ORACLE

Oracle Database 11g: Next Generation High Availability

Juan Loaiza
Senior Vice President, Oracle Corporation

Traditional Database HA

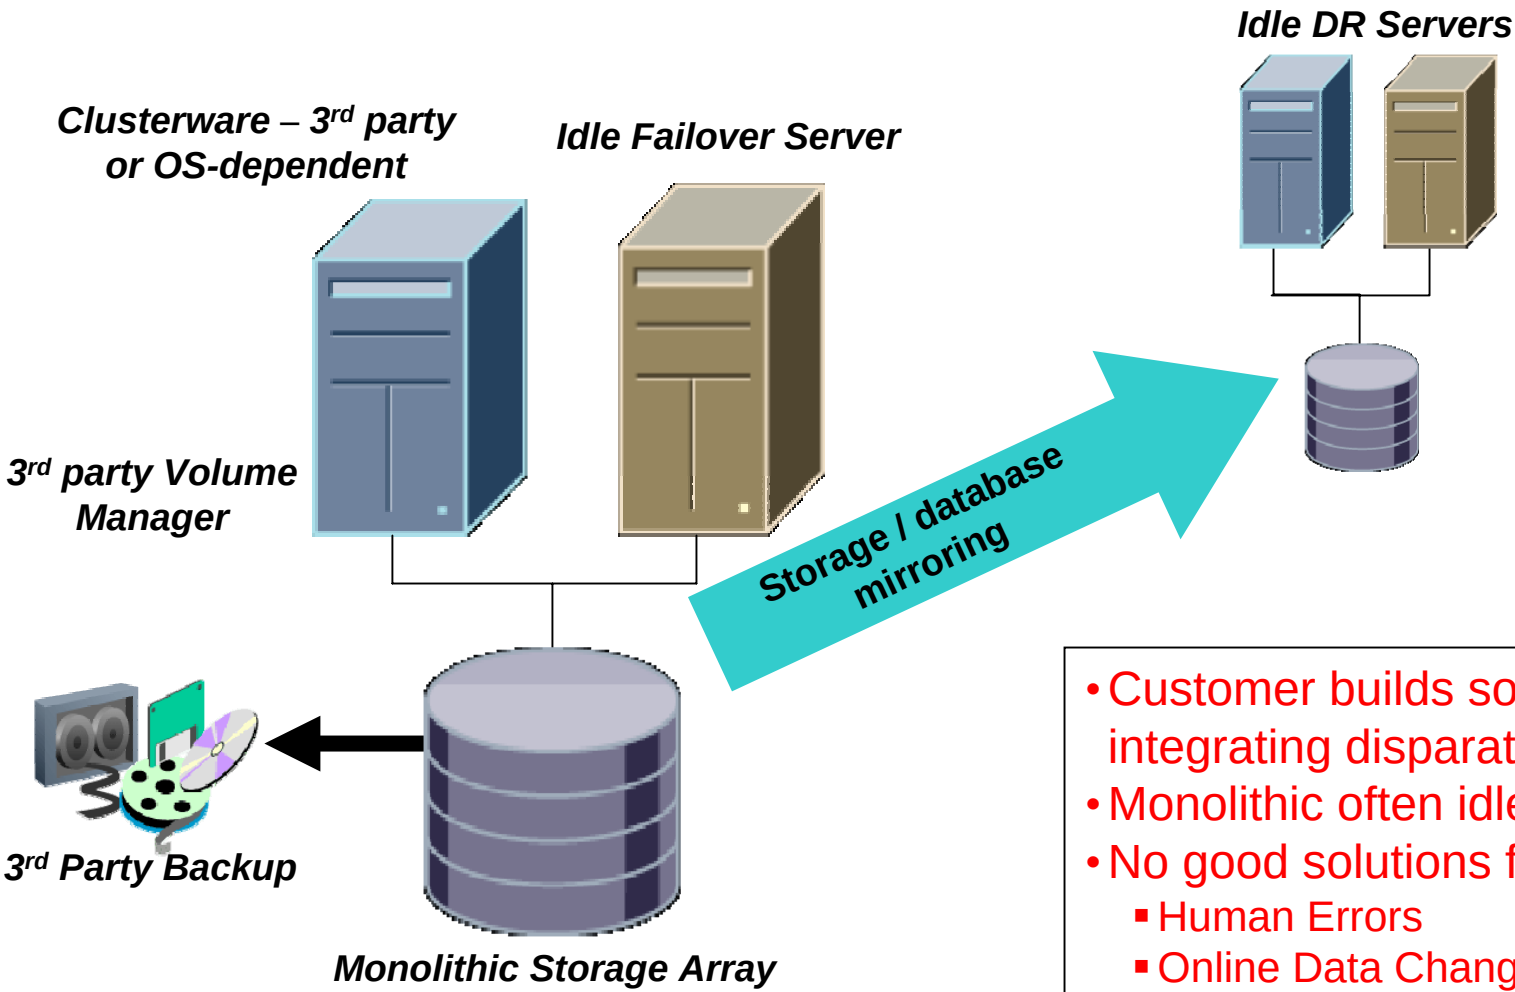

- Customer builds solution by integrating disparate components
- Monolithic often idle hardware
- No good solutions for:
  - Human Errors
  - Online Data Changes
  - Software Upgrades

The Scale-Out Era

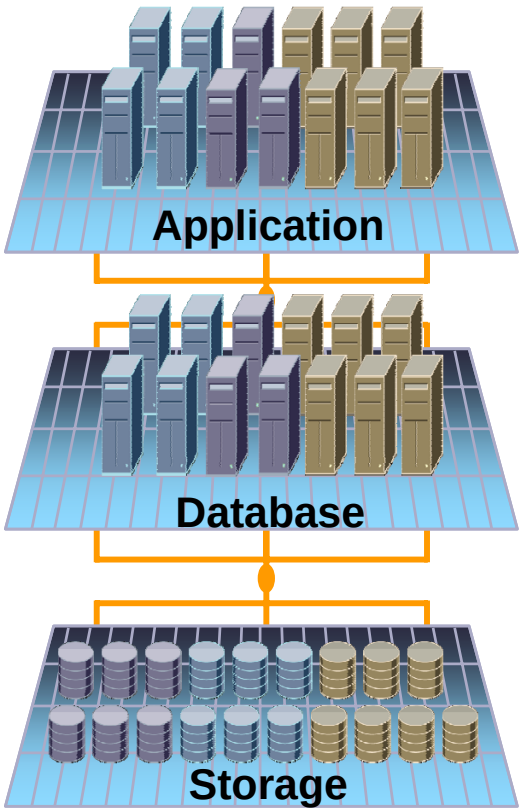

- Scale-Out architecture
  - High volume hardware building blocks
  - Inherently highly scalable & redundant
- Scalability & Availability responsibility moves out of hardware/OS to scale-out savvy software
  - First Web & Application server tiers
    - Application servers
  - Then DB tier
    - Shared disk and shared nothing databases
  - Then storage tier
    - Scale-out savvy storage software

Oracle's Innovative Approach Breaks Tradeoff Between Availability and Cost

Best Availability AND Lowest Cost

- Better than Mainframe Availability
- PC Economics
- Seamless and Simple to Use

Best Server Protection - at Lowest Cost

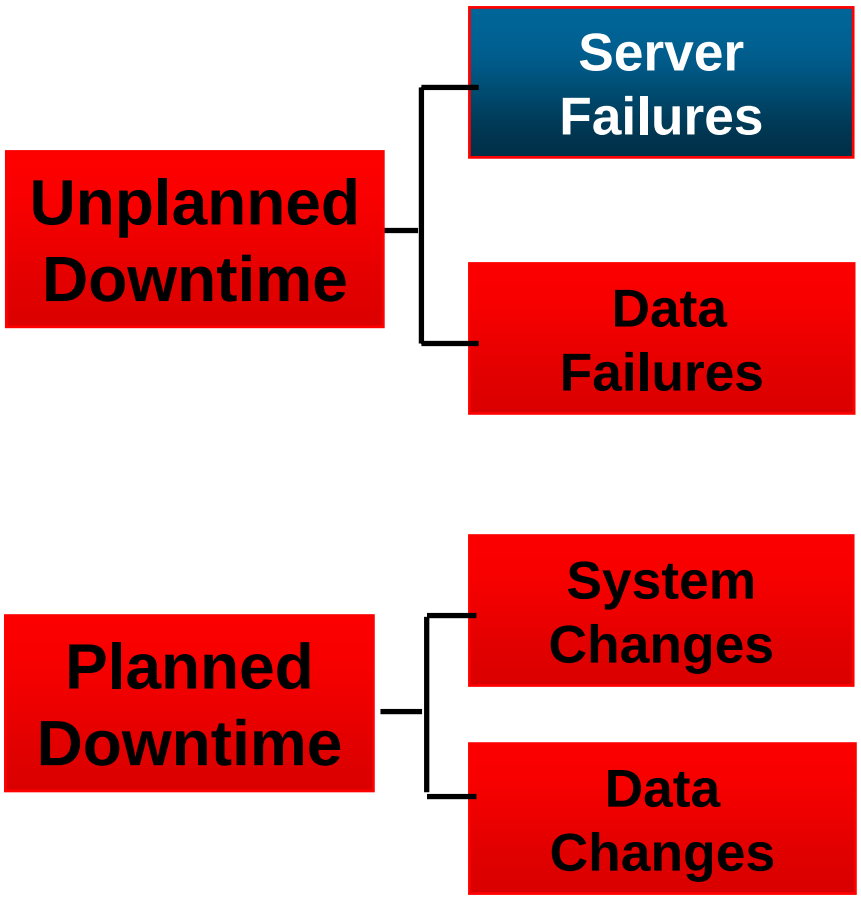

Real Application Clusters

Server Scale-Out with RAC

- RAC pools standard low cost servers
- Great Scalability & Availability
  - No Idle Resources
- Integrated clusterware for free
- Runs commercial applications
  - Oracle Applications, SAP, etc.
- Thousands of production customers

ORACLE 11g DATABASE

- 11g fine tunes performance, scaling, failover, management

Designed to Tolerate Server Failures

Best Data Protection - at Lowest Cost

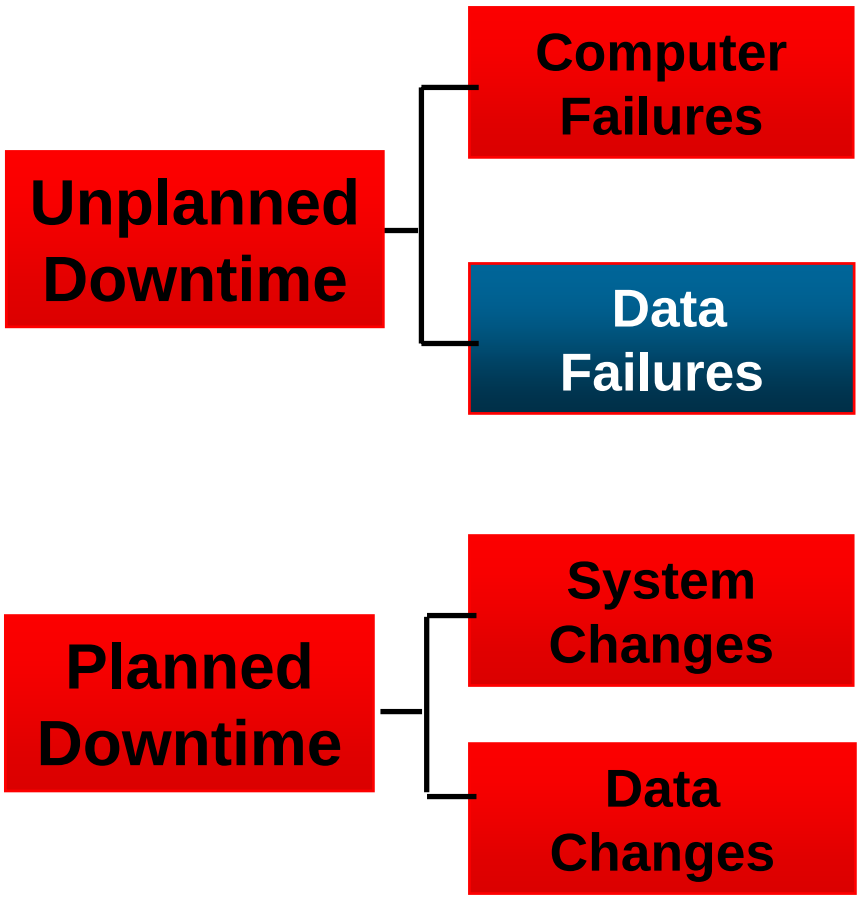

Best Storage Protection - at Lowest Cost

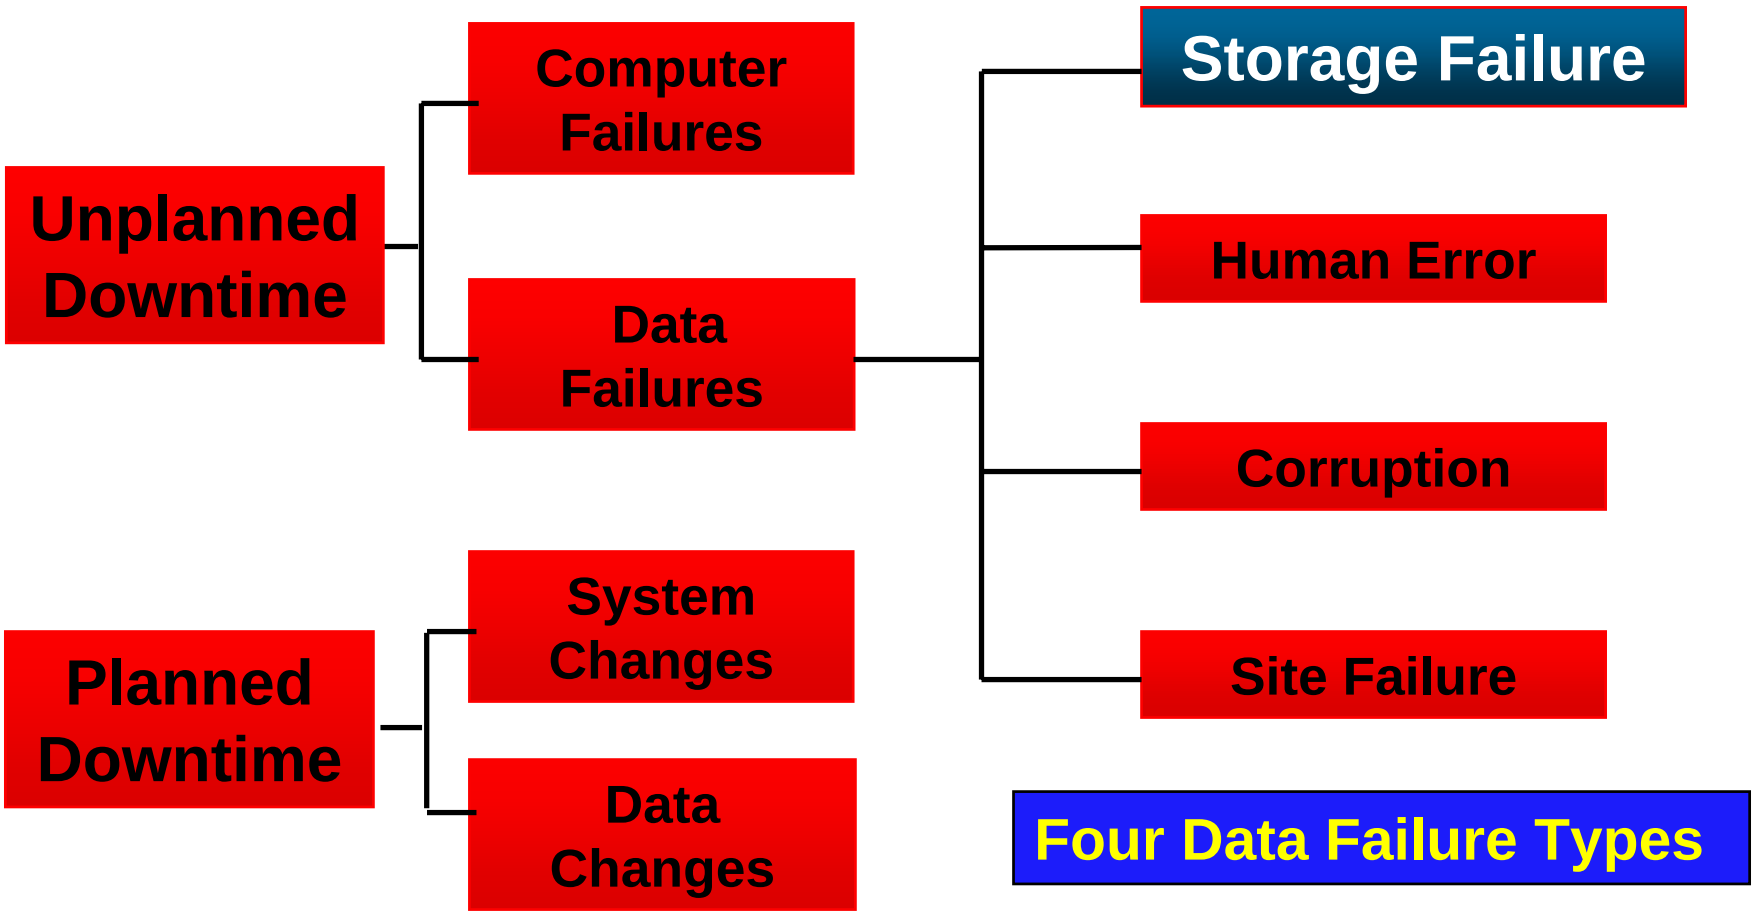

Storage Scale-Out with ASM

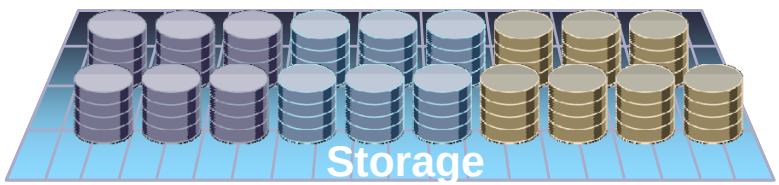

- ASM mirrors data across low cost modular storage arrays
  - Automatically remirrors when disk or array fails
- ASM is free

ORACLE 11g DATABASE

- 11g fine tunes recovery from corrupt blocks and crashed storage arrays

Designed to Tolerate Storage Array Failure

Best Human Error Protection - at Lowest Cost

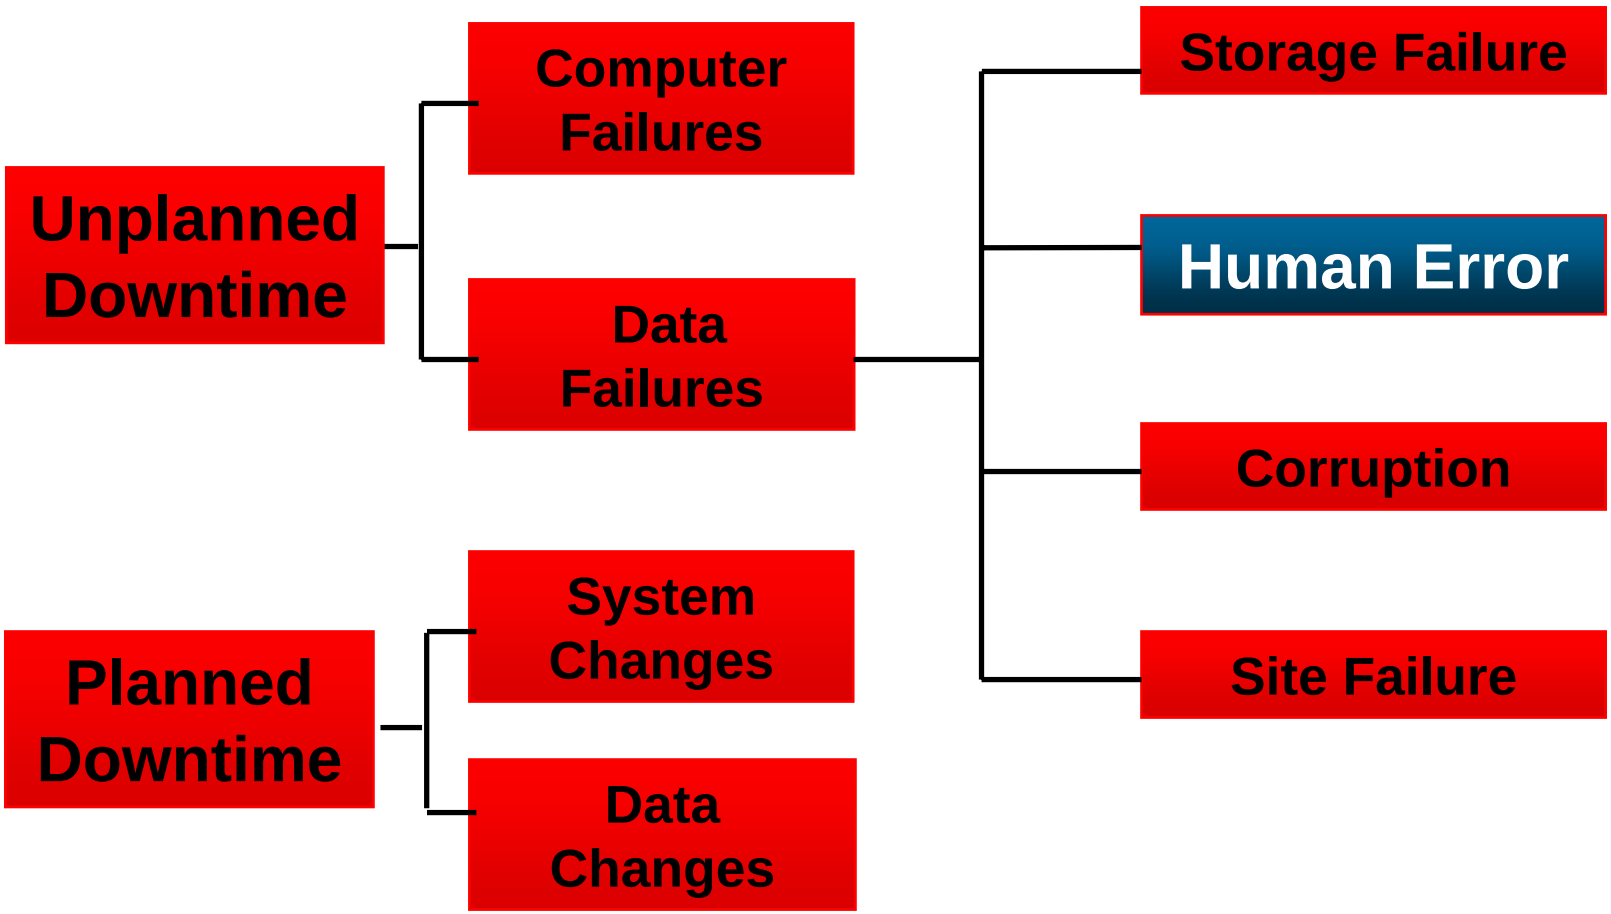

Error Investigation Using Flashback

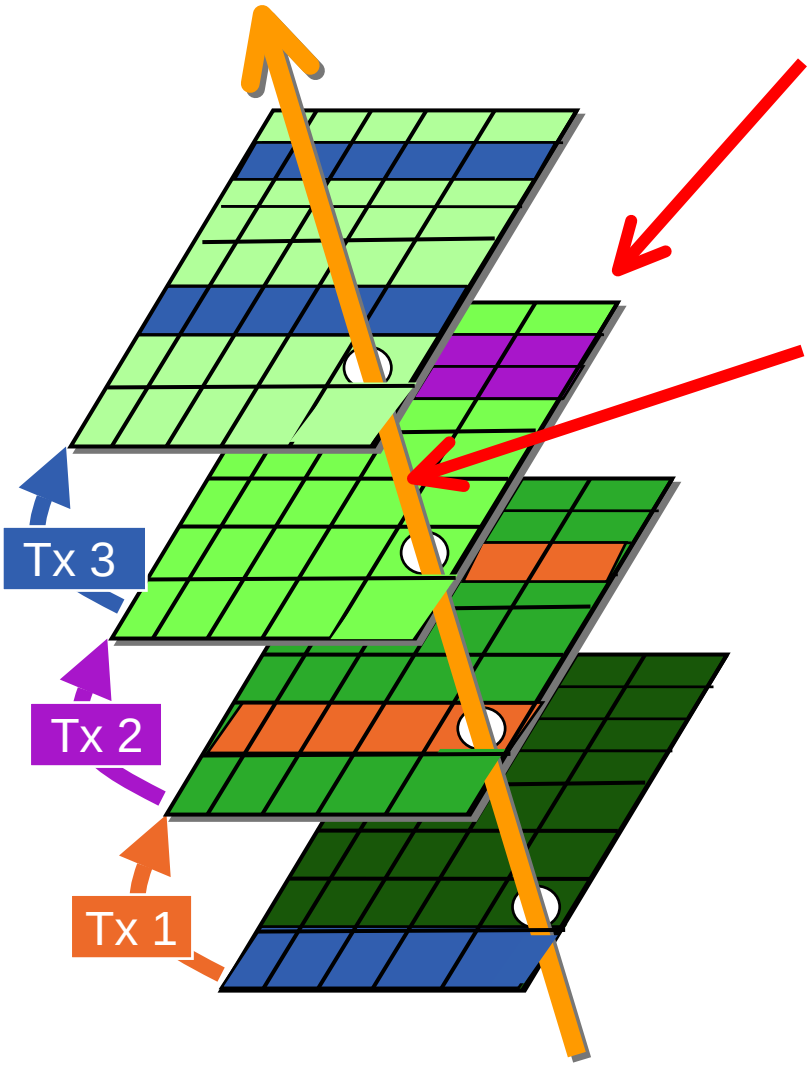

- Flashback Query
  - Query all data at point in time

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

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

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

- Flashback Query is free

Error Correction Using Flashback

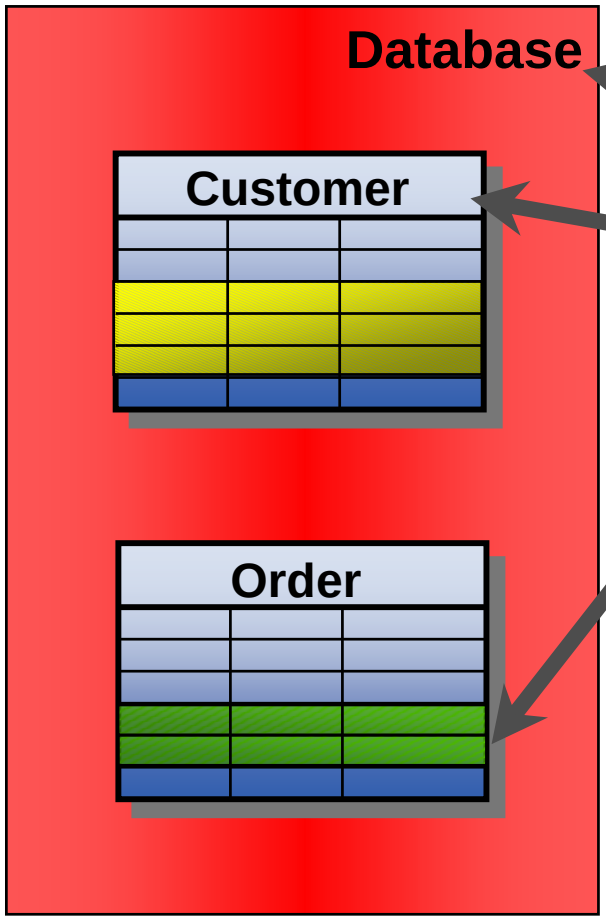

- Correct errors at any level
- Flashback Database - restore database to time
- Flashback Table - restore contents of tables to time
- Flashback Transaction – back out transaction and all subsequent conflicting transactions

- Flashback Correction is free with EE

Great for Testing Also

Revolution in Recovery

- Flashback Revolutionizes Error Recovery
  - Operates on just changed data
  - No restore from backup

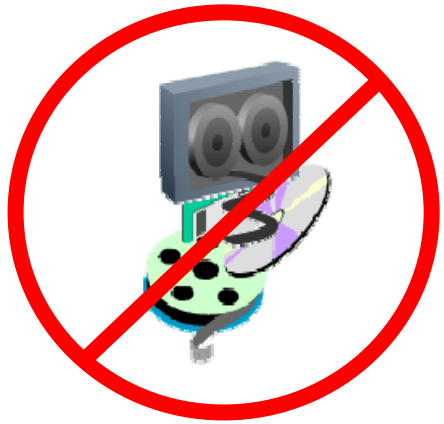

Correction Time = Error Time + ~~f(DB_SIZE)~~

- Flashback is **Easy**
  - Single command instead of complex procedure

Flashback Data Archive

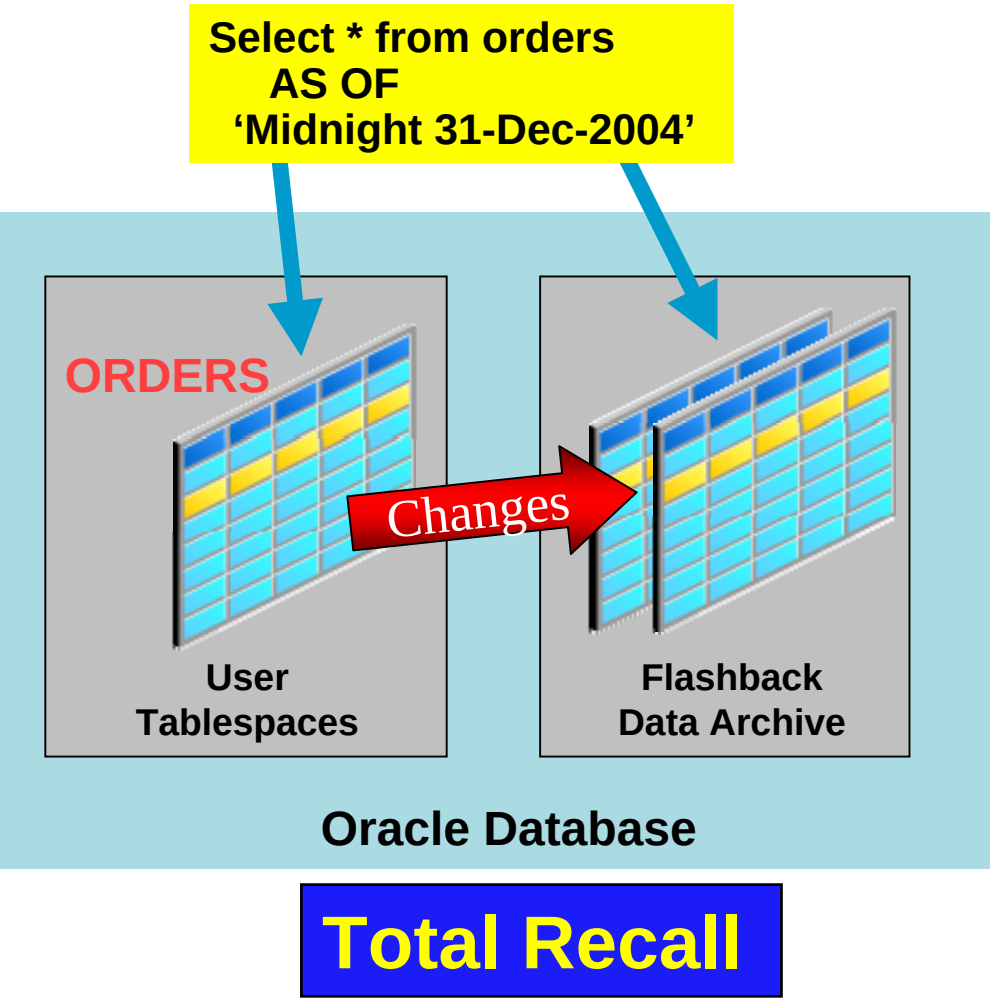

- Automatically stores all changes to selected tables
  - Archive cannot be modified
- View table as of any time
- Uses:
  - Change Tracking
  - ILM
  - Long term history - years
  - Auditing
  - Compliance

Want to Know More?

Technical Session

Total Recall using Flashback Data Archive

Date/Time: Tuesday, 11/13, 4:45 - 5:45 PM

Location: Moscone South 304

Demonstration

Oracle DEMOgrounds

Moscone West Exhibition Hall

Best Corruption Protection - at Lowest Cost

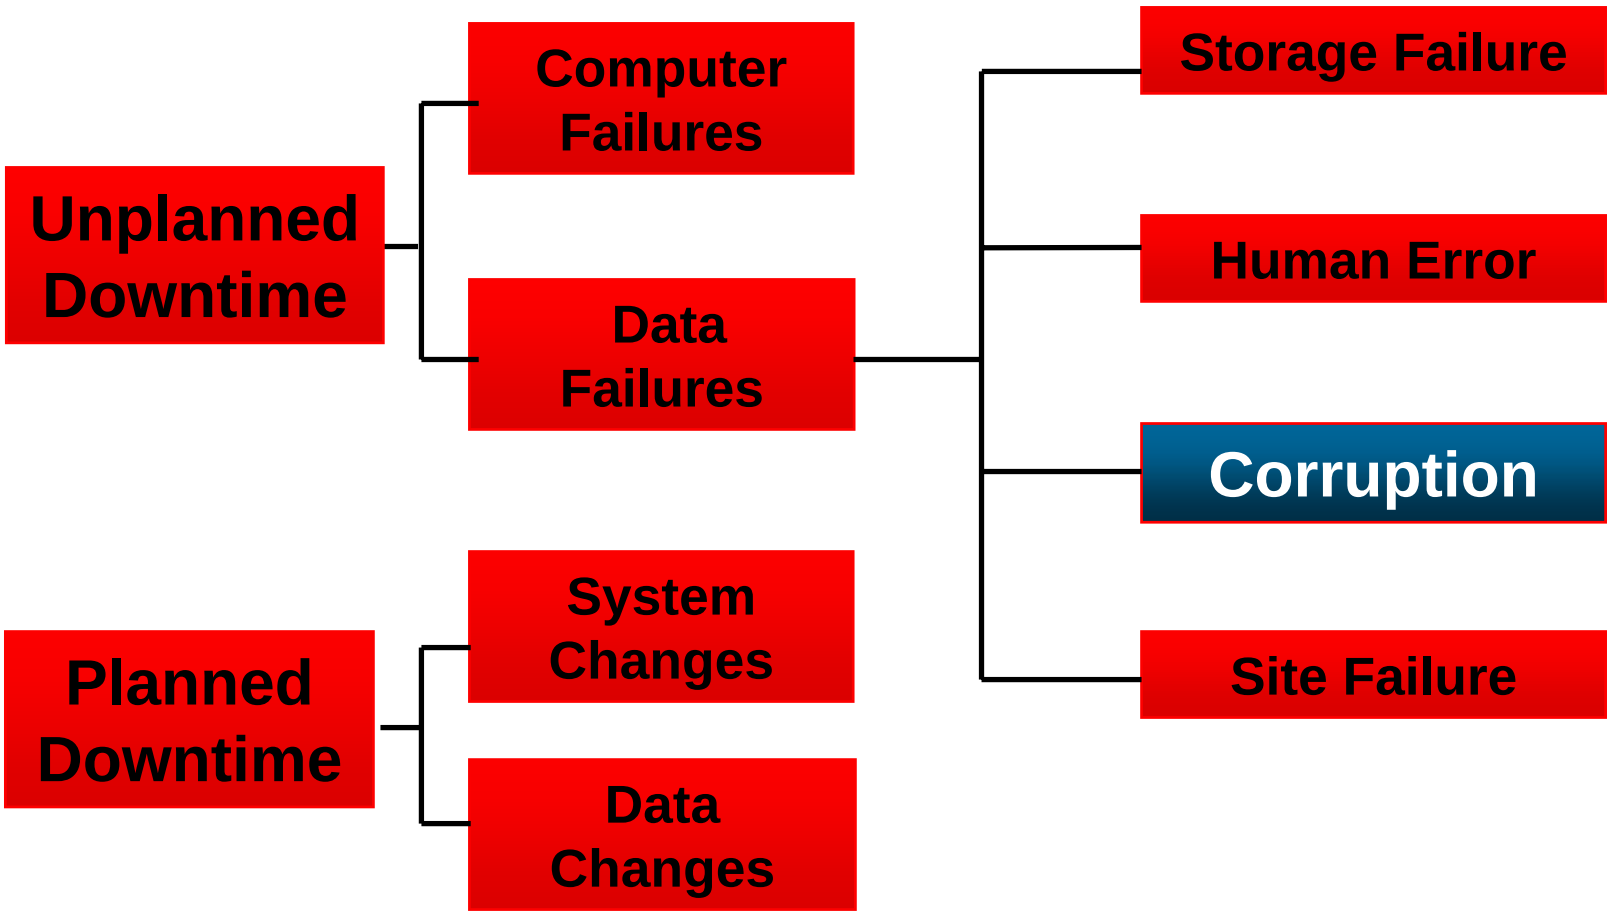

Automated Disk Backup

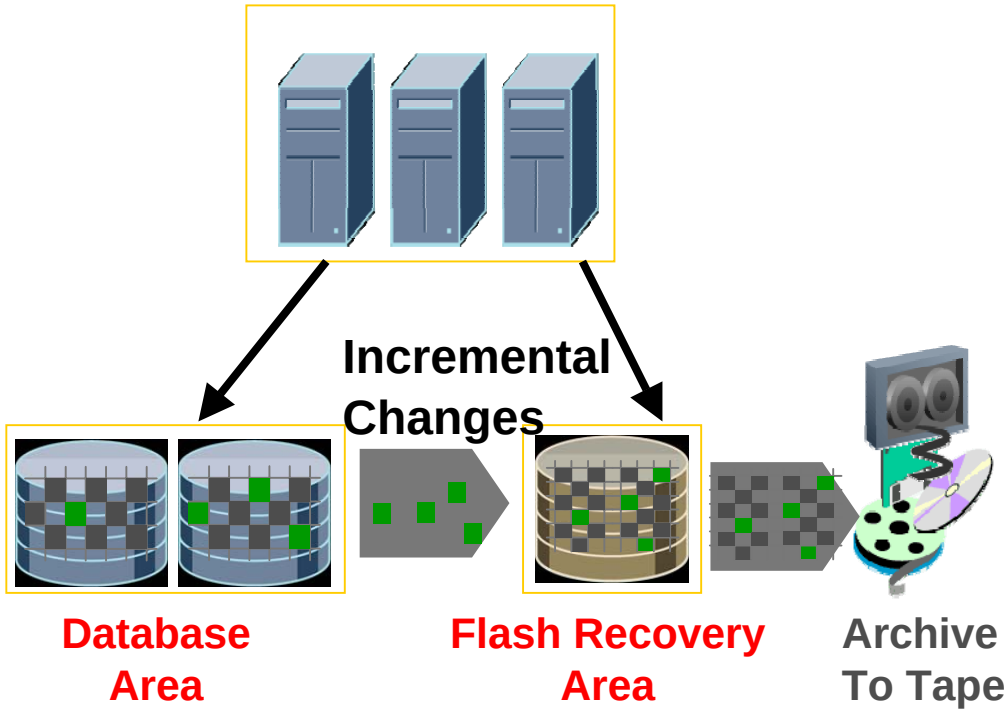

- Oracle fully automates DB backup and recovery to disk
- Nightly incremental backup rolls forward on-disk backup
  - Much faster than tape
  - Blocks validated to prevent corruption of backup
- Restore is instantaneous
- Use low cost ATA disk array for recovery area
- Free

Integrated Backup Tiering for Free

Oracle Secure Backup
Integrated Tape Backup Management

ORACLE SECURE BACKUP

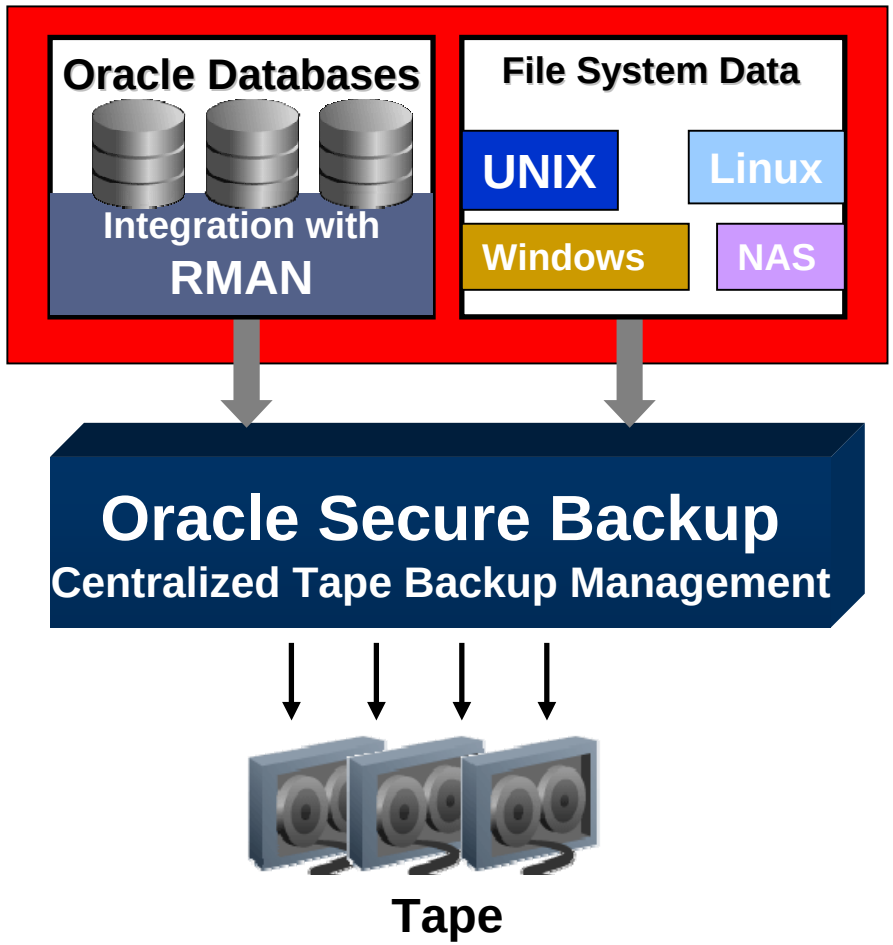

- Protects
  - Oracle DB (9*i* forward)
  - Application files
- Fastest database tape backup
- Lowest cost
- Advanced features such as
  - End-to-End encryption
  - Vaulting
  - StorageTek ACSLS
  - Database backup
  - NDMP filer backup
  - Backup over SAN
- Single Vendor Advantage

World's Fastest Database Backup

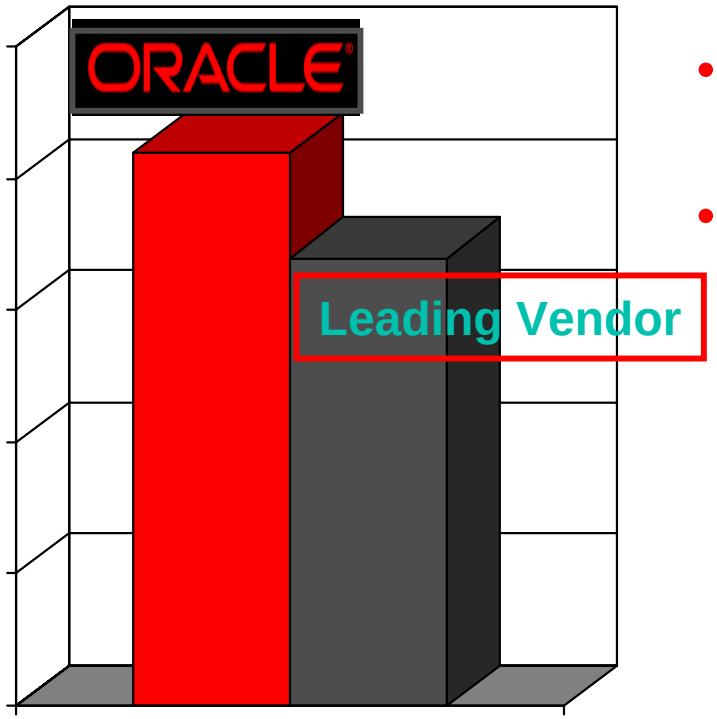

Speed of Database Backup

- Oracle Secure Backup is fast
  - 15% to 30% faster than competition
- Only backup that is tightly integrated with Database kernel
  - Less overhead with direct calls into DB engine
  - Eliminates backup of unused space
    - DB 10.2
  - Eliminates backup of committed UNDO
    - DB 11g

Lowest Price

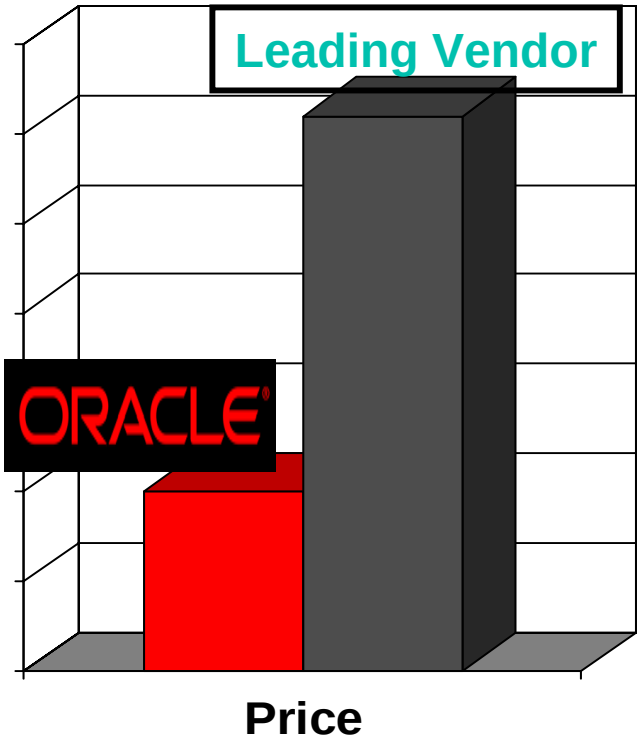

Feature	Leading Vendor	Oracle
Tape Drive	$ 3,000	$ 3,000
Client Host	$ 2,500	**Free**
Oracle Agent	$ 1,000	**Free**
Shared Drive	$ 2,000	**Free**
Media Server	$ 5,000	**Free**
NFS Filer-NDMP	$ 3,500	**Free**
Vaulting	$ 2,000	**Free**
VTL (per TB)	$ 1,000	**Free**

- Oracle Secure Backup price is just $3000 per tape drive
  - **Free** Express Edition allows one direct attached tape drive per DB server

Data Recovery Advisor

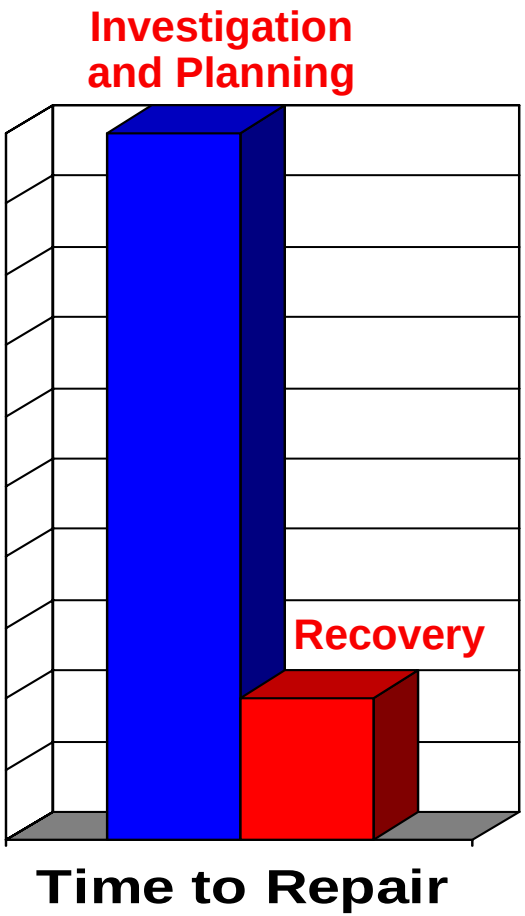

- In an outage, uncertainty and confusion greatly increase downtime
- Data Recovery Advisor
  - Automates problem investigation
  - Intelligently determines plan for recovery
  - Can automatically apply recovery plan

Reduces Downtime by Eliminating Confusion

Best Disaster Protection - at Lowest Cost

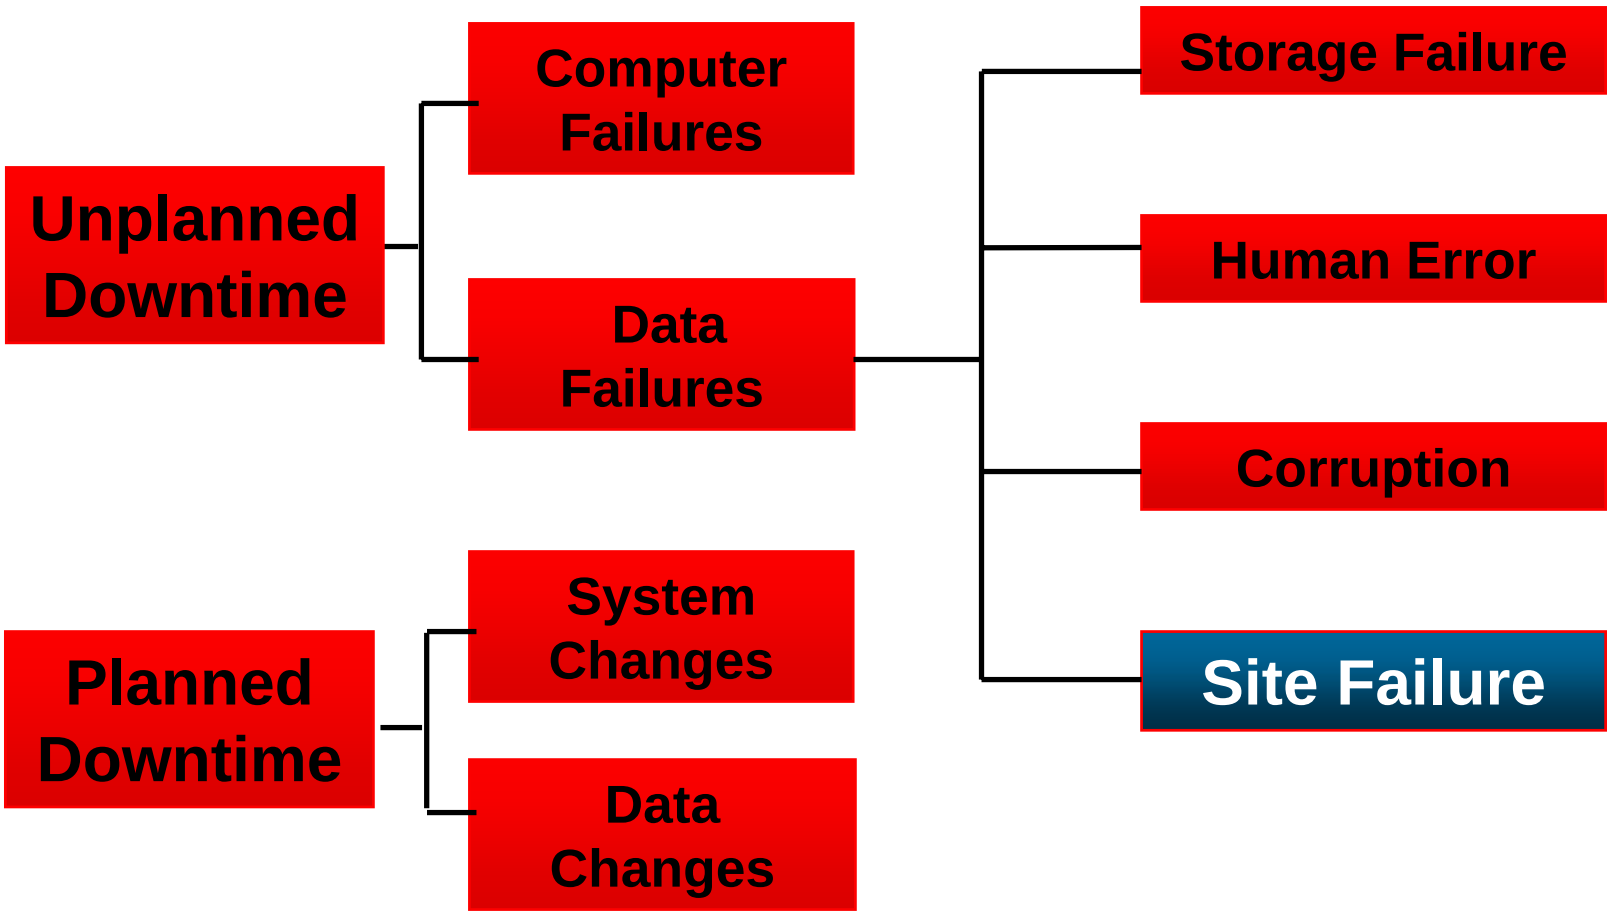

DR Realities

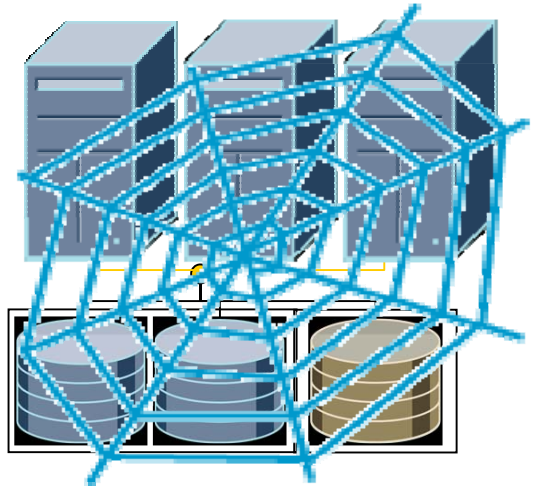

- Majority of customers never benefit from DR investment
  - Expensive – choose no DR, or under-configure DR
  - Loses Data – causes problems with interconnected systems
  - Slow – prefer to try to fix problems instead of using DR
  - Limited – site failures only
- Users are afraid to use DR because it may not work

DR Systems Gather Dust

Data Guard - Best Site Failure Protection - at Lowest Cost

- Synchronous or asynchronous log shipping
  - Corruptions don't propagate
- Compared to remote mirroring, Data Guard provides better:
  - Data protection, network usage, hardware utilization
- Free with EE and works on low cost servers and storage
  - Thousands of production customers

Zero Data Loss over Long Distance

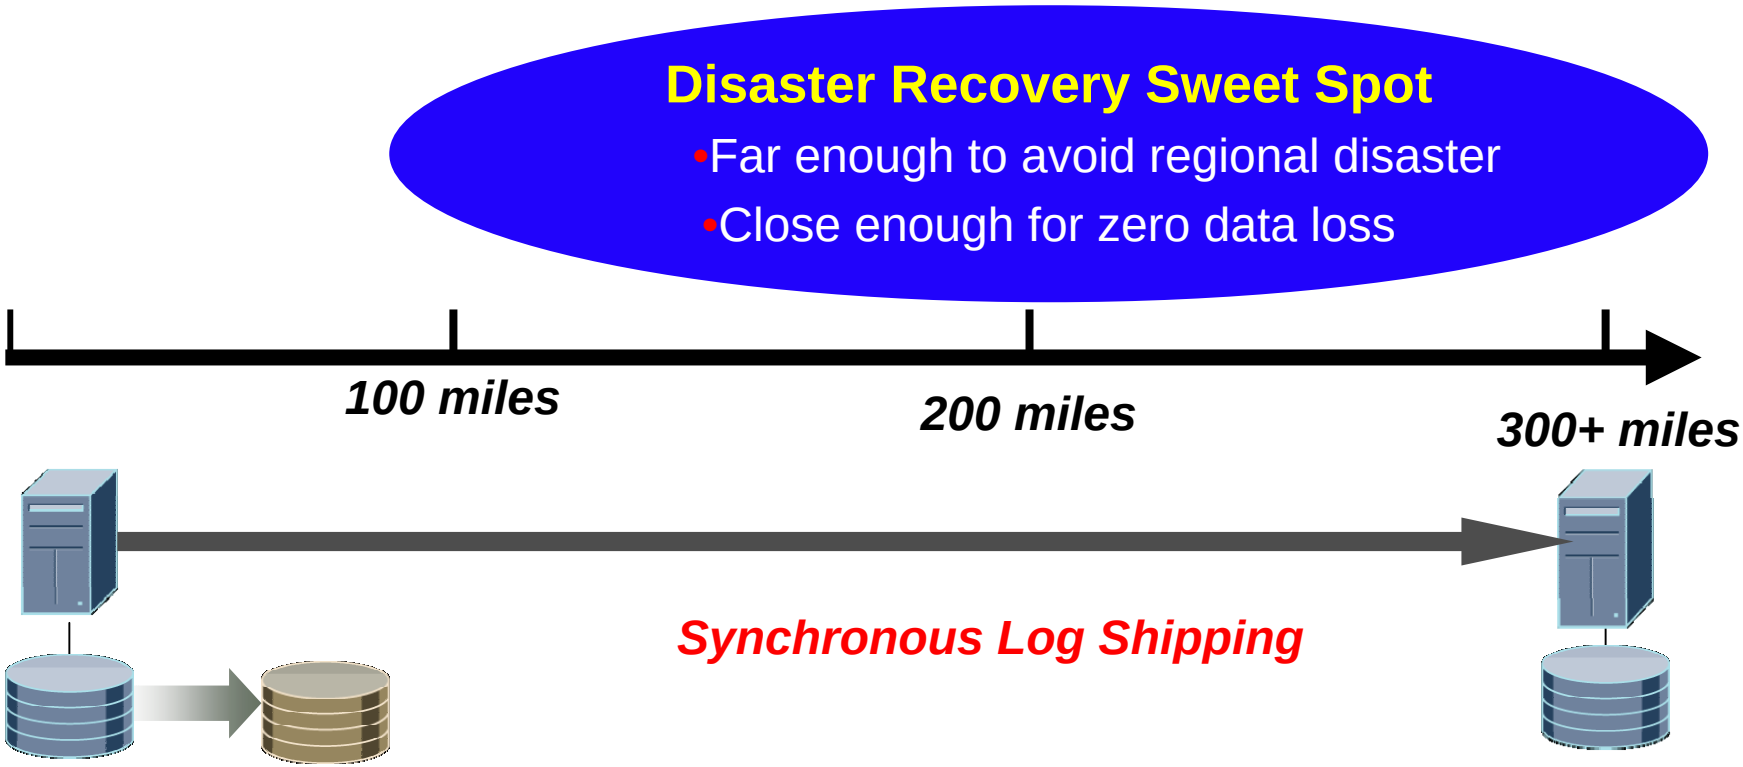

- Log shipping uses order of magnitude less network messaging than remote mirroring
  - Enables zero data loss at hundreds of miles

Automatic Failover to Standby (10.2)

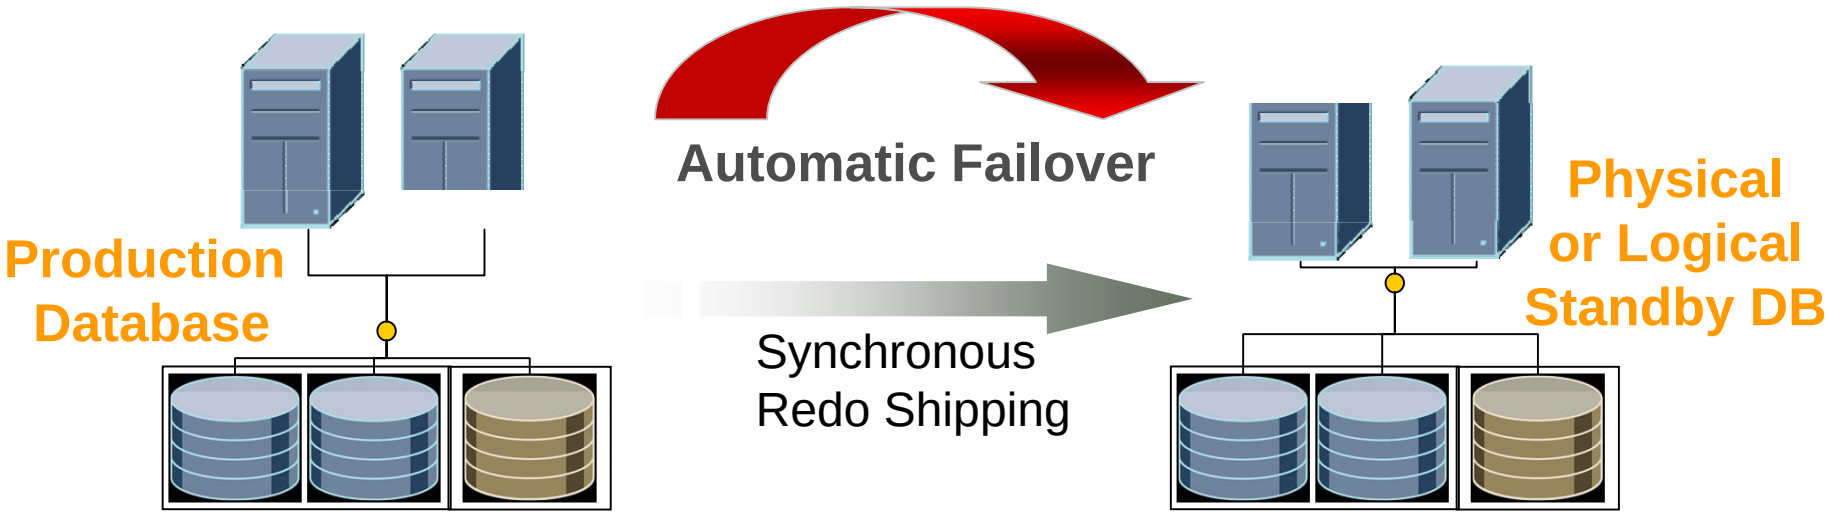

- Fast-Start Failover implements automatic failover to standby database
  - Zero data loss
  - Triggered by failure of Site, Hosts, Storage, Network, or Database
- Failover occurs in seconds
- Original site automatically rejoins the configuration after recovery

Snapshot Standby

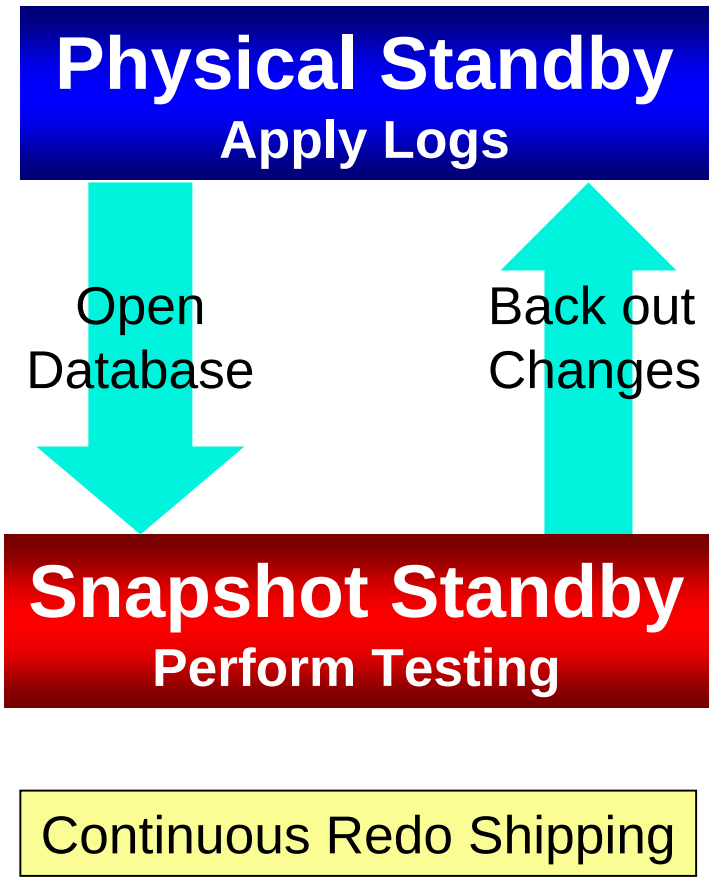

- Use Standby Database for testing and development
  - Eliminates cost of DR
  - Justify full scale test and DR environments by combining into one
- Preserves zero data loss while in test/dev mode
  - But no real time query or fast failover
- Similar to storage snapshots, but:
  - Provides DR at the same time
  - Single copy of storage

Active Data Guard

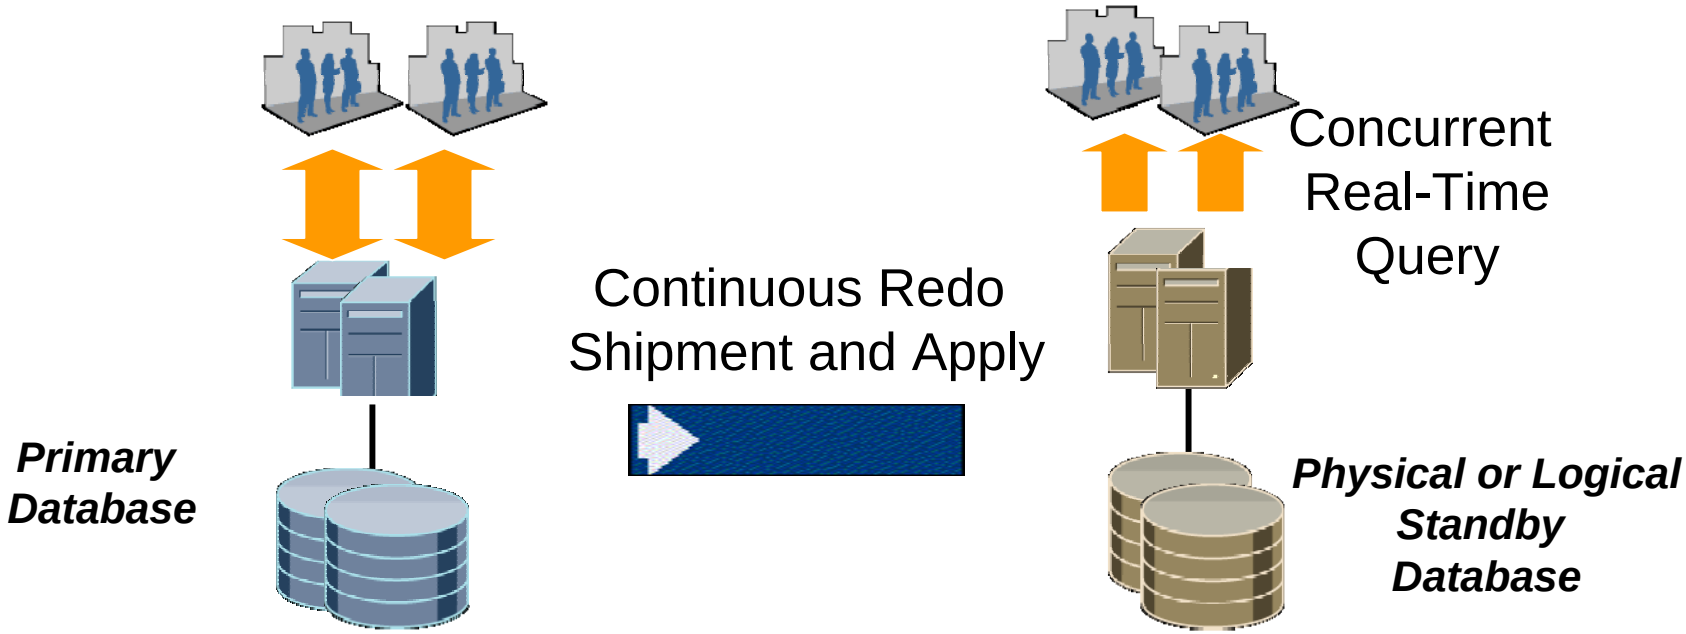

- Queries on standby database concurrent with apply
- Previously available with Logical Standby
- Available with Physical Standby in 11.1
  - Handles all data types, very fast, operationally simple
- Eliminates cost of DR
  - All hardware used for production load

Oracle Streams Replication

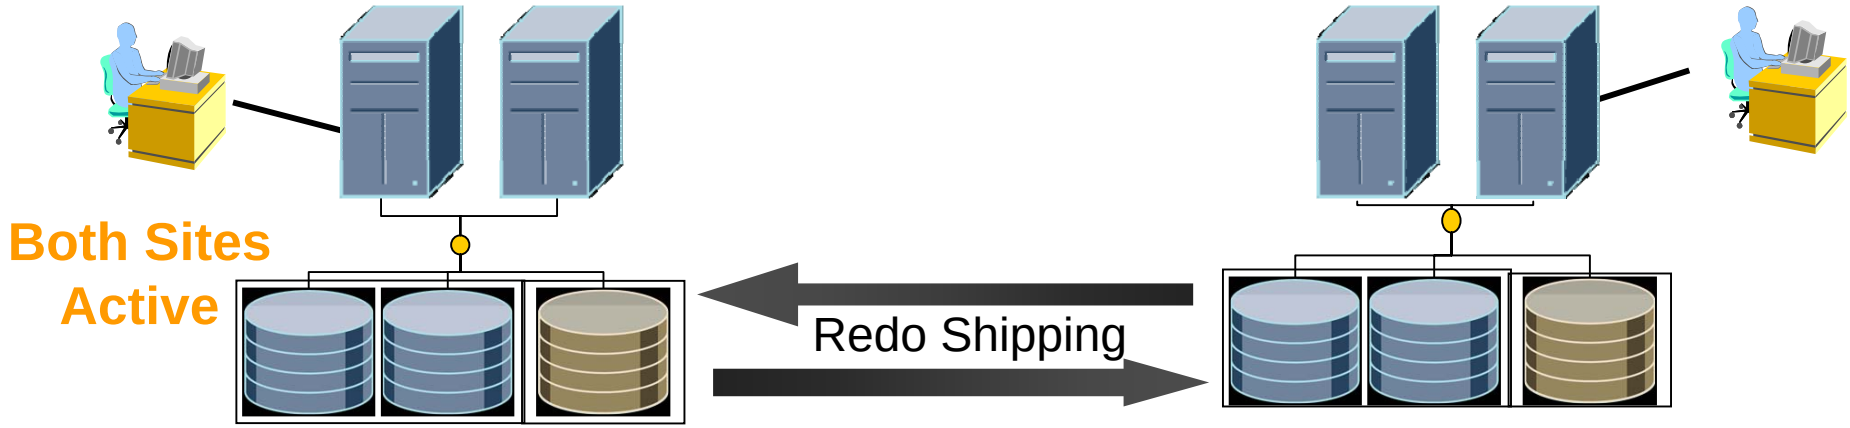

- Similar to Data Guard Logical
  - Shares redo transport, log mining, SQL apply infrastructure
  - Allows zero data loss synchronous transport
- Allows queries and updates on all sites
- Provides great availability for custom applications where update conflicts can be avoided or managed

Best Online System Changes

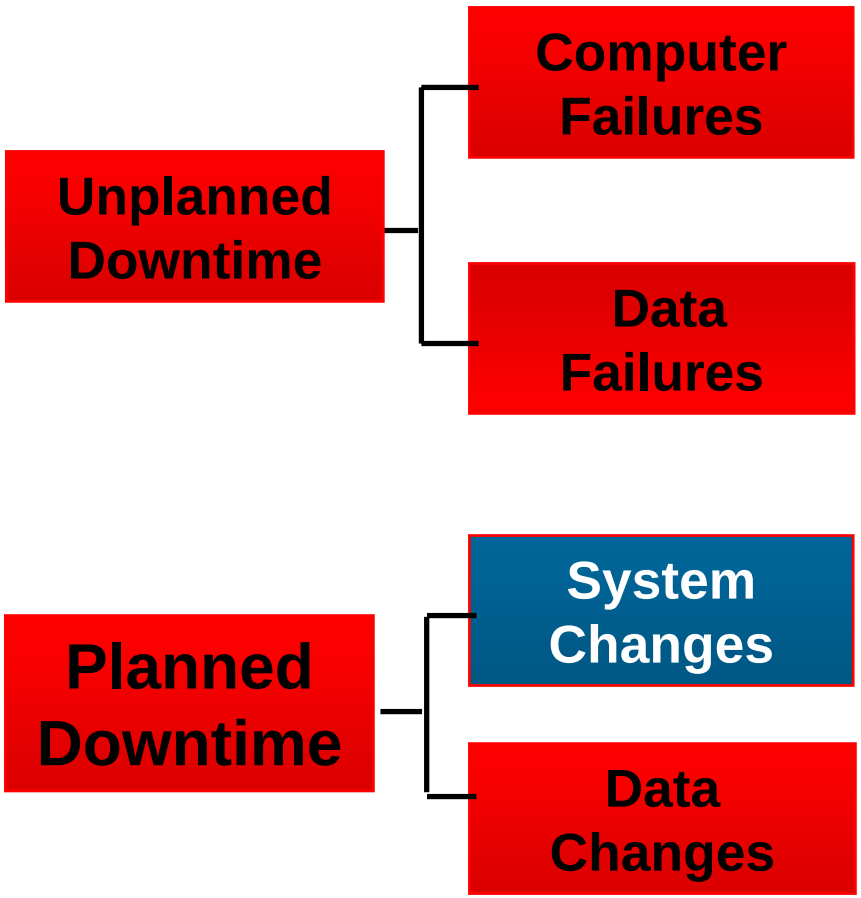

Scaling on-Demand

Online Upgrade

Scaling on Demand

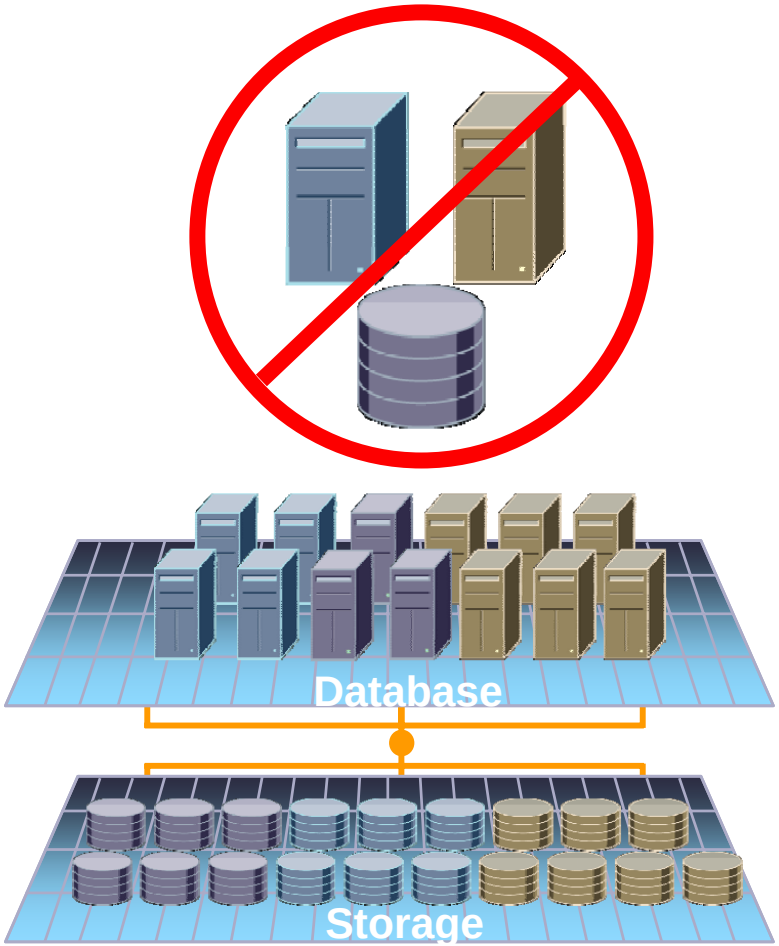

- Servers
  - Add/Remove RAC nodes online
  - No data movement needed
- Storage
  - Add/Remove ASM disks or arrays online
  - Automatically rebalance minimal data

Online Patching of One-off Patches

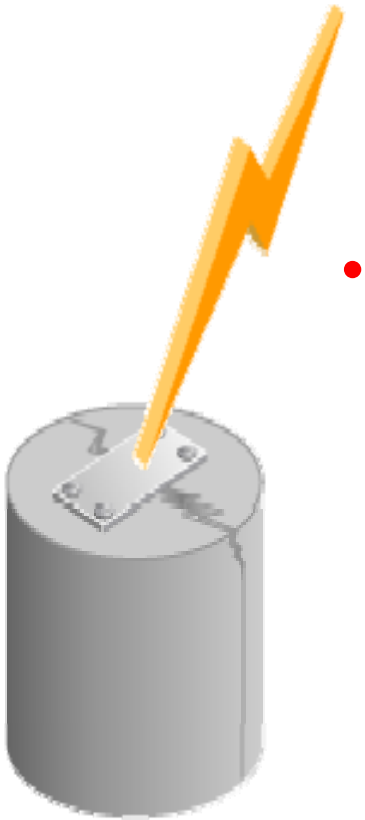

- Simple one-off patches can be applied to a running Oracle instance
  - Initially on Linux x86 and Solaris Sparc
  - Support for HP-UX Itanium and AIX coming
  - Others platforms to follow

Online Patching using RAC

- Roll database patch in one node at a time while other nodes keep running
  - More complex one-off patches
  - Critical Patch Updates

Online Database Upgrades using Data Guard

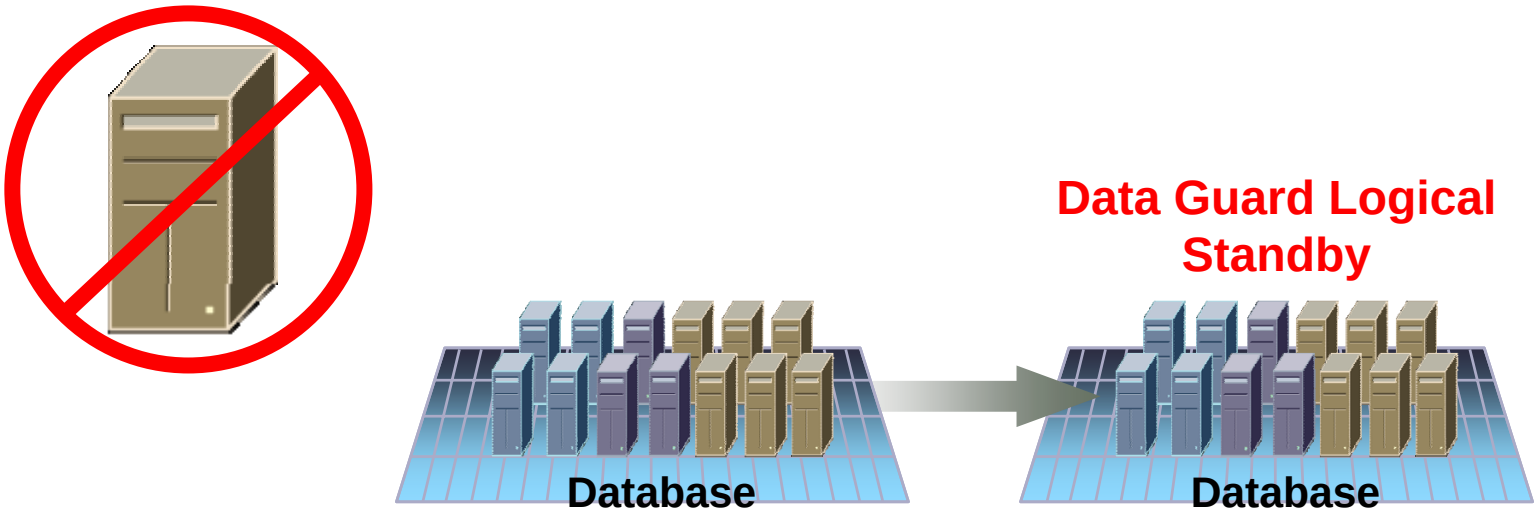

- Upgrade standby site first, then switch over users
  - Patch sets, major releases

Best Data Changes

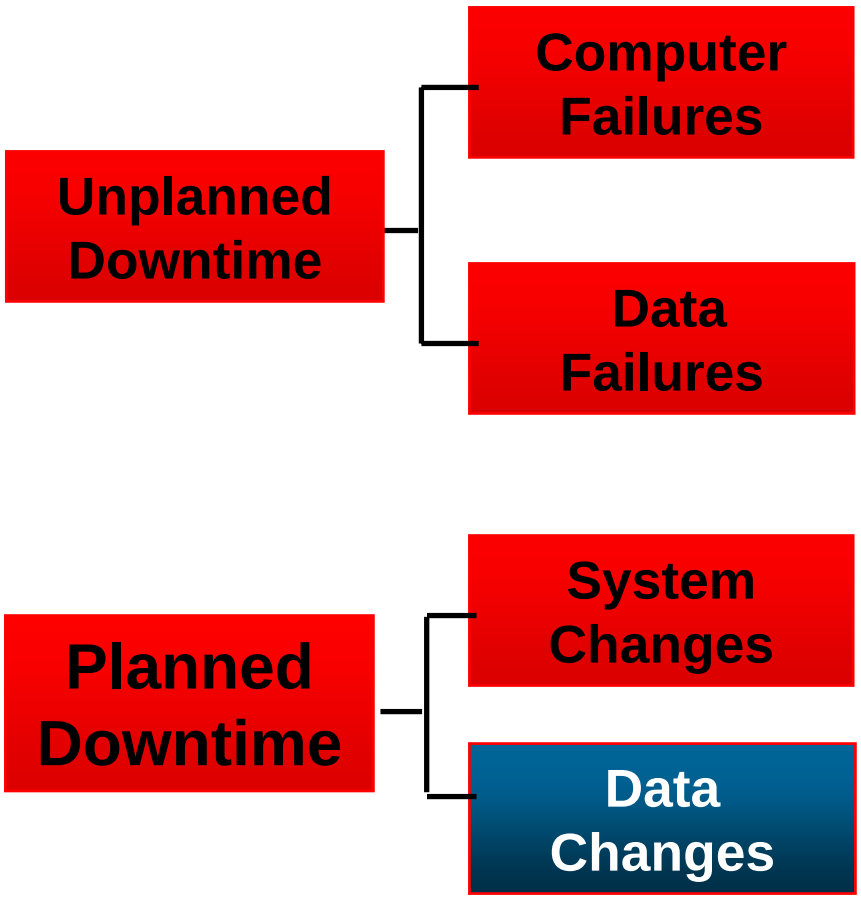

Online Redefinition

Online Table & Index Redefinition

- All index changes can be done online
- Tables can be Reorganized & Redefined online
  - Allows changing location, table type, partitioning, columns, column types
  - Contents can be transformed as they are copied
  - Online redefinition is free with EE

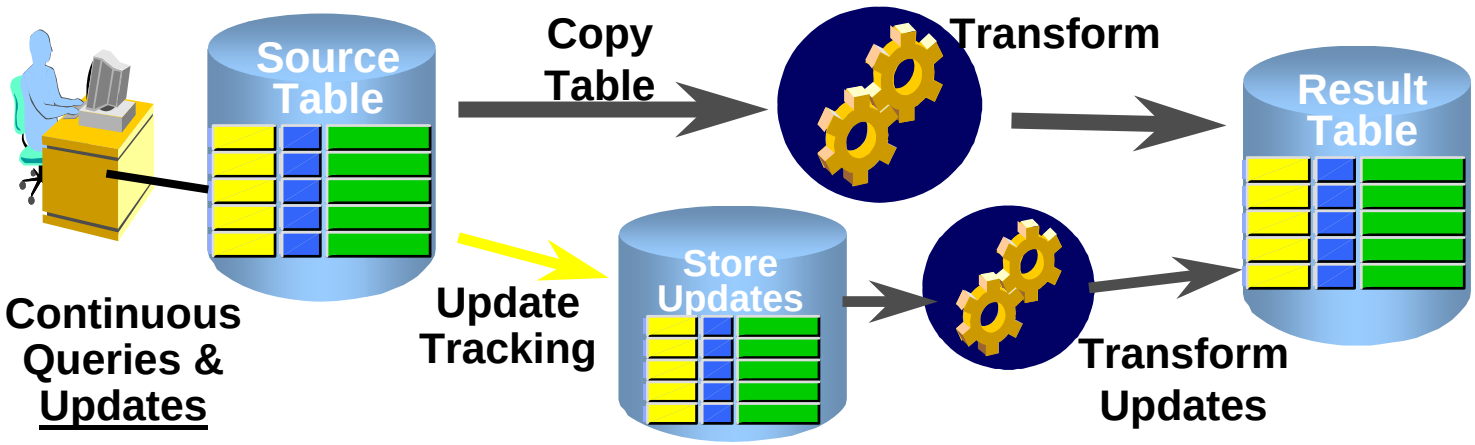

Maximum Availability Architecture (MAA)

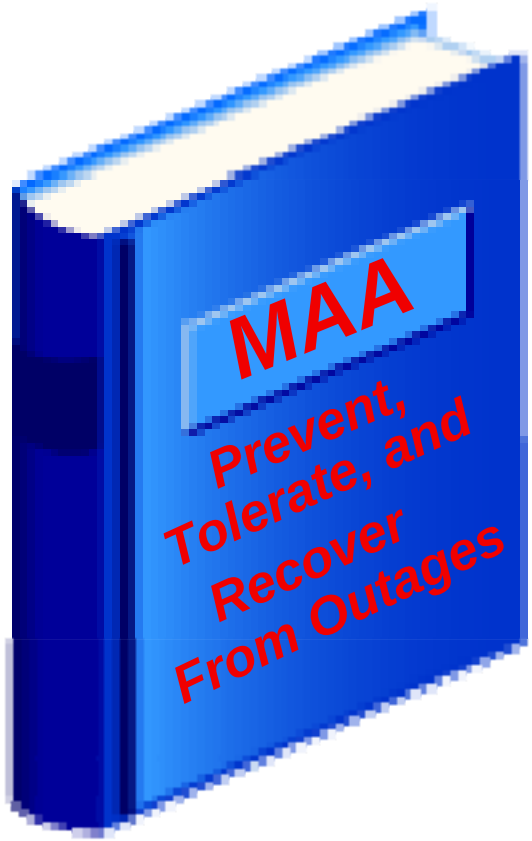

- Technology alone is not enough
- MAA has tested, validated, and documented best practices
  - Database, Storage, Cluster, Network
  - 40 person year effort
    - Not marketing fluff

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

Oracle Maximum Availability Architecture

Integrated suite of best-of-breed HA technologies
- Each is scale-out, fully active, data centric

Best Availability AND Lowest Cost

Real Application Clusters & Clusterware
Fault Tolerant
Server Scale-Out

Online Upgrade
Upgrade Hardware and Software Online

Data Guard
Fully Active
Failover Replica

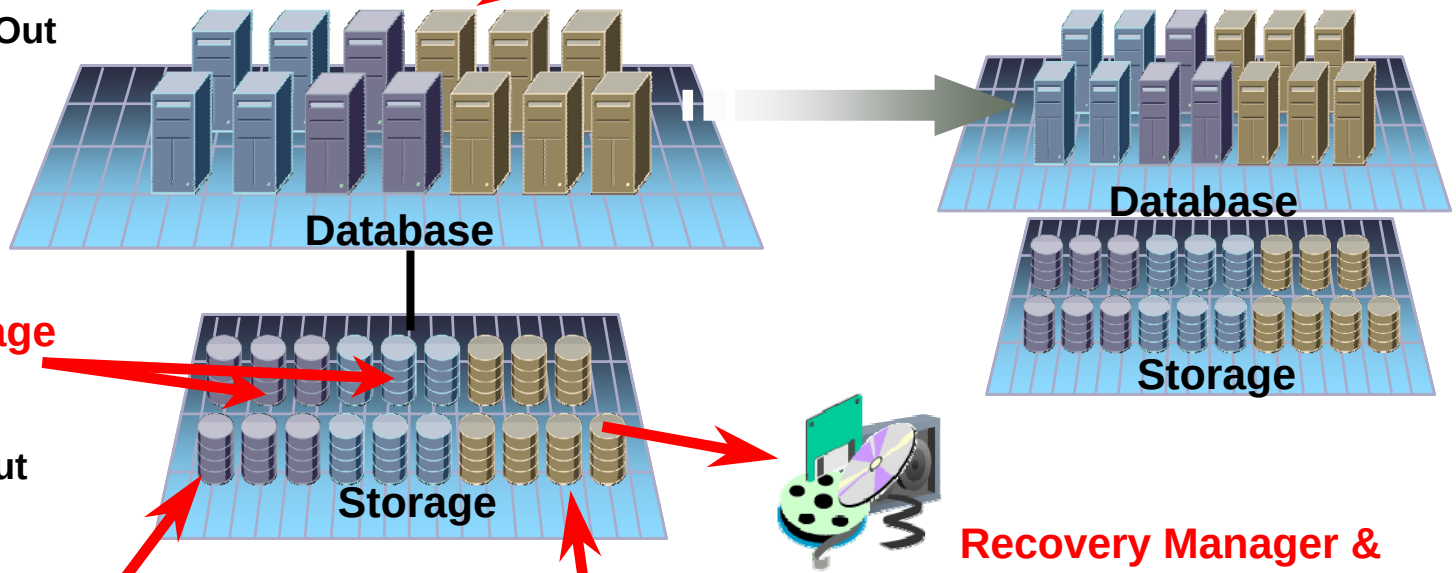

Automatic Storage Management
Fault Tolerant
Storage Scale-Out

Flashback
Correct Errors by Moving Back in Time

Online Redefinition
Redefine Tables Online

Recovery Manager & Oracle Secure Backup
Low Cost High Performance Data Protection & Archival

Database HA Sessions From Oracle Development

Monday, Nov 12

- S291483 - The Fastest and the Most Cost-Effective Backup for Oracle Database: What's New in Oracle Secure Backup 10.2, 11:00 am - 12:00 pm, Moscone South 304
- S291492 - Oracle Database 11*g*: Next-Generation High Availability, 12:30 - 1:30 pm, Moscone South 103
- S291923 - Implementing Oracle Maximum Availability Architecture (MAA) at Allstate Insurance Using Oracle 10*g* RAC, ASM, Oracle Data Guard and Oracle Grid Control, 3:15 - 4:15 pm, Moscone South 304
- S291484 - Oracle Database 11*g* Data Repair Technologies: Comprehensive, Intelligent Recovery, 4:45 - 5:45 pm, Moscone South 304

Tuesday, Nov 13

- S290710 - Maximum Availability Architecture Best Practices: Oracle E-Business Suite 12, 12:15 - 1:15 pm, Marriott Salon 10 & 11

Wednesday, Nov 14

- S291915 - What's New in Oracle Data Guard 11*g*: Revolutionizing Data Protection and Availability, 9:45 - 10:45 am, Moscone South 304

Database HA Sessions From Oracle Development

Wednesday, Nov 14

- S291487 - Backup and Recovery Best Practices for Very Large Databases (VLDB), 11:15 am - 12:15 pm, Moscone South 304
- S291920 - Oracle Active Data Guard: How to Utilize Your Standby Databases for Production Workload - What They Didn't Print in the Manuals, 3:00 - 4:00 pm, Moscone South 304
- S291917 - Oracle Data Guard Tips and Tricks: Direct From Oracle Development, 4:30 - 5:30 pm, Moscone South 102

Thursday, Nov 15

- S291495 - Oracle Streams Replication and Advanced Queuing (AQ): What's New in Oracle Database 11*g*, 8:30 - 9:30 am, Moscone South 304
- S291499 - Best Practices for Implementing Replication with Oracle Streams in Oracle Database 10*g* and 11*g*, 10:00 - 11:00 am, Moscone South 304
- S291525 - Maximum Availability Architecture (MAA) Best Practices: Online Patching, Rolling Upgrades and Planned Maintenance with Minimal Downtime with Oracle Database, 11:30 am - 12:30 pm, Moscone South 104
- S290542 - Maximum Availability Architecture (MAA) Best Practices for Siebel 8.0, 2:30 pm - 3:30 pm, Marriott Salon 10 & 11

Database HA Demos From Oracle Development

Monday, Nov 12 – Thursday, Nov 15
Oracle DEMOgrounds, Moscone West

Oracle Active Data Guard

Oracle Streams: Replication and Advanced Queuing

Oracle Secure Backup

Recovery Manager (RMAN) and Flashback Technologies

Maximum Availability Architecture

Appendix

HA Features by Version

New Oracle Database 11.1 HA Features

Clusters

- Better performance, failover, management
- Seamless integration with XA and Microsoft Transaction Server
- Optimized Cache Fusion protocols
- Transaction tagging

ASM

- Temporarily freeze I/O to unresponsive storage and replay writes when it returns
- Rolling Upgrade for ASM instances

Data Guard - Generic

- Snapshot standby
- Fast-Start Failover in secs
- Fast-Start Failover for asynchronous mode
- Fast Redo Transport
- Redo Transport Compression for gaps
- SYS user no longer required for redo shipping
- Standby across Linux/Windows on x86

Flashback

- Flashback Transaction
- Flashback Data Archive

Data Guard – SQL Apply

- Support XMLType data type (CLOB)
- Support Transparent Data Encryption (TDE)
- Apply performance improved significantly
- Support Fine Grained Auditing
- Support Virtual Private Database

Data Guard – Redo Apply

- Real-time query
- New parallel media recovery doubles redo apply performance
- Fast Incremental Backup
- Better RMAN integration

New Oracle Database 11.1 HA Features

Backup & Recovery

- Data Recovery Advisor
- Intra-file parallel backup and restore
- Faster backup compression
- Duplicate database over the network
- Integration with Windows Volume Shadow Copy Services (VSS) API
- Integrated protection through ultra safe mode

Upgrades

- Online patching
- Transient logical standby

Online Redefinition

- Fast add column with default value
- Invisible Indexes prevent premature use of newly created indexes
- Online index build with NO pause to DML
- No recompilation of dependent objects
- Easier to execute table DDL operations online

Streams

- Support XMLType data type (CLOB)
- Support Transparent Data Encryption (TDE)
- Apply performance improved significantly
- Support AWR, ADDM, ASH
- Streams Performance Advisor
- Synchronous capture
- Table Data Comparison
- JMS/AQ performance improvements

New Oracle Secure Backup 10.2 Features

- Enhanced security architecture implementing an NDMP data service
- Policy driven backup encryption
- Fully integrated Oracle Enterprise Manager Grid Control 11g for file system and database tape backups
- Integrated backup/restore for Oracle applications
- Eliminates backup of undo tablespace
- Optimization of SBT-allocation buffers
- Media lifecycle management - automates rotation of backup tapes between multiple locations
- User-defined policies automate duplication of backup tapes
- Support migration of virtual tapes from VTL to physical tape

10gR2 Availability Features

New Oracle Database 10.2 HA Features

Clusters

- Increase Max Instances to 1055 on all platforms
- Clusterware Open API
- Clusterware File Redundancy for OCR and voting disk
- Cluster Verification

Data Guard SQL Apply

- Support for IOTs with LOBs and Overflow
- Improved Monitoring for SQL Apply
- Optimized Creation of Logical Standby DB

Data Guard Generic

- Automatic Failover to Standby
- Fast Failover to Standby (seconds)
- Optimized Asynchronous Redo Transmission
- Automatic Deletion of Applied Archive Logs
- Conversion of Physical Standby DB to Reporting DB and back

Diagnostics

- Much faster db_block_cache_protect
- Improved block validation

Online Redefinition

- Online Move of a Partition of a Partitioned Table
- Support Clustered Tables
- Support Materialized View Logs
- Support AQ tables
- Support Abstract Data types
- Clone Statistics, Check and Not-null Constraints

New Oracle Database 10.2 HA Features

Oracle Backup

- Tape Management Support
- RMAN Integration
- NDMP
- Fibre Channel Drives
- Centralized Administration
- Great Database Backup Performance

Tuning

- Faster Instance Startup for Large Memory Configurations

Backup & Recovery

- Temporary Datafiles are Re-created on RMAN Recovery
- Database Transport Across Same Endian Platforms
- Dynamic RMAN Channel Allocation for Backing Up and Recovering RAC Databases
- RMAN HARD compliant backups
- Enhanced RMAN Backup Job Views
- Incremental Roll Forward of Database Copy (standby db or reporting db)

Flashback

- Flashback Database Through Resetlogs
- Named Restore Points
- Guaranteed Restore Points
- Flashback Across Data Guard Switchovers

10gR1 Availability Features

New Oracle Database 10*gR1* HA Features

Clusters

- Portable Clusterware
- Cluster file system for Linux & Windows
- Automated Patching

Data Guard SQL Apply

- Support for Longs
- Support for multi-byte CLOBs and NCLOBs
- Support for Index Organized Tables
- Functional Indexes
- Simplified zero data loss failover
- Real time apply allows real time reporting
- Zero downtime instantiation

Data Guard Generic

- Data Guard Broker support for RAC
- Named Data Guard Configurations
- Real Time Apply
- Flashback Standby Database
- Flashback Reinstantiation
- Improved Recovery Parallelism
- Role based parameter settings

Rolling Upgrades

- Rolling Upgrades Using Data Guard SQL Apply

Online Redefinition

- Support of Unique Indexes
- One Step Cloning of Dependent Objects
- Columns can be Populated Using Sequences & Sysdate
- Signature Based Dependency Tracking Using Synonyms
- Online Segment Shrink

New Oracle Database 10*gR1* HA Features

Automated Disk Based Backup & Recovery

- Automated Management of B&R Disk Space
- Simplified Backup Using Image Copy
- Change Aware Incremental Backups
- Incrementally Updated Backups

Tuning

- Improved Recovery Parallelism
- Faster Instance Startup & Cache Warm

Backup & Recovery

- Simplified Recovery Through Resetlogs
- Restore Tolerates Missing Backups
- Proxy Backup of Archives
- Automated TSPITR Instantiation
- Full DB Begin Backup
- Automated Backup Channel Failover
- Simplified RMAN cataloging of backup files
- Automated File Creation during Recovery
- Drop Database
- Rename Tablespace

Flashback

- Flashback Drop
- Flashback Row History
- Flashback Table
- Flashback Transaction History
- Flashback Database
- Better map of time to SCN for flashback

LogMiner

- Automated Specification of Logs to Mine
- Support for Shared Server Configurations
- Fine Grained Supplemental Logging

9*i* Availability Features

New Oracle 9*i*R2 High Availability Features

Online Operations

- Rename Column
- Rename Constraint
- Online Redefinition without primary key

Human Error

- Flashback time in FROM clause of SQL statement
- Flashback retention time for LOB data

Logminer

- Support for LOBs, Longs, PDML
- Multi-versioned dictionary
- Real time mining

Data Guard

- SQL Apply
- Data Guard Broker failover automation
- Better handling of network glitches
- Multiple synchronous standbys
- Disallow unlogged DML option
- Resend corrupt or missing logs
- RAC support in all modes

Hardware Assisted Resilient Data

Faster DB Upgrades

- Do nothing on load of unchanged package, view, synonym
- Change to invalid package won't invalidate dependents
- Load everything first, then compile
- Fast loading of source
- Parallel Recompilation

Miscellaneous

- Better space handling when restoring archives
- Automatically backup control file after structural DB change

New Oracle 9*i*R1 High Availability Features

Data Recovery

- Block level media recovery
- Trial Recovery
- Tolerate corrupt redo logs
- Self-describing backups
- Policy based automated backup and recovery
- Stored backup configurations
- Resumable backup and restore

Online Operations

- Online indexing operations
- Online table redefinition and reorganization
- Dynamic buffer cache/shared pool resizing
- Online ANALYZE VALIDATE
- Online add and remove CPU

Human Error

- Flashback Query

Miscellaneous

- Quiesce DB for maintenance
- Online add column/site for replication groups
- Offline Diagnostics
- Maintain global indexes during partition DDLs
- Resumable Space Allocation

New Oracle 9*i*R1 High Availability Features

Fast Fault Recovery

- Minimal I/O crash recovery
- Time-based limit on crash recovery
- Resumable space allocation

Log Analysis

- Support clustered tables, chained rows, DDL, direct loads
- Query by content of change

Data Protection

- Zero data loss standby
- Push-Button standby automation
- Delayed apply standby
- Network outage tolerance
- Near real-time reporting
- Tolerate corrupt logs

Cluster Recovery

- Non-disruptive cluster reconfiguration
- Disk heartbeat validates network heartbeat
- Integrated Oracle Parallel Fail Safe
- Multi-node Fail Safe for Windows 2000

Other Oracle DB HA Related Features

- Tablespace Point-in-Time Recovery
- Transparent Application Failover
- On Demand Rollback
- Incremental Checkpoints
- Unlimited Row Level Locking
- Read Consistency
- Range Partitioned Tables
- Read-Only Tablespaces
- Parallel “Everything”
- Advanced Queuing
- Online IOT move
- Database Resource Manager
- Oracle Enterprise Manager