18th Annual International zSeries Oracle SIG Conference Present:

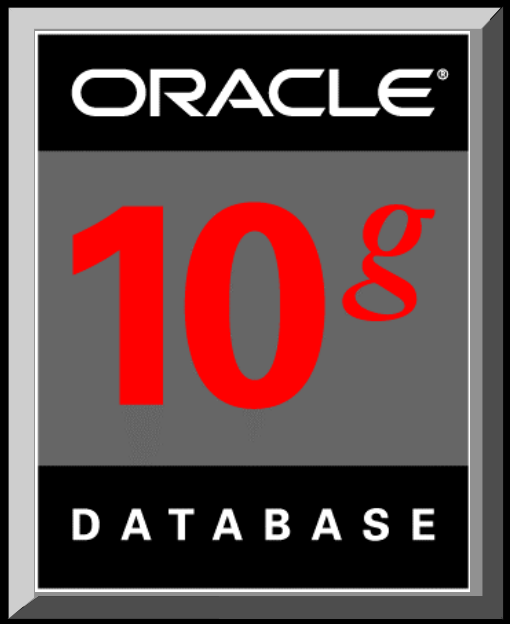


# Oracle Database Backup & Recovery, Flashback* Whatever, & Data Guard

**Tammy Bednar**
**Tammy.Bednar@oracle.com**
**Manager, HA Solutions & Backup / Recovery**
**Server Technologies**
**Oracle Corporation**

**Ashish Ray**
**Ashish.Ray@oracle.com**
**Manager, HA Solutions & Data Guard**
**Server Technologies**
**Oracle Corporation**

ORACLE

# Agenda

- Recovery Manager Overview
  - Oracle Database 10g Features
- Flashback *
  - Granular Human Error Correction
- Data Guard
  - Overview
  - Enterprise Manager Integration
  - Best Practices for HA
- Questions and Answers

# Recovery Manager: Oracle's Backup & Recovery Utility

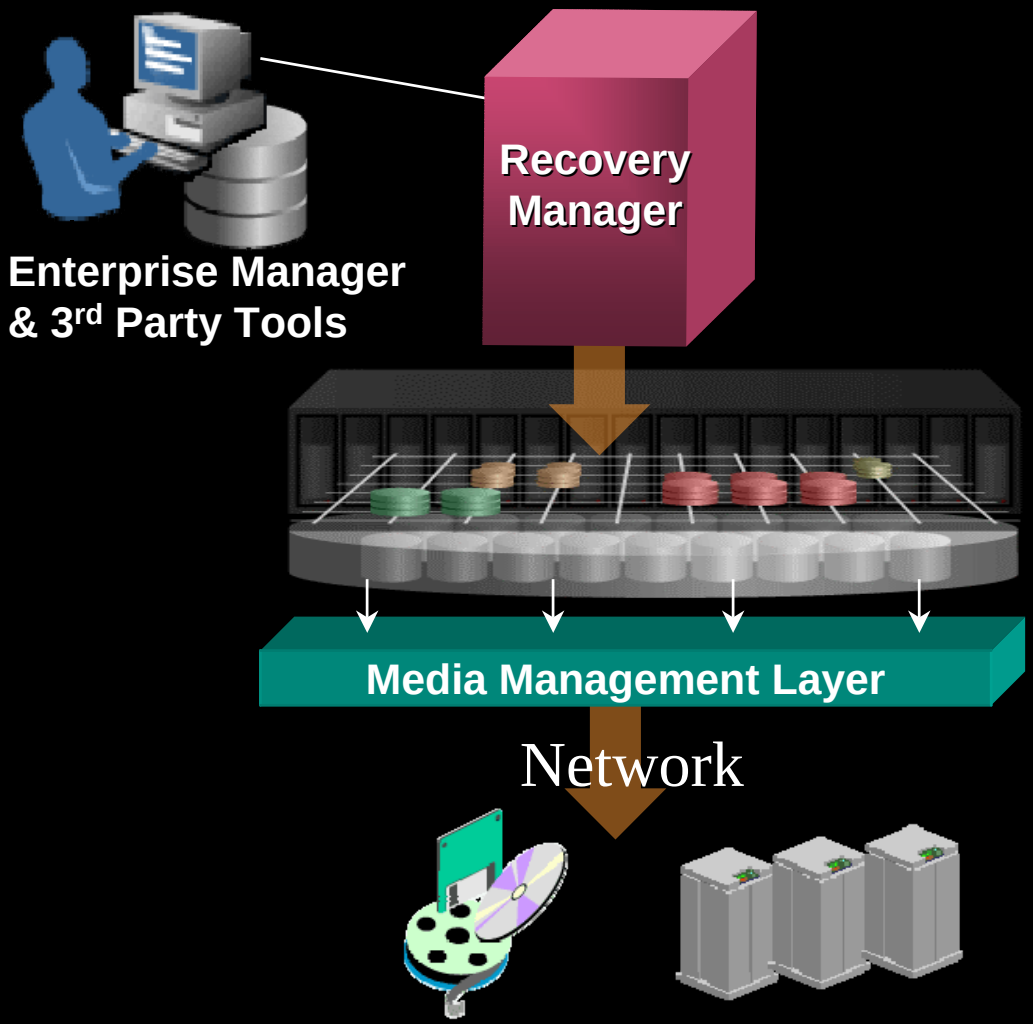


- Intimate knowledge of database file formats and recovery procedures
- Manages and automates the backup, restore, and recovery process
- Creates and maintains backup policies
- Catalogs all backup and recovery activities
- Operates on-line and in parallel for fast processing
- Corrupt block detection during backup and restore and the ability to validate backups
- Integrated with Enterprise Manager & 3rd party network backup products

# Recovery Manager

- Request backup at database, tablespace, or datafile level
- Incremental backups (up to 4 levels)
- Backup to tape through third party media manager software
- Comprehensive reporting
- Stored scripts that automate backup and recovery procedures
- Automatic parallelization of backup, restore, and recovery
- Backups can be restricted to limit reads per file, per second to avoid interfering with OLTP work
- No generation of extra redo during online database backups
- Proxy Copy Backup Accelerator allows fast copy technology at the storage subsystem level

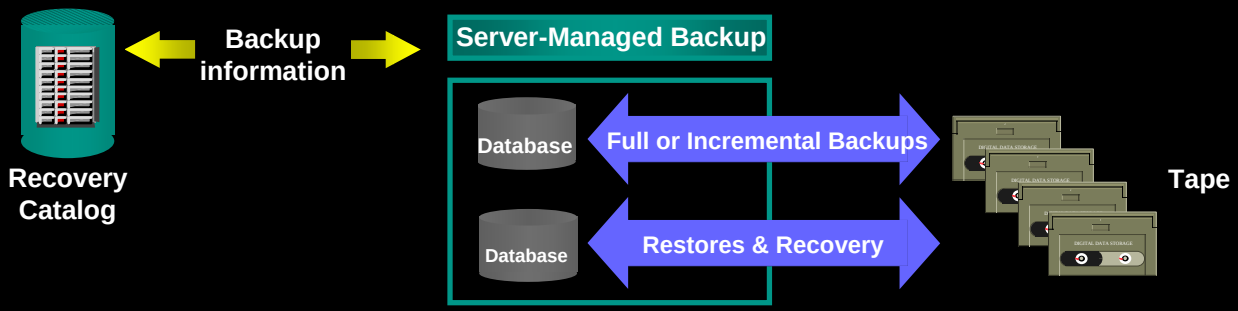

# Flash Recovery Area

- Unified storage location for all recovery files and recovery related activities in an Oracle Database.
  - Centralized location for control files, online redo logs, archive logs, flashback logs, backups
  - A flash recovery area can be defined as file system or ASM disk group
  - A single recovery area can be shared by more than one database
- Minimize the number of initialization parameters to set when you create a database
  - Define a database area and flash recovery area location
  - Oracle creates and manages all files using OMF

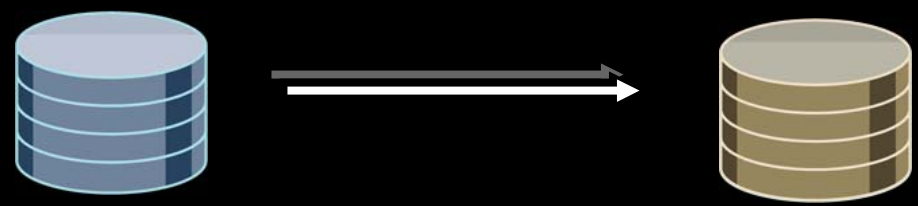

Database Area Flash Recovery Area

# Flash Recovery Area Space Management

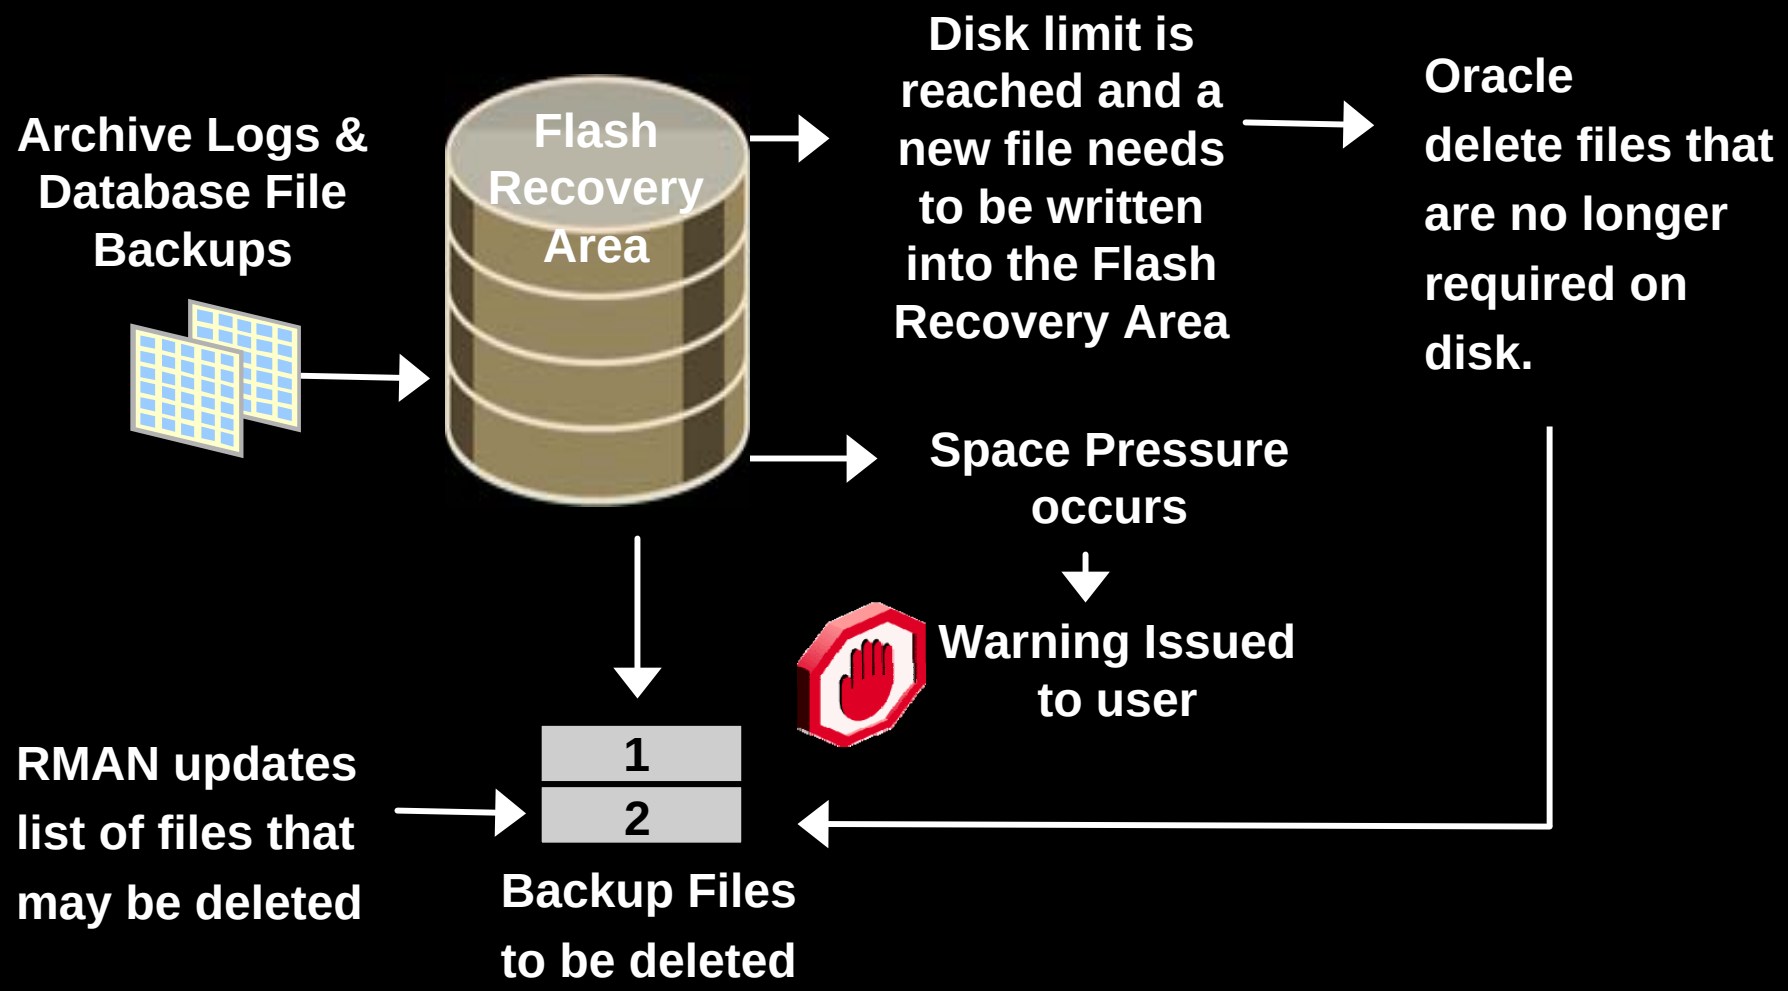

# Change Tracking File

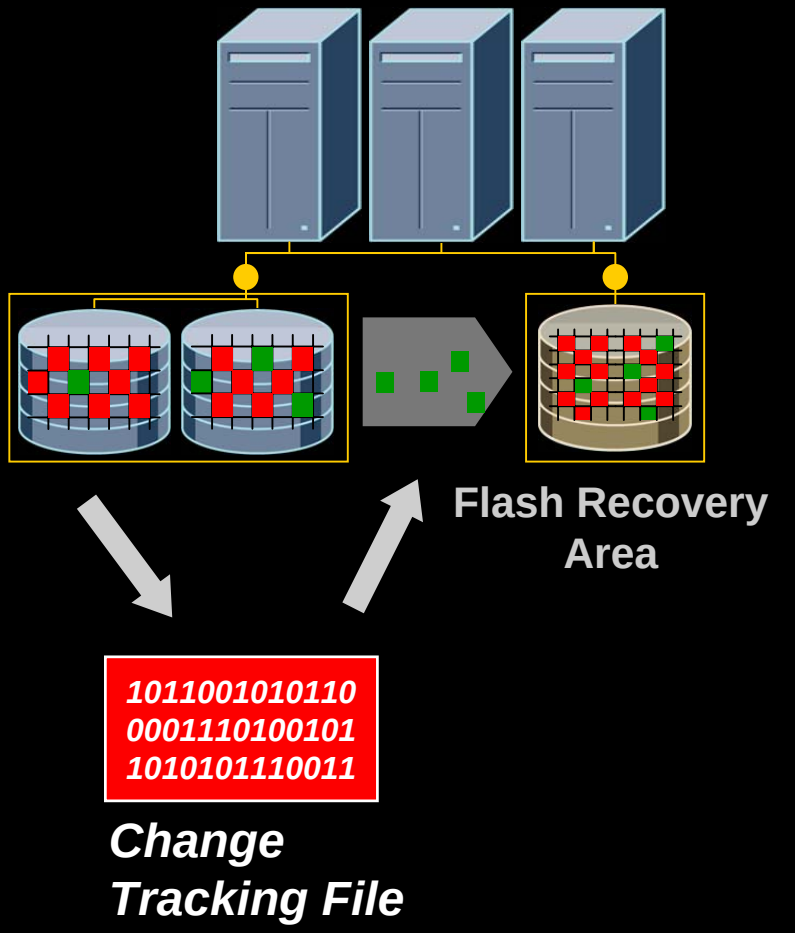


- Optimizes incremental backups
  - Track which blocks have changed since last backup
- Integrated change tracking file
  - Changed blocks are tracked as redo is generated
  - RMAN backup automatically uses changed block list

# Incrementally Updated Backups

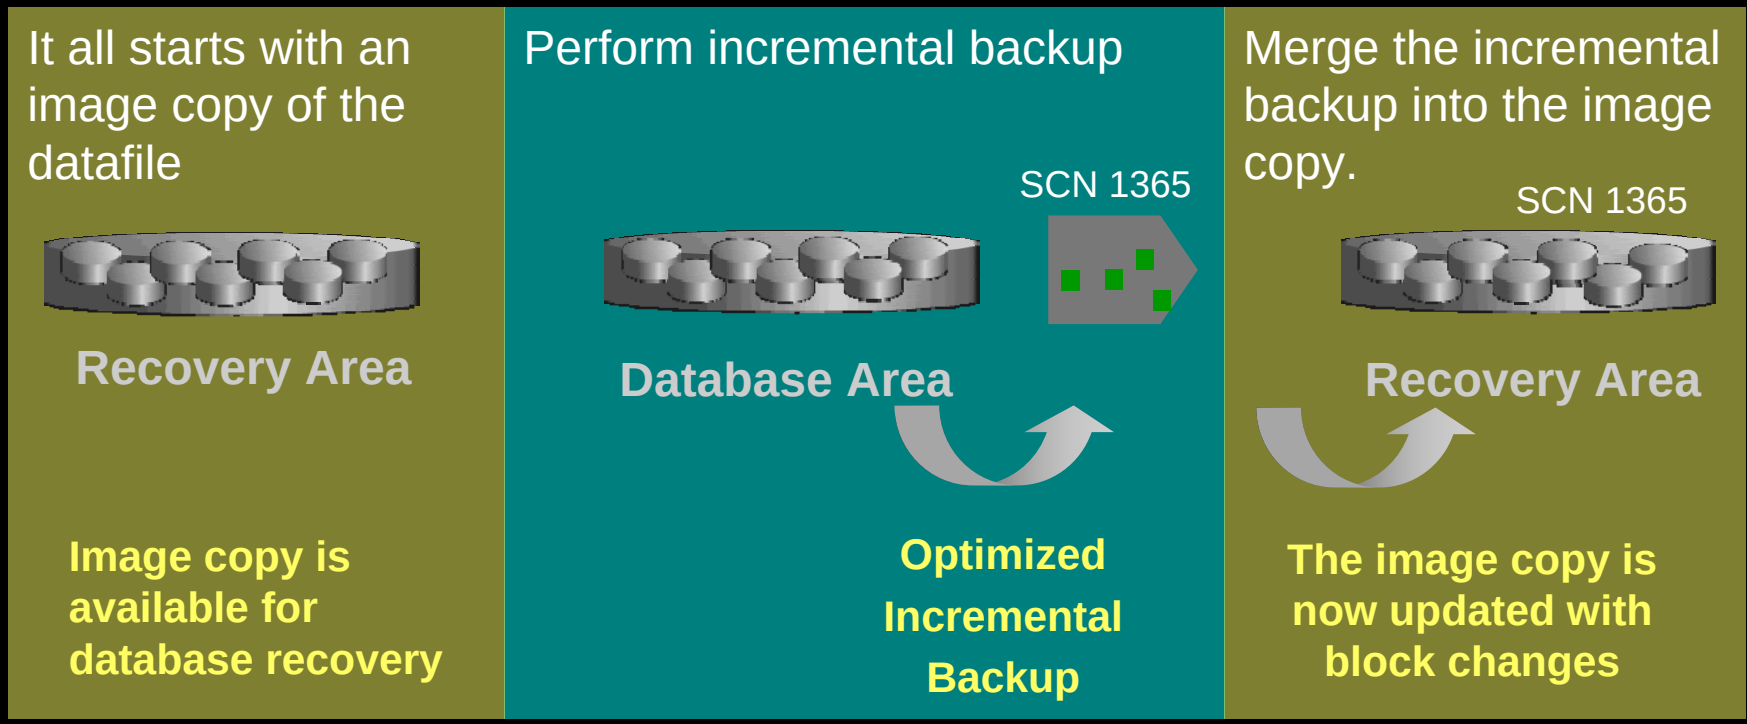


- **Eliminate the need to perform a whole database backup.**
- **Reduce the time required for media recovery since the image copy is updated with the latest block changes.**

# Eliminate Shrinking Backup Window Syndrome!

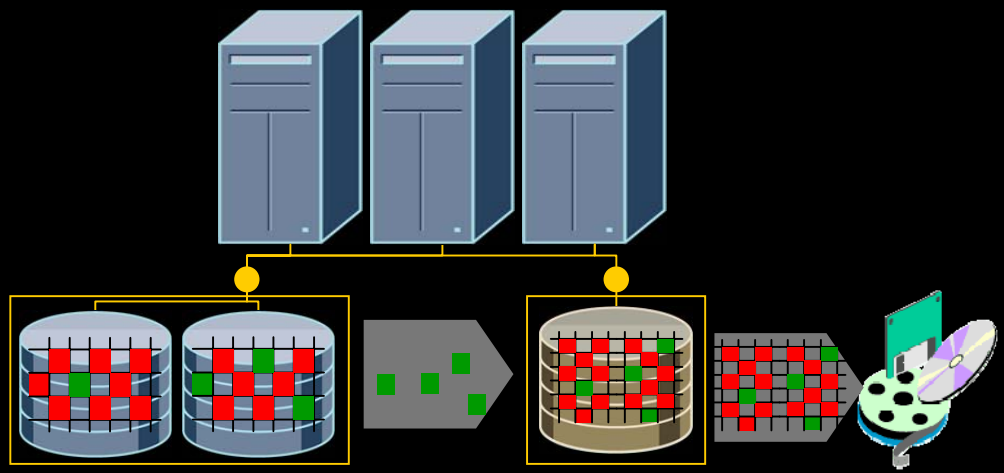


**Two Independent Disk Systems**

- Fully automatic disk based backup and recovery
  - Set it and Forget it
- Nightly incremental backup rolls forward recovery area backup
  - Changed blocks are tracked in production DB
- Full scan is never needed
  - Dramatically faster (20x)
  - Blocks validated to prevent corruption of backup copy
- Use low cost ATA disk array for recovery area

# Oracle Backup – The Lowest Cost Tape Backup Manager

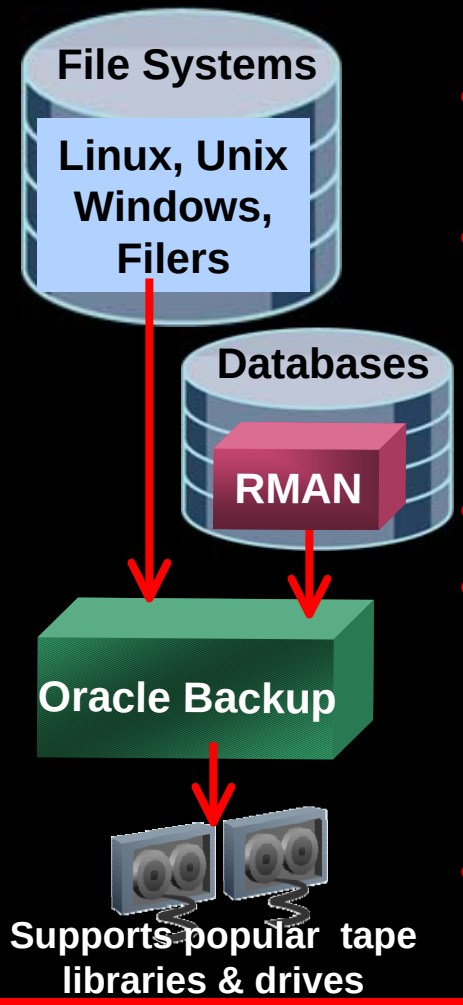


- Oracle Backup is ideal for customers seeking a low cost alternative to complex backup products
- Best integrated end-to-end backup of Oracle Databases
  - Media manger for RMAN backup and recovery of Oracle9*i* and 10*g* databases to tape
  - Fastest Database Backup on the market
- Backup Oracle Home, App Server and other file systems
- Oracle Backup includes:
  - Centralized management of network backups
  - Scalability to low 100's of servers, 10's of millions of files
  - Easy management through Enterprise Manager
- Bundled with Oracle Database – replaces LSSV
  - Single vendor support

# Human Error

- Estimated to be the biggest single cause of downtime
- Need to quickly determine what happened and fix it
  - Localized damage
    - Needs surgical detection and repair
    - Example – removed wrong person named ‘Smith’
  - Widespread damage
    - Requires drastic action to avoid long downtime
    - Example – batch job deletes this month’s orders
- Analysis and correction using traditional recovery is slow and complex
  - Restore database to point in time and extract data
- Oracle Database 10*g* is a breakthrough release for human error correction

# Flashback Time Navigation

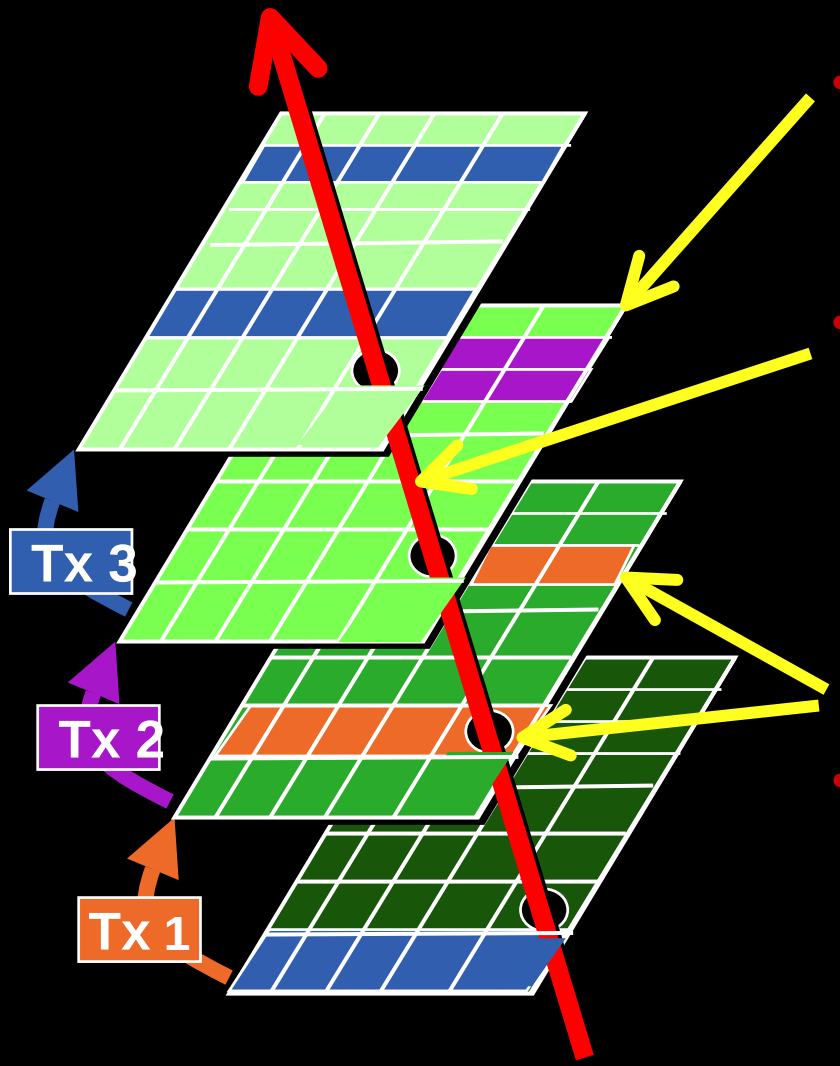


- **Flashback Query**
  - Query all data at point in time

```
Select * from Emp AS OF '2:00 P.M.' where …
```

- **Flashback Versions Query**
  - See all versions of a row between two times
  - See transactions that changed the row

```
Select * from Emp VERSIONS BETWEEN
 '2:00 PM' and '3:00 PM' where …
```

- **Flashback Transaction Query**
  - See all changes made by a transaction

```
Select * from FLASHBACK_TRANSACTION_QUERY
 where xid = '000200030000002D';
```

# How Does Flashback Time Navigation Work?

- Leverages Oracle's unique multi-version read consistency architecture
  - The data image is saved in the undo tablespace (or Rollback Segments) before being modified
  - Flashback Query uses the data saved in the undo tablespace to recreate an image of the data as it existed at a time in the past.
- Oracle's Automatic Undo Management feature allows administrators to specify how long they wish to retain the undo data
  - DBAs can control how far back a Flashback Query can go

# Flashback Query

A Time Machine for Your Data

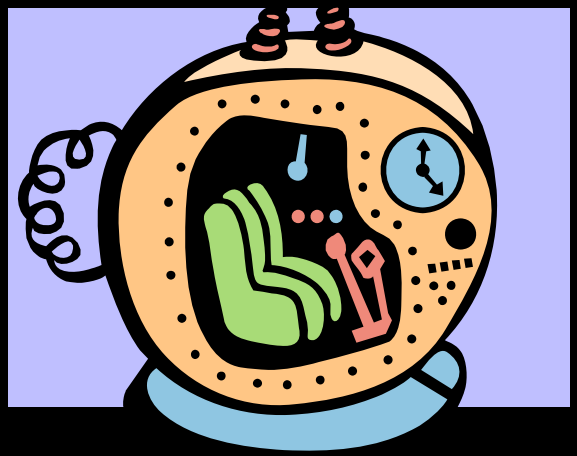

- **Flashback Query allows viewing data as it was before a mistake**
  - Query data at a time of your choosing
  - Standard SQL interface simplifies deployment
  - Self-service means faster, cheaper, and easier
  - Flashback Query is a fast operation to enable self service

Mistake

```
Delete from Emp
  where Ename='Smith';
```

Correction

```
Insert into Emp
  select * from Emp AS OF yesterday
    where Ename='Smith';
```

# Build Self Error Correcting Application

**Oracle Collaboration Suite utilizes Flashback Query's built in functionality!**

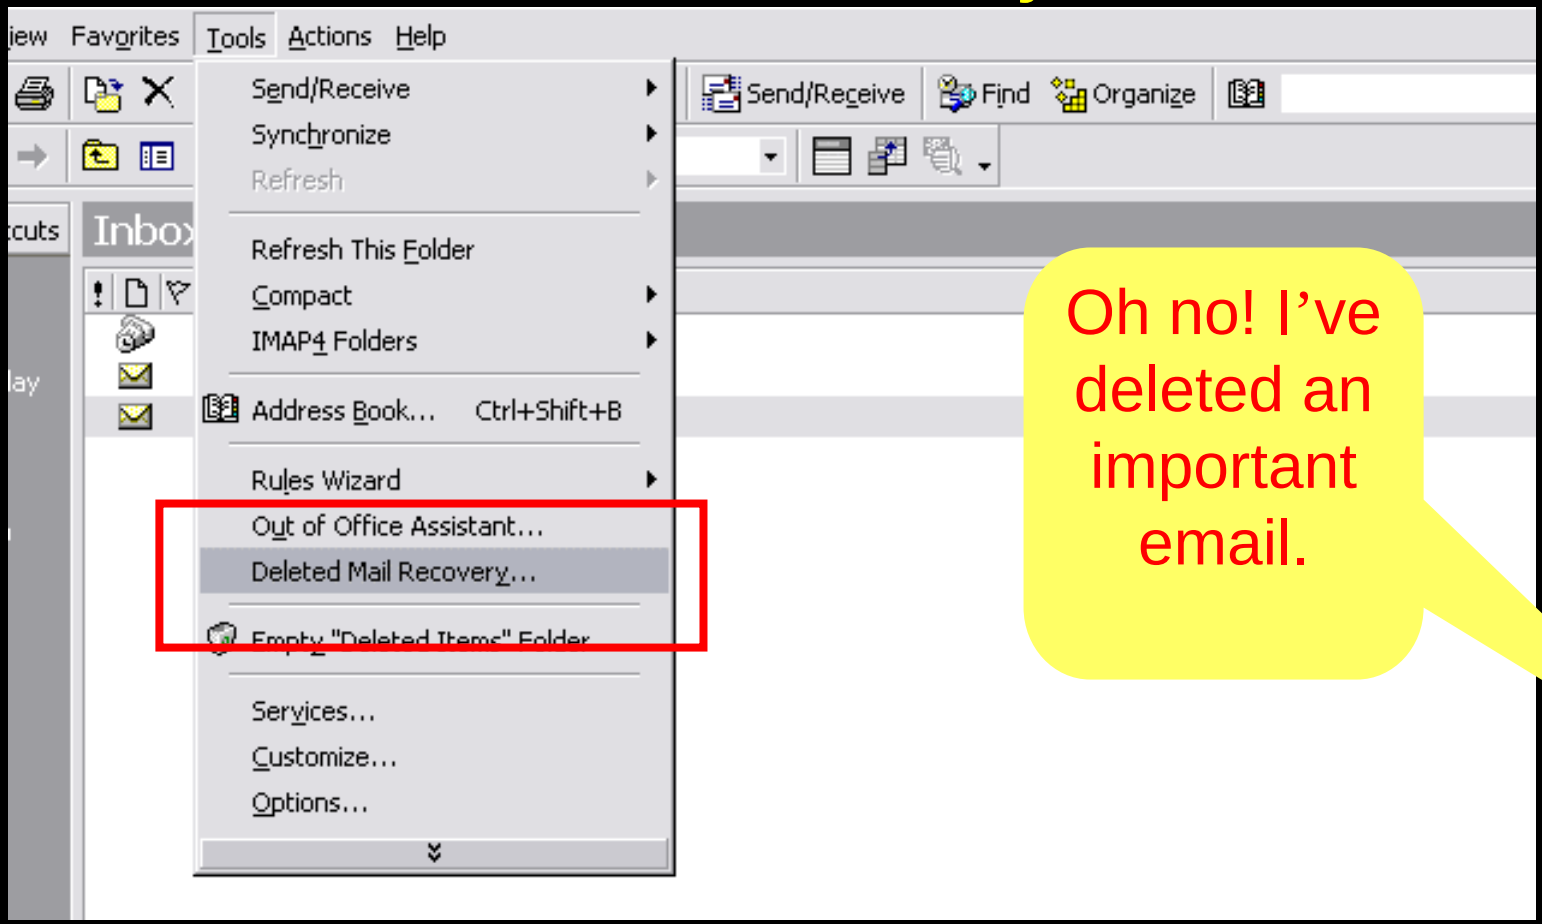

# Flashback Versions Query

- Provides a way to audit the rows of a table and retrieve information about the transactions that changed the rows.
- Retrieve all committed versions of the rows that exist or ever existed between the time the query was issued and a point in time in the past
- Use the transaction ID to perform transaction mining using LogMiner or Flashback Transaction Query to obtain additional information about the transaction.

# Flashback Versions Query

## View data changes over time

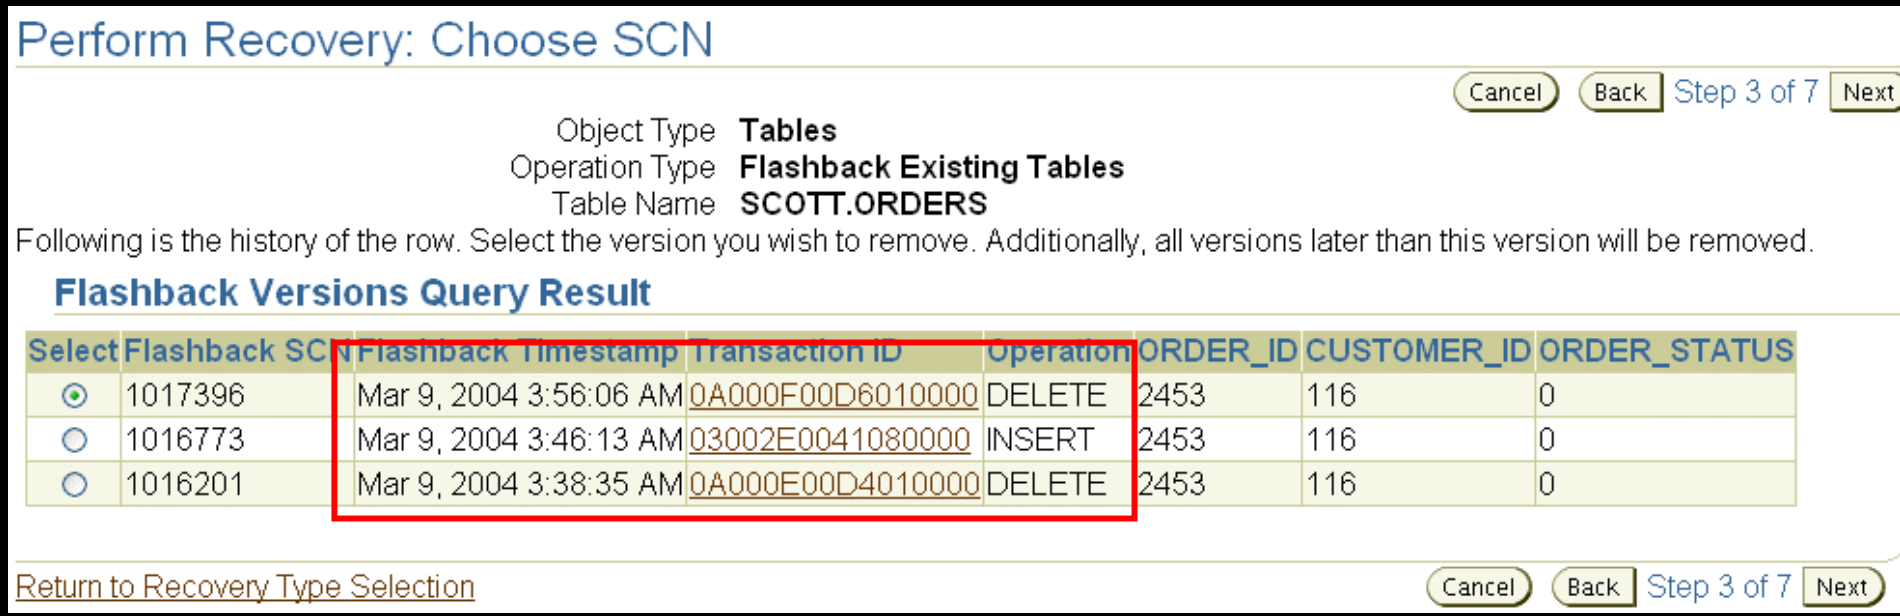

Perform Recovery: Choose SCN

Cancel Back Step 3 of 7 Next

Object Type **Tables**
Operation Type **Flashback Existing Tables**
Table Name **SCOTT.ORDERS**

Following is the history of the row. Select the version you wish to remove. Additionally, all versions later than this version will be removed.

**Flashback Versions Query Result**

| Select | Flashback SCN | Flashback Timestamp | Transaction ID | Operation | ORDER_ID | CUSTOMER_ID | ORDER_STATUS |
|---|---|---|---|---|---|---|---|
| ◉ | 1017396 | Mar 9, 2004 3:56:06 AM | 0A000F00D6010000 | DELETE | 2453 | 116 | 0 |
| ○ | 1016773 | Mar 9, 2004 3:46:13 AM | 03002E0041080000 | INSERT | 2453 | 116 | 0 |
| ○ | 1016201 | Mar 9, 2004 3:38:35 AM | 0A000E00D4010000 | DELETE | 2453 | 116 | 0 |

Return to Recovery Type Selection

Cancel Back Step 3 of 7 Next

- Fast and online access to data changes
- Utilizes the database undo and requires no additional overhead
- You control how far back in time data can be accessed

# Flashback Transaction Query

- Provides a way for you to view changes made to the database at the transaction level
- When used in conjunction with Flashback Versions Query, it allows you to easily recover from user or application errors.
- Benefits
  - Increase online diagnosability of problems in your database
  - Perform analysis and audits of transactions
  - Fast recovery at the transaction level

# Flashback Transaction Query
## View Transaction Details

- View all objects affected by a single transaction
- Using the UNDO SQL, quickly recover from the erroneous transaction

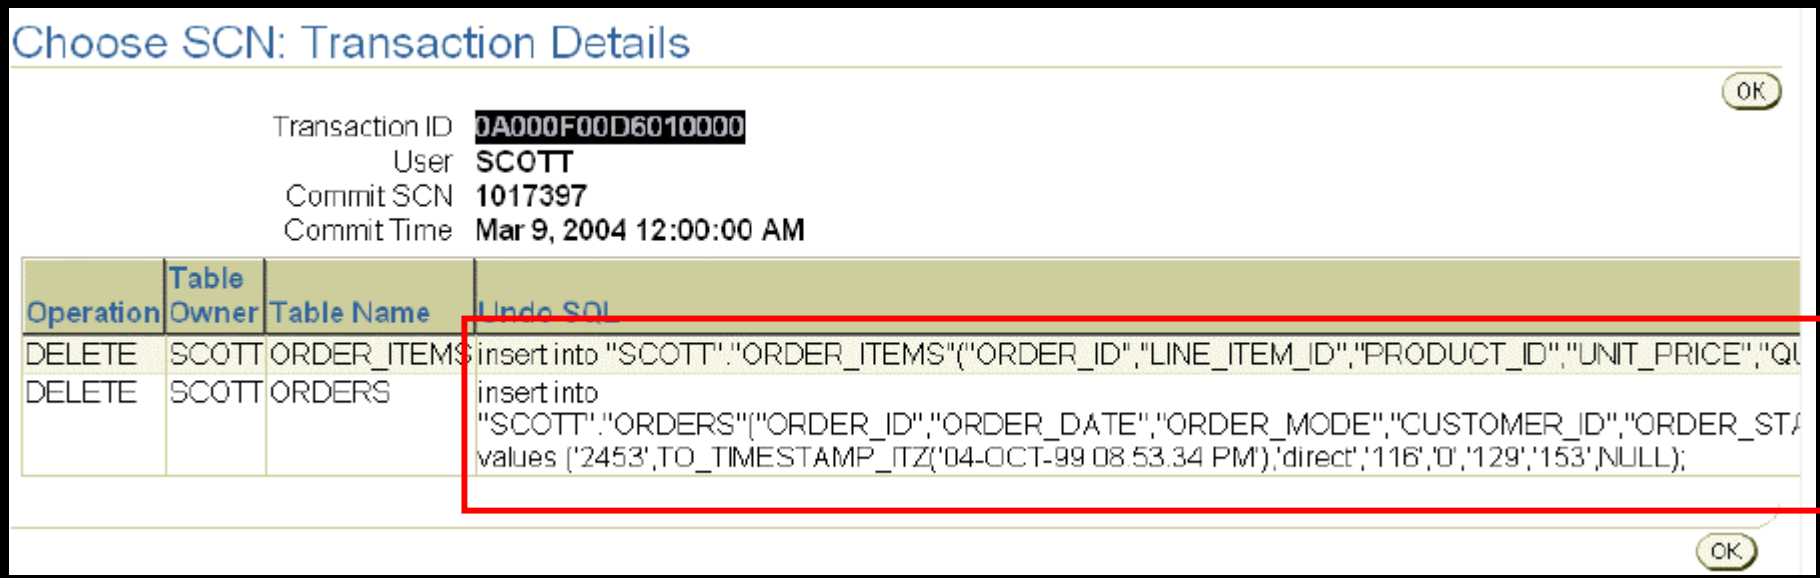

# Flashback Error Correction

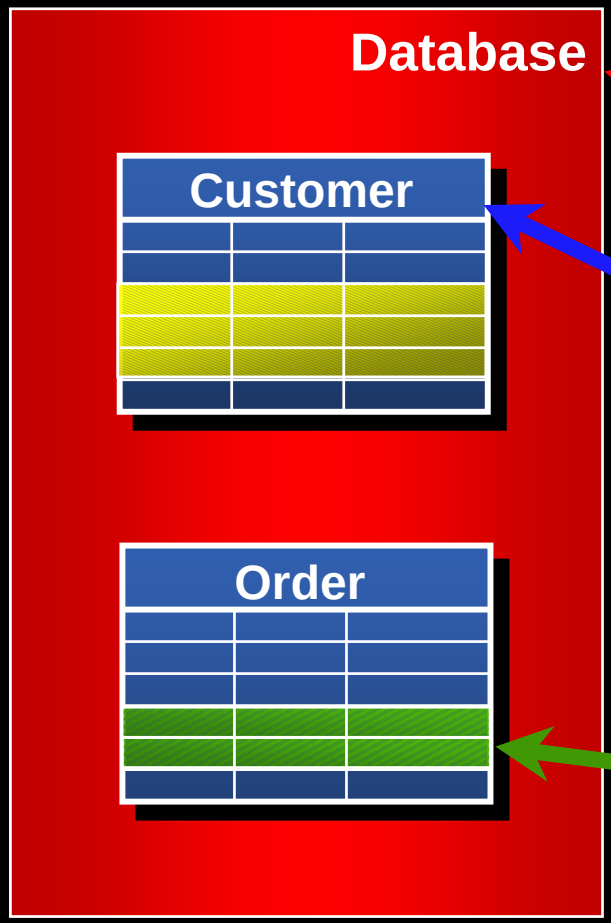


- Recovery at all levels
- Database Level
  - Flashback Database restores the whole database to time
    - Uses Flashback Logs
- Table Level
  - Flashback Table restores rows in a set of tables to time
    - Uses UNDO in database
  - Flashback Drop restores a dropped table or a index
    - Recycle bin for DROPs
- Row Level
  - Flashback Query restores rows to time

# Flashback Database

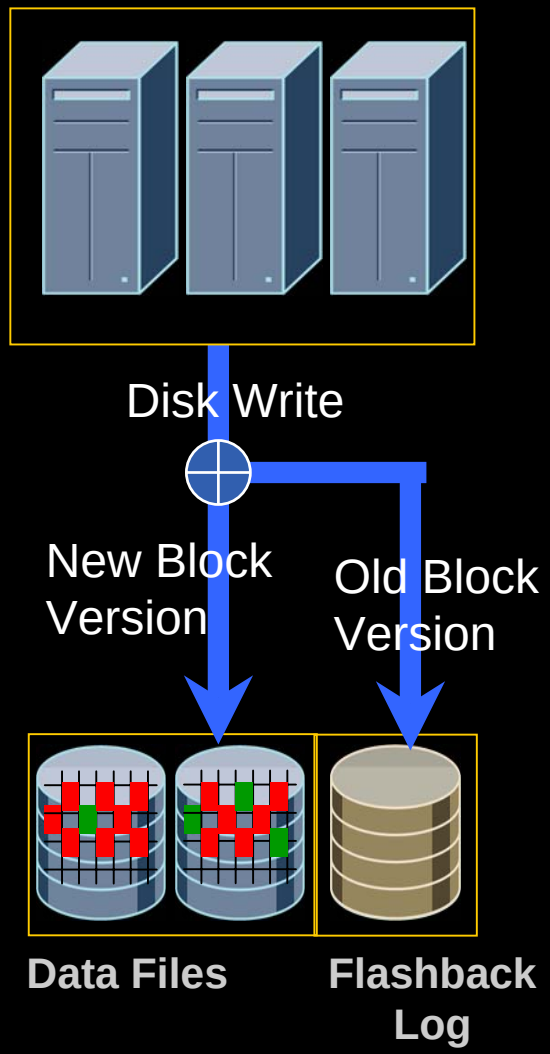


- A new strategy for point in time recovery
- Eliminate the need to restore a whole database backup
- Integrated seamlessly with RMAN
  - Think of it as a continuous backup
  - Restores just changed blocks
  - Replay log to restore DB to time
- It's fast - recover in minutes, not hours
- It's easy - single command restore

  Flashback Database to '2:05 PM'

**"Rewind" button for the Database**

# Flashback Database versus Classic Point-In-Time Recovery

***Recovery is 100 times faster with Flashback***

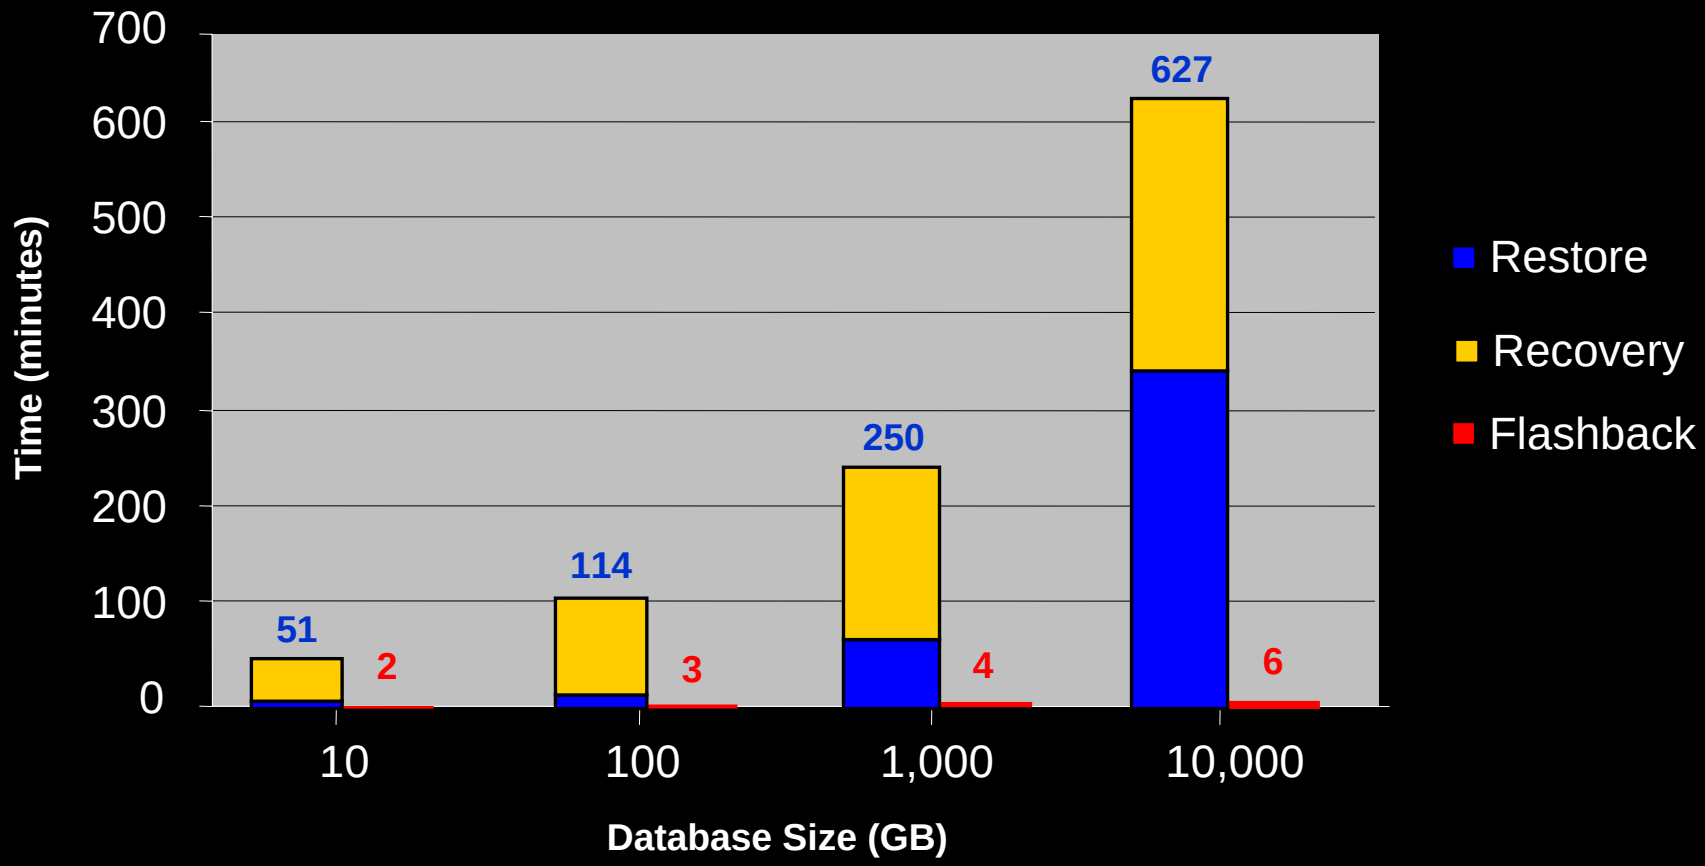

# Flashback Drop

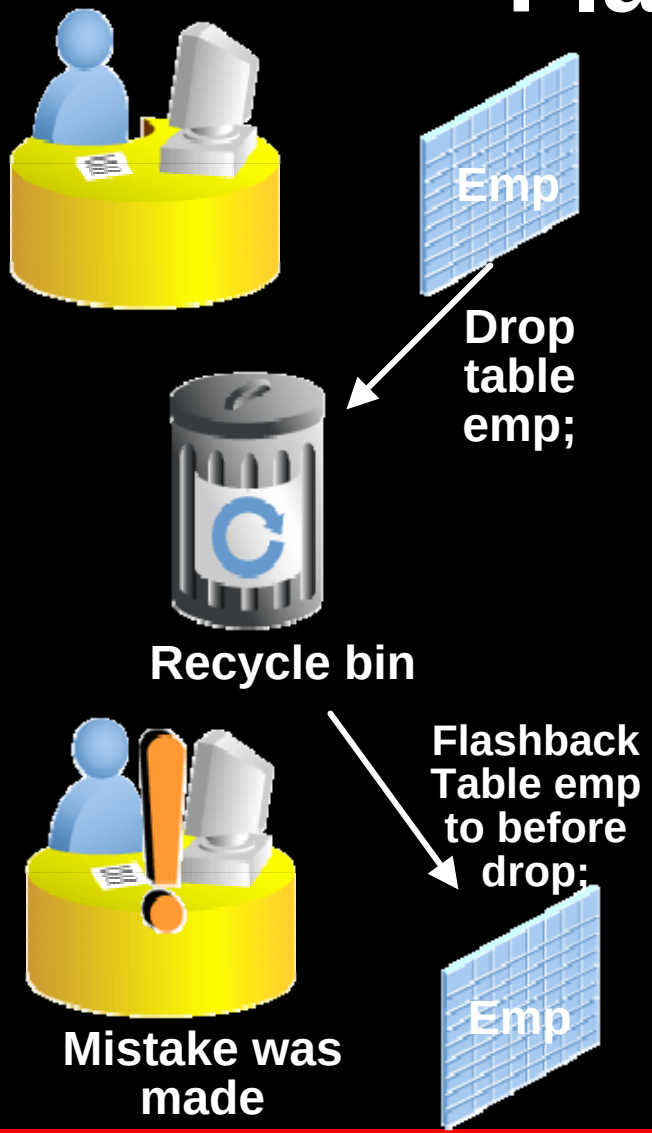


- Quickly recover dropped objects
  Provides self-service recovery
- Eliminate the need for TSPITR
- Virtual Recycle Bin
  - Objects remain in the recycle bin until you permanently drop them with the PURGE command or recover them with the Flashback Table command.
  - Objects will remain in the recycle bin until there is no room in the tablespace for new rows or updates to existing rows or until the tablespace needs to be extended
  - Objects are purged in the order they were dropped.

# Flashback Table

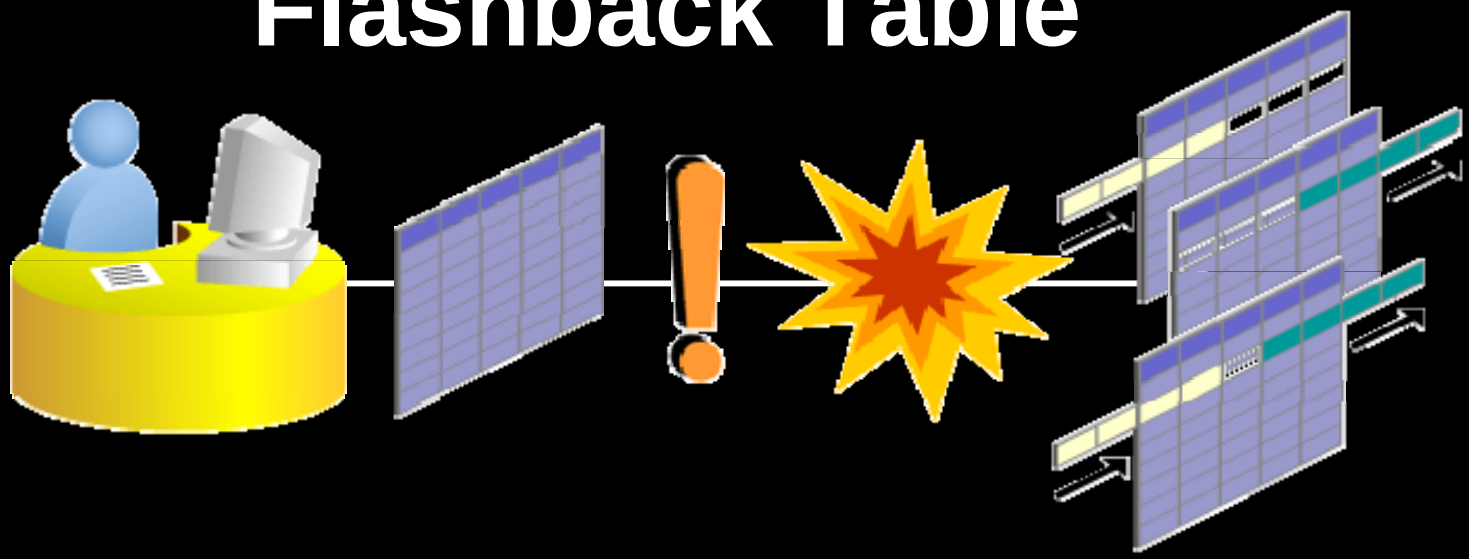

- Recover a table or tables to a specific point in time without restoring a backup
- Provides a way for users to easily and quickly recover from accidental modifications without DBA involvement
- In-place and online recovery of a table to a point in time in the past
- Eliminate traditional restores and clone instances to recover a table or tables to a specific point in time
- Data in the tables and all associated objects (indexes, constraints, triggers, etc.) are restored

# Revolution in Recovery

- Flashback Revolutionizes Recovery
  - Operates on just the changed data
  - Time to correct error equals time to make error
    - Minutes instead of hours

**Correction Time = Error Time + ~~*f*(DB_SIZE)~~**

- Flashback is Easy
  - Single command instead of complex procedure

# What is Oracle Data Guard?

- Oracle's disaster recovery solution for Oracle data
- Feature of Oracle Database Enterprise Edition (EE)
- Automates the creation and maintenance of one or more transactionally consistent copies (standby) of the production (or primary) database
- If the primary database becomes unavailable (disasters, maintenance), a standby database can be activated and assume the primary role

# Oracle Data Guard Focus

- Data Failures & Site Disasters:
  - Data Protection
  - Data Availability
  - Data Recovery

*All 3 are important!*

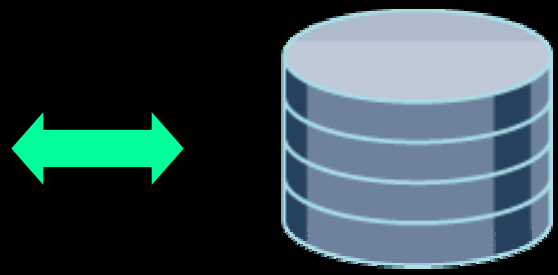

*Data is the core asset of the enterprise!*

- Also addresses human errors & planned maintenances

# Data Guard Configuration

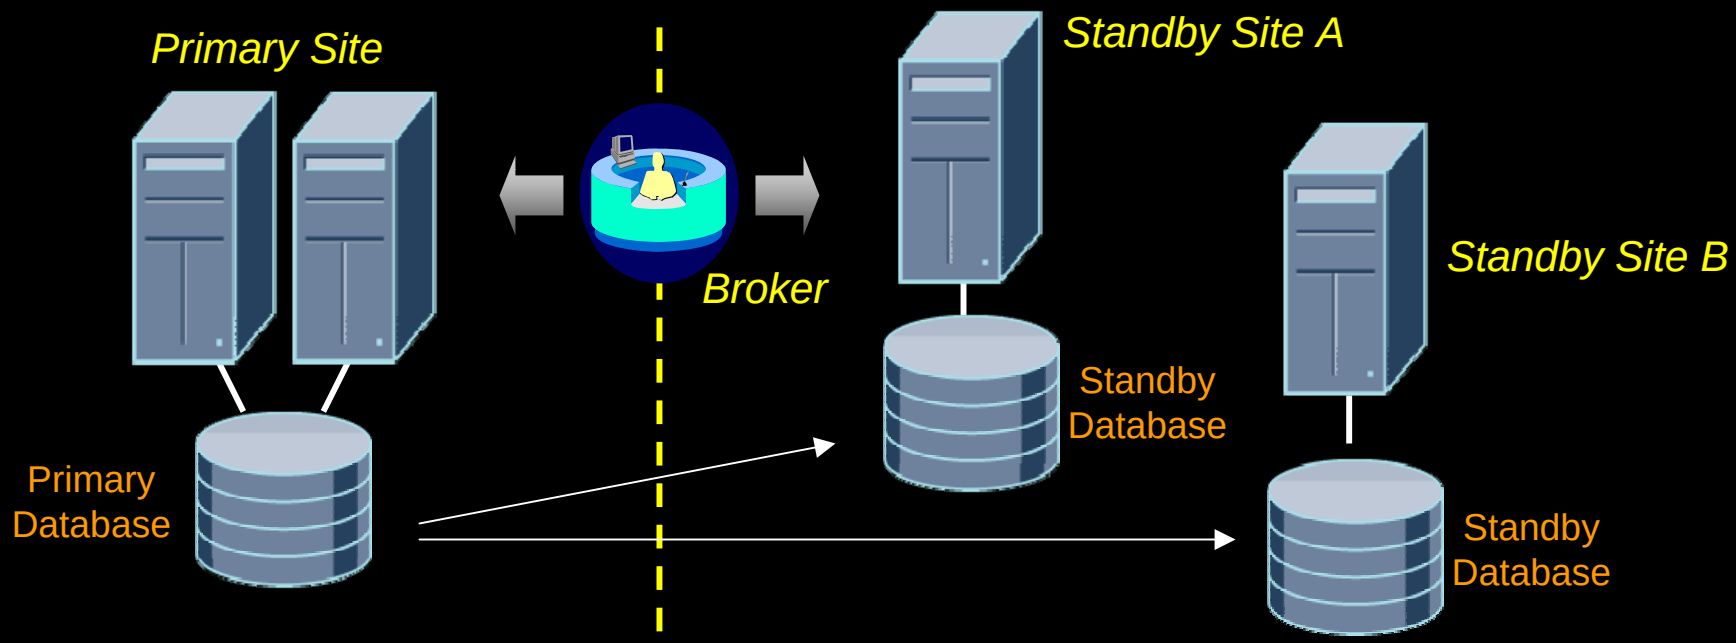


- Managed as a single configuration
- Primary and standby databases can be Real Application Clusters or single-instance Oracle
- Up to nine standby databases supported in a single configuration

# Oracle Data Guard Architecture

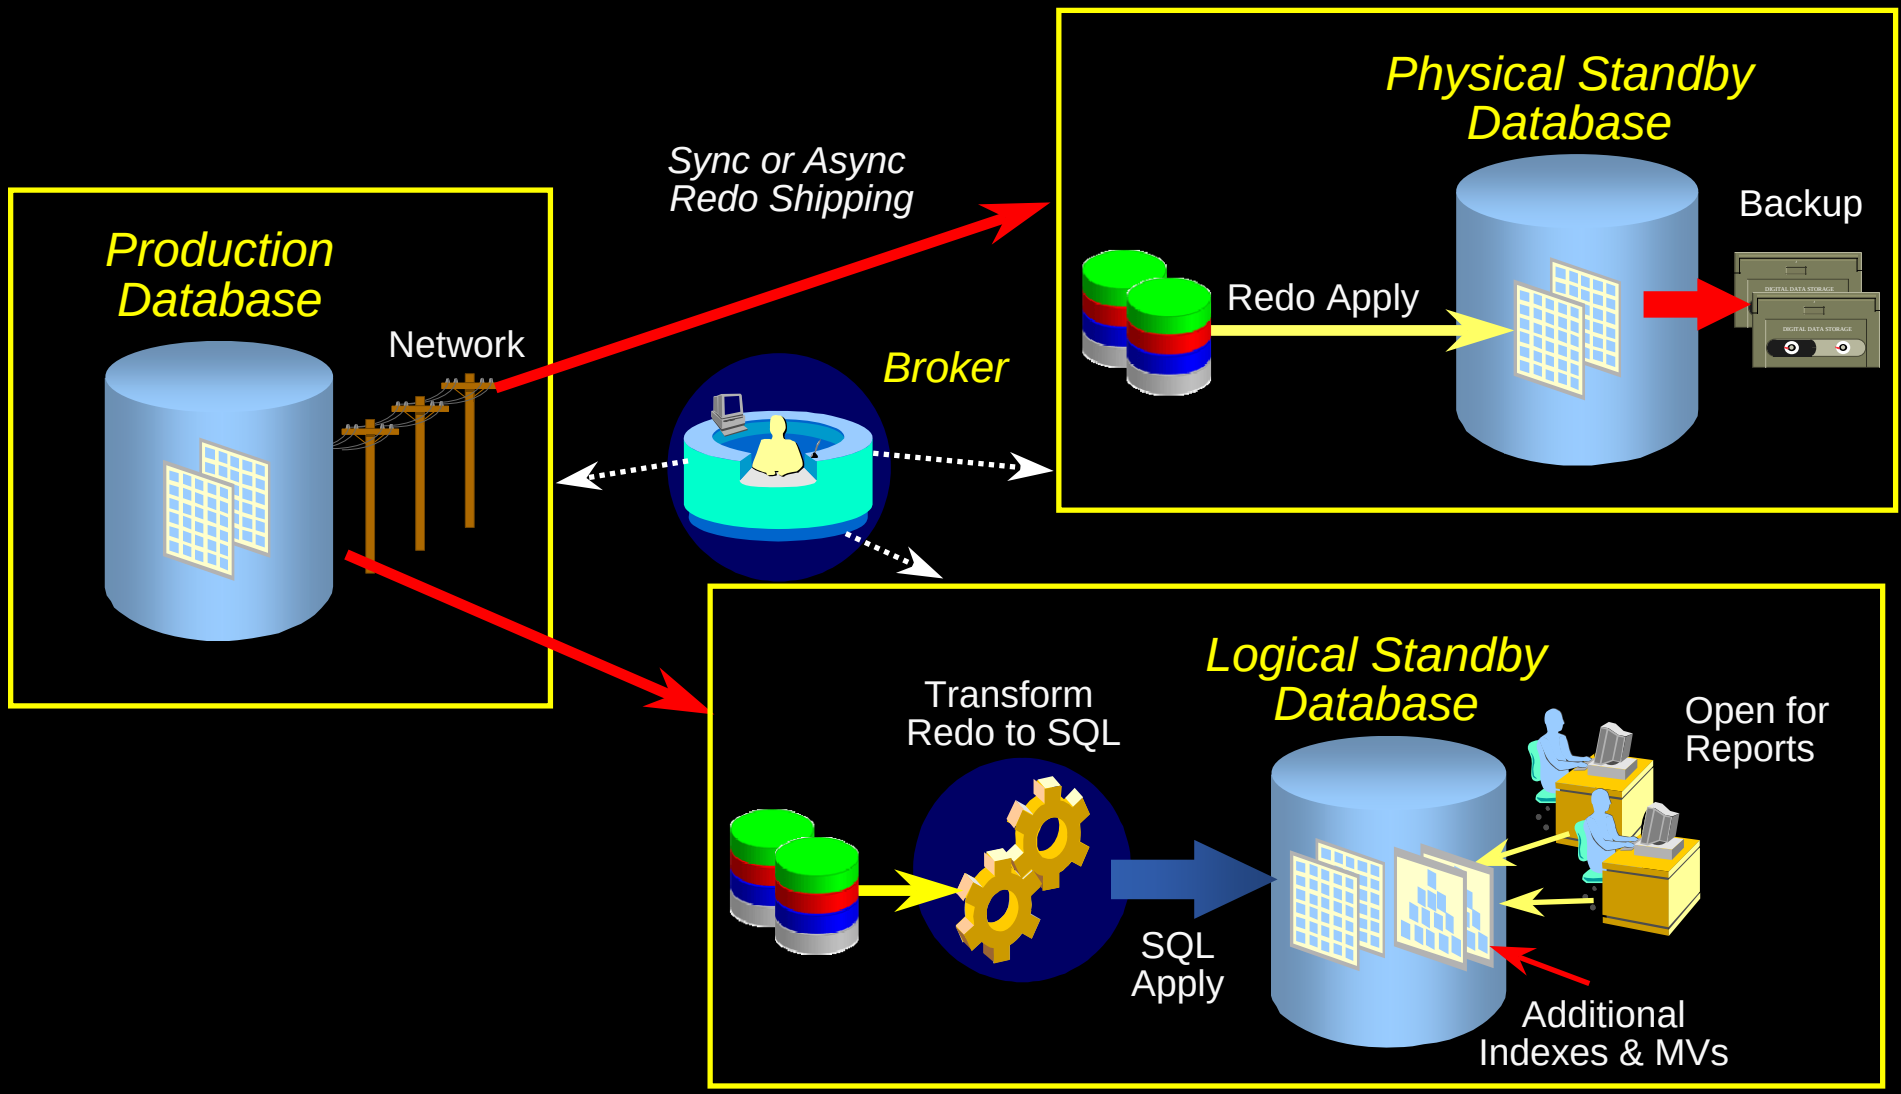

# Data Guard Redo Apply

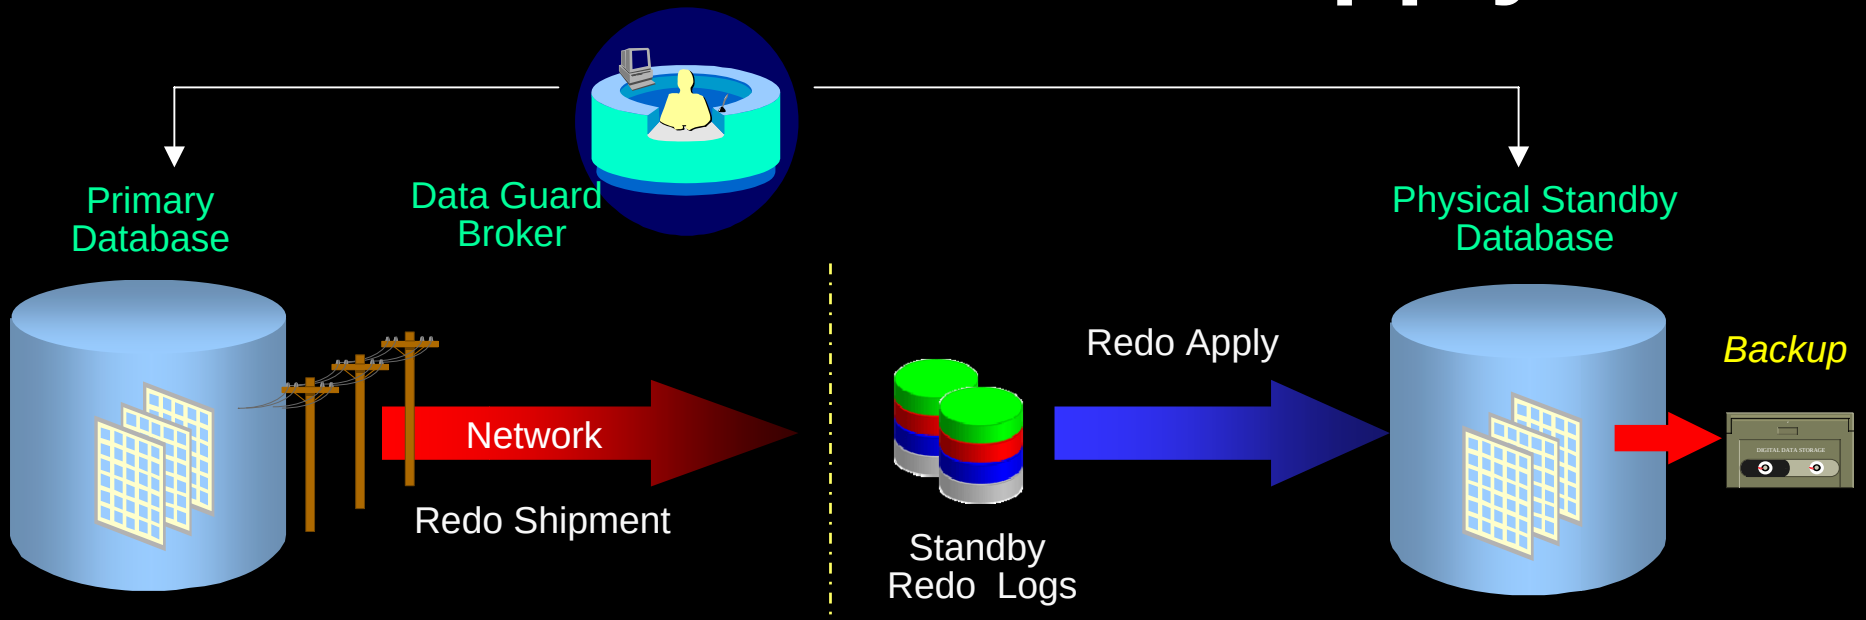


- Physical Standby Database is a block-for-block copy of the primary database
- Uses the database recovery functionality to apply changes
- Can be opened in read-only mode for reporting/queries
- Can also be used for backups, offloading production database

# Data Guard SQL Apply

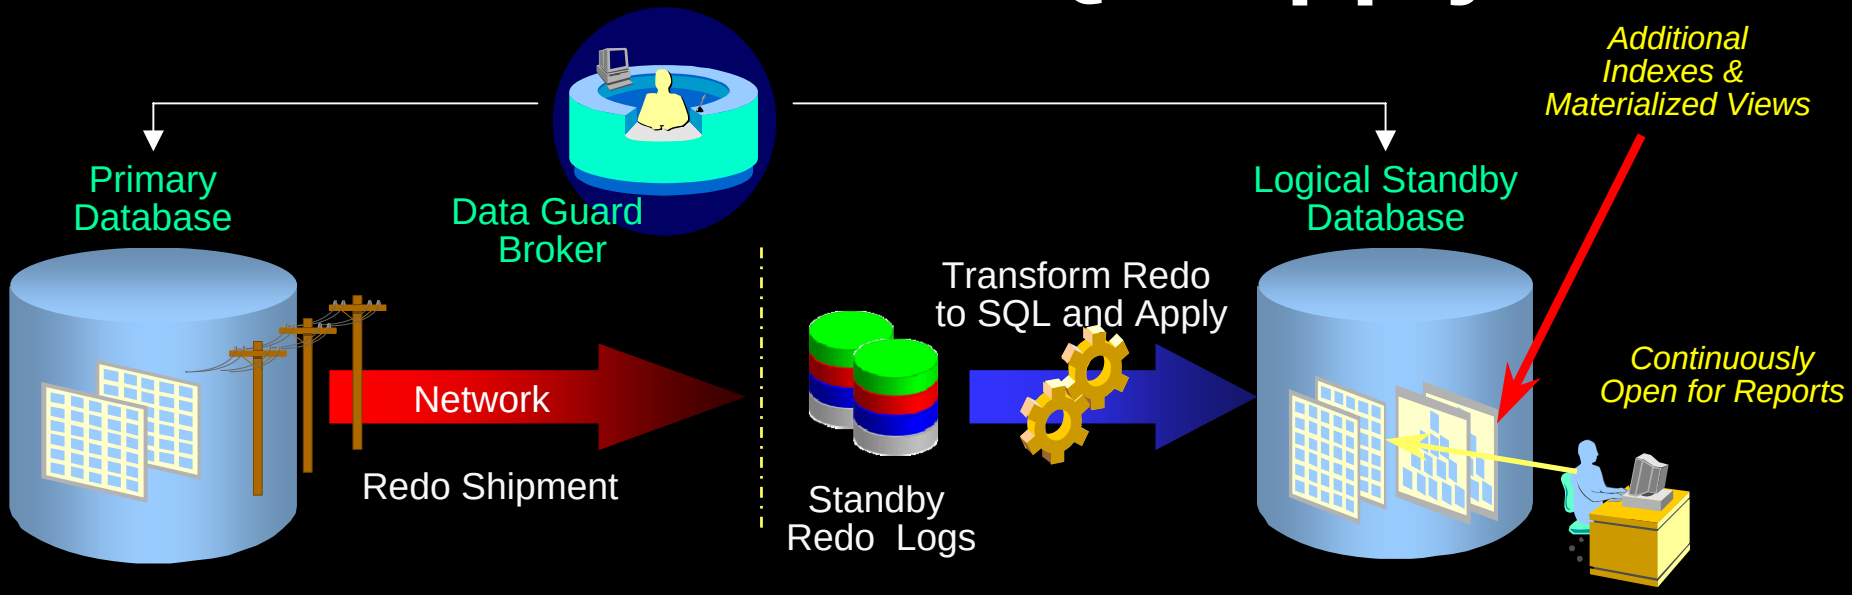


- Logical Standby Database is an open, independent, active database
  - Contains the same logical information (rows) as the production database
  - Physical organization and structure can be very different
  - Can host multiple schemas
- Can be queried for reports while logs are being applied via SQL
- Can create additional indexes and materialized views for better query performance

# Standby Databases Are Not Idle

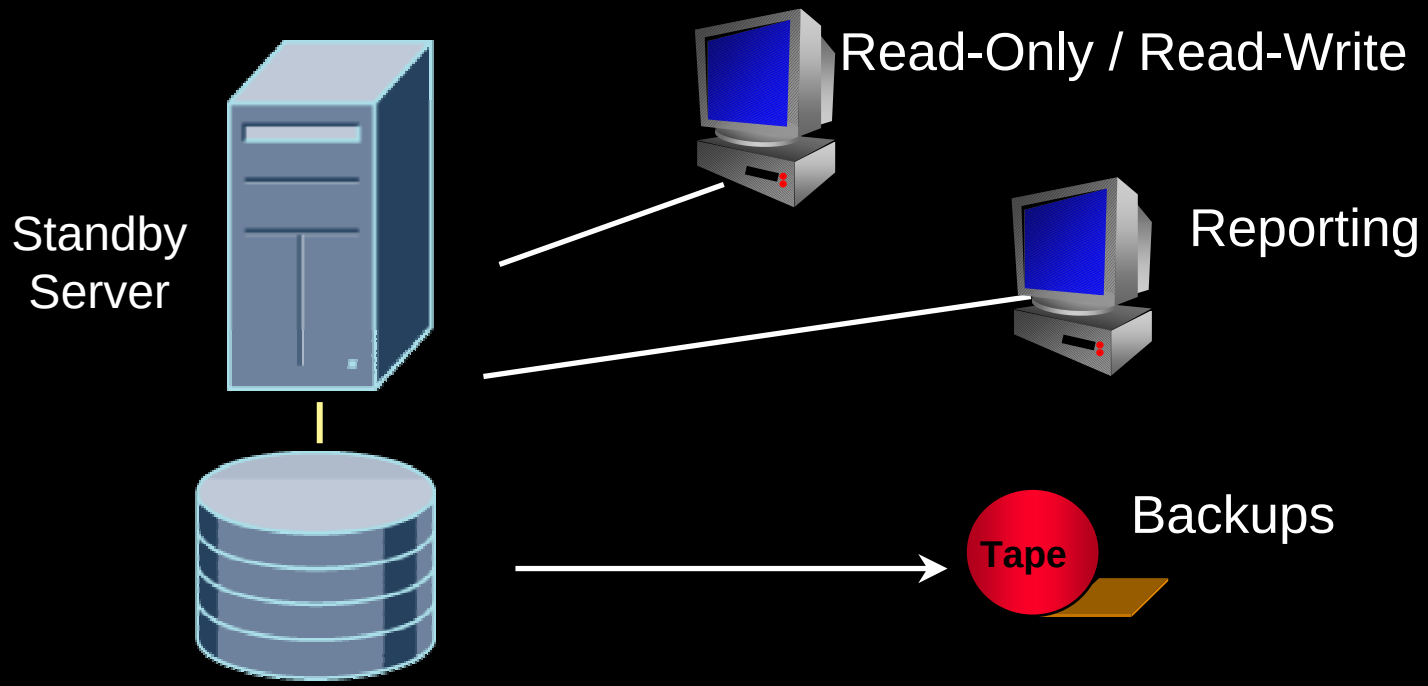


*Standby database can be used to offload the primary database, increasing the ROI*

# Protection from Human Errors and Data Corruptions

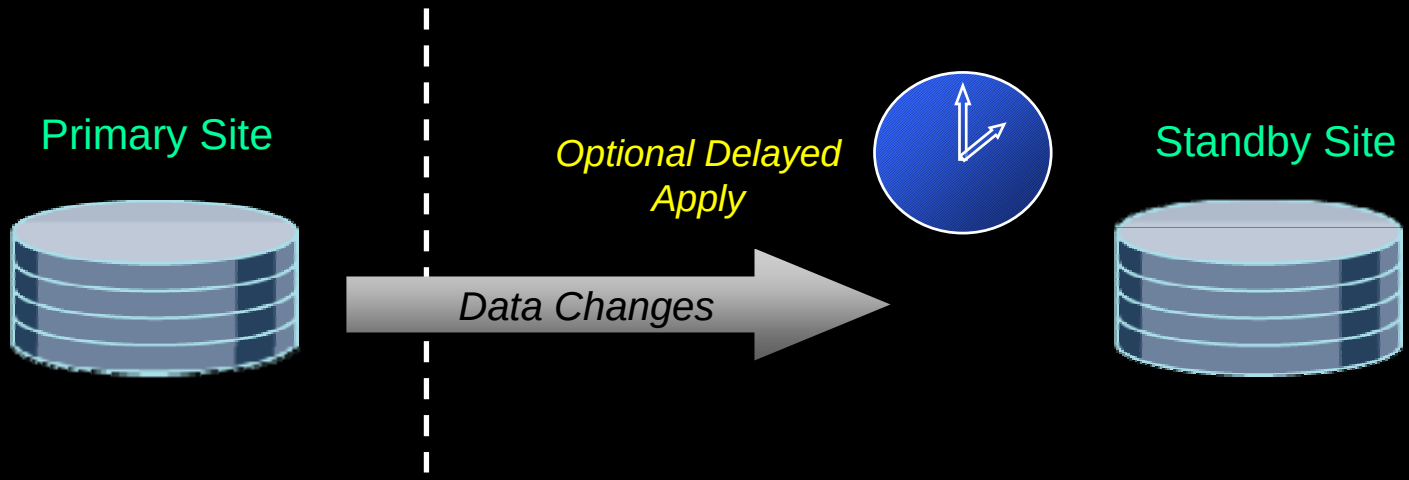


- Application of changes received from the primary can be delayed at standby to allow for the detection of user errors and prevent standby to be affected
- Administrators may choose not to configure any delay – if both primary and standby are affected, then they can be simply flashed back *[10g]*
- The apply process also revalidates the log records to prevent application of any log corruptions

# Switchover and Failover

- Primary and Standby role transitions
- Switchover
  - Planned role reversal
  - No database reinstantiation required
  - Used for maintenance of OS or hardware
- Failover
  - Unplanned failure (e.g. disasters) of primary
  - Primary database must be reinstantiated / flashed back *[10g]*
- Initiated using simple SQL / GUI interface
- Data Guard automates the processes involved

# Flexible Data Protection Modes

| Protection Mode | Risk of Data Loss | Redo Shipment |
|---|---|---|
| Maximum Protection | Zero Data Loss<br>Double Failure Protection | Synchronous redo shipping to 2 sites |
| Maximum Availability | Zero Data Loss<br>Single Failure Protection | Synchronous redo shipping |
| Maximum Performance | Minimal data loss – usually 0 to few seconds | Asynchronous redo shipping |

*Balance cost, availability, performance, and transaction protection*

# Automatic Resynchronization

- Network connectivity problems may occur
- Data Guard automatically resynchronizes standbys after network connectivity restored
  - Implicit
    - ARCH process idling away on the primary ‘pings’ all standbys on a regular basis to see if they are missing any redo data
    - If so it sends them the missing redo data
  - Explicit
    - Gap discovered during apply process in physical standby
    - Based on FAL_SERVER and FAL_CLIENT settings, primary notified, and it sends missing redo data

# Enhanced DR with Flashback Database

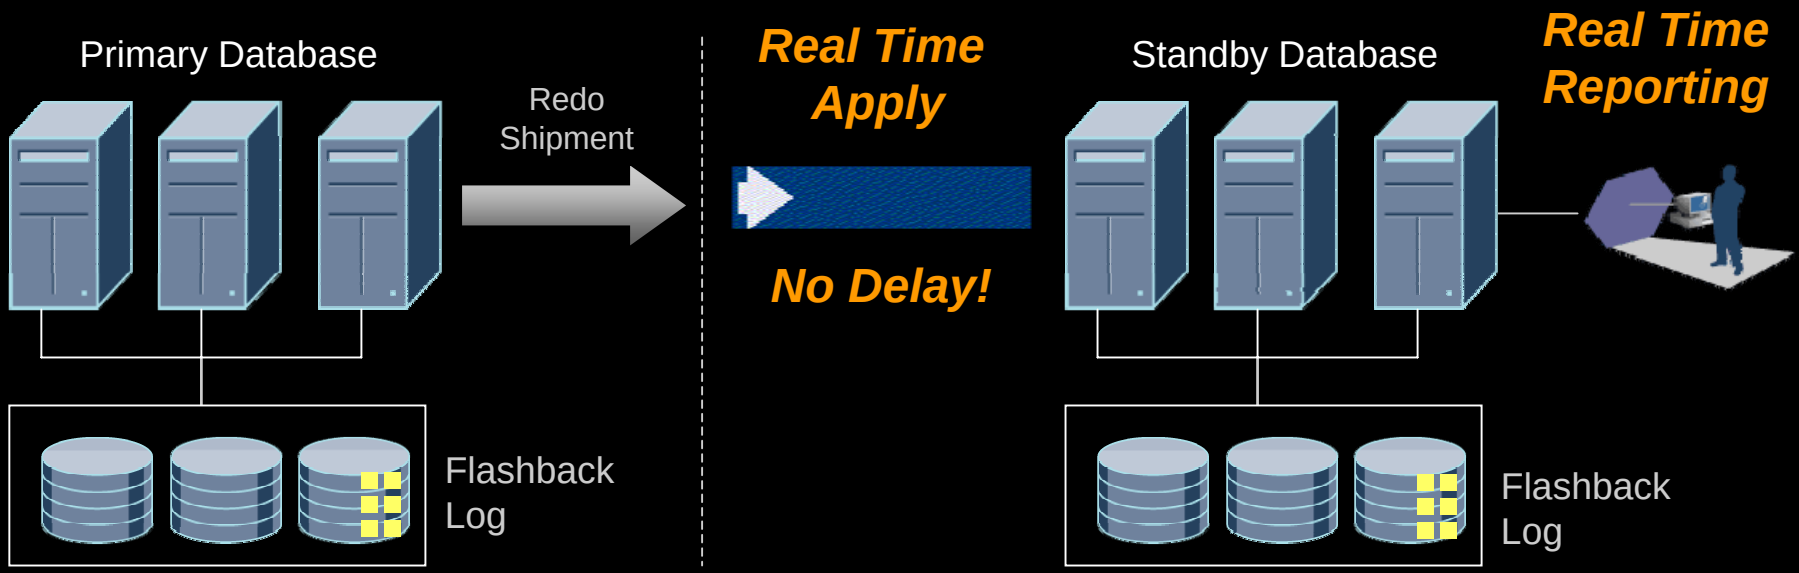


***Primary: No reinstantiation after failover!***

- Flashback DB removes the need to delay application of logs
- Flashback DB removes the need to reinstantiate primary after failover
- Real-time apply enables real-time reporting for logical standbys

# SQL Apply – Rolling Database Upgrades

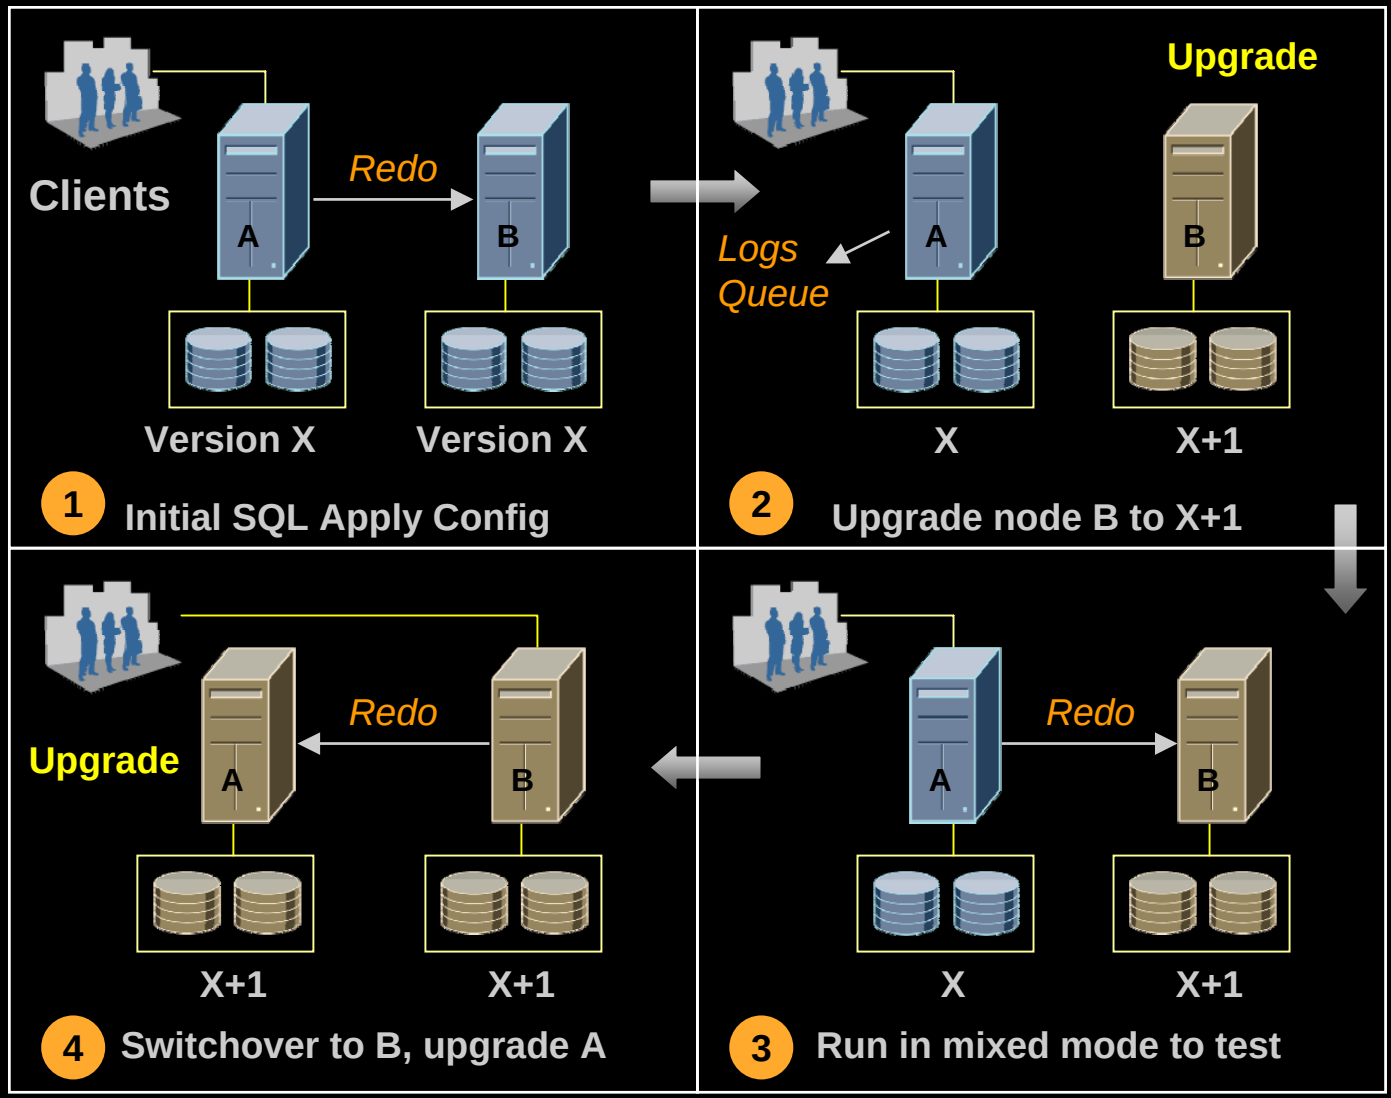


Patch Set Upgrades

Major Release Upgrades

Cluster Software & Hardware Upgrades

ORACLE 10g DATABASE

Oracle Enterprise Manager (SYSMAN) - Data Guard - Microsoft Internet Explorer

File Edit View Favorites Tools Help

ORACLE Enterprise Manager

Setup Preferences Help Logout

Home | Targets | Configuration | Alerts | Jobs | Management System

Hosts | Databases | Application Servers | Web Applications | Groups | All Targets

Cluster: drcluster > Cluster Database: San Francisco > Data Guard

# Data Guard

Page Refreshed August 1, 2003 5:27:25 PM EDT

## Overview

Overall Status ✔ **Normal**
Protection Mode Maximum Performance

## Primary Cluster Database

Name San Francisco
Cluster drcluster
Status ✔ Normal
Current Log Multiple Threads
Related Link Edit

*RAC Primary*

## Standby Progress Summary

This chart shows the amount of data that each standby has not yet received and applied.

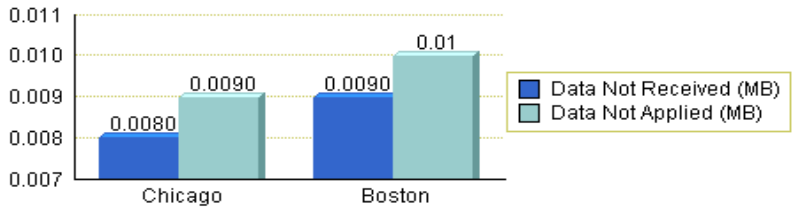


## Standby Databases

*Two standby dbs*

Add Standby Database | Edit | Remove | Switchover | Failover

| Select | Name | Host | Status | Role | Last Received Log | Last Applied Log |
|---|---|---|---|---|---|---|
| ◉ | Chicago | drlab3 | ✔ Normal | Physical Standby | Multiple Threads | Multiple Threads |
| ○ | Boston | drlab4 | ✔ Normal | Logical Standby | Multiple Threads | Multiple Threads |

## Performance

Performance Overview
Log File Details

## Additional Administration

Verify
Remove Data Guard Configuration

Home | **Targets** | Configuration | Alerts | Jobs | Management System | Setup | Preferences | Help | Logout


About Oracle Enterprise Manager

ORACLE 10g DATABASE

ORACLE

# Example – Ease of Use

- Switchover using Enterprise Manager is now literally two mouse clicks

Oracle Enterprise Manager (SYSMAN) - Data Guard - Microsoft Internet Explorer

File Edit View Favorites Tools Help

ORACLE Enterprise Manager

Setup Preferences Help Logout

Home Targets Configuration Alerts Jobs Management System

Hosts Databases Application Servers Web Applications Groups All Targets

Cluster: drcluster > Cluster Database: San Francisco > Data Guard

# Data Guard

Page Refreshed August 1, 2003 5:27:25 PM EDT

## Overview

Overall Status ✔ **Normal**
Protection Mode Maximum Performance

## Primary Cluster Database

Name San Francisco
Cluster drcluster
Status ✔ Normal
Current Log Multiple Threads
Related Link Edit

## Standby Progress Summary

This chart shows the amount of data that each standby has not yet received and applied.

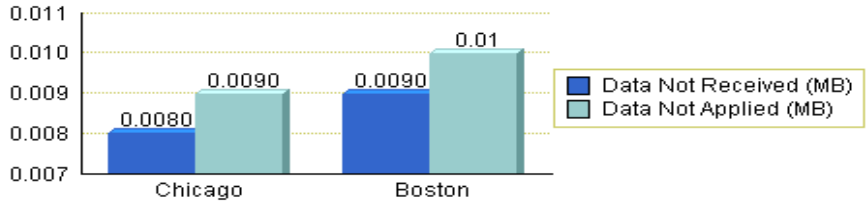


## Standby Databases

Add Standby Database | Edit | Remove | Switchover | Failover

| Select | Name | Host | Status | Role | Last Received Log | Last Applied Log |
|---|---|---|---|---|---|---|
| ◉ | Chicago | drlab3 | ✔ Normal | Physical Standby | Multiple Threads | Multiple Threads |
| ○ | Boston | drlab4 | ✔ Normal | Logical Standby | Multiple Threads | Multiple Threads |

## Performance

Performance Overview
Log File Details

## Additional Administration

Verify
Remove Data Guard Configuration

Home | **Targets** | Configuration | Alerts | Jobs | Management System | Setup | Preferences | Help | Logout


About Oracle Enterprise Manager


Oracle Enterprise Manager (SYSMAN) - Processing: Switchover - Microsoft Internet Explorer
File Edit View Favorites Tools Help
ORACLE
Enterprise Manager
Setup Preferences Help Logout
Home Targets Configuration Alerts Jobs Management System
Hosts Databases Application Servers Web Applications Groups All Targets
Oracle Enterprise Manager (SYSMAN) - Confirmation: Switchover to Chicago
A switchover will cause the primary and standby databases to switch roles. Since Chicago is a physical standby database, the primary and standby databases will be shutdown and restarted. The switchover operation cannot be cancelled.
Any active sessions connected to the primary database will be automatically closed during the switchover operation.
Browse Primary Database Sessions
TIP Standby databases not involved in the switchover will continue to function normally after the switchover.
Are you sure you want to switchover to Chicago?
No
Yes
Home | Targets | Configuration | Alerts | Jobs | Management System | Setup | Preferences | Help | Logout
Copyright © 1996, 2003, Oracle. All rights reserved.
About Oracle Enterprise Manager


Oracle Enterprise Manager (SYSMAN) - Processing: Switchover - Microsoft Internet Explorer
File Edit View Favorites Tools Help
ORACLE Enterprise Manager
Setup Preferences Help Logout
Home Targets Configuration Alerts Jobs Management System
Hosts Databases Application Servers Web Applications Groups All Targets
Processing: Switchover
Switching over to Chicago
This process will take some time. The page will automatically forward to the overview page upon completion.
Click on the alert log link to view progress details in a new browser window. View alert log: San Francisco Chicago
Performing role change.
Restarting databases.
Waiting for switchover to complete.
TIP This process cannot be cancelled. It will continue even if the browser window is closed.
Home | Targets | Configuration | Alerts | Jobs | Management System | Setup | Preferences | Help | Logout
Copyright © 1996, 2003, Oracle. All rights reserved.
About Oracle Enterprise Manager


Oracle Enterprise Manager (SYSMAN) - Processing: Switchover - Microsoft Internet Explorer
File Edit View Favorites Tools Help
ORACLE Enterprise Manager
Setup Preferences Help Logout
Home Targets Configuration Alerts Jobs Management System
Hosts Databases Application Servers Web Applications Groups All Targets
Processing: Switchover
Switching over to Chicago
This process will take some time. The page will automatically forward to the overview page upon completion.
Click on the alert log link to view progress details in a new browser window. View alert log: San Francisco Chicago
Performing role change.
Restarting databases.
Waiting for switchover to complete.
TIP This process cannot be cancelled. It will continue even if the browser window is closed.
Home | Targets | Configuration | Alerts | Jobs | Management System | Setup | Preferences | Help | Logout
Copyright © 1996, 2003, Oracle. All rights reserved.
About Oracle Enterprise Manager


Oracle Enterprise Manager (SYSMAN) - Processing: Switchover - Microsoft Internet Explorer
File Edit View Favorites Tools Help
ORACLE
Enterprise Manager
Setup Preferences Help Logout
Home Targets Configuration Alerts Jobs Management System
Hosts Databases Application Servers Web Applications Groups All Targets
Processing: Switchover
Switching over to Chicago
This process will take some time. The page will automatically forward to the overview page upon completion.
Click on the alert log link to view progress details in a new browser window. View alert log: San Francisco Chicago
Performing role change.
Restarting databases.
Waiting for switchover to complete.
TIP This process cannot be cancelled. It will continue even if the browser window is closed.
Home | Targets | Configuration | Alerts | Jobs | Management System | Setup | Preferences | Help | Logout
Copyright © 1996, 2003, Oracle. All rights reserved.
About Oracle Enterprise Manager

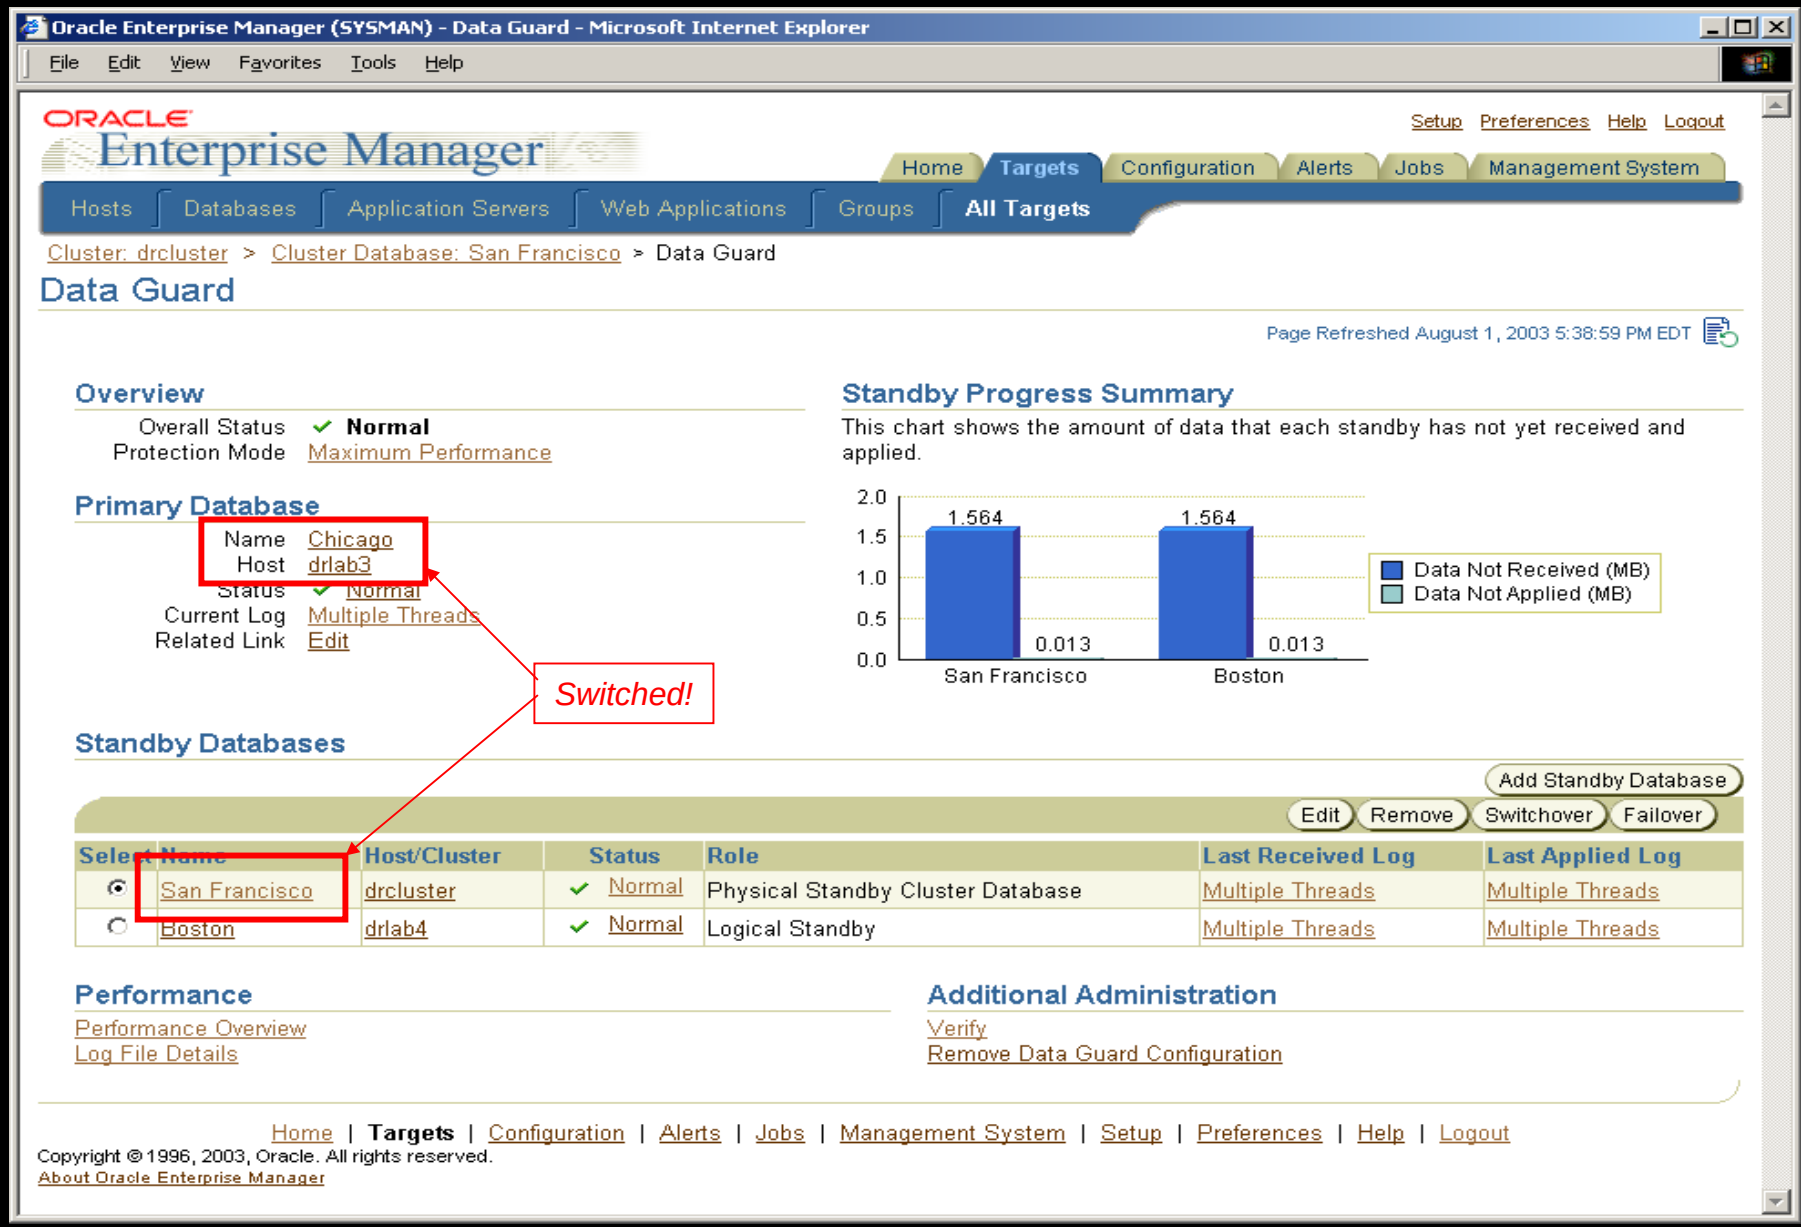
Oracle Enterprise Manager (SYSMAN) - Data Guard - Microsoft Internet Explorer
File Edit View Favorites Tools Help
ORACLE Enterprise Manager
Setup Preferences Help Logout
Home Targets Configuration Alerts Jobs Management System
Hosts Databases Application Servers Web Applications Groups All Targets
Cluster: drcluster > Cluster Database: San Francisco > Data Guard
Data Guard
Page Refreshed August 1, 2003 5:38:59 PM EDT
Overview
Overall Status Normal
Protection Mode Maximum Performance
Primary Database
Name Chicago
Host drlab3
Status Normal
Current Log Multiple Threads
Related Link Edit
Switched!
Standby Progress Summary
This chart shows the amount of data that each standby has not yet received and applied.
2.0
1.5
1.0
0.5
0.0
1.564
0.013
1.564
0.013
San Francisco
Boston
Data Not Received (MB)
Data Not Applied (MB)
Standby Databases
Add Standby Database
Edit Remove Switchover Failover
Select Name Host/Cluster Status Role Last Received Log Last Applied Log
San Francisco drcluster Normal Physical Standby Cluster Database Multiple Threads Multiple Threads
Boston drlab4 Normal Logical Standby Multiple Threads Multiple Threads
Performance
Performance Overview
Log File Details
Additional Administration
Verify
Remove Data Guard Configuration
Home | Targets | Configuration | Alerts | Jobs | Management System | Setup | Preferences | Help | Logout
Copyright © 1996, 2003, Oracle. All rights reserved.
About Oracle Enterprise Manager

# Data Guard Customers

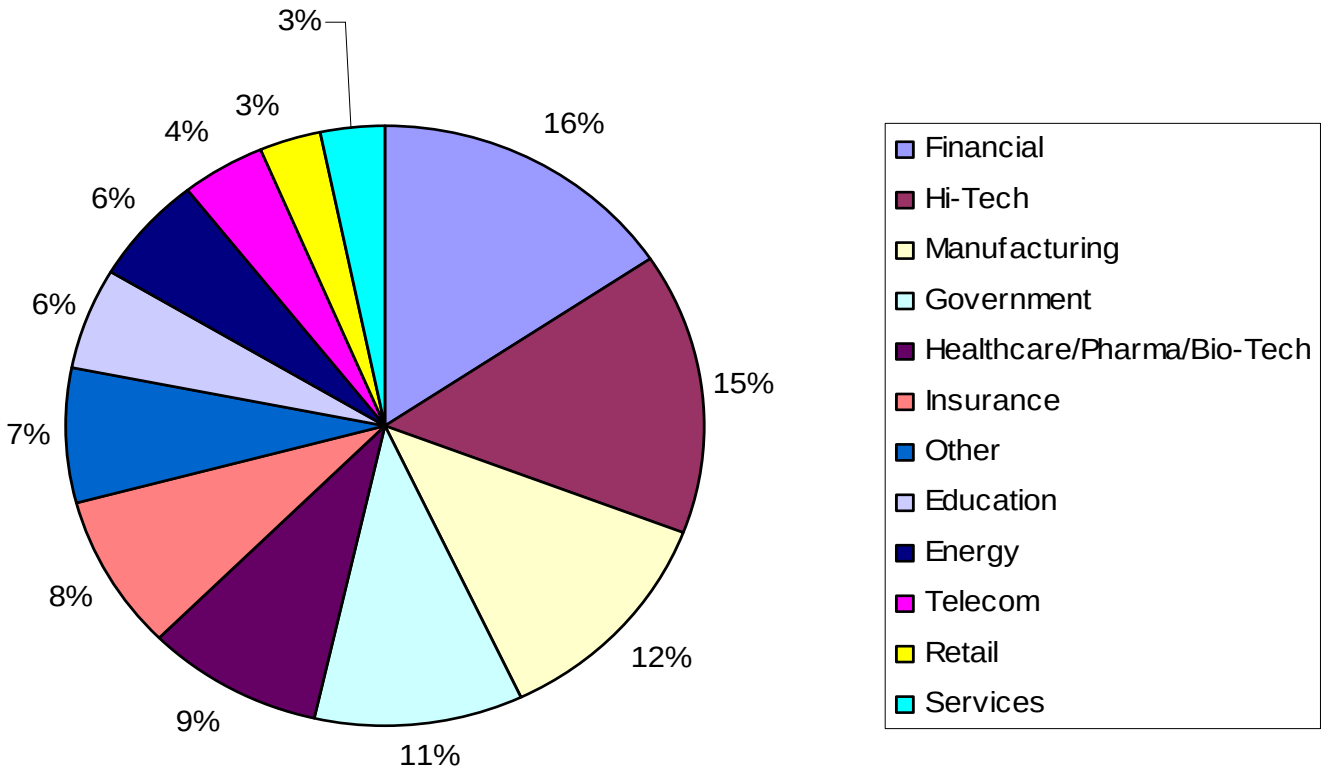

# Data Guard Technical Case Studies

- **ADT Security Services** - Using Data Guard SQL Apply Across a Wide Area Network
- **Amadeus** - Using Data Guard for Disaster Recovery & Rolling Database Upgrades
- **Fannie Mae** - Supporting 835 transactions per second & Zero Data Loss Protection in Oracle Database 10g
- **First American Real Estate Solutions** - Using Oracle9i Data Guard and Planning ahead for Data Guard in Oracle Database 10g
- **Ohio Savings Bank** - Maximum Availability Architecture & Zero Data Loss with Oracle Database 10g
- **Oracle Global IT** - Oracle E-Business Suite with Data Guard over a WAN
- **Swedish Post** - SQL Apply
- **VP Bank** - SQL Apply

*Ref. http://www.oracle.com/technology/deploy/availability/htdocs/HA_CaseStudies.html for latest updates*

# Oracle's Integrated HA Solution Set

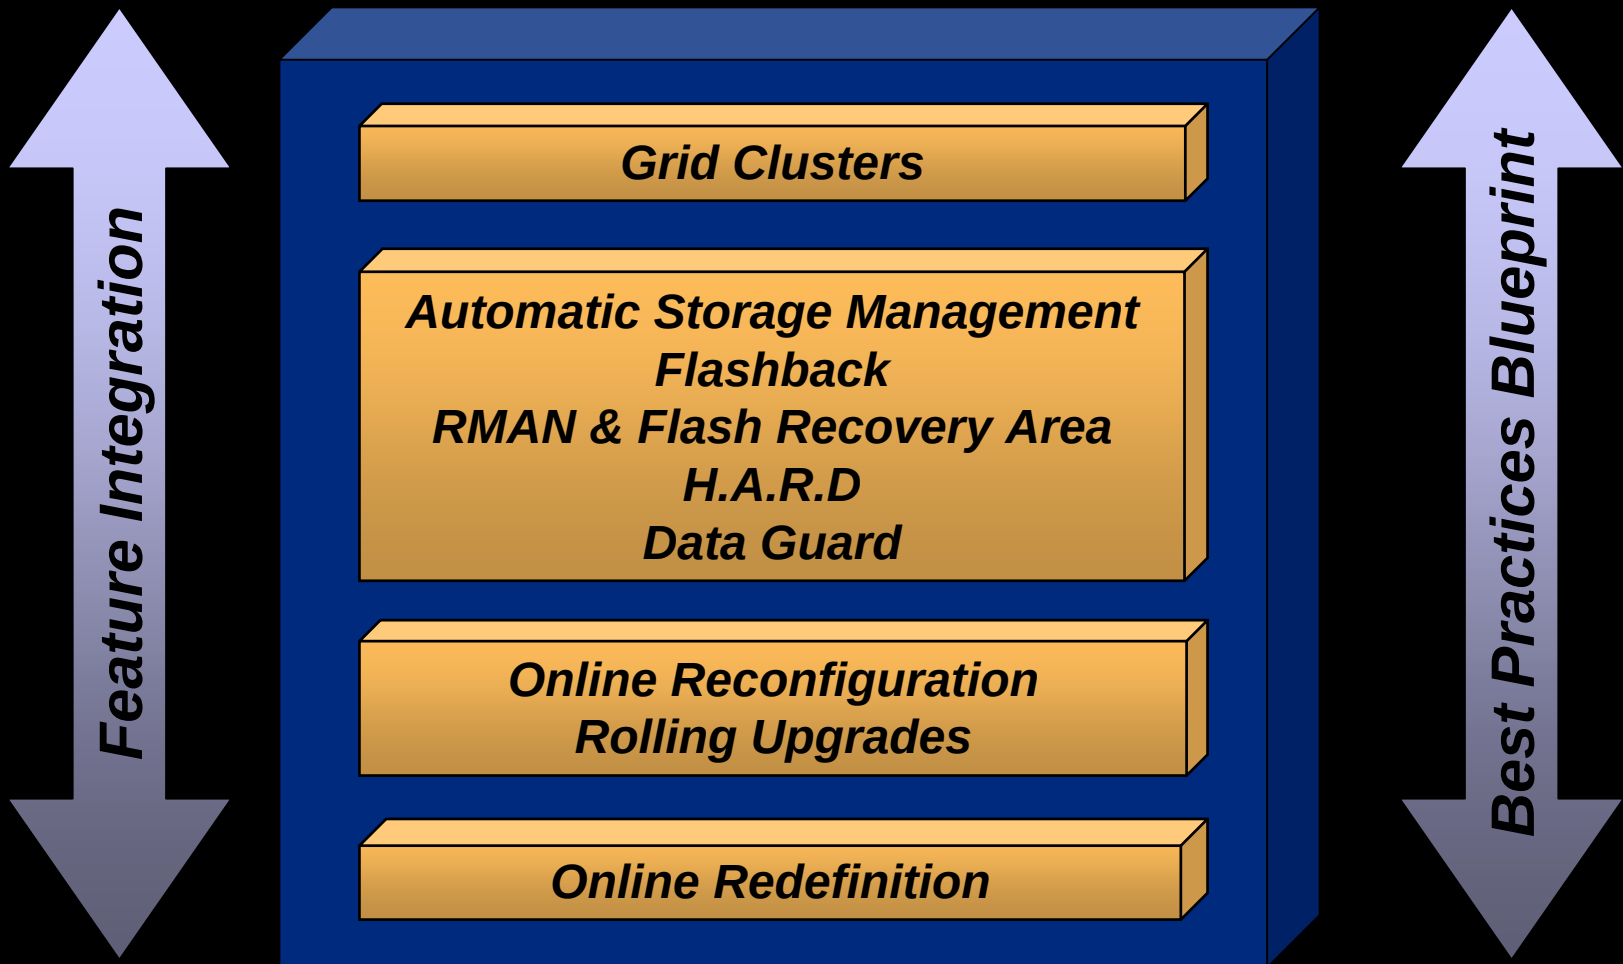

# MAA Best Practice Publications

- Best Practices on:
  - ✓ RAC/ Data Guard configuration
  - ✓ Redo data transport mechanisms
  - ✓ Instance Recovery
  - ✓ Switchover/Failover
  - ✓ Media recovery
  - ✓ SQL Apply configuration
  - ✓ Network configuration
  - ✓ Integration of HA technologies

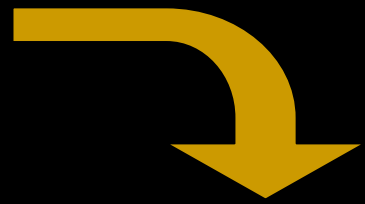

- White papers[1]:
  - ✓ MAA – detailed
  - ✓ Media Recovery
  - ✓ Site/Network configuration
  - ✓ Fast-Start Checkpointing
  - ✓ SQL Apply Best Practices
  - ✓ Role Management

*1. Ref. http://otn.oracle.com/deploy/availability/htdocs/maa.htm for latest updates*

# QUESTIONS & ANSWERS

For more information on Oracle database High Availability, Disaster Protection, Backup & Recovery, and Storage Management technology

http://otn.oracle.com/deploy/availability/