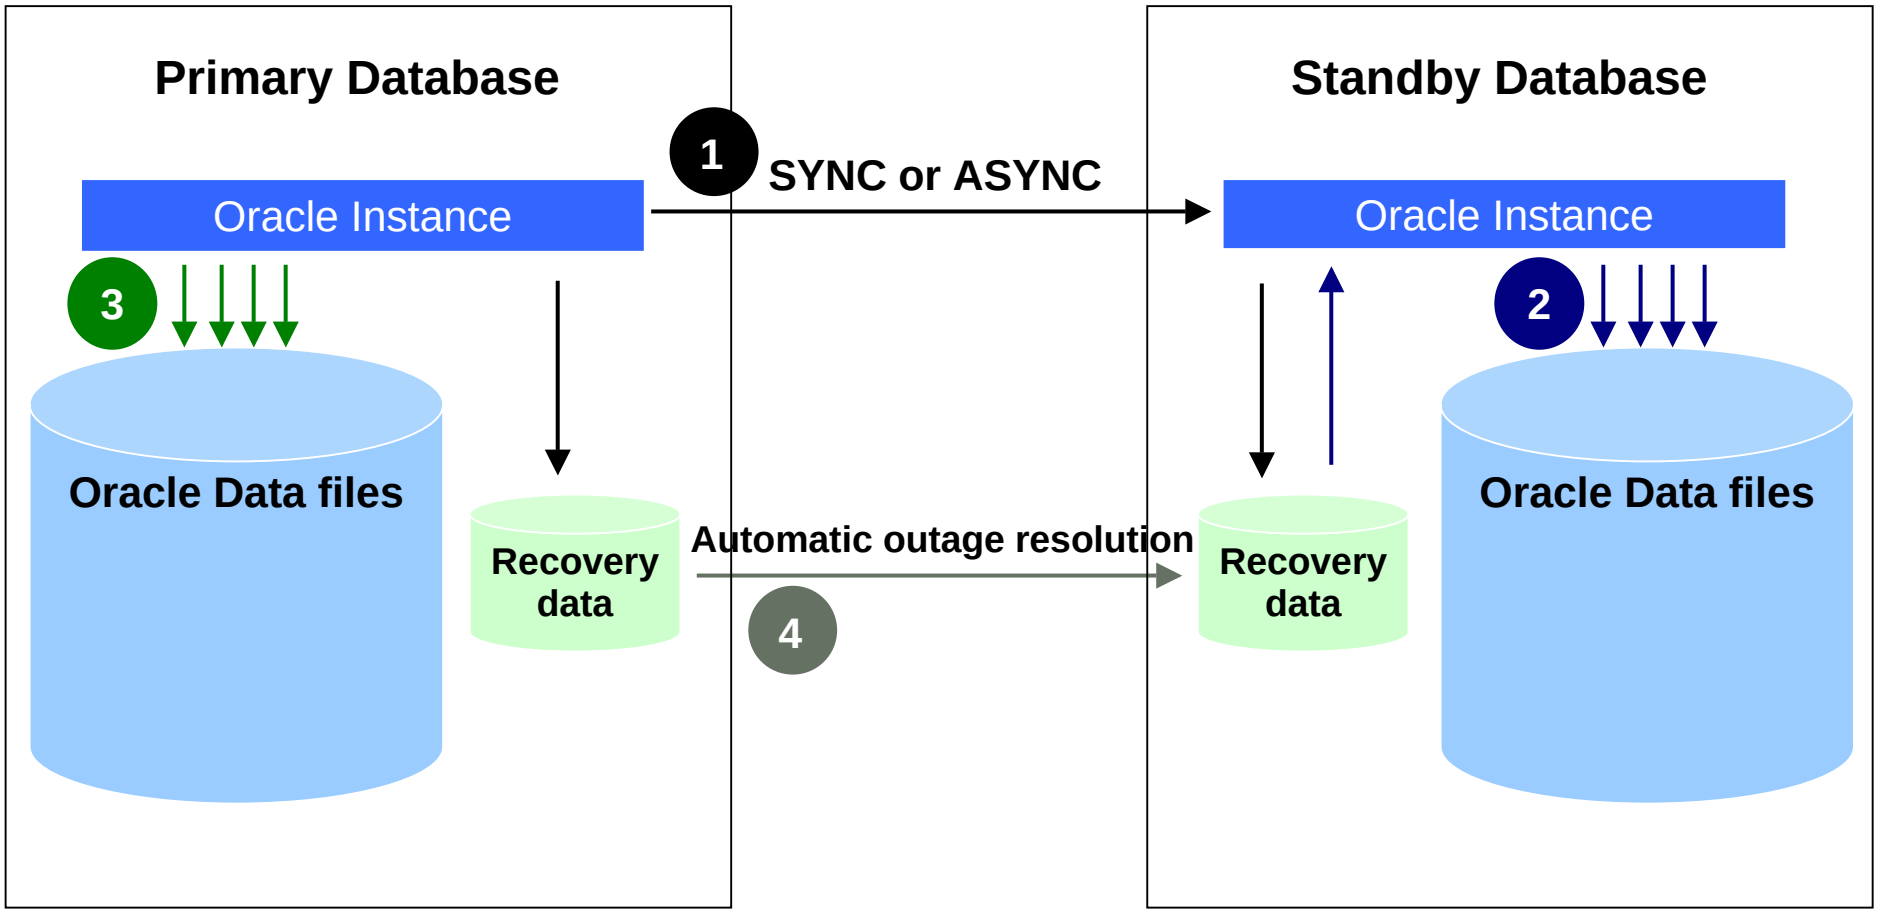

# HA Attribute: Redundancy with Isolation

## Data Integrity

- Primary changes transmitted directly from SGA
  - Isolates standby from I/O corruptions
- Software code path on standby different than primary
  - Isolates standby from firmware and software errors
- Multiple Oracle corruption detection checks
  - Data applied to the standby is logically and physically consistent
- Standby detects silent corruptions that occur at primary
  - Hardware errors and data transfer faults that occur after Oracle receives acknowledgment of write-complete
- Known-state of standby database
  - Oracle is open, ready for failover if needed

# HA Attribute: Zero Data Loss

## Synchronous redo transport

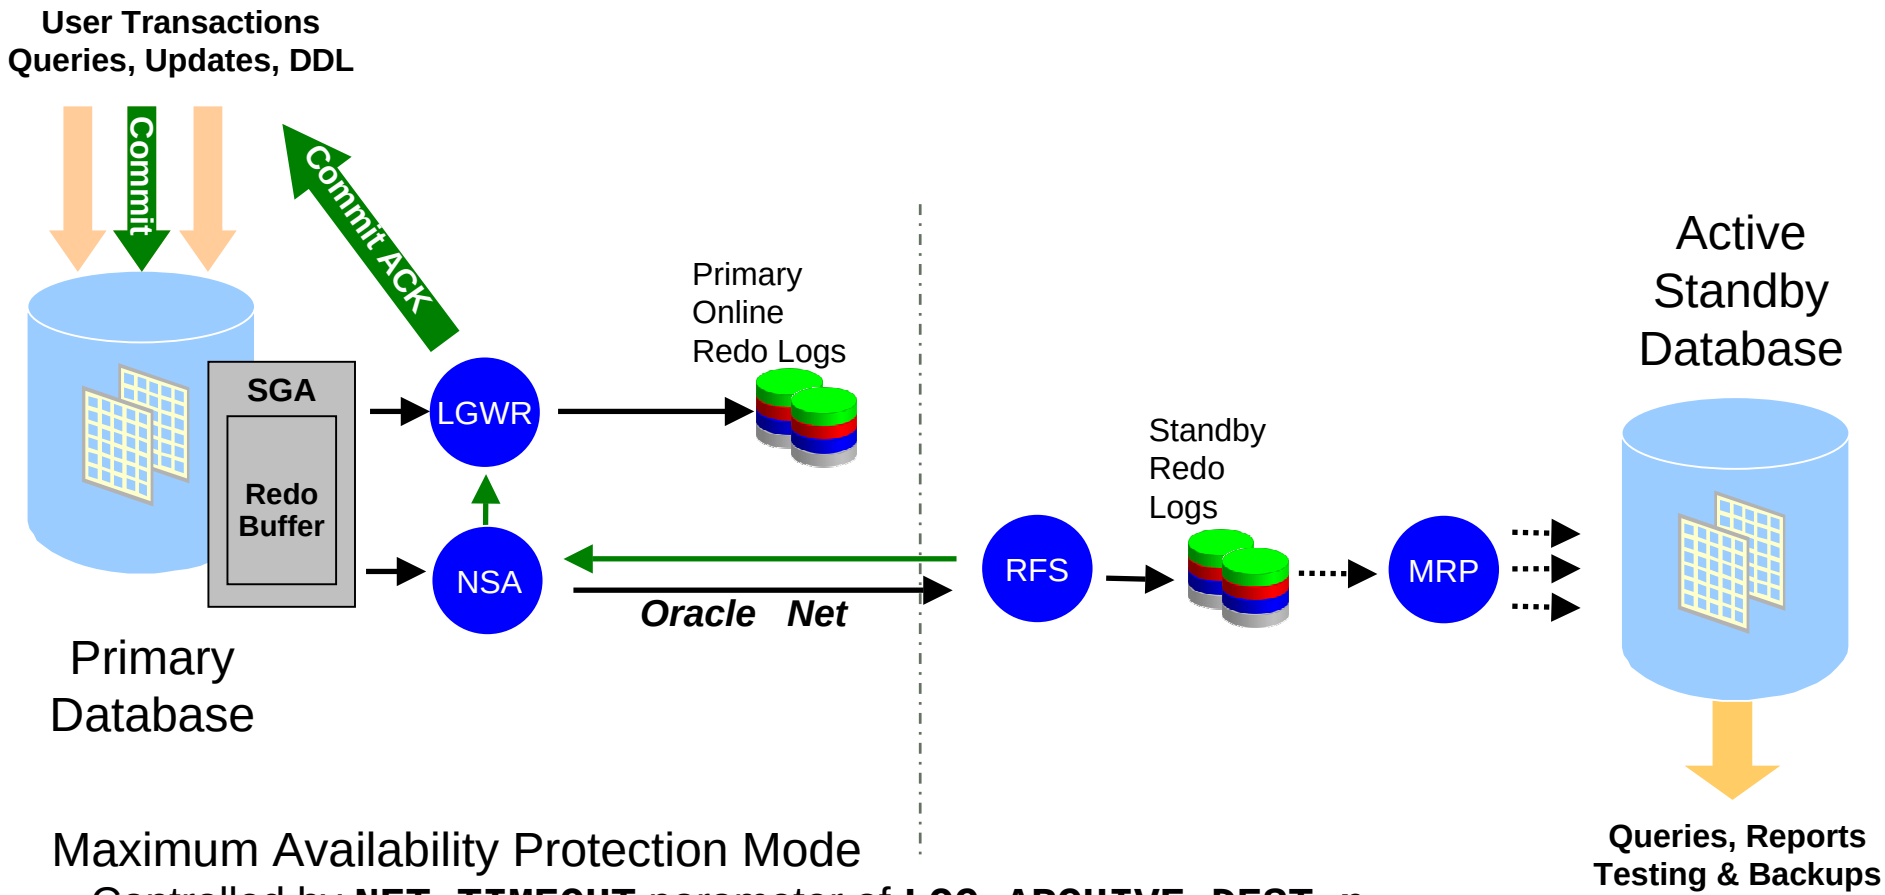


Maximum Availability Protection Mode

- Controlled by **`NET_TIMEOUT`** parameter of **`LOG_ARCHIVE_DEST_n`**
- Default value 30 seconds in Data Guard 11g

# HA Attribute: Automatic Failover

## Database

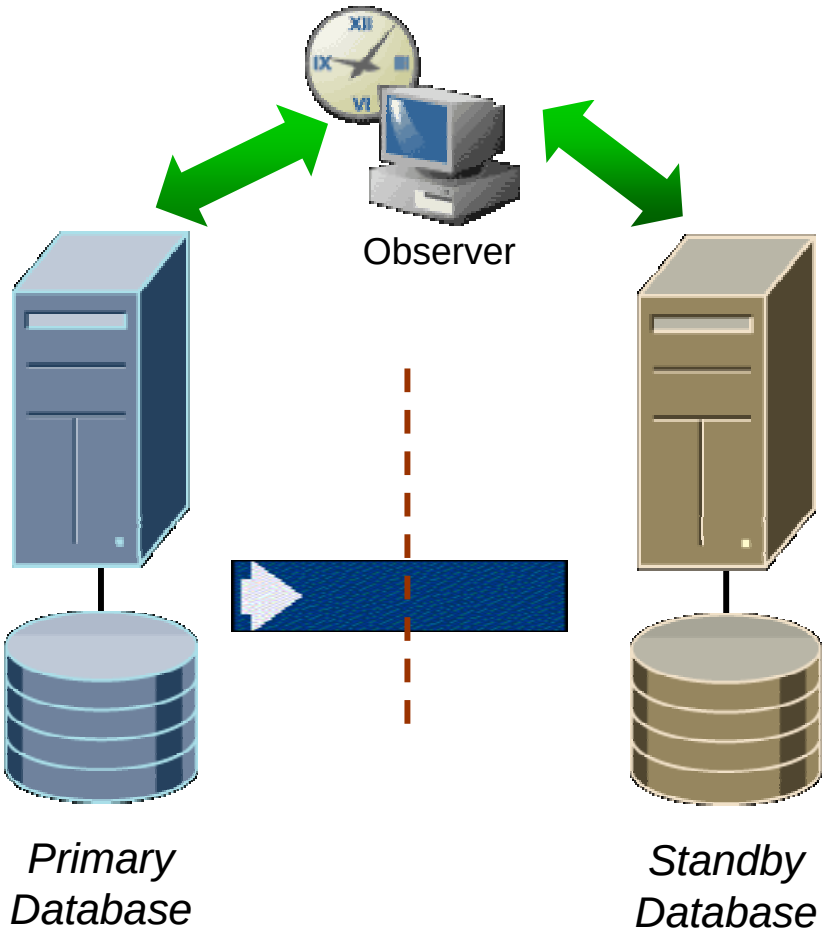


### Data Guard Fast-Start Failover

- Automatic failover
  - Database down
  - Designated health-check conditions
  - Or at request of an application
- Failed primary automatically reinstated as standby database
- All other standby's automatically synchronize with the new primary

# HA Attribute: Automatic Failover

## Database

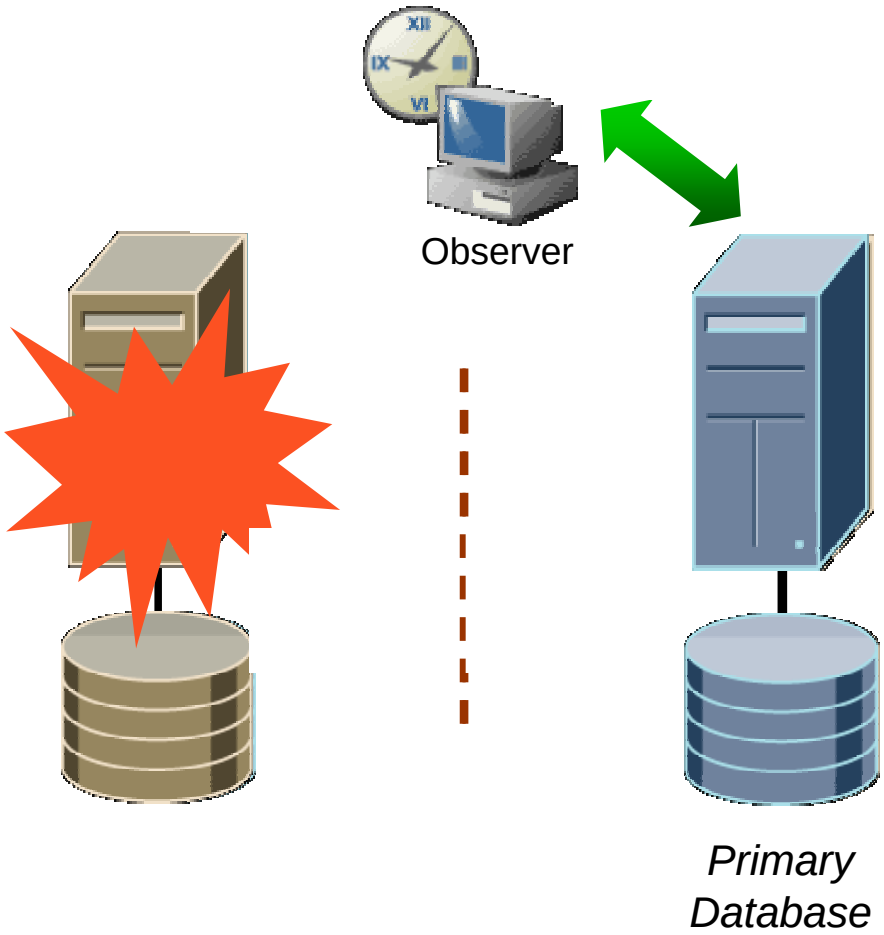


### Data Guard Fast-Start Failover

- Automatic failover
  - Database down
  - Designated health-check conditions
  - Or at request of an application
- Failed primary automatically reinstated as standby database
- All other standby's automatically synchronize with the new primary

# HA Attribute: Automatic Failover

## Database

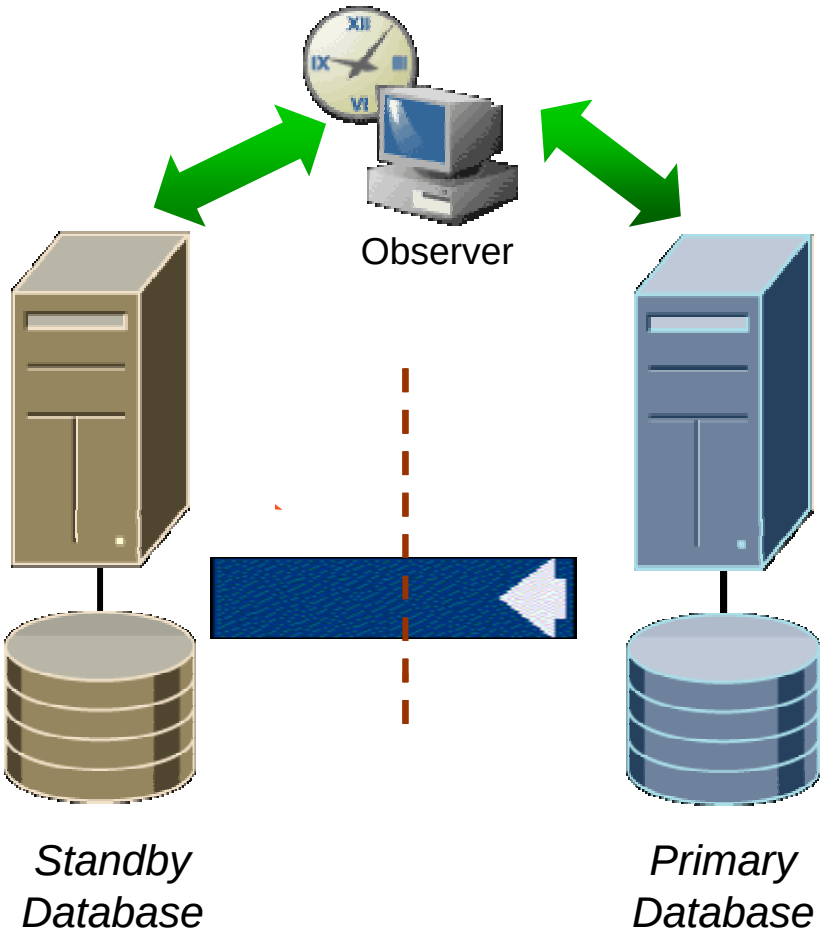


### Data Guard Fast-Start Failover

- Automatic failover
  - Database down
  - Designated health-check conditions
  - Or at request of an application
- Failed primary automatically reinstated as standby database
- All other standby's automatically synchronize with the new primary

# HA Attribute: Automatic Failover

## Applications

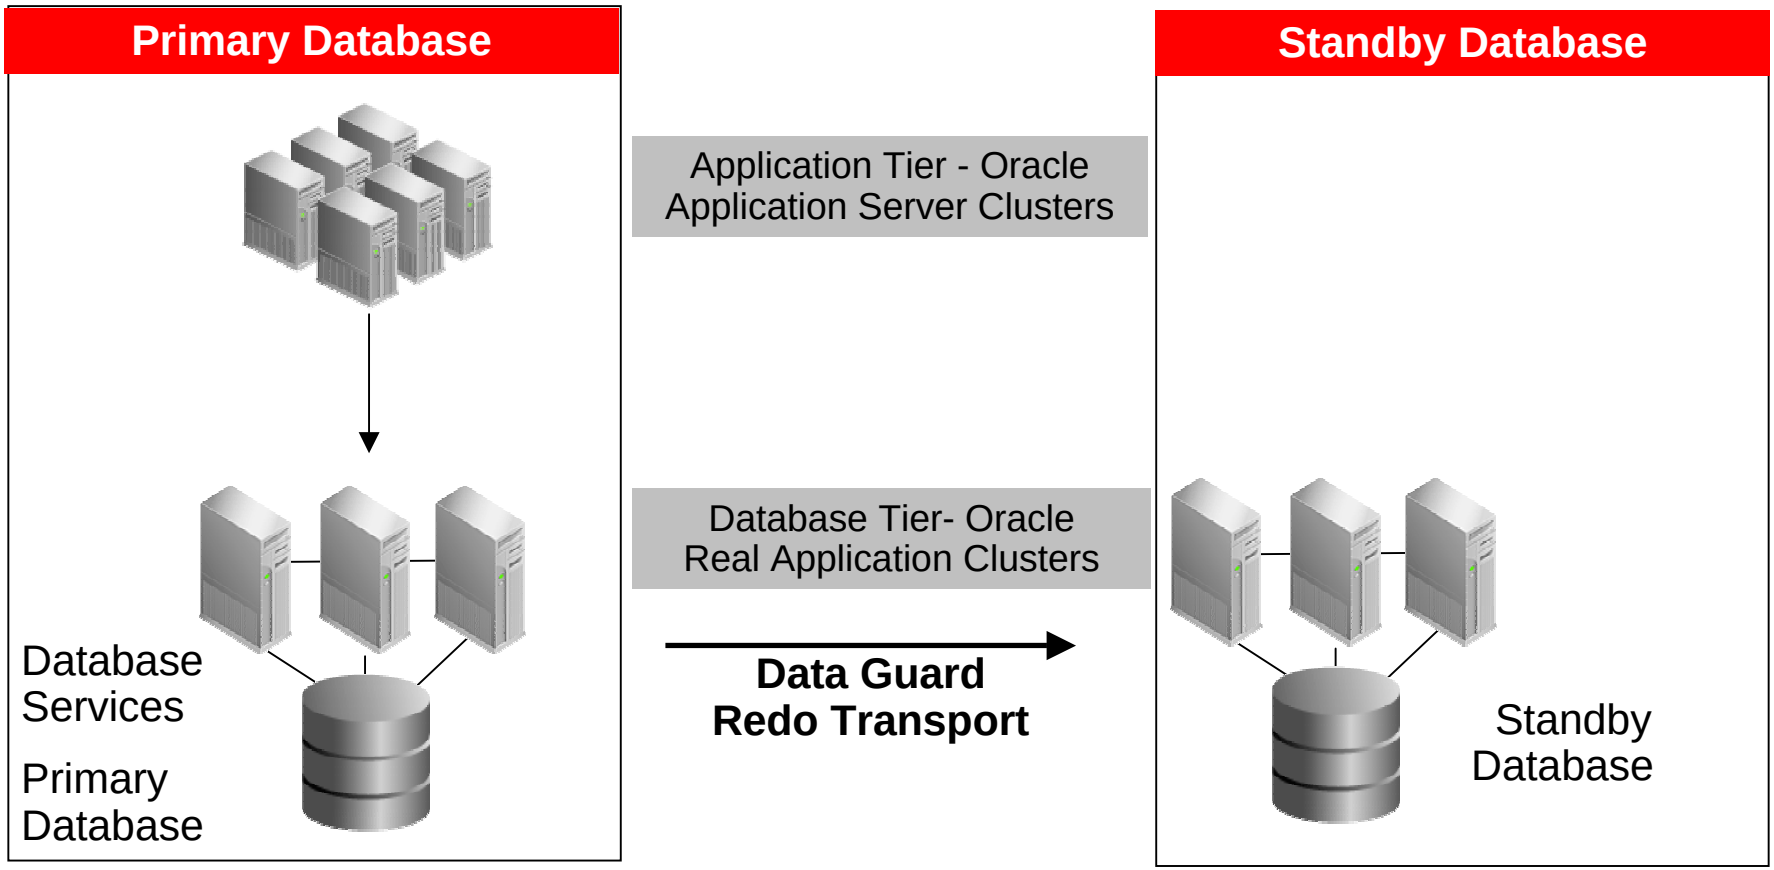

# HA Attribute: Automatic Failover

## Applications

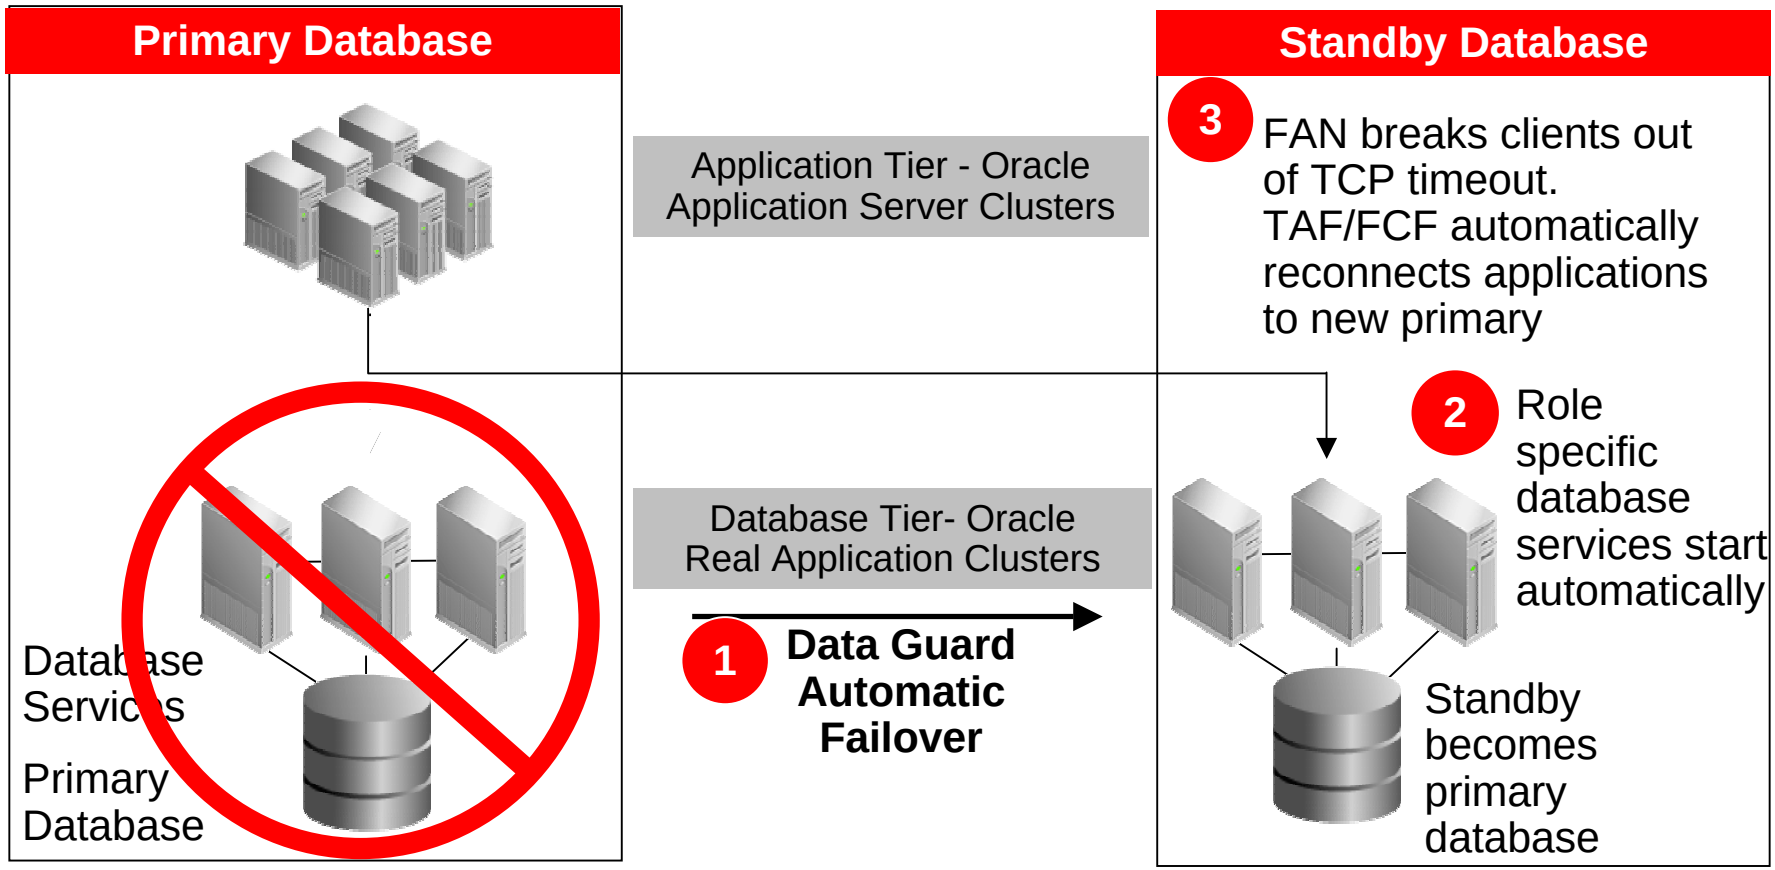

# HA Attribute: Extreme Performance

## Primary Database

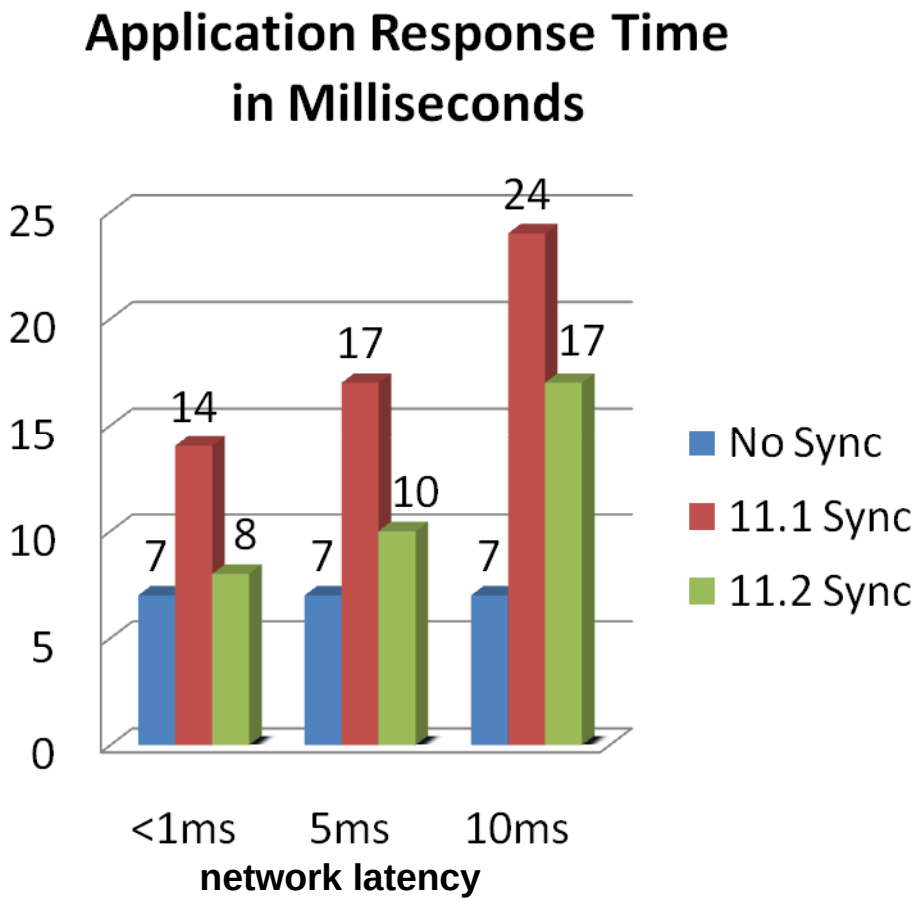


- Data Guard 11.2 SYNC
  - Redo shipped in parallel with LGWR write to local online log file
- Little to no impact on response time when using SYNC in low latency network
  - 40% improvement over 11.1 on low latency LAN

# HA Attribute: Extreme Performance

## Standby Database

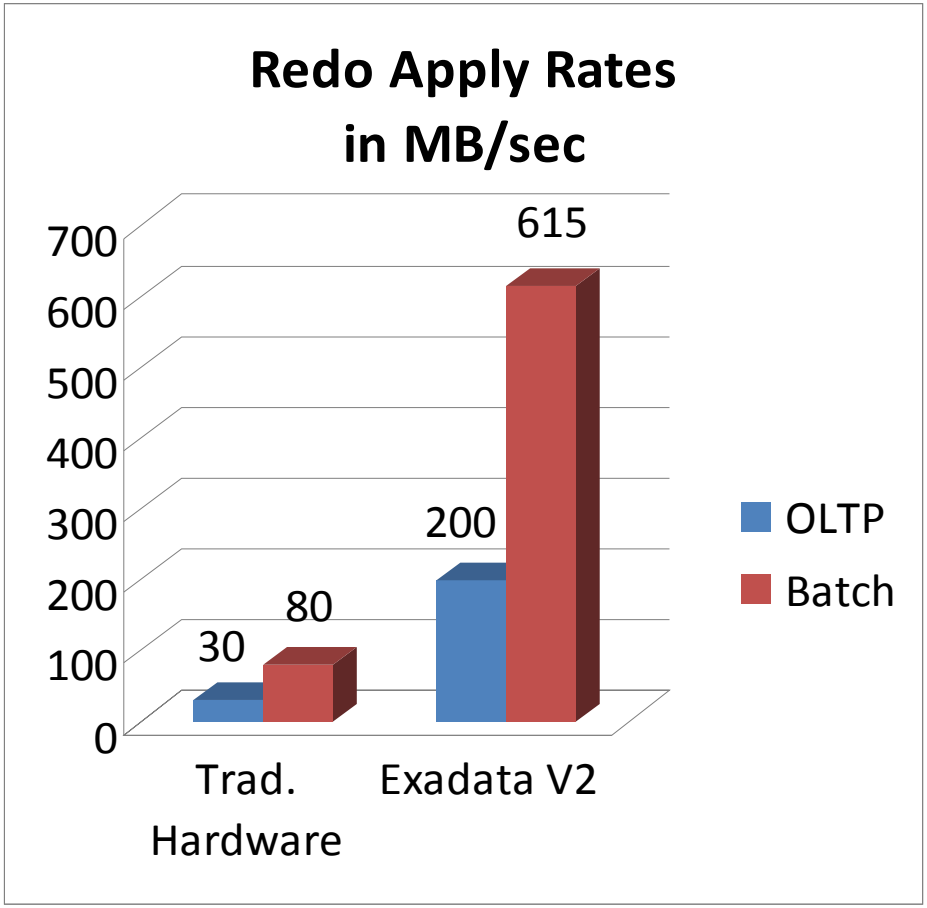


- Data Guard 11.2 Redo Apply
  - Across the board increase in apply rates
  - High query load on active standby does not impact apply
  - Redo Apply is optimized to utilize Exadata I/O bandwidth
  - Improved “Apply Lag” stat allows for finer grained monitoring of standby progress

# HA Attribute: Full Systems Utilization

## Active Data Guard

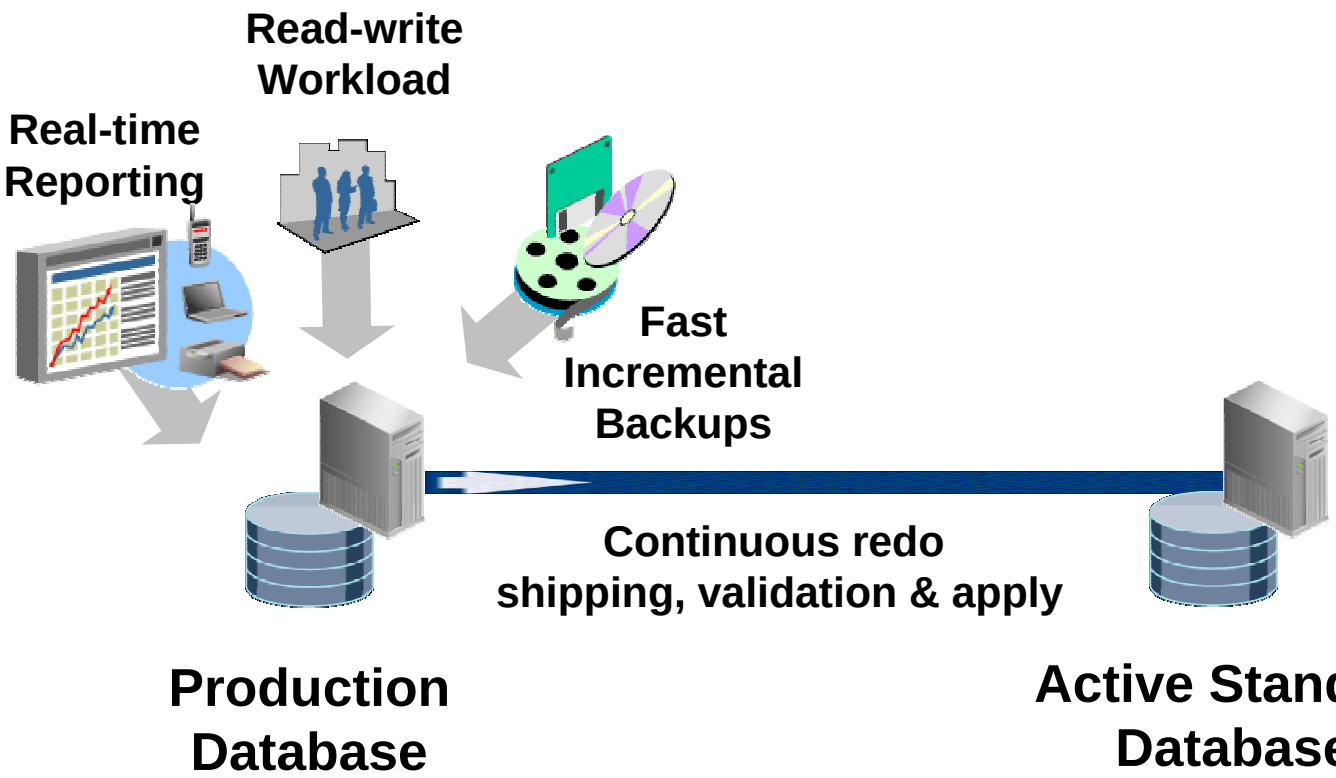

# HA Attribute: Full Systems Utilization

## Active Data Guard

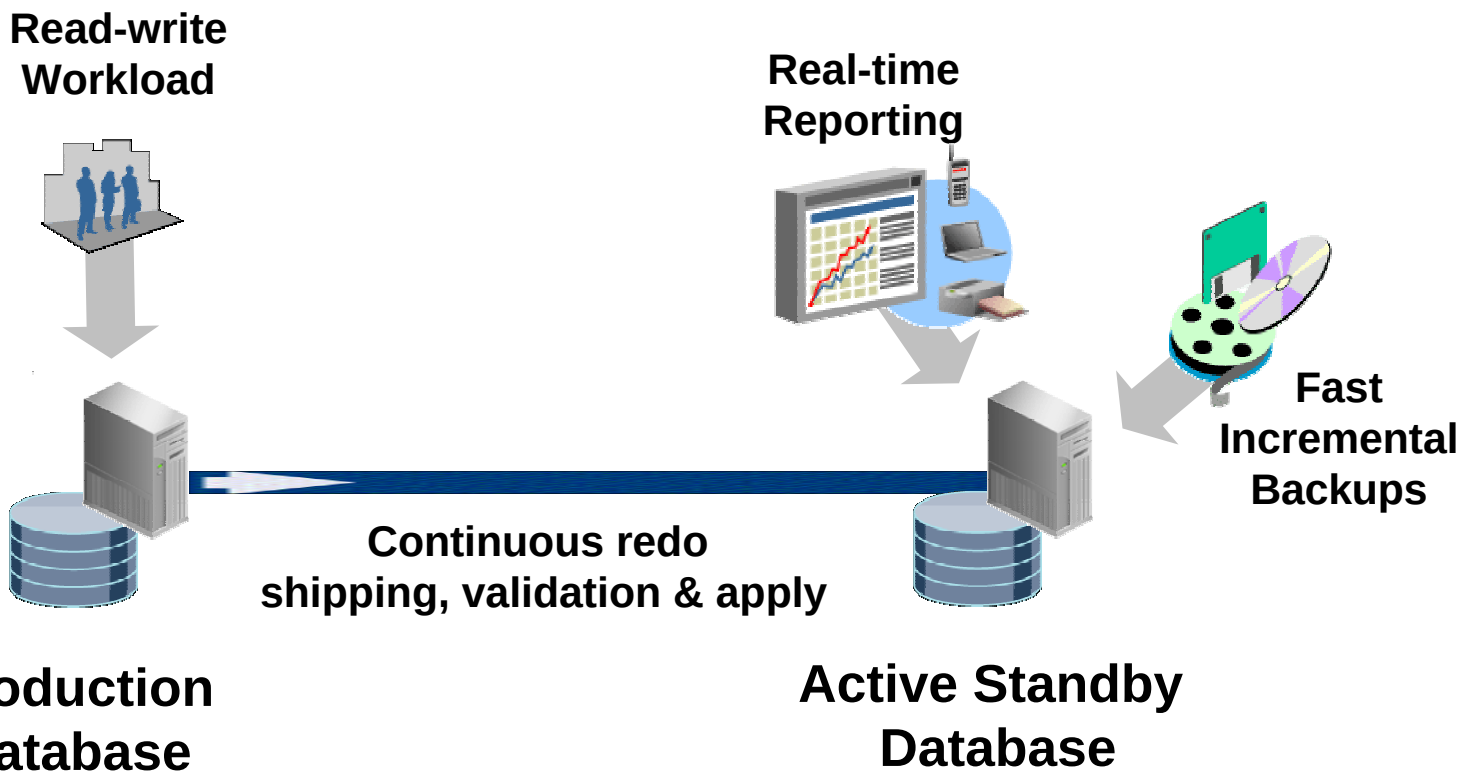


- Offload read-only queries to an up-to-date physical standby
- Use fast incremental backups on a physical standby – up to 20x faster

# Standby is used as Production System

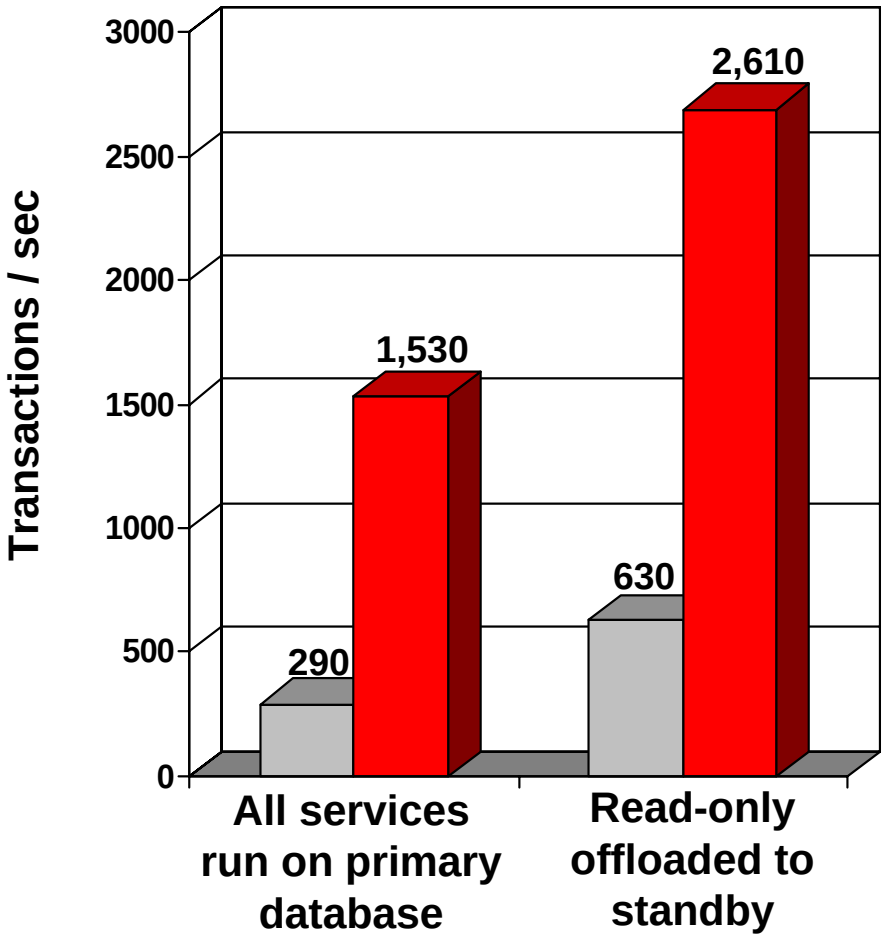


- More scalable
- Better performance
  - Eliminate contention between read-wite and read-only workload
  - Simplify performance tuning

**+ 117%** Read-write service

**+ 70%** Read-only service

# Standby is used to Reduce Planned Downtime

- Database rolling upgrades
  - Transient Logical Standby
- Migrations to ASM and/or RAC
- Technology refresh – servers and storage
- Windows/Linux migrations *
- 32bit/64bit migrations*
- Implement major database changes in rolling fashion
  - e.g. ASSM, initrans, blocksize
- Implement new database features in rolling fashion
  - e.g. Advanced Compression, SecureFiles, Exadata Storage

* see Metalink Note 413484.1

# Standby is used to Eliminate Risk

## Data Guard Snapshot Standby – Ideal for Testing

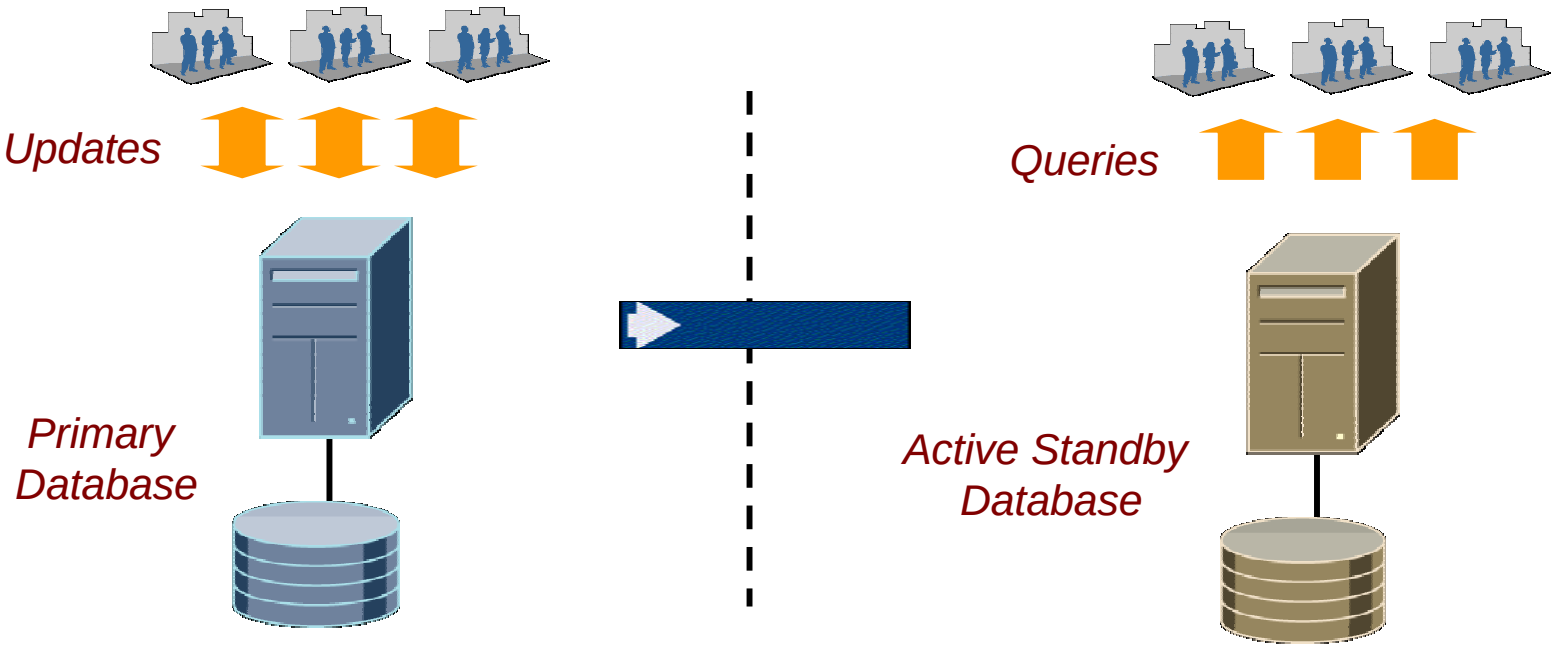

# Standby is used to Eliminate Risk

## Data Guard Snapshot Standby – Ideal for Testing

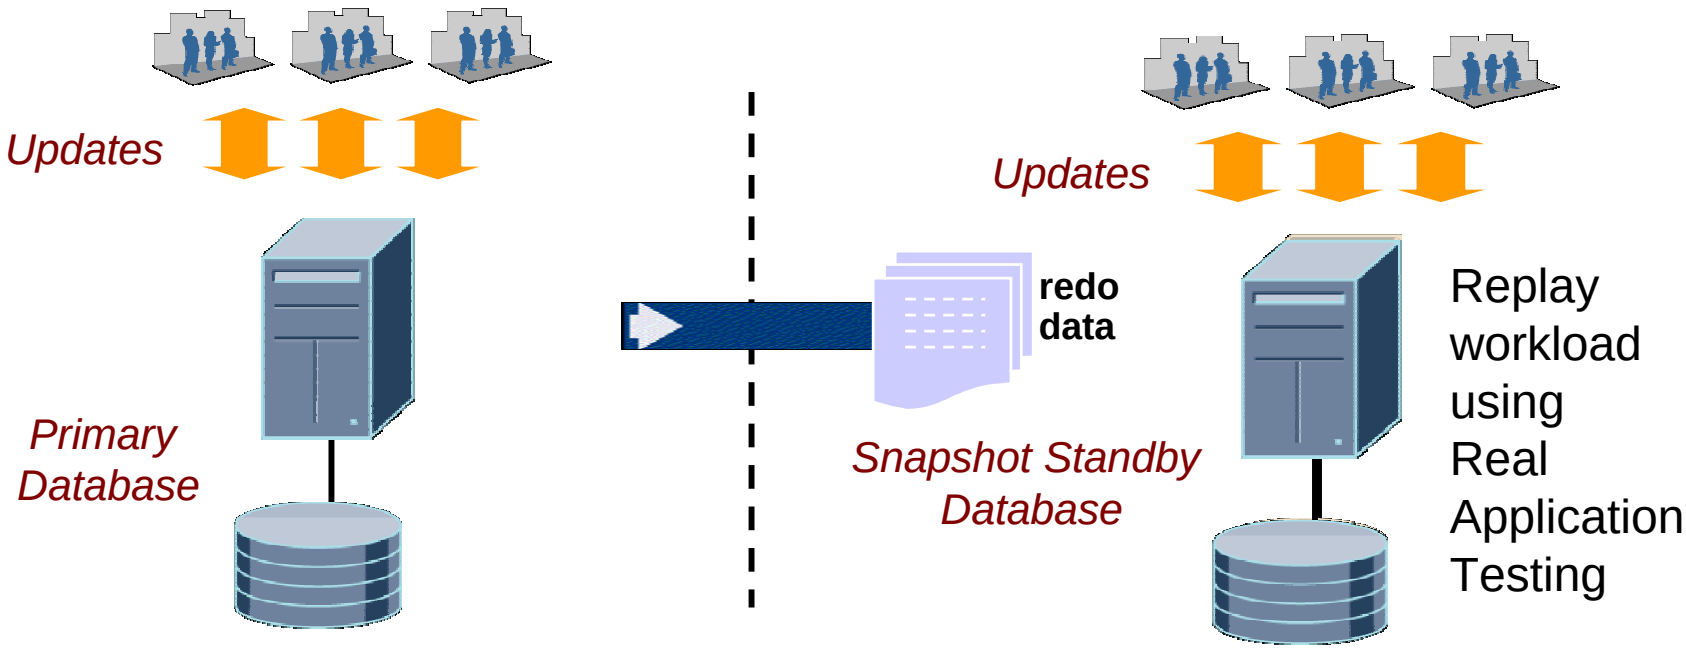


```
DGMGRL> convert database <name> to snapshot standby;
```

# Standby is used to Eliminate Risk

## Data Guard Snapshot Standby – Ideal for Testing

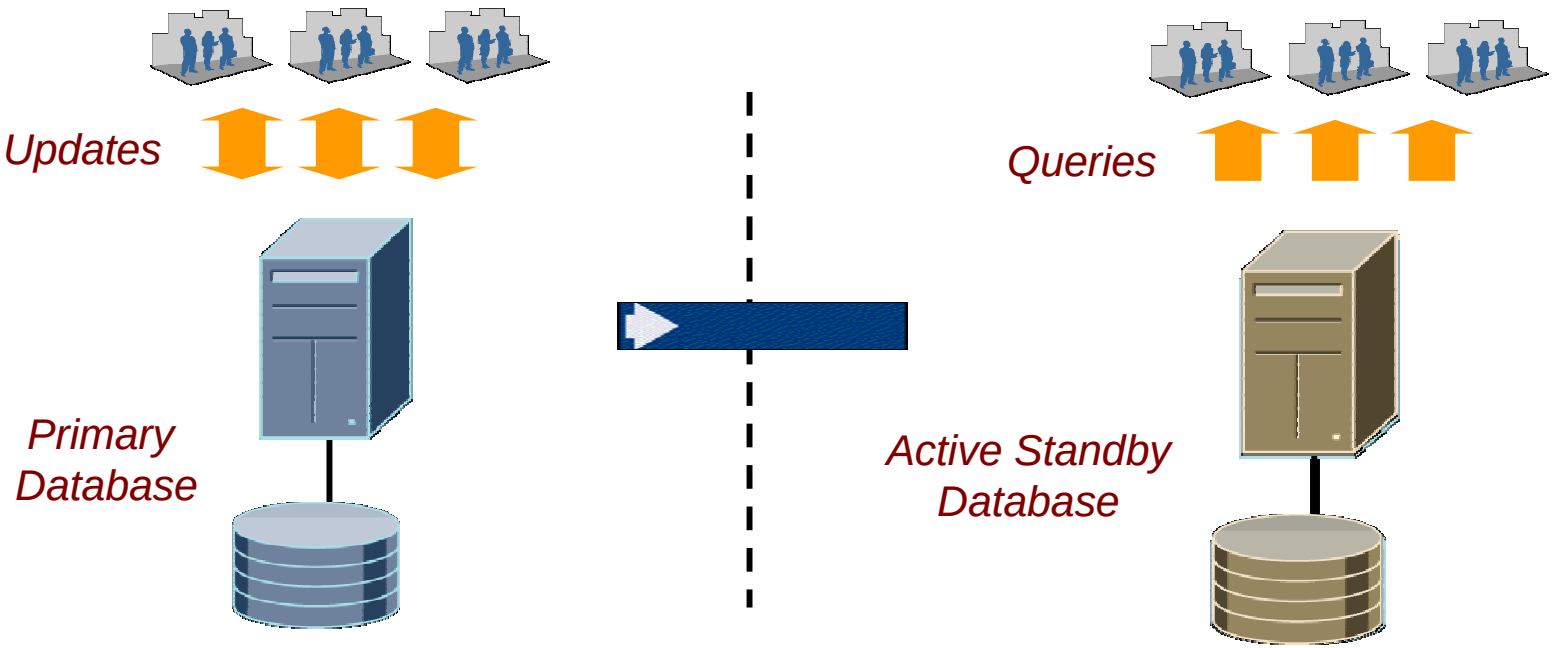


```
DGMGRL> convert database <name> to snapshot standby;
DGMGRL> convert database <name> to physical standby;
```

# HA Attribute: Simple to Manage

Active Data Guard

- All data types
- All storage attributes
- All DDL
- Fewest moving parts
- Based on media recovery – mature technology
- Highest performance
- Guaranteed EXACT replica of production

# HA Attribute: Simple to Manage

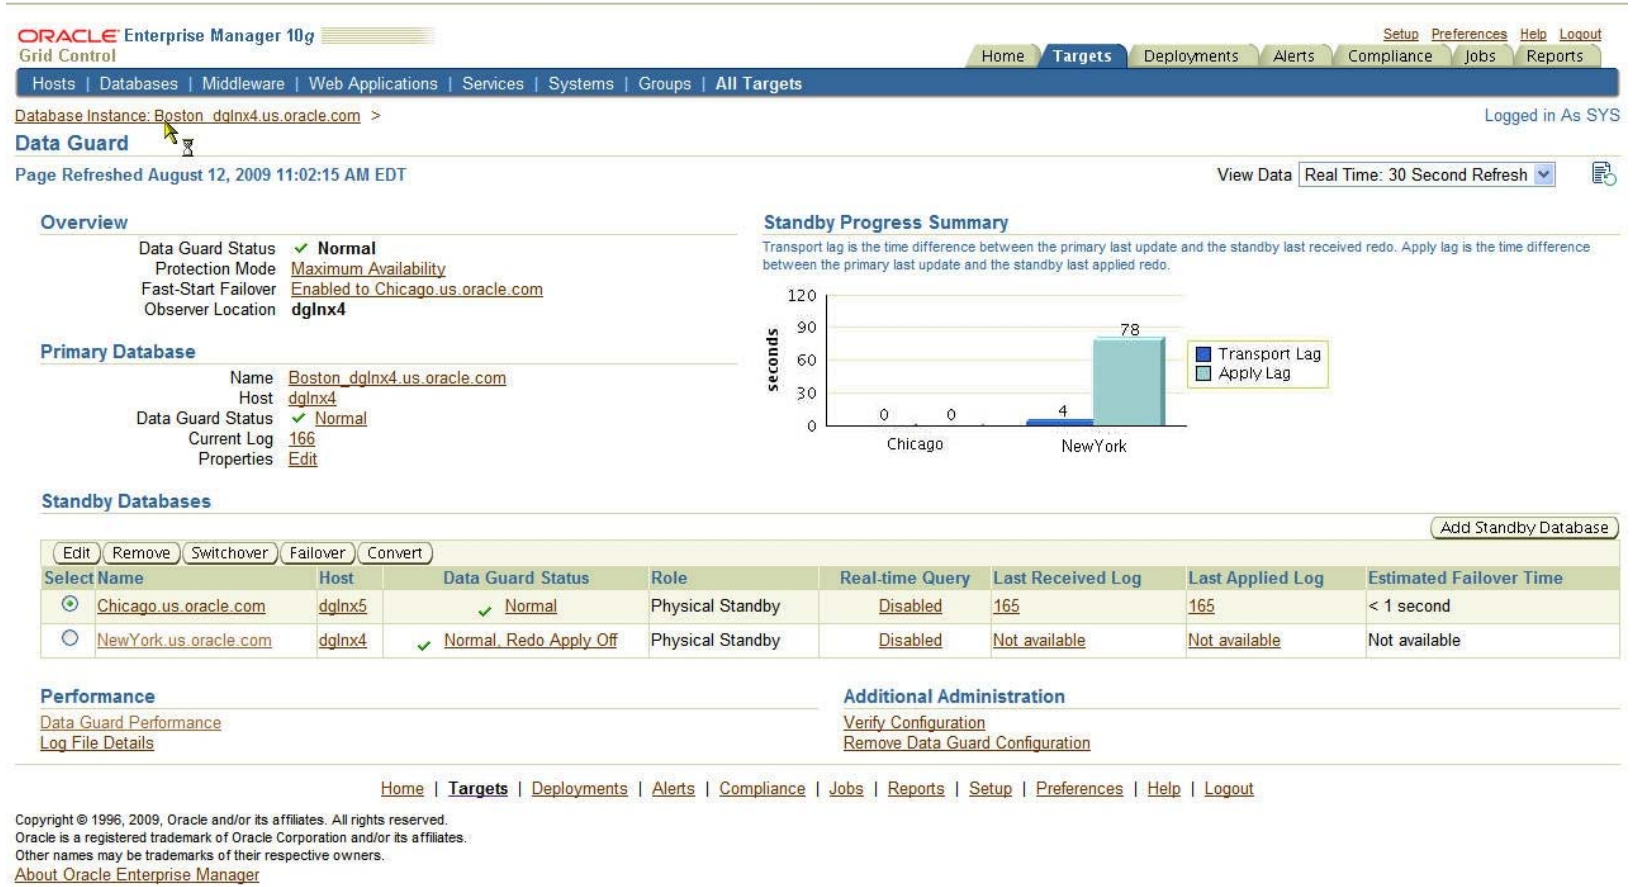

# Program

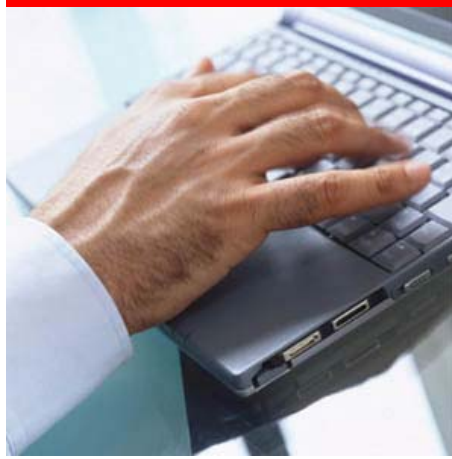

- Traditional approach to HA
- The ultimate HA solution
- Active Data Guard 11.2
- Implementation
- Resources

# Adding a Local Data Guard Standby Database

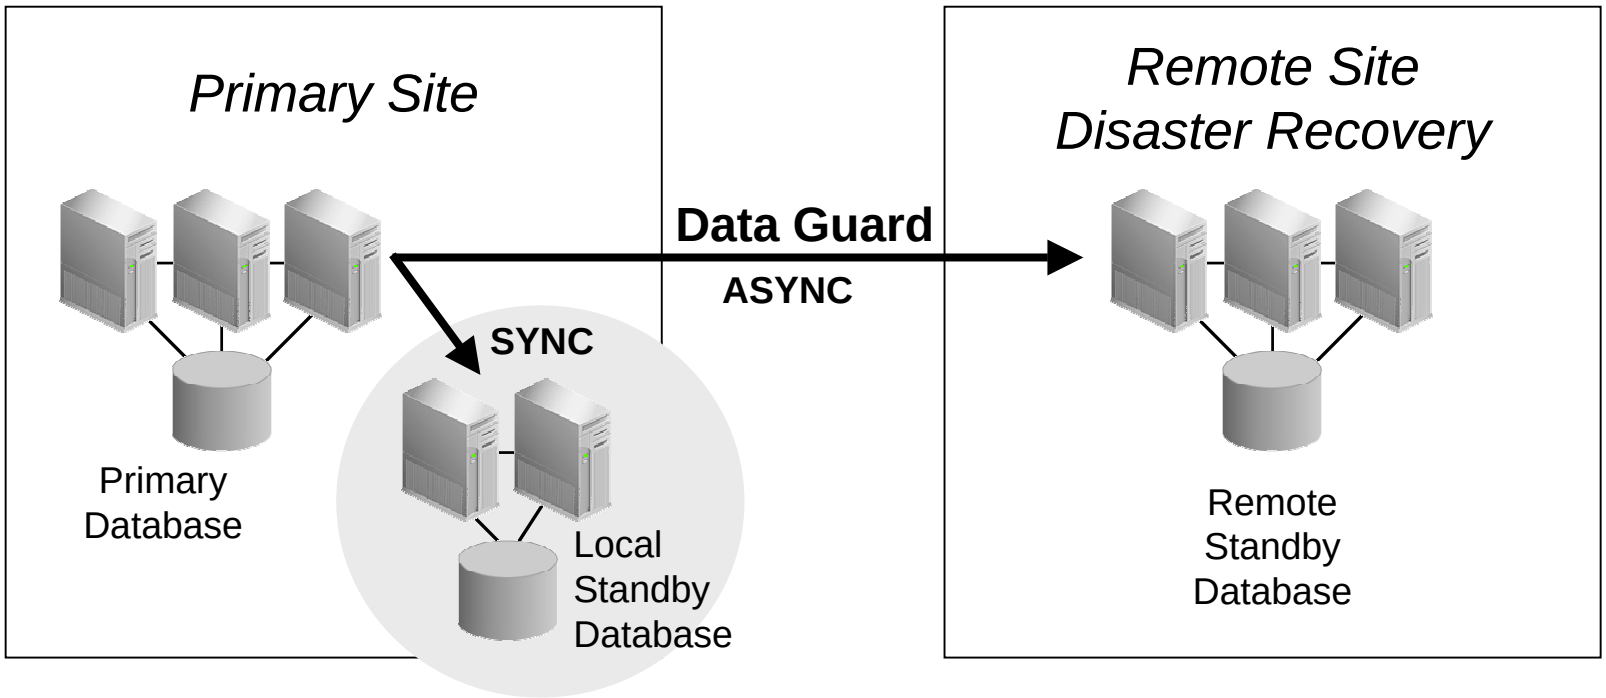

# Key Components

- Local physical standby – Maximum Availability
- Active Data Guard
- Data Guard Broker
- Data Guard Observer and Fast-Start Failover
- Flashback Database
- Fast Application Failover

# Implementation Considerations

## Data Guard Transport Tuning and Configuration

- Local Standby
  - Low latency network (ideally less than 5ms)
  - Maximum Availability Mode with SYNC transport
  - Set `NET_TIMEOUT` to 10 seconds from default of 30
  - Standby redo logs on fast storage
- Remote Standby
  - High network latency
  - ASYNC transport
  - Potentially increase log_buffer to ensure LNS reads from memory instead of disk (MetaLink Note 951152.1)
  - Tune TCP socket buffer sizes and device queues
    - Value is a function of bandwidth and latency
    - See HA Best Practices

# Implementation Considerations

## Basic Configuration

- Flashback Database
  - Configure on all databases in the configuration
  - Appropriately size Flash Recovery Area
  - **`FLASHBACK_RETENTION_PERIOD`** minimum of 60 minutes
  - See MetaLink Note 565535.1 for performance best practices
- Data Guard Broker
  - Required for Fast-Start Failover
  - Required for auto-restart of role specific database services (11.2)
  - Required for Fast Application Notification
  - Close integration with RAC (ie apply instance failover)
  - Simplified role transitions when using multiple standbys
  - Check MetaLink for Data Guard Broker bundled patch
    - E.g. 10.2.0.4 bundle has backports of several Broker 11.1 features

# Implementation Considerations

## Fast-Start Failover

- Data Guard Observer
  - Local standby is the Fast-Start Failover Target
  - Deploy Observer on 3rd host, independent of primary/standby
  - Set **`FastStartFailoverThreshold`**
    - 10 seconds for single instance databases
    - 20 seconds plus time for node eviction for Oracle RAC
  - Use Oracle Enterprise Manager for Observer HA
    - Auto restart of Observer on new host

# Implementation Considerations

## Configuring Client Failover

- Role based services (11.2)
  - Application service only runs on primary database
- All primary and standby hostnames in ADDRESS_LIST / URL
- Outbound connect timeout
  - Limits amount of time spent waiting for connection to failed resources
- Application notification
  - Break clients out of TCP with Fast Application Notification events
- Pre Data Guard 11.2 please refer to Client Failover Best Practices

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

# The Result

## An HA architecture built on the assumption that eventually something will fail

# Ultimate High Availability

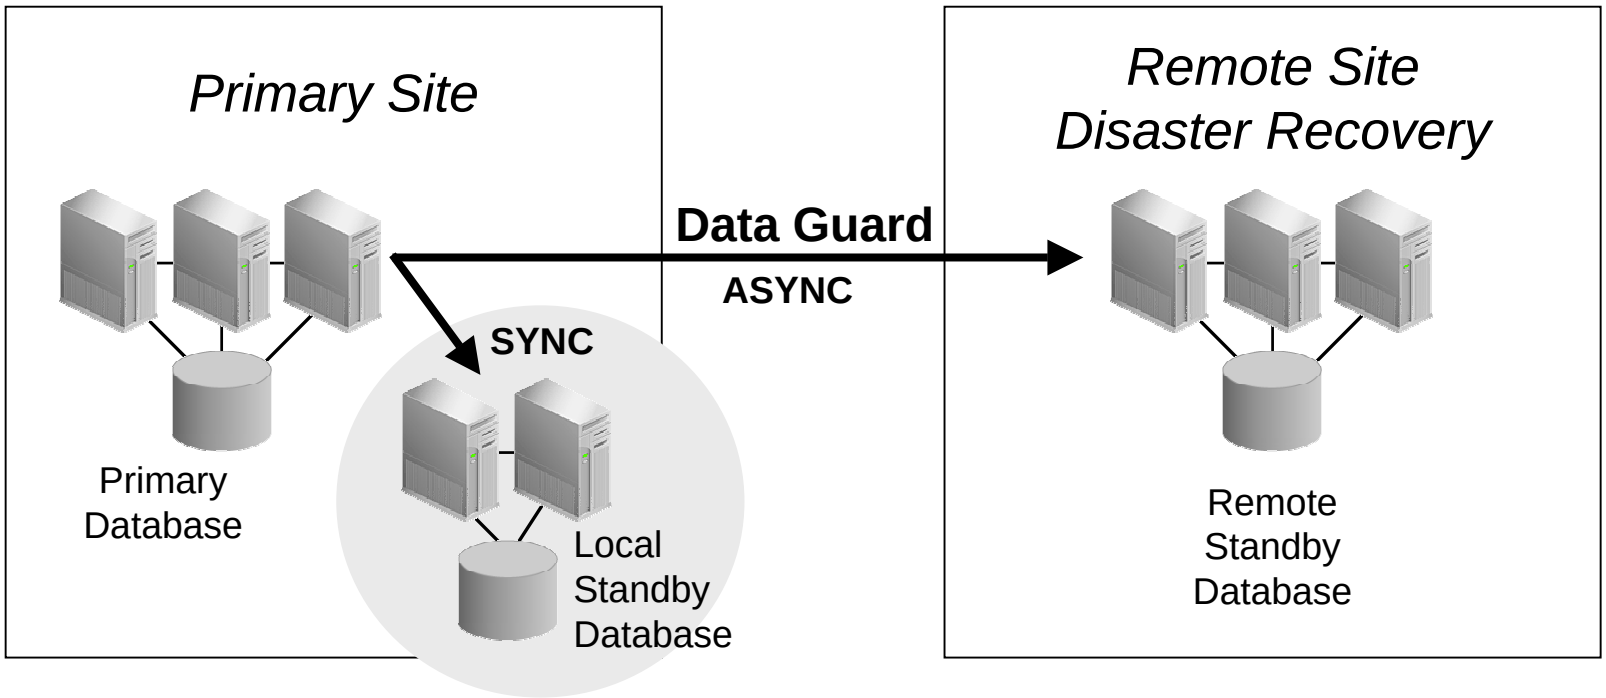

# Ultimate High Availability

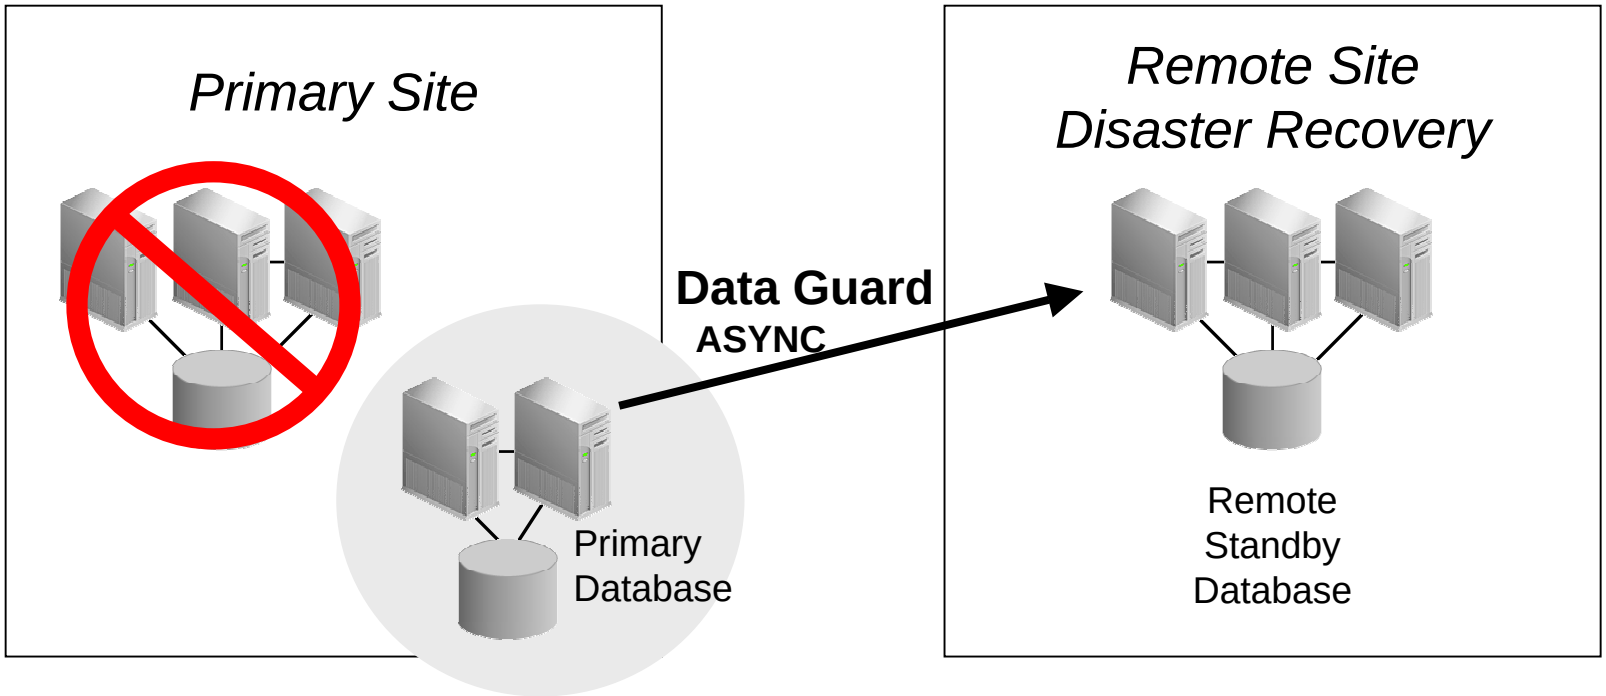


| | |
|---|---|
| ✓ Redundancy with isolation | ✓ Automatic failover |
| ✓ Zero data loss | ✓ Full systems utilization |
| ✓ Extreme performance | ✓ Management simplicity |

# Start Here

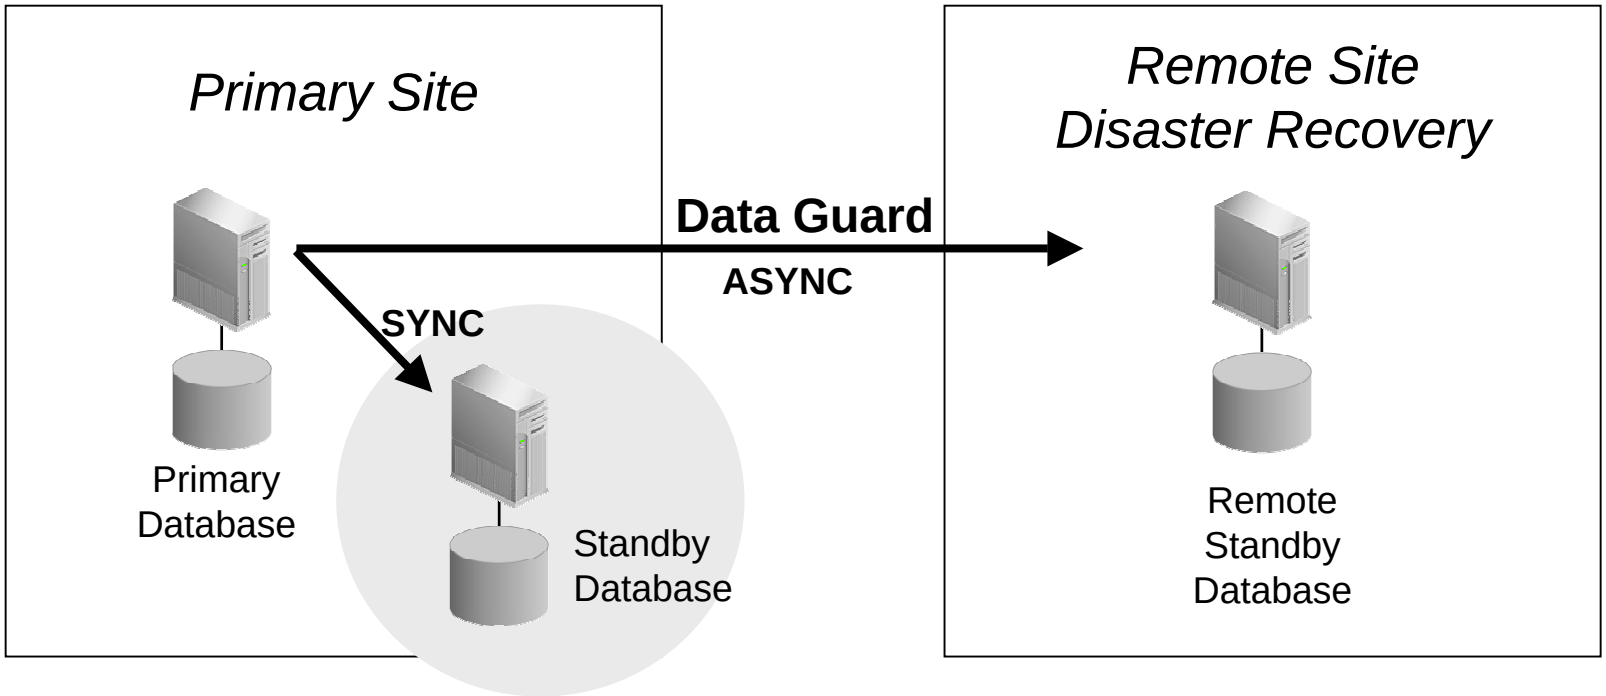


| | |
|---|---|
| ✓ Redundancy with isolation | ✓ Automatic failover |
| ✓ Zero data loss | ✓ Full systems utilization |
| ✓ Extreme performance | ✓ Management simplicity |

# Key Best Practices Documentation

- HA Best Practices
  http://www.oracle.com/pls/db111/portal.portal_db?selected=14&frame=
- Active Data Guard and Redo Apply
  http://www.oracle.com/technology/deploy/availability/pdf/maa_wp_11gr1_activedataguard.pdf
- Data Guard Redo Transport
  http://www.oracle.com/technology/deploy/availability/pdf/MAA_WP_10gR2_DataGuardNetworkBestPractices.pdf
- Data Guard Fast-Start Failover
  http://www.oracle.com/technology/deploy/availability/pdf/MAA_WP_10gR2_FastStartFailoverBestPractices.pdf
- Automating Client Failover (Data Guard 10g and 11gR1)
  http://www.oracle.com/technology/deploy/availability/pdf/MAA_WP_10gR2_ClientFailoverBestPractices.pdf
- Managing Data Guard Configurations with Multiple Standby Databases
  http://www.oracle.com/technology/deploy/availability/pdf/maa10gr2multiplestandbybp.pdf
- Using your Data Guard Standby for Real Application Testing
  http://www.oracle.com/technology/deploy/availability/pdf/oracle-openworld-2008/298770.pdf
- S307560 Active / Active Configurations with Oracle Active Data Guard
  http://www.oracle.com/technology/deploy/availability/pdf/oracle-openworld-2009/307560.pdf

# HA Sessions, Labs, & Demos by Oracle Development

**Sunday, 11 October – Hilton Hotel Imperial Ballroom B**

**3:45p** Online Application Upgrade

**Monday, 12 October – Marriott Hotel Golden Gate B1**

**11:30a** Introducing Oracle GoldenGate Products

**Monday, 12 October – Moscone South**

**1:00p** Oracle's HA Vision: What's New in 11.2, Room **103**

**4:00p** Database 11g: Performance Innovations, Room **103**

**2:30p** Oracle Streams: What's New in 11.2, Room **301**

**5:30p** Comparing Data Protection Solutions, Room **102**

**Tuesday, 13 October – Moscone South**

**11:30a** Oracle Streams: Replication Made Easy, Room **308**

**11:30a** Backup & Recovery on the Database Machine, Room **307**

**11:30a** Next-Generation Database Grid Overview, Room **103**

**1:00p** Oracle Data Guard: What's New in 11.2, Room **104**

**2:30p** GoldenGate and Streams - The Future, Room **270**

**2:30p** Backup & Recovery Best Practices, Room **104**

**2:30p** Single-Instance RAC, Room **300**

**4:00p** Enterprise Manager HA Best Practices, Room **303**

## Hands-on Labs Marriott Hotel Golden Gate B2

Monday **11:30a-2:00p** Oracle Active Data Guard, Parts I & II

Thursday **9:00a-11:30a** Oracle Active Data Guard, Parts I & II

**Tuesday, 13 October – Marriott Hotel Golden Gate B1**

**11:30a** GoldenGate Zero-Downtime Application Upgrades

**1:00p** GoldenGate Deep Dive: Architecture for Real-Time

**Wednesday, 14 October – Moscone South**

**10:15a** Announcing OSB 10.3, Room **300**

**11:45a** Active Data Guard, Room **103**

**5:00p** Exadata Storage & Database Machine, Room **104**

**Thursday, 15 October – Moscone South**

**9:00a** Empowering Availability for Apps, Room **300**

**12:00p** Exadata Technical Deep Dive, Room **307**

**1:30p** Zero-Risk DB Maintenance, Room **103**

## Demos Moscone West DEMOGrounds

Mon & Tue **10:30a - 6:30p;** Wed **9:15a - 5:15p**

Maximum Availability Architecture (MAA), **W-045**

Oracle Streams: Replication & Advanced Queuing, **W-043**

Oracle Active Data Guard, **W-048**

Oracle Secure Backup, **W-044**

Oracle Recovery Manager & Flashback, **W-046**

Oracle GoldenGate, **3709**

# For More Information

search.oracle.com

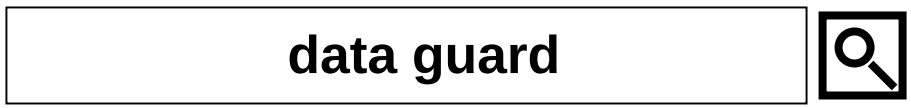


**or**

**oracle.com/ha**


ORACLE®

ORACLE IS THE INFORMATION COMPANY