

ORACLE



Exadata Database Machine: Maximum Availability Architecture (MAA)

Platinum Tier Focused

—
April 2020

Safe harbor statement

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, timing, and pricing of any features or functionality described for Oracle's products may change and remains at the sole discretion of Oracle Corporation.




Oracle Maximum Availability Architecture(MAA) Solution Options



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 (f2)
Disasters: corruptions and site failures	Hours to days. RPO since last backup or near zero with ZDLRA
Planned Maintenance	
Software/hardware updates	Minutes (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




SILVER

Prod/Departmental

Bronze +

- Real Application Clustering (RAC)
- Application Continuity



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

Checklist found in MAA OTN

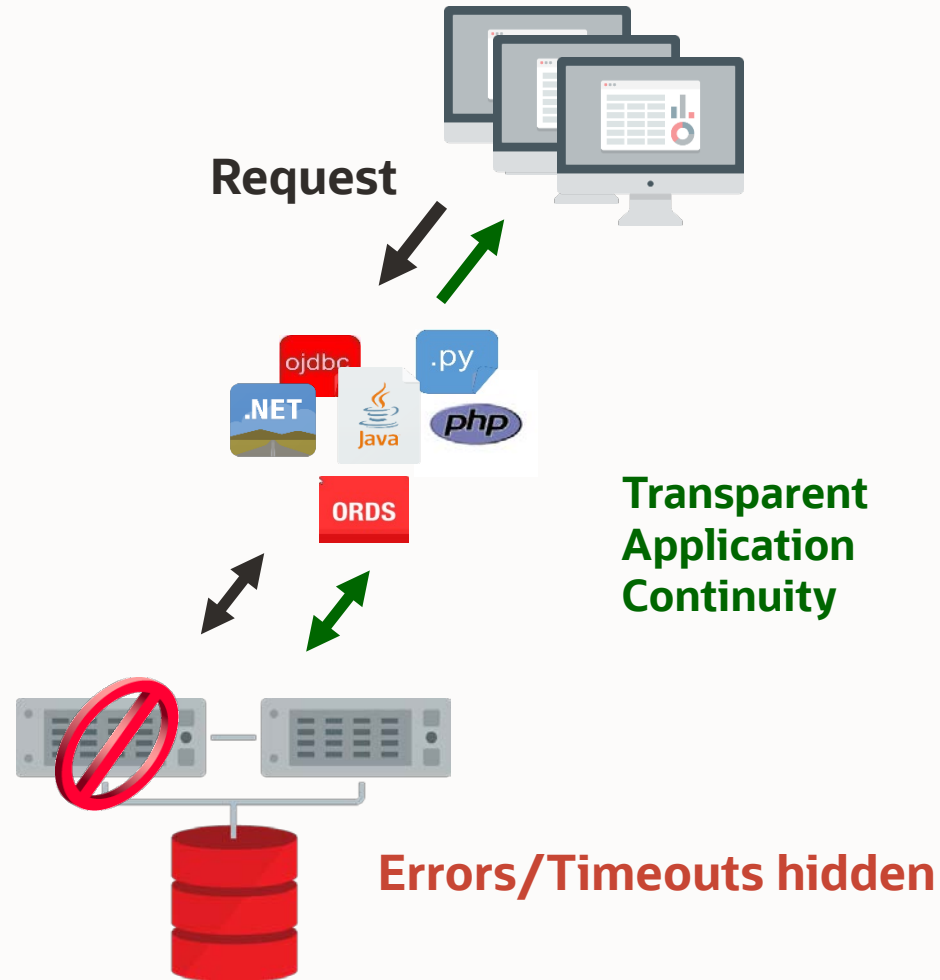
<https://www.oracle.com/technetwork/database/options/clustering/applicationcontinuity/adb-continuousavailability-5169724.pdf>





Transparent Application Continuity (TAC)

Application does not see errors during outages



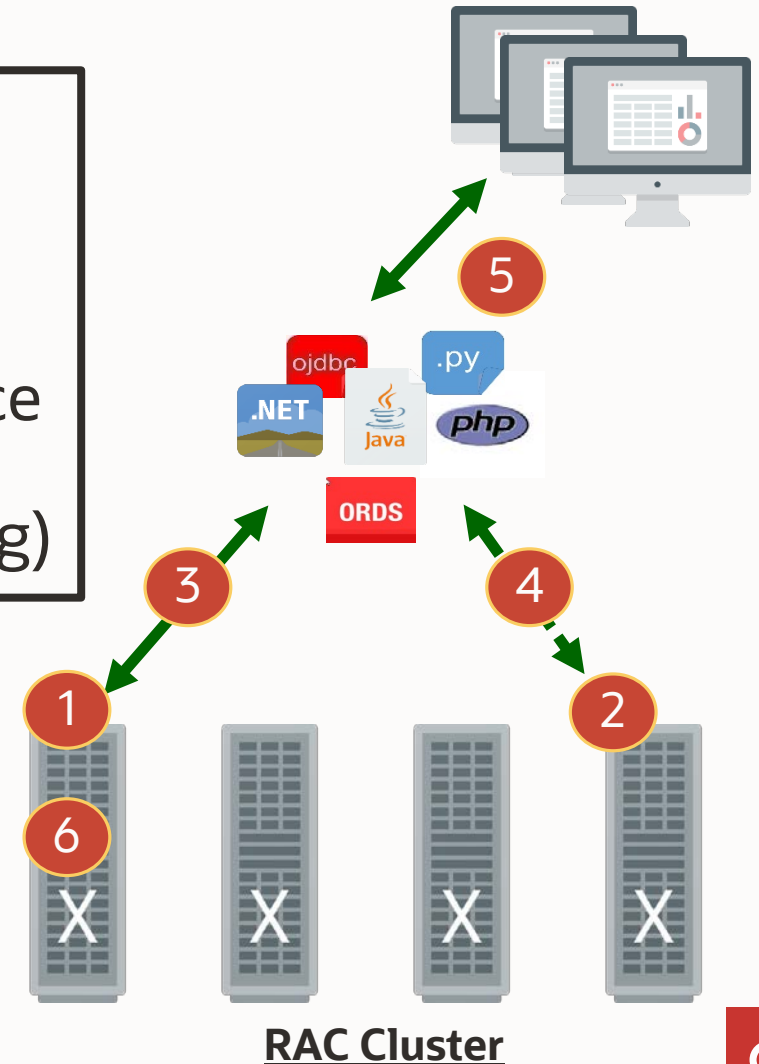
- Uses Application Continuity and Oracle Real Application Clusters
- Transparently tracks and records session information in case there is a failure
- Built inside of the database, so it works without any application changes
- Rebuilds session state and replays in-flight transactions upon unplanned failure
- 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 Outages!):

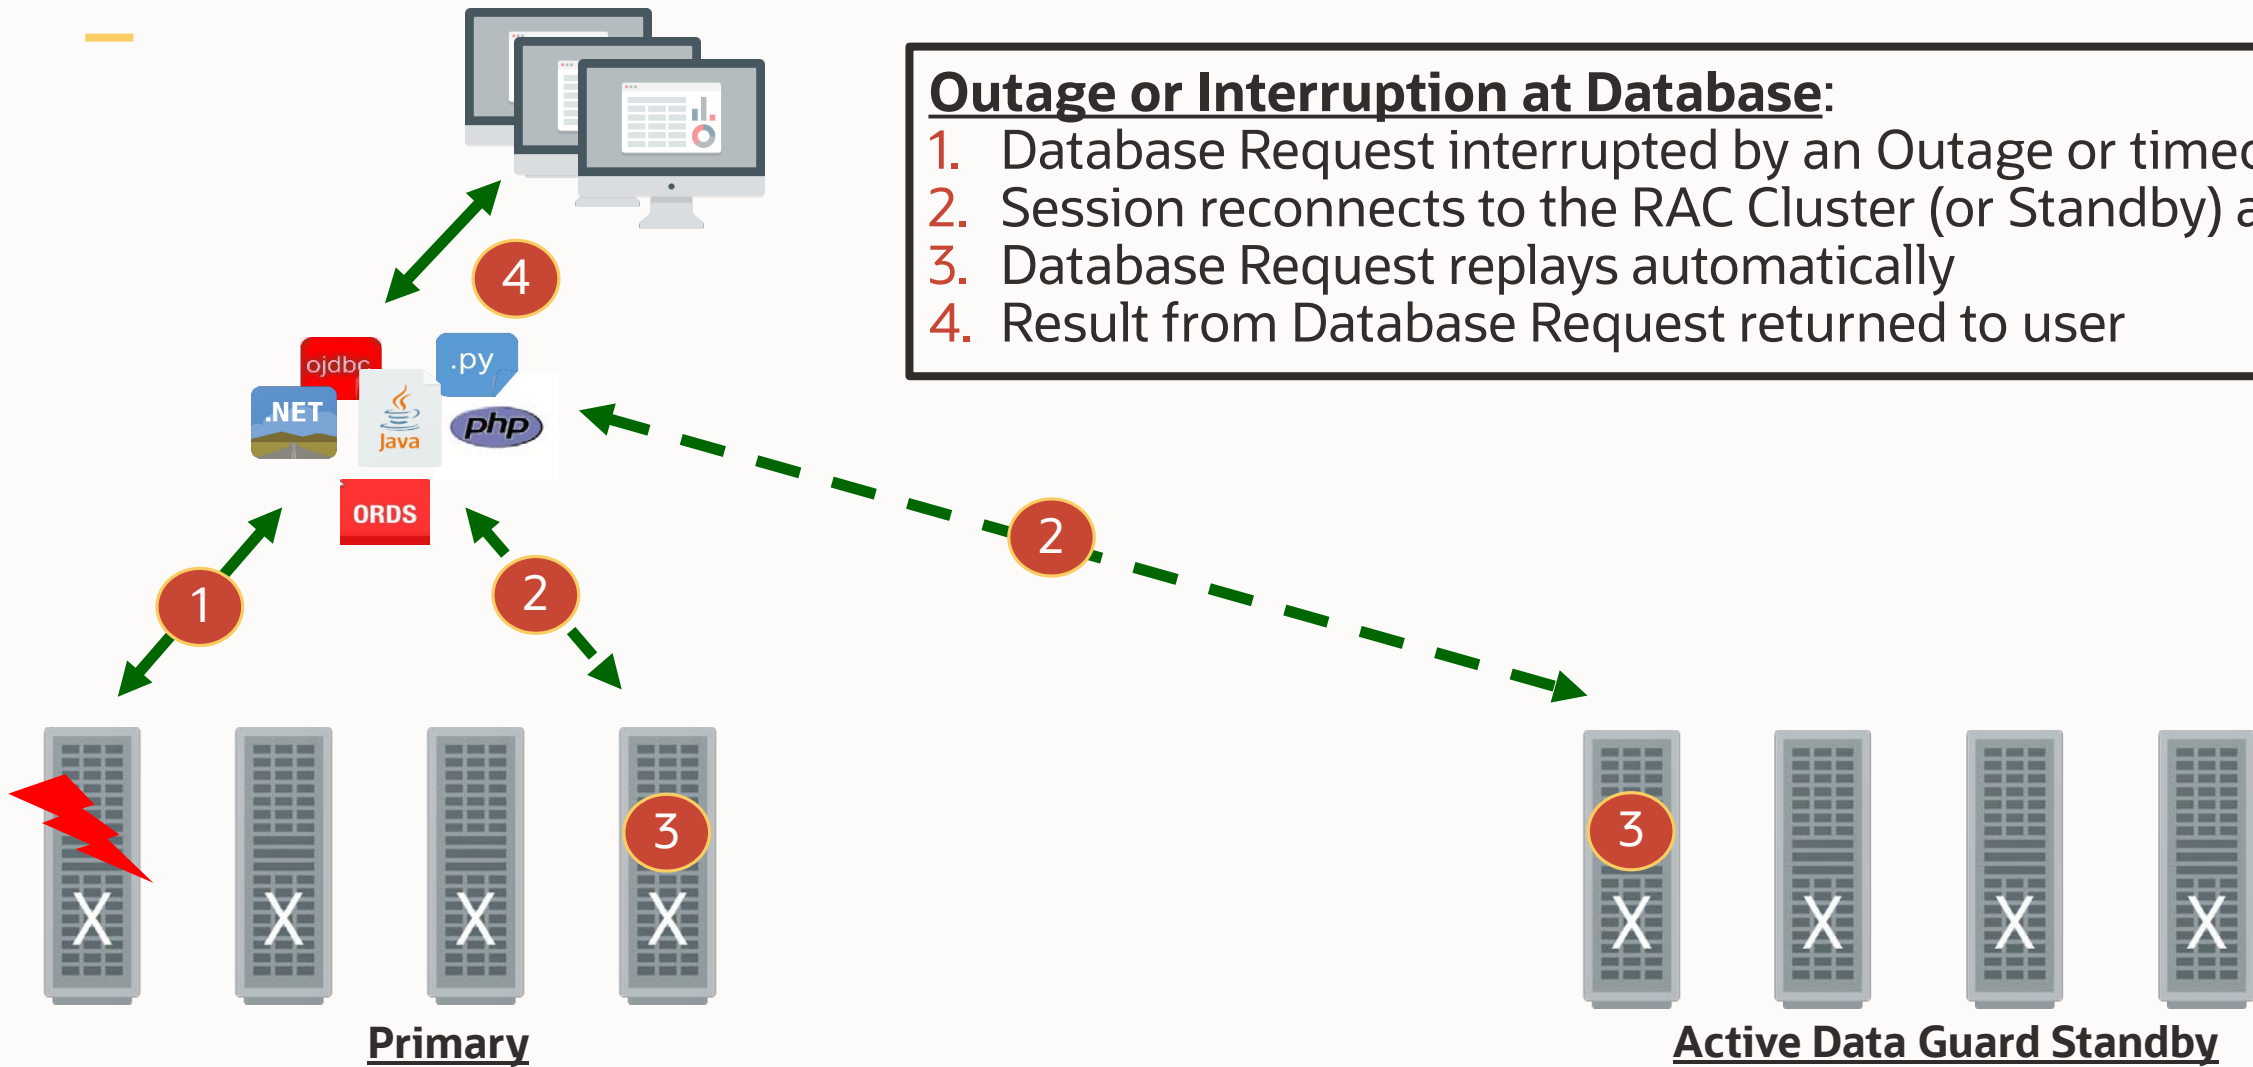
1. Database Service is relocated or stopped
2. Service starts on another RAC instance
3. Sessions connected to the service are drained
4. New sessions connect to Service on another instance
5. Results from Database Request returned to user
6. Maintenance activities can start on first node (rolling)



Unplanned Outages, without Impact

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





Checklist for Achieving Zero Application Downtime

1. Use Oracle Clusterware Service (never use default service)
2. Use Recommended Connection String
3. Configure FAN for Connection Pool
4. Drain your service
5. Use Application Continuity or Transparent Application Continuity

1) MAA Whitepaper: [Application Checklist for Continuous Service for MAA Solutions](#)

**2) Using RHPHelper to Minimize Downtime During Planned Maintenance on Exadata
([MOS 2385790.1](#))**

3. Fleet Patch and Provisioning incorporates [MAA practices](#)



GOLD

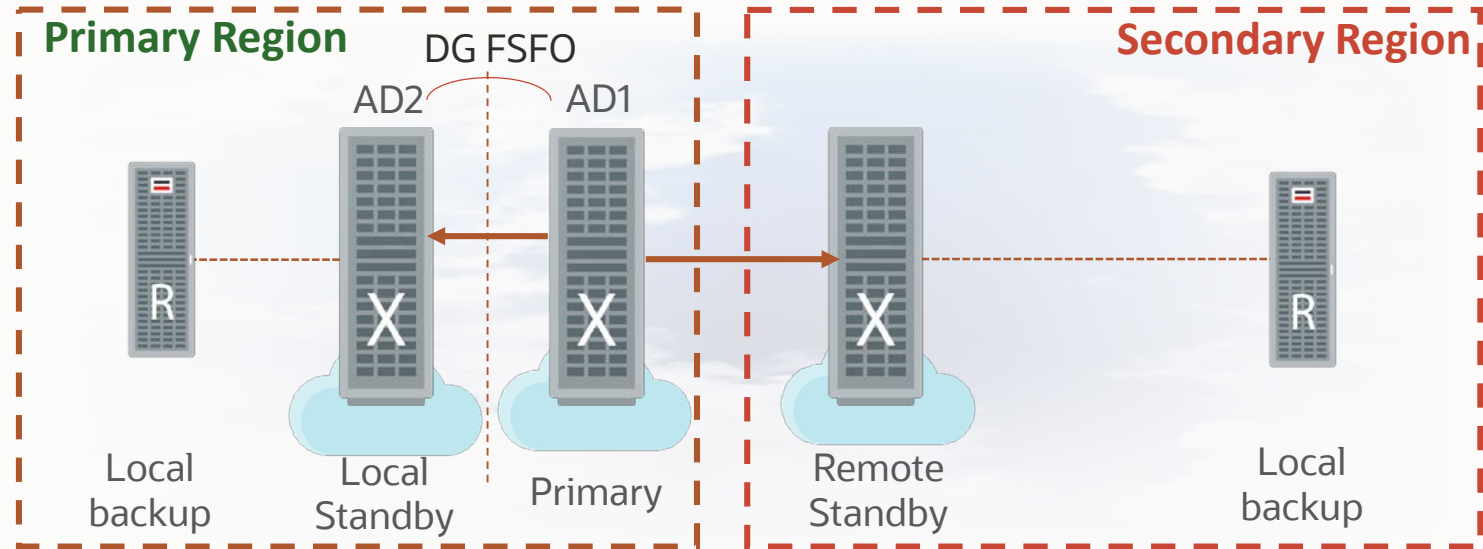
Mission Critical

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






Active Data Guard Overview

**Primary
Open Read-Write**



**Standby
Open Read-Only**



DML Redirection



Zero Data Loss at any Distance



Automatic Block Repair



Offload read only or read mostly workloads to the standby database

- Synchronous zero data loss replication
- Database rolling upgrade to reduce downtime for planned maintenance
- Automatic failover for High Availability

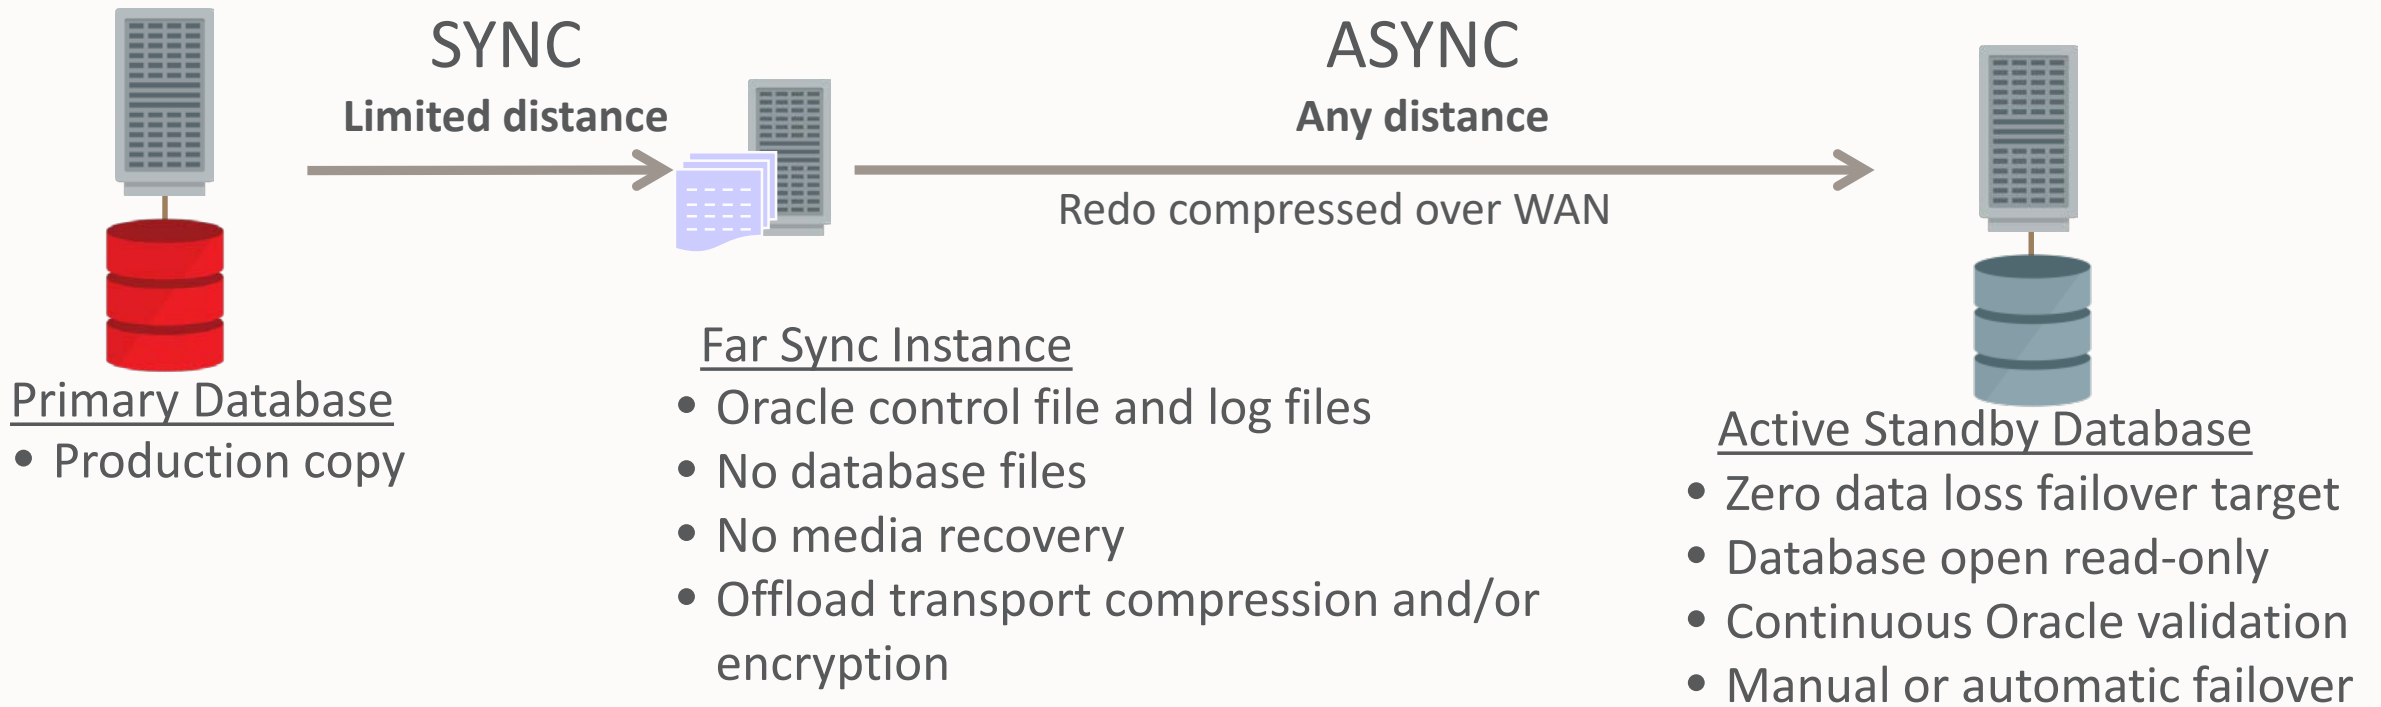
Multi-instance Redo Apply for RAC
(In Memory supported)





Active Data Guard Far Sync

Zero Data Loss Protection at Any Distance



PLATINUM

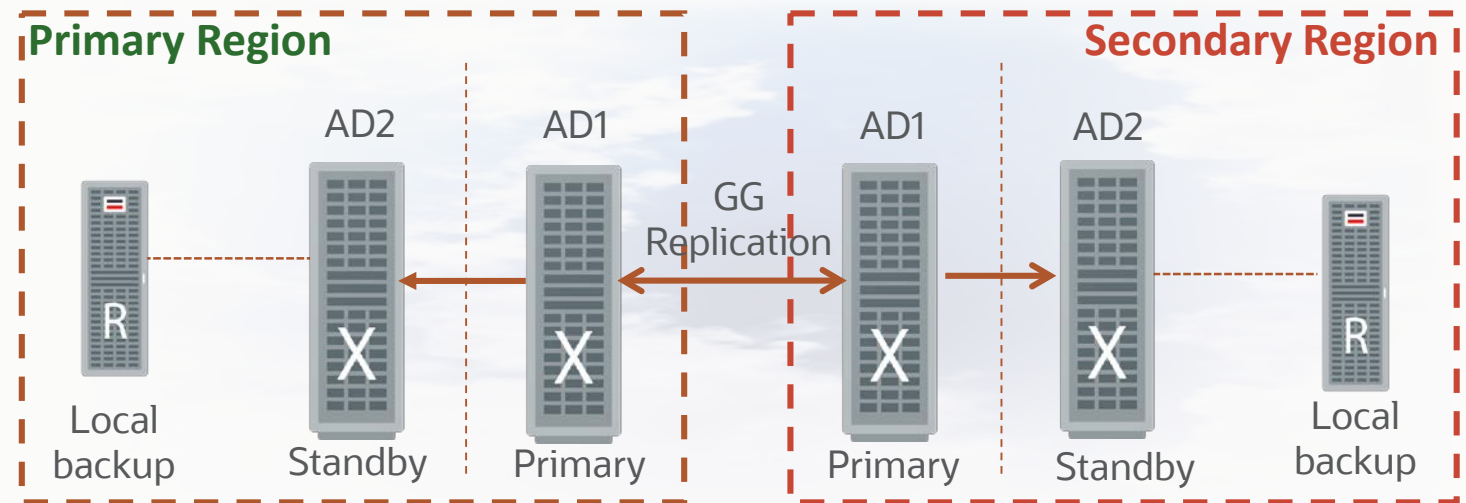
Extreme Critical

Gold +

- GoldenGate Active/Active Replication
- Optional Sharding & Editions Based Redefinition

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 & Editions Based Redefinition for zero downtime application upgrade
- Sharding for scalability and fault isolation
- 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

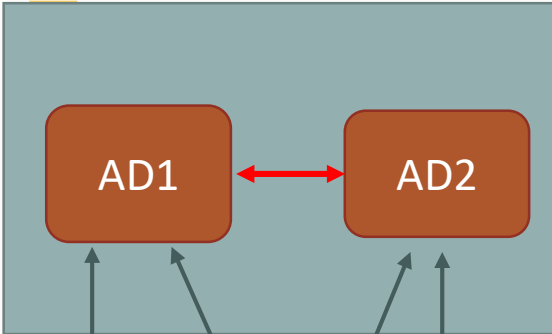
f2: To achieve zero downtime or lowest impact, apply application checklist best practices

f3: Application failover is custom or with Global Data Services

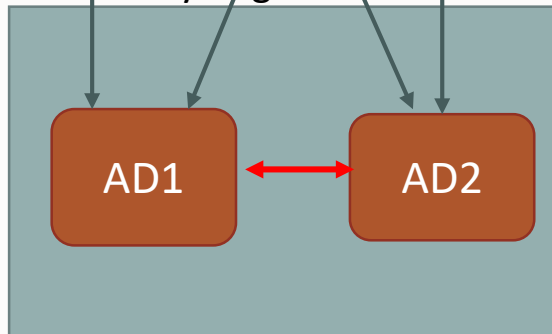


Data Center Architecture & Requirements

Primary Region – West NAS



Secondary Region – East NAS



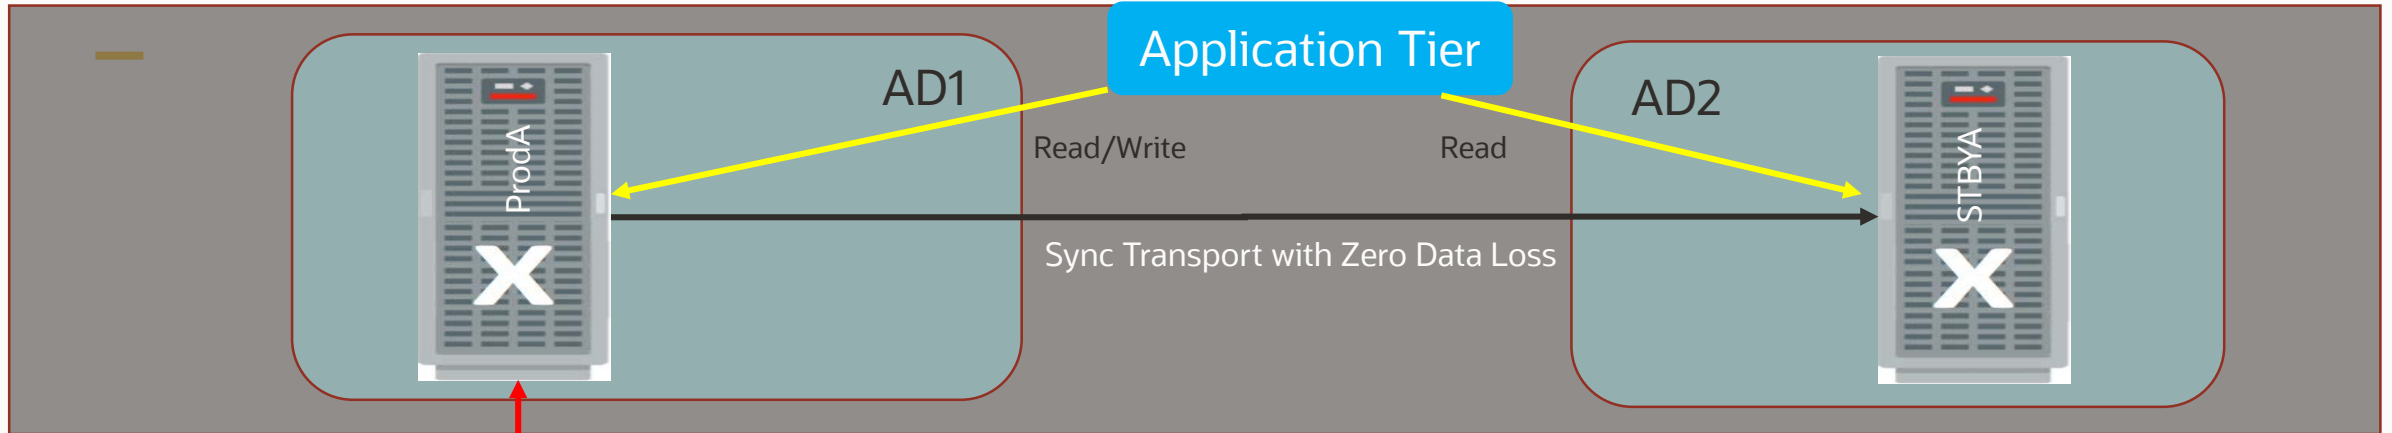
- A minimum of 2 Regions for Disaster Recovery Failover
 - Region is a localized geographic area
 - West Coast NAS – Primary example
 - East Coast NAS – Secondary example
- Each Region should have a minimum of 2 Availability Domains (AD)
- Availability Domain Characteristics
 - AD's are isolated from each other & fault tolerant
 - AD's do not share infrastructure such as power, cooling or AD Network
 - A failure of one AD does not effect other AD's.
 - AD's within a Region are connected via high speed network within same geographical area.

↔ High Speed with < 1ms Latency

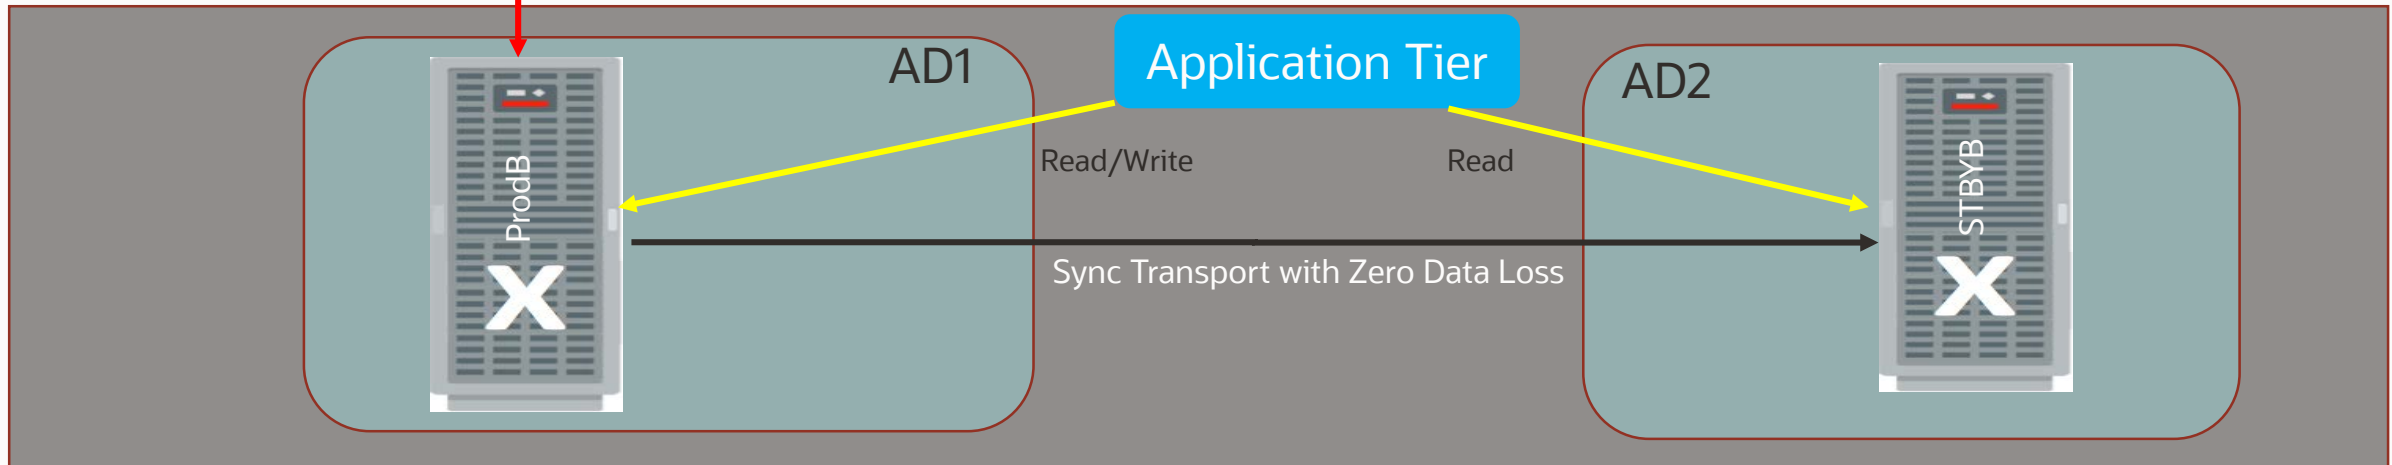


Platinum Reference Architecture

Primary Region 1 – West US



Secondary Region – East US



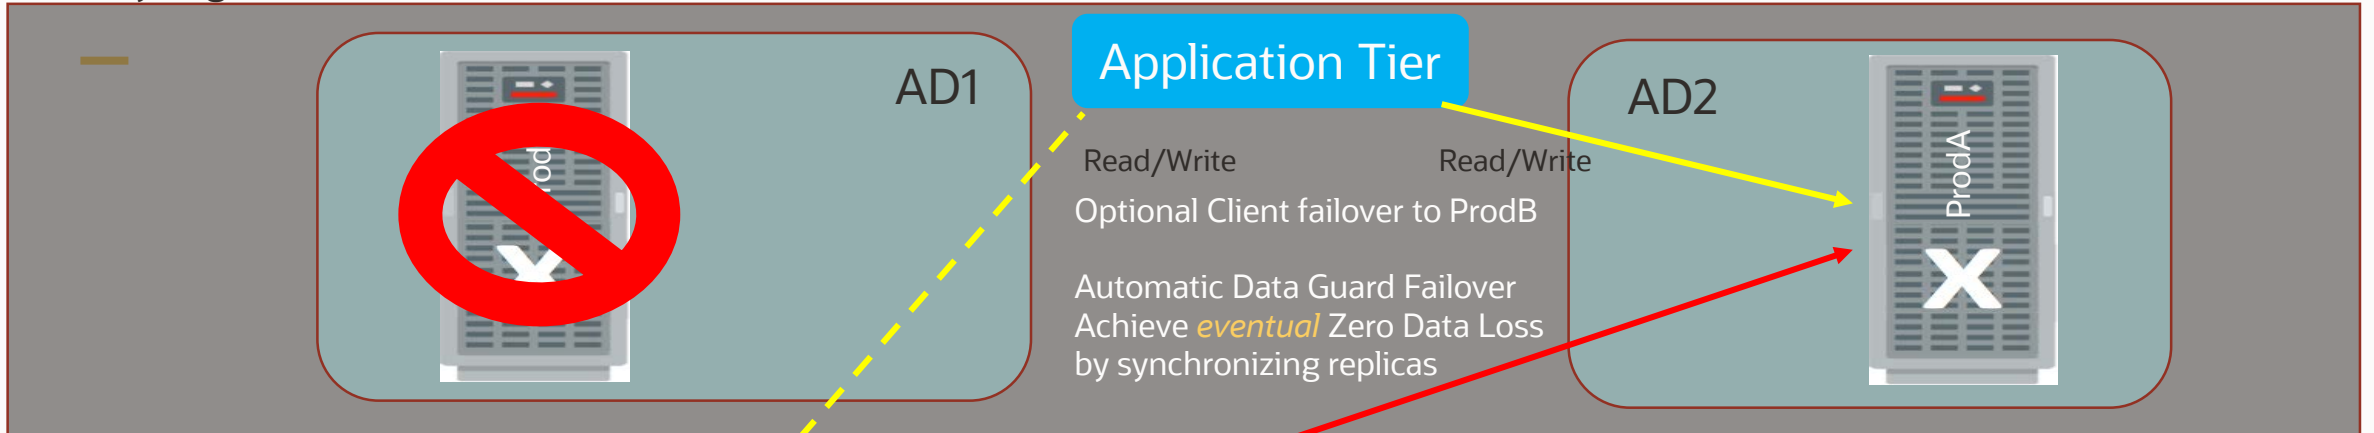
➔ Active Data Guard Fast-Start Failover,

↔ Oracle GoldenGate Replication

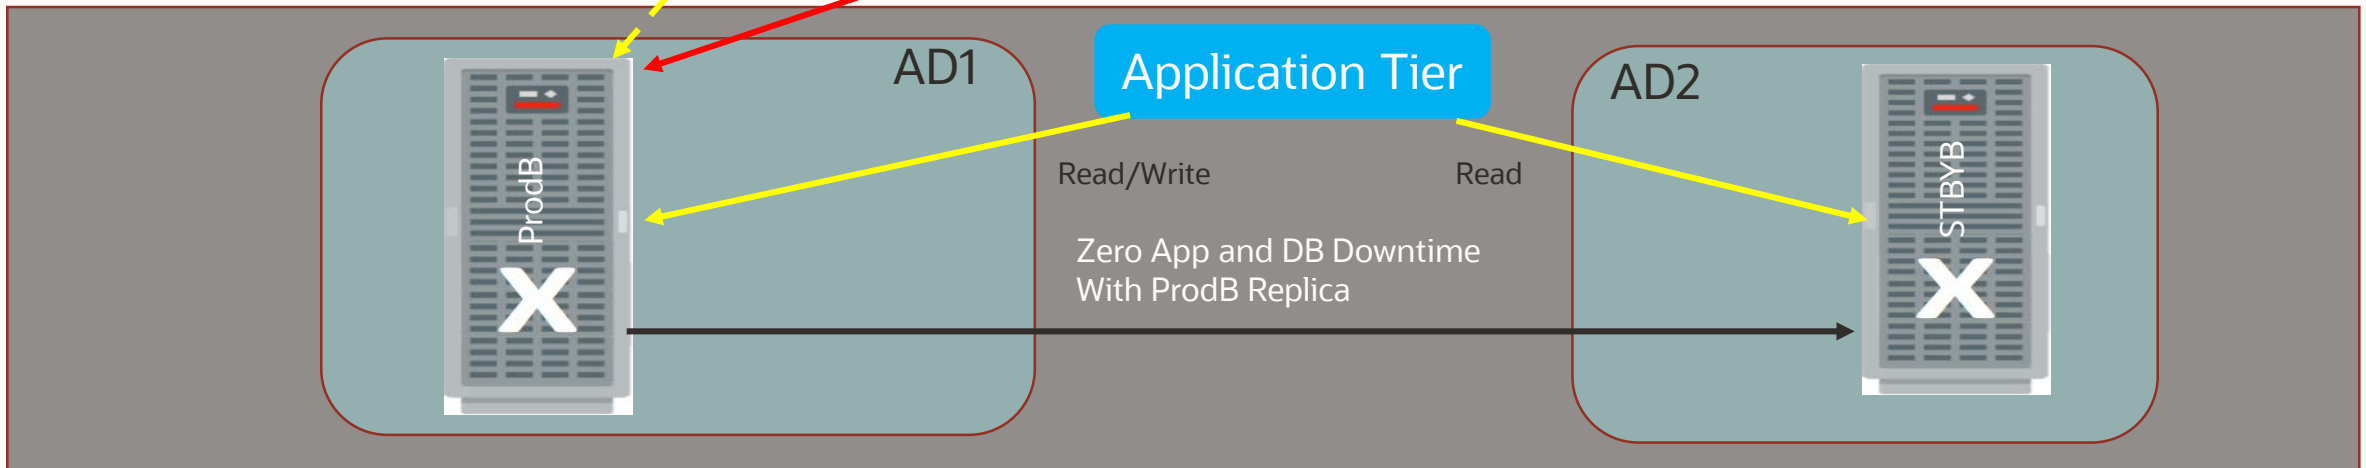
Reference Architecture – Zero App Downtime and Zero Data Loss

(Disaster Scenario: Loss of Entire Data Center)

Primary Region 1 – West US



Secondary Region – East US



➔ Active Data Guard Fast-Start Failover,

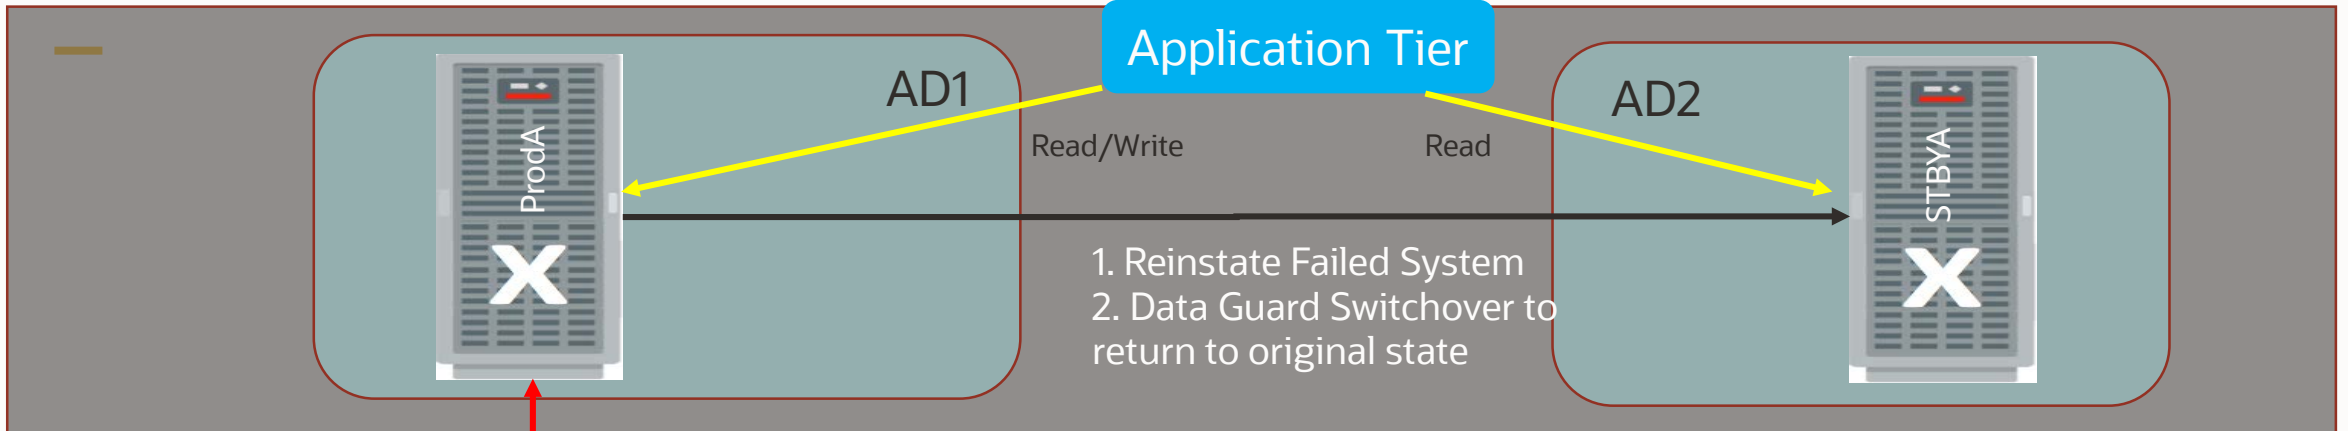
↔ Oracle GoldenGate Replication



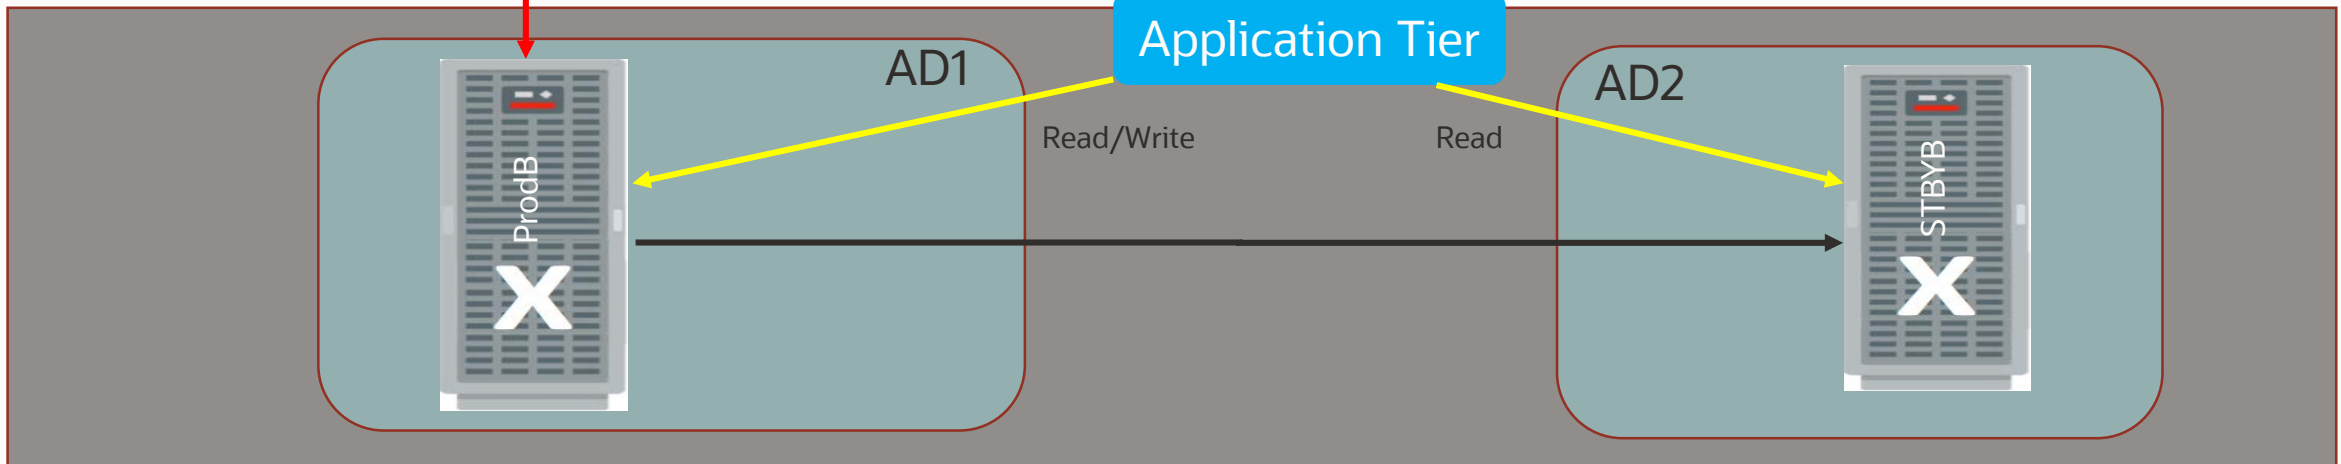
Reference Architecture – Switching Back

ProdA returns to Primary and STBYA+ to Standby

Primary Region 1 – West US



Secondary Region – East US



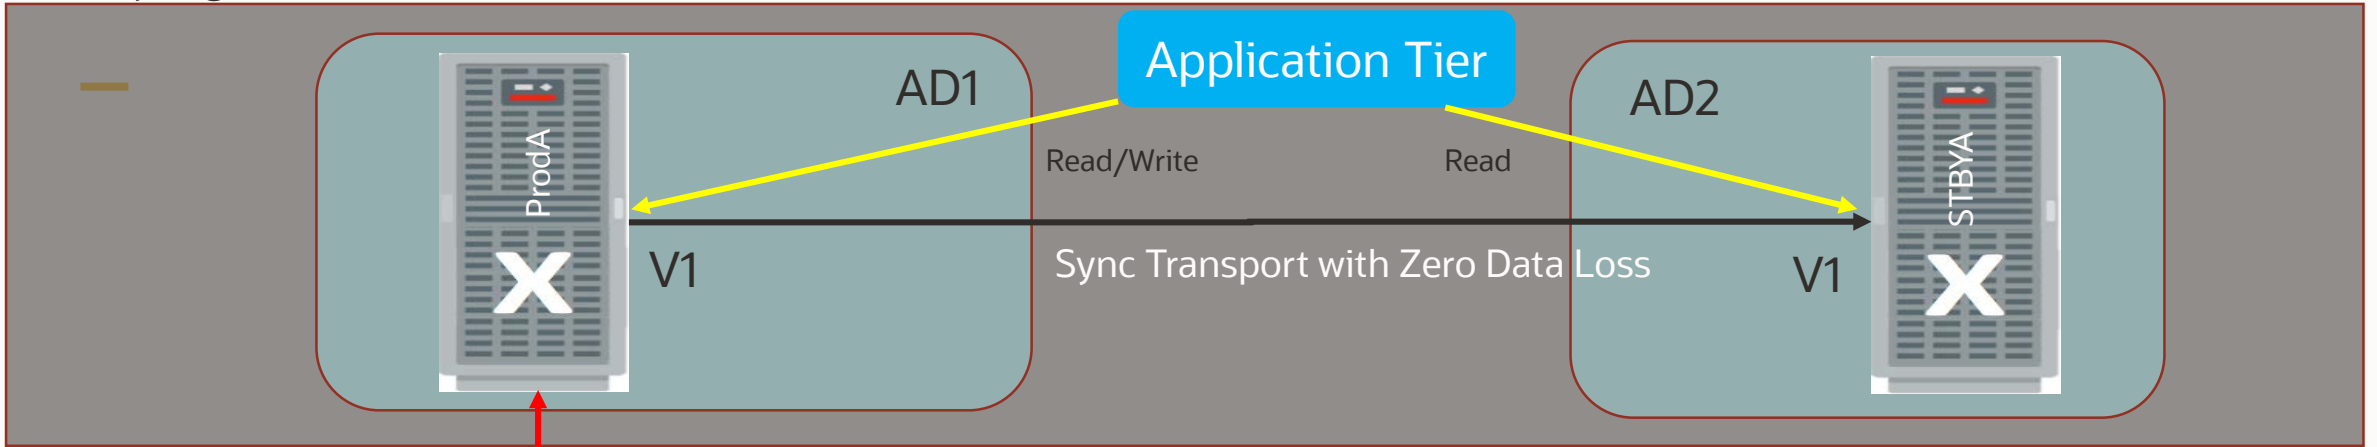
➔ Active Data Guard Fast-Start Failover,

↔ Oracle GoldenGate Replication

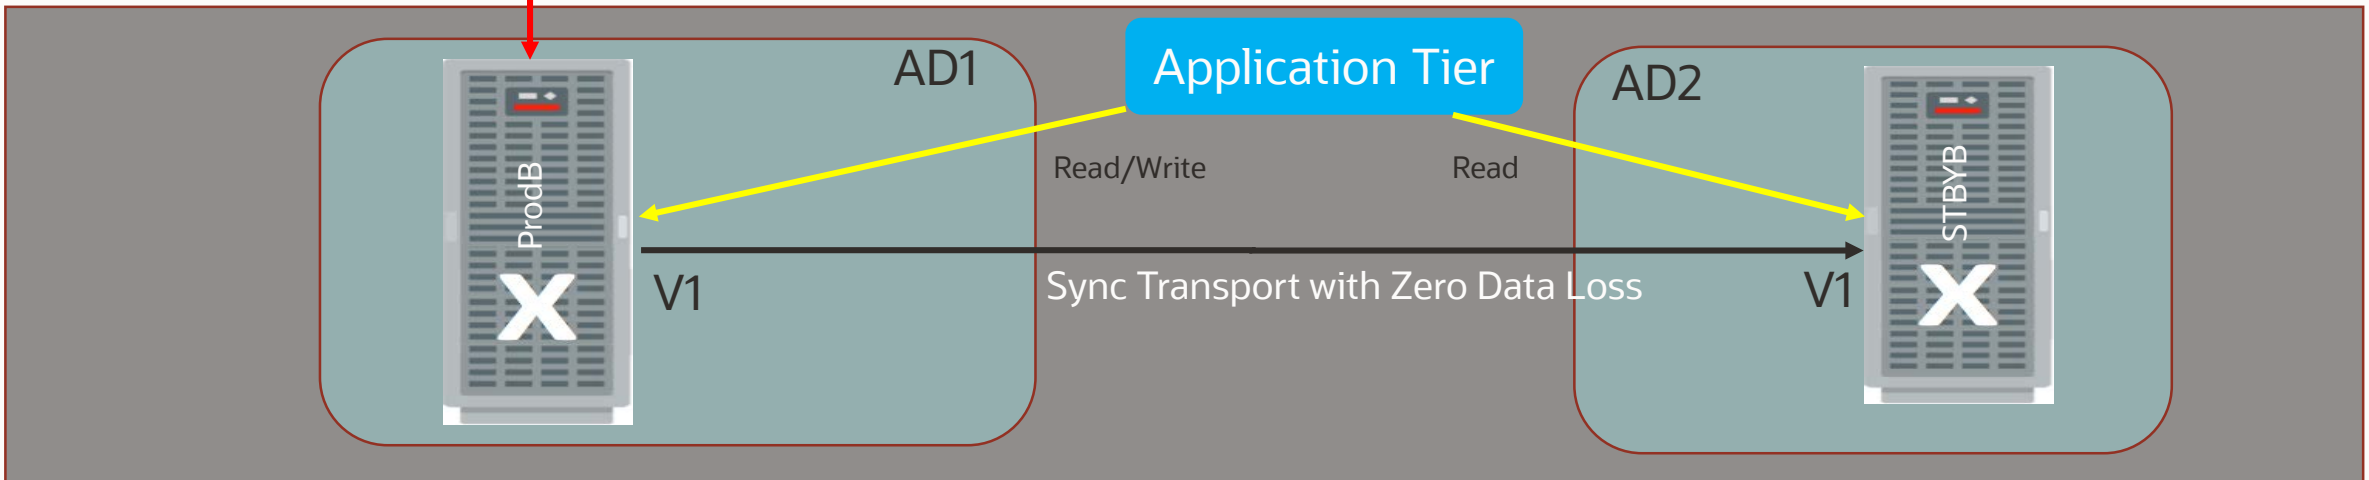


Reference Architecture – Upgrade Scenario

Primary Region 1 – West US



Secondary Region – East US



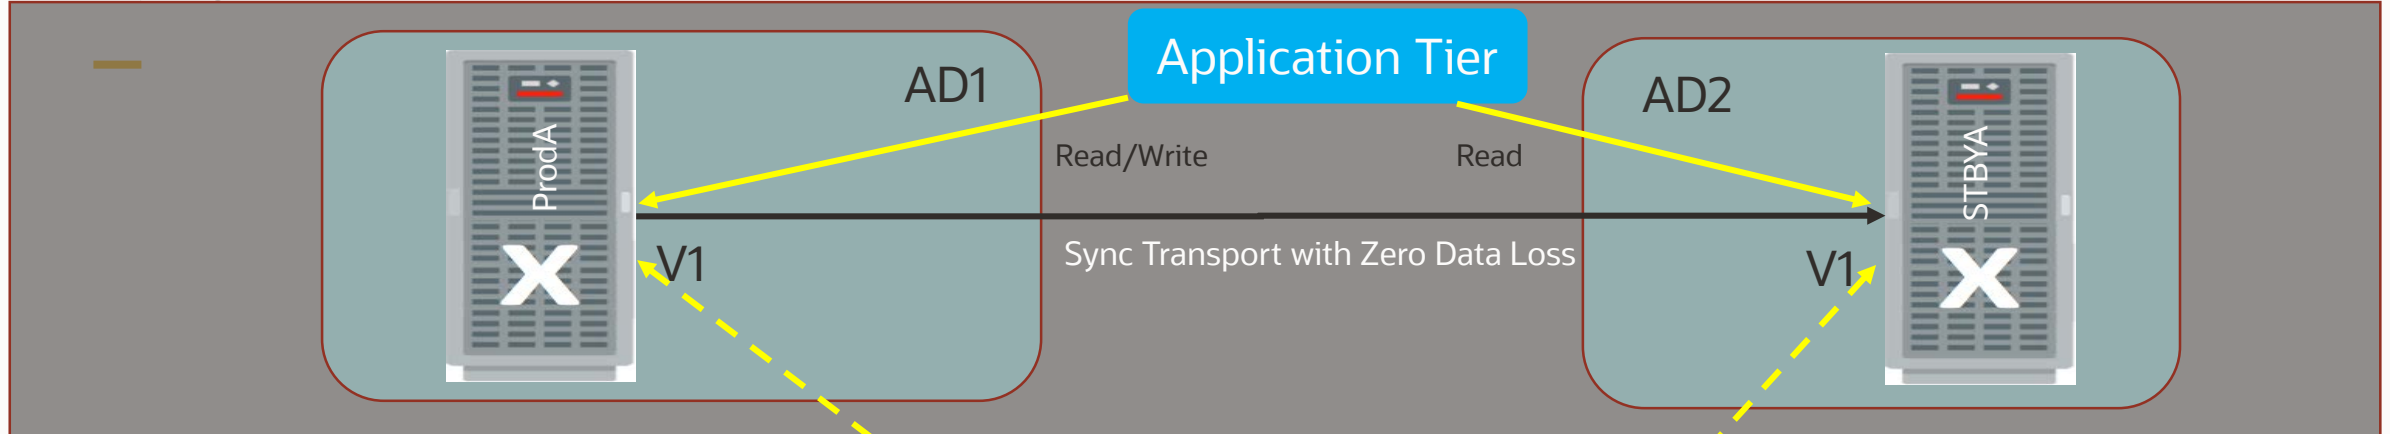
➔ Active Data Guard Fast-Start Failover,

↔ Oracle GoldenGate Replication

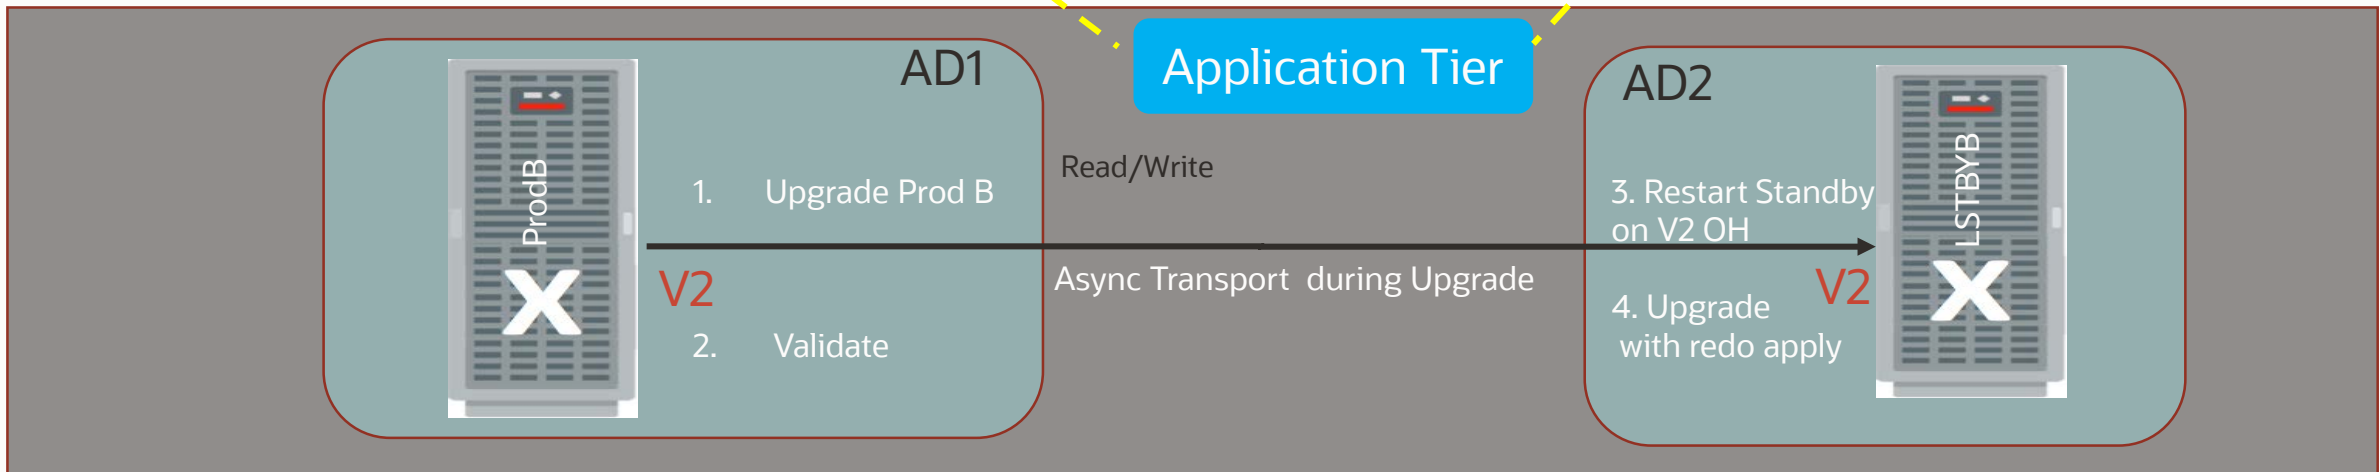


Upgrade Scenario Step 1: Upgrade Prod B and Standby

Primary Region 1 – West US



Secondary Region – East US



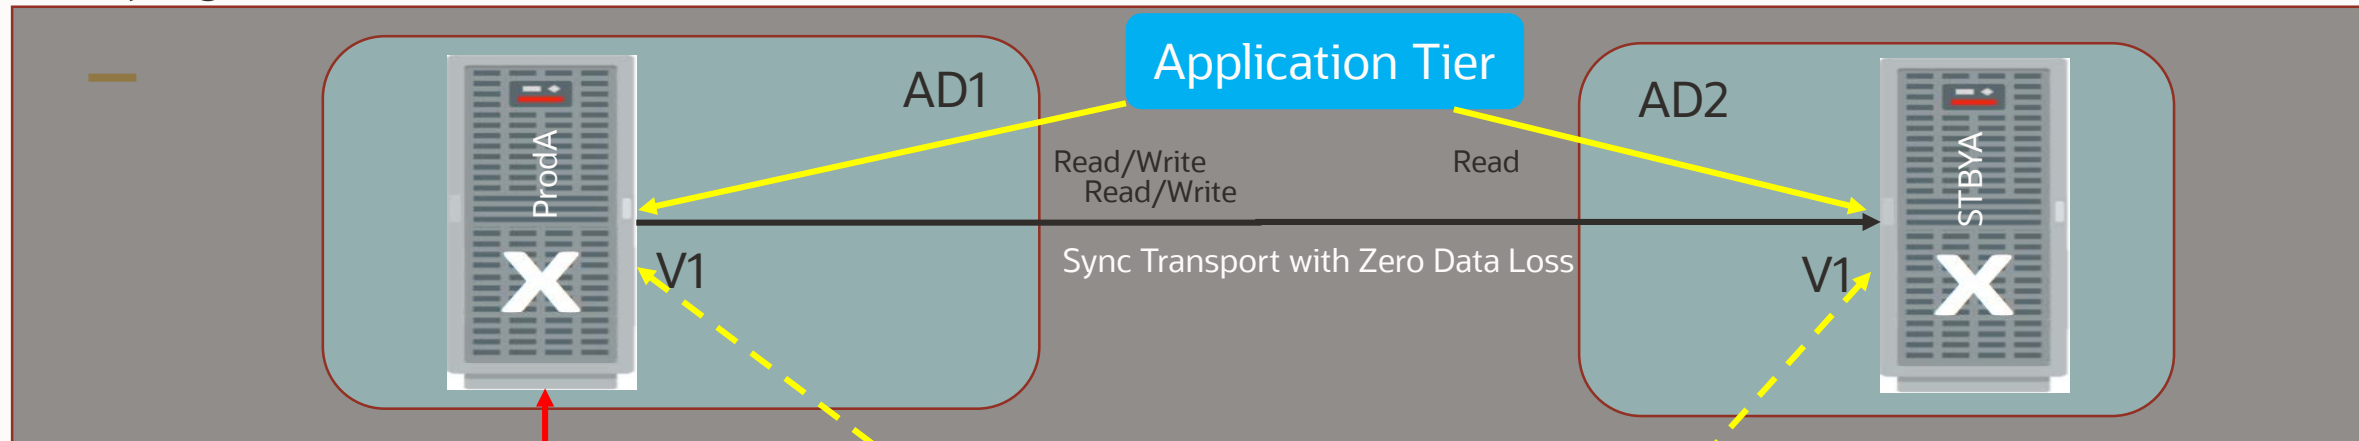
➔ Active Data Guard Fast-Start Failover,

↔ Oracle GoldenGate Replication

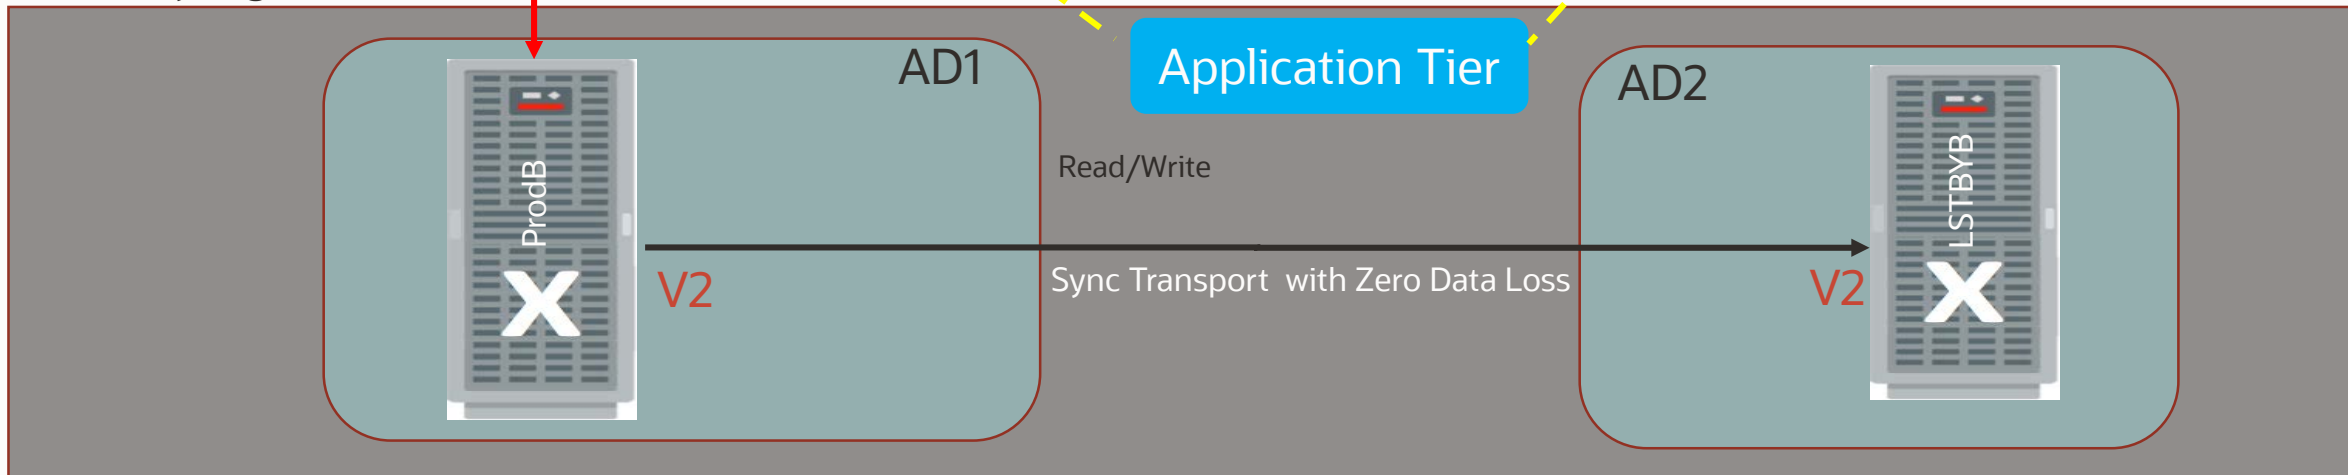


Upgrade Scenario Step 2: Synchronize GG Replicas

Primary Region 1 – West US



Secondary Region – East US



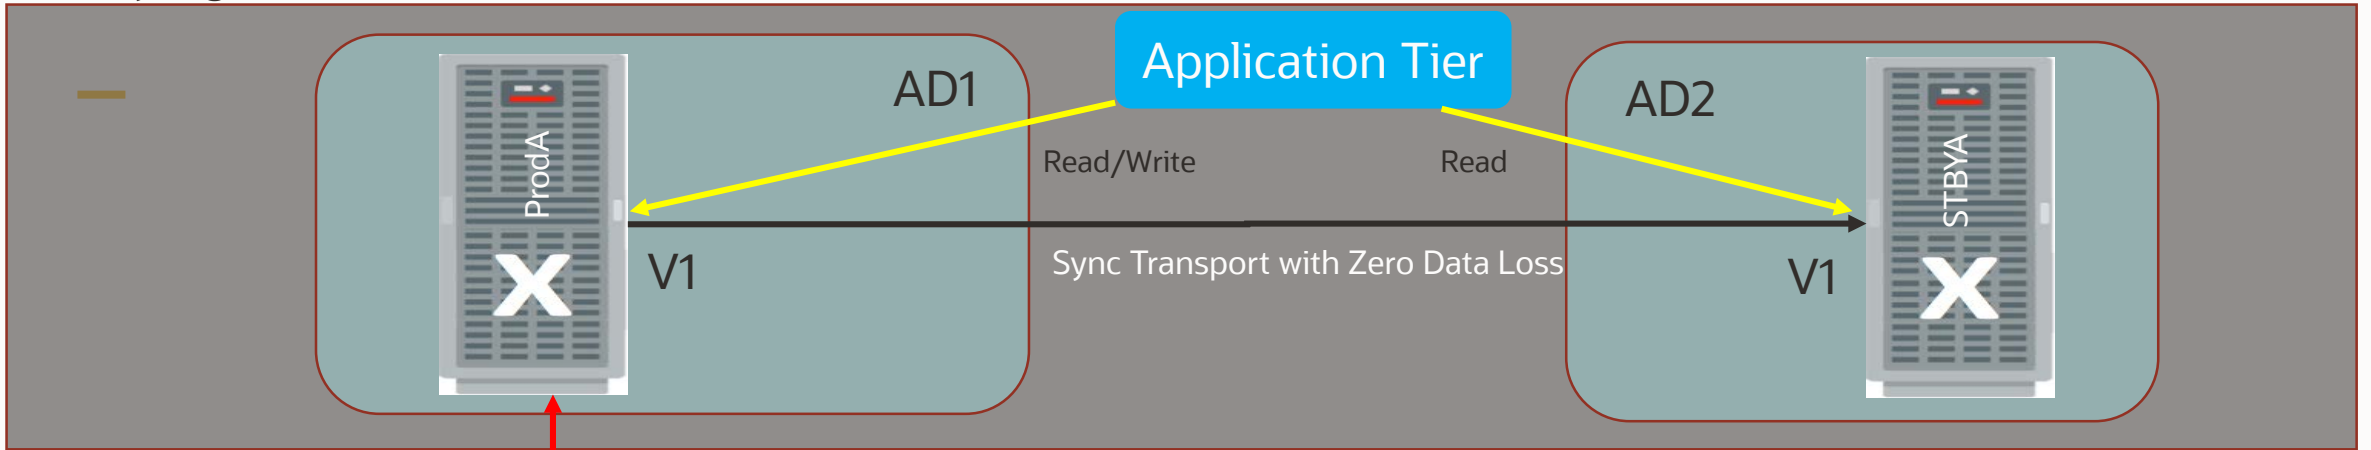
➔ Active Data Guard Fast-Start Failover,

↔ Oracle GoldenGate Replication

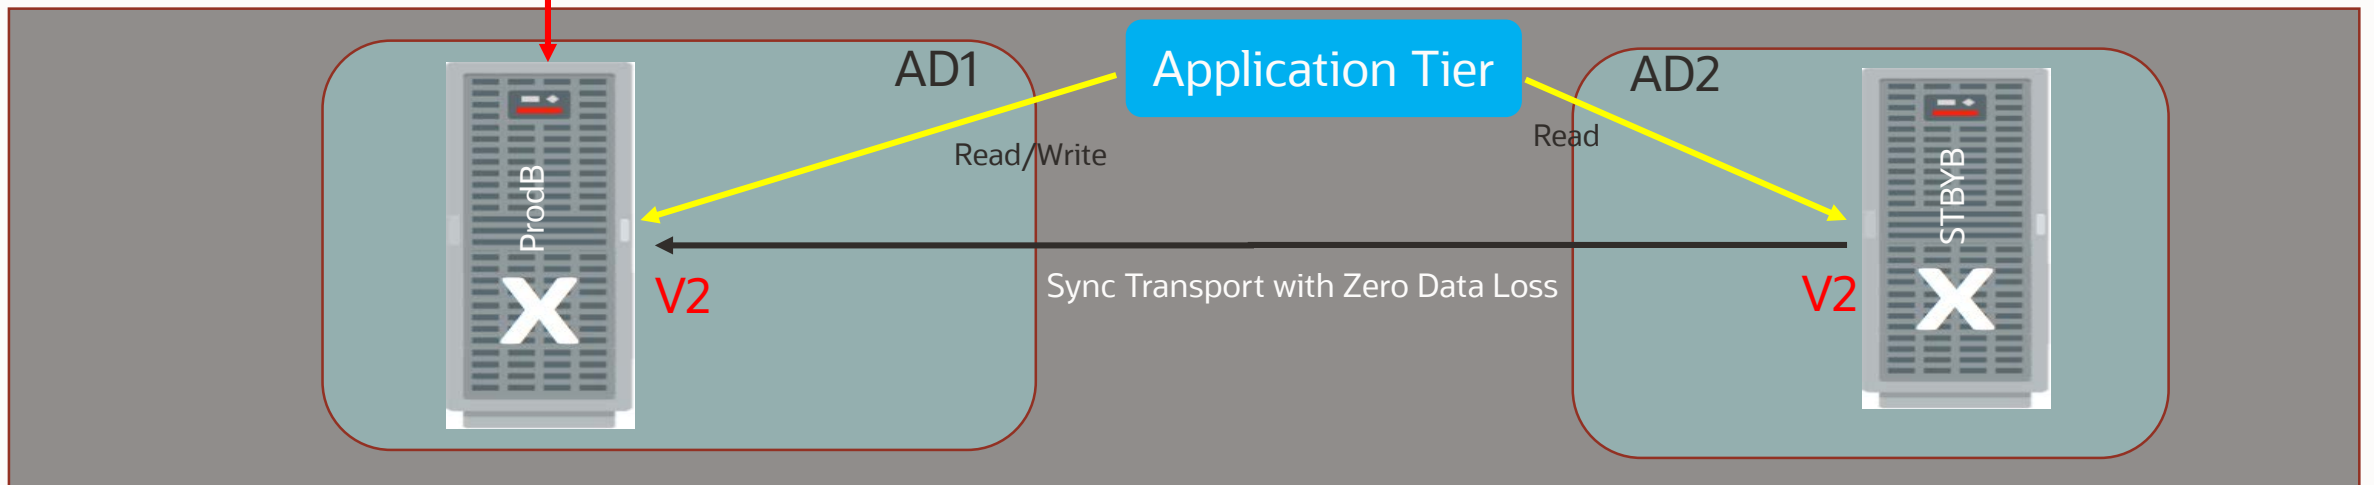


Upgrade Scenario Step 3: Co-Exist with V1 and V2

Primary Region 1 – West US



Secondary Region – East US



➔ Active Data Guard Fast-Start Failover,

↔ Oracle GoldenGate Replication



Platinum Advantages for Upgrade

Final Decision Point

Benefits

1. Zero Downtime and Zero Data Loss
2. Evaluate V1 and V2 at the same time
3. GoldenGate replication between V1 and V2 provides simple switchover and fallback

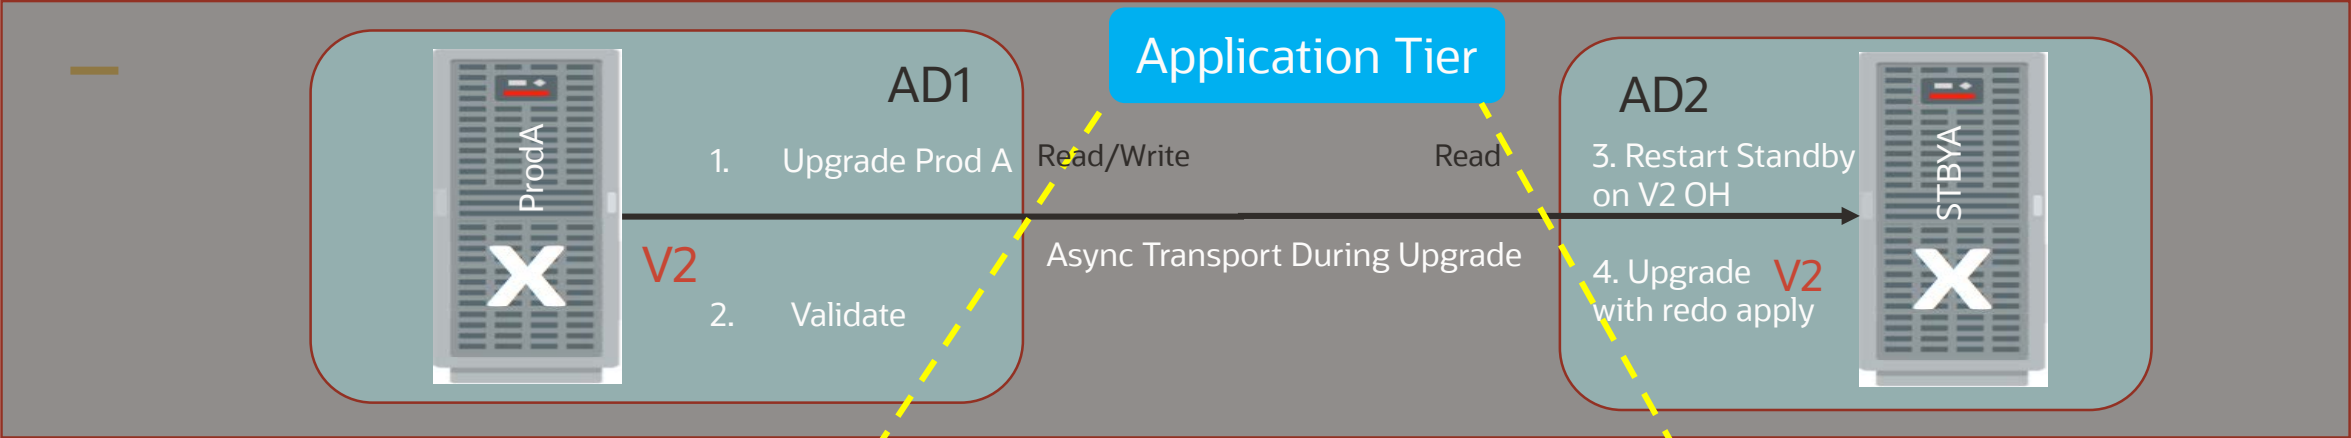
Once V2 has been validated and deemed acceptable, then:

- Repeat process and upgrade both V1 primary and standby at the same time

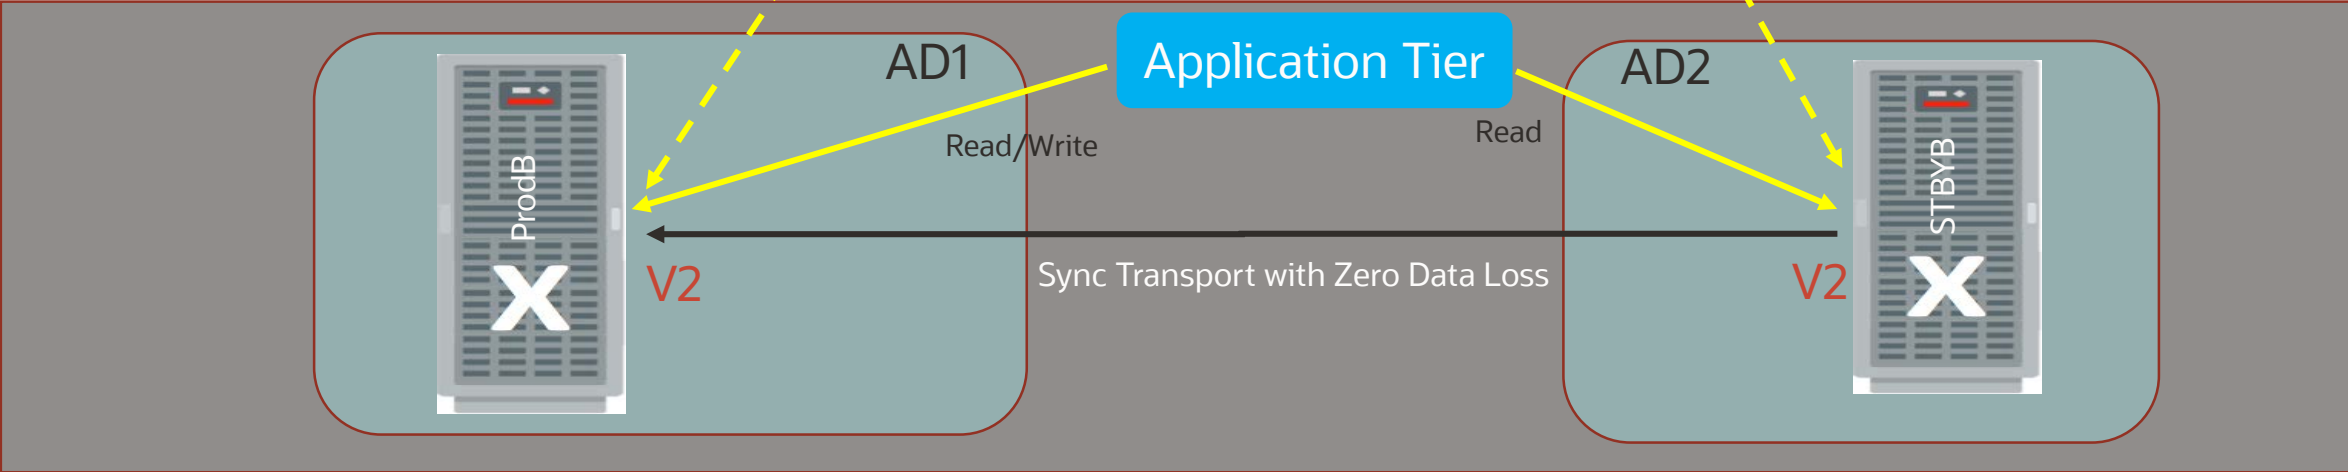


Upgrade Scenario Step 4: Upgrade Prod A and Standby B to V2

Primary Region 1 – West US



Secondary Region – East US



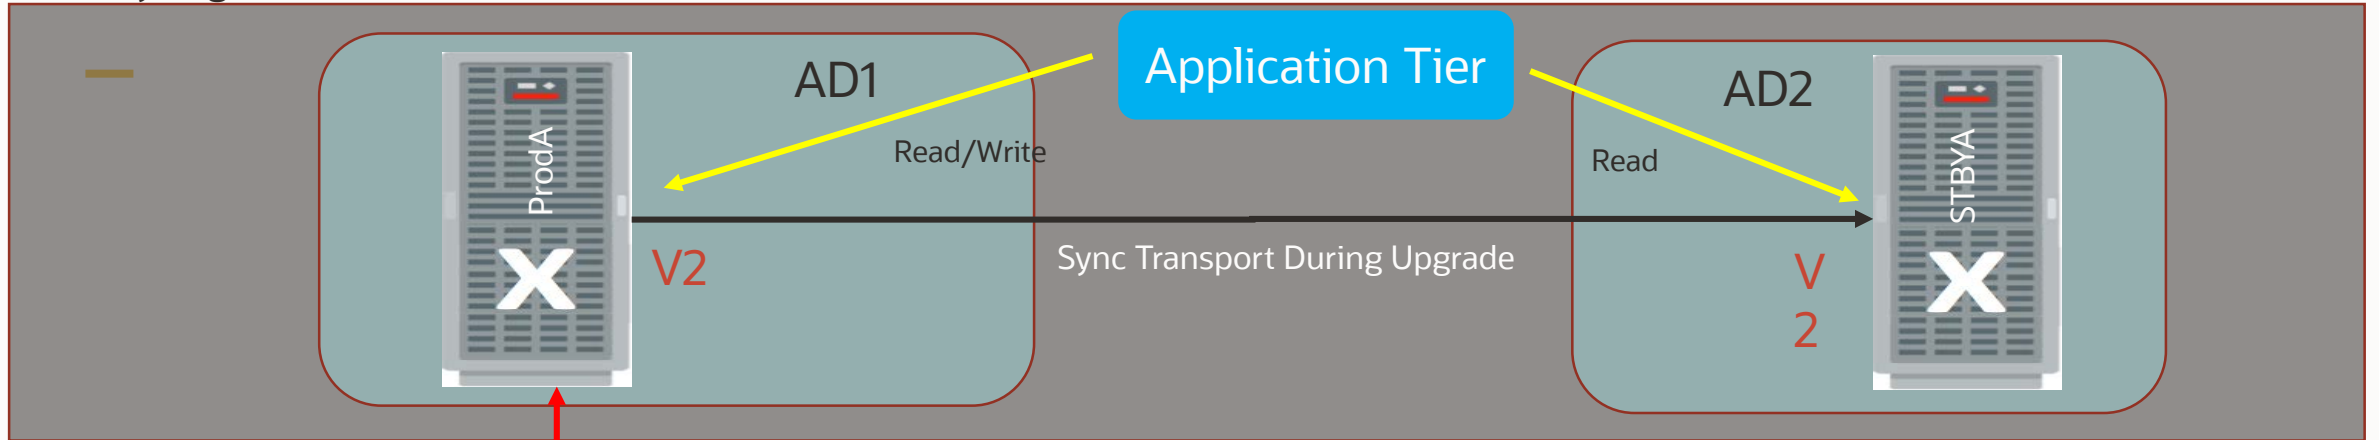
➔ Active Data Guard Fast-Start Failover,

↔ Oracle GoldenGate Replication

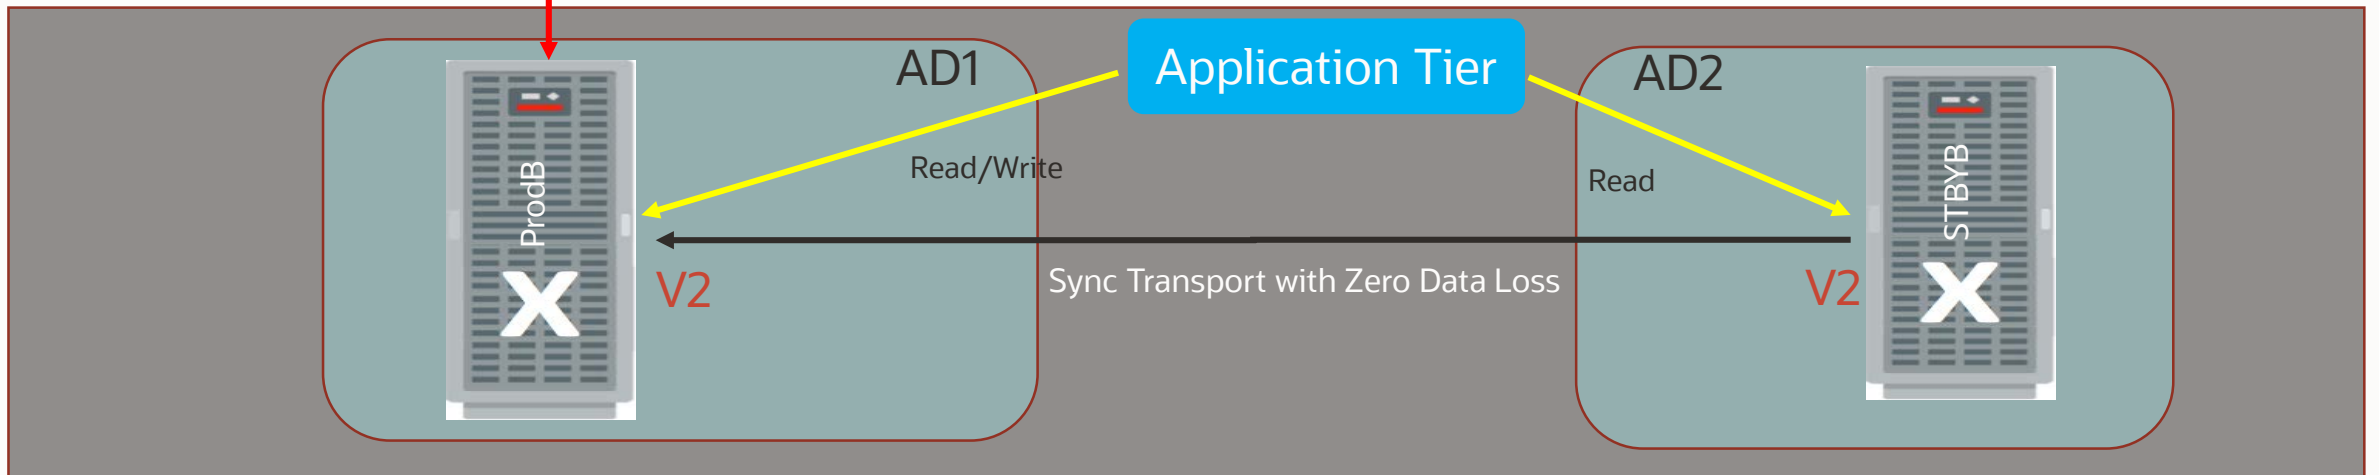


Upgrade Scenario Steps 5/6: Synchronize and Back to Normal

Primary Region 1 – West US



Secondary Region – East US



➔ Active Data Guard Fast-Start Failover,

↔ Oracle GoldenGate Replication



Unplanned Outages for Platinum MAA with Exadata

Unplanned Outages	Database Downtime (RTO)	Application Impact	Data Loss (RPO)	Key Enablers
Exadata Cluster Network Fabric or Storage Failures	Zero	Zero or Near Zero	Zero	Exadata ASM Disk Groups in High Redundancy
RAC Instance or Node Failures	Zero	Single Digit Seconds	Zero	Exadata, RAC Application Continuity with MAA Checklist
Data Corruptions	Zero	Zero or Isolated Failure	Zero or Isolated Logical Impact	Active Data Guard MOS 1302539.1 Flashback Technologies ZDLRA
Disasters including database, cluster or site failures	Zero since GG replica is available	Zero or Near Zero Single Digit Seconds with GDS	Eventual Zero	Oracle GoldenGate Data Guard Fast-Start Failover Custom App Failover or Global Data Services or Site Guard



Planned Maintenance for Platinum MAA with Exadata

Planned Maintenance	Database Downtime (RTO)	Application Impact	Key Enablers
Exadata Infrastructure SW or HW Updates	Zero	Zero or Near Zero	Exadata Platform ASM Disk Groups in High Redundancy
Database and Grid Infrastructure Software Updates	Zero	Zero	RAC Application Continuity <u>Continuous Availability - Application Checklist for Continuous Service for MAA Solutions</u>
Database Upgrades or non-Rolling Updates	Zero	Zero or Near Zero	GoldenGate Custom Application failover or Global Data Services



ORACLE

Our mission is to help people
see data in new ways, discover insights,
unlock endless possibilities.

