ORACLE

Multitenant MAA Solutions

Blueprints for reduced planned and unplanned downtime for the Multitenant Oracle Database Architecture

November 2022

Program Agenda

1. Multitenant Overview and Benefits
2. Exadata and Resource Management Benefits
3. Choosing the Best Exadata Multitenant MAA Solution to Meet SLAs
4. Exadata Multitenant MAA Reference Architectures and Key Features

Maximum Availability Architecture (MAA)

Multitenant Overview and Benefits

Advantages of Multitenant Architecture

Isolation and agility with economies of scale

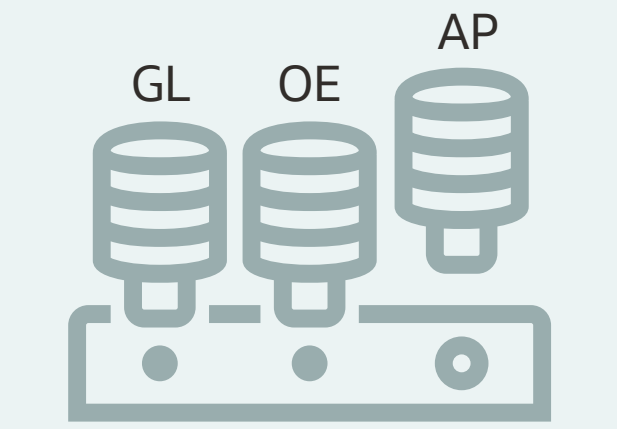

Complementary to VMs

Self-contained PDB for each application
- Applications run unchanged
- Rapid provisioning (via clones)
- Portability (via pluggability)

Common operations performed at CDB level
- Manage many as one (upgrade, HA, backup)
- Granular control when appropriate

Shared memory and background processes
- More applications per server

Oracle Multitenant Features

Rapid cloning and provisioning

- Local clones and remote clones
- Snapshot clones
- Refreshable PDBs

Manage many as one

- Database consolidation
- Improve productivity
- Maintain granular control

Improve agility for development teams

- Pre-configured service level agreement
- Compatibility
- Interface

Enhance security

- Separation of duties
- Data security
- Resource isolation

Integration with Oracle RAC

- High availability
- Scalability
- Flexibility

Exadata and Resource Management

- Exadata is best database consolidation platform.
- Consolidate databases with similar SLAs, planned maintenance windows and DR requirements within the same container (CDB).
- Resource Management manages and isolates resources for applications.
- Critical Success Factors for Best Oracle Database Consolidation

Consolidation Architecture Best Practices

Multitenant reduce overall costs

Minimize OpEx

- Manage many databases as one (Less administration)
- One upgrade, one HA solution, one backup, possibly one DR solution (more on this later)

Enable higher consolidation densities

- Share resources: CPU, memory, background processes

Mission Critical Applications/databases may require more isolation

- Non-CDBs are supported up to 19c
- Avoiding oversubscribing PDBs to prioritize predictable performance, HA, and DR
- PDBs within the same CDB should have the same/similar SLAs and planned maintenance windows (e.g. DR tests)

And you may need multiple CDBs!

- For multiple Oracle versions, e.g., 18c vs 19c
- For varying Oracle options, e.g., partitioning, RAC, Data Guard
- For different performance, HA & DR, security and consolidation density requirements
- For different planned maintenance window or DR testing requirements

Best Practice Consolidation Architecture

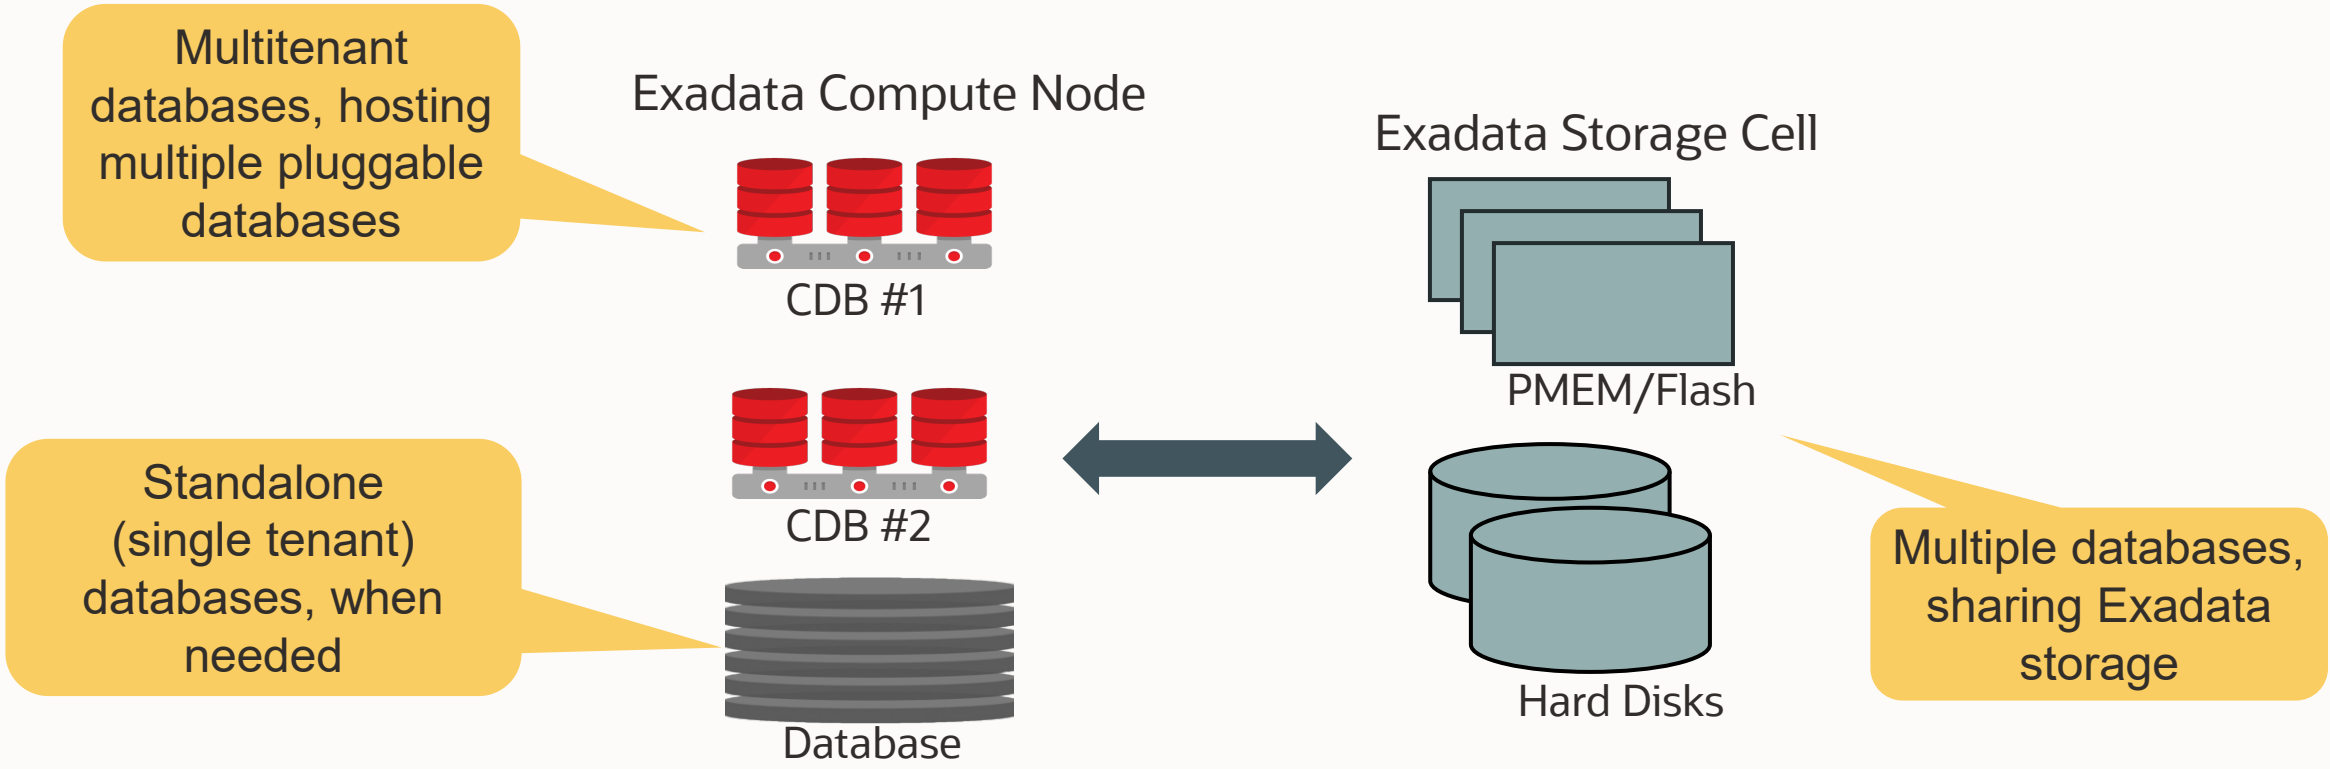

Best practice and most prevalent consolidation architecture:
Multiple multitenant databases on Exadata!

Best Practice Consolidation Architecture

Why Exadata?

Smart storage offloads database servers

- Smart scan, columnar features, storage indexes decrease the server and network load
- Improves consolidation density!

Built-in consolidation support

- Network prioritization for OLTP
- Resource management of PMEM and flash cache space, flash I/Os, disk I/Os
- Enables consolidation of mixed workloads
- Enables consolidation of production and dev/test databases
- Not available from any other storage vendor!

Resource Management

Consolidation Goals

Noisy neighbors

- One database's high loads should not disturb the others
- Data warehouses must co-exist with OLTP databases

Fair access to resources

- Databases need configurable, guaranteed access to CPU, memory, I/O, etc.

Options to capitalize on excess or idle resources

- Allow databases to "burst" into unused resources
- Or limit usage for "pay for performance" clouds or to ensure predictable performance

Dynamic resize of database resource allocations

- Change resource allocations without restarts

Resource Management

Which Resources?

For PDBs:
- CPU
- SGA
- PGA
- Sessions
- Parallel servers

For Databases:
- CPU
- SGA
- PGA
- RAC network

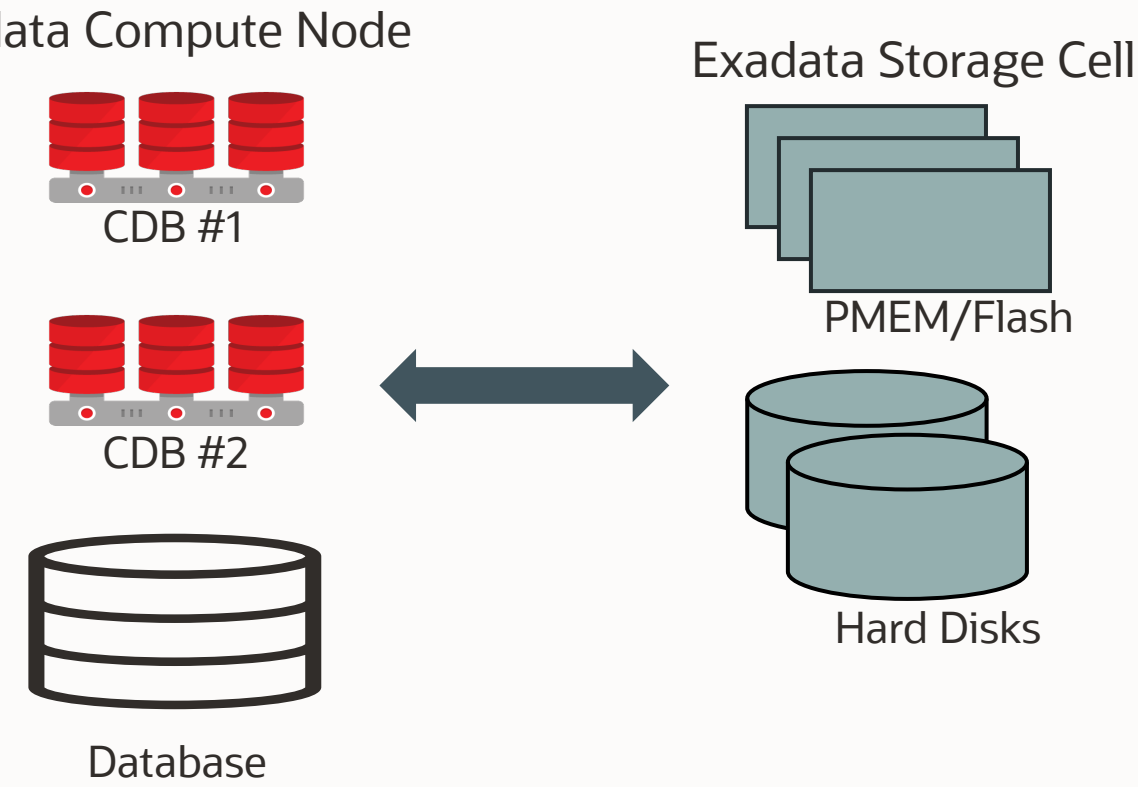

For Exadata cells
PMEM & flash space
Flash bandwidth
Disk bandwidth
Storage network

Resource Manager controls access to all critical resources.

Maximum Availability Architecture (MAA)

Choosing Exadata Multitenant MAA Architecture to Meet SLAs

Multitenant MAA Considerations & Solution Flow

1. Key Questions to determine SLAs

2. Understands MAA solution offerings that meets SLAs

3. Customer picks best reference architecture in the cloud or on-premise

4. Implements Best Practices

Key Questions

1. *What's the maximum RTO SLA for local failures?*
   - Pick one: Near Zero or Secs (HA) or hour(s)
   - If "hour(s)", skip all questions and choose Bronze MAA. Otherwise, minimally Silver MAA.
2. *What's the maximum RTO for disasters?*
   - Pick one: Near Zero, < 2 mins, < 15 mins, Hours
   - Hours: Bronze or Silver MAA
   - Less than 15 mins, Gold MAA or "Aurous" Option strategy (introduced later in deck)
   - Less than 2 mins, Gold MAA option strategy (possible with restricted number of databases and services)
   - Near Zero: Platinum MAA
3. *What's the maximum data loss tolerated (RPO SLA) in case of disaster?*
   - Pick one: Zero, Seconds, < 5 minutes, < 30 minutes
   - Less than 30 minutes: Available in all MAA tiers
   - Less than 5 minutes: Gold MAA or "Aurous" Option strategy (introduced later in deck)
   - Near Zero or Seconds: Gold MAA with async transport with strategy
   - Zero: Gold or Platinum MAA with sync transport

Key Questions

4. Does the primary and standby resources have to be identical to preserve same application performance after a role transition?
   - Yes: Symmetric Standby with same resource allocation
   - No: Asymmetric standby is allowed
5. What is the minimum isolation and distance for DR solution if required?
   - Meters: Fault Domain isolated system resources and power but in same data center
   - KMs: Availability Domain is different data center in the same region (e.g., < 25Kms)
   - 100s to 1000s KMs: Region in different cities or countries
6. Can the application fail over to remote database and still meet performance requirements?
   - Yes – implement transparent app failover following MAA best practices
   - No – A separate app tier may be required for DR. For DR, a site failover may include app and database failover

Picking MAA Reference Architecture

MAA Reference Architecture	RTO for HA / Downtime for Software Updates	RTO for DR / Downtime for DB Upgrades	RPO for DR	On-Premise	Oracle Cloud
Bronze MAA	Hour(s)	Hours	< 15 or 30 mins* Secs with ZDLRA*	Yes	Yes
Silver MAA	Zero or Secs	Hours	< 15 or 30 mins* Secs with ZDLRA*	Yes – Exadata, RAC	Yes- ExaDB, ExaCC, ADB
Gold MAA	Zero or Secs	< 2 mins	Zero or secs Secs with ZDLRA*	Yes RAC & ADG	Yes – ExaDB ExaCC, ADB-D ***
Platinum MAA	Zero or Secs	Zero or secs	Zero or secs	Yes RAC, ADG,+GG	Yes – ExaDB/ExaCC ***

All RTO and RPO values are service level **objectives** that have been validated. To achieve these values specific MAA config and operational practices are in place. Note all DR operations are fully automatic today. The RTO for DR calculation is based on after the failover is automatically or manually triggered.

* Based on archive backup frequency. For ADB, it's < 15 mins. With ZDLRA, seconds. ZDLRA is available for On-premise and ExaCC

** Based on online redo log push

*** Data Guard Fast-Start Failover (FSFO) is manual in the cloud except for ADB-D which includes FSFO automation. Oracle GoldenGate setup is manual

Pluggable Database Placement Considerations

- Container should not mix databases with different SLAs
  - Test/Dev/Production that needs HA only should be Silver Container
  - Production Databases that need HA and DR should be in Gold Container with ADG
  - Production Databases that need HA, DR and logical replication should be in Platinum Container with ADG & GG
- All pluggable databases within a container should have
  - Similar planned maintenance windows (relevant for all tiers)
  - Same database version and upgrade schedules (relevant for all tiers)
  - Same DR test requirements (relevant for Gold and Platinum)
  - Same DR target destination requirements (relevant for Gold and Platinum)
- To reduce DR failover times
  - Minimize number of pluggable databases within a container for Gold/Platinum (<25 PDBs)
  - Minimize number of Clusterware managed services within a cluster (< 250 services/CDB)
  - Ensure sufficient system resources and symmetric standby
  - Test and tune using MAA practices

Pluggable Database Placement Considerations

PDB Open on Database Startup

All application and user access to a PDB should utilize user-defined services

- Applications should *never* utilize default PDB services

Controlled PDB open on database startup with Oracle Clusterware

- Configure preferred and available RAC instances and assigned database role to route services according for different outage cases
- Do NOT use AFTER STARTUP triggers or ALTER PLUGGABLE DATABASE.. SAVE STATE

Maximum Availability Architecture (MAA)

MAA Multitenant: Blueprints for Oracle Database HA & DR

Oracle Maximum Availability Architecture (MAA)

Standardized Reference Architectures for Never-Down Deployments

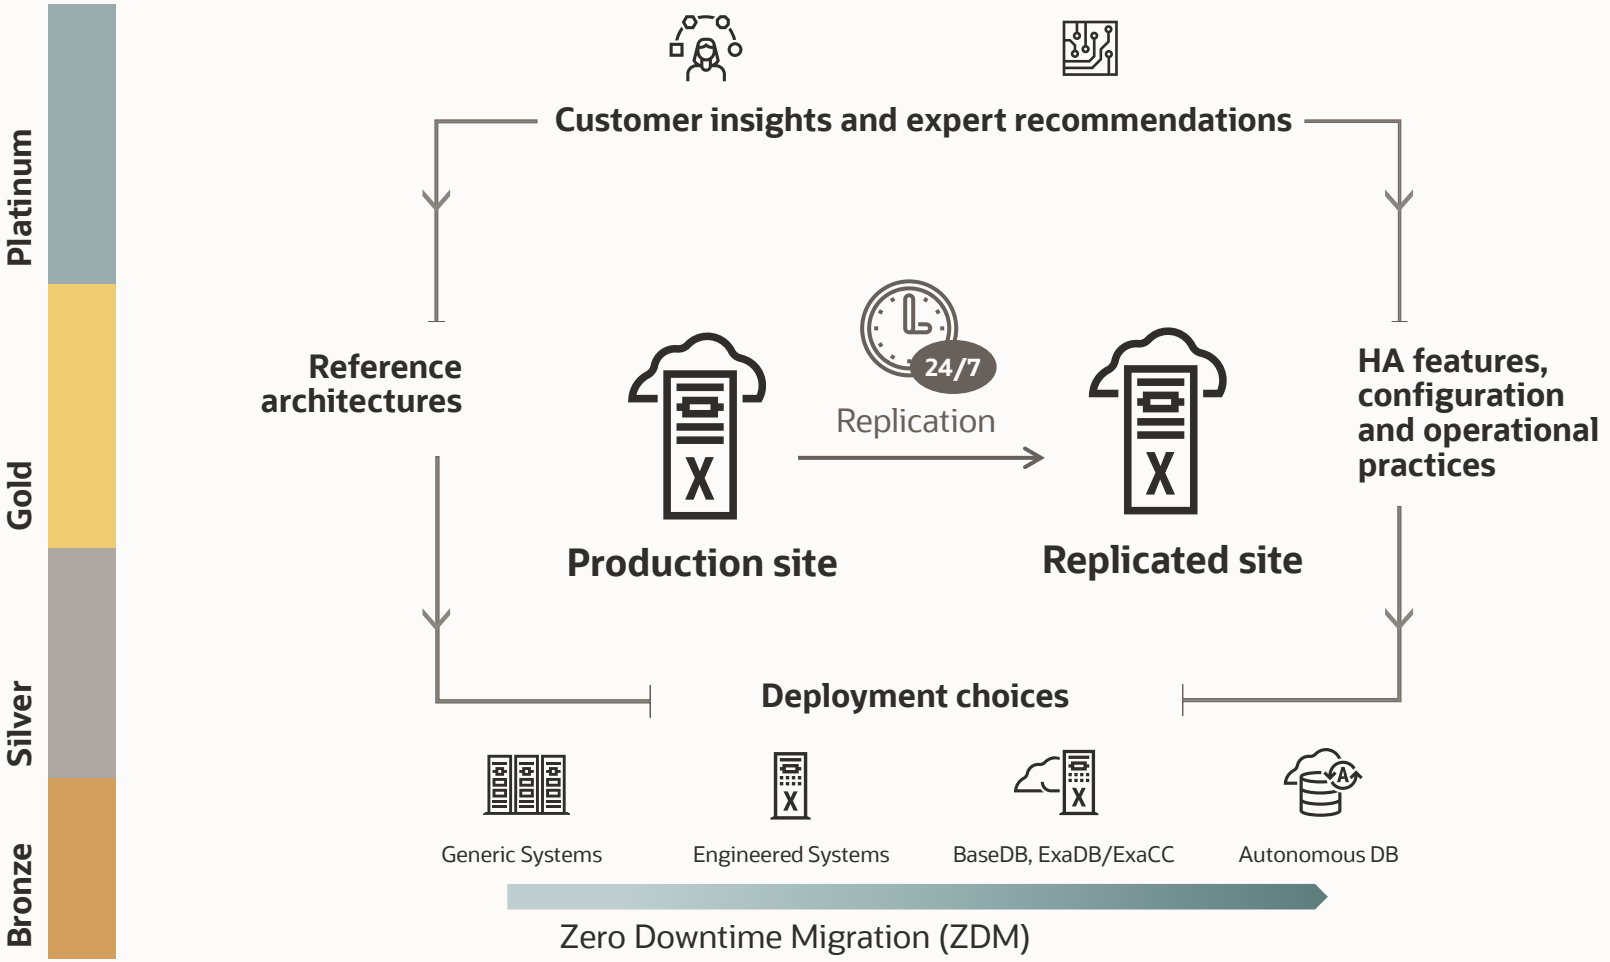

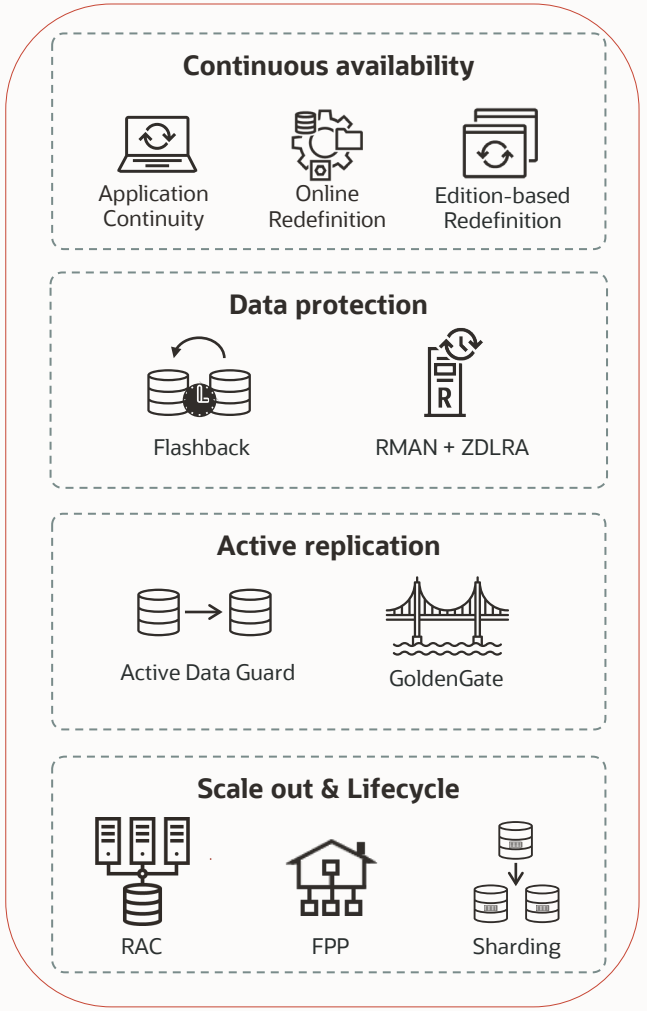

BRONZE

Dev, Test, Prod - Single Instance or Multitenant Database with Backups

- Single Instance with Clusterware Restart
- Advanced backup/restore with RMAN
  - Optional ZDLRA with incremental forever and near zero RPO
- Storage redundancy and validation with ASM
- Multitenant Database/Resource Management with PDB features
- Online Maintenance
- Some corruption protection
- Flashback technologies

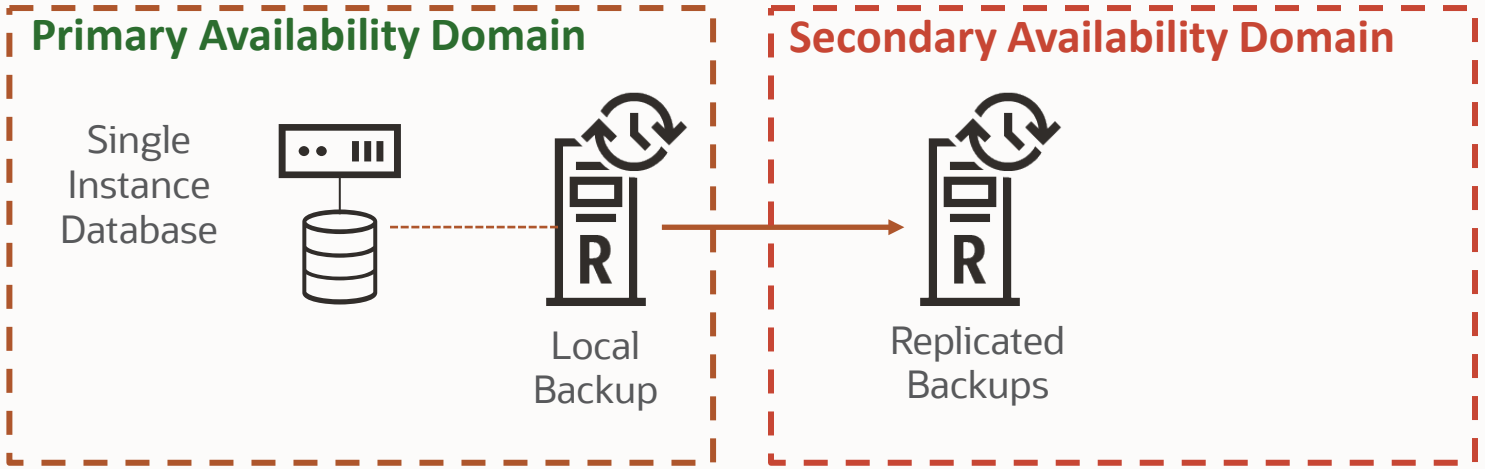

Outage Matrix

Unplanned Outage	RTO / RPO Service Level Objectives (f1)
Recoverable node or instance failure	Minutes to hour (f2)
Disasters: corruptions and site failures	Hours to days. RPO since last backup or near zero with ZDLRA
Planned Maintenance	
Software/hardware updates	Minutes to hour (f2)
Major database upgrade	Minutes to hour

f1 : RPO=0 unless explicitly specified
f2 : Exadata systems has RAC but Bronze Exadata configuration with Single Instance database running with Oracle Clusterware has highest consolidation density to reduce costs

Oracle Clusterware for Automatic Restart

1. Oracle Clusterware is available for all Oracle Databases
2. Enables HA capabilities and resource management:
   - Automatic Restart of database instances, listeners and other resources
   - Fleet patching
   - Service management including restarting service after failure
   - Automatic Storage Management (ASM) for HA, data protection and ease of use

- Trade off: additional software maintenance for Grid Infrastructure

Oracle Recovery Manager - RMAN

Database Integrated Backup and Recovery

- Unique knowledge of database fil formats and recovery procedures
  - Oracle block validation
  - Online block-level recovery
  - Native encryption, compression
  - Table/partition-level recovery
  - Oracle Multitenant support
- Tape and cloud backups
- Unified Management

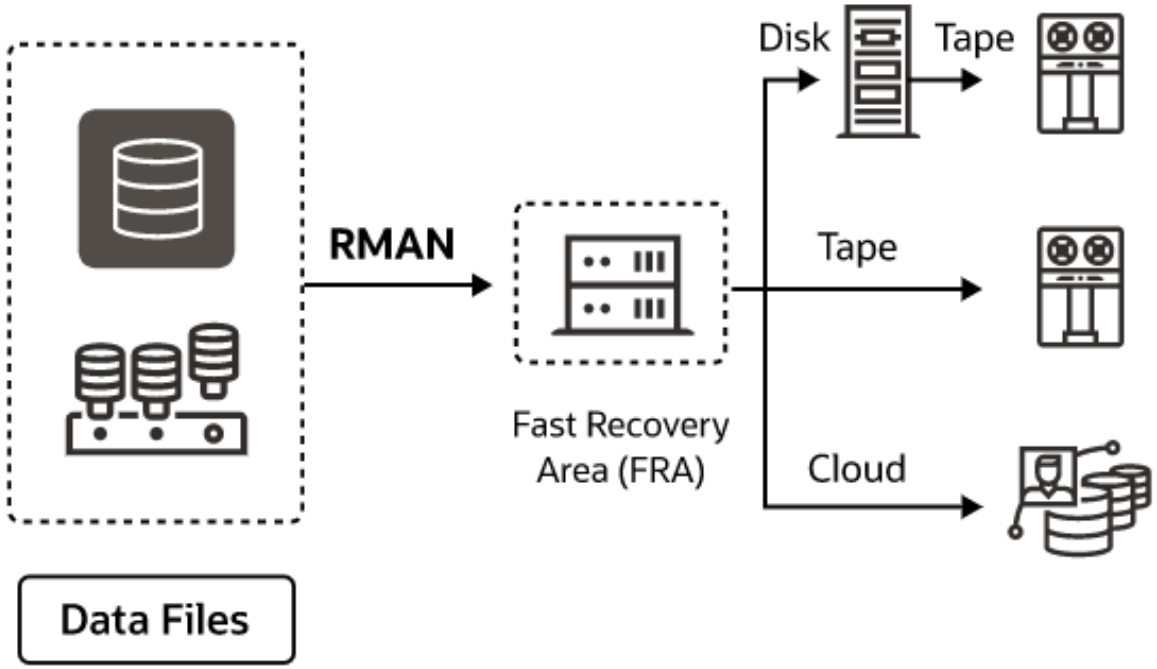

Recovery Appliance Recommended

EM Real-Time Protection Status & Space Monitoring

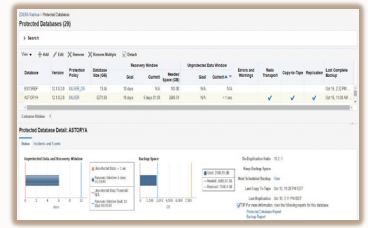

Databases

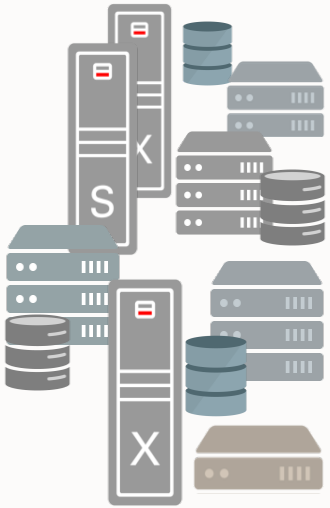

Transactional Block Changes

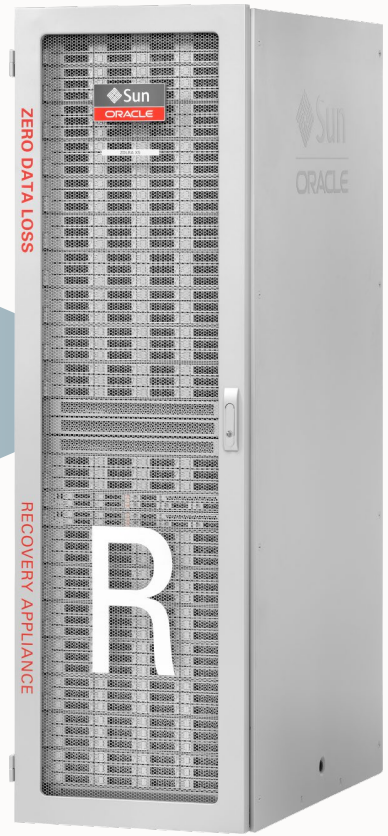

Day 1 Full

Day 2 Changes

Day N Changes

Day N State

Virtual Full Backup

Cloud Storage

Remote Replica

Tape

Oracle DB 12c-21c on Any Platform

No More Full Backups, Incremental Forever

End-to-End Oracle Recovery Validation
Near Zero Data Loss for DR

Database and Exadata Health Checks

Assessment Report

- Health Score, Summary, Findings

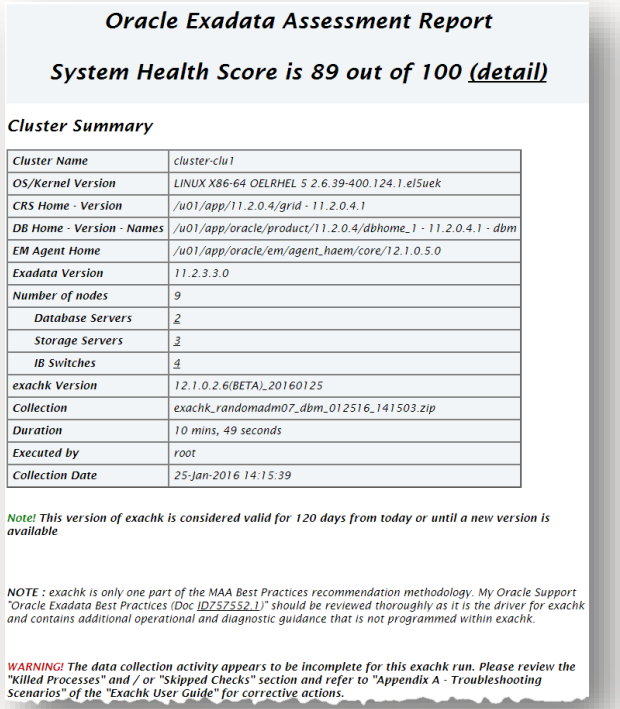

Oracle Exadata Assessment Report

System Health Score is 89 out of 100 (detail)

Cluster Summary

Cluster Name	cluster-clu1
OS/Kernel Version	LINUX X86-64 OELRHEL 5 2.6.39-400.124.1.el5uek
CRS Home - Version	/u01/app/11.2.0.4/grid - 11.2.0.4.1
DB Home - Version - Names	/u01/app/oracle/product/11.2.0.4/dbhome_1 - 11.2.0.4.1 - dbm
EM Agent Home	/u01/app/oracle/em/agent_haem/core/12.1.0.5.0
Exadata Version	11.2.3.3.0
Number of nodes	9
Database Servers	2
Storage Servers	3
IB Switches	4
exachk Version	12.1.0.2.6(BETA)_20160125
Collection	exachk_randomadm07_dbm_012516_141503.zip
Duration	10 mins, 49 seconds
Executed by	root
Collection Date	25-Jan-2016 14:15:39

Note! This version of exachk is considered valid for 120 days from today or until a new version is available

NOTE : exachk is only one part of the MAA Best Practices recommendation methodology. My Oracle Support "Oracle Exadata Best Practices (Doc ID757552.1)" should be reviewed thoroughly as it is the driver for exachk and contains additional operational and diagnostic guidance that is not programmed within exachk.

WARNING! The data collection activity appears to be incomplete for this exachk run. Please review the "Killed Processes" and / or "Skipped Checks" section and refer to "Appendix A - Troubleshooting Scenarios" of the "Exachk User Guide" for corrective actions.

Findings & Recommendations

- How to Solve the problem?

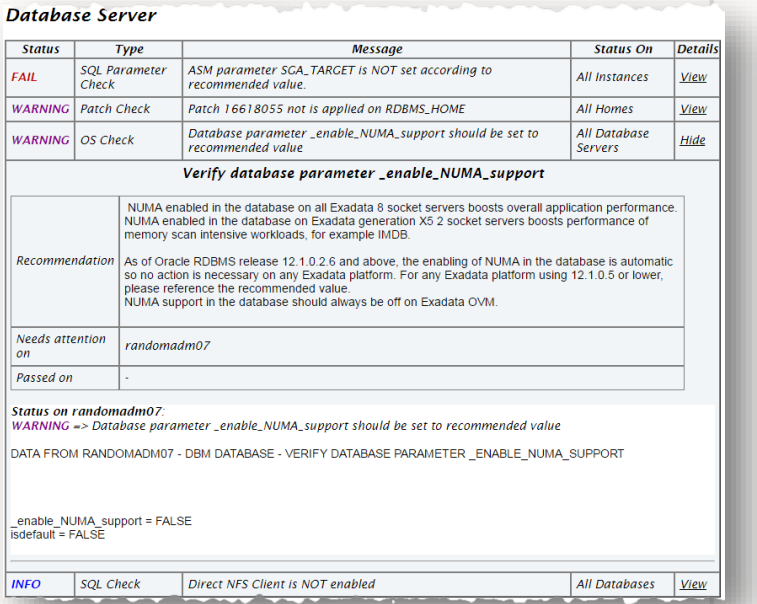

Database Server

Status	Type	Message	Status On	Details
FAIL	SQL Parameter Check	ASM parameter SGA_TARGET is NOT set according to recommended value.	All Instances	View
WARNING	Patch Check	Patch 16618055 not is applied on RDBMS_HOME	All Homes	View
WARNING	OS Check	Database parameter _enable_NUMA_support should be set to recommended value	All Database Servers	Hide

Verify database parameter _enable_NUMA_support

Recommendation	NUMA enabled in the database on all Exadata 8 socket servers boosts overall application performance. NUMA enabled in the database on Exadata generation X5 2 socket servers boosts performance of memory scan intensive workloads, for example IMDB. As of Oracle RDBMS release 12.1.0.2.6 and above, the enabling of NUMA in the database is automatic so no action is necessary on any Exadata platform. For any Exadata platform using 12.1.0.5 or lower, please reference the recommended value. NUMA support in the database should always be off on Exadata OVM.
Needs attention on	randomadm07
Passed on	-

Status on randomadm07:
WARNING => Database parameter _enable_NUMA_support should be set to recommended value

DATA FROM RANDOMADM07 - DBM DATABASE - VERIFY DATABASE PARAMETER _ENABLE_NUMA_SUPPORT

_enable_NUMA_support = FALSE
isdefault = FALSE

INFO	SQL Check	Direct NFS Client is NOT enabled	All Databases	View

MAA Score Card

- MAA architectural readiness and configuration practices

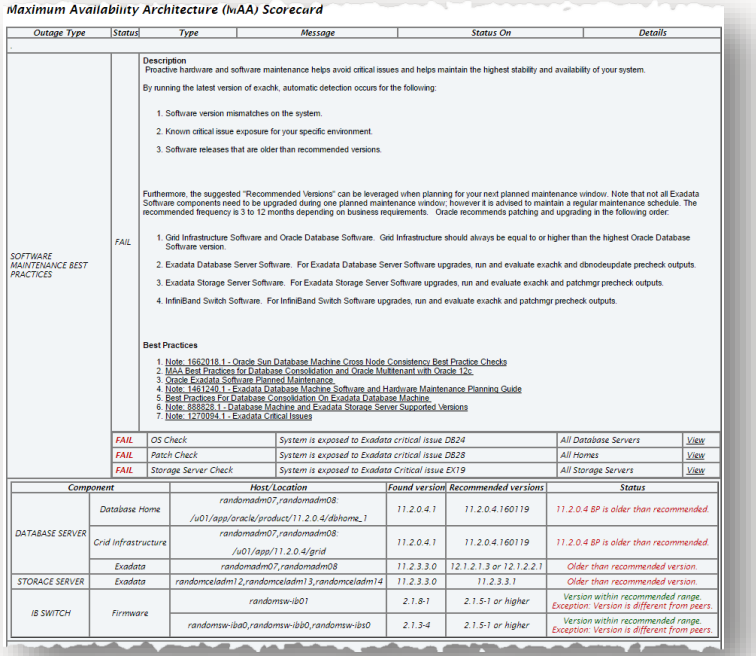

Maximum Availability Architecture (MAA) Scorecard

Outage Type	Status	Type	Message	Status On	Details
SOFTWARE MAINTENANCE BEST PRACTICES	FAIL		Description Proactive hardware and software maintenance helps avoid critical issues and helps maintain the highest stability and availability of your system. By running the latest version of exachk, automatic detection occurs for the following: 1. Software version mismatches on the system. 2. Known critical issue exposure for your specific environment. 3. Software releases that are older than recommended versions. Furthermore, the suggested "Recommended Versions" can be leveraged when planning for your next planned maintenance window. Note that not all Exadata Software components need to be upgraded during one planned maintenance window; however it is advised to maintain a regular maintenance schedule. The recommended frequency is 3 to 12 months depending on business requirements. Oracle recommends patching and upgrading in the following order: 1. Grid Infrastructure Software and Oracle Database Software. Grid Infrastructure should always be equal to or higher than the highest Oracle Database Software version. 2. Exadata Database Server Software. For Exadata Database Server Software upgrades, run and evaluate exachk and dbnodeupdate precheck outputs. 3. Exadata Storage Server Software. For Exadata Storage Server Software upgrades, run and evaluate exachk and patchmgr precheck outputs. 4. InfiniBand Switch Software. For InfiniBand Switch Software upgrades, run and evaluate exachk and patchmgr precheck outputs. Best Practices 1. Note: 1662018.1 - Oracle Sun Database Machine Cross Node Consistency Best Practice Checks 2. MAA Best Practices for Database Consolidation and Oracle Multitenant with Oracle 12c 3. Oracle Exadata Software Planned Maintenance 4. Note: 1461240.1 - Exadata Database Machine Software and Hardware Maintenance Planning Guide 5. Best Practices For Database Consolidation On Exadata Database Machine 6. Note: 888828.1 - Database Machine and Exadata Storage Server Supported Versions 7. Note: 1270094.1 - Exadata Critical Issues		
	FAIL	OS Check	System is exposed to Exadata critical issue DB24	All Database Servers	View
	FAIL	Patch Check	System is exposed to Exadata critical issue DB28	All Homes	View
	FAIL	Storage Server Check	System is exposed to Exadata Critical issue EX19	All Storage Servers	View

Component		Host/Location	Found version	Recommended versions	Status
DATABASE SERVER	Database Home	randomadm07,randomadm08: /u01/app/oracle/product/11.2.0.4/dbhome_1	11.2.0.4.1	11.2.0.4.160119	11.2.0.4 BP is older than recommended.
	Grid Infrastructure	randomadm07,randomadm08: /u01/app/11.2.0.4/grid	11.2.0.4.1	11.2.0.4.160119	11.2.0.4 BP is older than recommended.
	Exadata	randomadm07,randomadm08	11.2.3.3.0	12.1.2.1.3 or 12.1.2.2.1	Older than recommended version.
STORAGE SERVER	Exadata	randomceladm12,randomceladm13,randomceladm14	11.2.3.3.0	11.2.3.3.1	Older than recommended version.
IB SWITCH	Firmware	randomsw-ib01	2.1.8-1	2.1.5-1 or higher	Version within recommended range. Exception: Version is different from peers.
		randomsw-ib0,randomsw-ib1,randomsw-ib2	2.1.3-4	2.1.5-1 or higher	Version within recommended range. Exception: Version is different from peers.

Note: Automated Orachk/Exachk Healthcheck MOS 107954.1 updated frequently

Online Operations

Online Redefinition Improvements

DBMS_REDEFINITION allows you to reorganize and redefine tables online

- Add/drop/rename/reorder columns
- Switch physical storage structures
- Reorganize & transform data while online

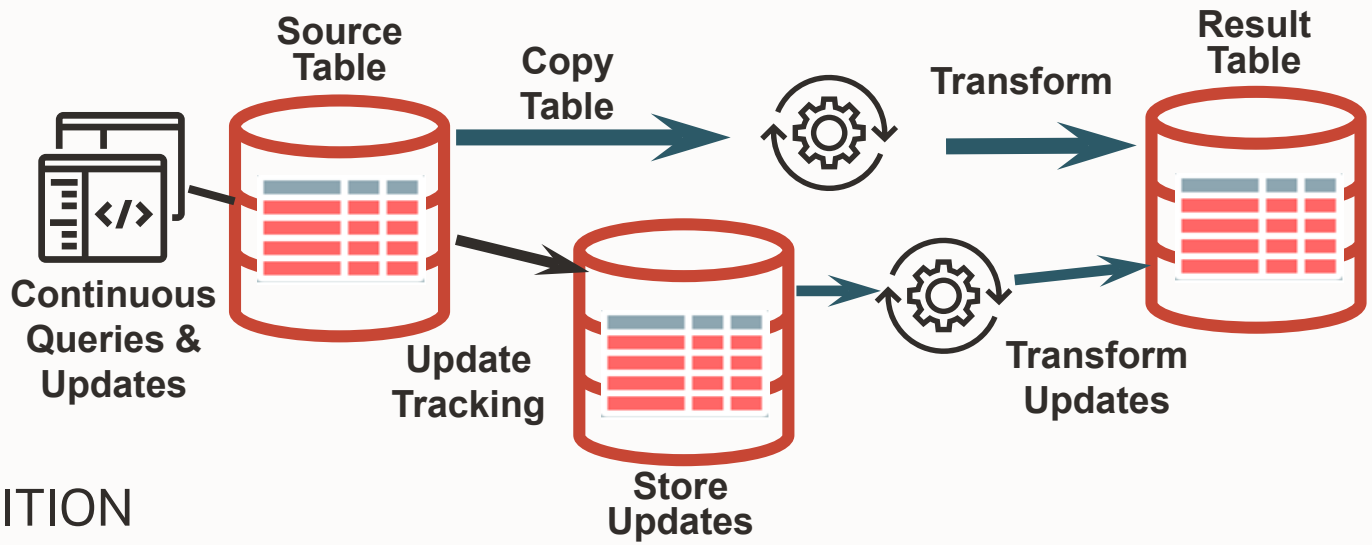

Additional Benefits of using DBMS_REDEFINITION

- Fault Tolerant (resume at point of failure) and track changes to enable fast rollback to prior definition
- Entire redefinition process runs without acquiring Exclusive DDL lock
- Monitor reorganization using V$online_redef

Flashback Technologies

Rewind Button for Oracle Databases

- Fast point-in-time recovery (PITI without expensive restore opera
- Error investigation
  - View data as of previous point in ti
- Error correction
  - Back-out a transaction
  - Incorrect table updates
  - Rewind the entire database

@T1	Col-1	Col-..	Col-n
Row-1	Abby	1234	officer
Row-2	Ben	8834	mgr
Row-3	Charlie	9837	officer
Row-n	Tom	8793	vp

@T2	Col-1	Col-..	Col-n
Row-1	Tom	1234	vp
Row-2	Ben	8834	vp
Row-3	Charlie	9837	vp
Row-n	Tom	8793	vp

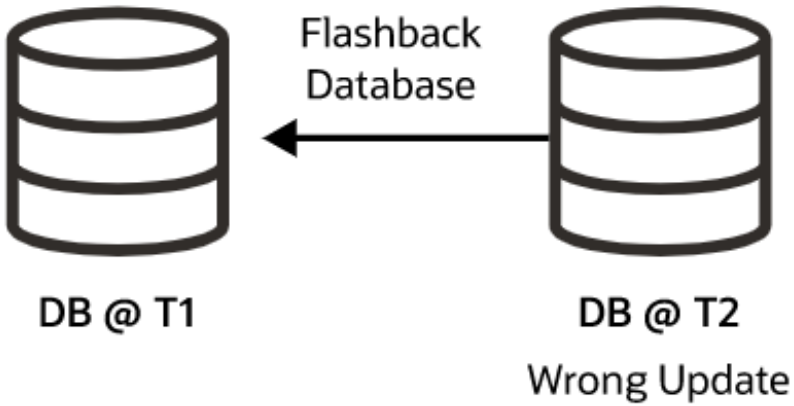

SILVER

Prod/Departmental

Bronze +

- Real Application Clustering (RAC)
- Application Continuity
- Sharding (Optional)
  - Provides fault isolation, scalability and geographical distribution

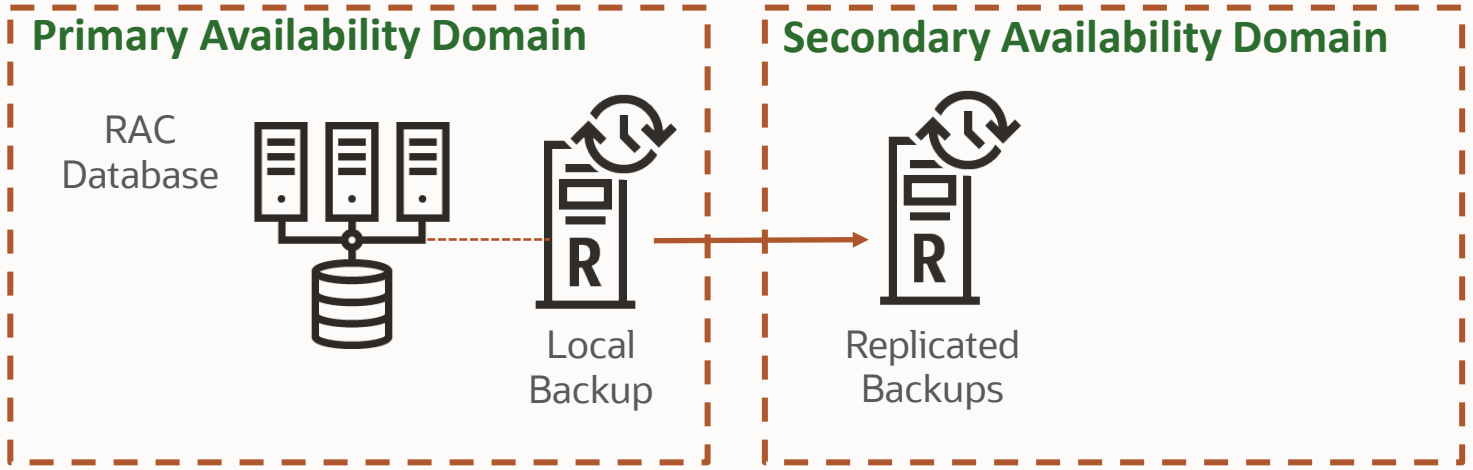

Outage Matrix

Unplanned Outage	RTO/RPO Service Level Objectives(f1)
Recoverable node or instance failure	**Single digit seconds** (f2)
Disasters: corruptions and site failures	Hours to days. RPO since last backup or near zero with ZDLRA
Planned Maintenance	
Software/Hardware updates	**Zero** (f2)
Major database upgrade	Minutes to hour

f1: RPO=0 unless explicitly specified
f2: To achieve zero downtime or lowest impact, apply application checklist best practices; Batch jobs should be deferred outside planned maintenance window.

Checklist found in MAA OTN

https://www.oracle.com/a/tech/docs/application-checklist-for-continuous-availability-for-maa.pdf

Oracle Real Application Clusters (Oracle RAC)

Node Failure, Instance Failure, Rolling Maintenance

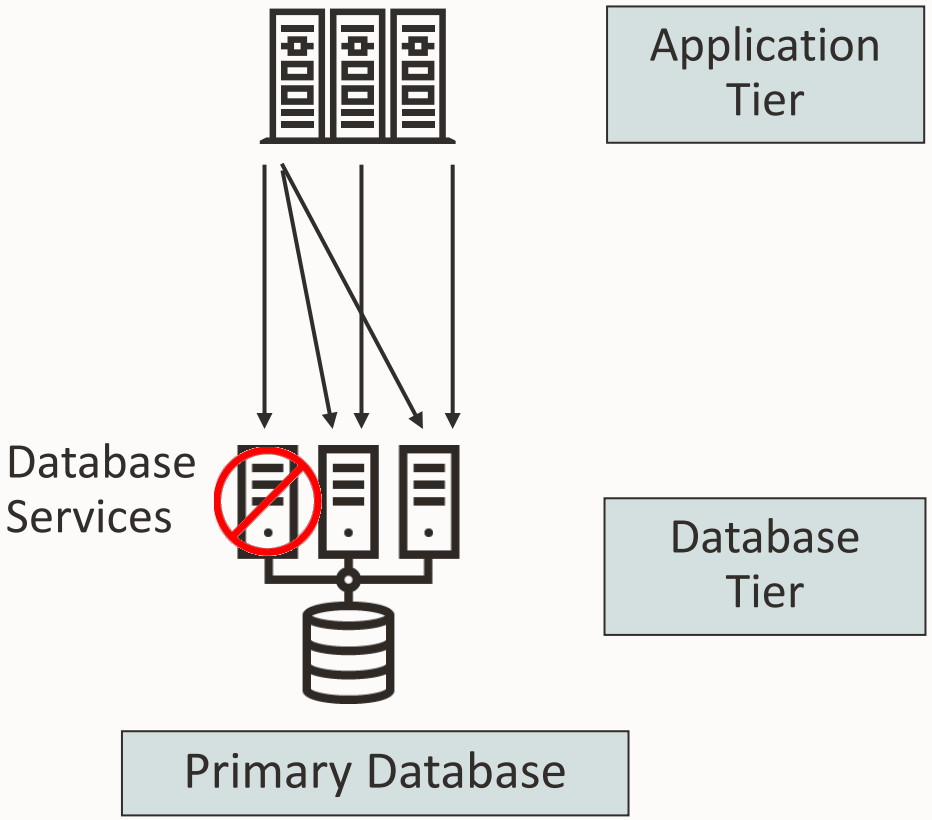

- Utilizes two or more instances of an Oracle Database concurrently
- Very Scalable
  - All instances active; Add capacity online; Ideal for database consolidation
- Highly Available
  - Auto-failover of services to an already running instance
  - Outage is transparent to user, in-flight transactions succeed
  - Zero downtime rolling maintenance

Multitenant and RAC: Affinity

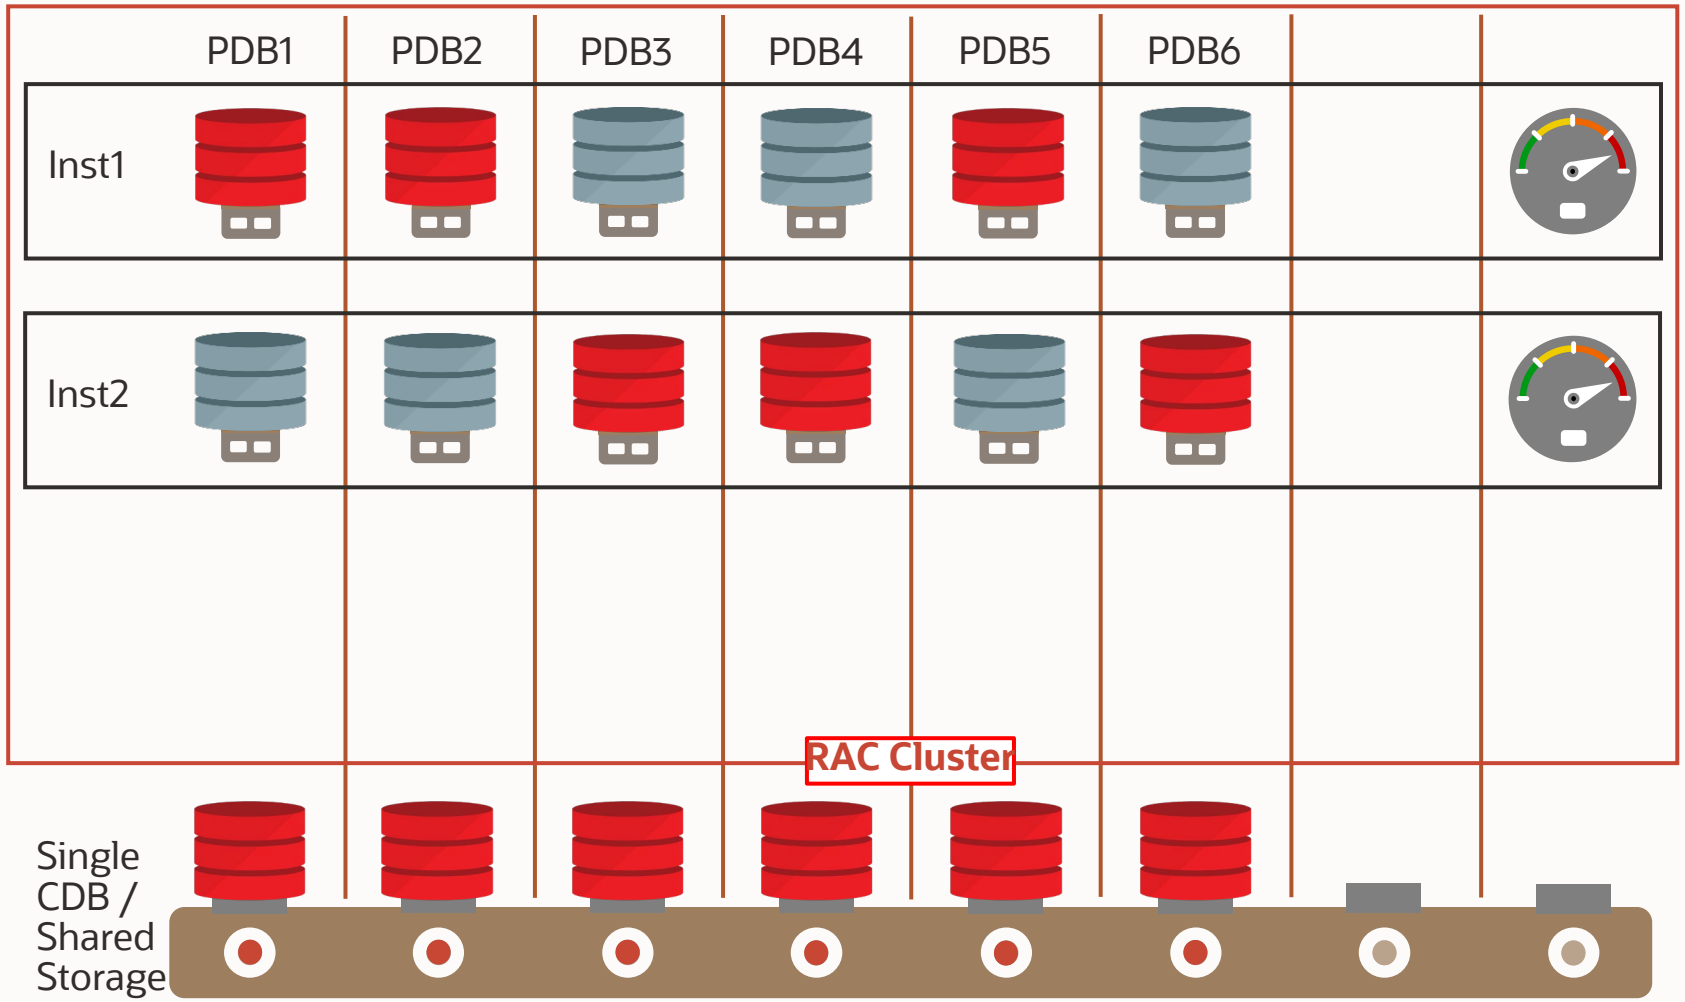

- Single CDB
- Two-node cluster
- Single instance per node
- PDB's affinity to a node defined by starting its services there
  - Present in "mounted" state in other nodes

Multitenant and RAC: Scalability and Agility

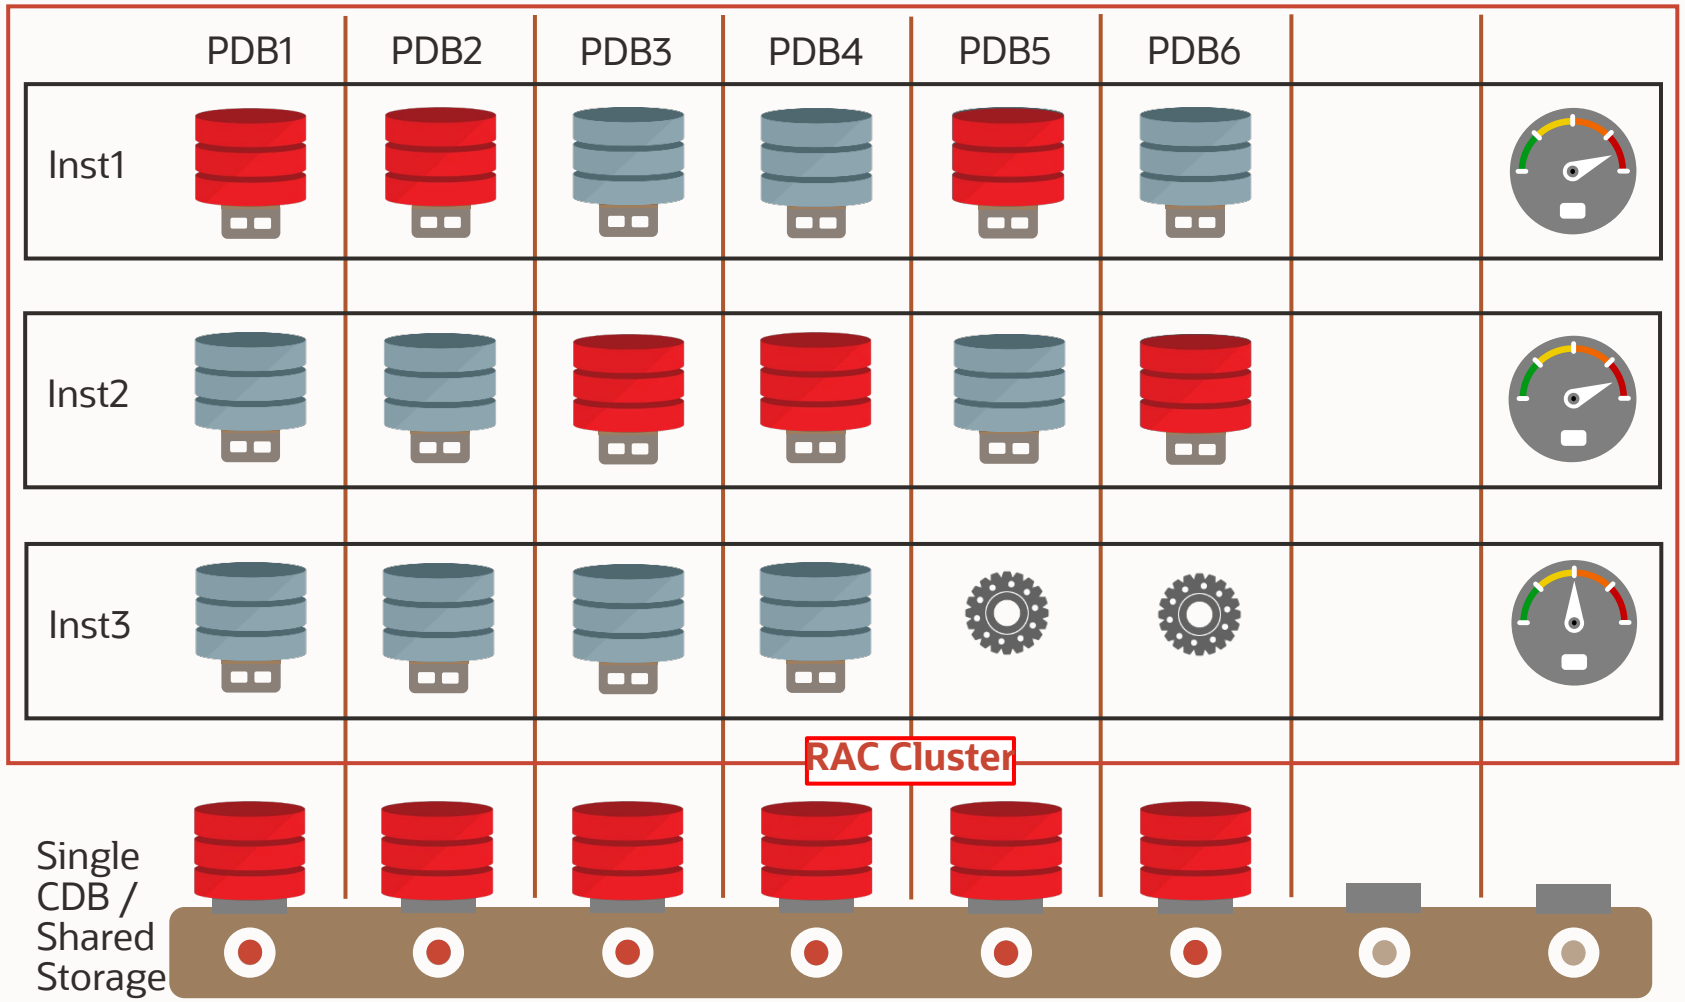

- Expand cluster
- Redistribute PDBs
  - Use Srvctl

Multitenant and RAC: Scalability and Agility

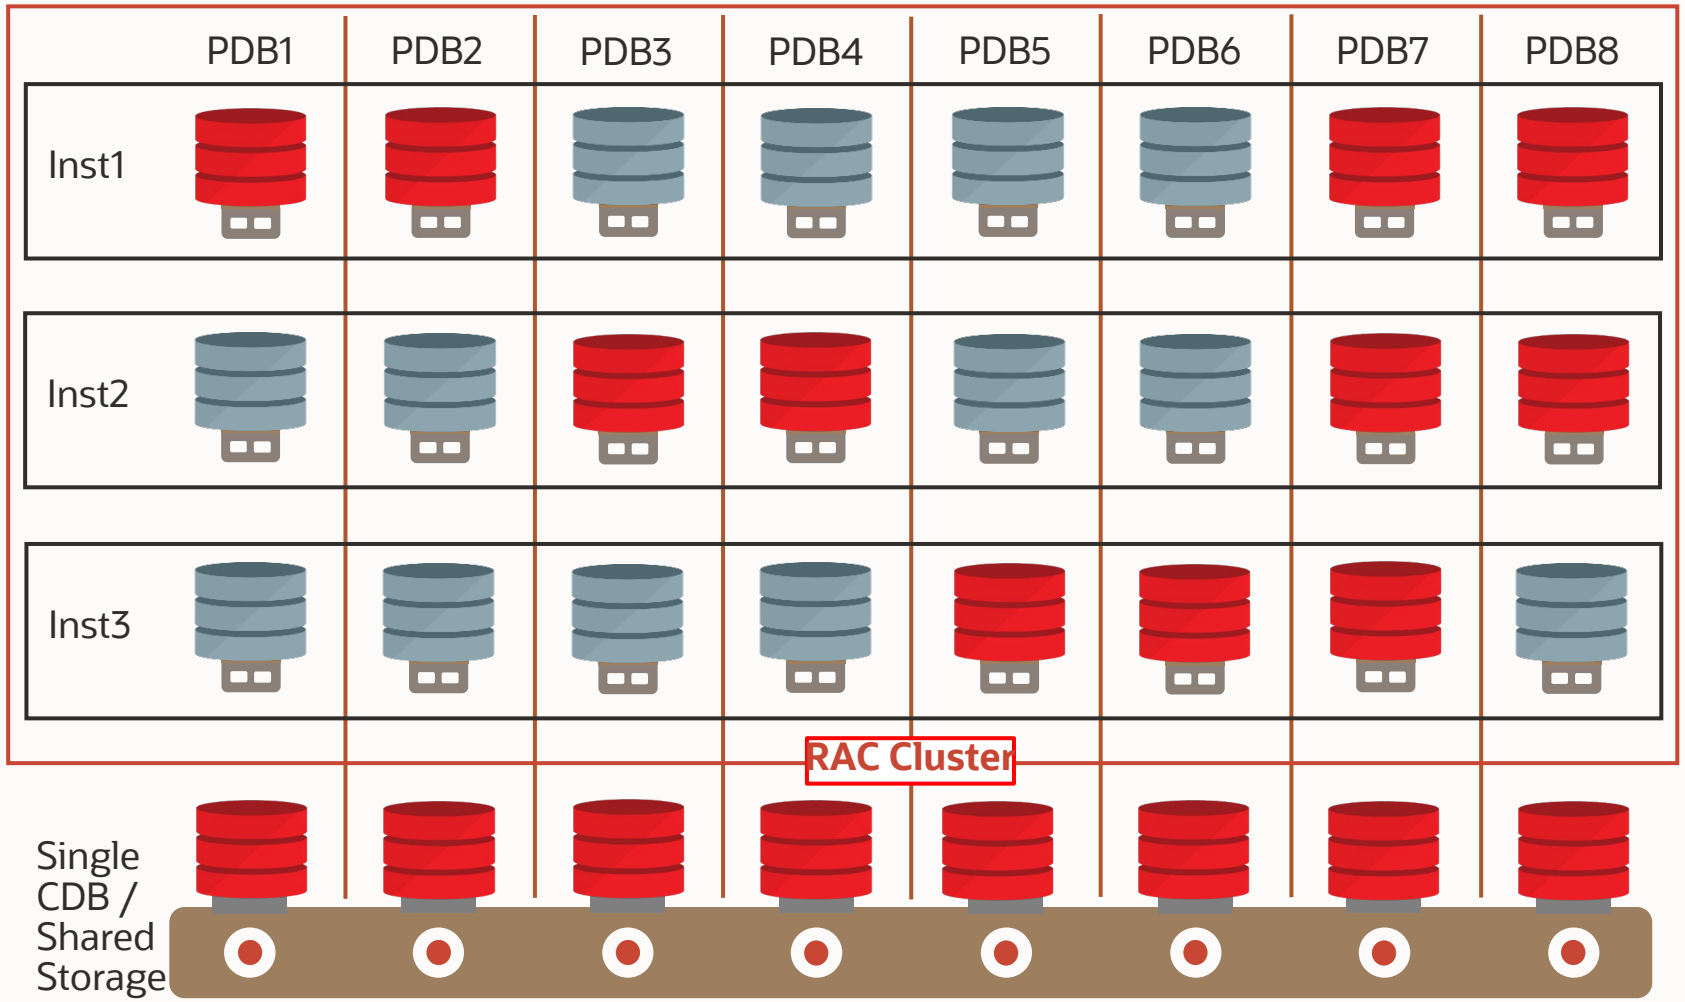

- Single CDB
- Single instance per node
- PDBs may be configured with "singleton" affinity to a specific node
  - Present in "mounted" state in other nodes
- PDBs may be open in multiple nodes

Singleton Affinity of PDB to Instance in RAC Cluster

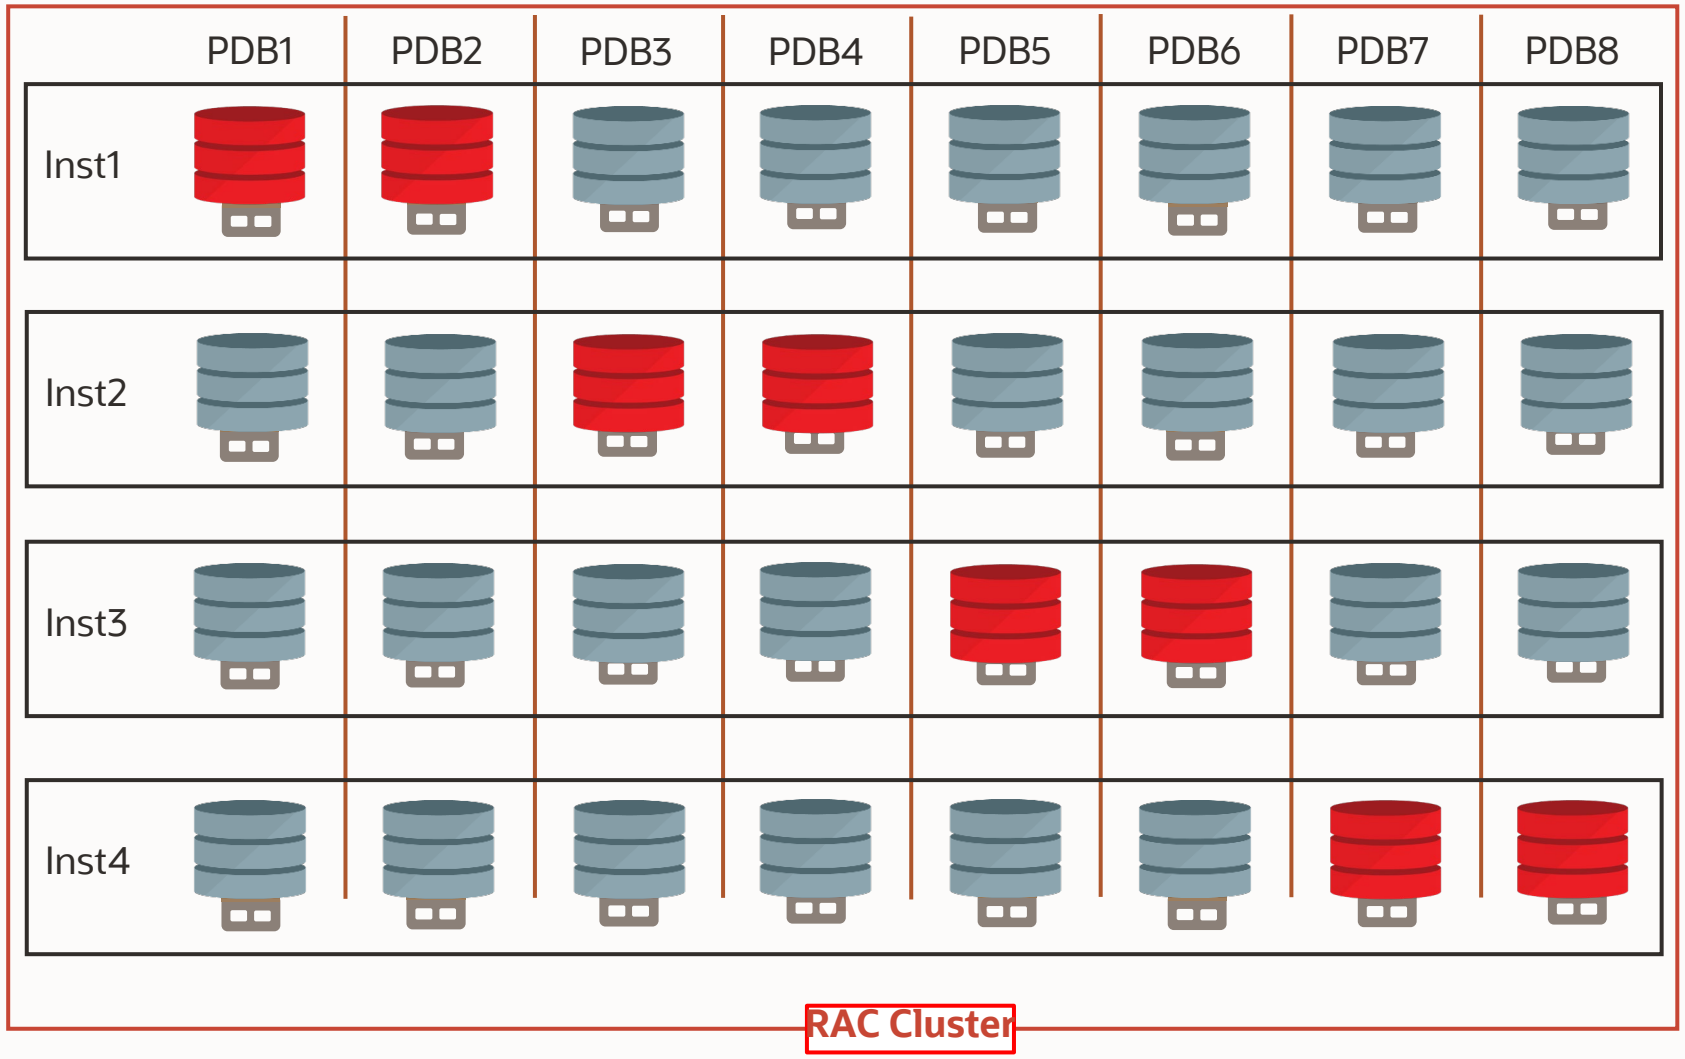

- Per-PDB lock domain provides strong isolation
- *Cache Fusion* following instance failure can ignore PDBs not open in affected nodes

Multitenant and RAC: Availability

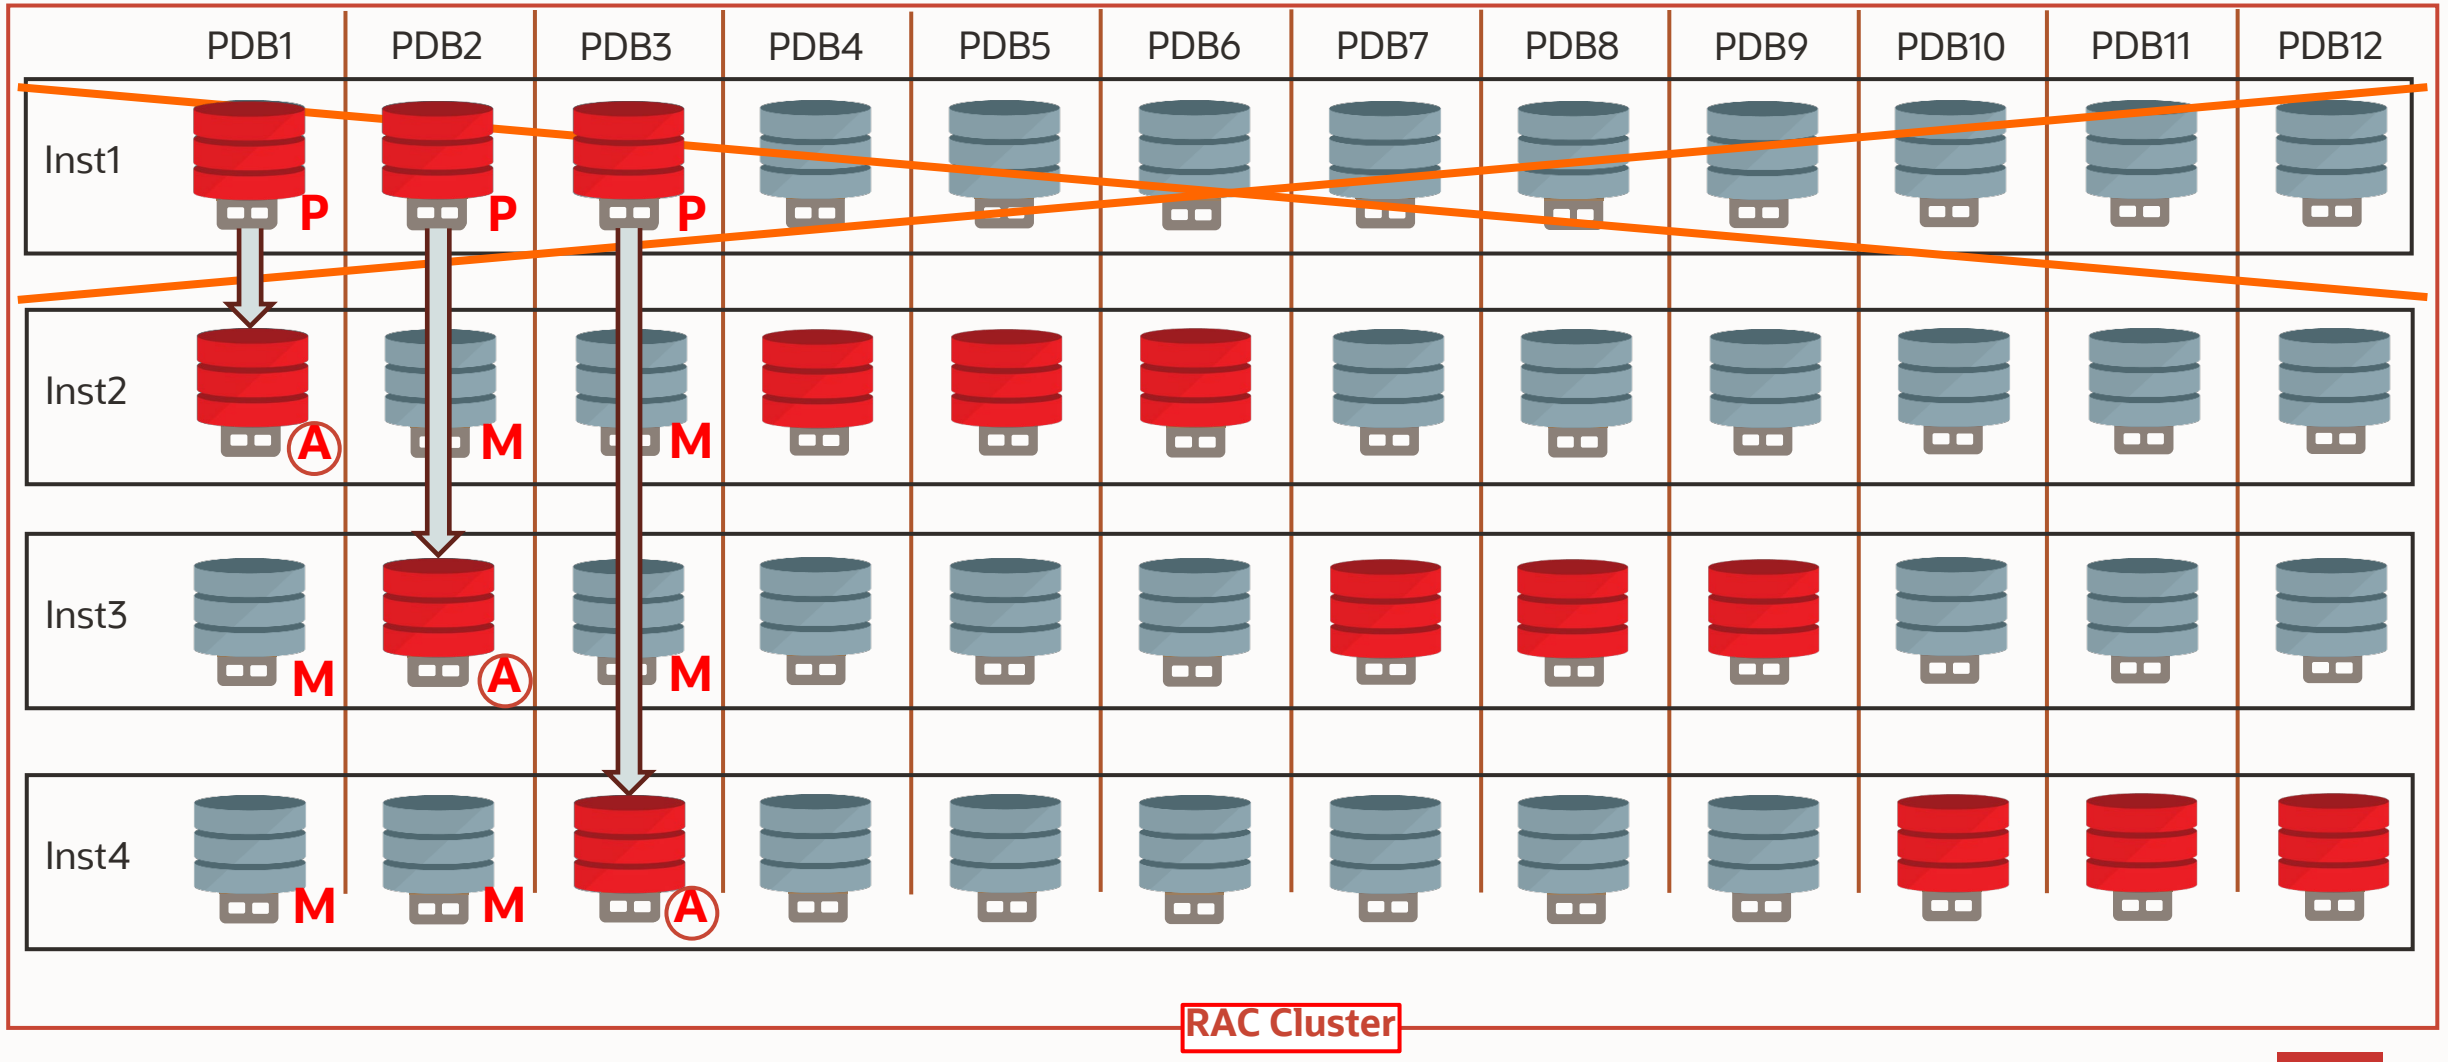

Transparent Application Continuity (TAC)

Application does not see errors during outages

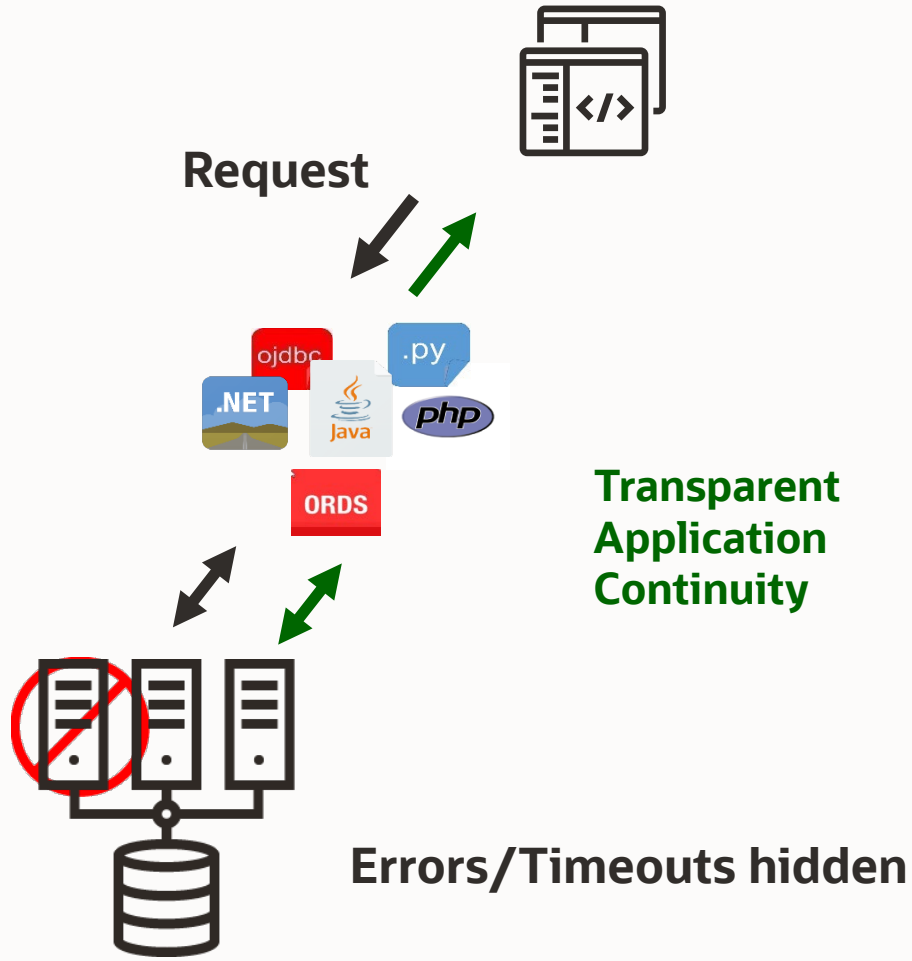

- Uses Application Continuity and Oracle Real Application Clusters
- Transparently tracks and records session information in case there is a failure
- Built inside of the database, so it works without any application changes
- Rebuilds session state and replays in-flight transactions upon unplanned failure
- Planned maintenance can be handled by TAC to drain sessions from one or more nodes
- Adapts as applications change: protected for the future

Planned Maintenance

Planned Maintenance (without the Outages!):

1. Database Service is relocated or stopped
2. Service starts on another RAC instance
3. Sessions connected to the service are drained
4. New sessions connect to Service on another instance
5. Results from Database Request returned to user
6. Maintenance activities can start on first node (rolling)

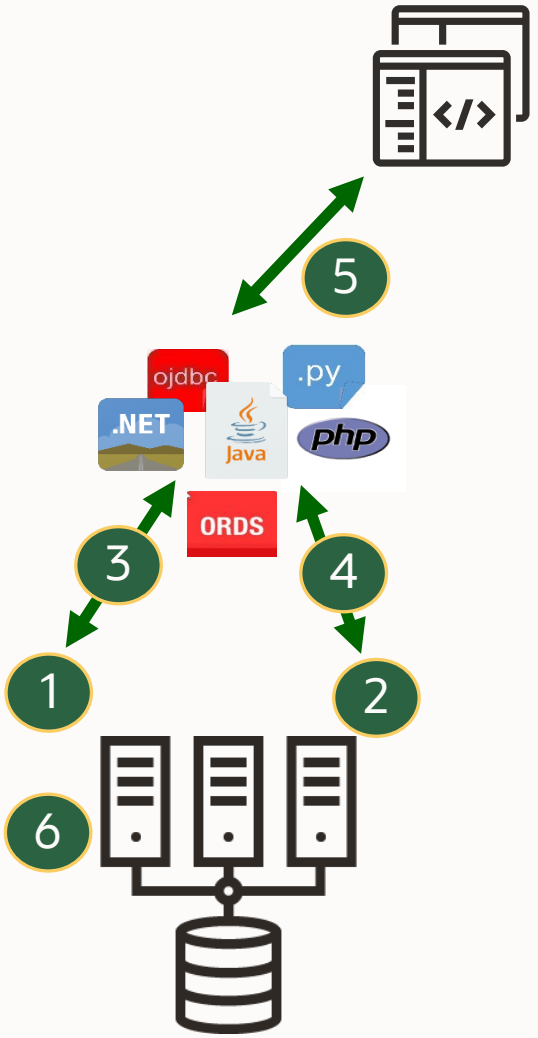

Unplanned Outages, without Impact

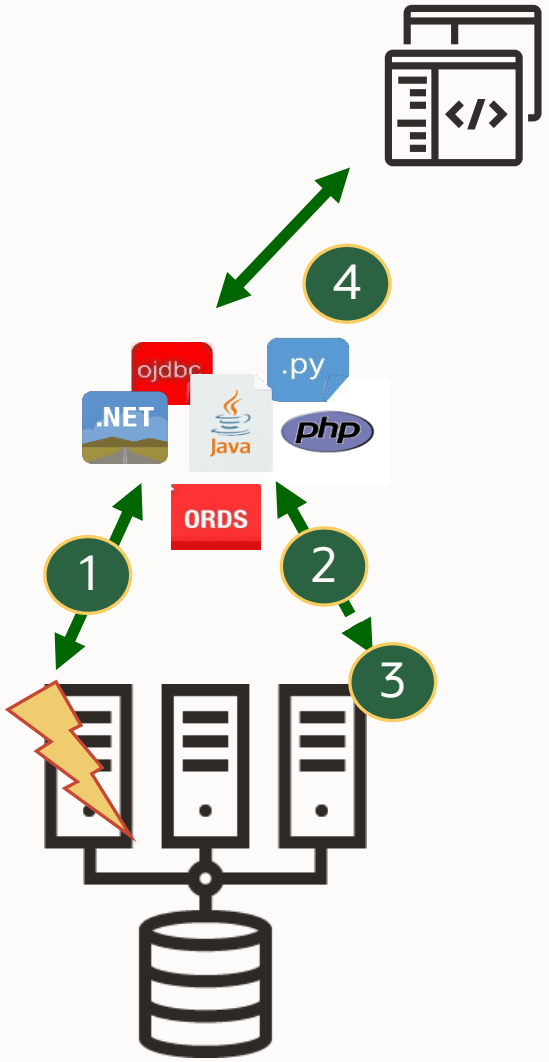

<u>Outage or Interruption at Database</u>:

1. Database Request interrupted by an Outage or timeout
2. Session reconnects to the RAC Cluster and
3. Database Request replays automatically
4. Result from Database Request returned to user

Introducing the “Aurous” MAA Option

Multitenant Architecture-based Disaster Recovery

Flexible PDB Placement using PDB refreshable clones

- Great HA and PDB flexibility. Good data protection & DR. Automation built-in to ADB-S
  - CDBs can host any PDB with varying SLAs
  - Non-Critical, Business Critical and Mission Critical PDBs can reside in the same CDB
  - Business critical PDBs can fail over to another RAC instance
  - Mission critical PDBs can fail over to remote PDB in another CDB
- Some advanced features not available such as RO standby, automatic DR failover for site failure, fast reinstate after role transition
- Capability is very innovating
- RTO=secs for HA and RTO < 10 min for DR (not automatic)
- RPO for DR < 5 min

Picking MAA Reference Architecture

MAA Reference Architecture	RTO for HA / Downtime for Software Updates	RTO for DR / Downtime for DB Upgrades	RPO for DR	On-Premise	Oracle Cloud
Bronze MAA	Hour(s)	Hours	< 15 or 30 mins* Secs with ZDLRA*	Yes	Yes
Silver MAA	Zero or Secs	Hours	< 15 or 30 mins* Secs with ZDLRA*	Yes – Exadata, RAC	Yes- ExaDB, ExaCC, ADB
"Aurous" Option	Zero or Secs	< 10 mins	< 5 mins **	No – refreshable clones can help	Yes – ADB-S only
Gold MAA	Zero or Secs	< 2 mins	Zero or secs Secs with ZDLRA*	Yes	Yes – ExaDB, ExaCC, ADB-D ***
Platinum MAA	Zero or Secs	Zero or secs	Zero or secs	Yes	Yes – ExaDB/ExaCC ***

All RTO and RPO values are service level **objectives** that have been validated. To achieve these values specific MAA config and operational practices are in place. Note all DR operations are fully automatic today. The RTO for DR calculation is based on after the failover is automatically or manually triggered.

* Based on archive backup frequency. For ADB-D, it's < 15 mins. With ZDLRA, seconds. ZDLRA is available for On-premise and ExaCC

** Based on online redo log push

*** Limitations described later but setup is mostly manual

The “Aurous” MAA Strategy

- CDB has mix of primary & DR copy PDBs
- Individual PDBs can transition roles independently
- Standby PDBs are mounted

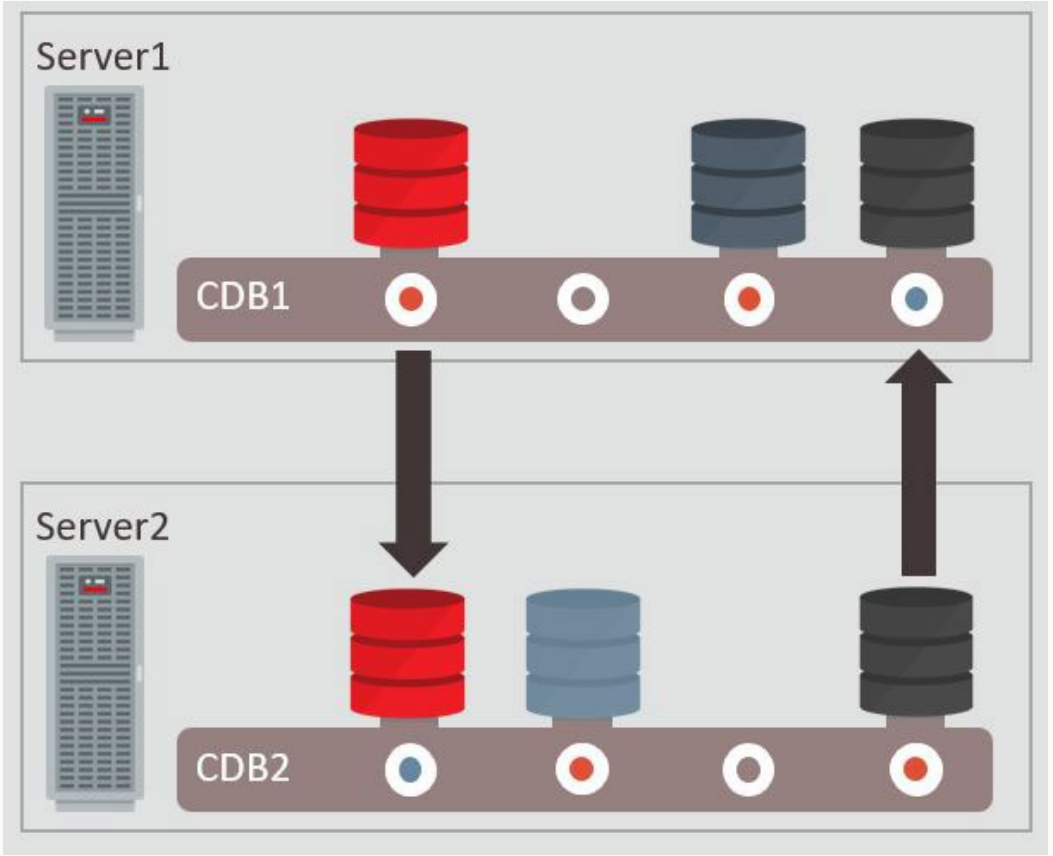

Refreshable PDBs as Replicas

Per-PDB replica with only two CDBs to manage!

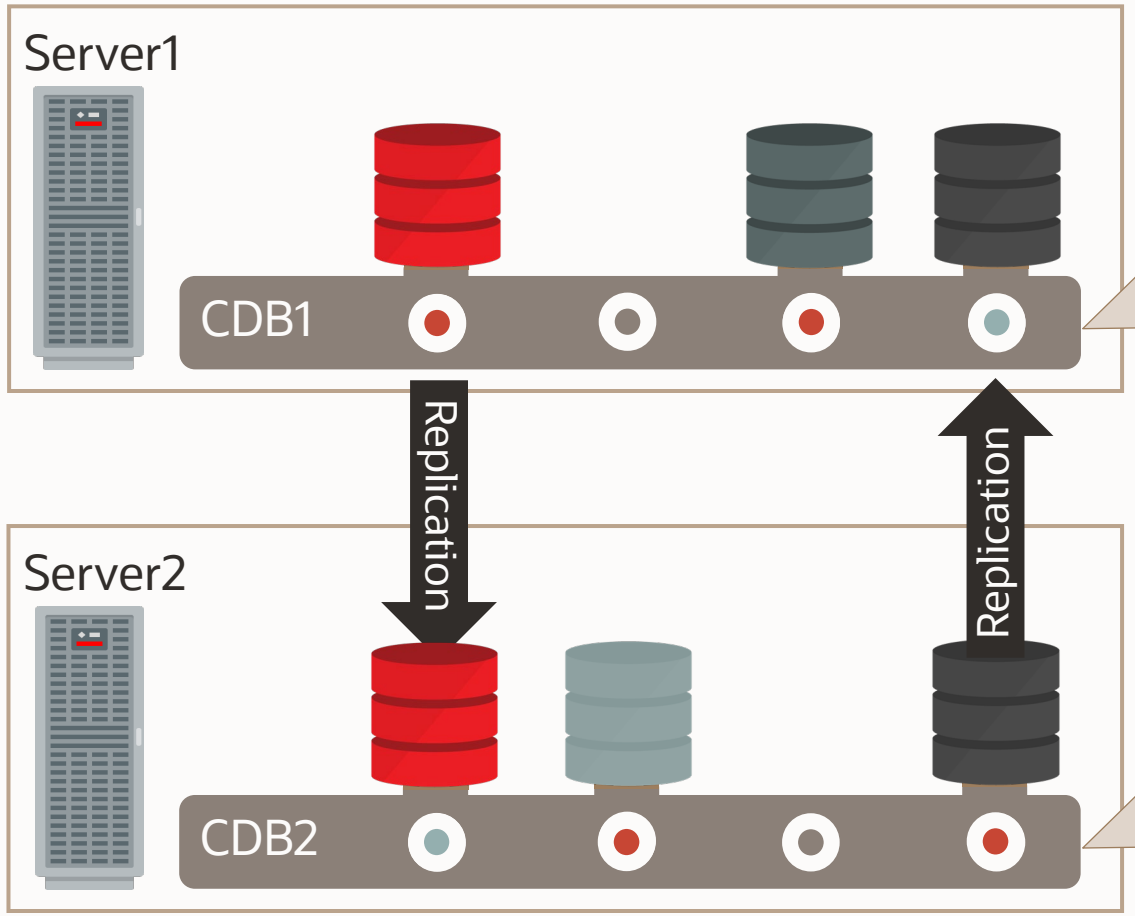


```
1. create pluggable database Red;
4. create pluggable database Brown;
6. create pluggable database Grey
   from Grey@CDB2_Link
   refresh mode every 2 minutes;
```

```
2. create pluggable database Red
   from Red@CDB1_Link
   refresh mode every 2 minutes;
3. create pluggable database Gold;
5. create pluggable database Grey;
```

Refreshable PDB Switchover

Planned switchover

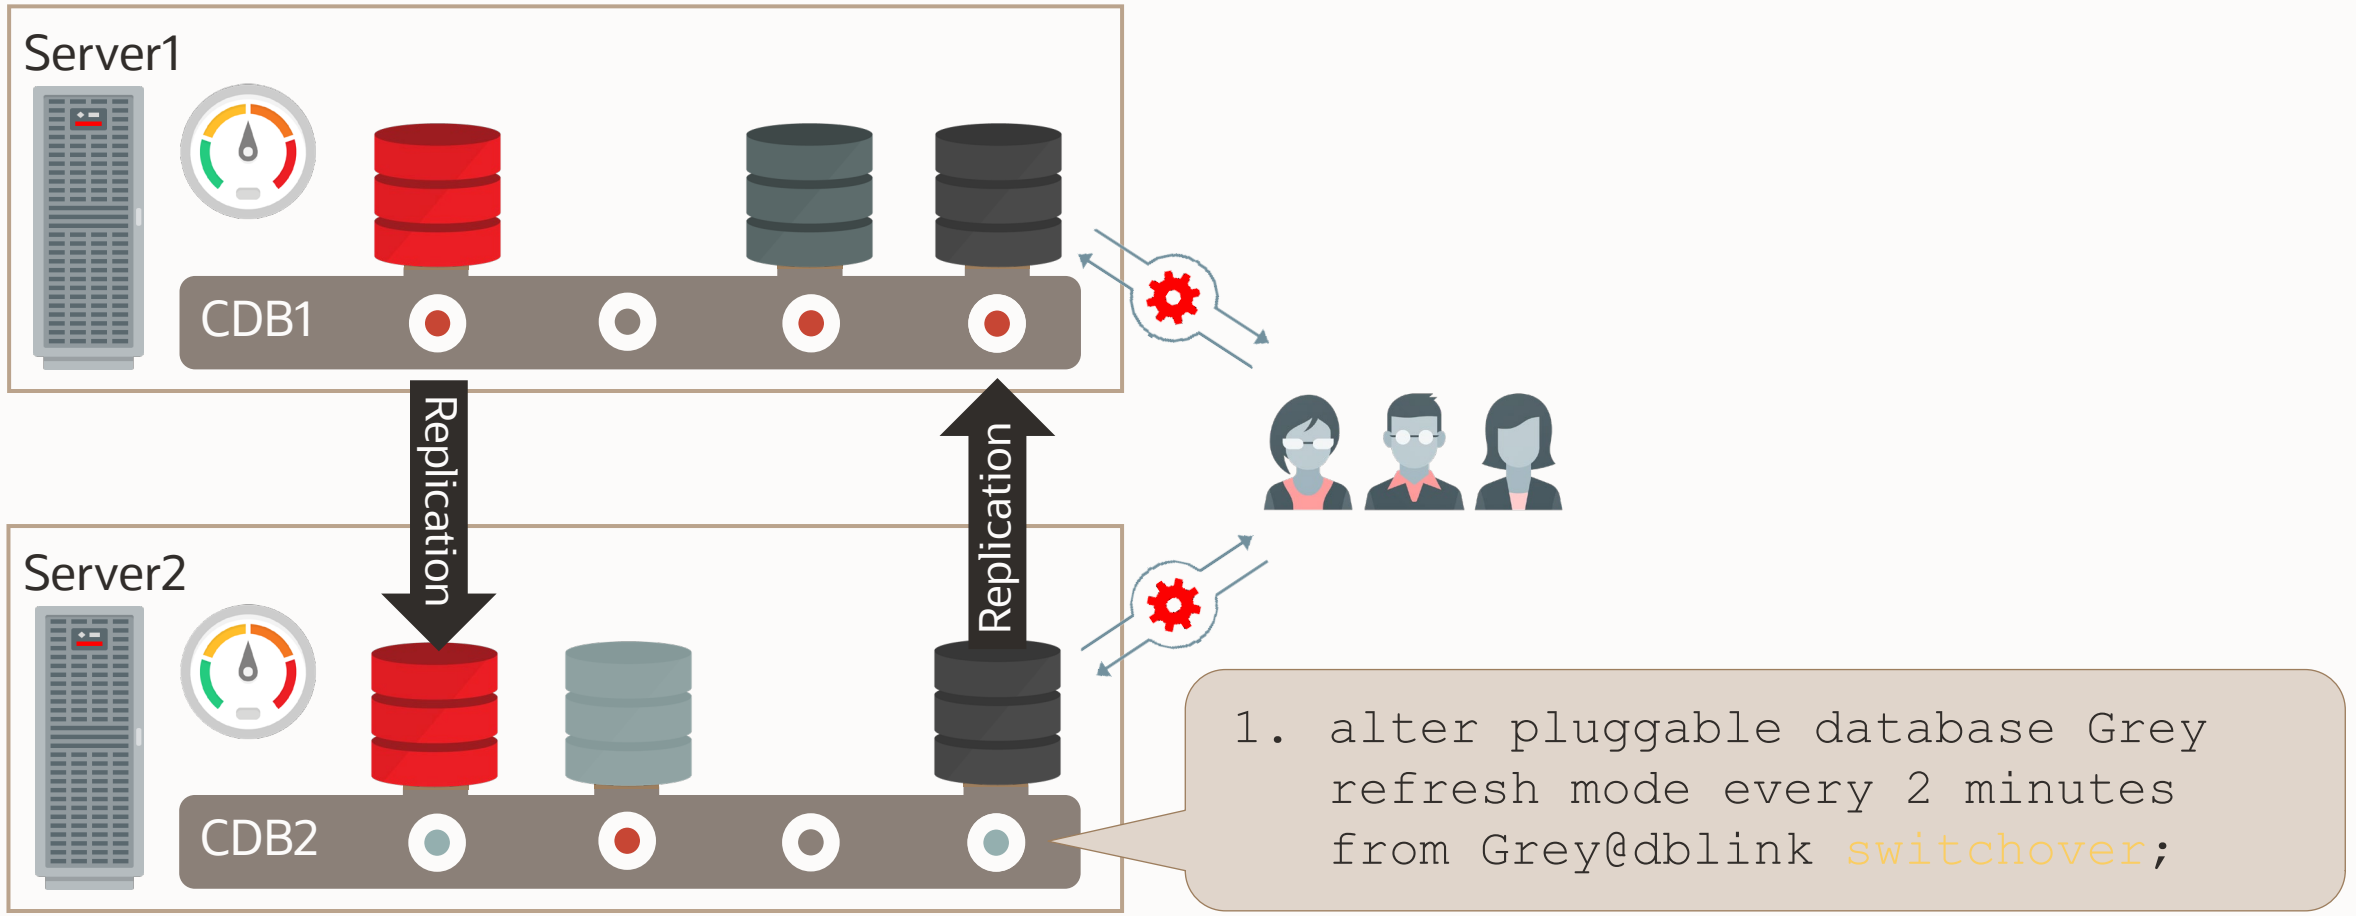

Refreshable PDB Switchover

Unplanned switchover

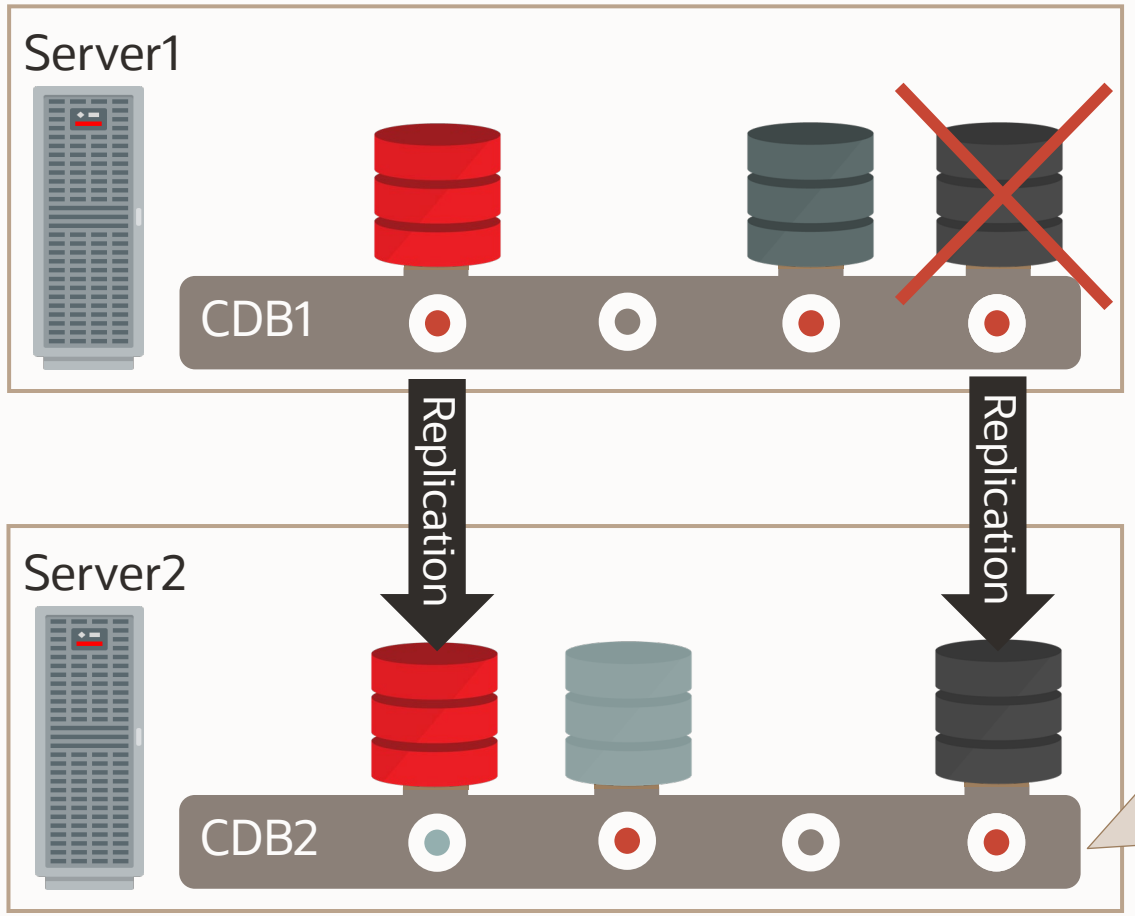


```
1. alter pluggable database Grey
   refresh;
2. alter pluggable database Grey
   refresh mode none;
3. alter pluggable database Grey
   open read write;
```

Recovery Time Objectives (RTOs)

How long does it take to recover?

Determining factors:

- Rate of redo generation at source
- Frequency of refresh of replica

Worst case: Failure of source just prior to next refresh

Best practices:

- Relatively high frequency of refresh (minutes rather than hours)
- Thorough testing with realistic transaction volumes

Multitenant “Aurous” MAA Strategy

Unplanned Outages	Key Features for Solution	RTO	RPO
Recoverable node or instance failure	Real Application Cluster (RAC) Application Continuity (AC/TAC)	Secs	Zero
Disasters: corruptions and site failures	Many Refreshable PDB Switchover operations	Not automatic (target < 15 mins after manual ops)	Since last refresh (min 15 mins)
PDB unrecoverable failure or “sick” PDB	Refreshable PDB Switchover		

Planned Maintenance	Solution	RTO
Software and hardware updates	RAC, AC or TAC	Zero
Major database upgrade (for all PDBs)	CDB offline upgrade	60+ Minutes
Migration to remote CDB	PDB Relocate	Mins
Migration to remote CDB (logical migration)	Data Pump and GoldenGate or ZDM	Potentially Zero
Migration plus upgrade (single PDB)	PDB Relocate + Upgrade	Mins

Using PDB Relocation to Move a Single PDB to Another CDB Without Upgrade (Doc ID 2771737.1)
Using PDB Relocation to Upgrade an Individual PDB Document (Doc ID 2771716.1)
ZDM stands for Zero Downtime Migration. Refer to www.oracle.com/goto/zdm

GOLD

Mission Critical

Silver +

- Active Data Guard
  - Comprehensive Data Protection

MAA Architecture:

- At least one standby required across AD or region.
- Primary in one data center(or AD) replicated to a Standby in another data center
- Active Data Guard Fast-Start Failover (FSFO)
- Local backups on both primary and standby

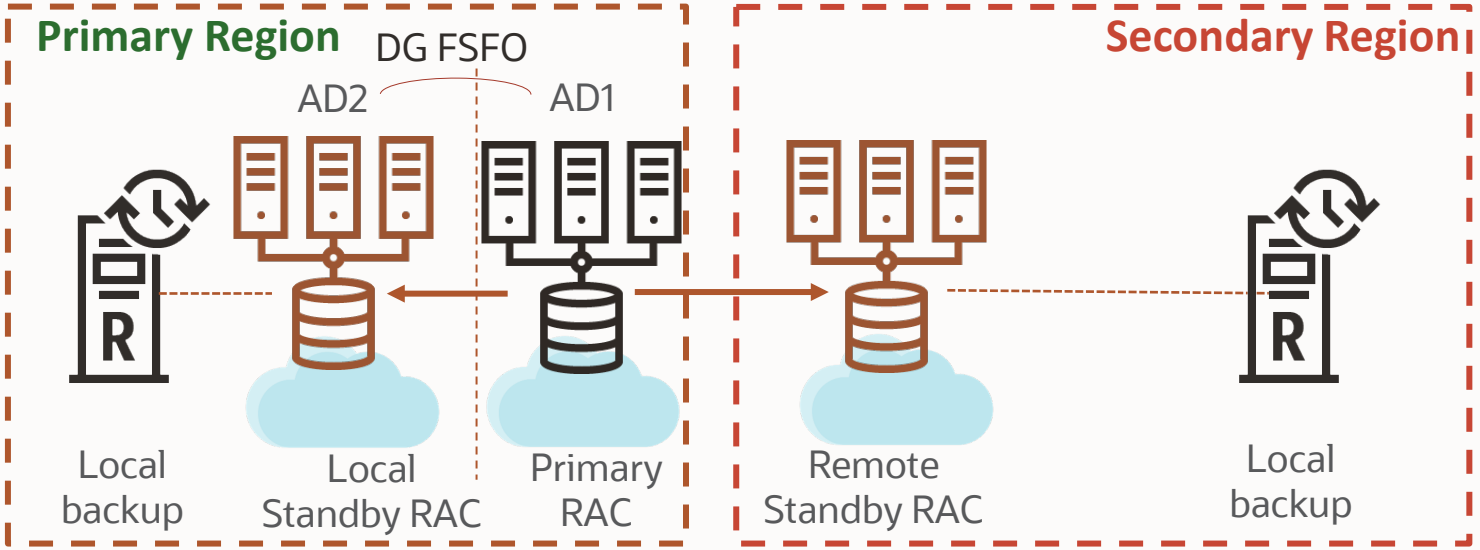

Outage Matrix

Unplanned Outage	RTO/RPO Service Level Objectives (f1)
Recoverable node or instance failure	Single digit seconds (f2)
Disasters: corruptions and site failures	**Seconds to 2 minutes. RPO zero or seconds**
Planned Maintenance	
Software/Hardware updates	Zero (f2)
Major database upgrade	**Less than 30 seconds**

f1: RPO=0 unless explicitly specified
f2: To achieve zero downtime or lowest impact, apply application checklist best practices; Batch jobs should be deferred outside planned maintenance window.

Storage Remote Mirroring Architecture

Generic - Must Transmit Writes to All Files

.... INCLUDING CORRUPTED BLOCKS OR BAD DATA

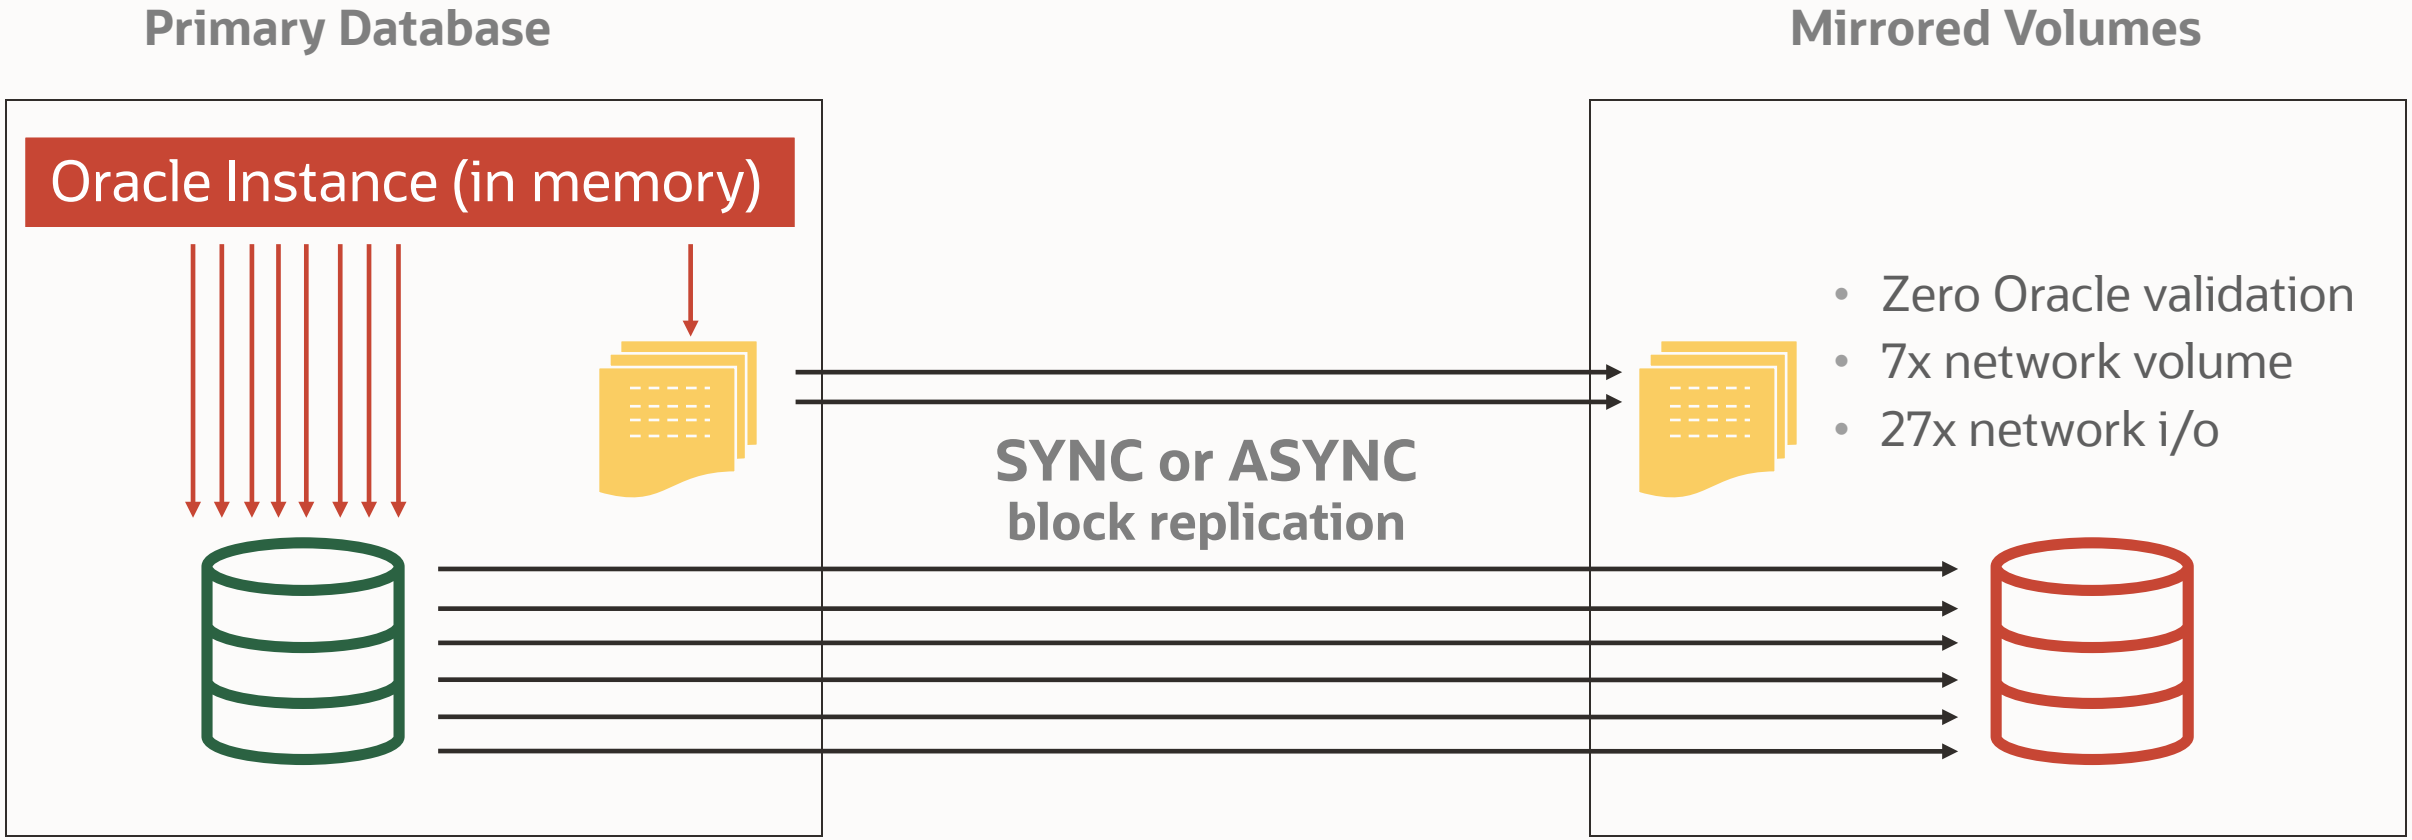

Oracle Data Guard (DG)

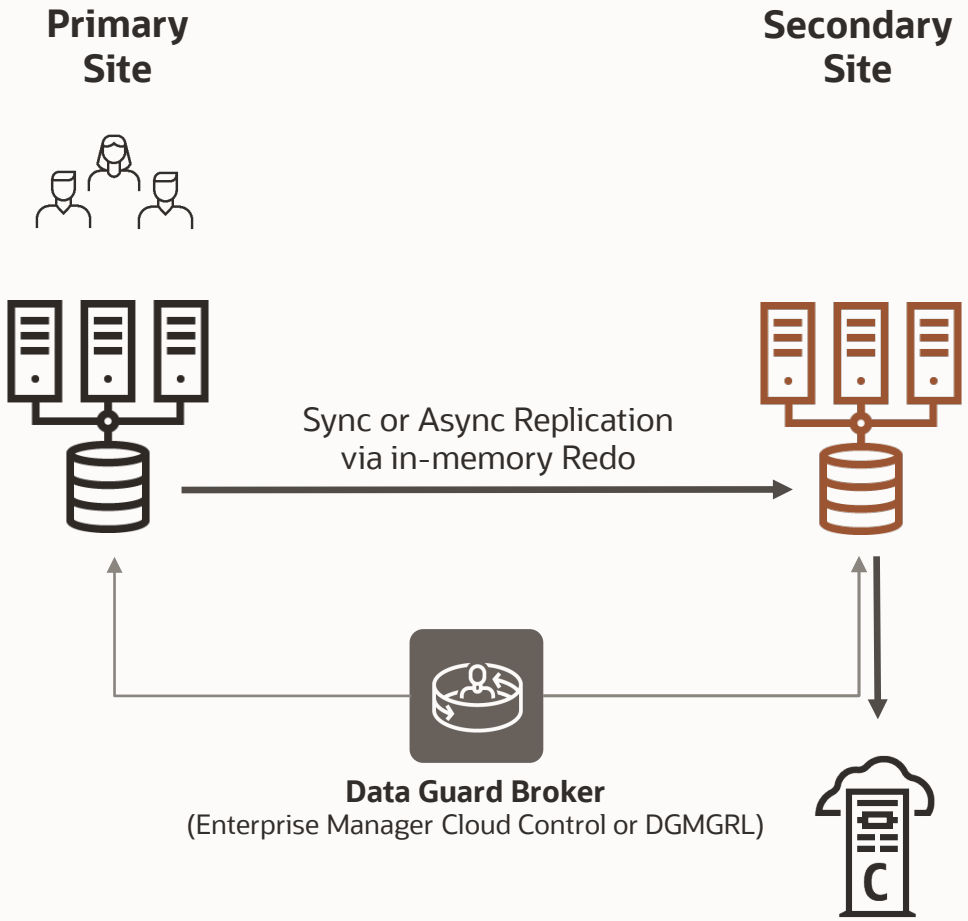

- **Basic DR (included with DB EE)**
  - License primary and secondary sites
- **Active-passive**
  - Standby is used only for failovers
- **Automatic failover to Standby site**
- **Zero / near-zero data loss**
- **Continuous data validation**
- **Simple migrations and upgrades**

https://www.oracle.com/database/technologies/high-availability/dataguard-activedataguard-demos.html

Oracle *Active* Data Guard (ADG)

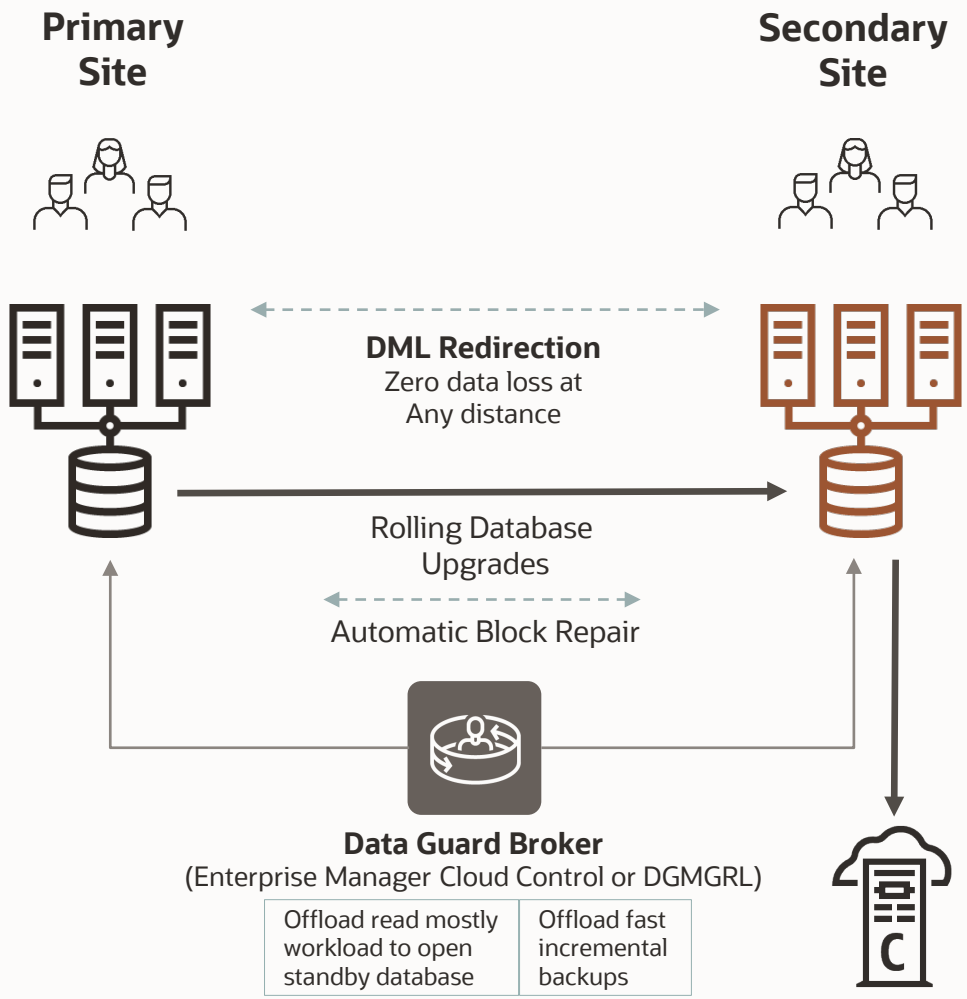

- **Advanced Disaster Recovery**
- **Active-active***
  - Queries, reports, backups
  - Occasional updates (19c)
  - Assurance of knowing system is operational
- **Automatic block repair**
- **Application Continuity**
- **Zero data loss across any distance**
- **Many other features**

https://www.oracle.com/database/technologies/high-availability/dataguard-activedataguard-demos.html

Gold MAA Strategy with Multitenant

- Active Data Guard of the CDB with automatic failover
  - Automatic smart failover when primary is inaccessible
  - Auto-block repair, offload read and reporting and backups
- PDB failover and relocate operations

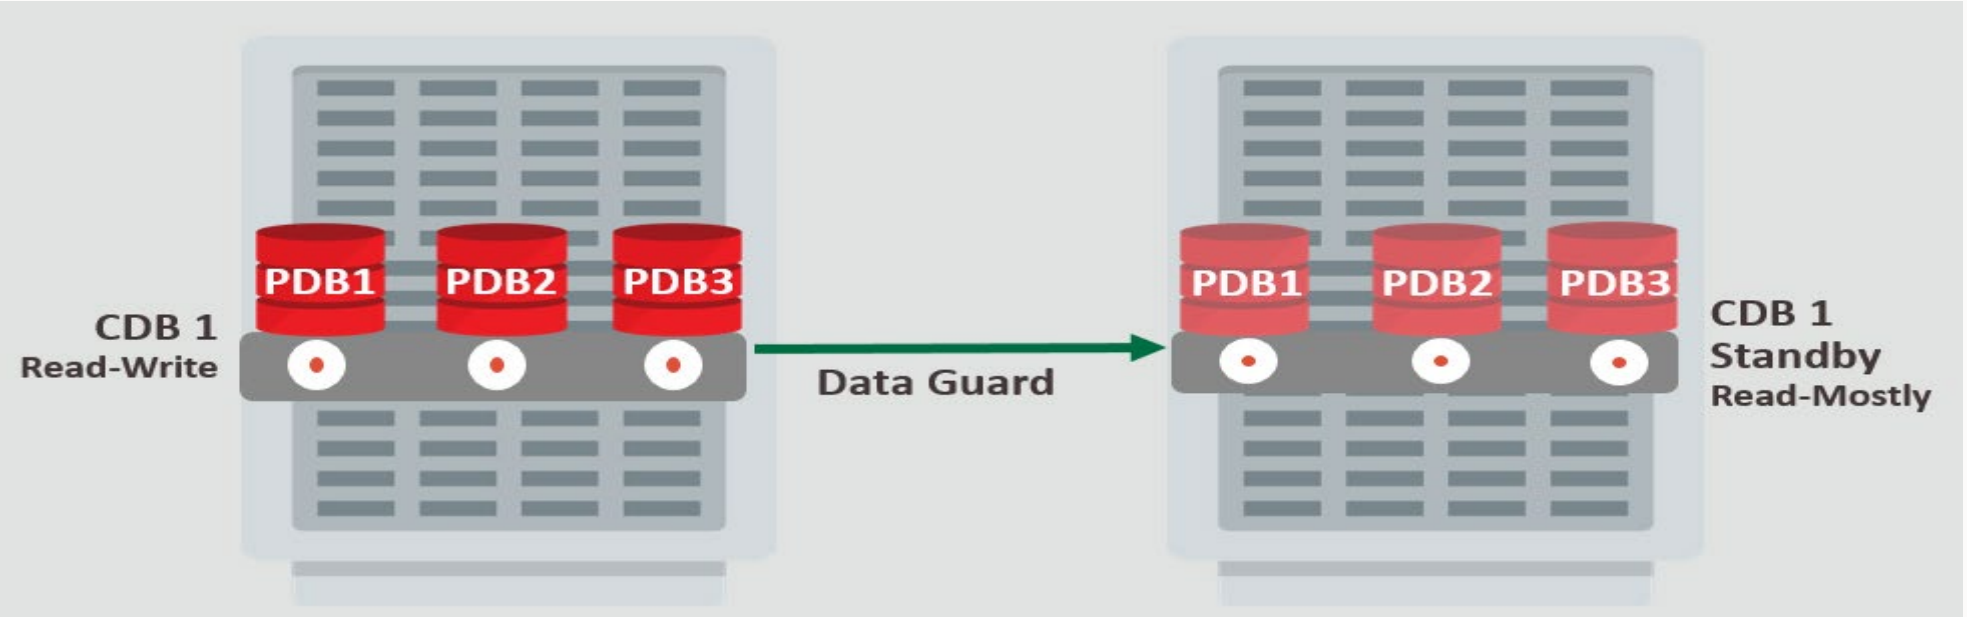

Active Data Guard Far Sync

Zero Data Loss Protection at Any Distance

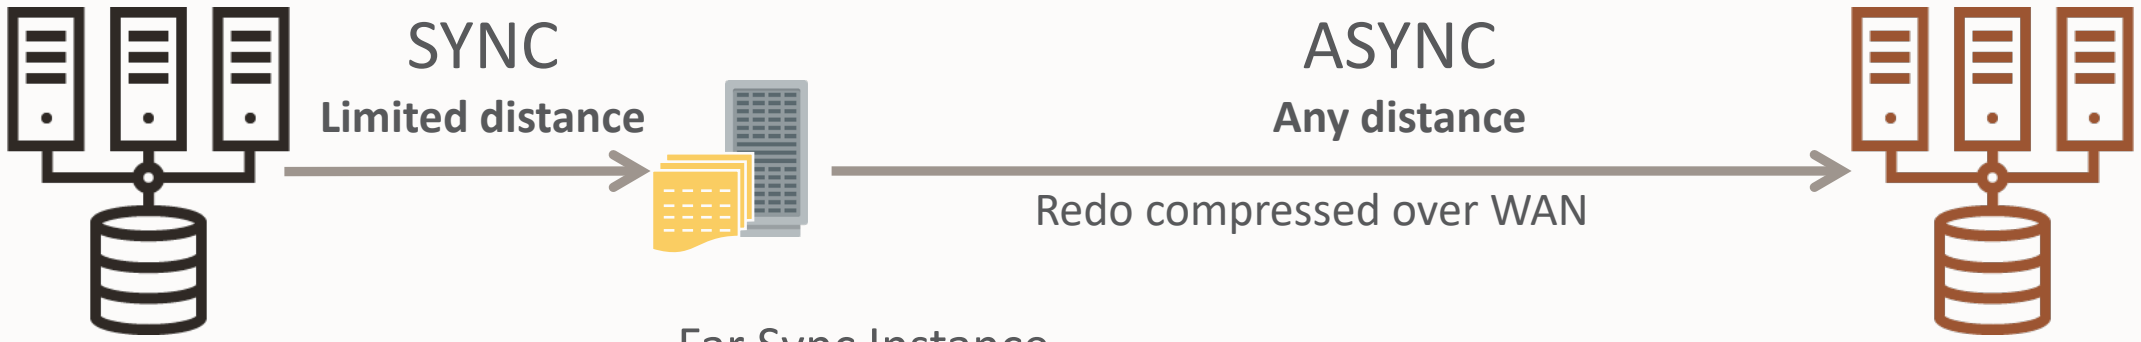

Primary Database
- Production copy

Far Sync Instance
- Oracle control file and log files
- No database files
- No media recovery
- Offload transport compression and/or encryption

Active Standby Database
- Zero data loss failover target
- Database open read-only
- Continuous Oracle validation
- Manual or automatic failover

Unplanned Outages, without Impact expanded to the Standby

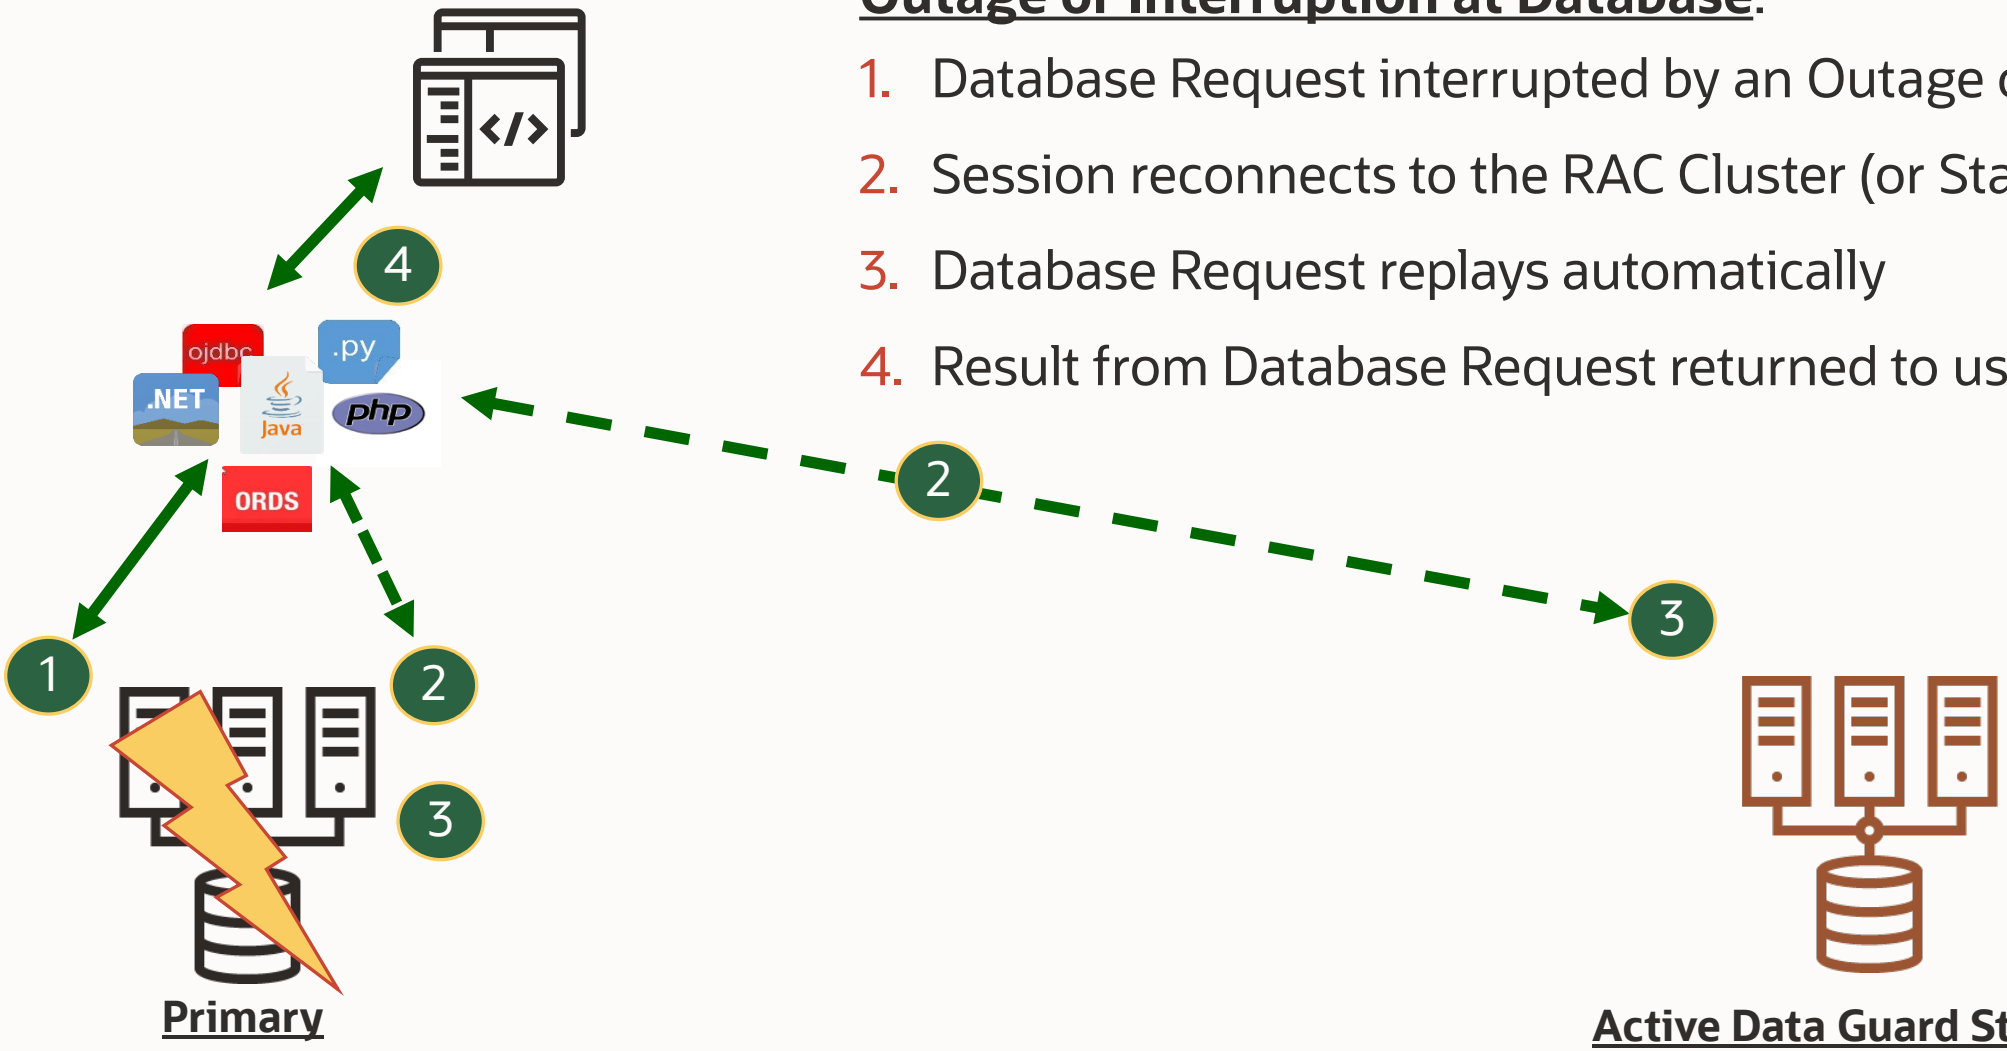

Outage or Interruption at Database:

1. Database Request interrupted by an Outage or timeout
2. Session reconnects to the RAC Cluster (or Standby) and
3. Database Request replays automatically
4. Result from Database Request returned to user

Extend Footprint of ADG Applications

Support for DML Re-direction

- DML Re-direction is automatically performed from an Active Data Guard standby to the primary (ACID uncompromised)
- New parameter `ADG_REDIRECT_DML` controls DML Redirection
- New `ADG_REDIRECT_DML and ADG_REDIRECT_PLSQL`

- "Read-Mostly, Occasional Updates" applications supported for Oracle Database 19c

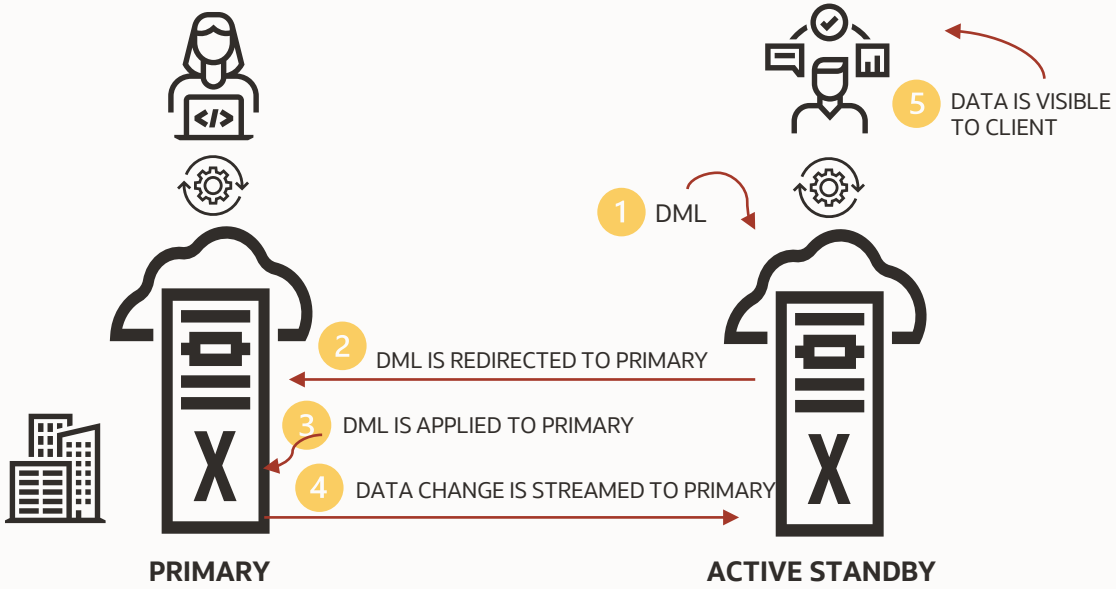

Multi-Instance Redo Apply Performance

Lower Latency Active Data Guard Standby Databases

- Utilizes all RAC nodes on the Standby database to parallelize recovery
- OLTP workloads on Exadata show great scalability

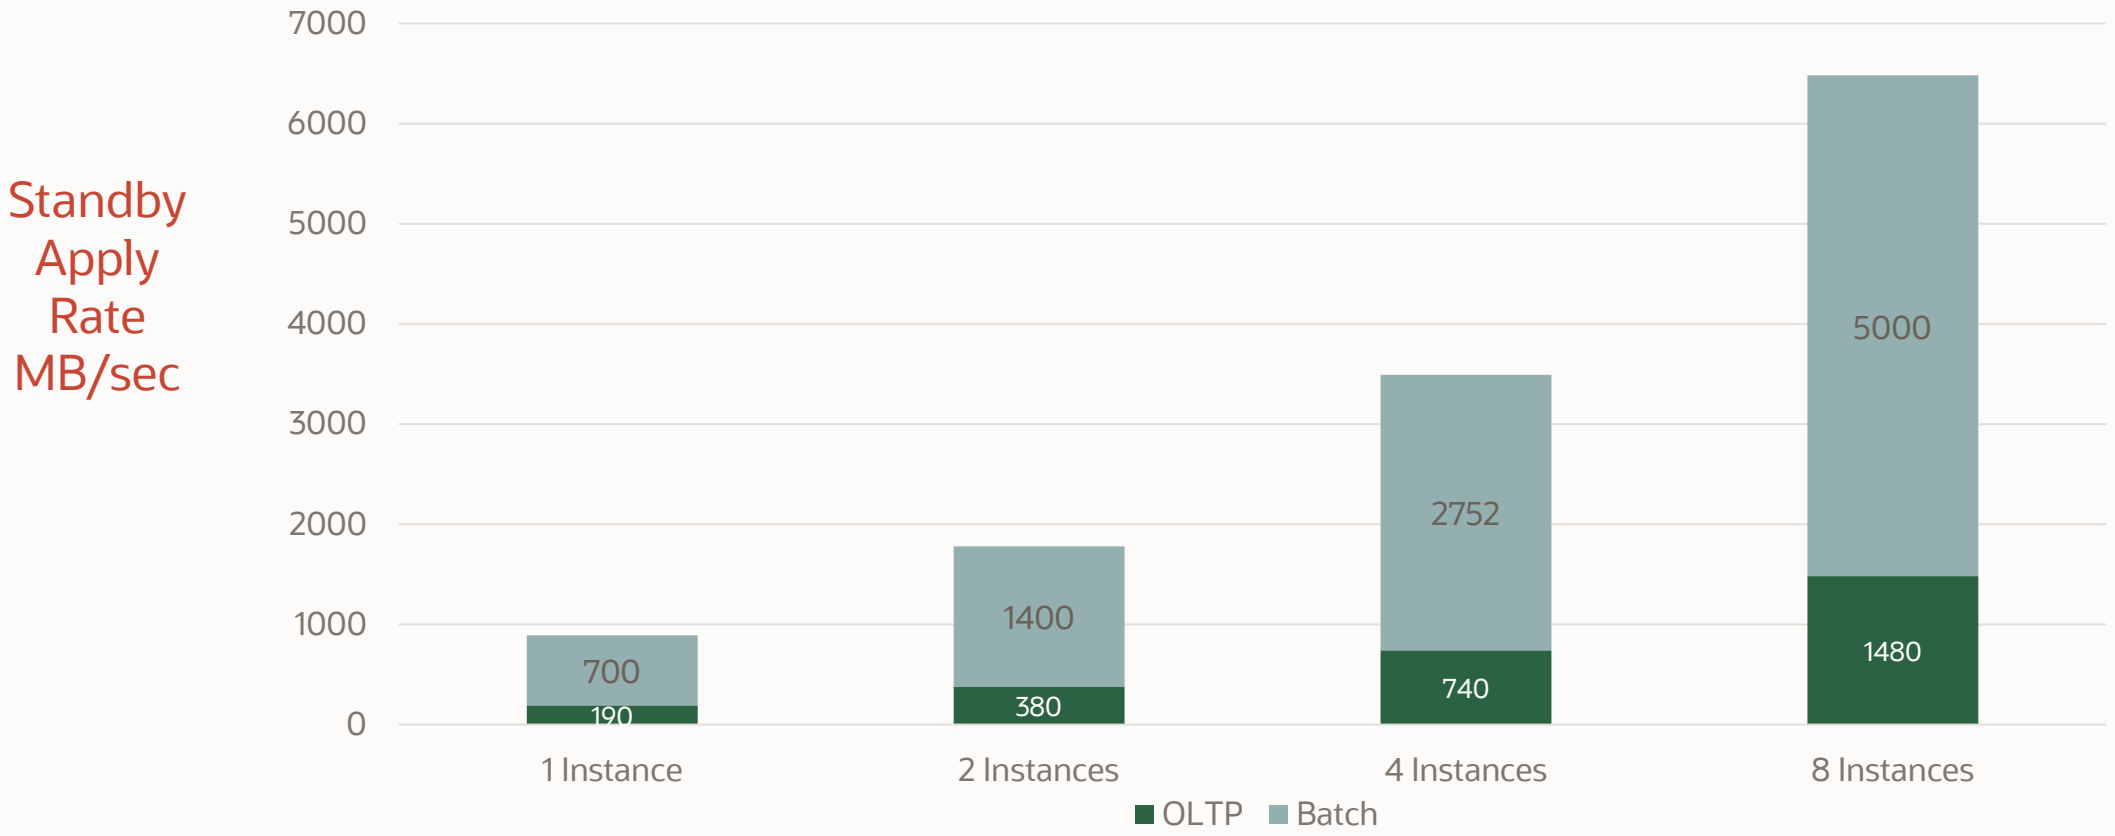

Database Rolling Upgrade to 19c

Database Rolling Upgrade with DBMS_ROLLING

- Pre-checks and early problem detection
- Fault tolerant, resumable and rollback capabilities
- Three Role Transition Steps: Start, Switchover, Finish
- Potential Maintenance Window: Hours
- Potential Database and Application Downtime: Seconds

Automated Database Upgrades using Oracle Active Data Guard and DBMS_ROLLING

Multitenant Replication Capabilities

Manage many as one with Data Guard

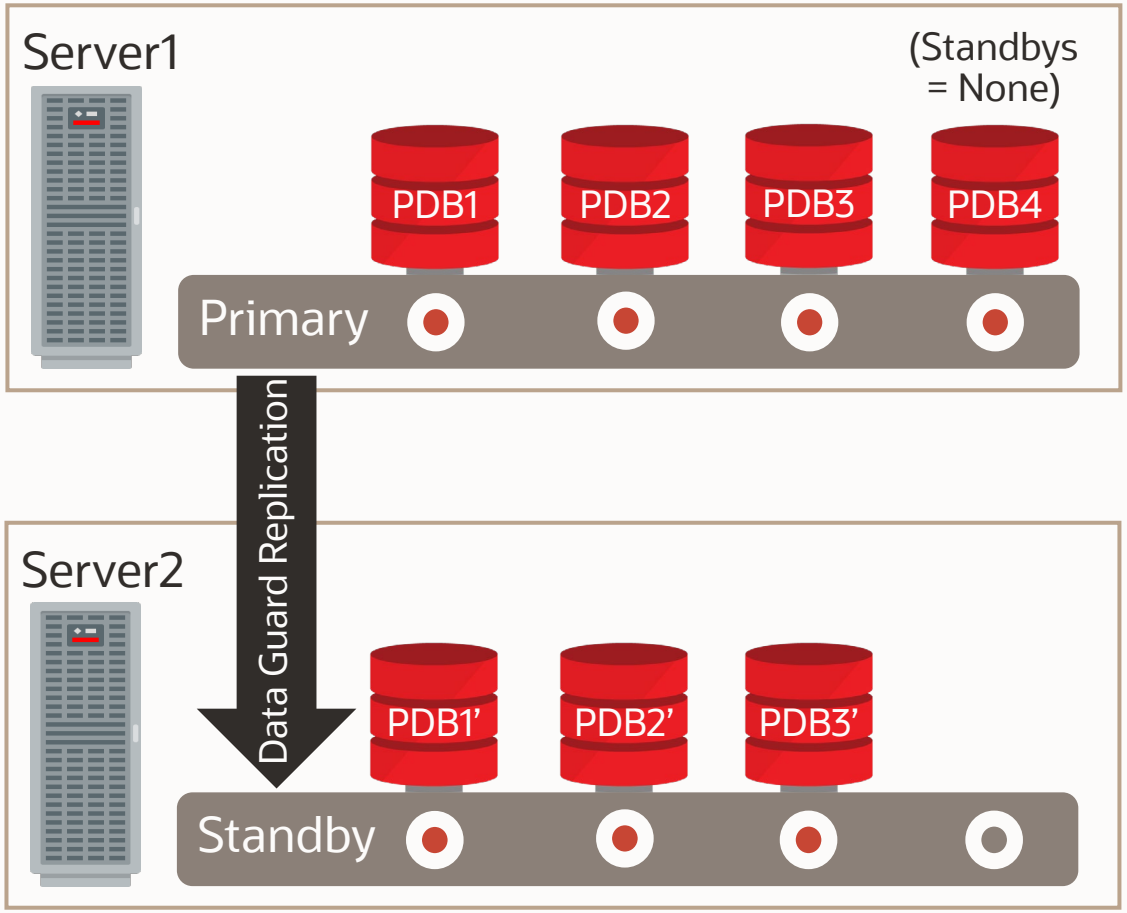

- When creating new local PDBs, the new PDBs are automatically replicated and represented on the standby.
- What happens when cloning a remote PDB or plugging in new PDBs?

What if only one PDB is down or needs to be relocated?

- PDB failover is possible by migrating to a new CDB.
  - This should be extremely rare case. With proper planning, all database (PDBs) in the same CDB should have same HA and DR requirements, planned maintenance schedules, DR testing schedules and data center requirements
  - A single failed PDB that is not restartable is extremely rare
- PDB relocate can be used move a PDB to another CDB (same release or different release)

PDB Failover: Normal Runtime

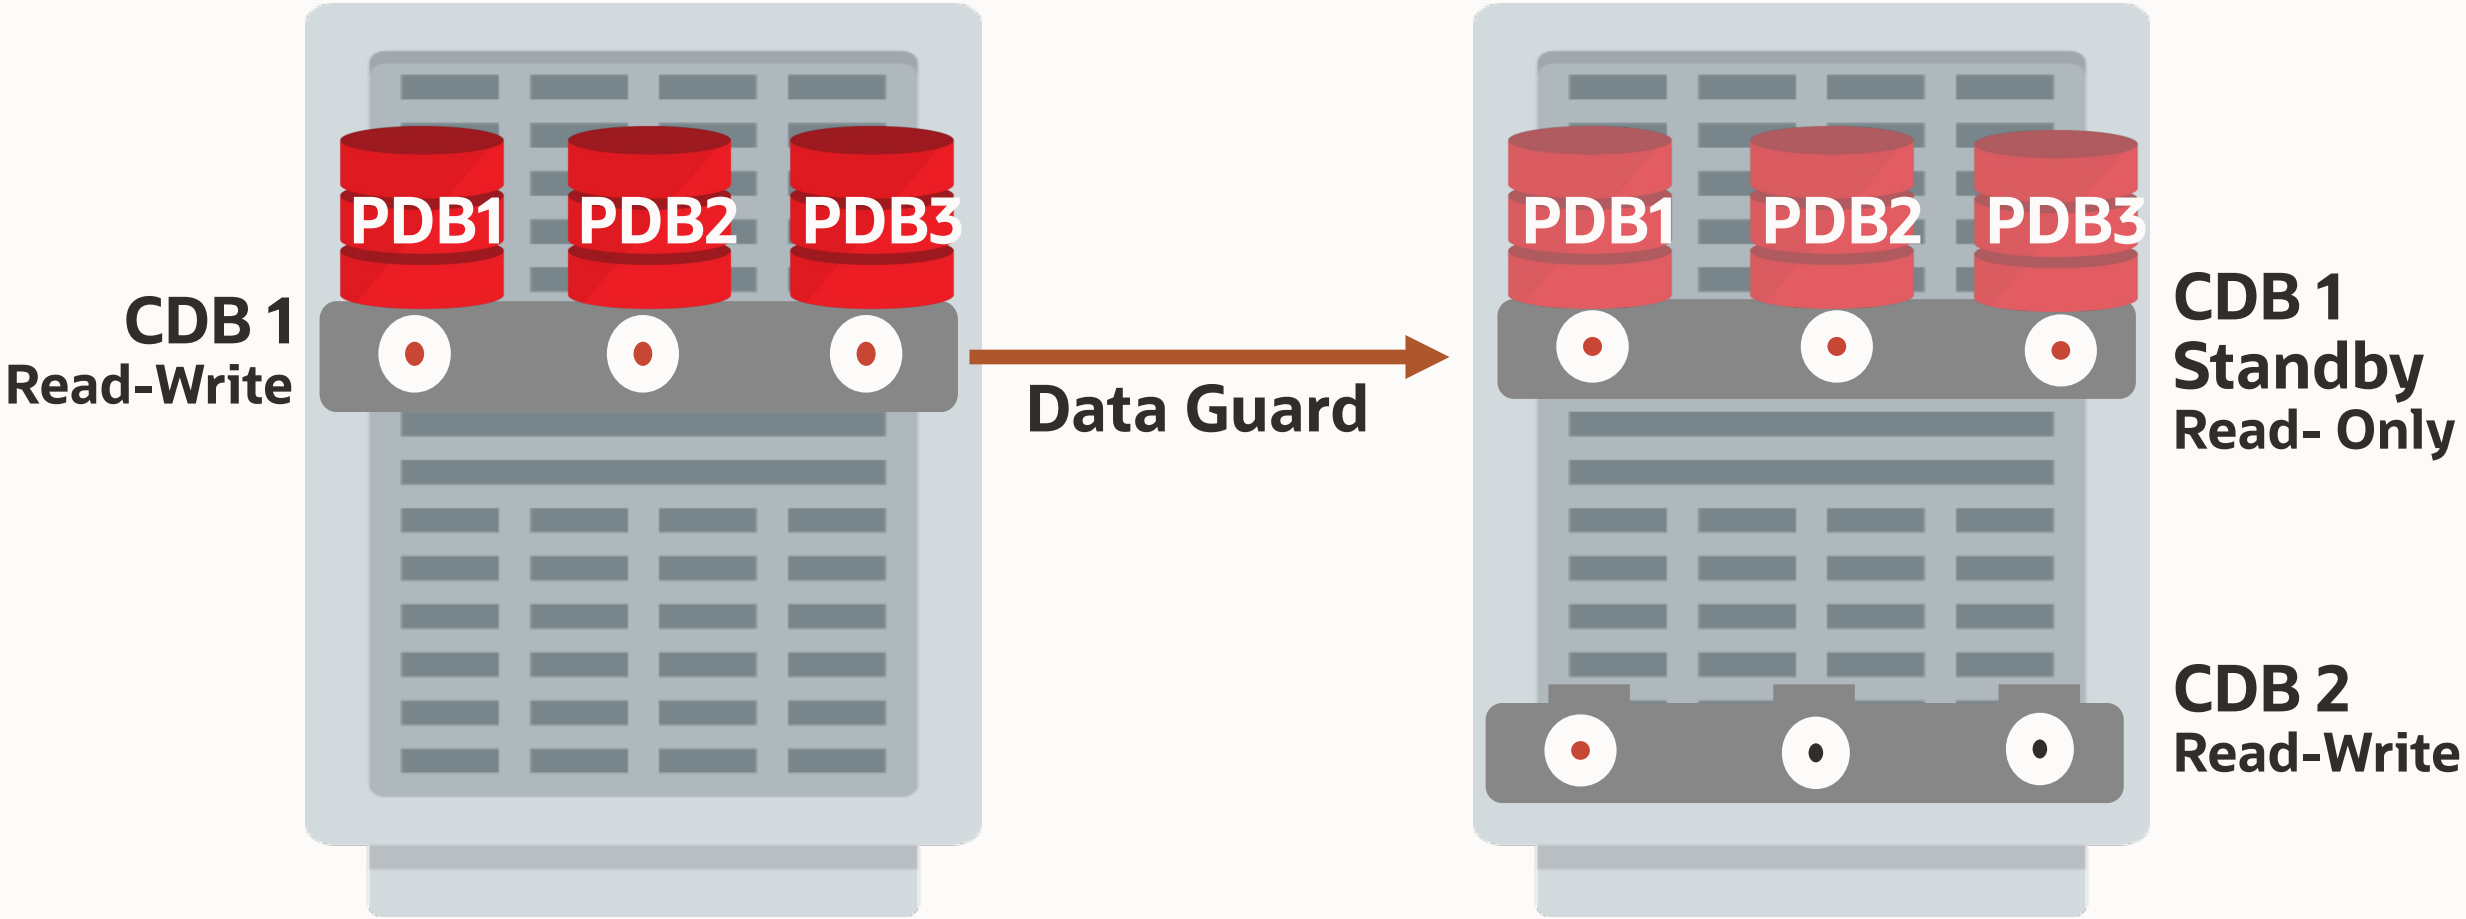

PDB Failover after PDB 2 Outage

Migrate PDB2 from CDB1 standby to empty CDB2 and failover application connections

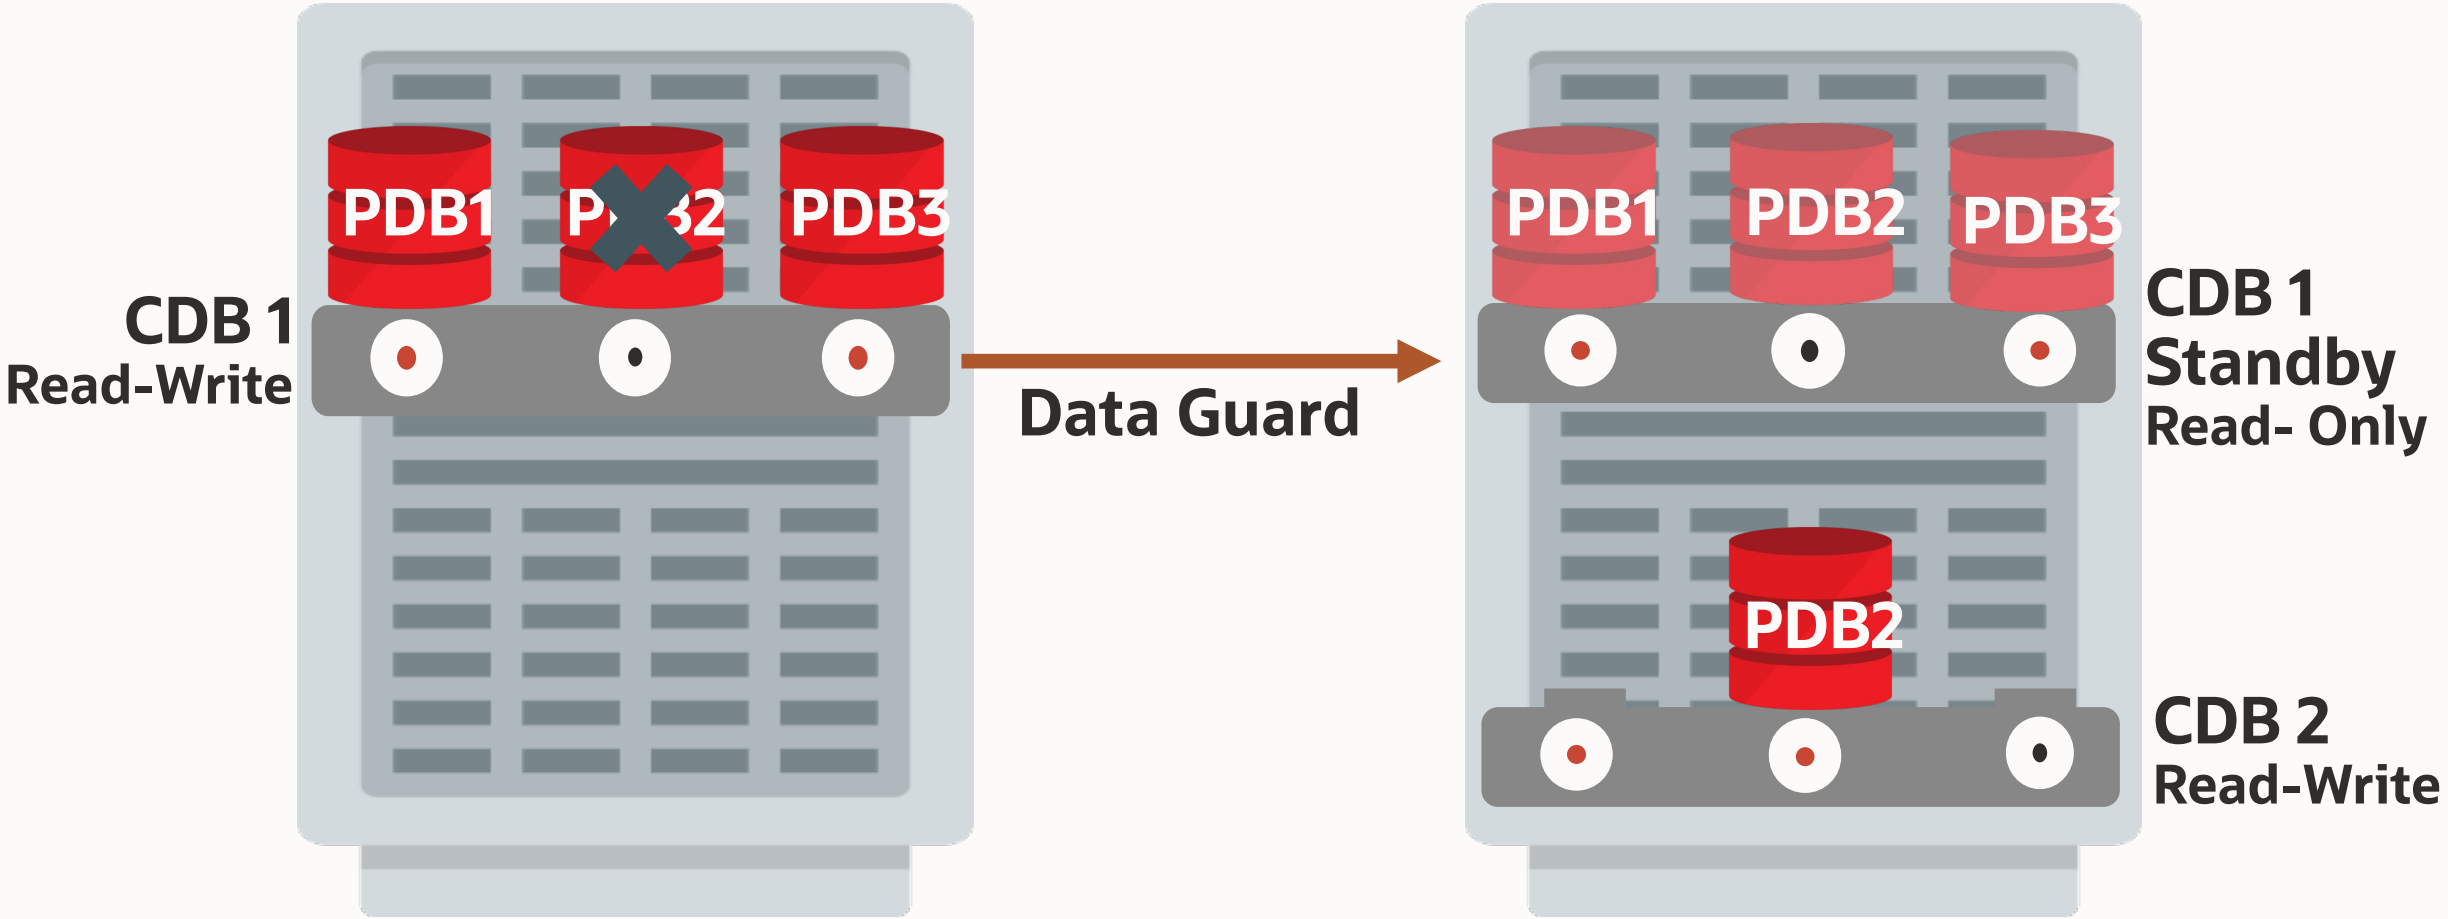

Online PDB Relocation

Online PDB Relocation

- Relocate with minimal downtime

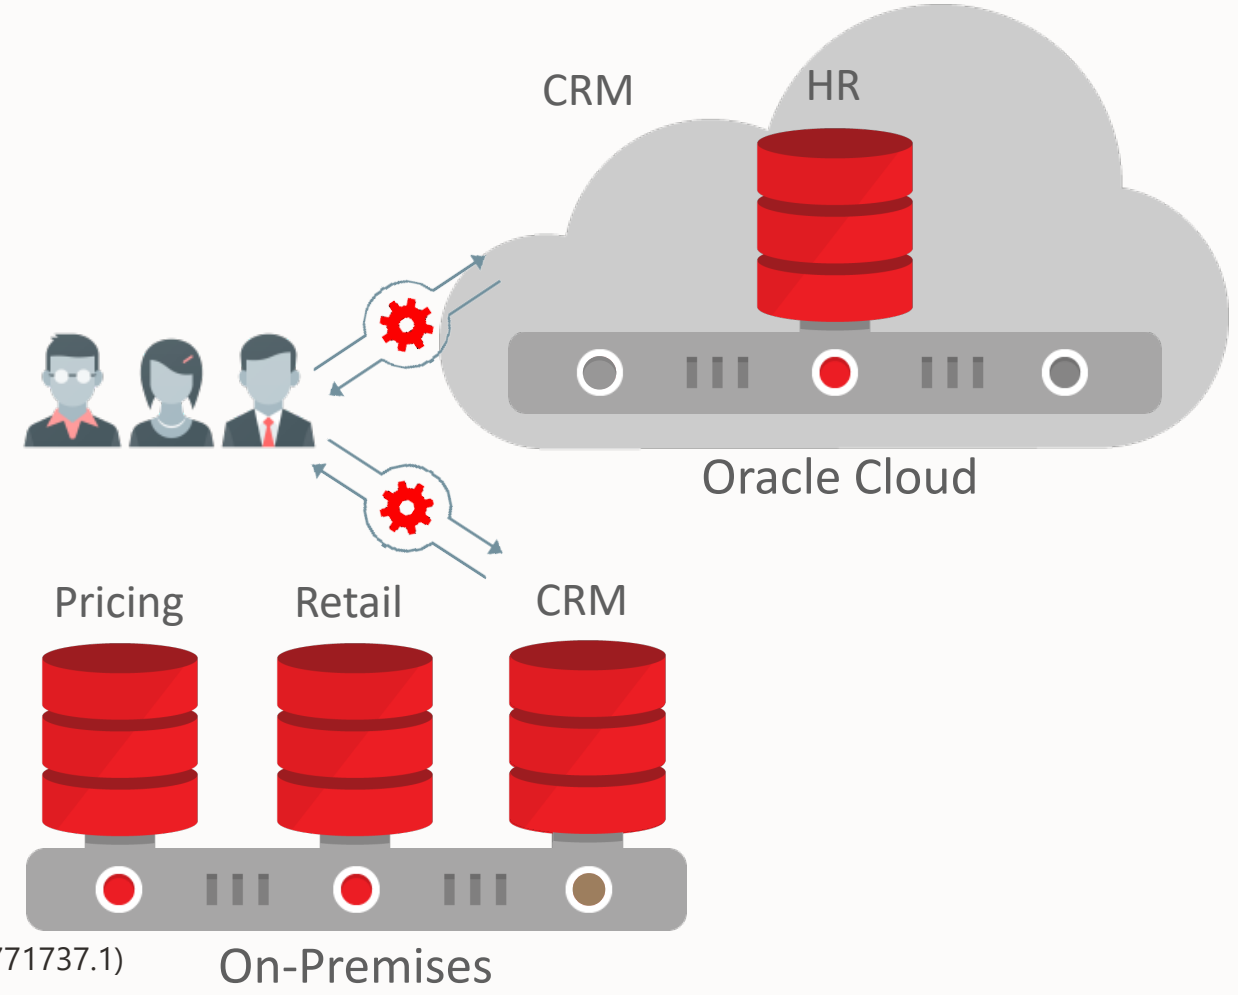

Using PDB Relocation to Move a Single PDB to Another CDB Without Upgrade (Doc ID 2771737.1)
Using PDB Relocation to Upgrade an Individual PDB Document 2771716.1.

Multitenant “Gold” MAA Strategy

Unplanned Outages	Key Features for Solution	RTO	RPO
Recoverable node or instance failure	Real Application Cluster (RAC) Application Continuity (AC/TAC)	Secs	Zero
Disasters: corruptions and site failures	Active Data Guard Fast-Start Failover	< 2 mins	Zero or Secs
PDB unrecoverable failure or “sick” PDB or “PDB DR test”	PDB Failover (unplug/plug) Another target CDB on the same cluster required (MOS 2088201.1)	< 2 mins	Zero or Secs

Planned Maintenance	Solution	RTO
Software and hardware updates	RAC, AC or TAC	Zero
Major database upgrade	Active Data Guard DBMS_ROLLING	Secs
PDB “DR Test” / PDB Switchover	Similar to MOS 2088201.1	Mins
Migration to remote CDB	PDB Relocate	Mins
Migration to remote CDB (logical migration)	Data Pump and GoldenGate or ZDM	Potentially Zero
Migration plus upgrade	PDB Relocate + Upgrade	Mins

Using PDB Relocation to Move a Single PDB to Another CDB Without Upgrade (Doc ID 2771737.1)
Using PDB Relocation to Upgrade an Individual PDB Document (Doc ID 2771716.1)
ZDM stands for Zero Downtime Migration. Refer to www.oracle.com/goto/zdm

PLATINUM

Extreme Critical

Gold +

- GoldenGate Active/Active Replication
- Edition-based Redefinition (Alternative)

MAA Architecture:

- Each GoldenGate "primary" replica protected by Exadata, RAC and Active Data Guard
- Primary in one data center (or AD) replicated to another Primary in remote data center (or AD)
- Oracle GG & Edition-based Redefinition for zero downtime application upgrade
- Local backups on both sites
- Achieve zero downtime through custom failover to GG replica

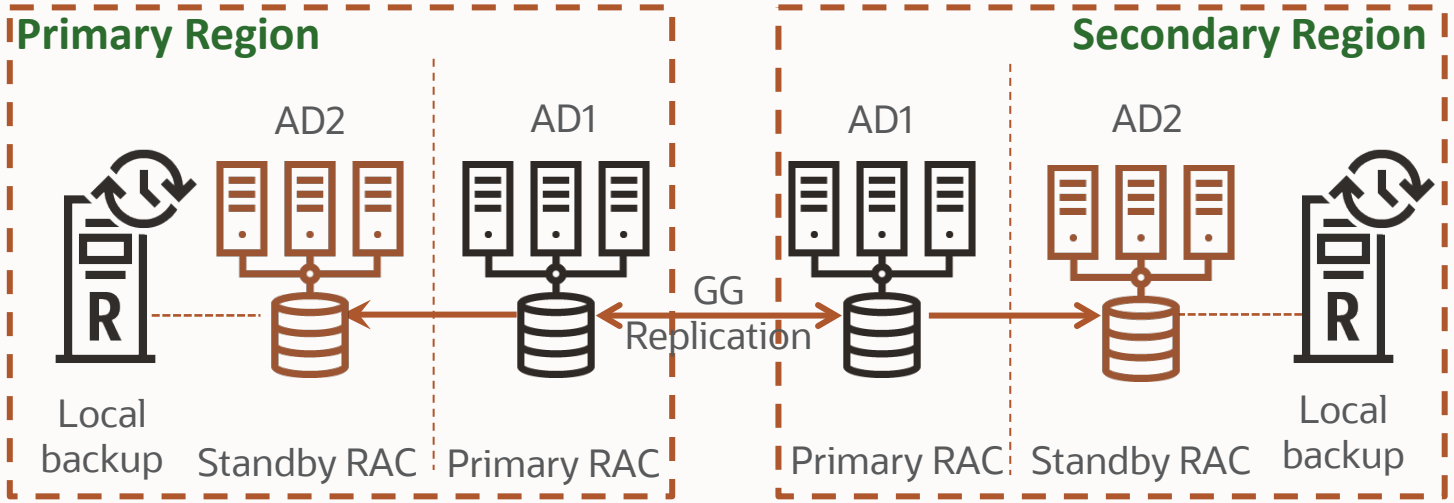

Outage Matrix

Unplanned Outage	RTO/RPO Service Level Objectives (f1)
Recoverable node or instance failure	**Zero or single-digit seconds (f2/f3)**
Disasters including corruptions and site failures	**Zero (f3)**
Planned Maintenance	
Most common software/hardware updates	Zero (f2)
Major database upgrade, application upgrade	**Zero (f3)**

f1: RPO=0 unless explicitly specified
f2: To achieve zero downtime or lowest impact, apply application checklist best practices
f3: Application failover is custom or with Global Data Services

External Resources

Maximum Availability Architecture

- Platinum MAA details:
  - Oracle Maximum Availability Architecture (MAA) - Platinum Tier
- MAA Home:
  - http://oracle.com/goto/maa
- On-Premise MAA:
  - https://www.oracle.com/database/technologies/high-availability/oracle-database-maa-best-practices.html
- Exadata MAA:
  - https://www.oracle.com/database/technologies/high-availability/exadata-maa-best-practices.html
- Cloud MAA:
  - https://www.oracle.com/database/technologies/high-availability/oracle-cloud-maa.html


ORACLE