




ORACLE  
**OPEN**  
WORLD

# Your. Open. World.

## The Future of Availability: Maximum Availability Architecture

Juan Loaiza  
Senior Vice President, Oracle

ORACLE®



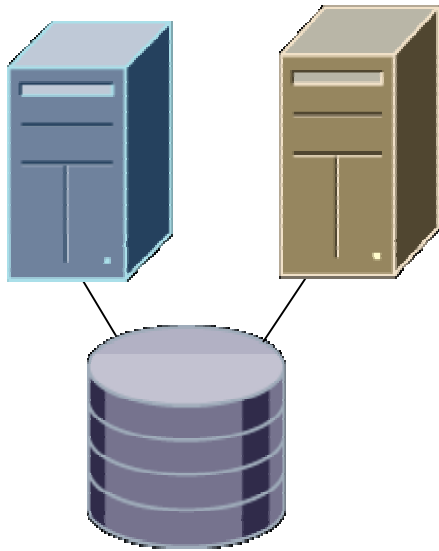
The following is intended to outline our general product direction. It is intended for information purposes only, and may not be incorporated into any contract. It is not a commitment to deliver any material, code, or functionality, and should not be relied upon in making purchasing decisions. The development, release, and timing of any features or functionality described for Oracle's products remains at the sole discretion of Oracle.

# Goals and Vision for Availability

- Fault-Tolerant Grid-Based Infrastructure
- Application Oriented
- Complete, Integrated, Automated, Simple
  - Yet Open, Standard, Pluggable
  - and Low Cost

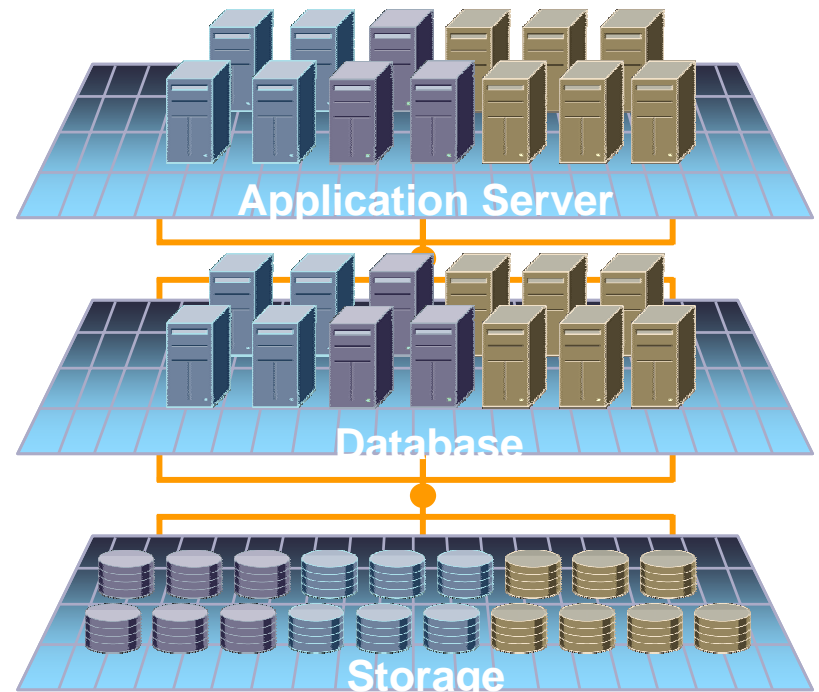
**Maximum Availability Architecture**

# Fault-Tolerant Grid-Based Infrastructure



## Traditional Approach

- Monolithic scale-up hardware
- Idle failover servers
- Idle DR systems
- Limited & high cost



## Grid Approach


- Scale-out standard hardware
- All servers active
- Active DR systems
- Unlimited & low-cost

# Grid Infrastructure Breaks Tradeoff Between Availability and Cost

**Best Availability AND Lowest Cost**

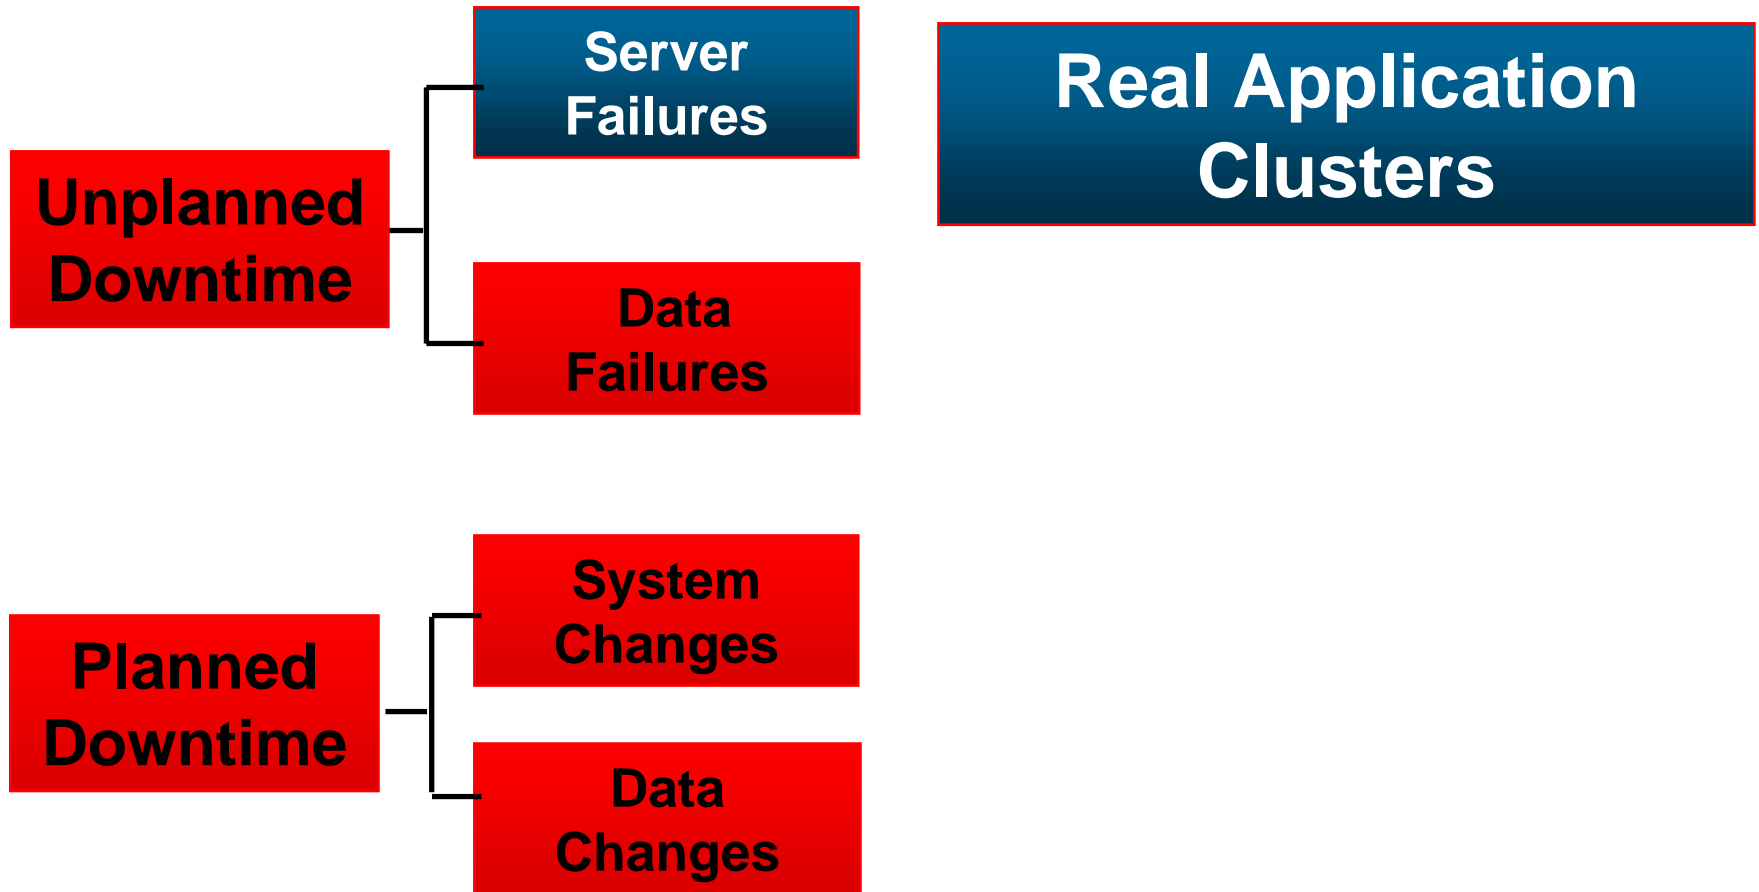
- Better than Mainframe Availability
- PC Economics
- Seamless and Simple to Use

# Application Oriented Availability



- Protect and Recover application objects & transactions
- Recover from Human Errors
  - System Admins, DBAs, Developers, and End Users
- Online Changes to Database and Applications

# Best Server Protection - at Lowest Cost



# Server Scale-Out with RAC



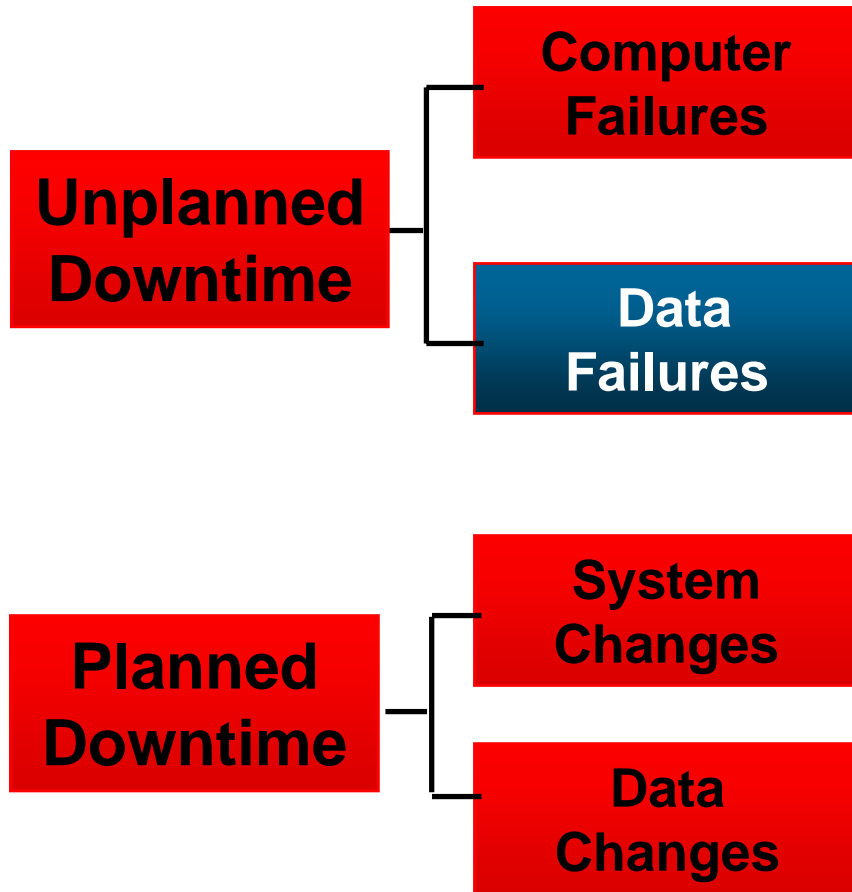
- RAC pools standard low cost servers
- Great Scalability & Availability
  - No Idle Resources
- Integrated clusterware for free
- Runs commercial applications
  - Oracle Applications, SAP, etc.
- Thousands of production customers



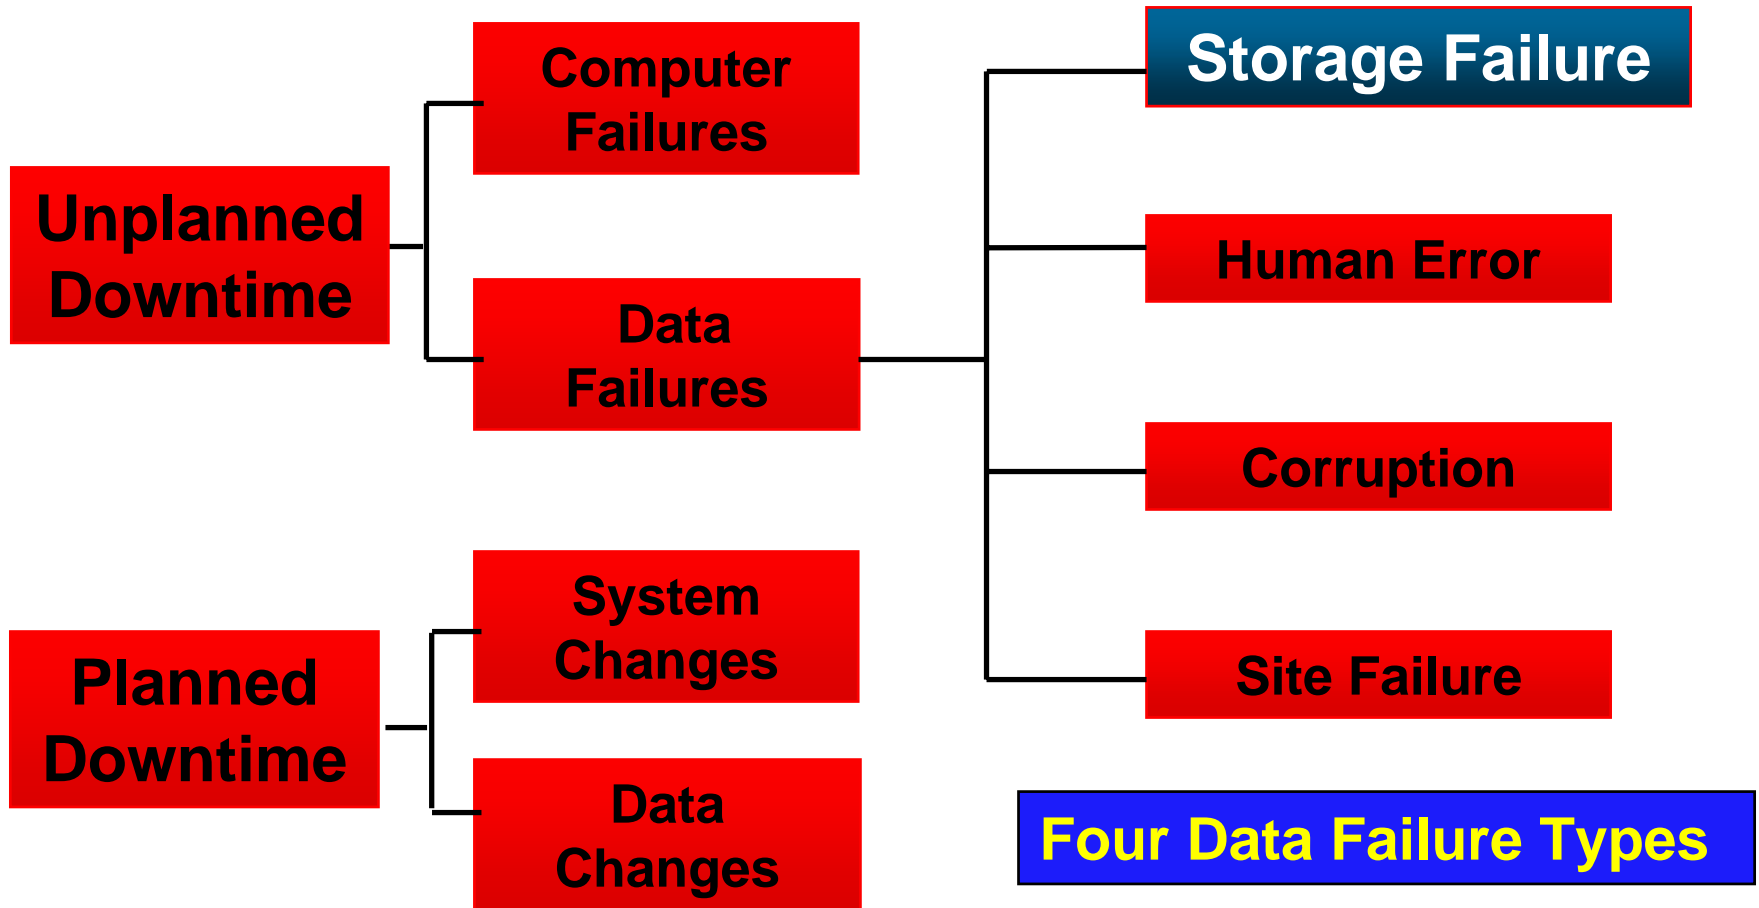
- 11g fine tunes performance, scaling, failover, management

**Designed to Tolerate Server Failures**



# Best Data Protection - at Lowest Cost



# Best Storage Protection - at Lowest Cost



# Storage Scale-Out with ASM



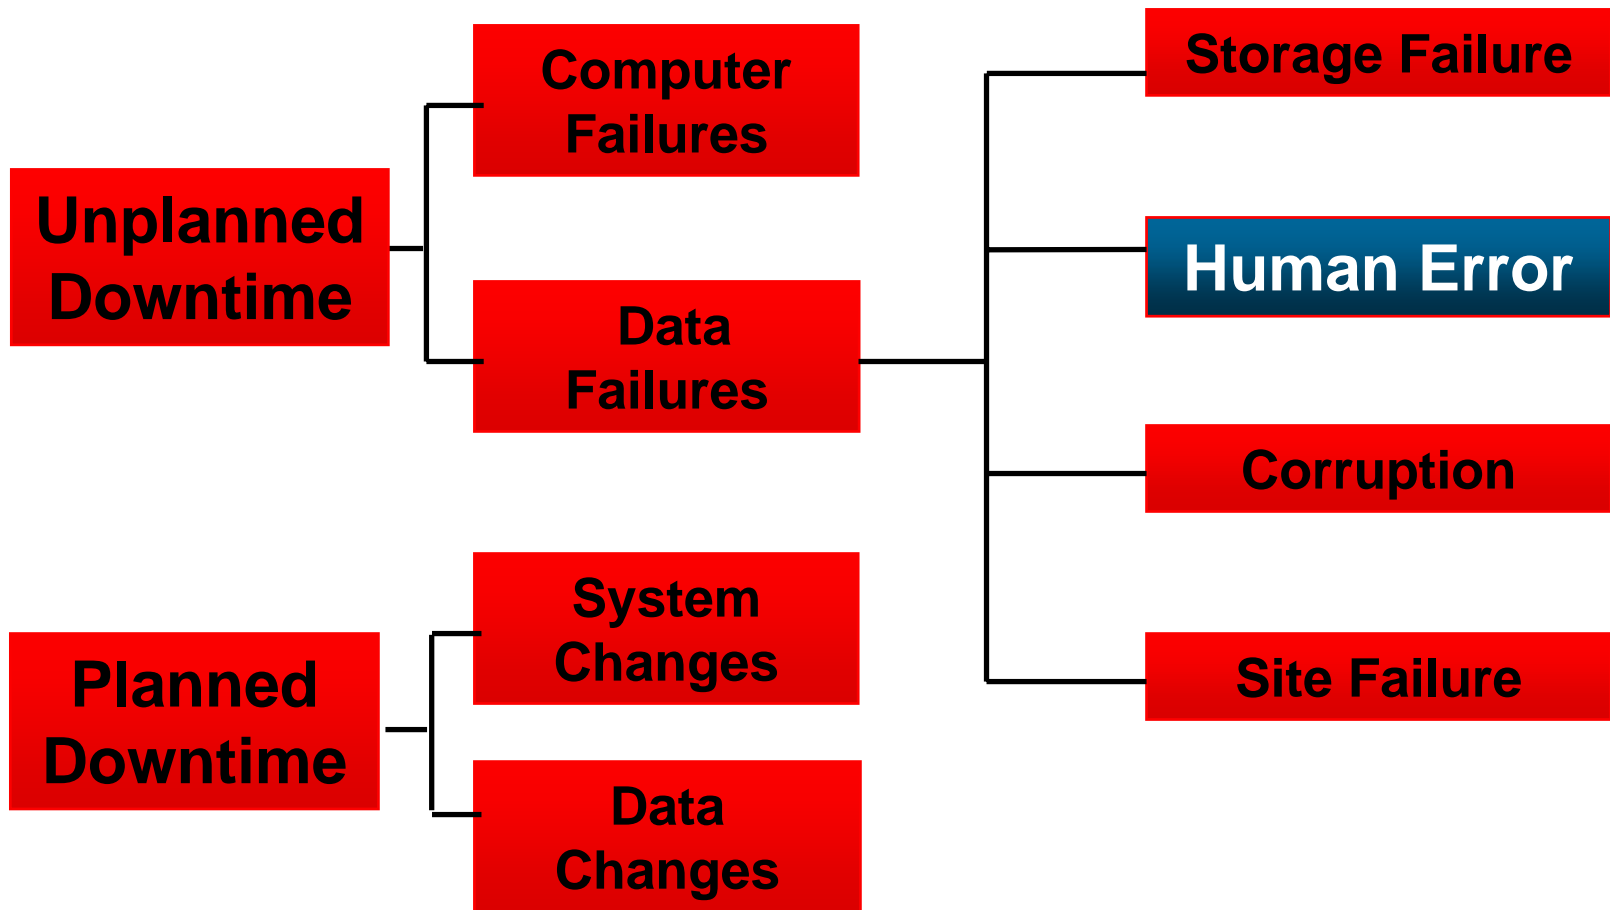
- ASM mirrors data across low cost modular storage arrays
  - Automatically remirrors when disk or array fails
- ASM is free



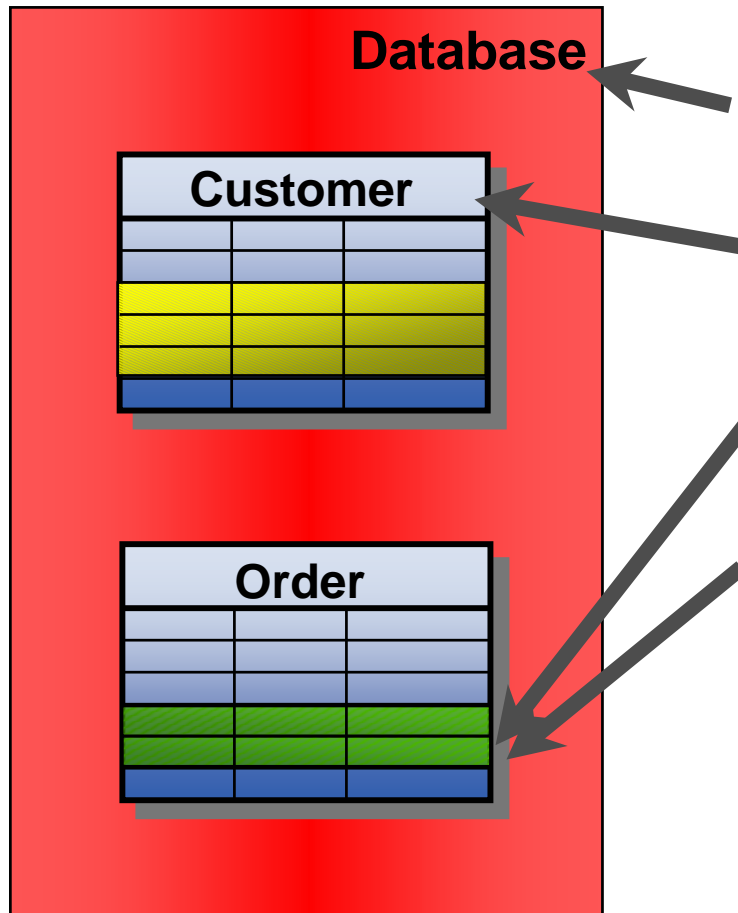
- 11g fine tunes recovery from corrupt blocks and crashed storage arrays

**Designed to Tolerate Storage Array Failure**

# Best Human Error Protection - at Lowest Cost



# Error 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 Query** - restore individual rows
- 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

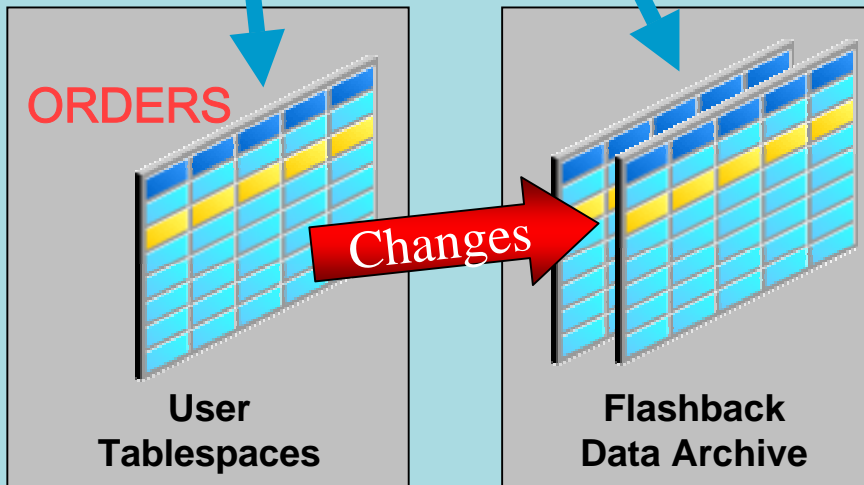


**Correction Time = Error Time + ~~f(DB\_SIZE)~~**

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

# Flashback Data Archive

Select \* from orders  
AS OF  
'Midnight 31-Dec-2004'

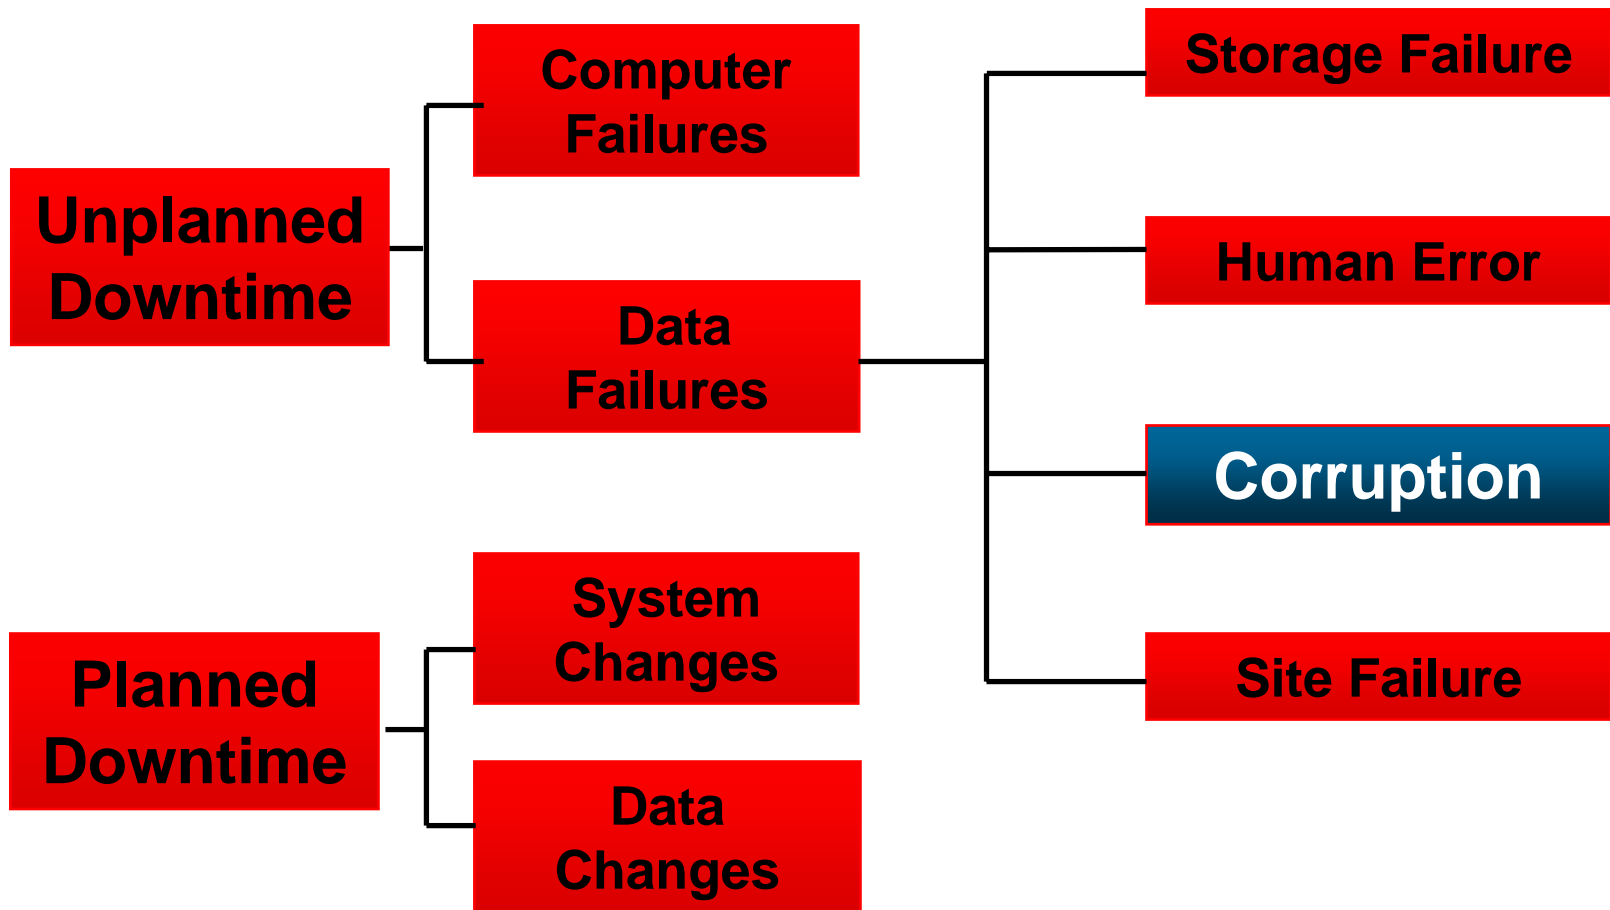


Oracle Database

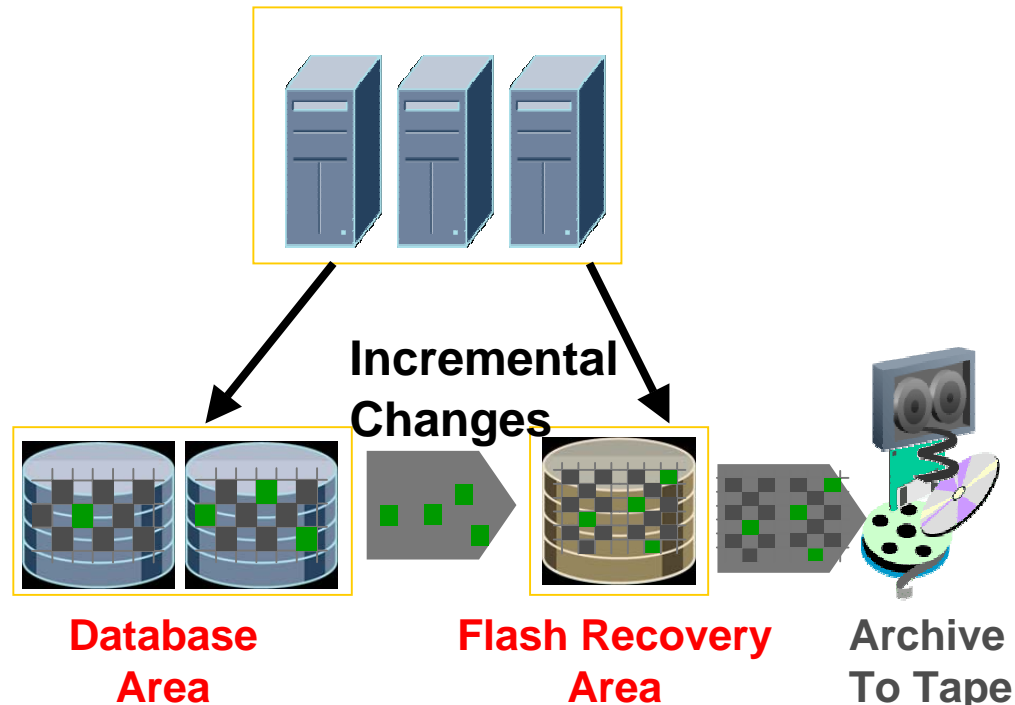
**Total Recall**

- 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

# Best Corruption Protection - at Lowest Cost



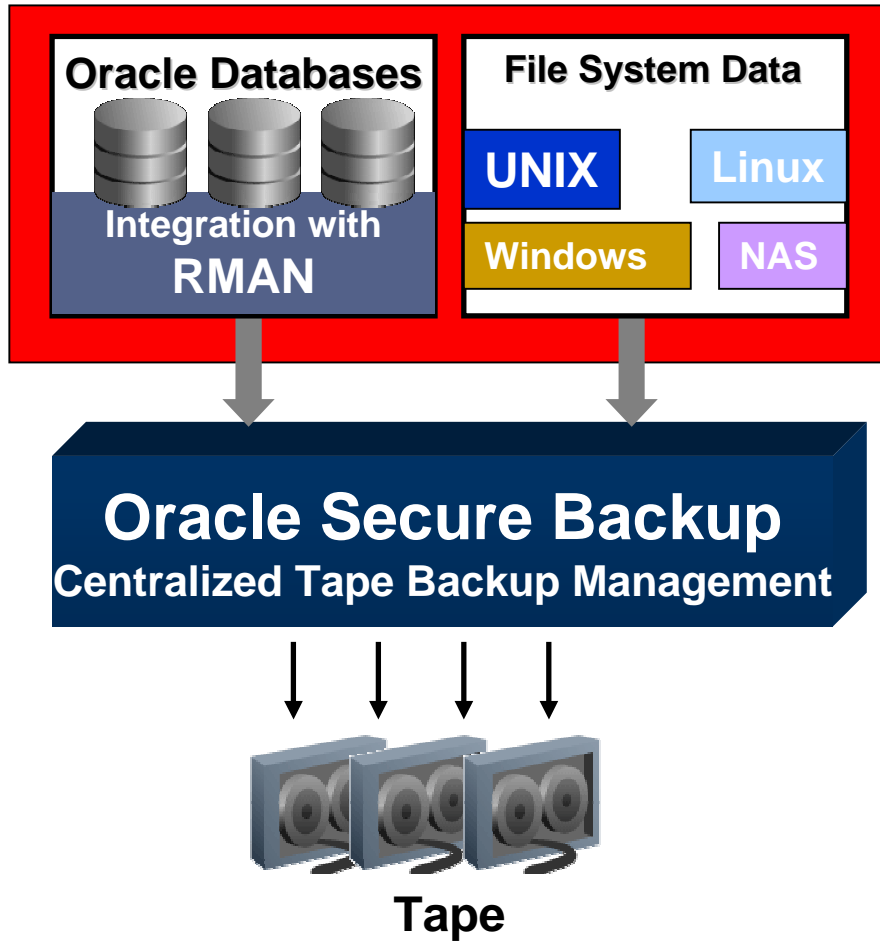
# Automated Disk Backup



- 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



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

# Offsite Backups in the Cloud


Database Files



RMAN



Local Disk Backup



RMAN Tape Interface



Secure Backup Cloud Module

Compression

Encryption



Amazon S3

- New Oracle Secure Backup module to move database backups to Amazon Cloud
  - Works with 9i and higher DB versions
- More Reliable than tapes, Faster Restores
- Eliminates tape backup and offsite tape management overhead
- No usage charge from Oracle
  - Pure software license model

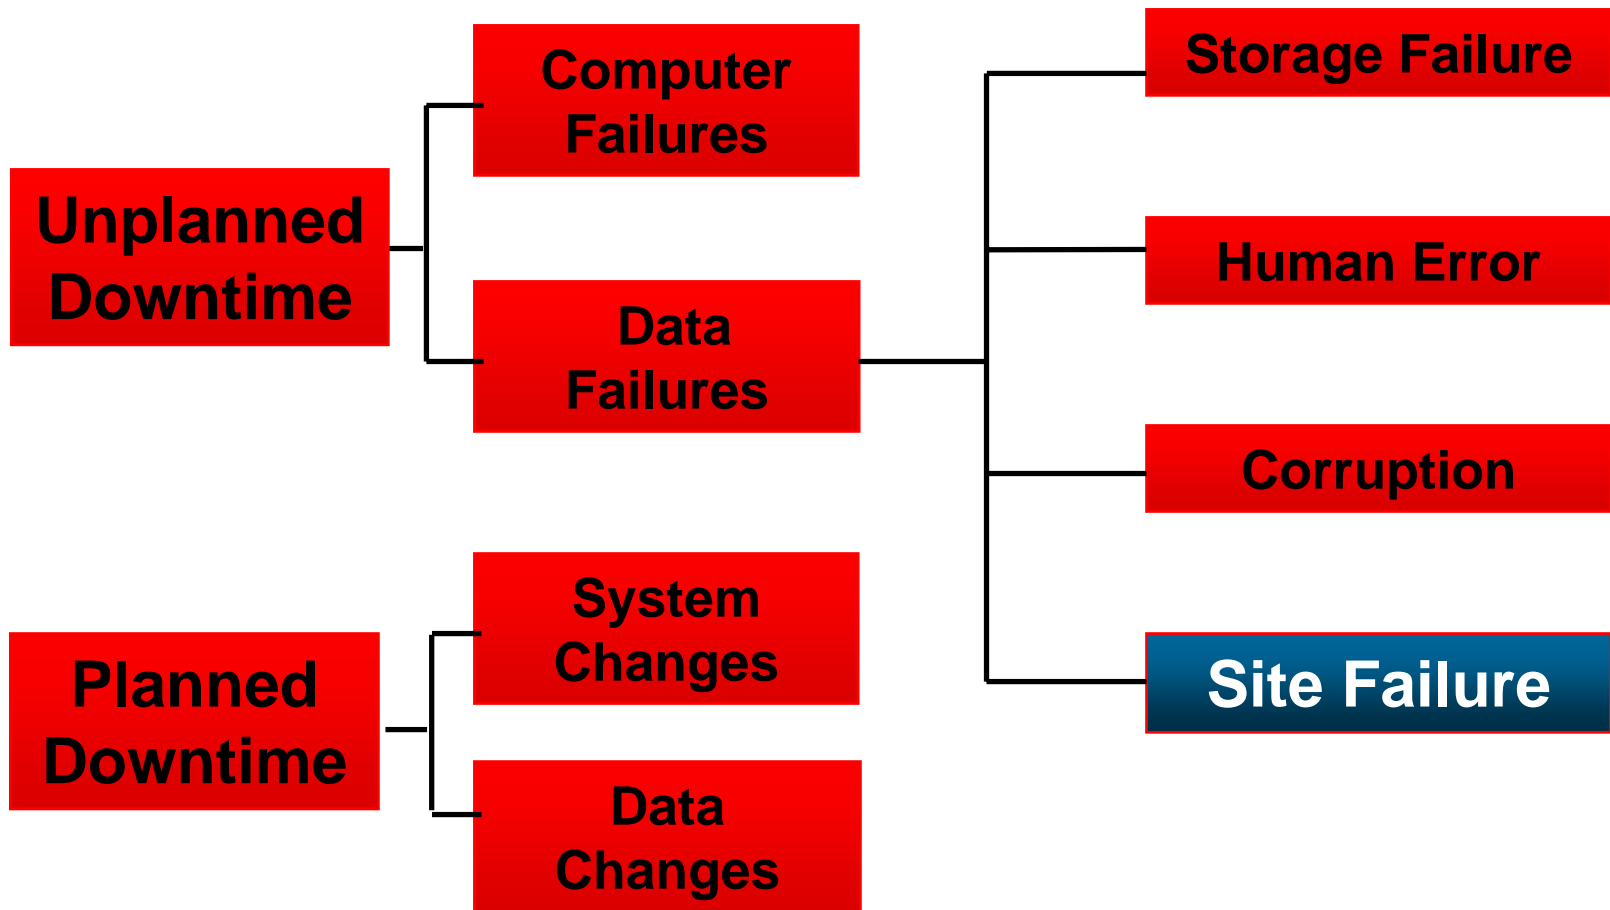
# Backup in the Cloud – Performance

**Viable Even For Large Databases**


DB Size (GB)	Full DB Backup Time	Incremental Backup Time	Monthly Amazon S3 Cost
500	4 Hours	30 Minutes	\$200
300	2 Hours	15 Minutes	\$120
100	40 Minutes	5 Minutes	\$40

**On-premise DB; Compressed Parallel Backups**

# 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



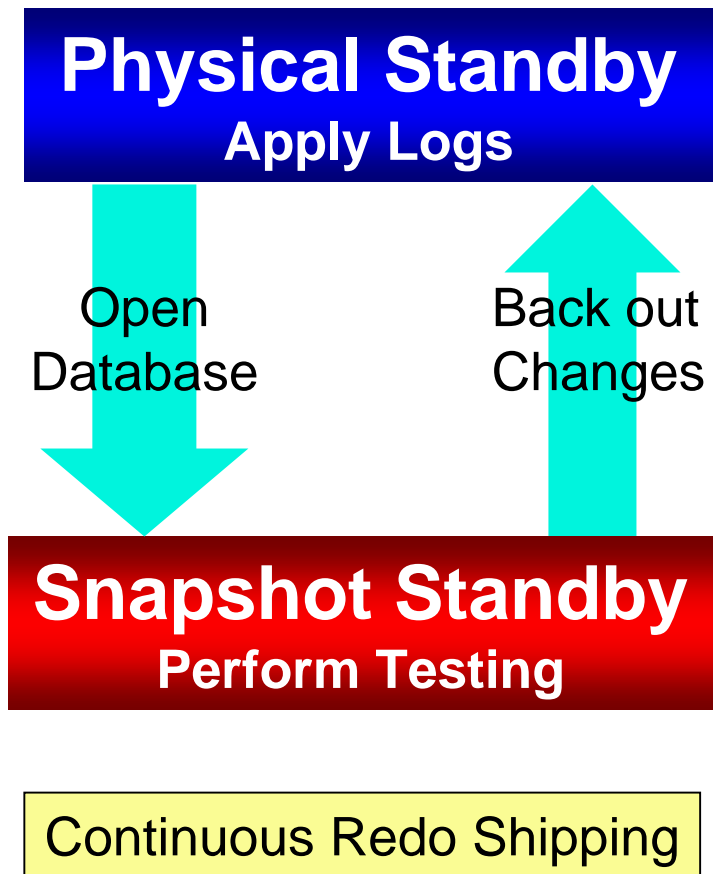
- Synchronous or asynchronous log shipping
  - Corruptions don't propagate
- Compared to remote mirroring, Data Guard provides better:
  - Data protection, network usage, hardware utilization
- Free with EE and works on low cost servers and storage
  - Thousands of production customers

# Automatic Failover to Standby (10.2)



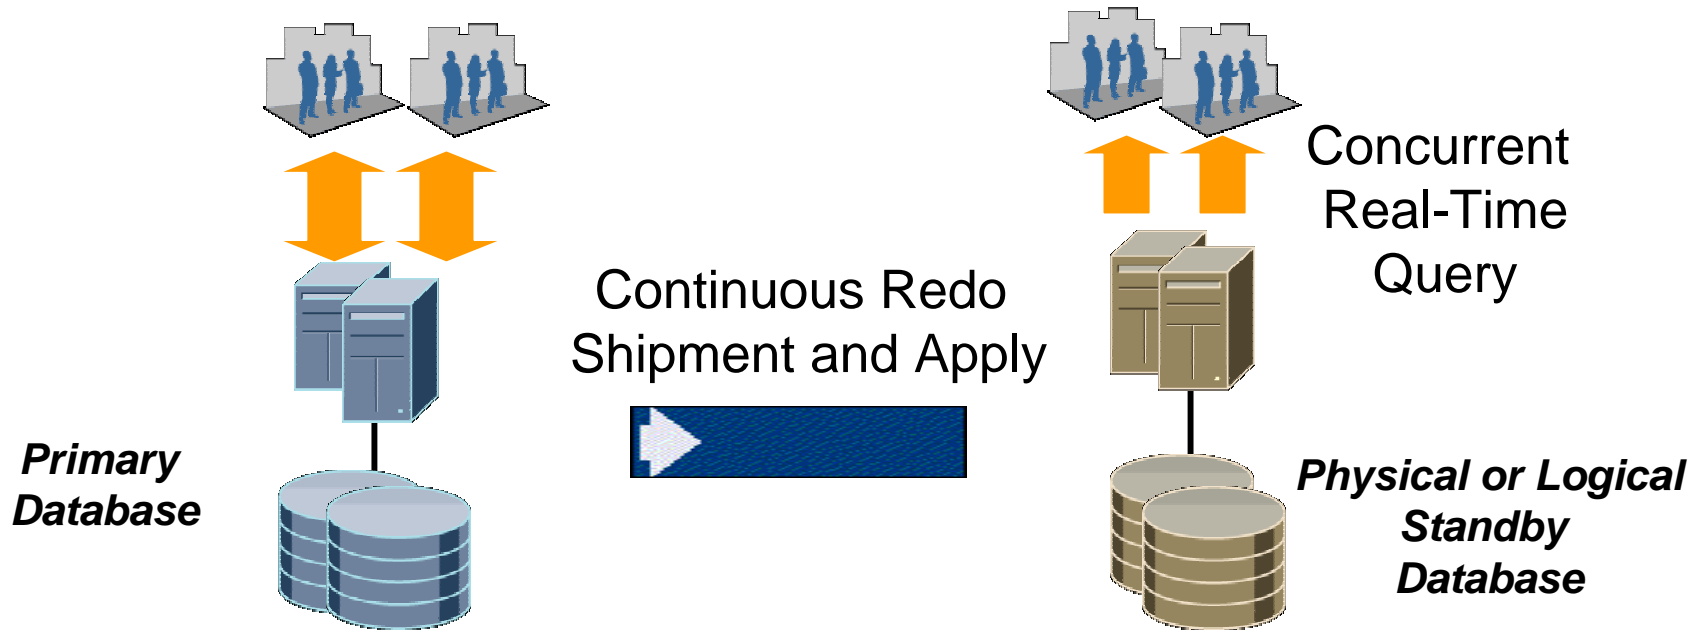
- **Fast-Start Failover** implements automatic failover to standby database
  - Zero data loss
  - Triggered by failure of Site, Hosts, Storage, Network, or Database
- Failover occurs in seconds

# Snapshot Standby



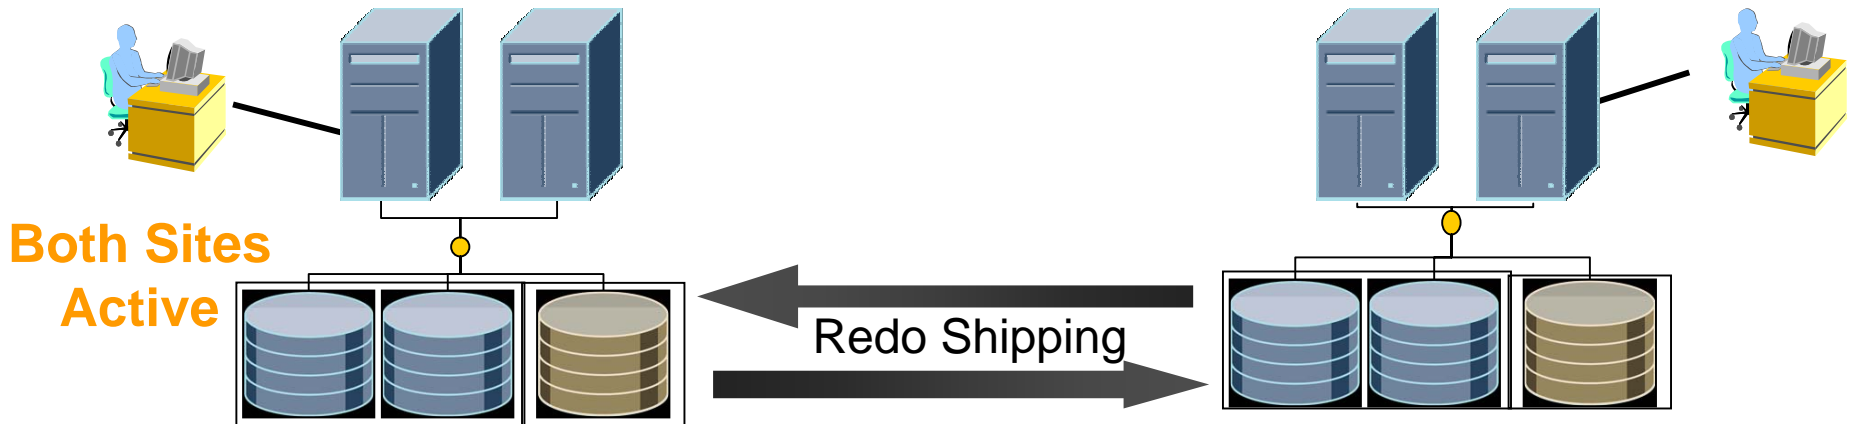
- Use Standby Database for testing and development
  - Eliminates cost of DR
  - Justify full scale test and DR environments by combining into one
- Preserves zero data loss while in test/dev mode
  - But no real time query or fast failover
- Similar to storage snapshots, but:
  - Provides DR at the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

# Active Data Guard



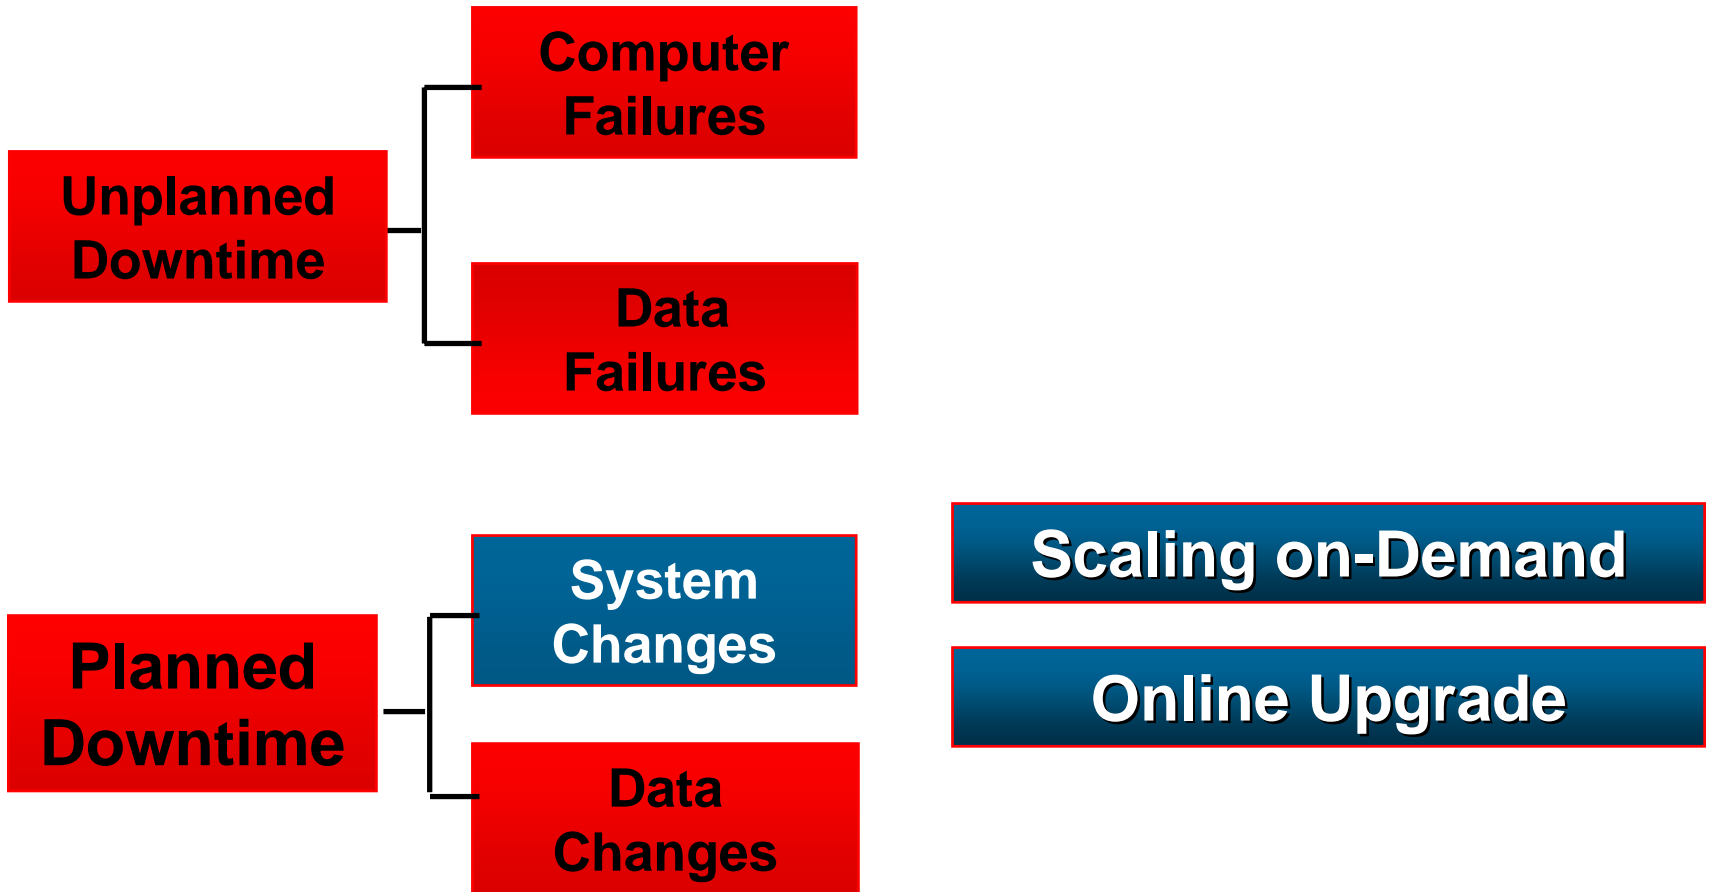
- Queries on standby database concurrent with apply
- Previously available with Logical Standby
- Available with Physical Standby in 11.1
  - Handles all data types, very fast, operationally simple
- **Eliminates cost of DR**
  - All hardware used for production load

# 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




- 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
    - OEL 4, RHEL 4, SUSE 10
  - Solaris 10 in 11.1.0.7
    - SunOS kernel patch 137111-04
  - HPUX in 11.1.0.7
    - 11iv3 (11.31) + [March 2008 Quality Pack + PHKL\_38038]
  - 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



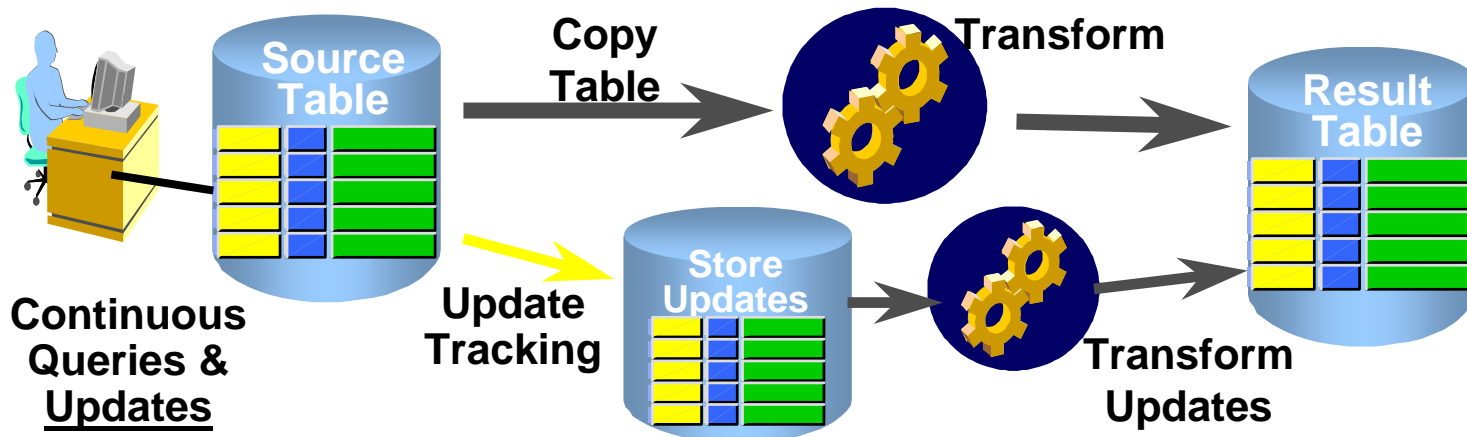
- Upgrade standby site first, then switch over users
  - Patch sets, major releases

# Best Data Changes




# Online 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, Application Server
    - And e-business suite and Siebel
  - 60 person year effort
    - Not marketing fluff

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

# Goals and Vision for Availability

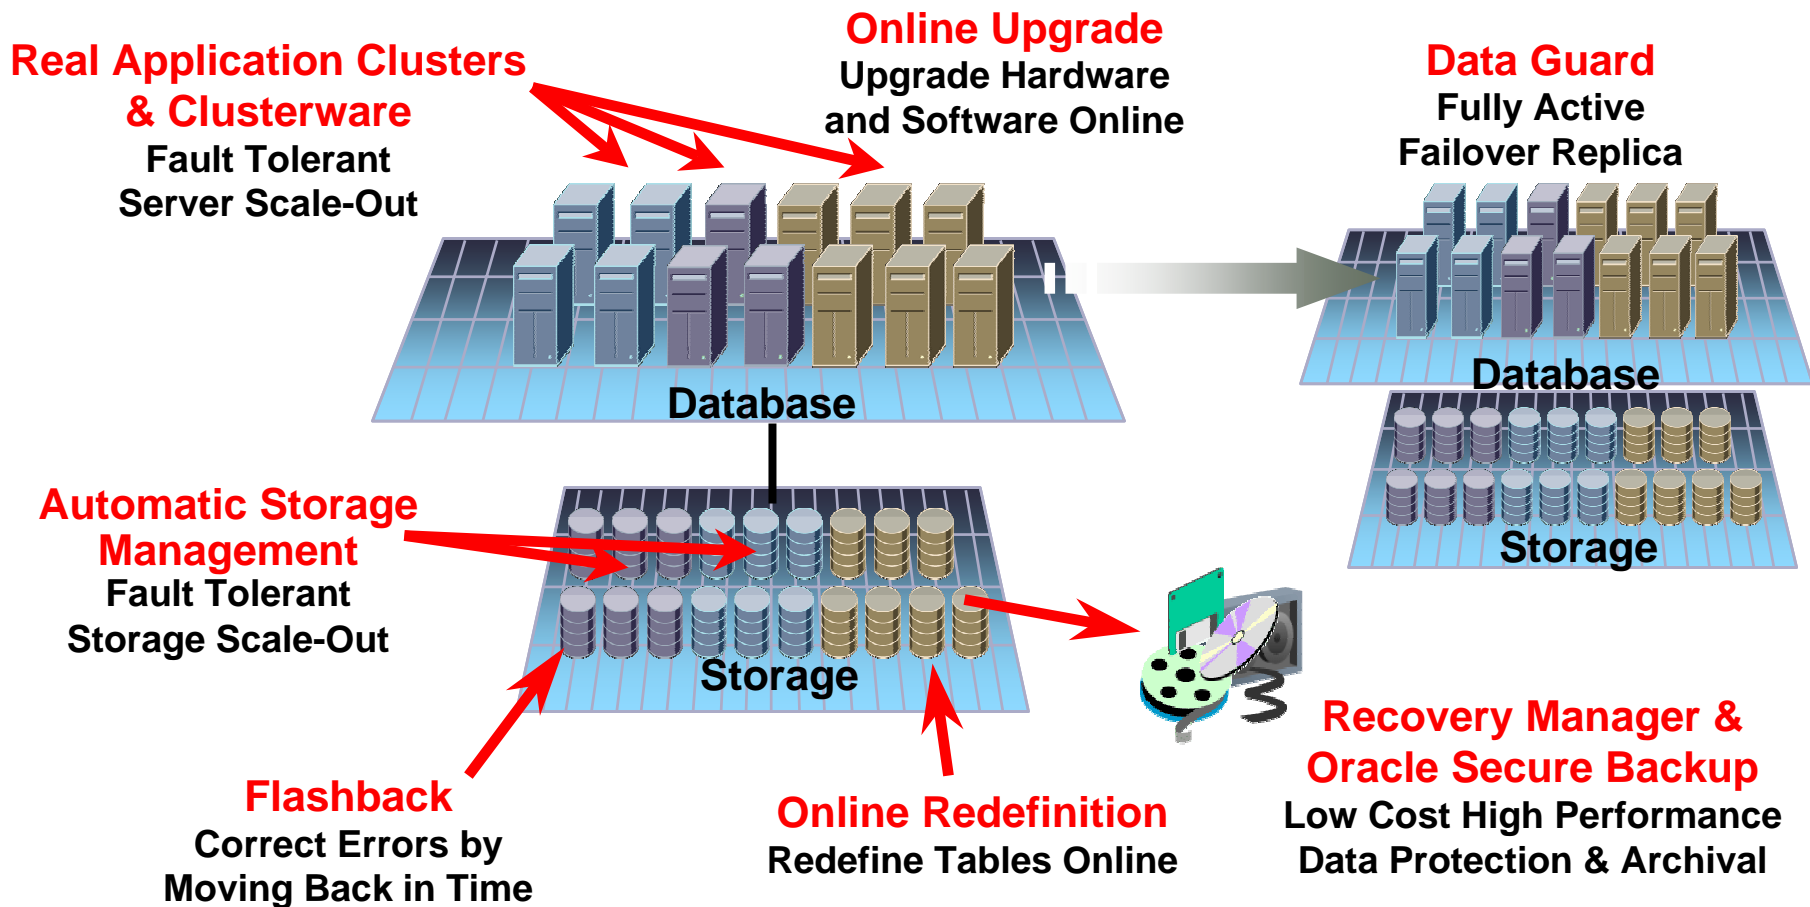
- Fault-Tolerant Grid-Based Infrastructure
- Application Oriented
- Complete, Integrated, Automated, Simple
  - Yet Open, Standard, Pluggable
  - and Low Cost

**Maximum Availability Architecture**

# 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**



# Allstate Insurance

## Oracle MAA Powers Next Generation Claims Applications


### Primary Site



Data Guard



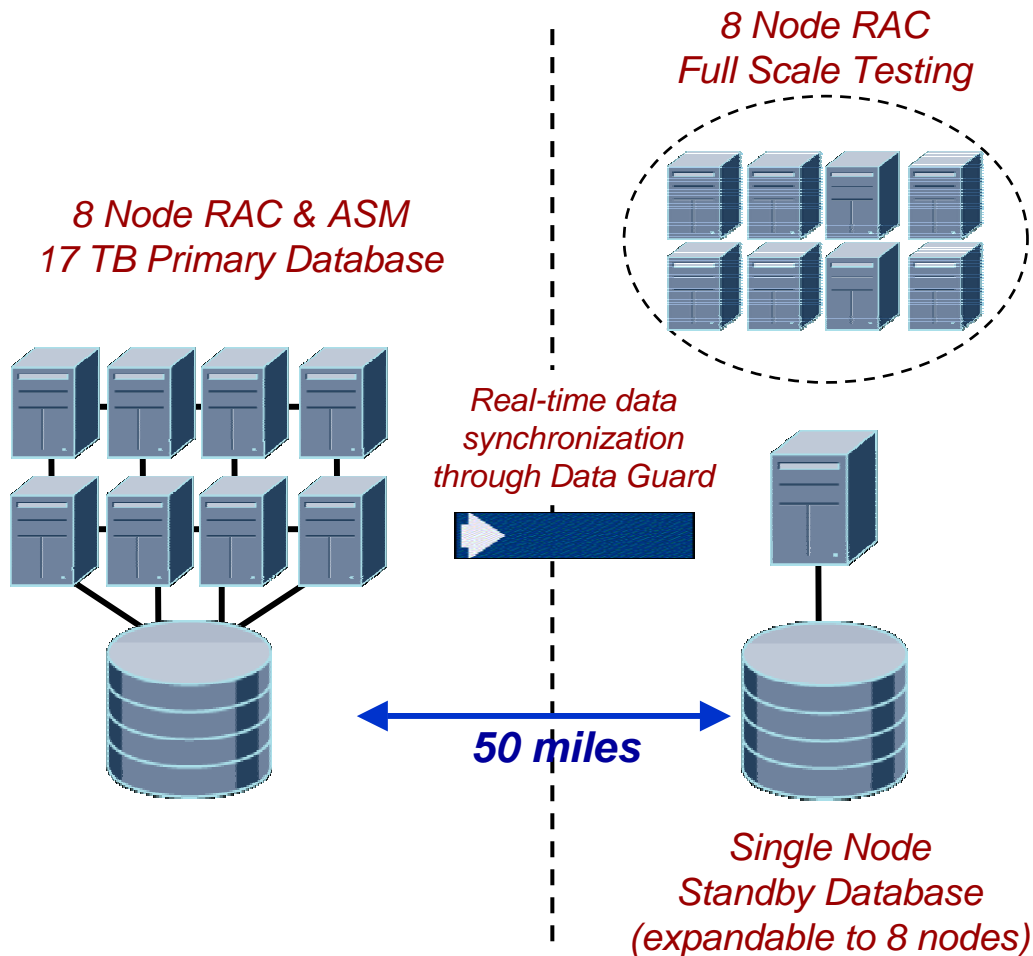
### Remote Recovery Site



- Next-generation architecture
  - Consolidation of 70+ legacy claims applications running on multiple mainframe & SQL Server databases
  - Based on Oracle MAA
- MAA benefits
  - Scales to handle peak workload
  - Most failures transparent
  - Fast recovery where not transparent
  - Best protection from physical corruptions
  - Offload operational reporting from OLTP
  - Rolling upgrades minimize or eliminate downtime for planned maintenance

# Pacific Gas & Electric (PG&E)

## Oracle MAA Provides Better Availability than Mainframe

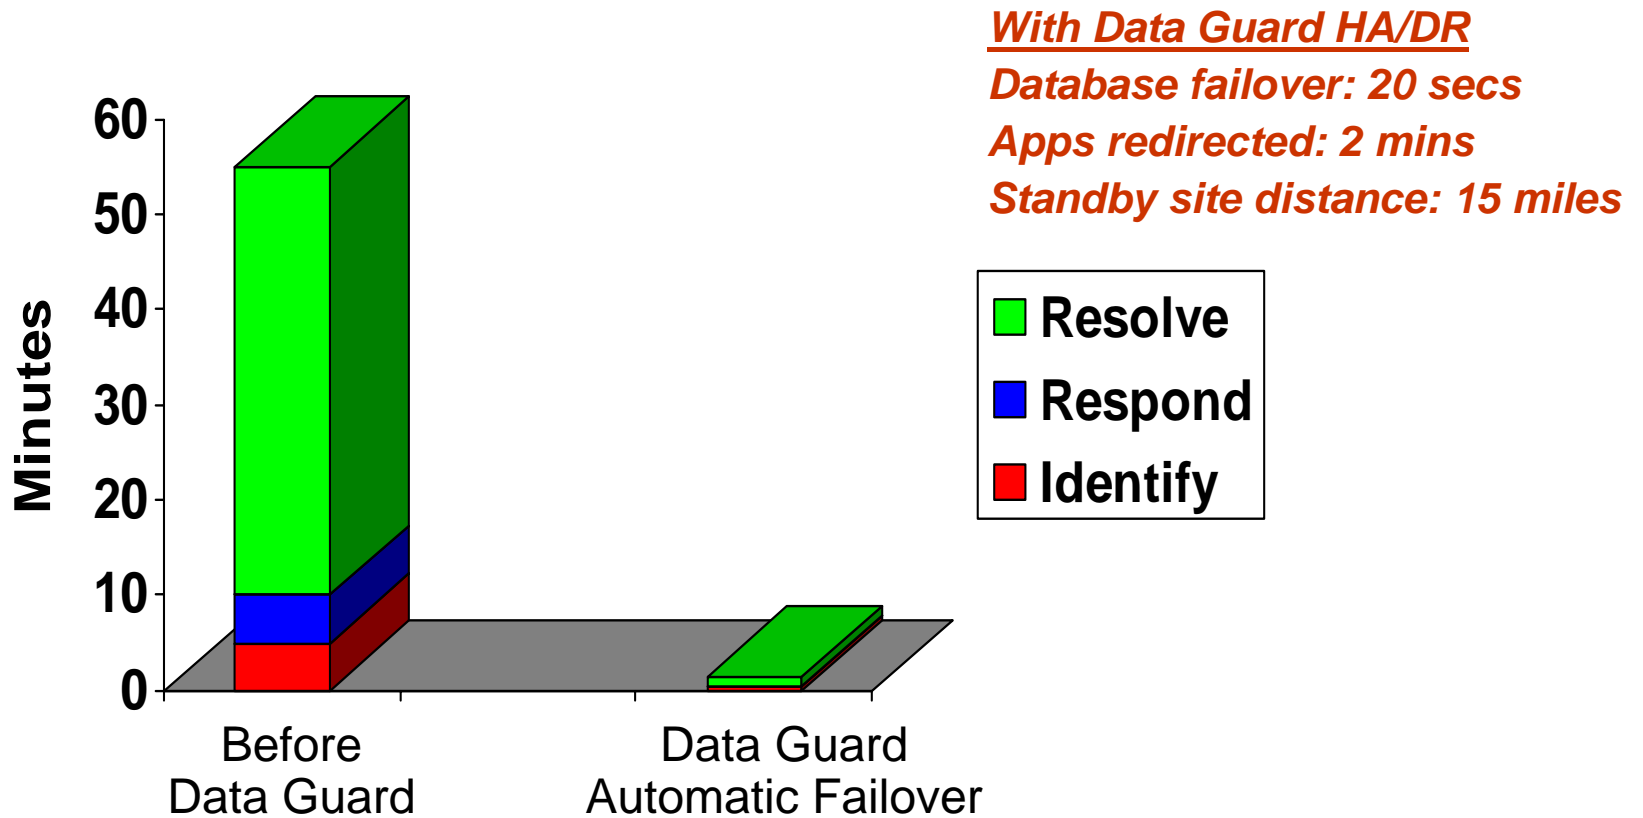


- #1 utility provider for Northern and Central California
- Customer Care & Billing System for 13 million users : DB2 / mainframe migration to Oracle MAA on commodity infrastructure
  - Daily load: 1+ billion select, 100 million insert/update/delete
  - 2300 concurrent online users, 300,000+ bills per day
- MAA components: **RAC, ASM, Data Guard, Flashback, RMAN, Grid Control:**
  - 28% increase in throughput
  - 24x7 uptime SLA
  - \$5M annual cost savings

# Amazon.com

## *High Availability Integrated with Disaster Recovery*

### End-to End Failover Time



# HA Sessions, Labs, Demos From Oracle Development

## Mon, Sep 22

- 2:30 pm - Database 11g: 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
- 5:00 pm - Data Guard & Real Application Testing, Moscone 102
- 5:00 pm - EM in Secure MAA, Moscone West 2001

## Wed, Sep 24 (contd.)

- 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



# 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 10gR1 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 10gR1 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



# ***9i* Availability Features**

# New Oracle 9iR2 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 9iR1 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 9iR1 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

# Which Model for Availability

**ORACLE®**

Vs.



**Database**

**Infrastructure**

- Protects and recovers business objects
- Comprehensive
  - Handles Human Errors, Application and Data Changes, Corruptions, Hardware Changes
- Efficient and Hardware Independent
  - Scale-out low on cost hardware
  - Full hardware utilization
- Integrated and Automated

- Works across applications
- Simple for simple failures

# Which Database for Availability

**ORACLE®**

Vs.



**Database**

**Other DBs**

- Oracle has Best-of-Breed technology
  - RAC is best Clustering
  - Data Guard is best Standby
  - Online Redefinition is best for Application Changes
  - Flashback is best Error Recovery
  - RMAN is best backup
  - Integration to application level
- Oracle is Lowest Cost
  - Scale-out servers
  - Scale-out storage
  - Fully active – no idle failover hardware