ORACLE®

Oracle Database 11*g*: Unbreakable Meets Best Practices

Tim Chien
Senior Product Manager, Database High Availability
Timothy.Chien@oracle.com

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 remain at the sole discretion of Oracle.

Agenda

- Oracle Database High Availability (HA)
- HA Enhancements in Oracle Database 11*g*
- Maximum Availability Architecture (MAA)

Traditional Database HA

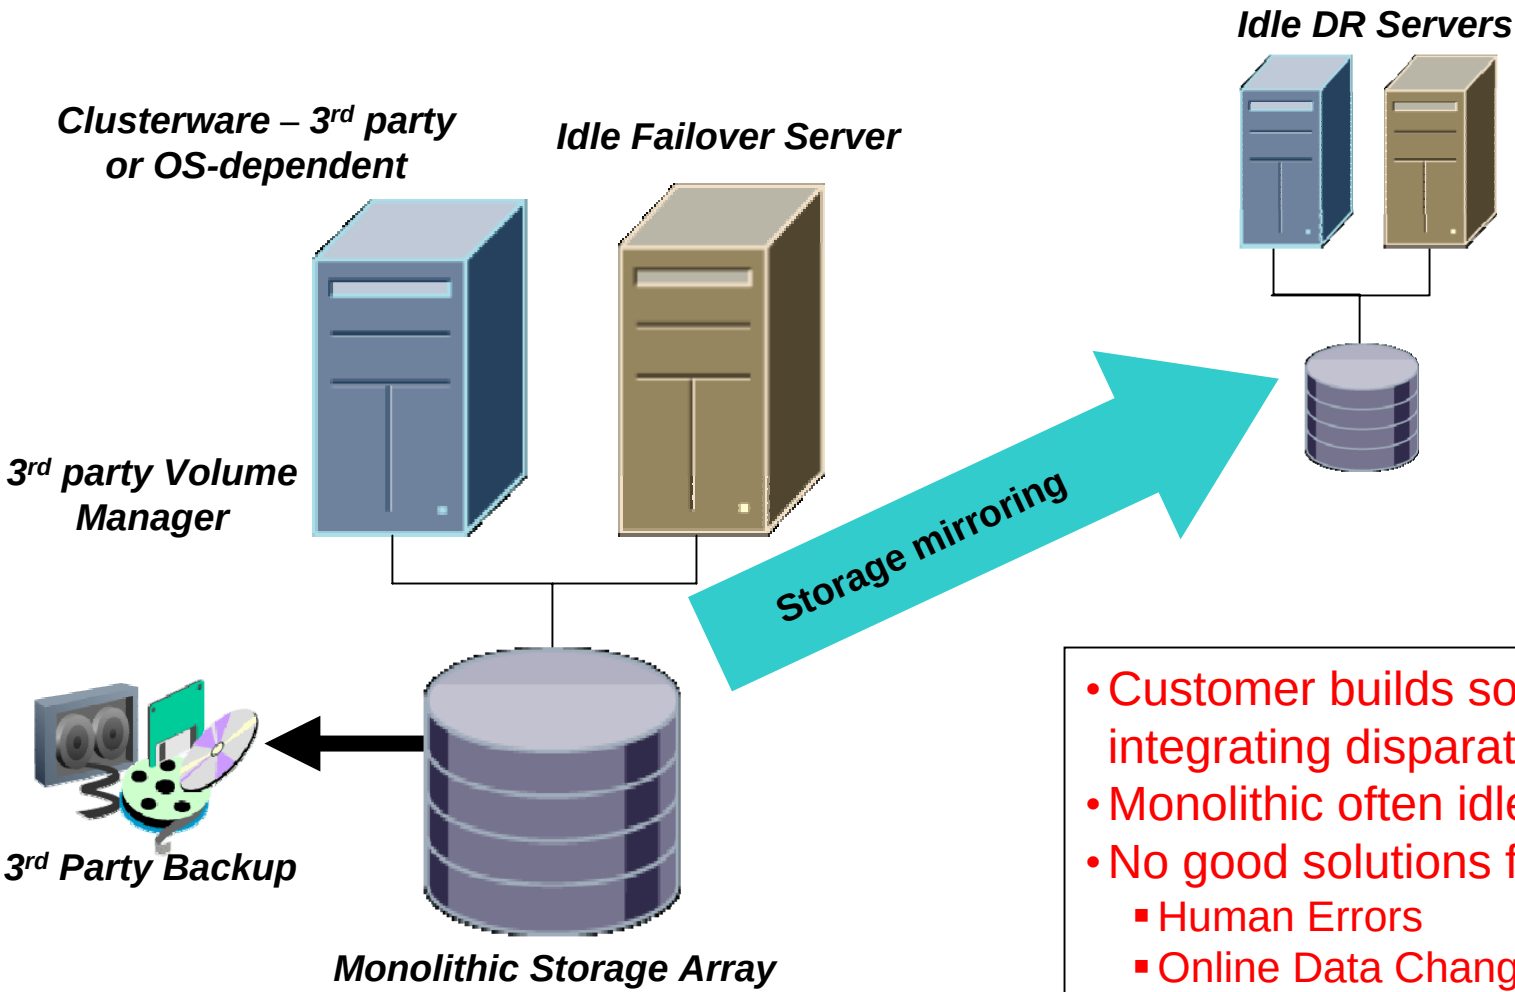

- Customer builds solution by integrating disparate components
- Monolithic often idle hardware
- No good solutions for:
  - Human Errors
  - Online Data Changes
  - Software Upgrades

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

Oracle: Great Fit for a Scale-Out Architecture

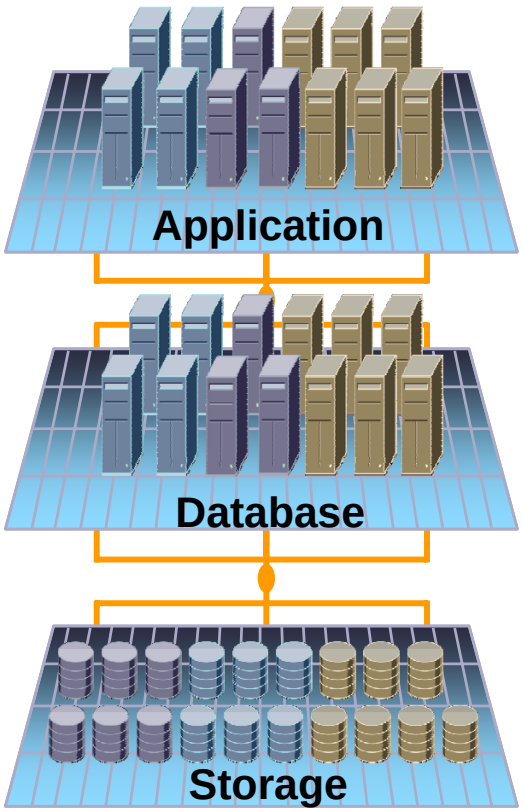

- Scale-Out architecture
  - Commodity hardware building blocks
  - Inherently highly scalable & redundant
- Scalability & Availability responsibility moves out of hardware/OS to scale-out savvy software
  - First Web & Application server tiers
    - Application servers
  - Then DB tier
    - Shared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

Oracle's Integrated HA Solution Set

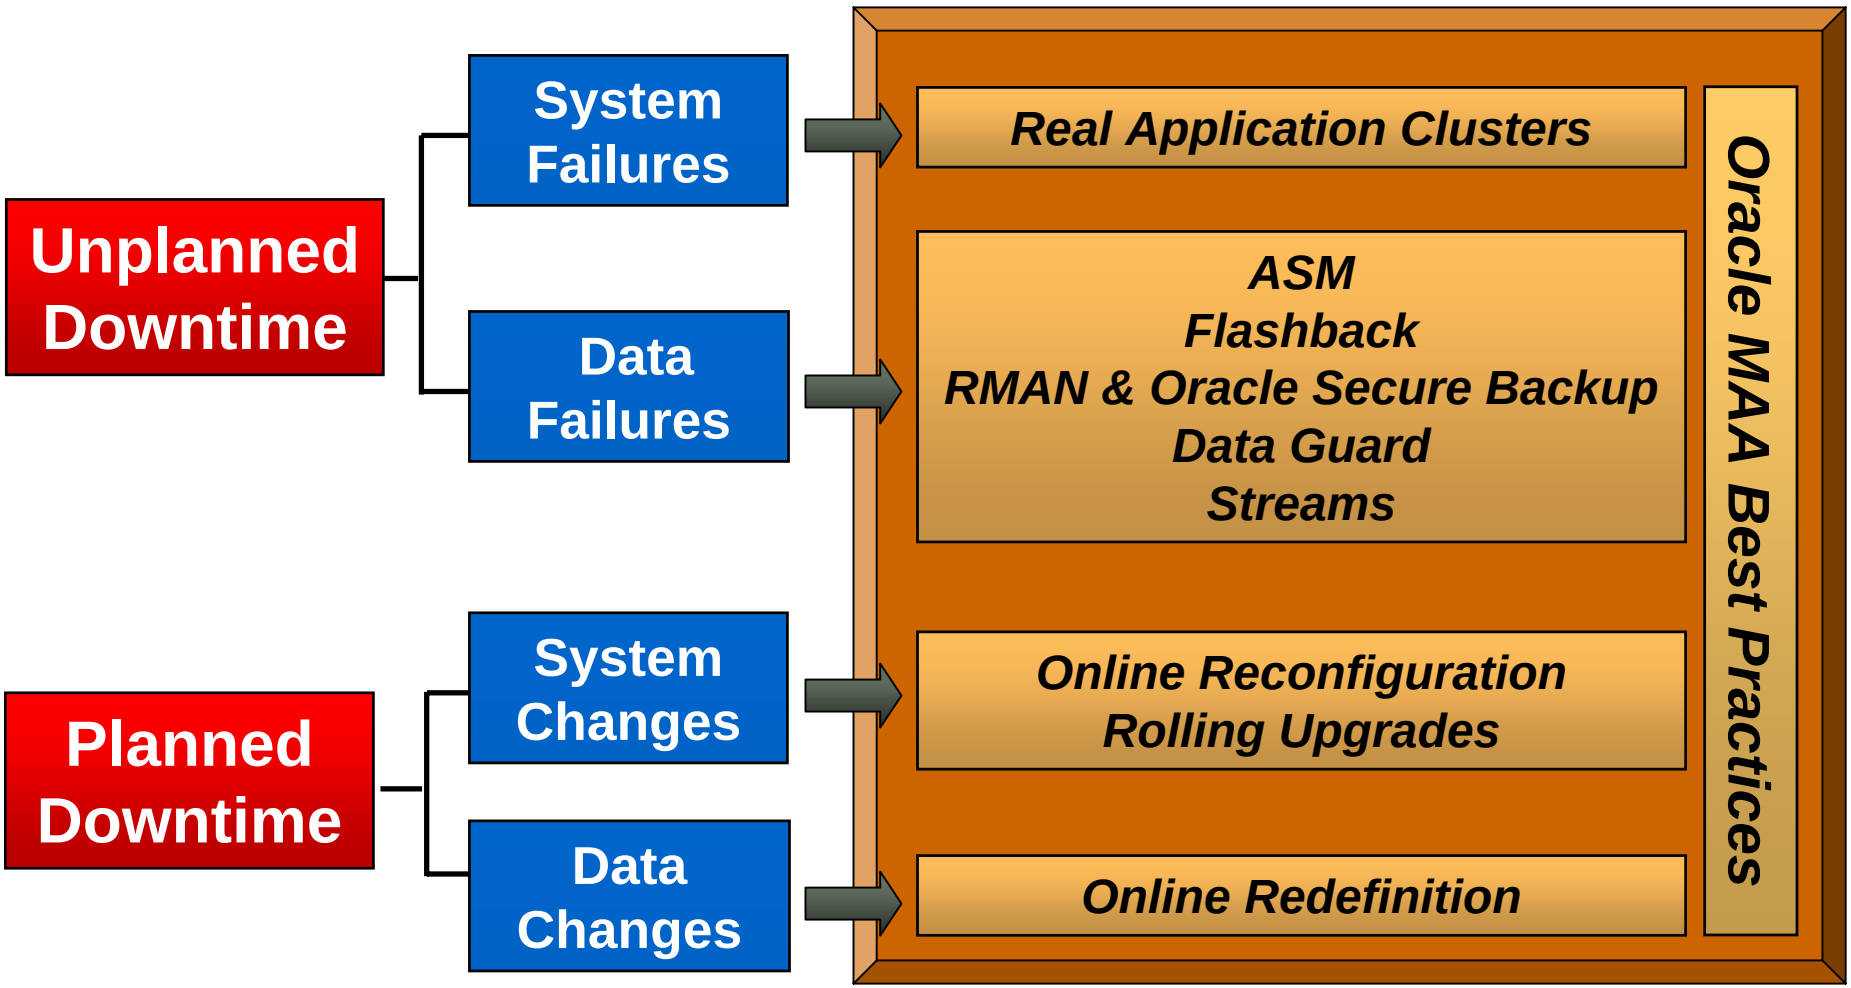

Oracle HA: Customer Success Stories

- ADT Security Services - Using Data Guard SQL Apply Across a Wide Area Network
- Amadeus - Using Data Guard for Disaster Recovery & Rolling Database Upgrades
- Amazon.com - Automatic Failover using Data Guard Fast-Start Failover
- Banknorth Group, Inc. - Using the Snapshot Capabilities of Flashback Technologies
- CGI - Helps Major North American Oil & Gas Company Save $500K with RMAN
- ChevronTexaco - RMAN DUPLICATE – DBA Time Saver to the Rescue
- Chicago Stock Exchange - Expects 171% ROI in Five Years from Oracle Enterprise Grid Computing
- Colgate-Palmolive - Increased Performance with RMAN
- CSX - Online RMAN Backups Protect over 16TB of Data
- Dell - Dell Consolidates European Support System with Oracle Enterprise Grid on Dell
- Fannie Mae - Supporting 835 transactions per second & Zero Data Loss Protection in Oracle Database 10g
- First American Real Estate - Using Data Guard
- Hartford - Incrementally Updating Transportable Tablespaces using RMAN
- Kemira GrowHow Ltd, UK - Replacing Outsourced Disaster Recovery Services with Oracle Data Guard
- KLM - KLM Royal Dutch Airlines Eliminates Costly Downtime with Grid Solution
- NeuStar - Synchronous Zero Data Loss Protection with Production and Standby Databases Separated by 300 Miles
- Ohio Savings Bank - Oracle Database 10g - Maximum Availability Architecture & Zero Data Loss
- Oracle Global IT - Oracle E-Business Suite with Data Guard over a WAN
- Purdue Pharma L.P. - Surviving Media Disaster with RMAN
- ReserveAmerica - Capitalizing on Oracle 10g Flashback Technologies
- Starwood Hotels - RMAN in Oracle Database 10g Best Practices for Maximum Benefit
- Swedish Post - Extending the DR system using reporting capabilities of Data Guard SQL Apply
- TALX Corporation - Increased Performance with RMAN and Oracle Database 10g
- Trilegiant - Online RMAN Backups Protect over 8TB of Data
- VP Bank - Using Data Guard SQL Apply to deploy content outside the corporate firewall

and many more* …

ORACLE

** http://www.oracle.com/technology/deploy/availability/htdocs/HA_CaseStudies.html*

Oracle Database HA in 11*g*

- Goals:
  - ✓ **Minimize** downtime
  - ✓ **Utilize** all resources
  - ✓ **Scale** for growth
- Achieve these with an integrated, best-of-breed HA architecture

Best-of-Breed Server Protection

At Lowest Cost

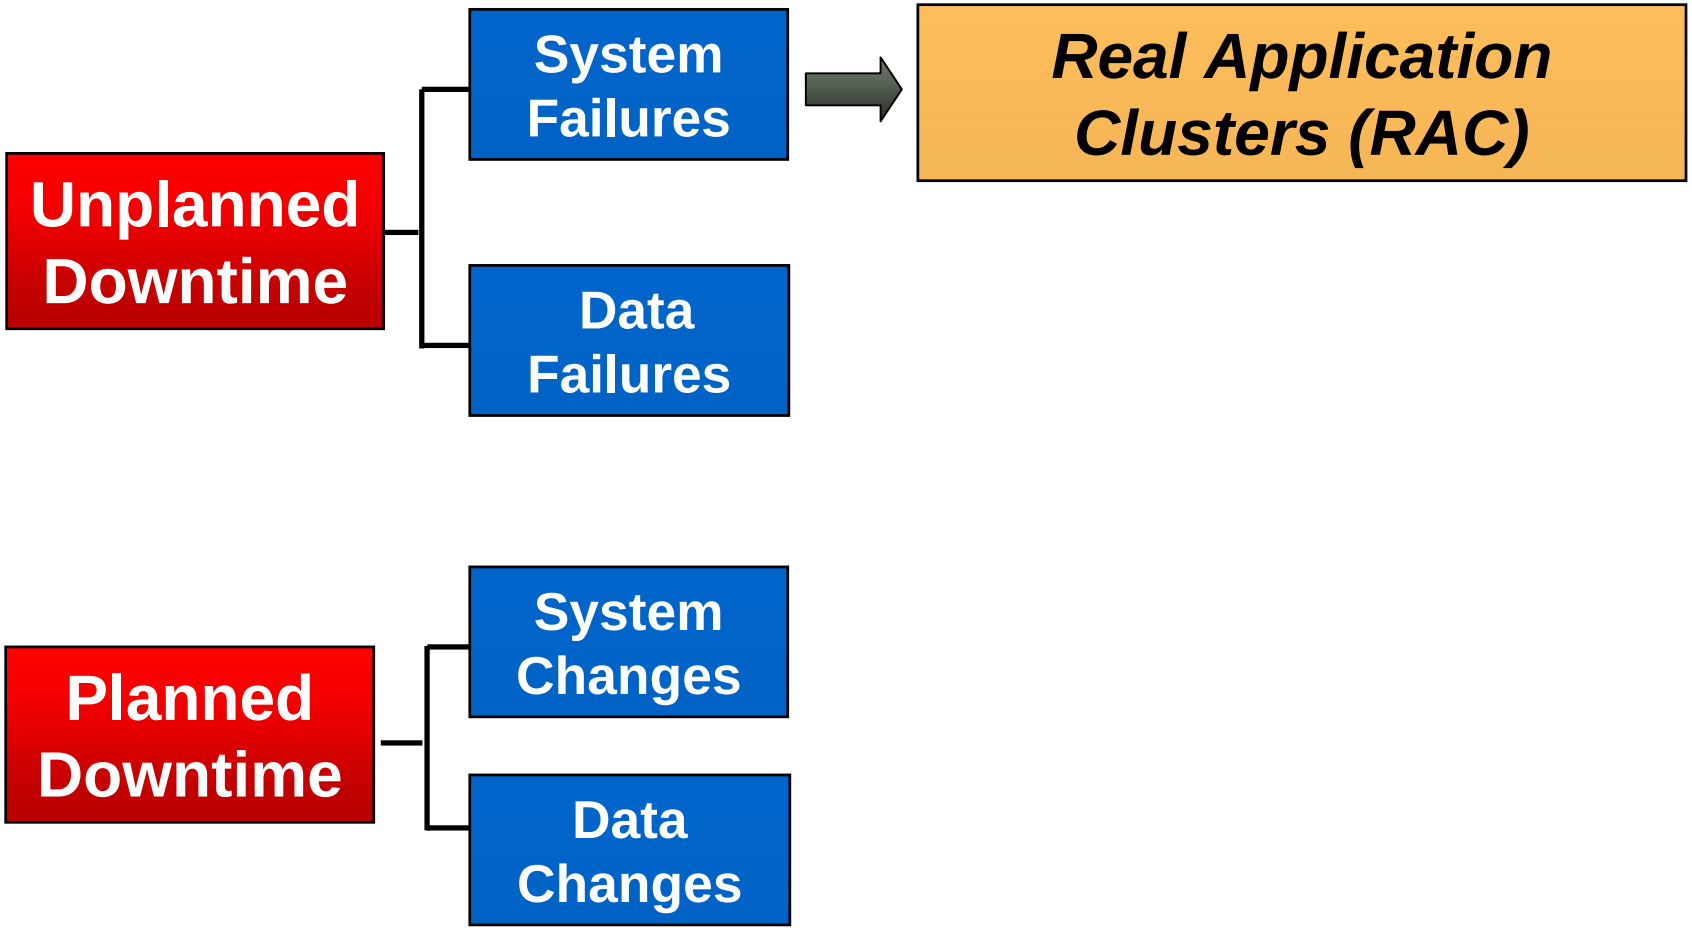

Server Scale-Out with RAC

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

ORACLE 11g DATABASE

- Fine-tuned performance, scaling, failover, management
- Enhanced, seamless integration with XA

Designed to Tolerate Server Failures

Best-of-Breed Storage Protection

At Lowest Cost

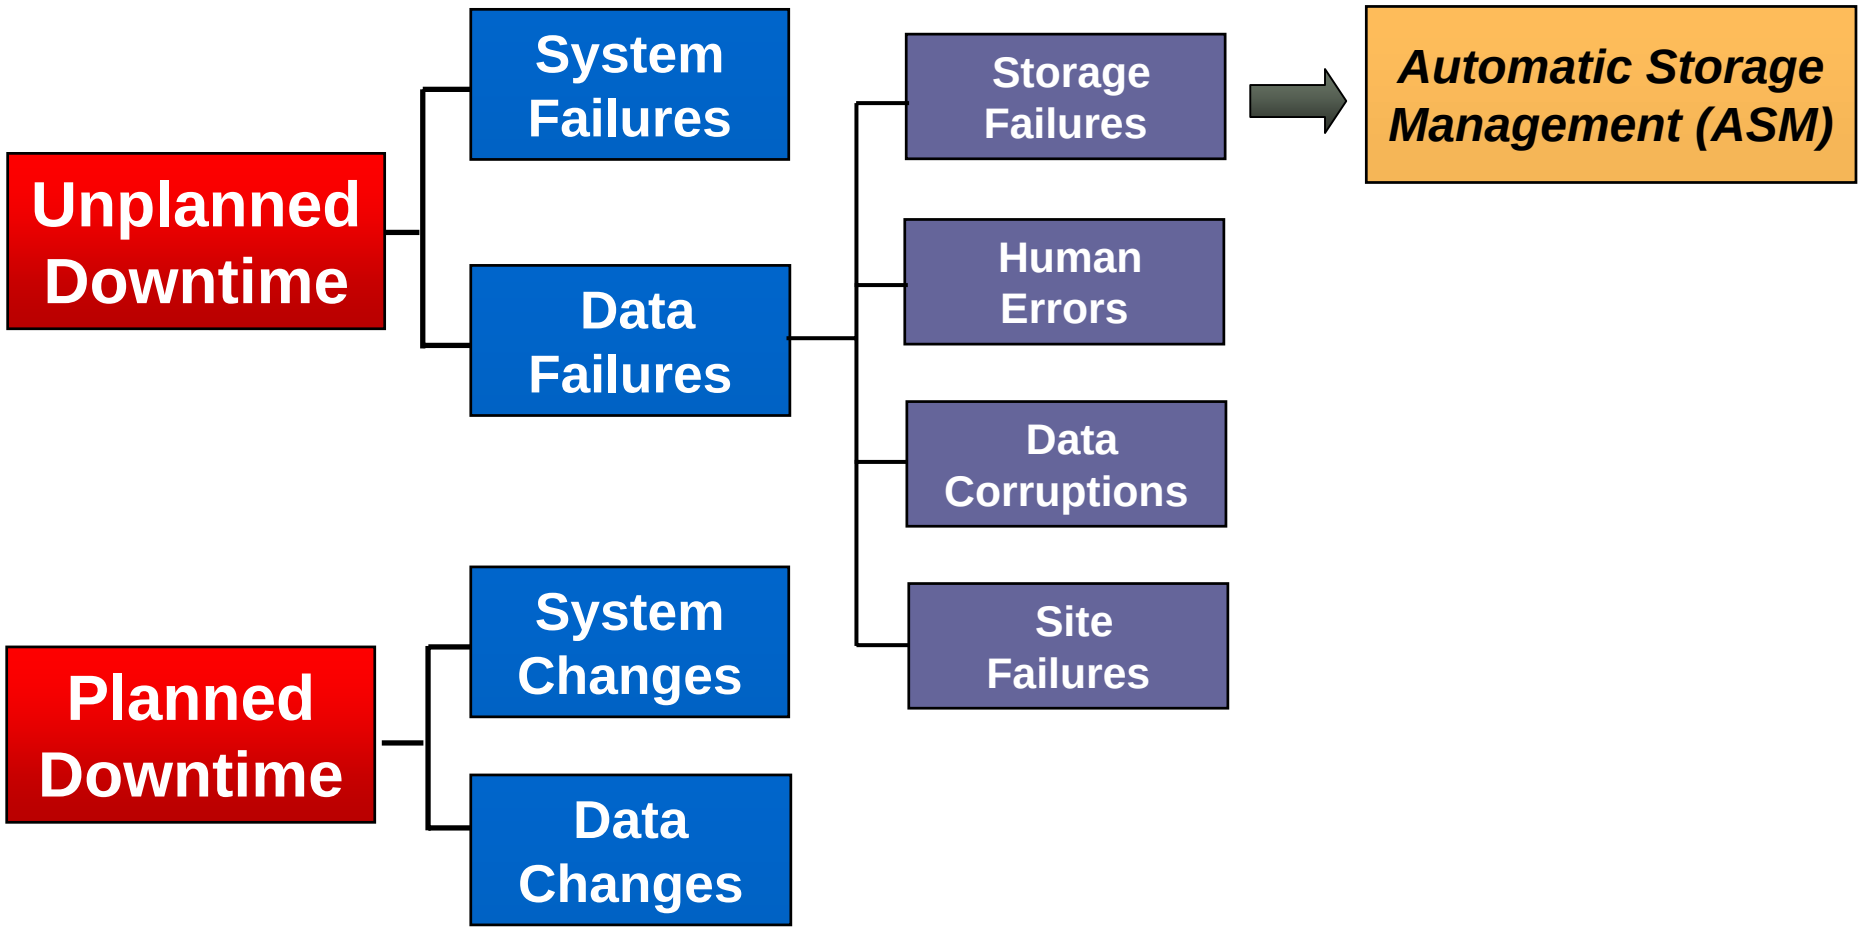

Data Mirroring with ASM

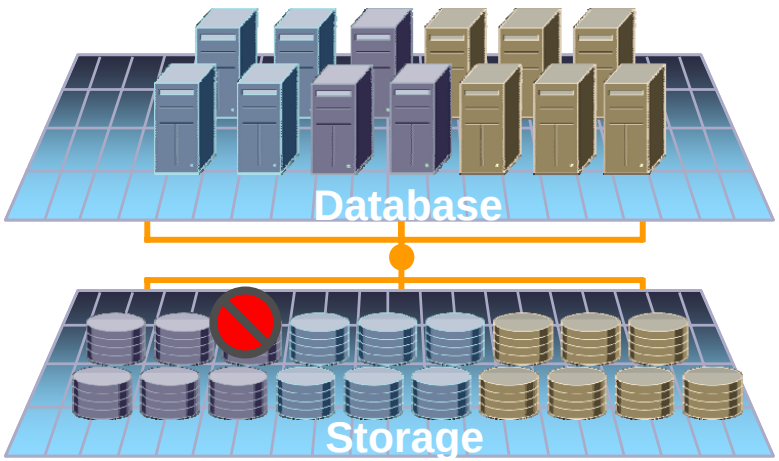

- ASM mirrors data across low cost modular storage arrays
  - Automatically remirrors when disk or array fails

ORACLE 11g DATABASE

- ASM Enhancements
  - Automatically repair corrupt blocks from mirror copy
  - Fast resync of mirror copy upon recovery from transient disk failures – uses only changed blocks
  - Rolling Upgrade for ASM instances

Designed to Tolerate Storage Array Failures

Best-of-Breed Human Error Protection

At Lowest Cost

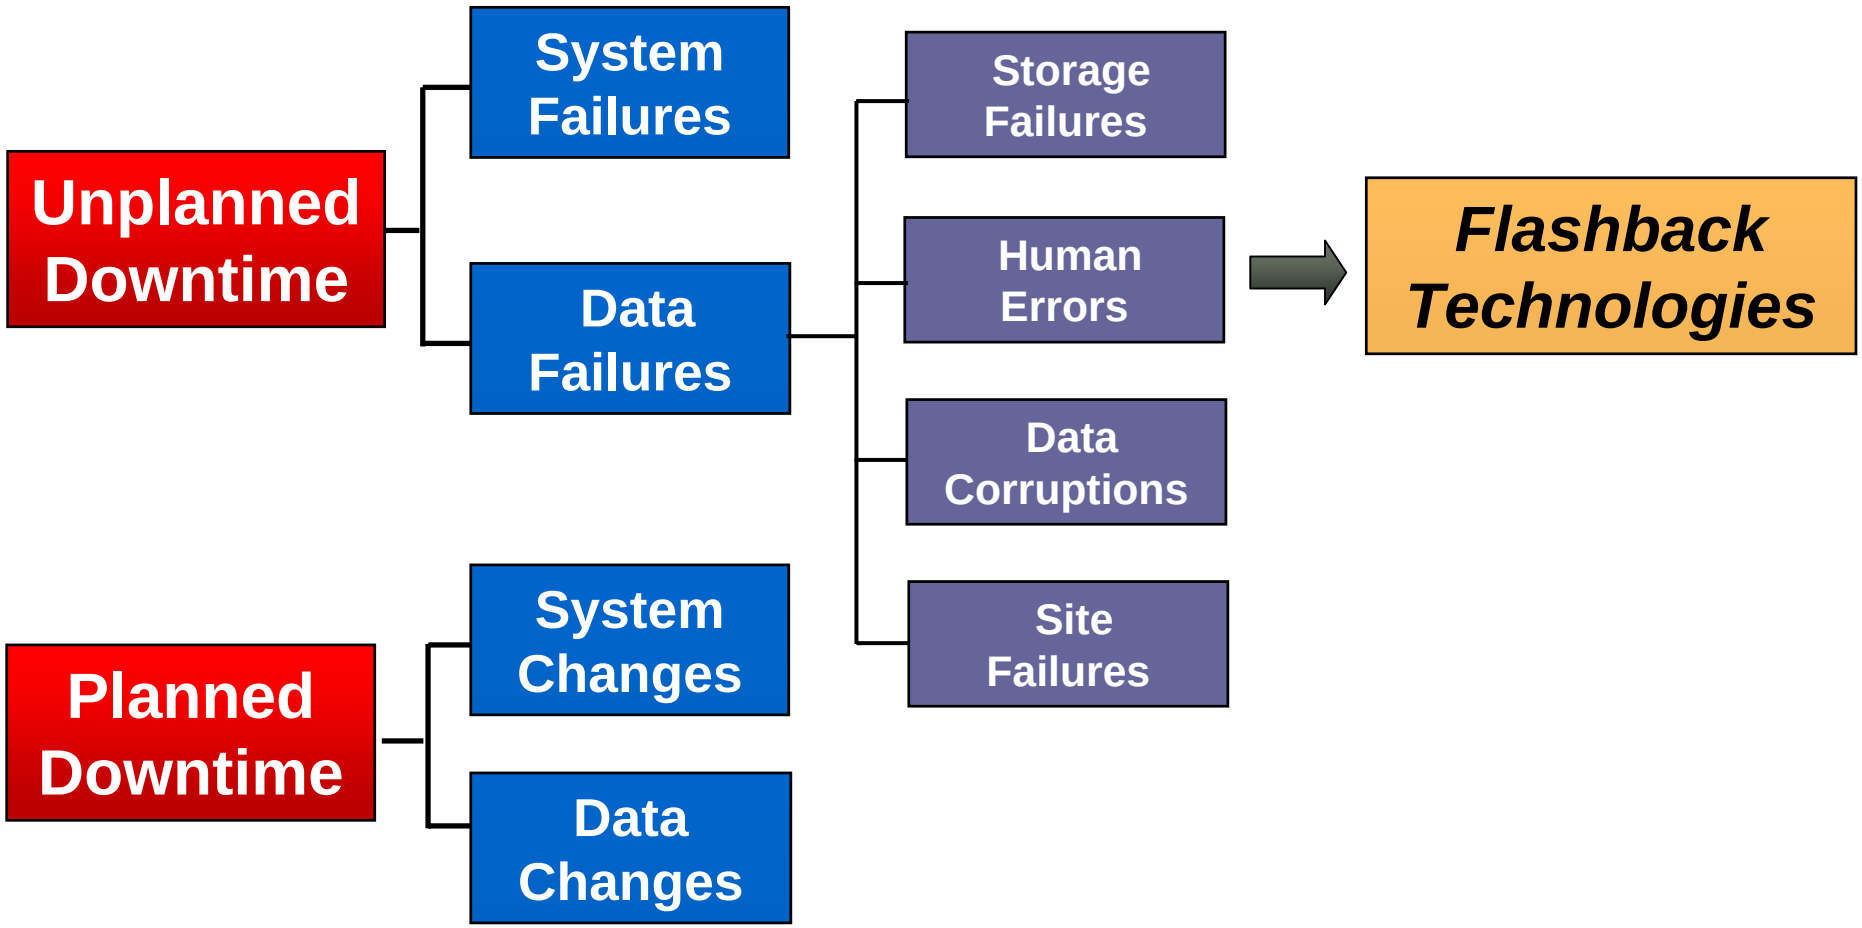

Revolution in Recovery

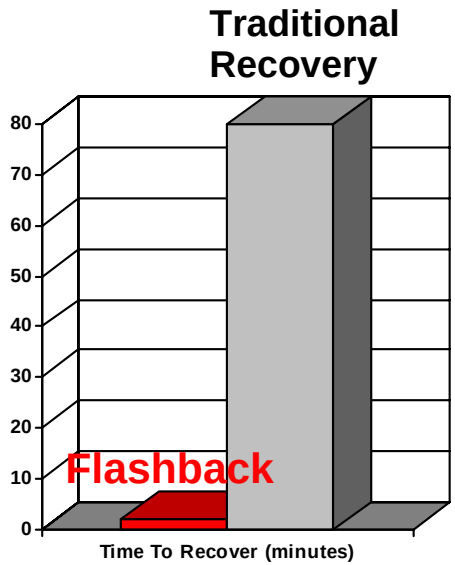

- Flashback Revolutionizes Error Recovery
  - Operates on just changed data
  - Time to correct error equals time to make error
    - Minutes instead of hours

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

- Flashback is Easy
  - Single command instead of complex procedure
- Very low performance overhead – less than 2%
- Great for testing also!

Error Investigation with Flashback

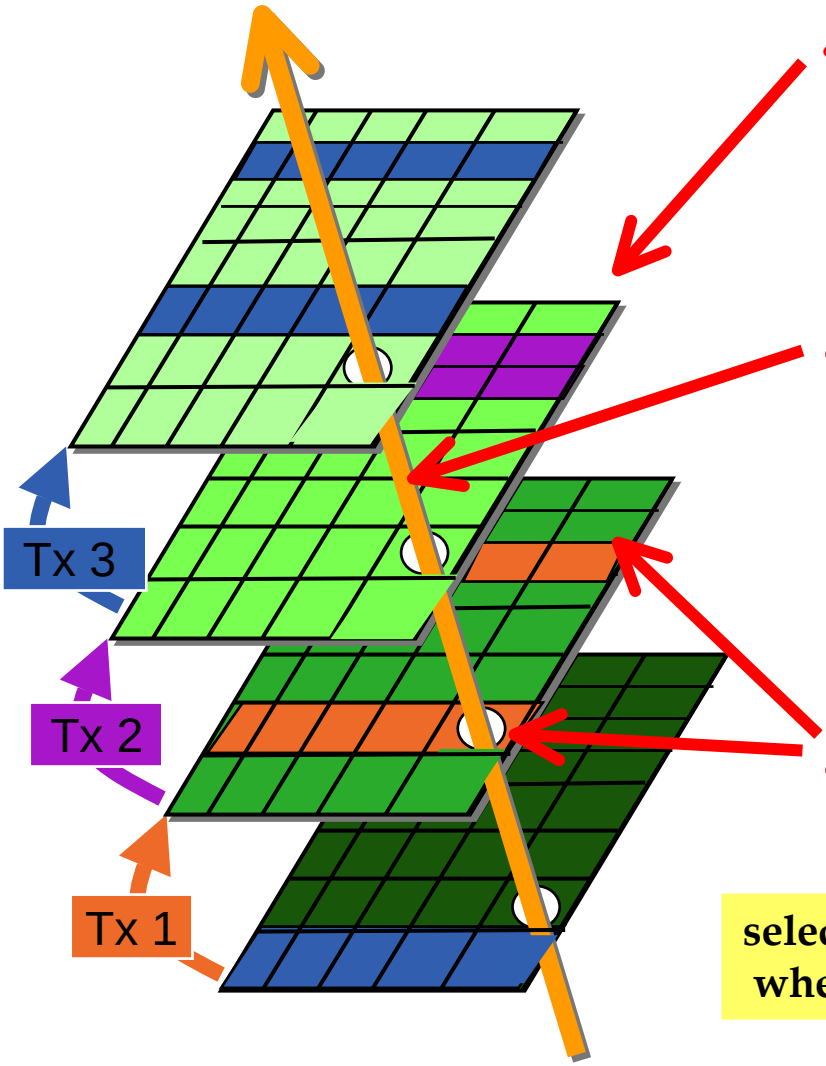

- **Flashback Query**
  - Query all data at point in time

```
select * from Salary AS OF '12:00 P.M.' where …
```

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

```
select * from Salary VERSIONS BETWEEN
 '12:00 PM' and '2:00 PM' where …
```

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

```
select * from FLASHBACK_TRANSACTION_QUERY
 where xid = '0002000300000002D';
```

Error Correction with Flashback

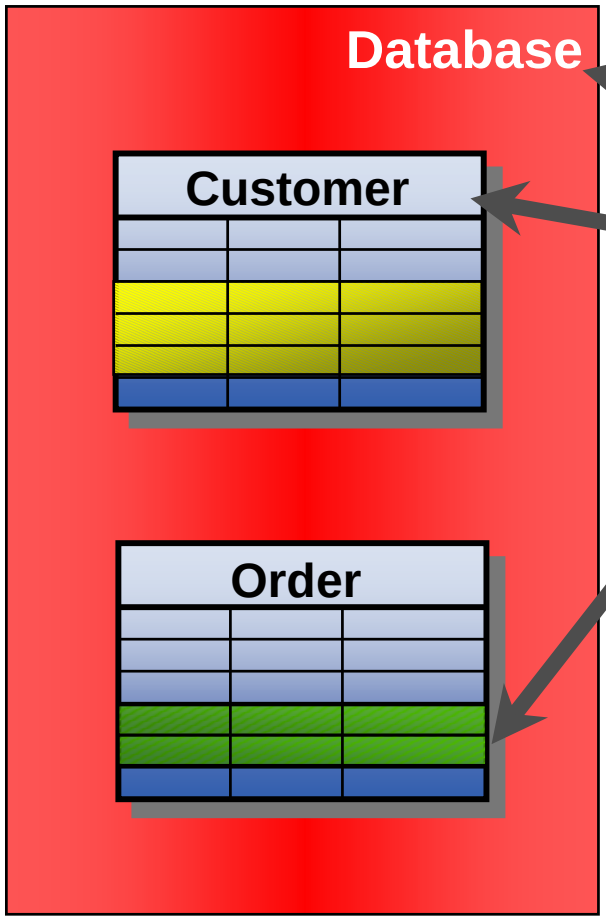

- Correct errors at any level
- Flashback Database – restore database to time
- Flashback Table – restore contents of tables to time

ORACLE 11g DATABASE

- Flashback Transaction – back out transaction and all subsequent conflicting transactions

Also Great for Testing

Flashback Data Archive

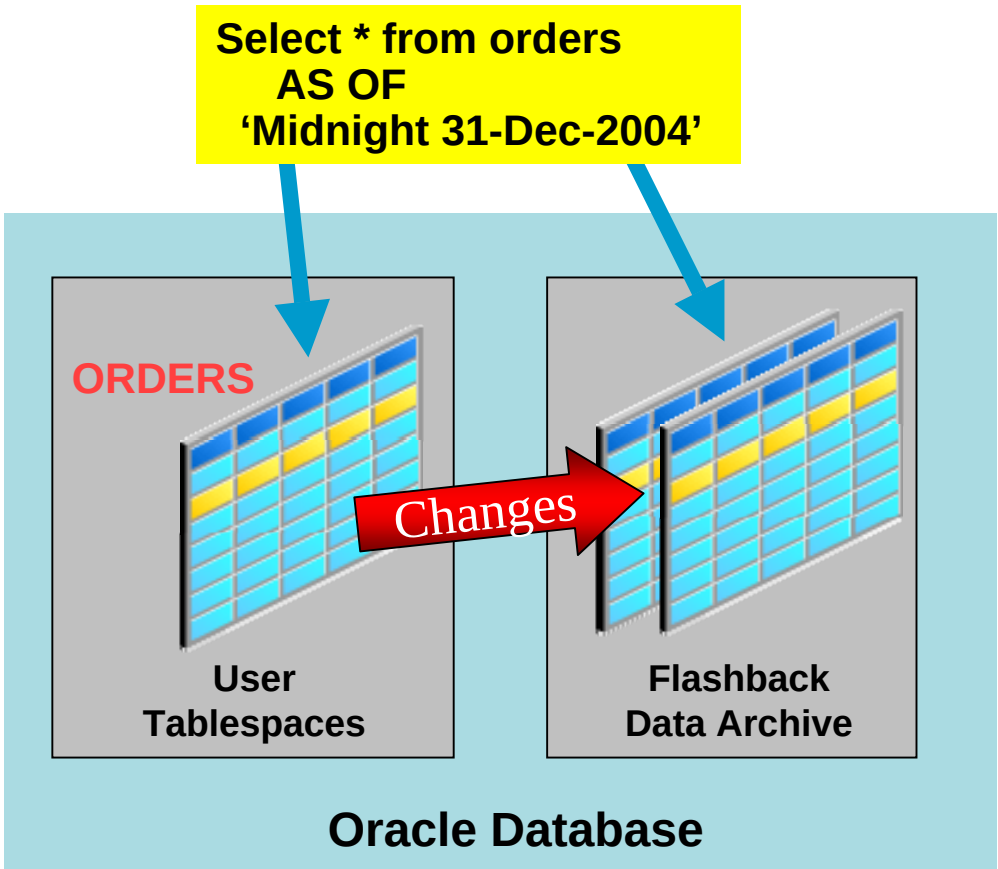

- Long term retention - years
- Automatically stores all changes to selected tables in Flashback Data Archive
  - Archive cannot be modified
  - Old data purged per retention policy
- View table contents as of any time using Flashback Query
- Uses
  - Change tracking/long term history
  - ILM
  - Auditing
  - Compliance

Best-of-Breed Data Corruption Protection
At Lowest Cost

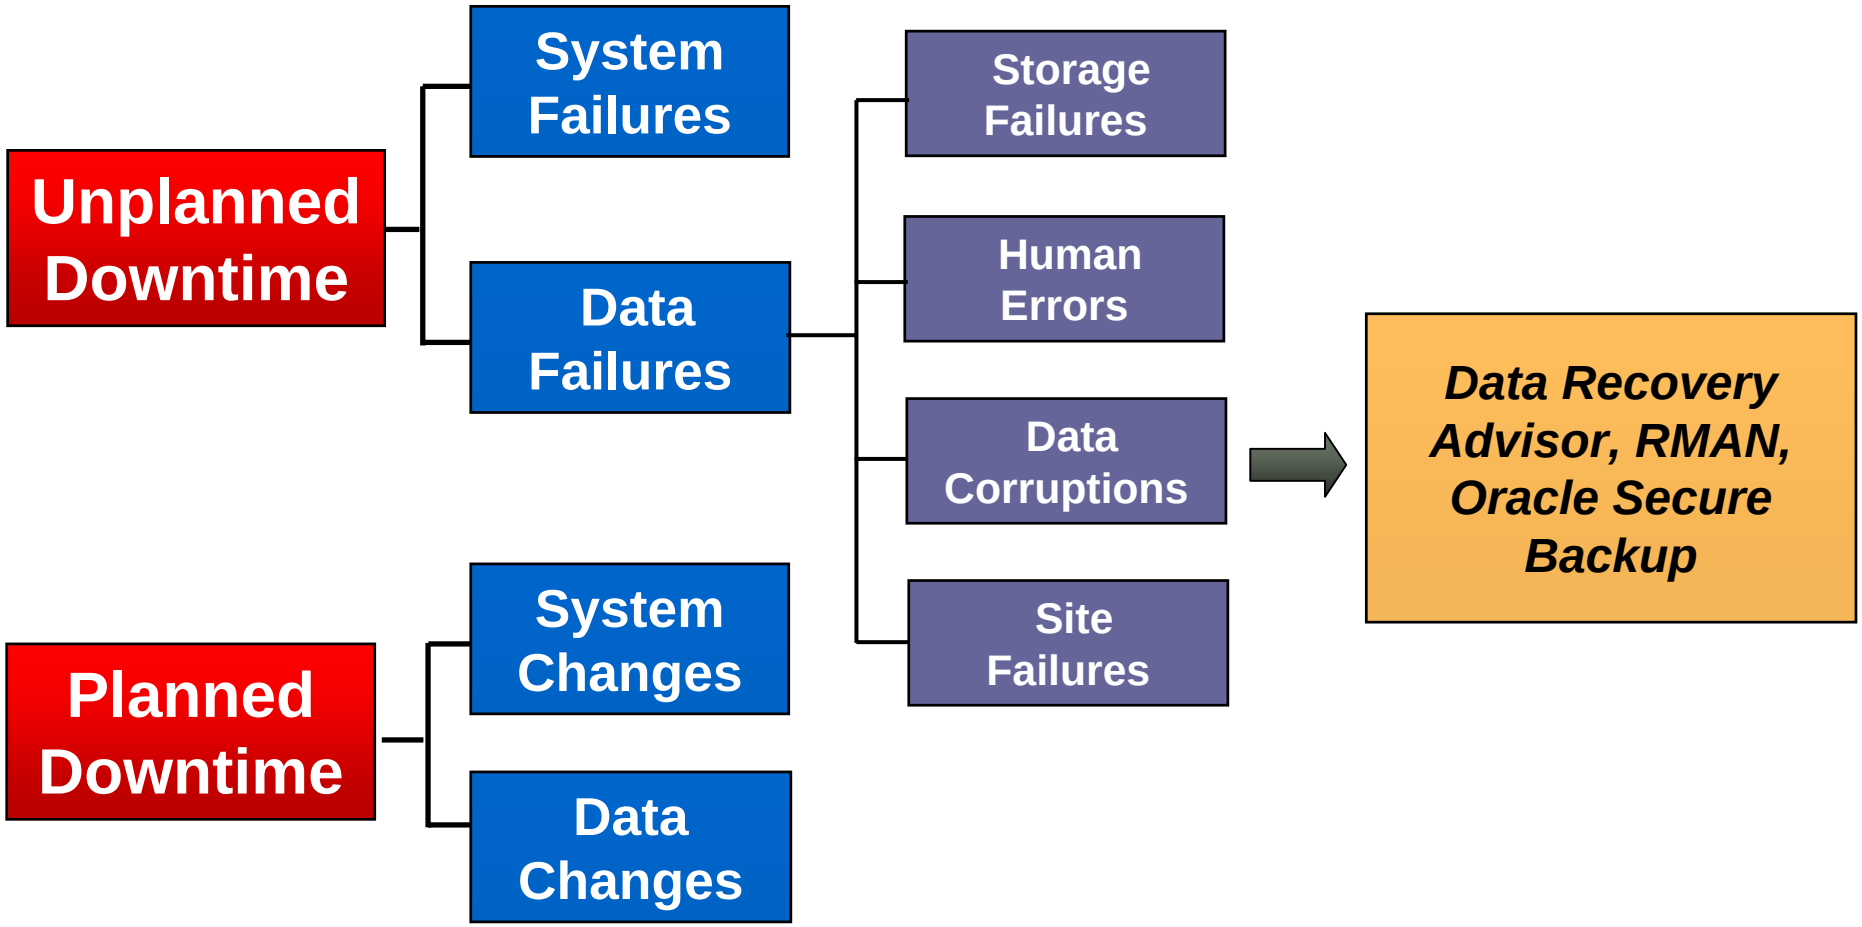

Automated Disk Backups

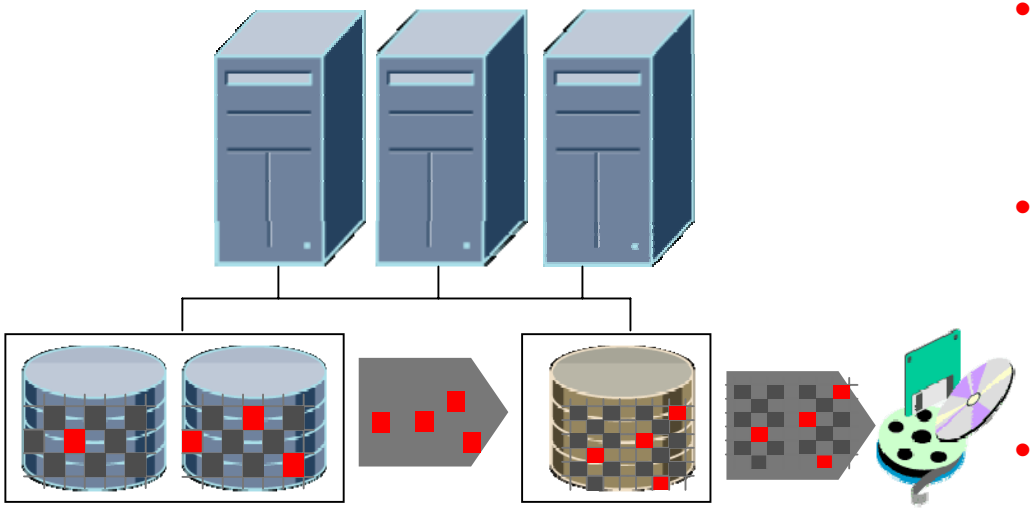

Integrated storage tiering within the database!

- Fully automatic disk-based backup and recovery
  - Set and Forget
- 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
- Low cost ATA disks can be used for recovery area

Data Recovery Advisor

The Motivation

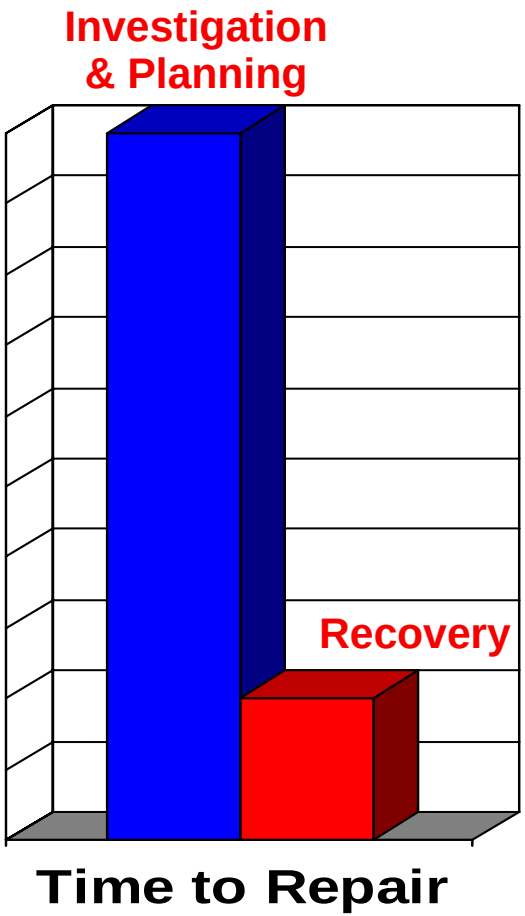

- Oracle provides robust tools for data repair:
  - ✓ RMAN – physical media loss or corruptions
  - ✓ Flashback – logical errors
  - ✓ Data Guard – physical or logical problems
- However, problem diagnosis and choosing the right solution can be error prone and time consuming
  - Errors more likely during emergencies

Data Recovery Advisor

ORACLE DATABASE 11g

- Oracle Database tool that automatically diagnoses data failures, presents repair options, and executes repairs at the user's request
- Determines failures based on symptoms
  - E.g. an “open failed” because datafiles f045.dbf and f003.dbf are missing
  - Failure Information recorded in diagnostic repository (ADR)
  - Flags problems before user discovers them, via automated health monitoring
- Intelligently determines recovery strategies
  - Consolidates failures for efficient recovery
  - Presents only feasible recovery options
  - Indicates any data loss for each option
- Can automatically perform selected recovery steps

Reduces downtime by eliminating confusion

RMAN Enhancements

ORACLE 11g DATABASE

- **Better performance**
  - Intra-file parallel backup and restore of single data files >= 1 GB (*multi-section backup*)
  - Faster backup compression (ZLIB, ~40% faster)
- **Better security**
  - Virtual Private Catalog - allows the catalog administrator to grant visibility of a subset of registered databases in the catalog to specific RMAN users
- **Lower space consumption**
  - Duplicate database or create standby database over the network, avoiding intermediate staging areas
- **Integration with Windows Volume Shadow Copy Services (VSS) API**
  - Allows database to participate in snapshots coordinated by VSS-compliant backup management tools and storage products
  - Database is automatically recovered upon snapshot restore via RMAN

Oracle Secure Backup

Integrated Tape Backup Management

ORACLE SECURE BACKUP

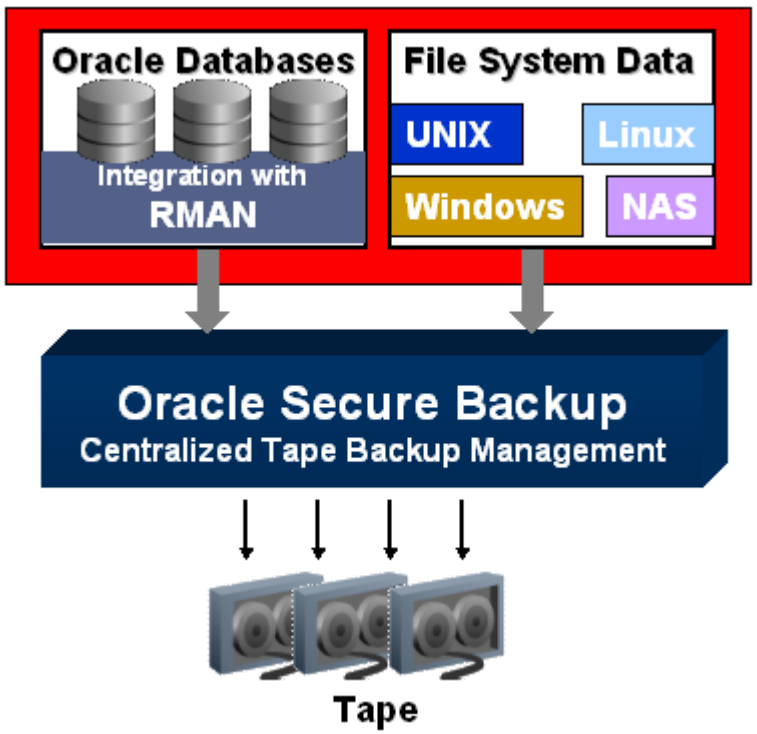

- Protects entire environment
  - Oracle9*i* forward
  - Application files
- Free Express edition bundled with the Oracle Database
- Low cost licensed edition
- Independent release schedule
  - Available: 10.1
  - Upcoming: 10.2
    - Beta planned for mid-summer

http://www.oracle.com/technology/products/secure-backup/index.html

Dare to Compare - Lowest Cost

ORACLE SECURE BACKUP

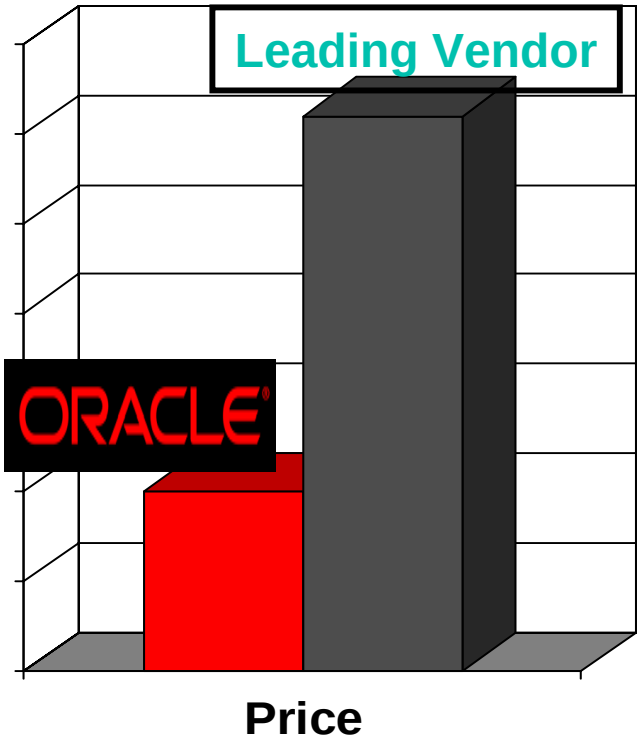

Feature	Leading Vendor	Oracle Secure Backup
Tape Drive	$3,000	$3,000
SAN Backup per drive	$2,000	Free
UNIX Client Host	$600	Free
UNIX Media Server	$13,000	Free
Linux Media Server	$5,500	Free
Oracle Agent	$7,500	Free
NAS Filer-NDMP	$6,500	Free
Advanced Features: Vaulting, Encryption etc.	$$$$$$$	Free

- Oracle Secure Backup price is just $3000 per tape drive
  - Backup to virtual tape device (disk) is **Free**
  - **Free** Express Edition protects one database server to one attached tape drive

Oracle Secure Backup 10.2 Enhancements

ORACLE SECURE BACKUP

- **Increased Security for data and backup domain**
  - Backup encryption for file systems and Oracle9*i* forward
- **Advanced media management**
  - Vaulting
  - Tape duplication
  - ACSLS support
- **Improved Manageability**
  - Automated backup of OSB catalog
  - Policy-based migration from VTL to tape
- **Performance improvements**
  - Strengthened RMAN and OSB Integration

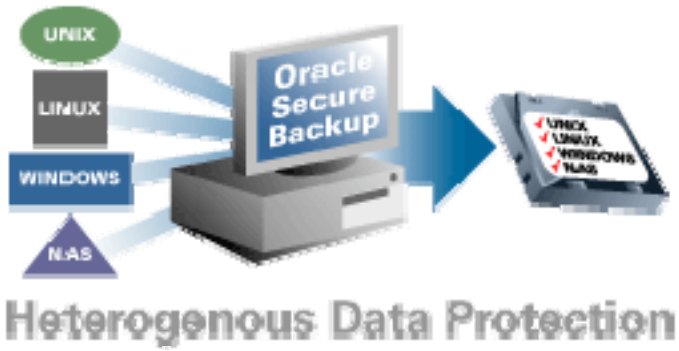

Advanced Functionality at NO Extra Cost!

Best-of-Breed Disaster Protection

At Lowest Cost

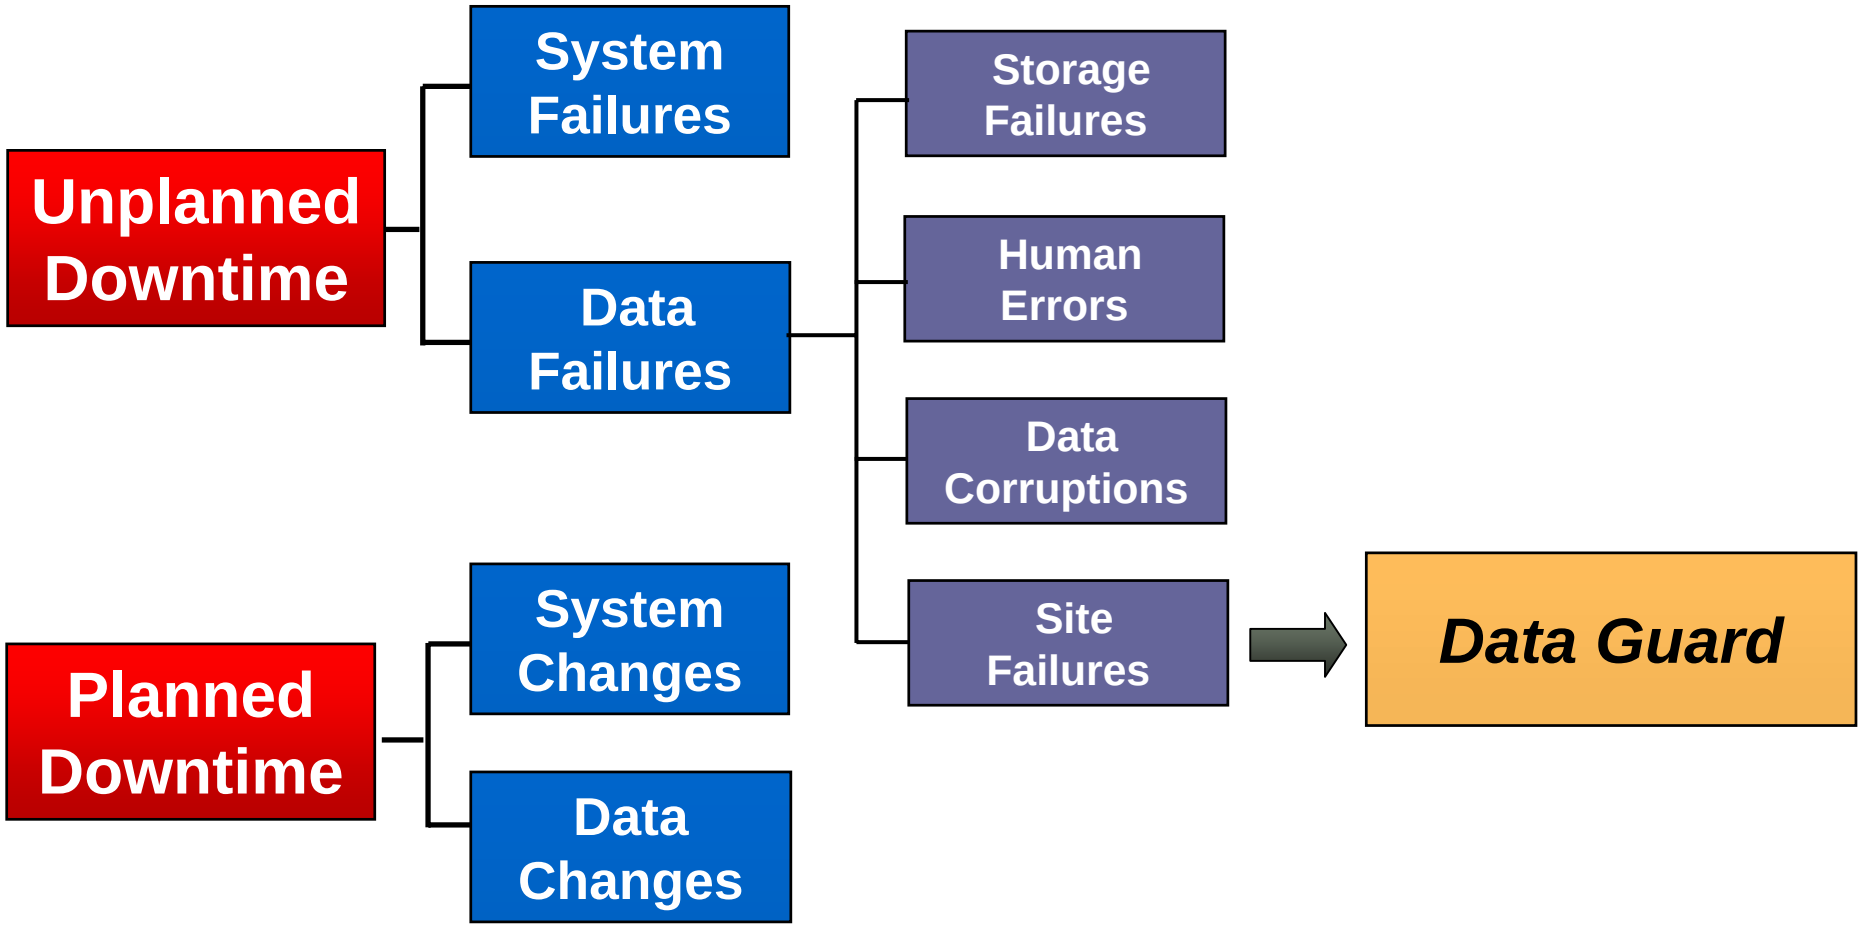

Disaster Recovery (DR) Realities

- **Customers don't benefit from DR investment**
  1. Expensive – choose no DR, or under-configure DR
  2. Idle systems – no productive use
  3. Rarely used – so no confidence failover will work
  4. Loses data – leads to downstream problems
  5. Slow – prefer to fix problems instead of using DR
  6. Limited protection – site failures only

- **Requirements for useful / ubiquitous DR**
  1. Cost-effective – hardware and software
  2. Efficient systems utilization
  3. Easy DR testing
  4. Fast automatic failover over long distances, with zero data loss
  5. Covers all common failures – not just site failures
  6. Application transparency
  7. Bonus – reduce planned downtime

- ***Need all of the above!***

Data Guard: Best Failure Protection at Lowest Cost

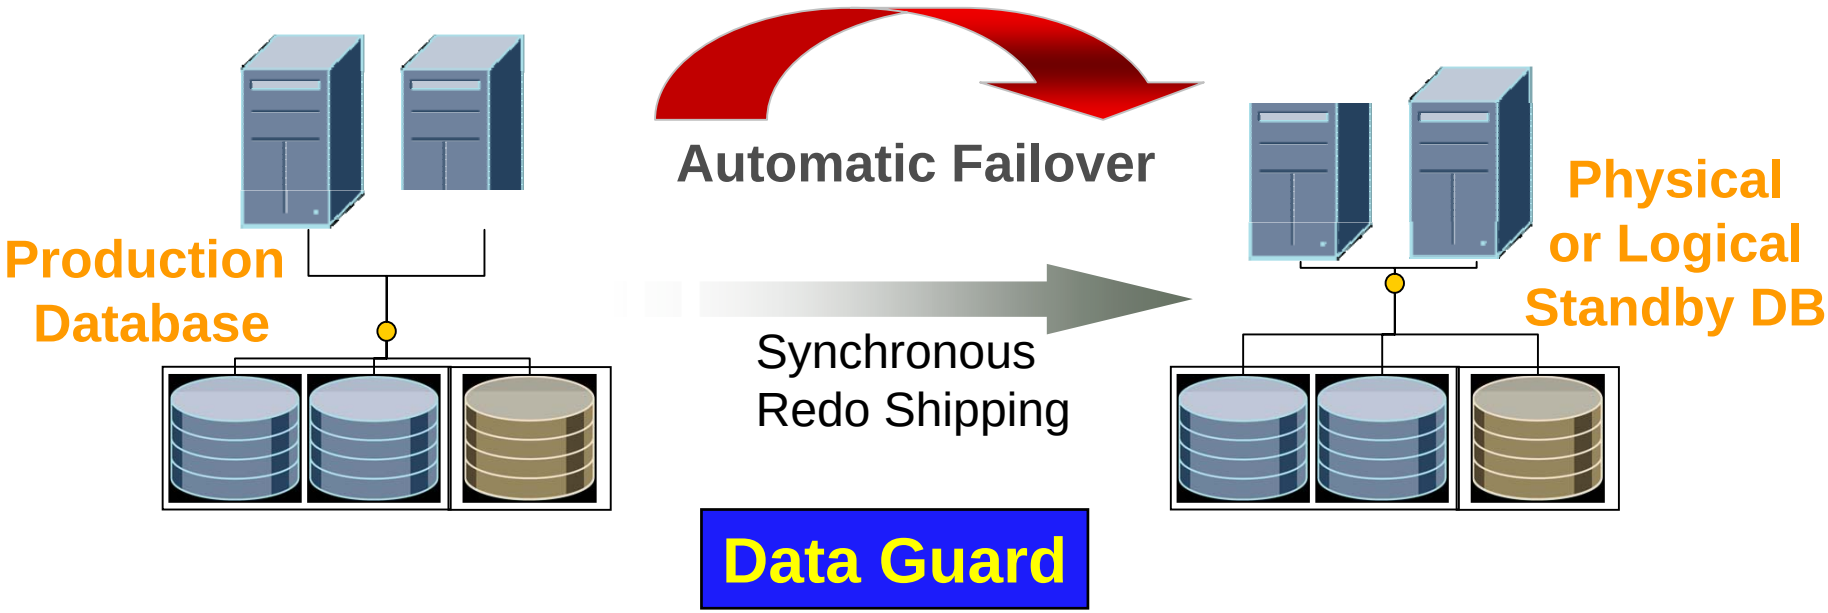

- Synchronous or asynchronous redo shipping
- Corruptions don't propagate
- Low cost servers and storage
- Data Guard is free with EE
- Thousands of production customers

Zero Data Loss over Long Distance

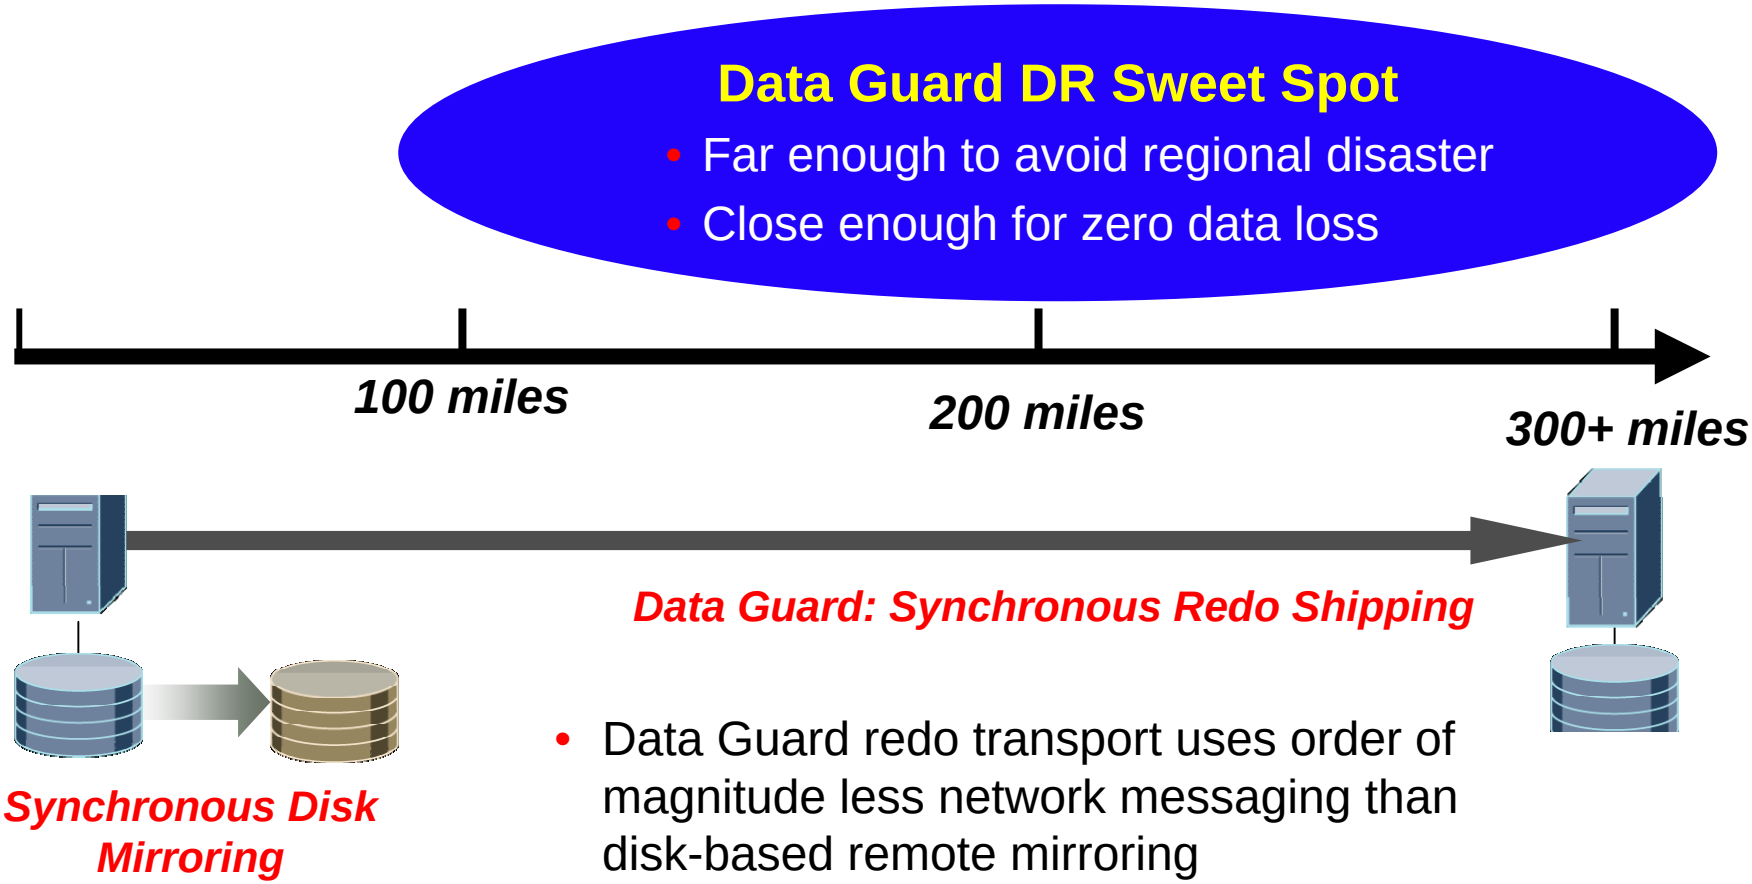

- Data Guard redo transport uses order of magnitude less network messaging than disk-based remote mirroring
  - Enables zero data loss at hundreds of miles

Data Guard Enhancements

- Better standby resource utilization
- Enhanced HA / DR functionality
- Improved performance

Data Guard becomes an integral part of IT operations

Physical Standby with Real-Time Query

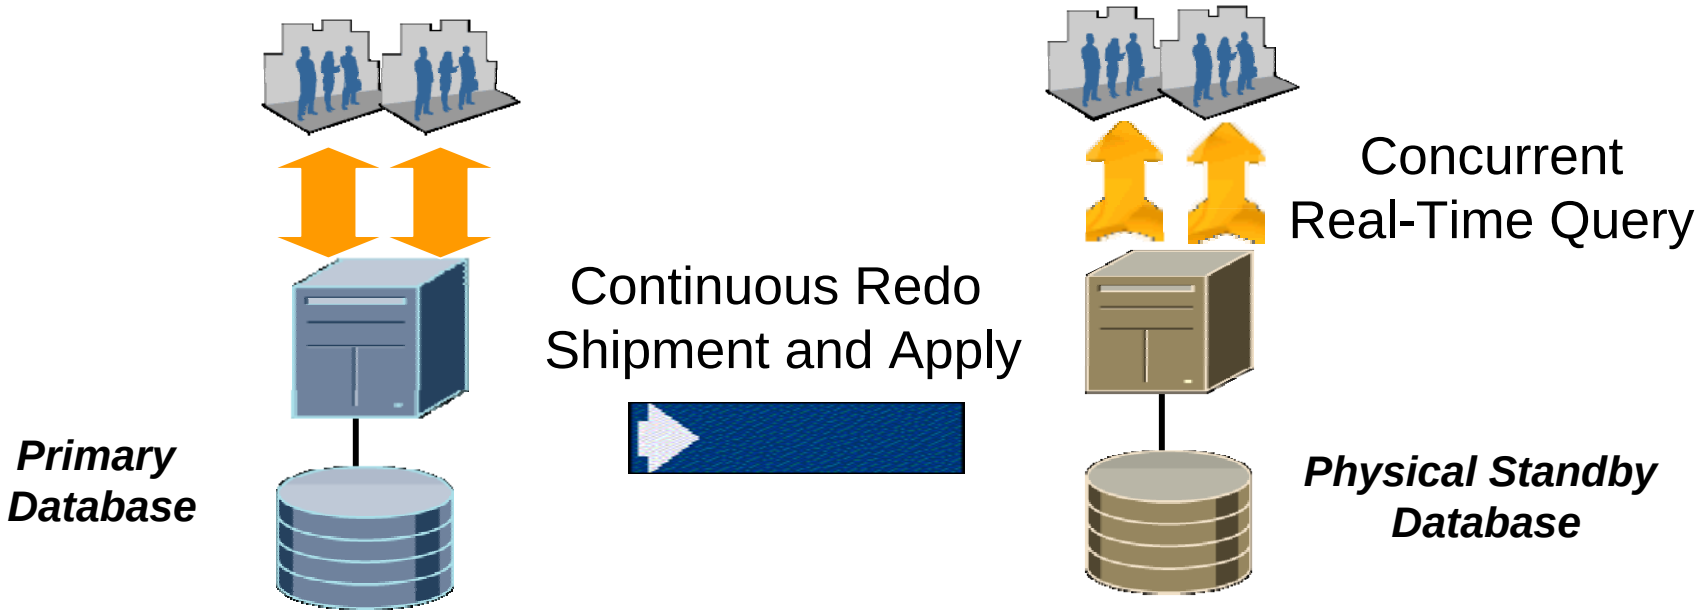

- Read-only queries on physical standby concurrent with redo apply
  - Supports RAC on primary / standby
  - Queries see transactionally consistent results
  - Handles all data types, but not as flexible as logical standby

Snapshot Standby – Leverage Standby Database for Testing

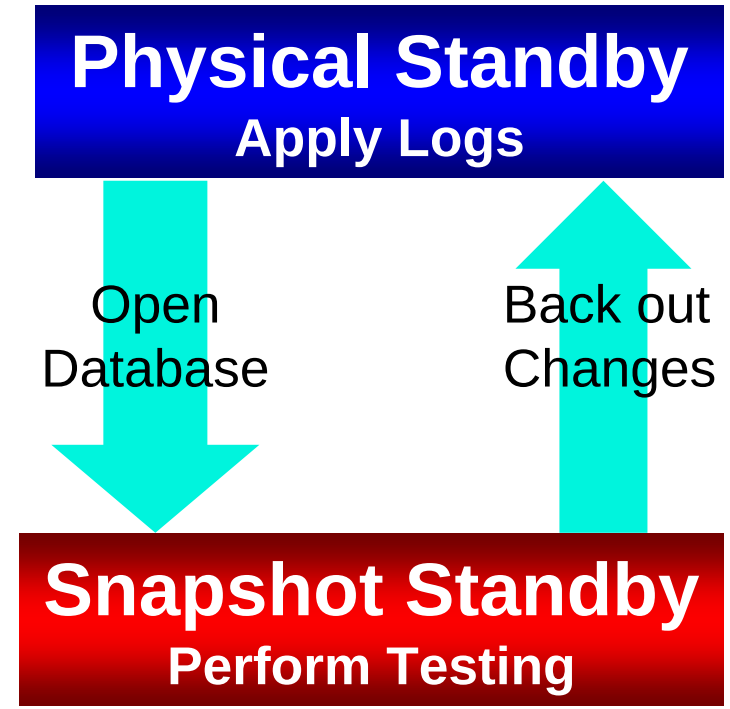

- Convert Physical Standby to Snapshot Standby and open for writes by testing applications
  - `ALTER DATABASE CONVERT TO SNAPSHOT STANDBY;`
- Discard testing writes and catch-up to primary by applying logs
  - `ALTER DATABASE CONVERT TO PHYSICAL STANDBY;`
- Preserves zero data loss
  - But no real time query or fast failover
- No idle resources
- Similar to storage snapshots, but:
  - Provides DR at the same time
  - Uses single copy of storage

Enhanced Fast-Start Failover

ORACLE 11g DATABASE

- Supports Maximum Performance (ASYNC) Mode
  - Automatic failover for long distance standby
  - Data loss exposure limited using Broker property **`FastStartFailoverLagLimit`** (*default = 30 secs*)
- Immediate fast-start failover for user-configurable health conditions
  - **`ENABLE FAST_START FAILOVER [CONDITION <value>];`**
    - Condition examples:
      - Datafile Offline
      - Corrupted Controlfile
      - Corrupted Dictionary
      - Inaccessible Logfile
      - Stuck Archiver
      - Any explicit ORA-xyz error
- Apps can request fast-start failover using **`DBMS_DG.INITIATE_FS_FAILOVER`**

Performance Improvements

ORACLE 11g DATABASE

- Faster Failover
  - Failover in seconds with Fast-Start Failover
- Faster Redo Transport
  - Optimized async transport for Maximum Performance Mode
  - Redo Transport Compression for gap fetching: new **`compression`** attribute for **`log_archive_dest_n`**
- Faster Redo Apply
  - Parallel media recovery optimization
- Faster SQL Apply
  - Internal optimizations
- Fast incremental backup on physical standby database
  - Support for block change tracking

Streams: Another Popular HA Solution

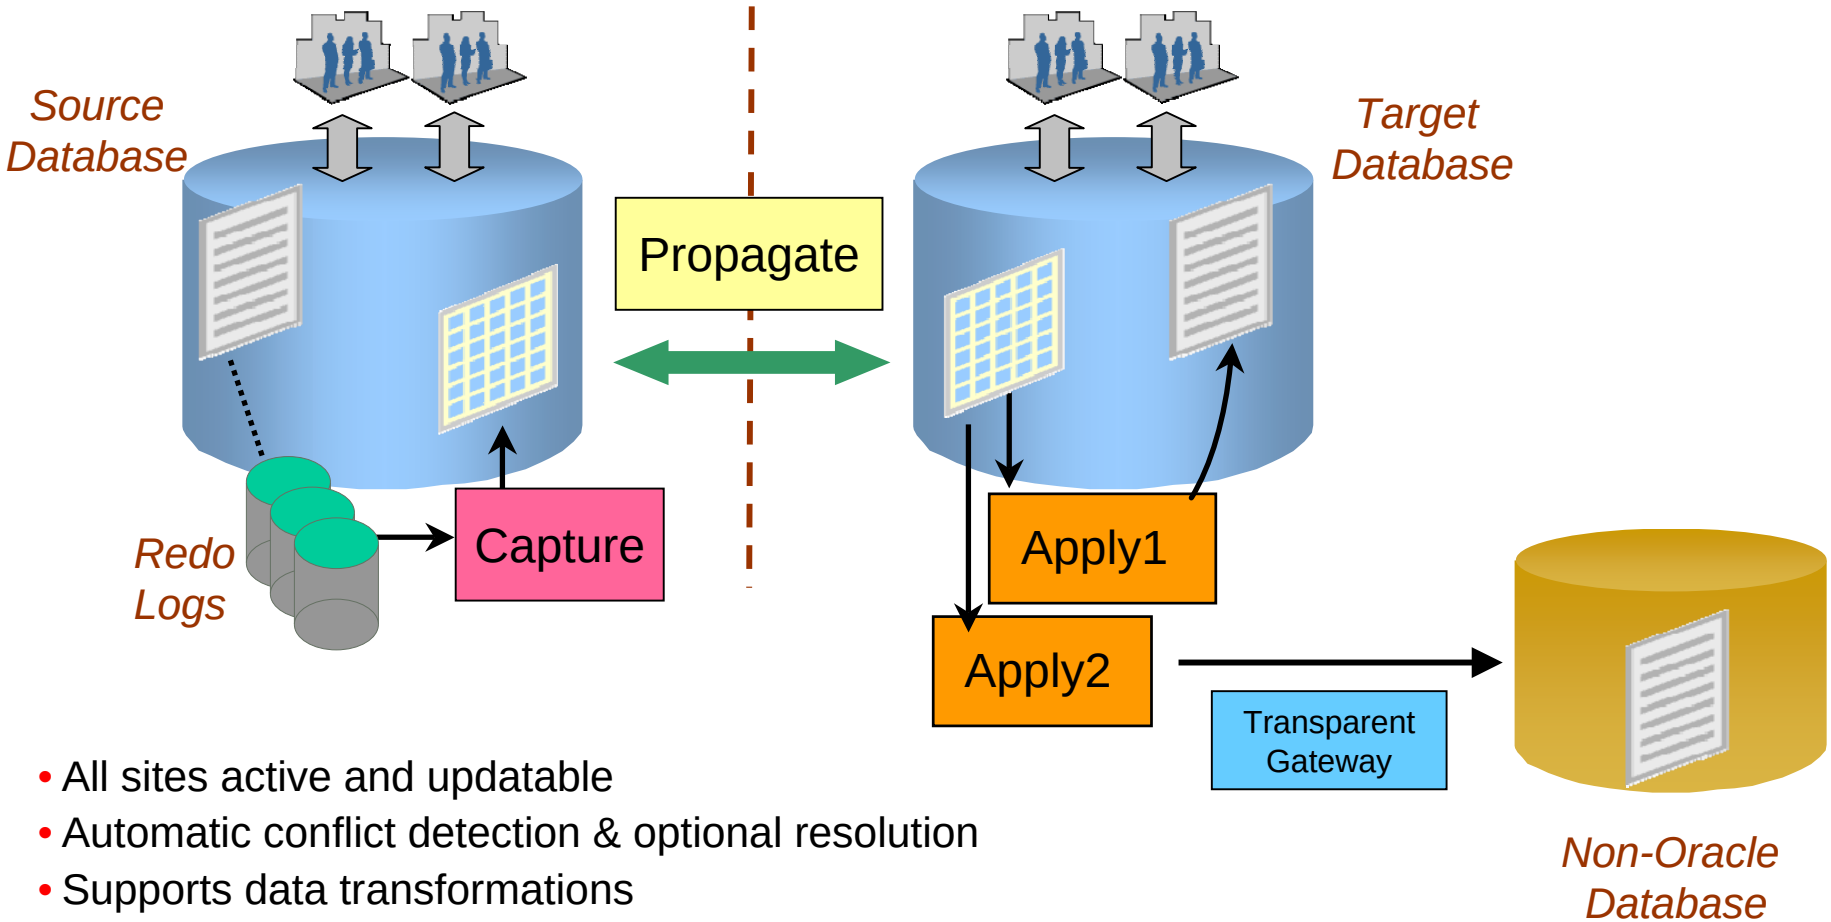

- All sites active and updatable
- Automatic conflict detection & optional resolution
- Supports data transformations
- Flexible configurations – n-way, hub & spoke, …
- Database platform / release / schema structure can differ
- Provides HA for custom apps where update conflicts can be avoided or managed

Streams Enhancements

ORACLE DATABASE 11g

- Streams Synchronous Capture
  - Available in all Editions of Oracle Database 11*g*
  - Efficient internal mechanism to immediately capture change
- Source and Target data compare & converge
- Streams Performance Advisor
- Split/Merge of Streams for Hub & Spoke replication
  - Maintains high performance for all replicas
  - Automated, fast “catch-up” for unavailable replica
- Cross-database LCR tracking
  - Trace Streams messages from start to finish in single view
- Performance optimizations

Best Online System Changes

At Lowest Cost

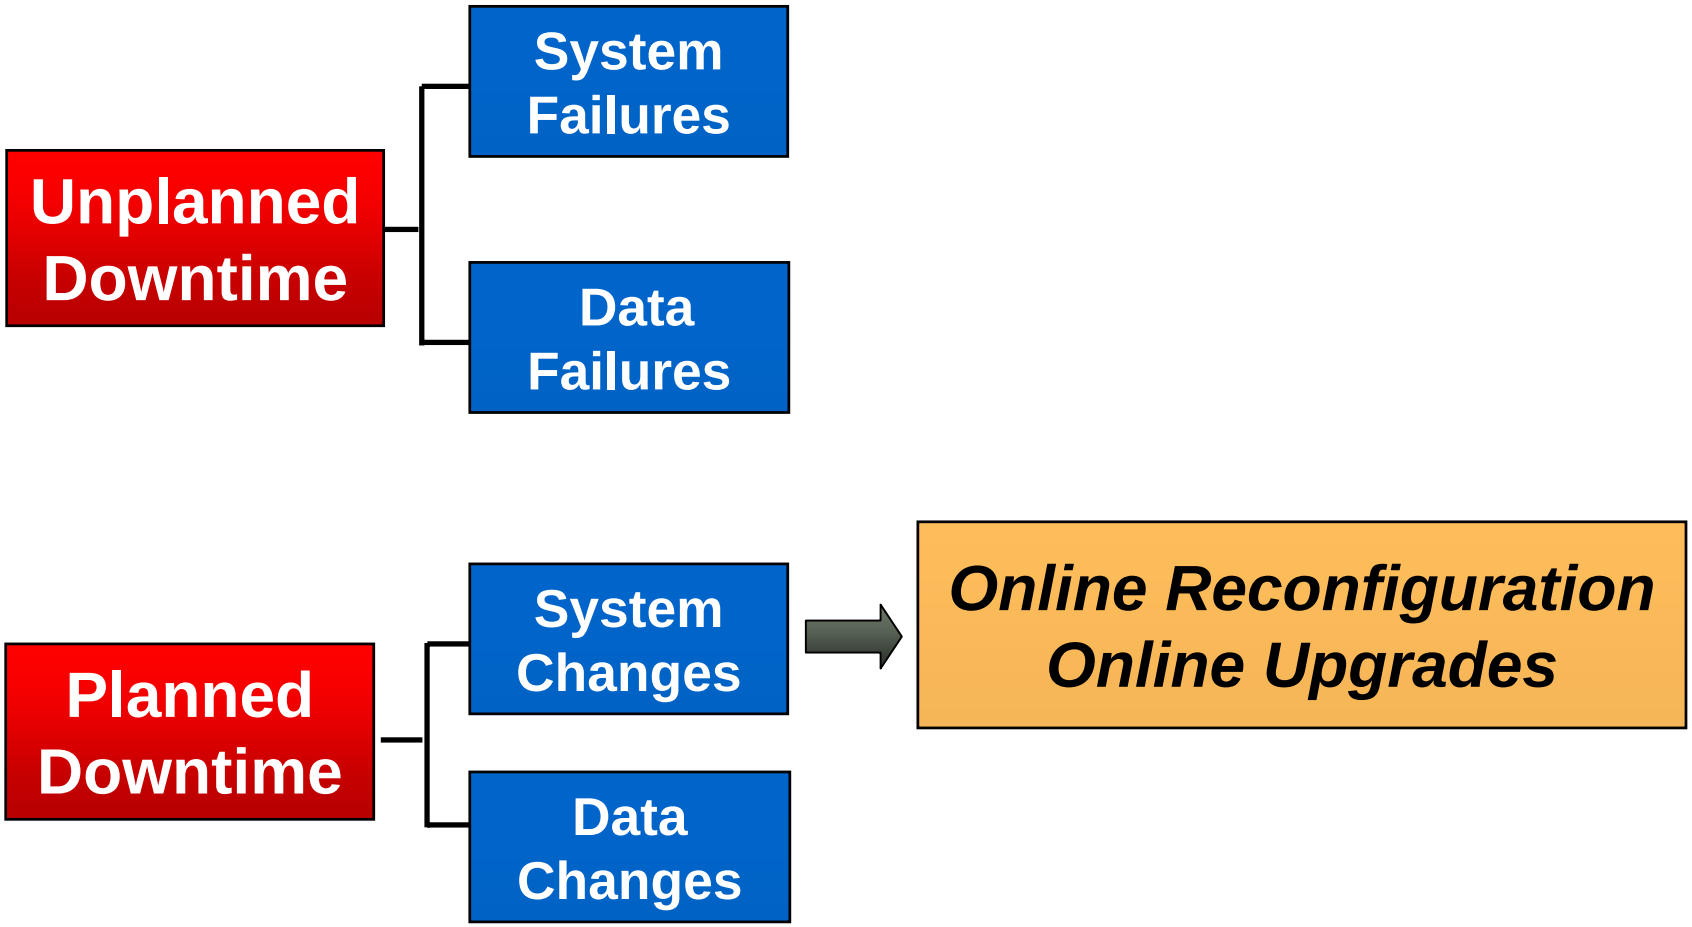

Online Reconfiguration – Scaling on Demand

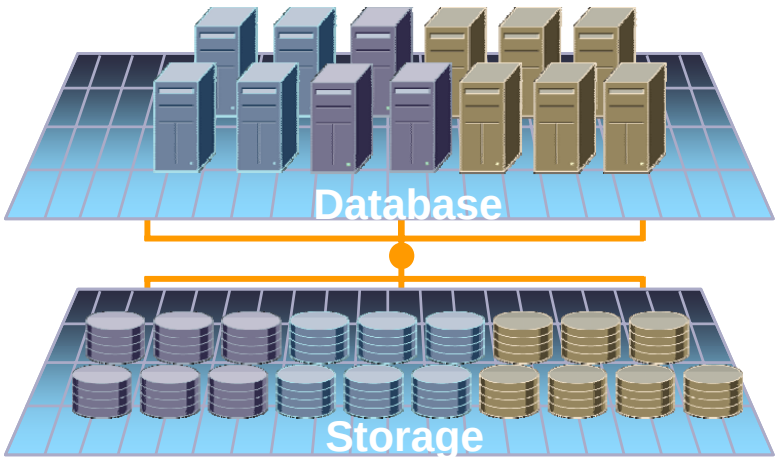

- CPU
  - Add/remove CPUs on SMP online
- Cluster Nodes
  - Add/remove RAC nodes online
  - No data movement needed
- Memory
  - Grow and shrink shared memory and buffer cache online
  - Auto tuning of memory online
- Disk
  - Add/remove ASM disks online
  - Automatically rebalance

Rolling Patch Update using RAC

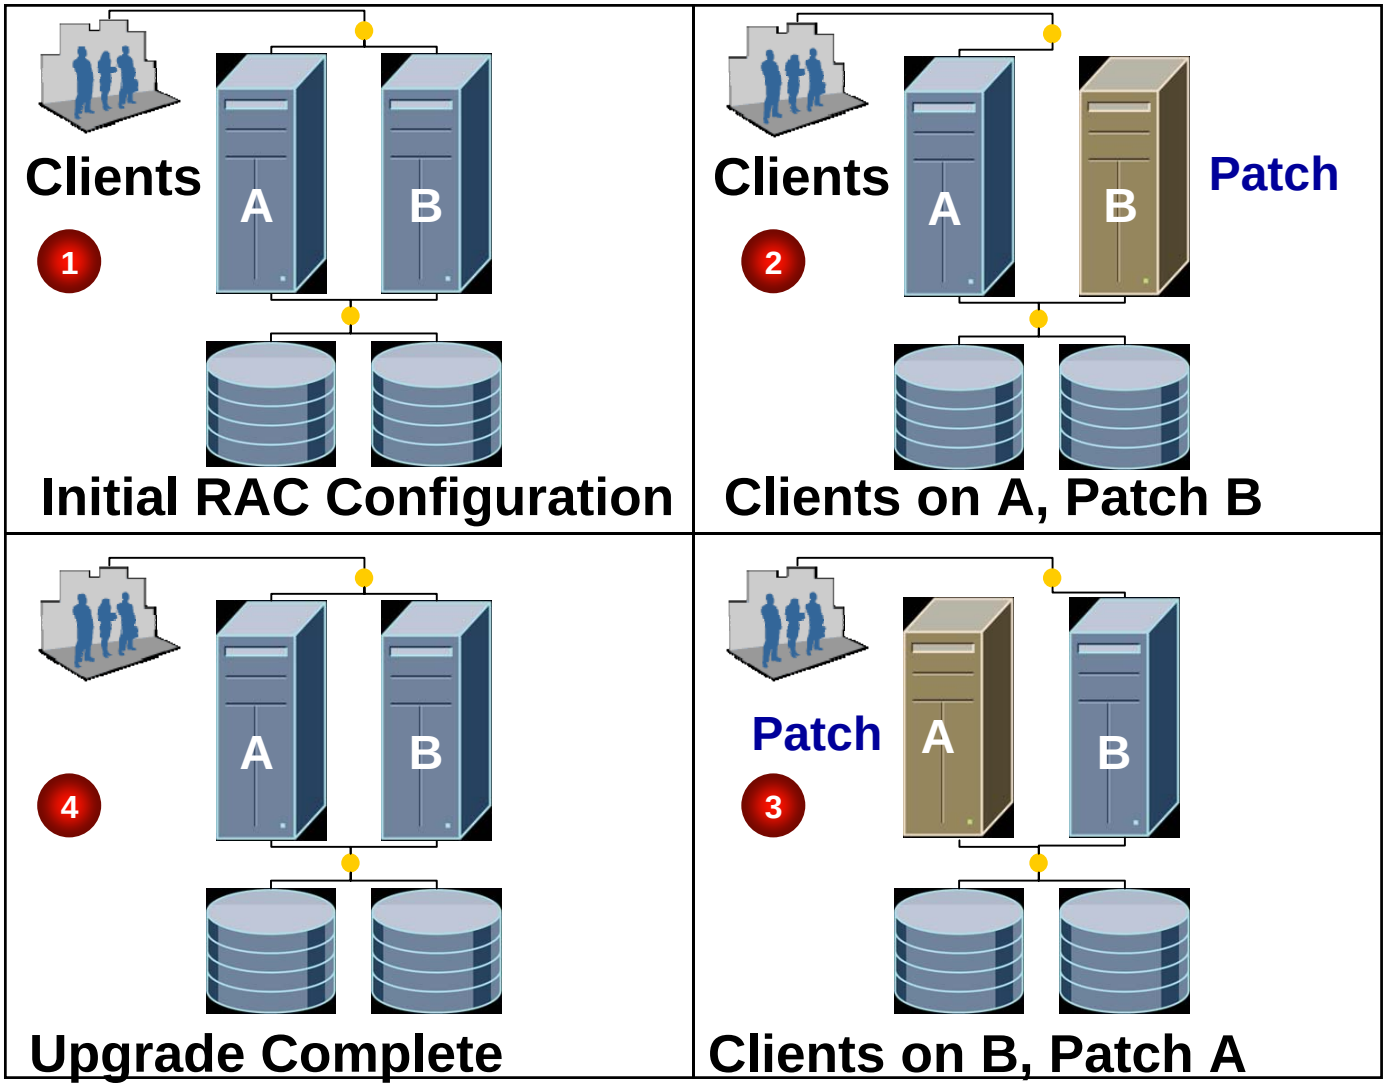

Oracle Patch Upgrades, including Critical Patch Updates (CPUs)

Operating System Upgrades

Hardware Upgrades

SQL Apply – Rolling Database Upgrades

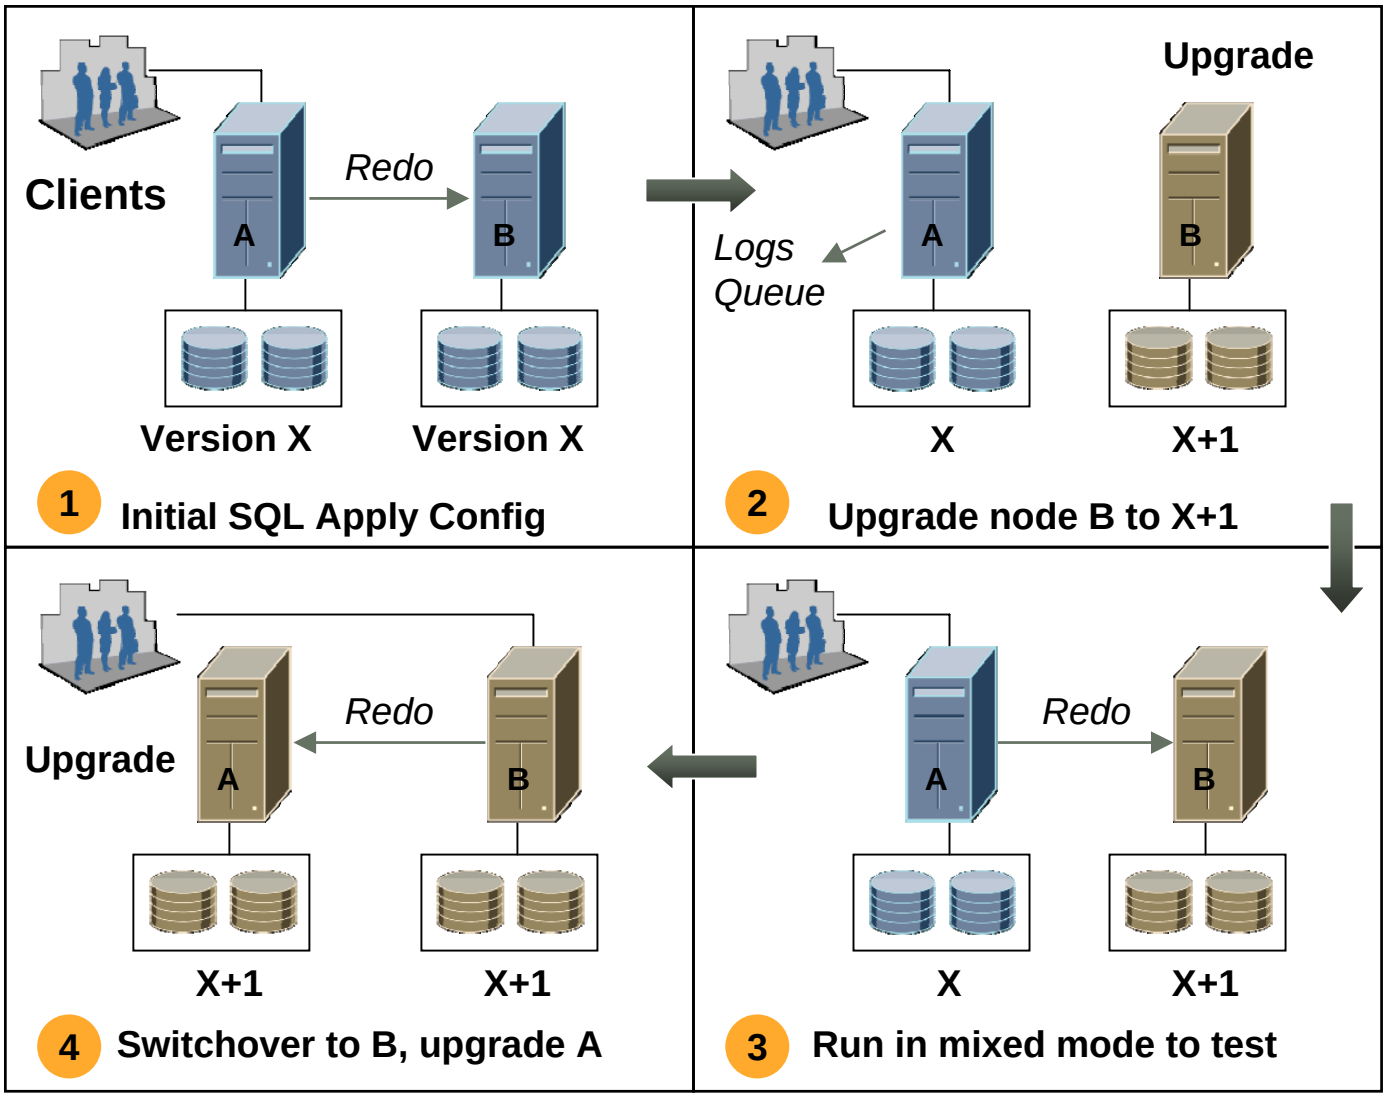

Patch Set Upgrades

Major Release Upgrades

Cluster Software & Hardware Upgrades

Online Patching of One-off Patches

- Ability to patch running Oracle executable
  - No downtime
  - No need to do rolling upgrades using RAC / Data Guard
  - Many one-off patches can be patched online
  - Great for diagnostic patches
    - E.g. debugging changes to better understand a problem before applying fix
- Supports enabling, disabling, de-installing patches with no downtime
- Integrated with Opatch
  - E.g. determine if a patch can be applied online:
    - `opatch query -is_online`
- Initially available on Linux (32 & 64-bit) and Solaris (64-bit)
- Long term goal is online patching of Critical Patch Updates (CPUs)

Rolling Database Upgrades Using Transient Logical Standby

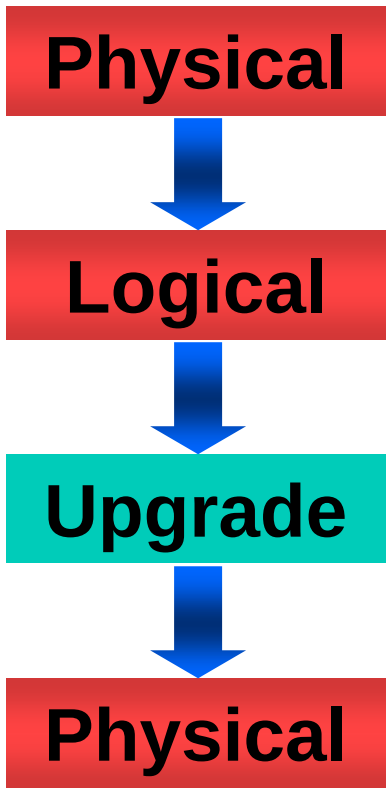

- Start rolling database upgrades with physical standbys
- Temporarily convert physical standby to logical to perform the upgrade
  - Data type restrictions limited to short upgrade window
- No need for separate logical standby for upgrade
- Also possible in 10.2 (more manual steps)

Leverage your physical standbys!

Best Online Data Changes

At Lowest Cost

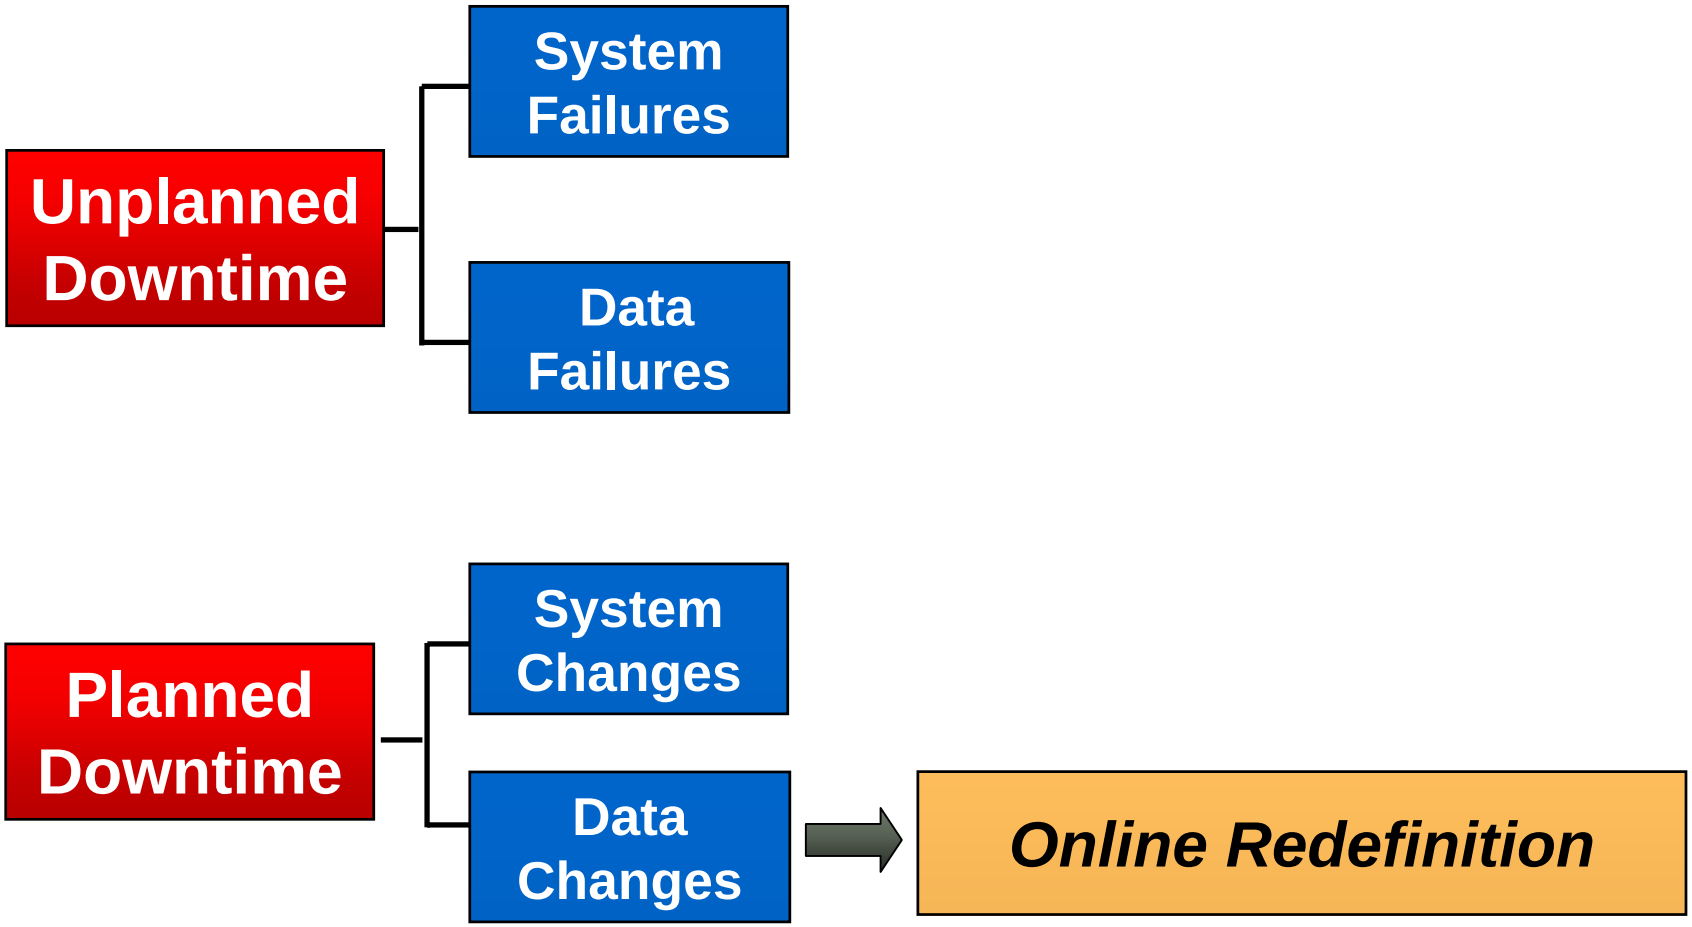

Online Redefinition

- All indexing operations can be done online
  - Create new index, move index, defragment index
- Tables can be Reorganized & Redefined online (**DBMS_REDEFINITION**)
  - Table contents are copied to a new table
    - Defragments and allows changing location, table type, partitioning
  - Contents can be transformed as they are copied
    - Can change columns, types, sizes - specified using SQL “Select”
- Updates and Queries can continue uninterrupted

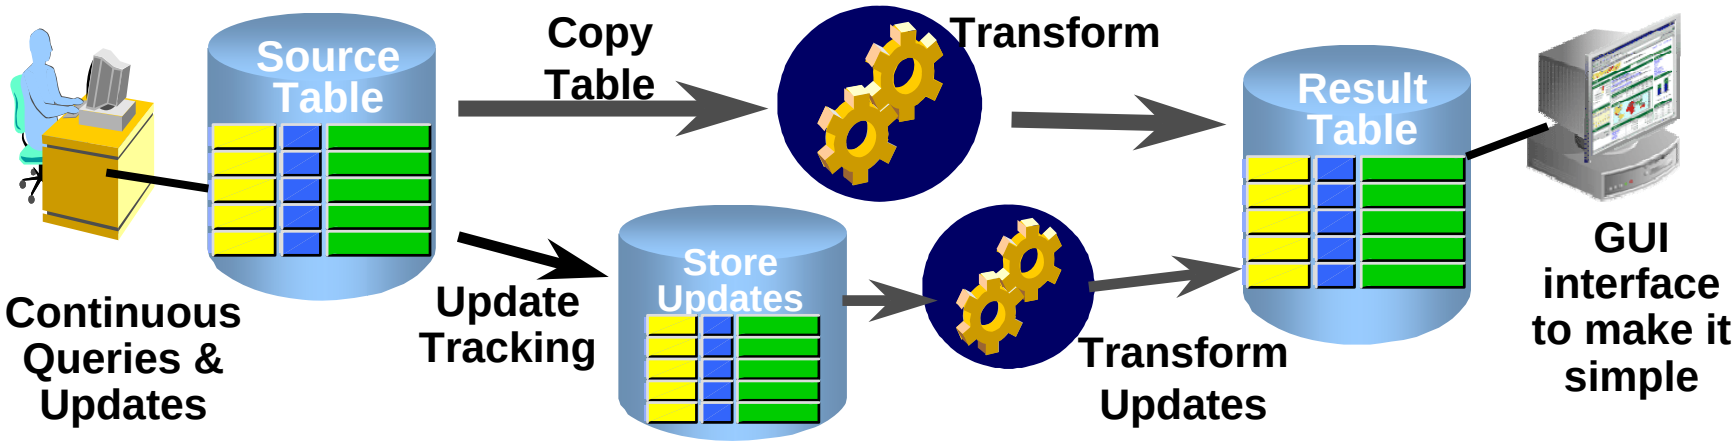

Online Operations & Redefinition Improvements

- Fast ‘add column’ with default value
- Invisible indexes speed application migration and testing
- No recompilation of dependent objects when Online Redefinition does not logically affect objects
- Support Online Redefinition for tables with Materialized Views
- Enhanced Online DDL execution
  - DDL operations now wait if underlying resource is busy (configured through **`DDL_LOCK_TIMEOUT`** parameter)
  - Some DDL operations (add/modify constraint, add column, Index create/rebuild) only require shared lock

Oracle Maximum Availability Architecture

Maximum Availability Architecture (MAA)

Integrated set of HA best practices

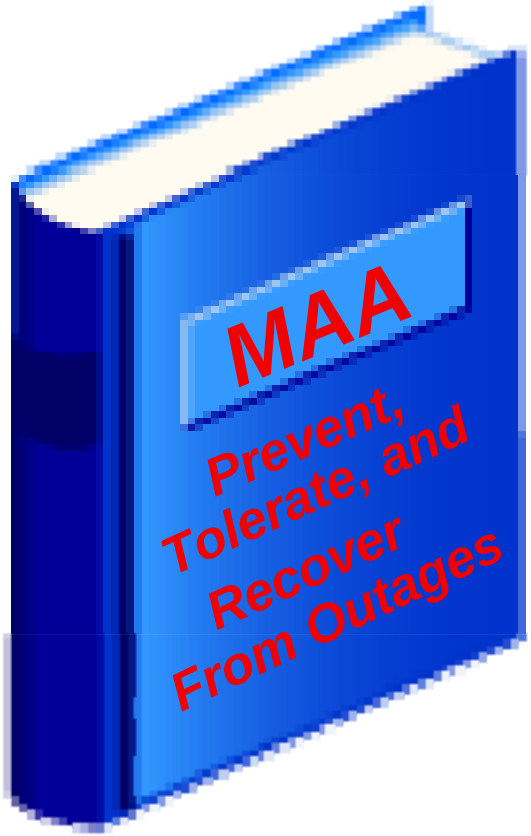

- Technology alone is not enough
- MAA is also a blueprint for achieving HA
  - Brings together all that has been discussed
  - Operational best practices
  - Prevent, tolerate, and recover
- Tested, validated, and documented
  - Database, Storage, Cluster, Network
  - 35 person year effort

otn.oracle.com/deploy/availability

Maximum Availability = Unbreakable Architecture + Best Practices

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

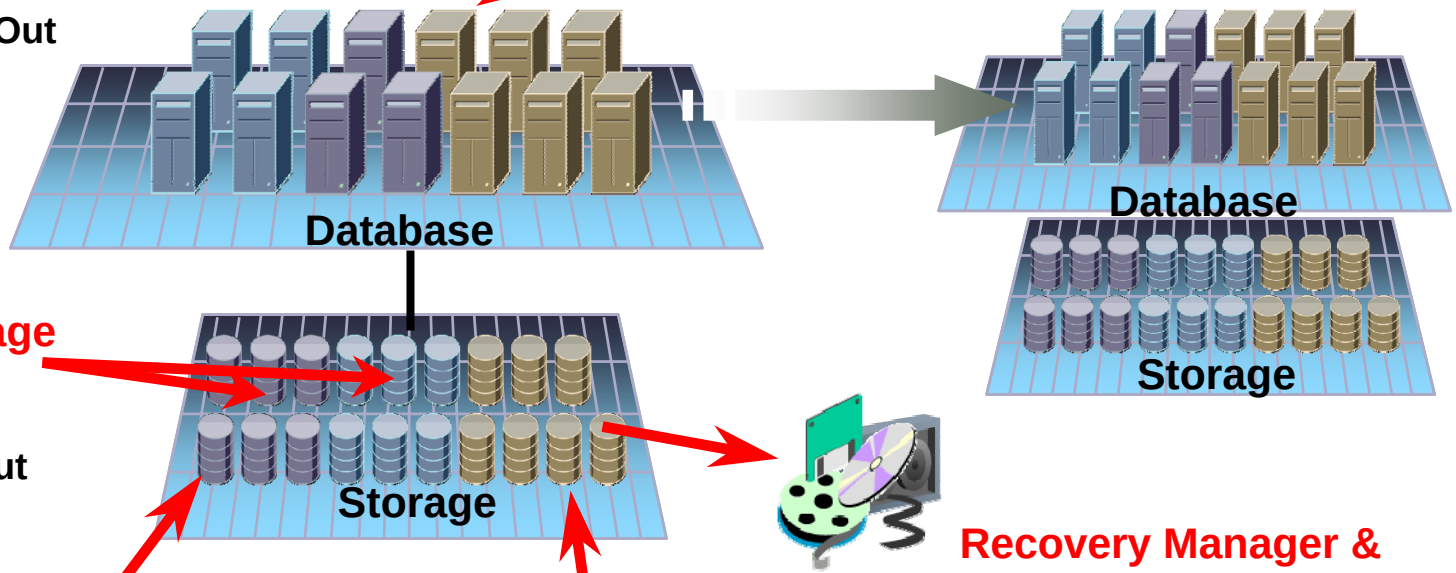

Automatic Storage Management
Fault Tolerant
Storage Scale-Out

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

Oracle MAA Changes Traditional HA/DR Paradigm

- Many businesses implement localized component level HA solutions
- DR is an afterthought, often implemented using mirroring technologies which do not offer adequate protection
  - Correlated failures, inter-component failures, software failures, upgrades, etc. remain significant vulnerabilities
  - Requires integration of disparate technologies
- MAA: integration of HA and DR
  - Data Guard standby database becomes an essential HA element of any systems architecture
    - Integrated with RAC for server HA
    - Provides highly effective fault isolation
    - Capable of failovers in seconds, with zero data loss
    - Standby database provides a productive computing resource

Resources

- **Maximum Availability Architecture white papers:**
  http://www.oracle.com/technology/deploy/availability/htdocs/maa.htm
- **Oracle HA Portal on OTN:**
  http://www.oracle.com/technology/deploy/availability/
- **Oracle HA Customer Success Stories on OTN:**
  http://www.oracle.com/technology/deploy/availability/htdocs/HA_CaseStudies.html


Q&A
QUESTIONS
ANSWERS