
ORACLE®

CON6761 Online Database Maintenance with No Impact to Applications

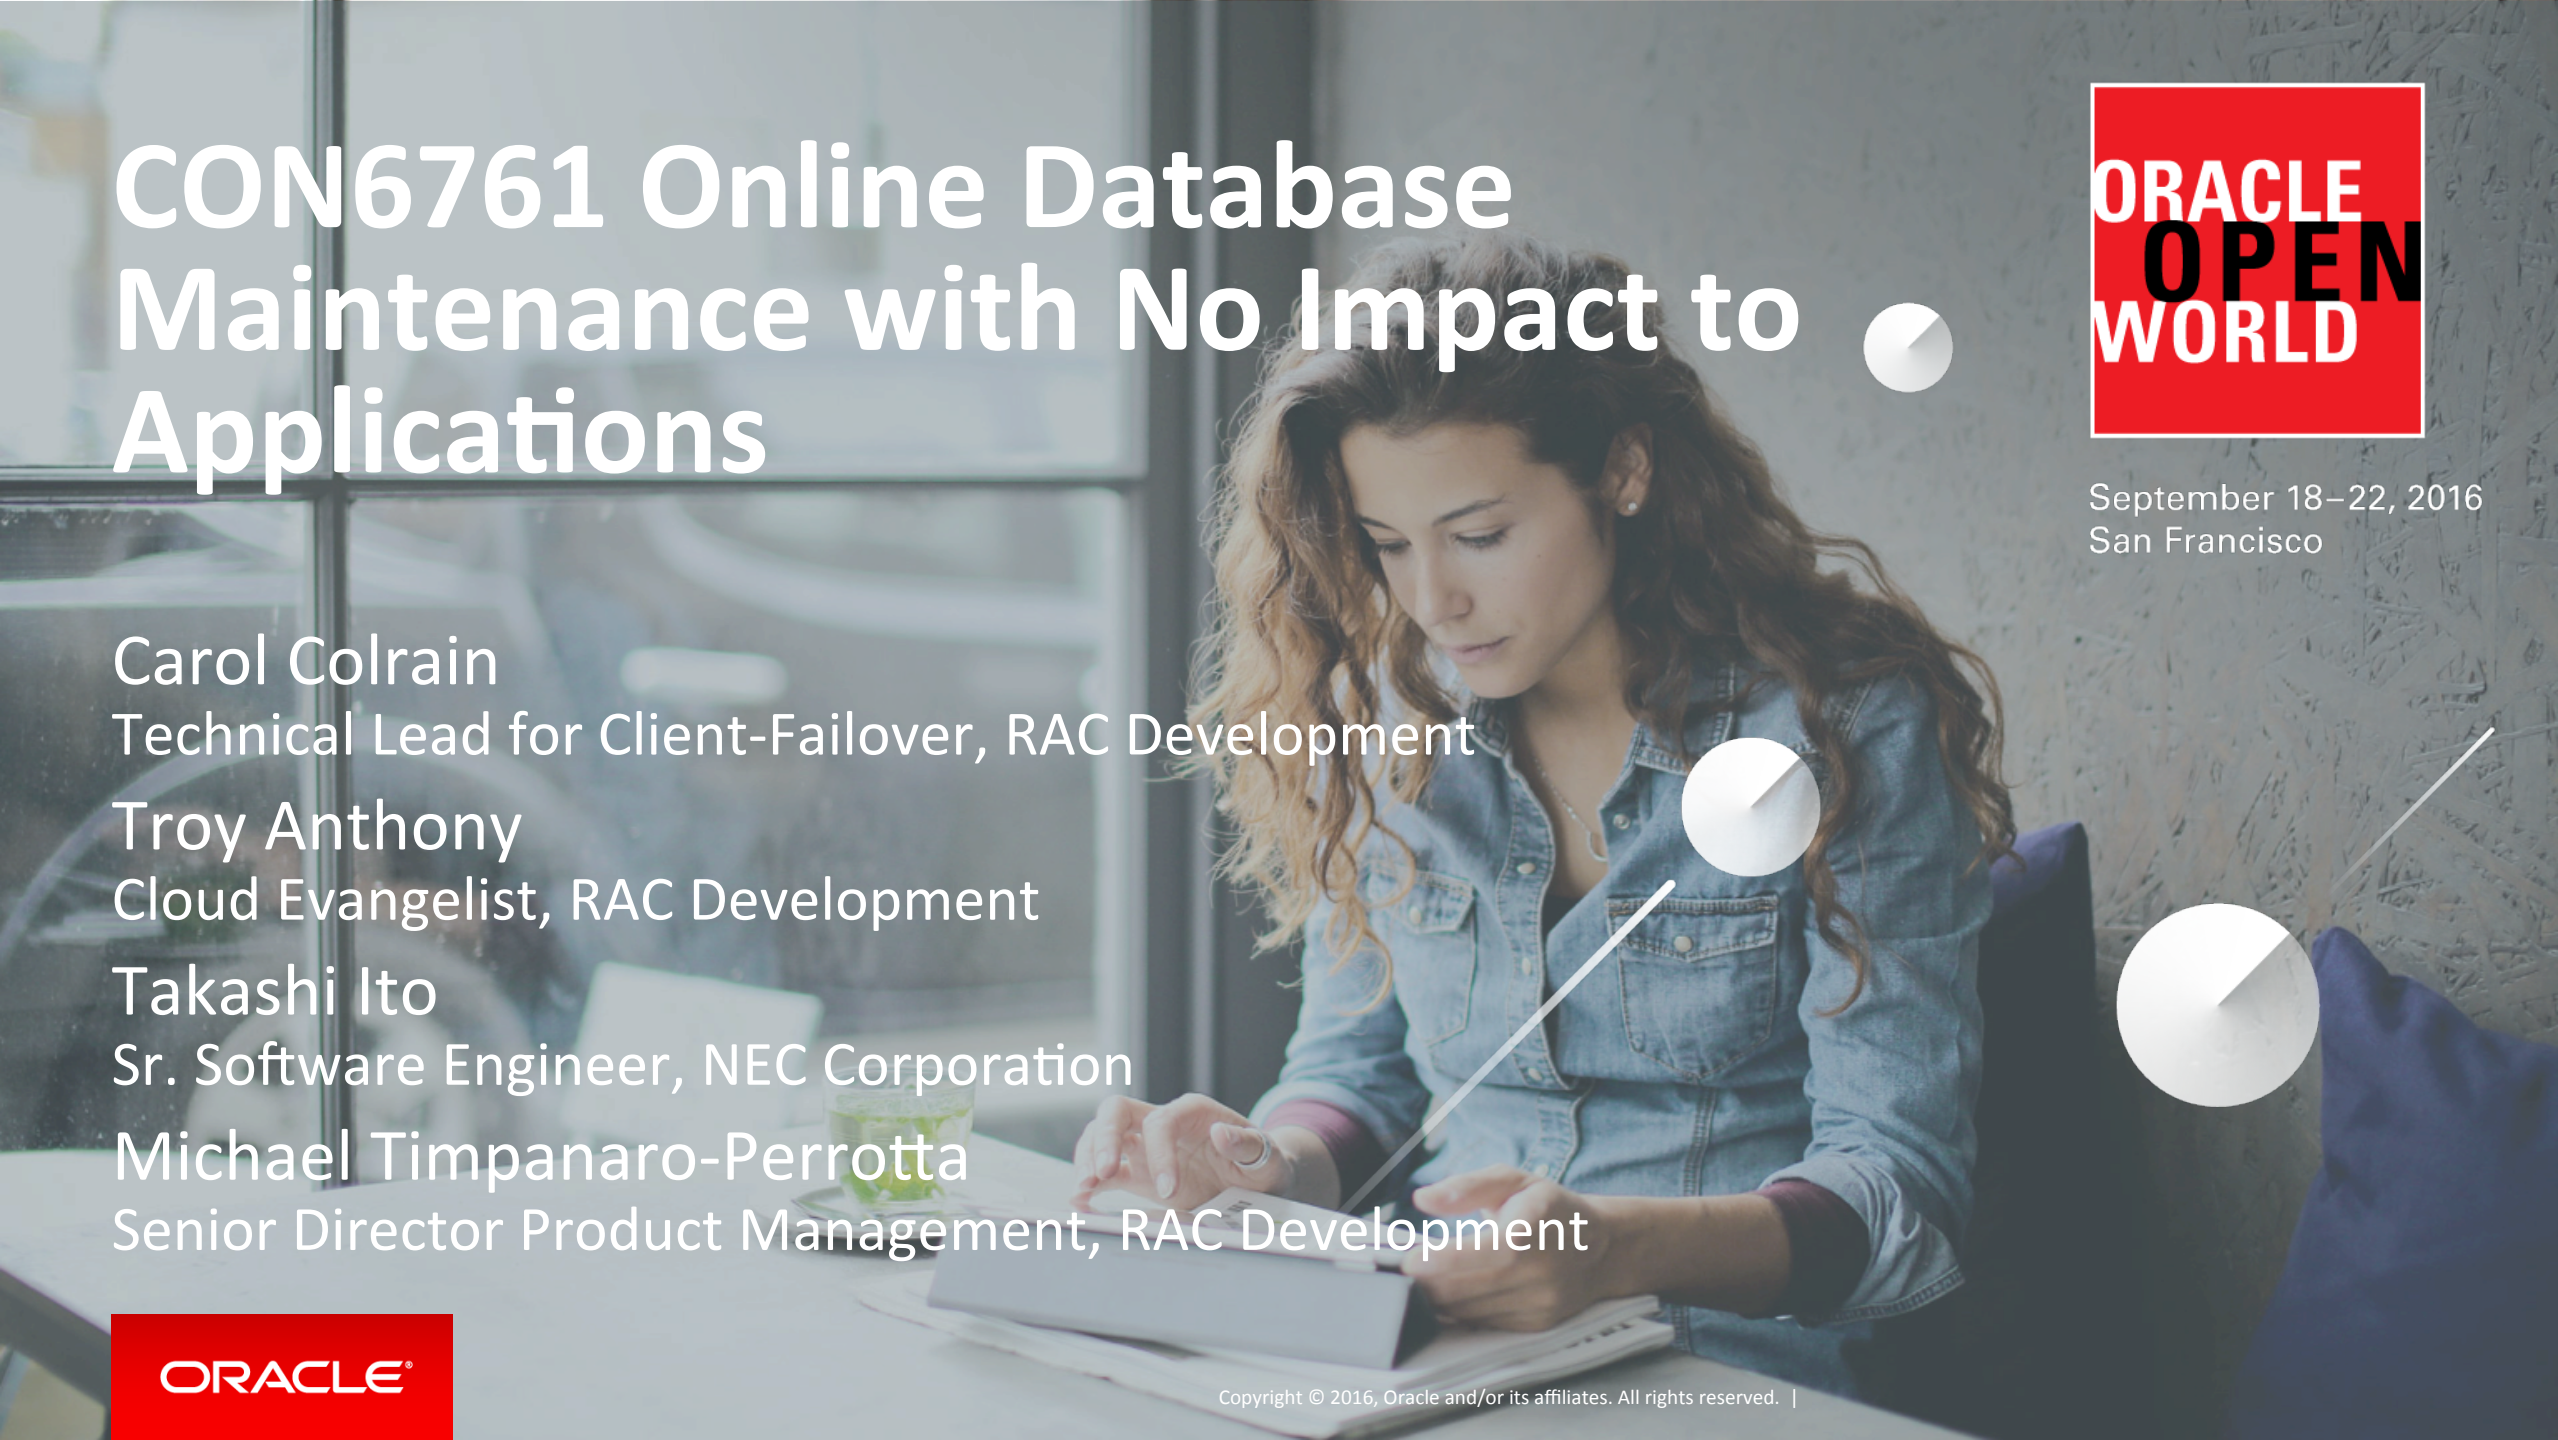

ORACLE OPEN WORLD

September 18–22, 2016
San Francisco

Carol Colrain
Technical Lead for Client-Failover, RAC Development

Troy Anthony
Cloud Evangelist, RAC Development

Takashi Ito
Sr. Software Engineer, NEC Corporation

Michael Timpanaro-Perrotta
Senior Director Product Management, RAC Development

ORACLE

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

Announcing Oracle Database 12c Release 2 on Oracle Cloud

- Available now
  - Exadata Express Cloud Service
- Coming soon
  - Database Cloud Services
  - Exadata Cloud Machine

Oracle is presenting features for Oracle Database 12c Release 2 on Oracle Cloud. We will announce availability of the On-Prem release sometime after Open World.

Program Agenda

1. How do we schedule server maintenance?
2. Are the building blocks in place ?
3. Draining Solutions
4. Older Style Applications
5. Set and Forget
6. NEC Success Story

1 How do we schedule server maintenance ?

Users should see no errors during maintenance

Preventable Situation

There is no reason for users to see downtime during scheduled database maintenance

- Service is unavailable
- Application owners unable to agree maintenance windows
- Long running jobs see errors
- DBA's and engineers work off hours
- Application and middleware components need to be restarted

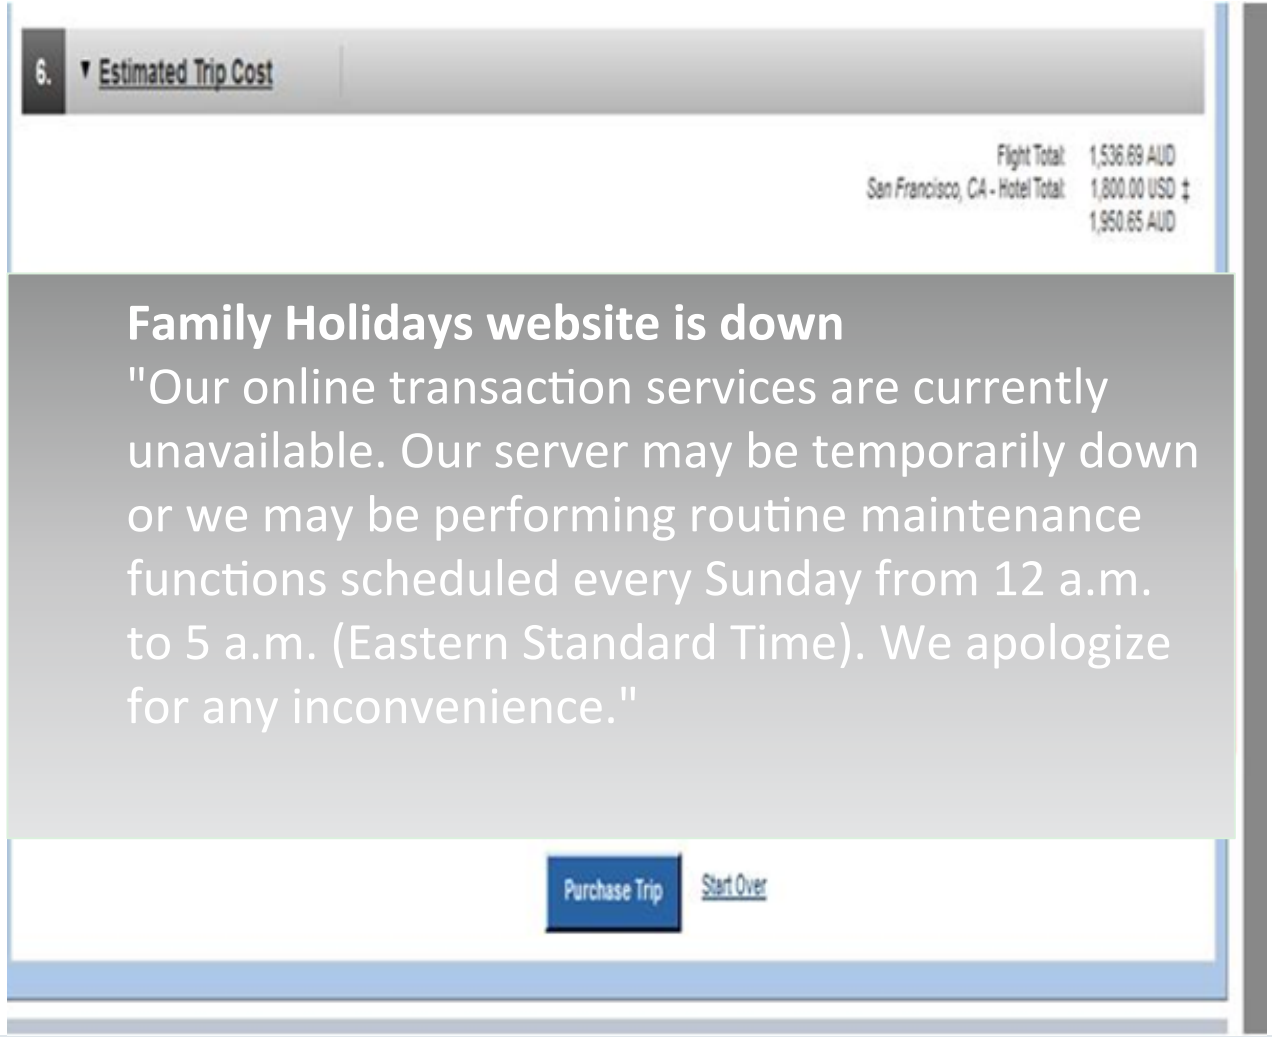

How do we solve for all your applications?

- Move work to different instance/database with no errors reported to applications
- Transparent to applications and mid-tiers
- Support all server side maintenance patches, PSUs, repairs, changes, major release, unplug/plug, migration, expansion, h/w replacement
- Configure once, same for all commands

High Availability by Patch Type

	One- Off	PSU/CPU	Bundle Patch	Patch Set
RAC Rolling	96%	All	Most	No
DG Standby First	98%	All	All	No
Online - Hot	82%*	No	No	No
GoldenGate	All	All	All	All

2 Are your building blocks in place ?

Affinity and Draining Locally - Switchover across Sites

Production Site

RAC
- *Online Rolling Maintenance*
- *Scalability*
- *Server HA*

RAC One
- *Online Rolling Maintenance*
- *Server HA*

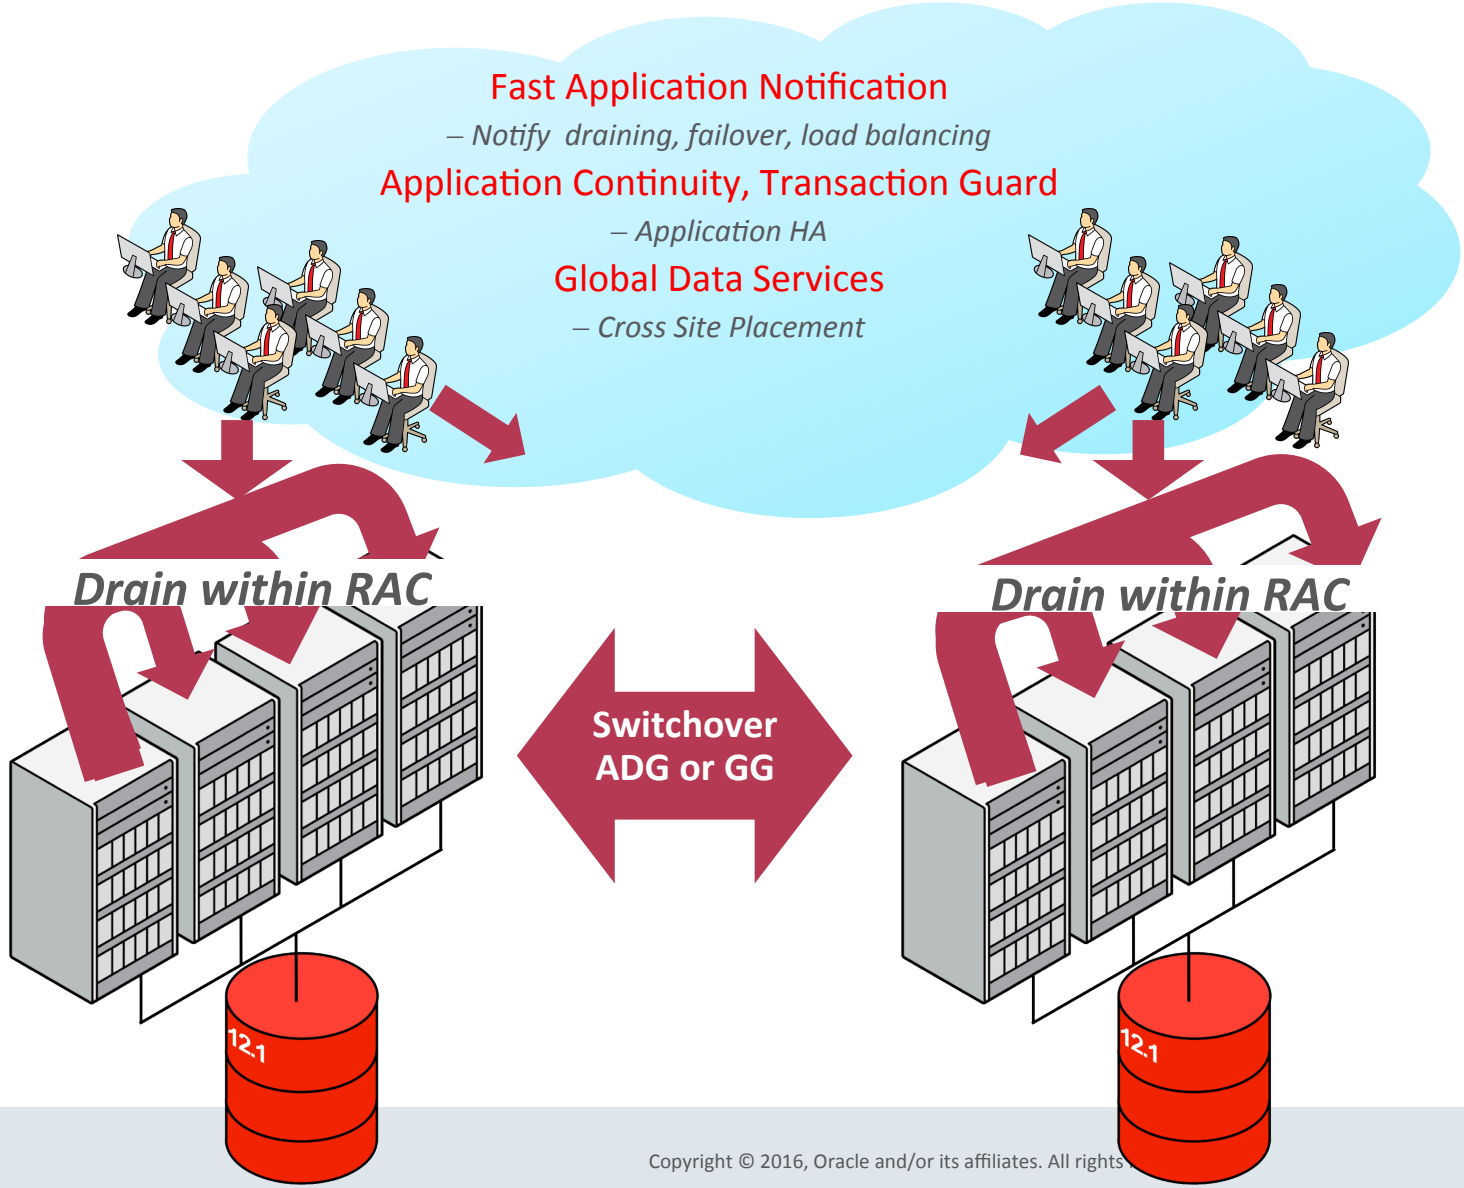

Replicas

Active Data Guard
- *Scheduled switchover*
- *Data Protection, DR*
- *Query Offload*

Data Guard
- *Scheduled switchover*
- *Data Protection, DR*

GoldenGate
- *Scheduled switchover*
- *Active-active replication*
- *Heterogeneous*

Sharding
- *Massive OLTP*
- *Scheduled switchover*
- *Active-active replication*
- *Heterogeneous*

Oracle RAC One Node and Oracle RAC

Drain services to another active instance gradually

Oracle RAC *One Node*: one instance per database running at a time, **both instances during maintenance**

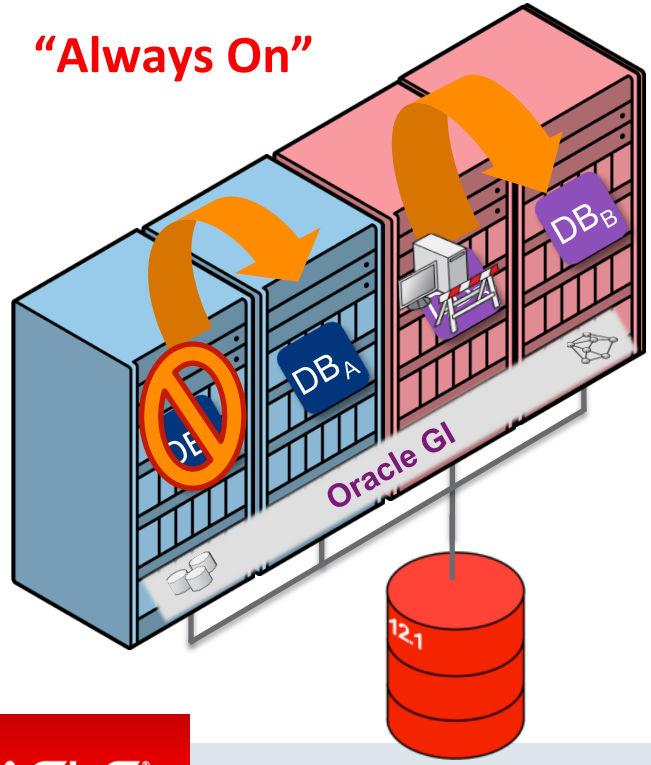

Oracle RAC: **multiple instances per database running concurrently**

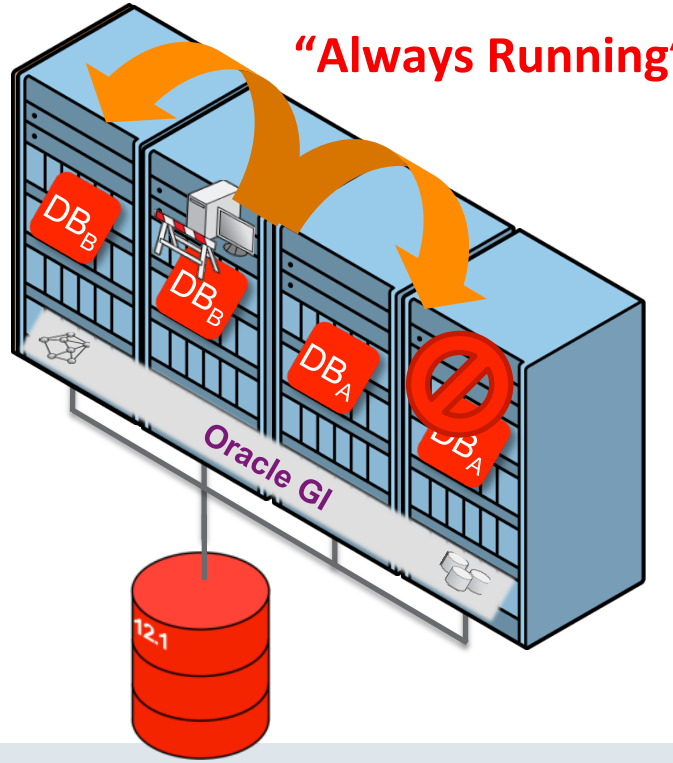

12cR2 Oracle Active Data Guard

Continuous Service Availability

Drain primary and switch

```
Switchover to
<db_resource_name> [wait]
```

ADG sessions survive standby role change

Application Continuity recovers stragglers

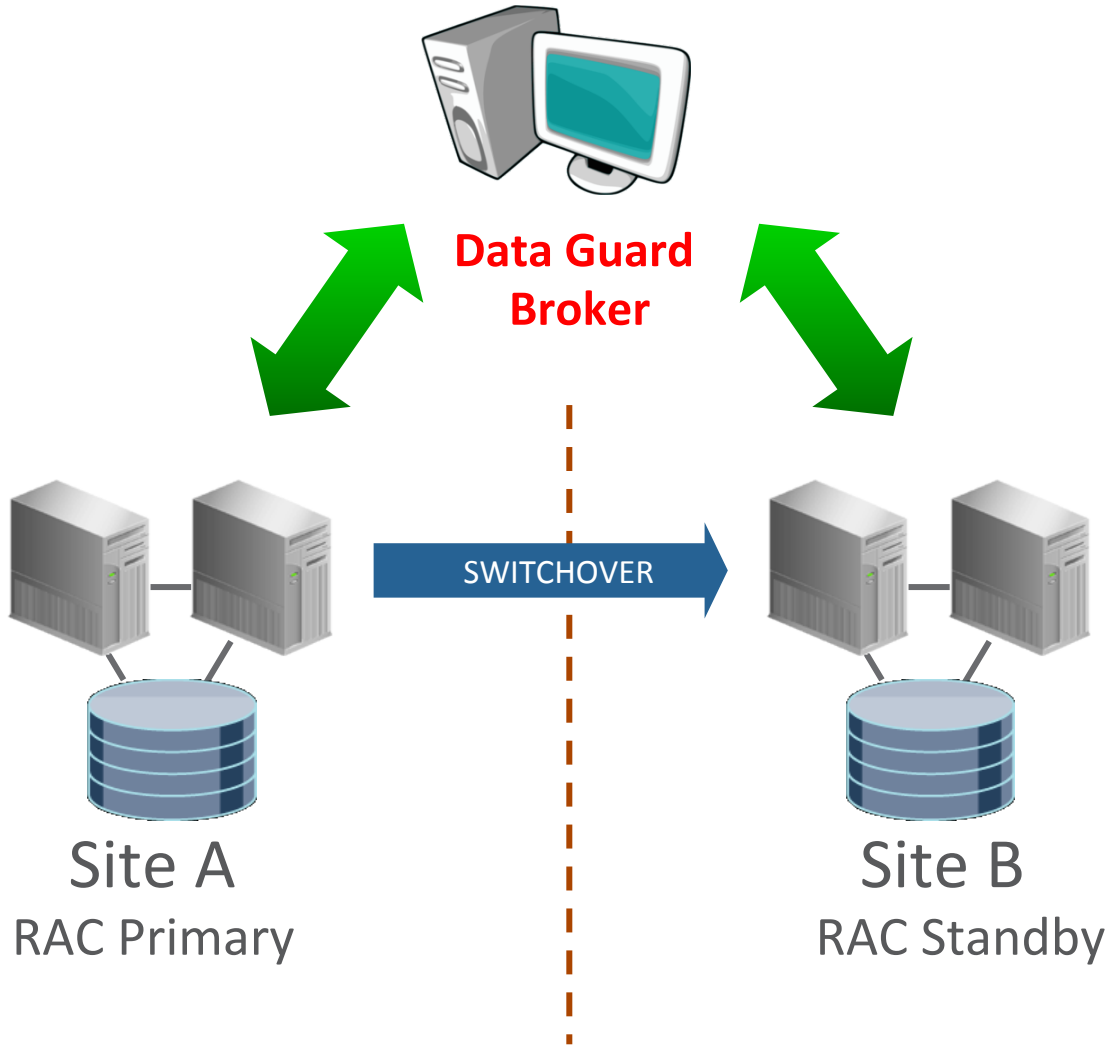

12cR2 Oracle Multitenant

Continuous Service Availability

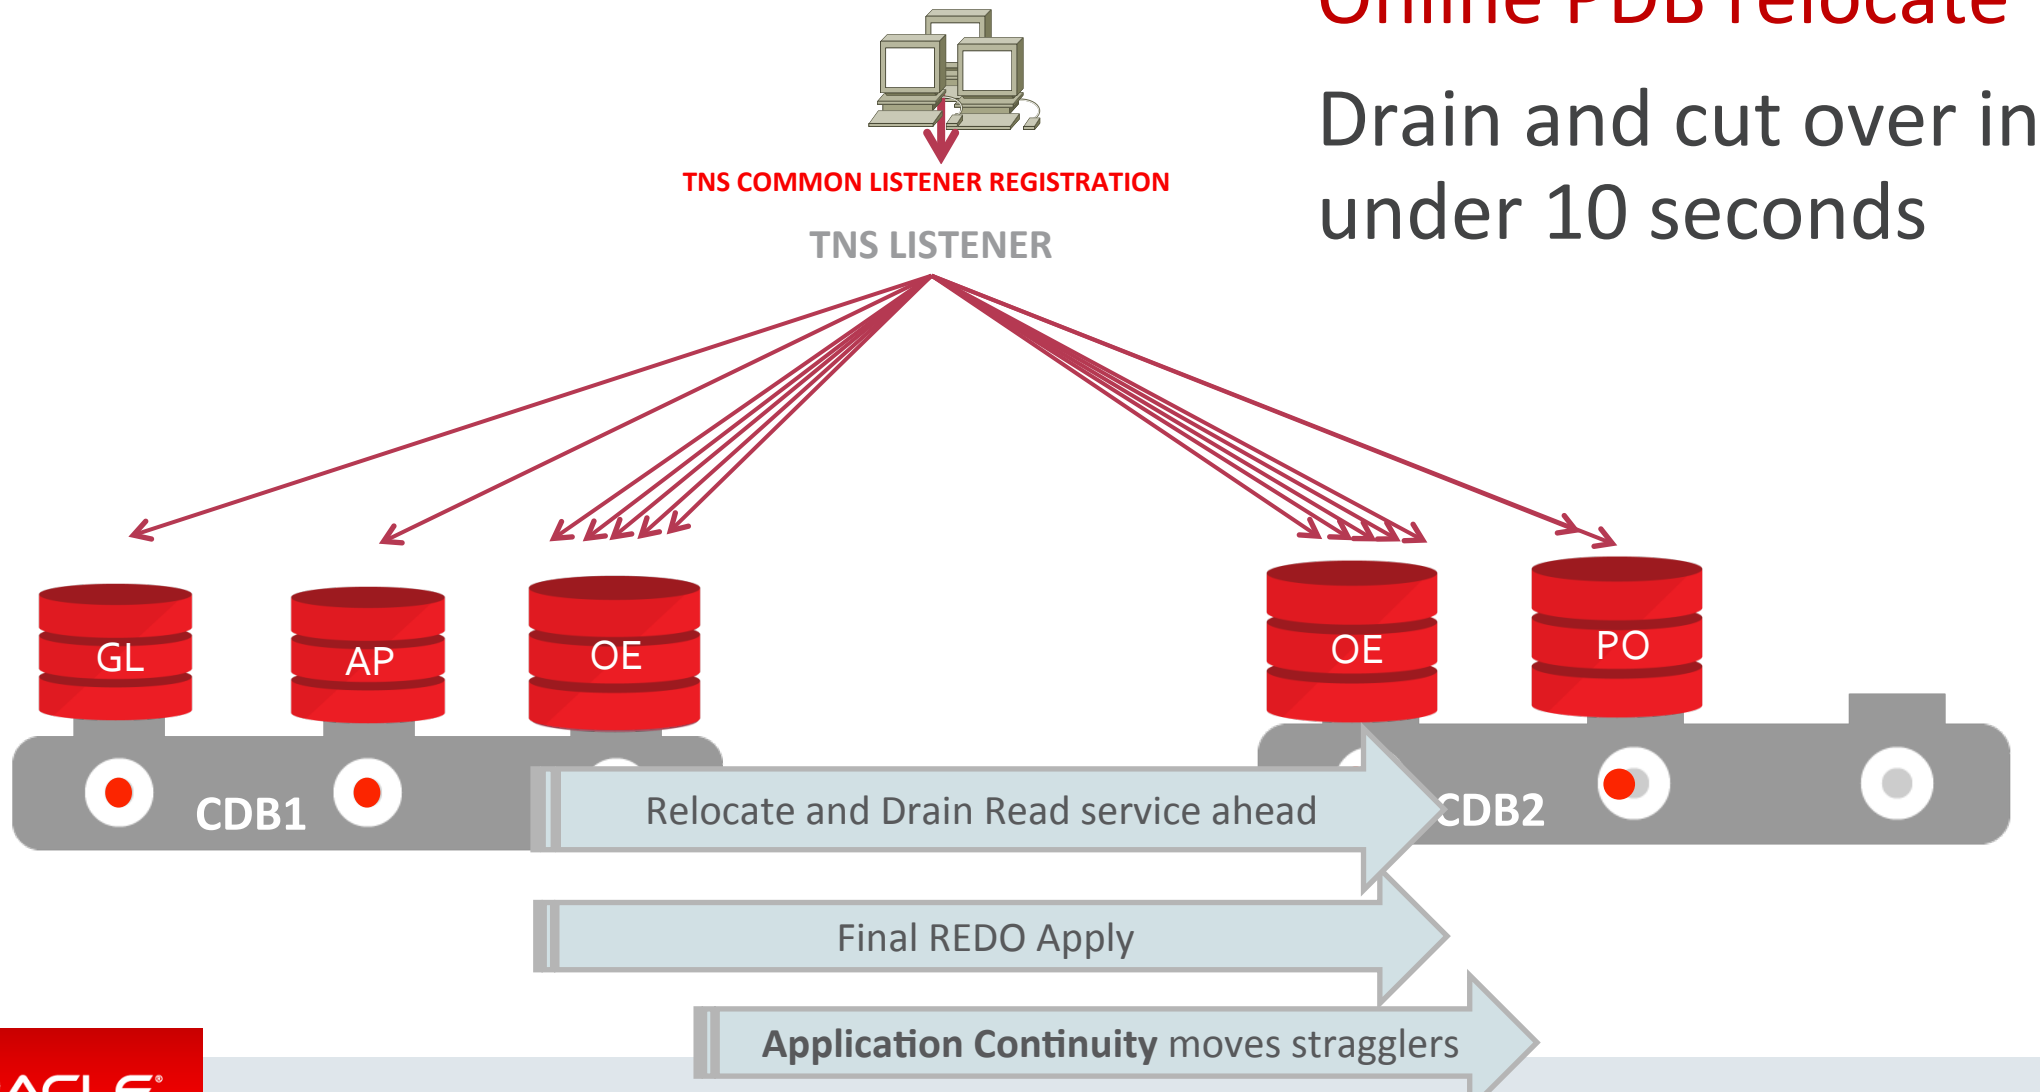

Online PDB relocate

Drain and cut over in under 10 seconds

Use Services for Location Transparency

Services provide a “dial in number” for your application

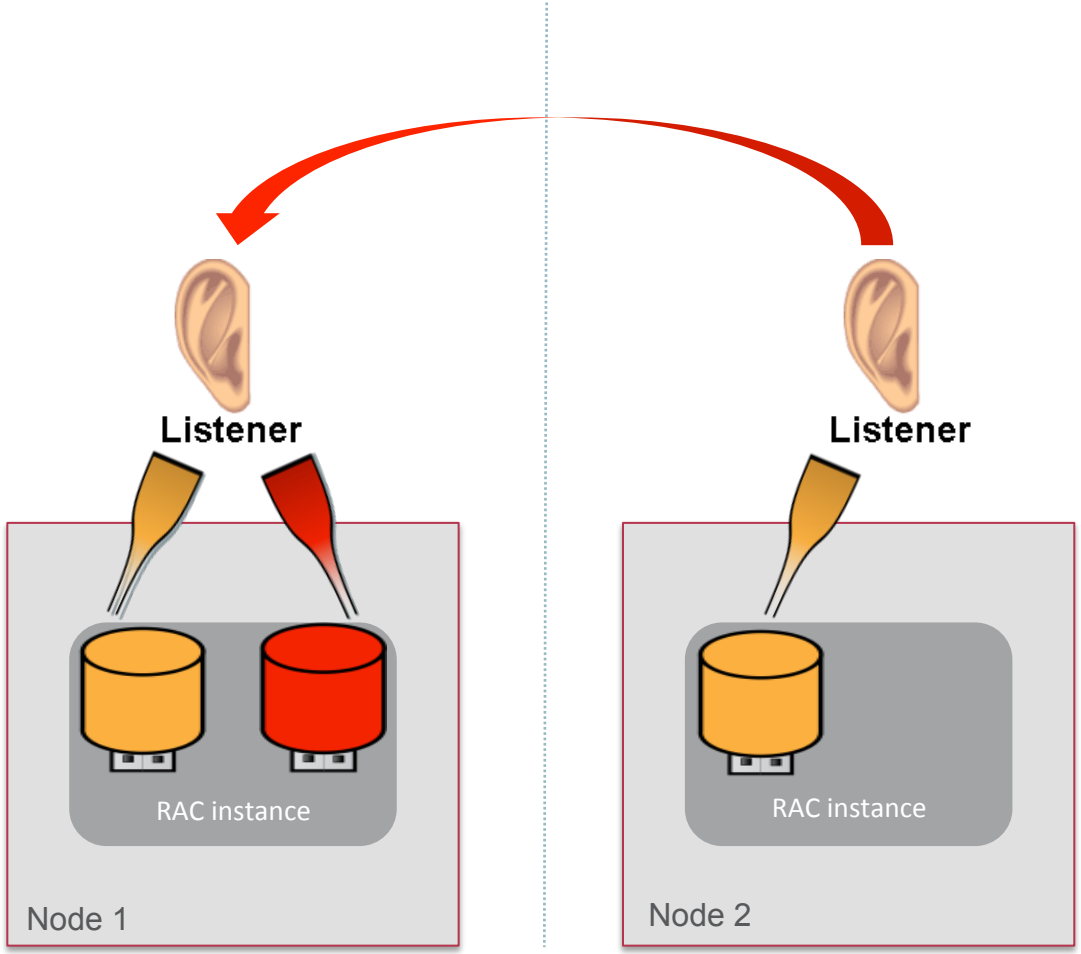

- Regardless of location, application keeps the name
- Moving, reshaping, prioritizing controls how a service is offered
- Batch and OLTP separated
- DB and PDB names for admin only

Are connections properly configured?

Automatic Retries

```
alias =(DESCRIPTION =
    (CONNECT_TIMEOUT=90) (RETRY_COUNT=20)(RETRY_DELAY=3)
        (TRANSPORT_CONNECT_TIMEOUT=3)
     (ADDRESS_LIST =
         (LOAD_BALANCE=on)
         ( ADDRESS = (PROTOCOL = TCP)(HOST=primary-scan)(PORT=1521)))
     (ADDRE
         (LO
         ( AD                                    can)(PORT=1521)))
    (CONNE
```

Configure in One Place
LDAP or TNS names

ALWAYS use a service that is NOT DB/PDB name

Align the Timeouts

NO interruption when Active/ Active RAC

Application Timeout

- **Drain + Switch to DG**

DG Broker

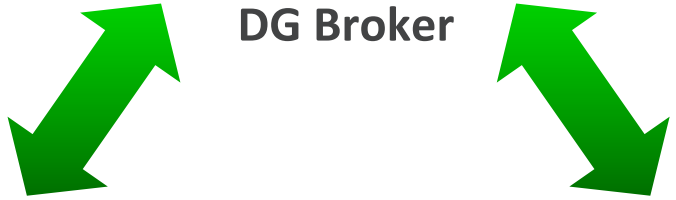

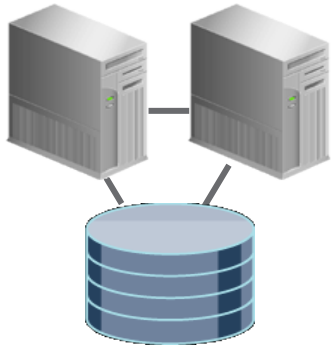

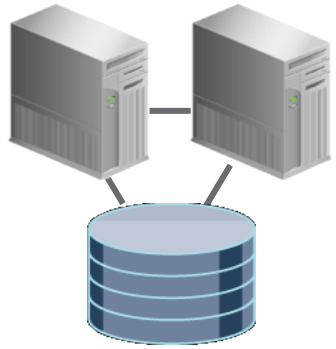

Stop or Relocate Service

Drain Work

RAC Primary

RAC Standby

Drain Work

Switchover

3 **You are now ready to drain the work**

Applications should see no errors during maintenance.

Roll Maintenance - Drain work at safe places

Move services in advance

Replace connections where applications don't notice

- Connection pools
- Connection tests
- Transaction boundaries

Continue on a new connection with states restored

DBA steps – Works Today

Repeat for each service allowing time to drain

- **Stop service (no –force)**

```
srvctl stop service –db .. -instance .. -service .. (omit service for all)
```

- or **Relocate service (no –force)**

```
srvctl relocate service –db .. -service .. –oldinst .. –newinst
srvctl relocate service –db .. -service .. –currentnode.. –targetnode
```

- **Wait for sessions and transactions to drain**
- **For remaining sessions, stop transactional**

```
exec dbms_service.disconnect_session( '… your service ..',
 DBMS_SERVICE.POST_TRANSACTION);
```

- **Stop the instances using your preferred tool; option to disable**

Draining with Oracle Connection Pools – Works Today

Applications using …	**Oracle – WebLogic Active GridLink, UCP, OCI, ODP.NET managed and unmanaged, OCI Session Pool, Tuxedo** 3rd party App Servers using UCP: IBM WebSphere, Apache Tomcat, NEC WebOTX
DBA Step	**srvctl [relocate|stop] service (no –force)**
Sessions Drain	Immediately new work is redirected Gradually Active sessions are released when returned to pools

FAN Planned

Pools drain sessions as work completes

FAN with other Java-Based App Servers – Works Today

Use UCP – a simple Data Source replacement

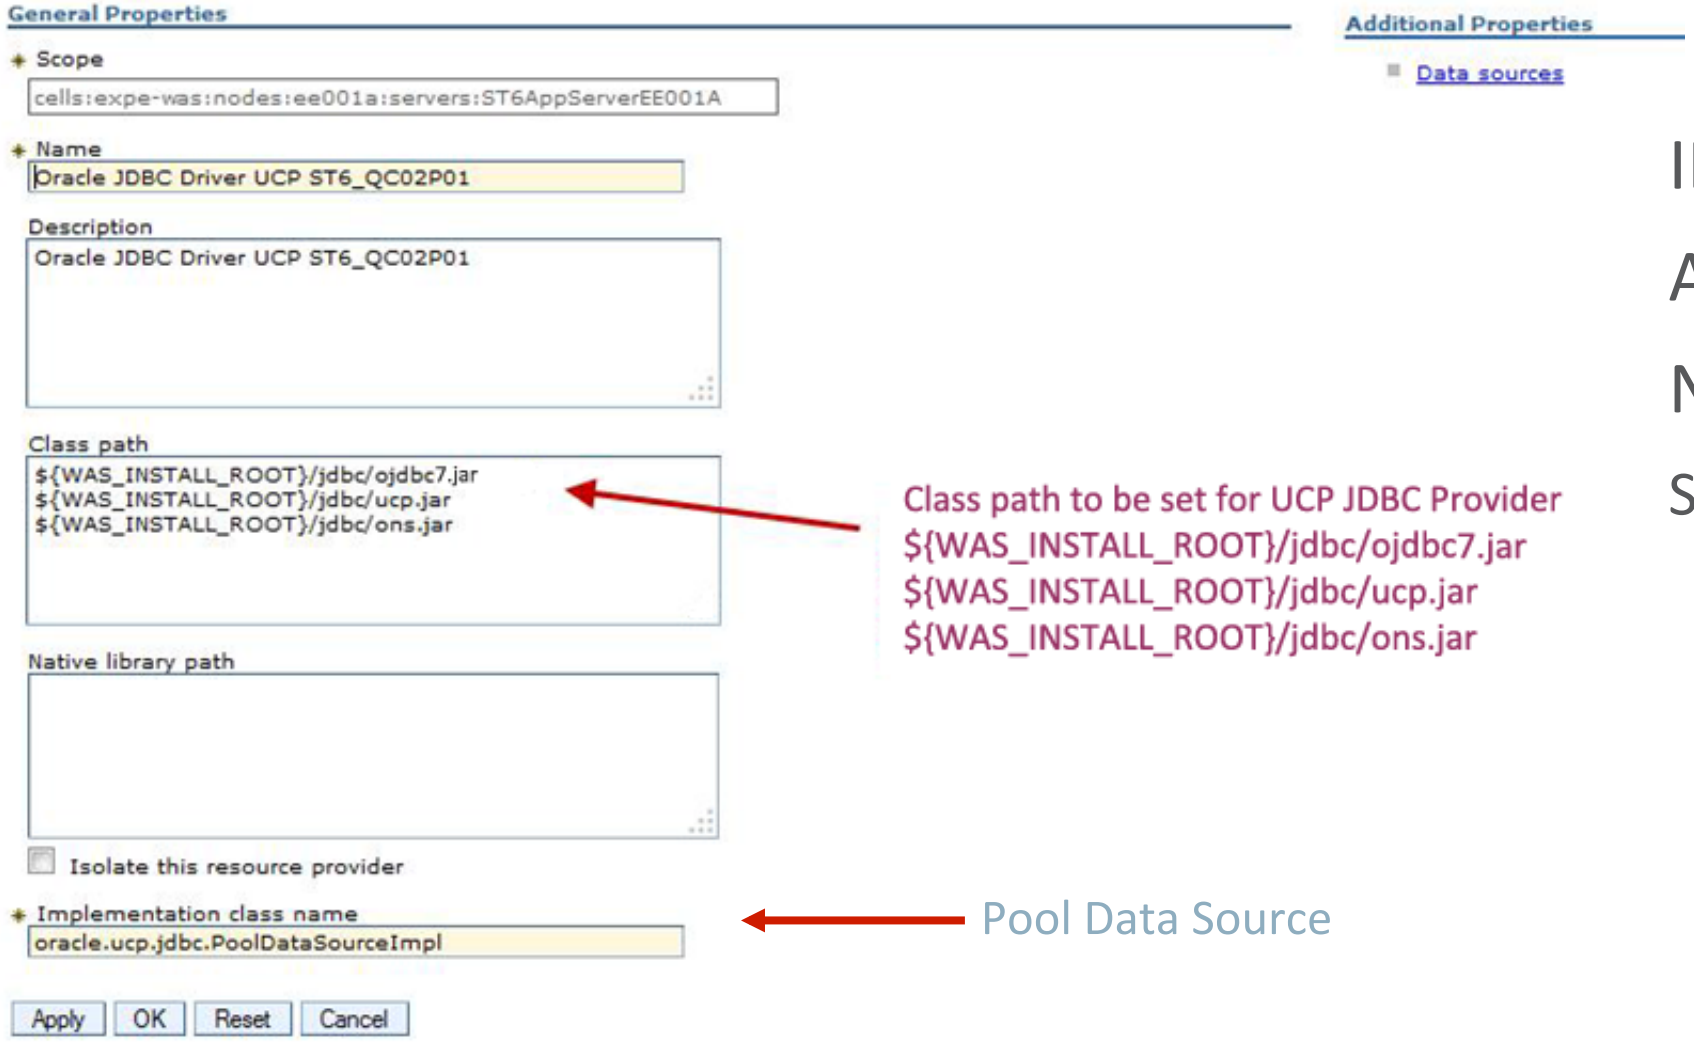

Class path to be set for UCP JDBC Provider
${WAS_INSTALL_ROOT}/jdbc/ojdbc7.jar
${WAS_INSTALL_ROOT}/jdbc/ucp.jar
${WAS_INSTALL_ROOT}/jdbc/ons.jar

Pool Data Source

IBM WebSphere

Apache Tomcat

NEC WebOTX

See OTN.

12*c* FAN: Standardized, Auto-Configured (**HIDDEN**)

All Oracle use ONS

- JDBC Universal Connection Pool
- JDBC Thin Driver **(12.2)**
- OCI/OCCI driver
- ODP.NET Unmanaged Provider (OCI)
- ODP.NET Managed Provider (C#)
- OCI Session Pool
- WebLogic Active GridLink
- Tuxedo
- Listener

Auto-Configured

```
DESCRIPTION =
    (CONNECT_TIMEOUT=90)
    (RETRY_COUNT=20)(RETRY_DELAY=3)
    (TRANSPORT_CONNECT_TIMEOUT=3)
    (ADDRESS_LIST =
        (LOAD_BAL
        ( ADDRESS = (PROTOCOL = TCP)
        (HOST=primary-scan) (PORT=1521)))
    (ADDRESS_LIST =
        (LOAD_BA
        ( ADDRESS = (PROTOCOL = TCP)
        (HOST=second-scan)(PORT=1521)))
(CONNECT_DATA=(SERVICE_NAME=gold)))
```

ONS Node Set 1

ONS Node Set 2

Keep it simple – Stagger draining services

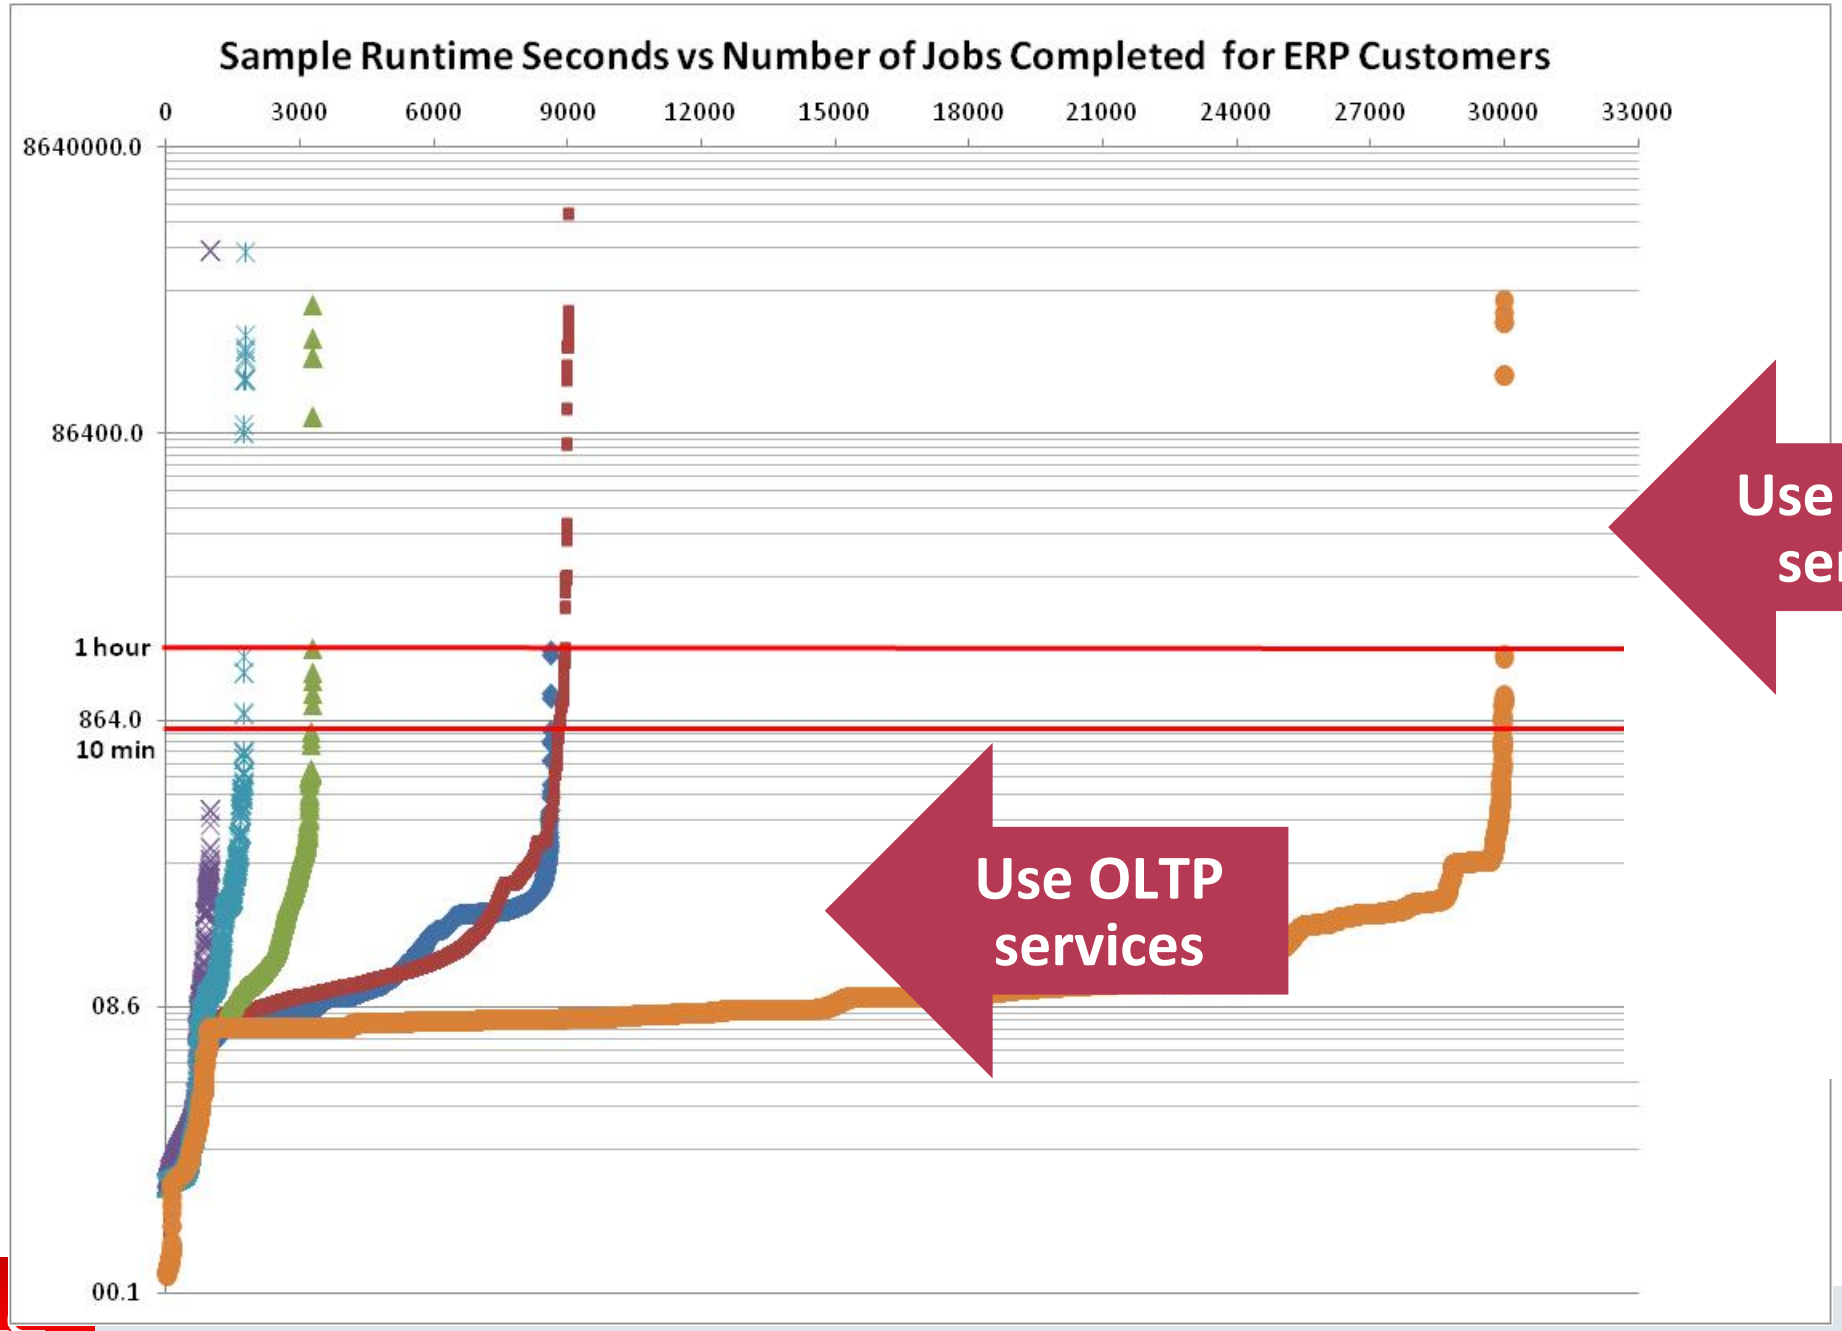

Stagger Draining – Batch then Online Services

- Use at least two application services – online and batch/backend
- Start draining batch early
- Online drains in seconds

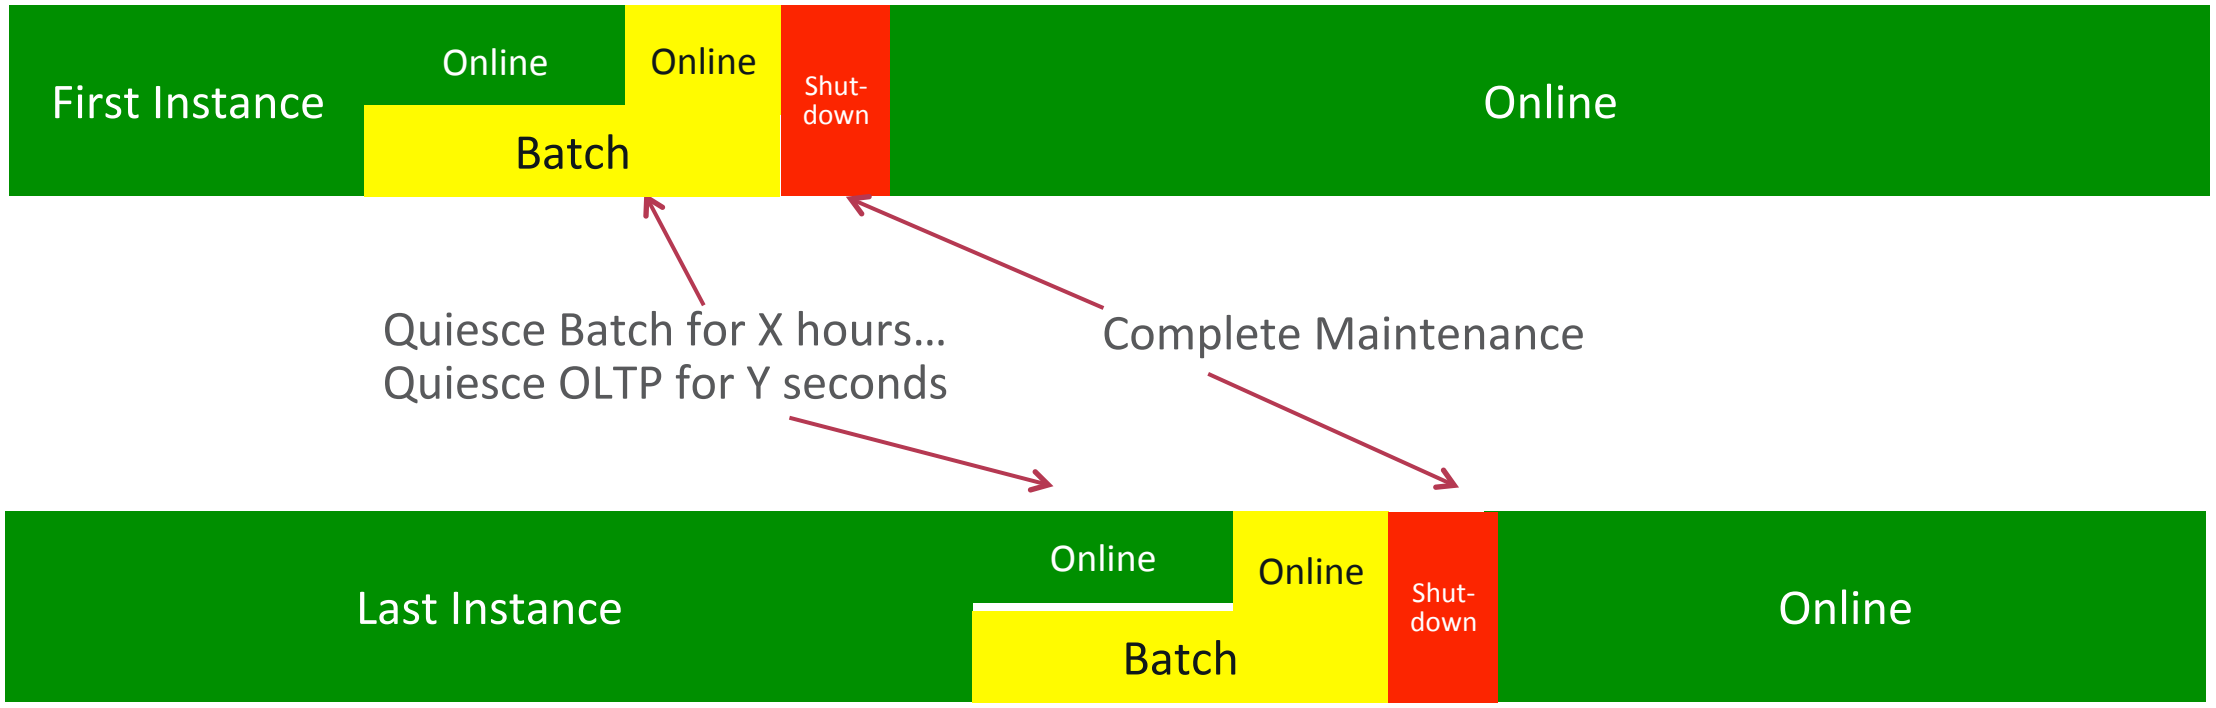

Eliminate Logon Storms

12c Drains Gradually

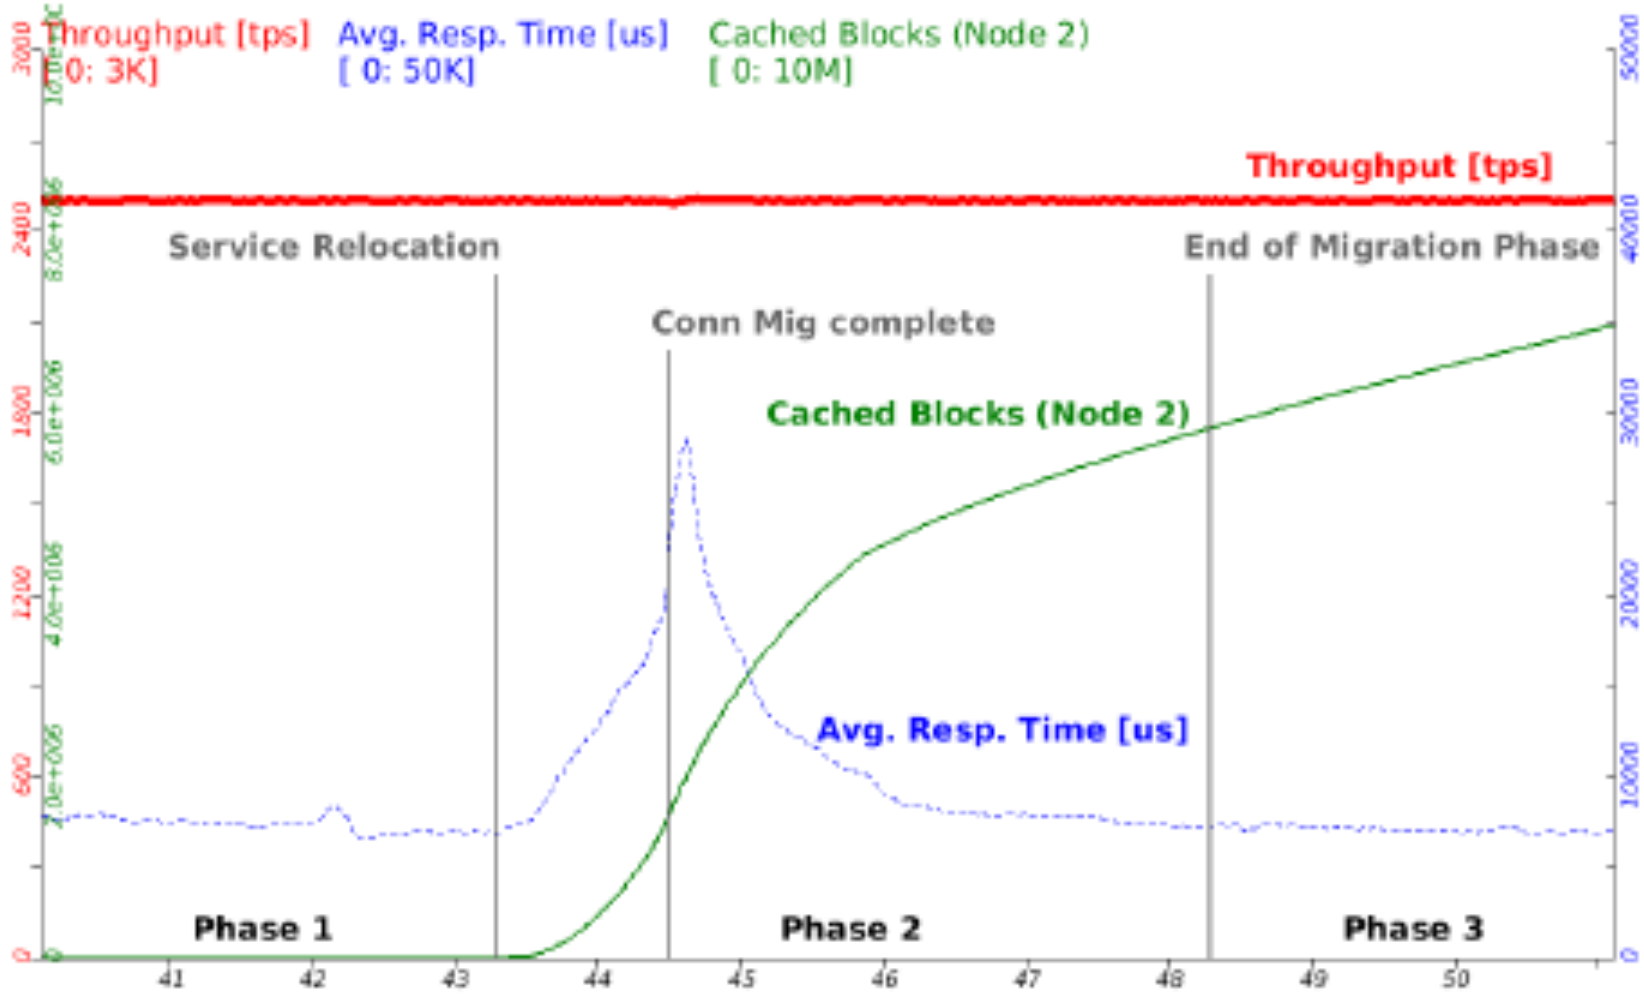

12cR2 FAN Use FAN to Invoke Draining

Supports 11.2, 12.1, and later databases

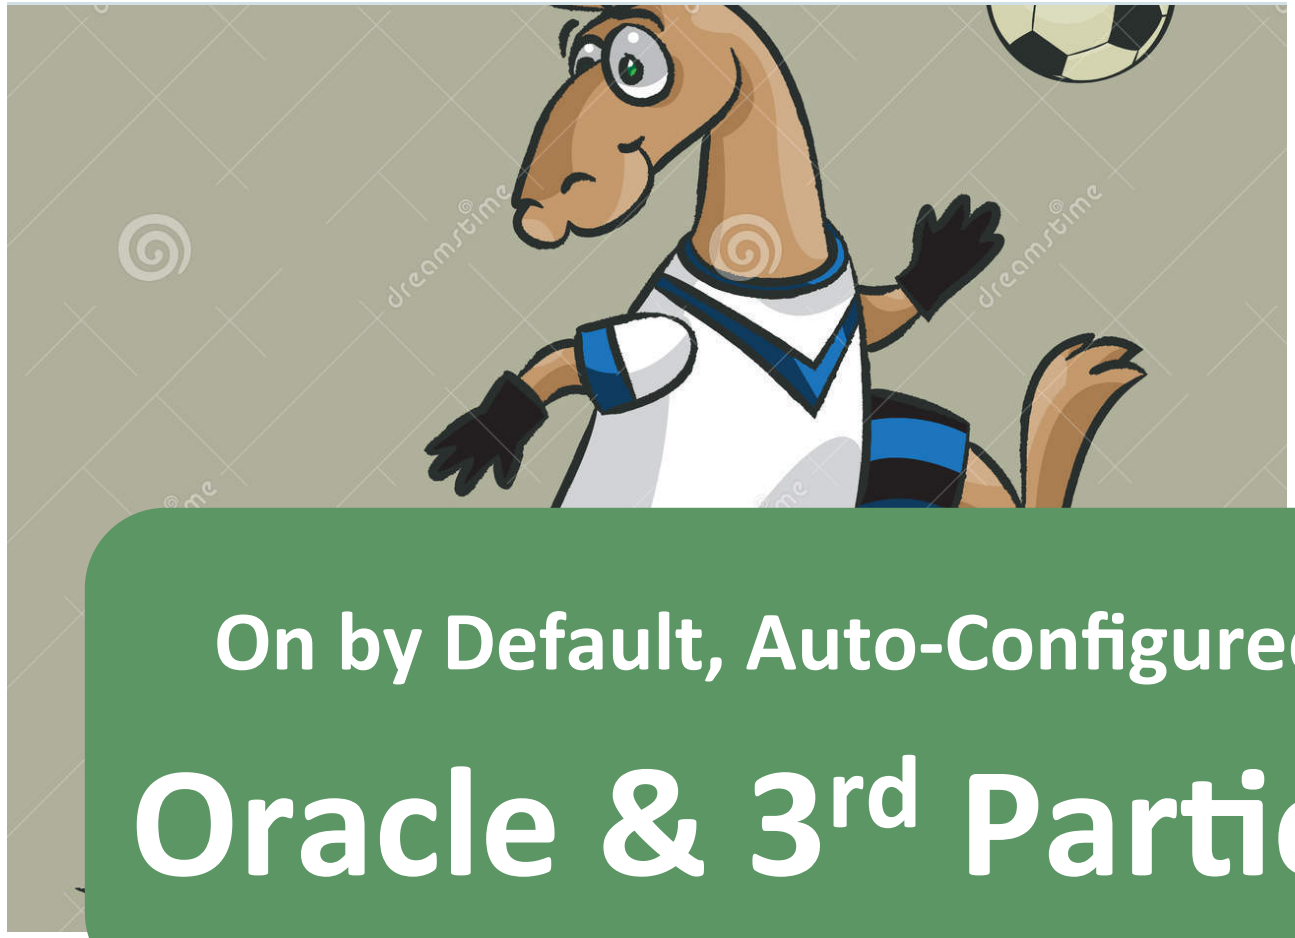

On by Default, Auto-Configured

Oracle & 3rd Parties

All Applications

Drivers detect, break-out, and drain

12cR2 FAN Enable Connection Tests to Drain

Supports 11.2, 12.1, and later databases

When asked,
" its not good"`

- isValid
- isUsable
- isClosed
- pingDatabase
- SQL Tests
- OCI_ATTR_SERVER_STATUS

Drain Application Servers at Safe Places

Application Server	Test Name	Connection Test to DB
Oracle WebLogic – Generic and Multi data sources	TestConnectionsOnReserve TestConnectionsOnCreate	isUsable SQL – SELECT 1 FROM DUAL
Oracle WebLogic Active GridLink	embedded	isUsable
IBM WebSphere	PreTest Connections	SQL - SELECT USER FROM DUAL
RedHat JBoss	check-valid-connection-sql	SQL - SELECT COUNT(*) FROM DUAL
Apache TomCat	TestonBorrow TestonRelease	SQL - SELECT 1 FROM DUAL

Drain Applications at Safe Places

Application	Condition	Connection Test to DB
eBusiness Suite	Connection borrowed from Weblogic	TestConnectionsOnReserve with "BEGIN NULL; END;"
Fusion Applications	Connection returned to WebLogic and C++ pools and checked	TestConnectionsOnReserve with isValid OCI_ATTR_SERVER_STATUS
Siebel	EAI (batch) checks status before starting	OCI_ATTR_SERVER_STATUS
Customer	Custom pool with Meta data table Checks status every 60s	OCI_ATTR_SERVER_STATUS

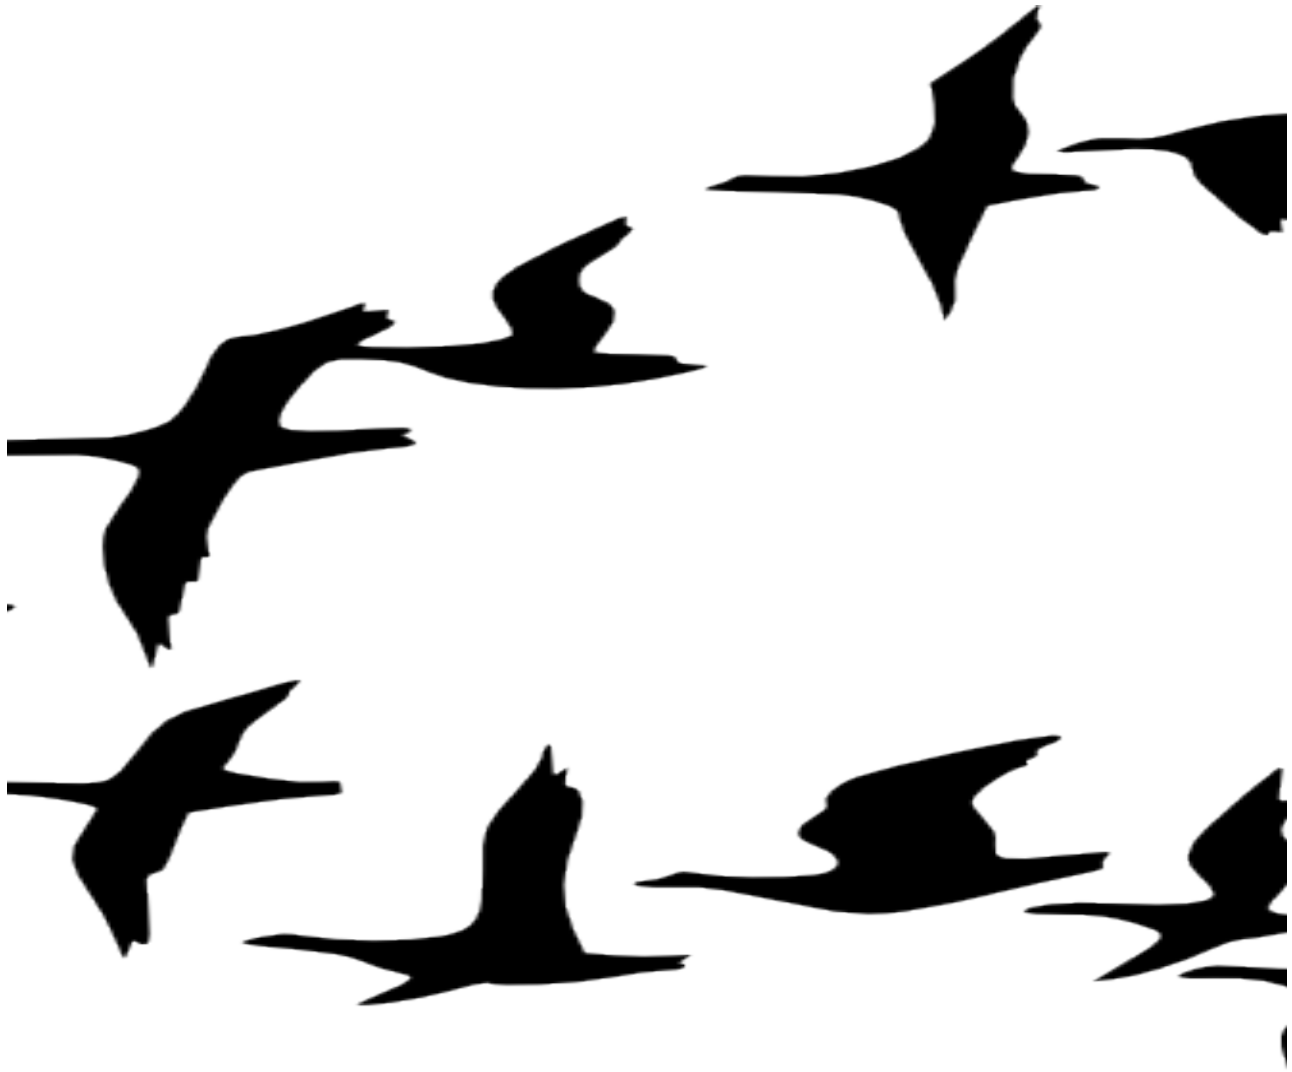

Planned Draining Demonstration

4 DBA Operation's Simplified

12cR2 One command for maintenance operations

Focus on user experience

- **Group operations** at pdb, instance, node and database for services
- New service attributes set once

  drain_timeout (seconds)

  stopoption (immediate, transactional)

5 Transactional Disconnect

For your older style apps

Some Applications have Read-Only Failover Built-In

Application	Set and Forget
Siebel	TNS + STOPOPTION = Transactional Drain_Timeout
PeopleSoft	
JD Edwards	
Informatica	

FAN on for unplanned

12cR2 TAF SELECT with Transaction Guard

Hides scheduled maintenance and unplanned for read-only apps that set state at start

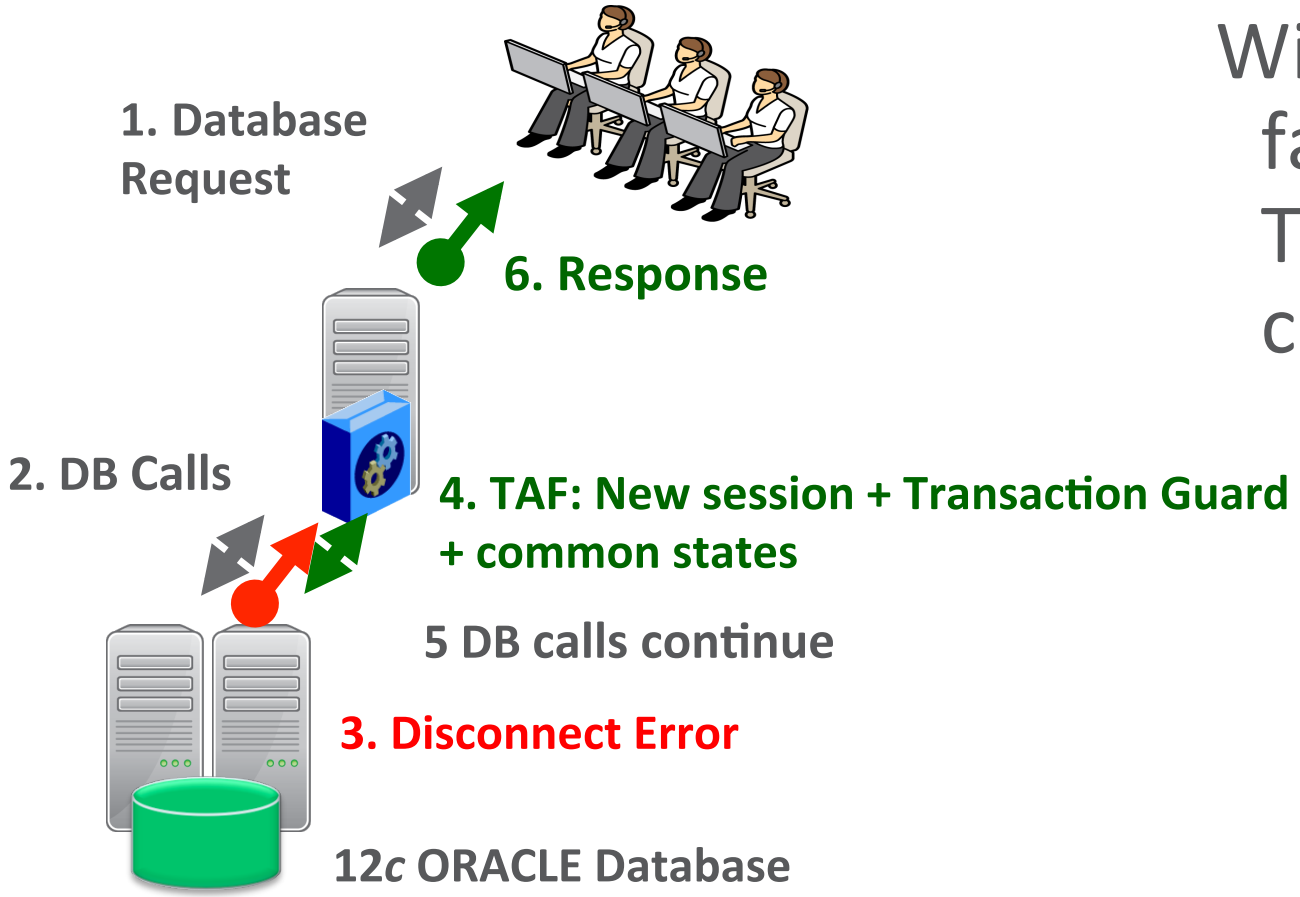

With stopoption transactional, fails over, validates with Transaction Guard, resets common states

12cR2 Application Continuity for Longer Running

- Applications that borrow a connection and do not return until completed
- Purges committed transactions as it goes
- Stop option Transactional, Drain Timeout “n
- Rules apply

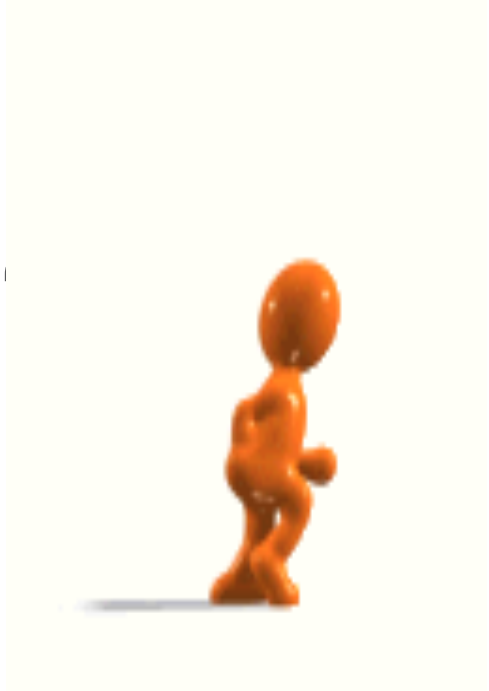

Application Continuity – Long Running

Hides scheduled maintenance and unplanned for static applications

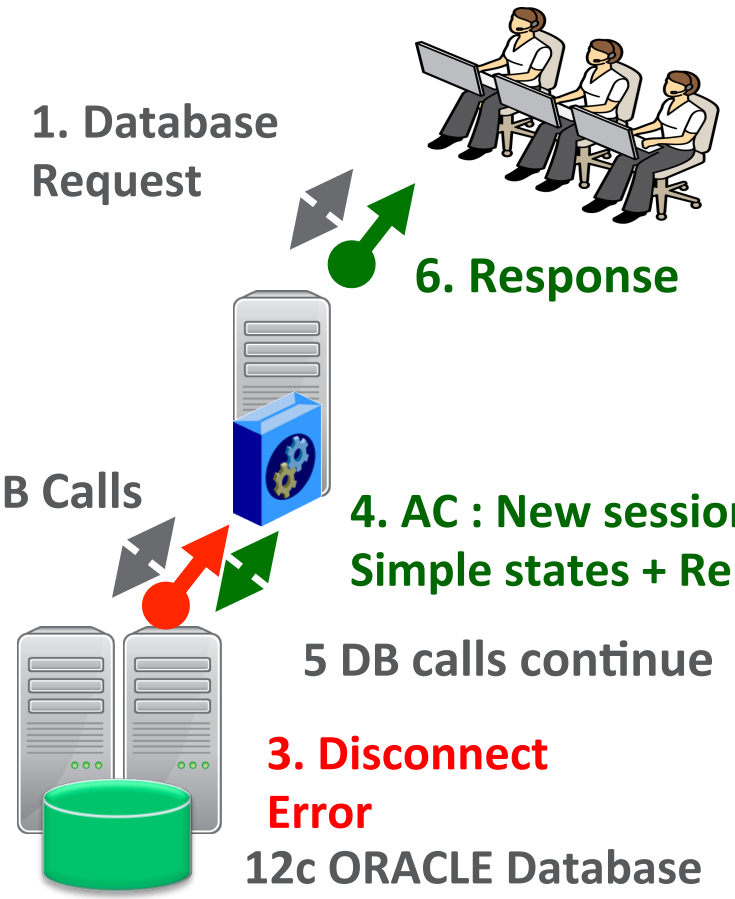

At disconnect transactional, fails over, validates with TG, syncs states, replays

Appears to applications as a delayed execution.

6

Set and Forget

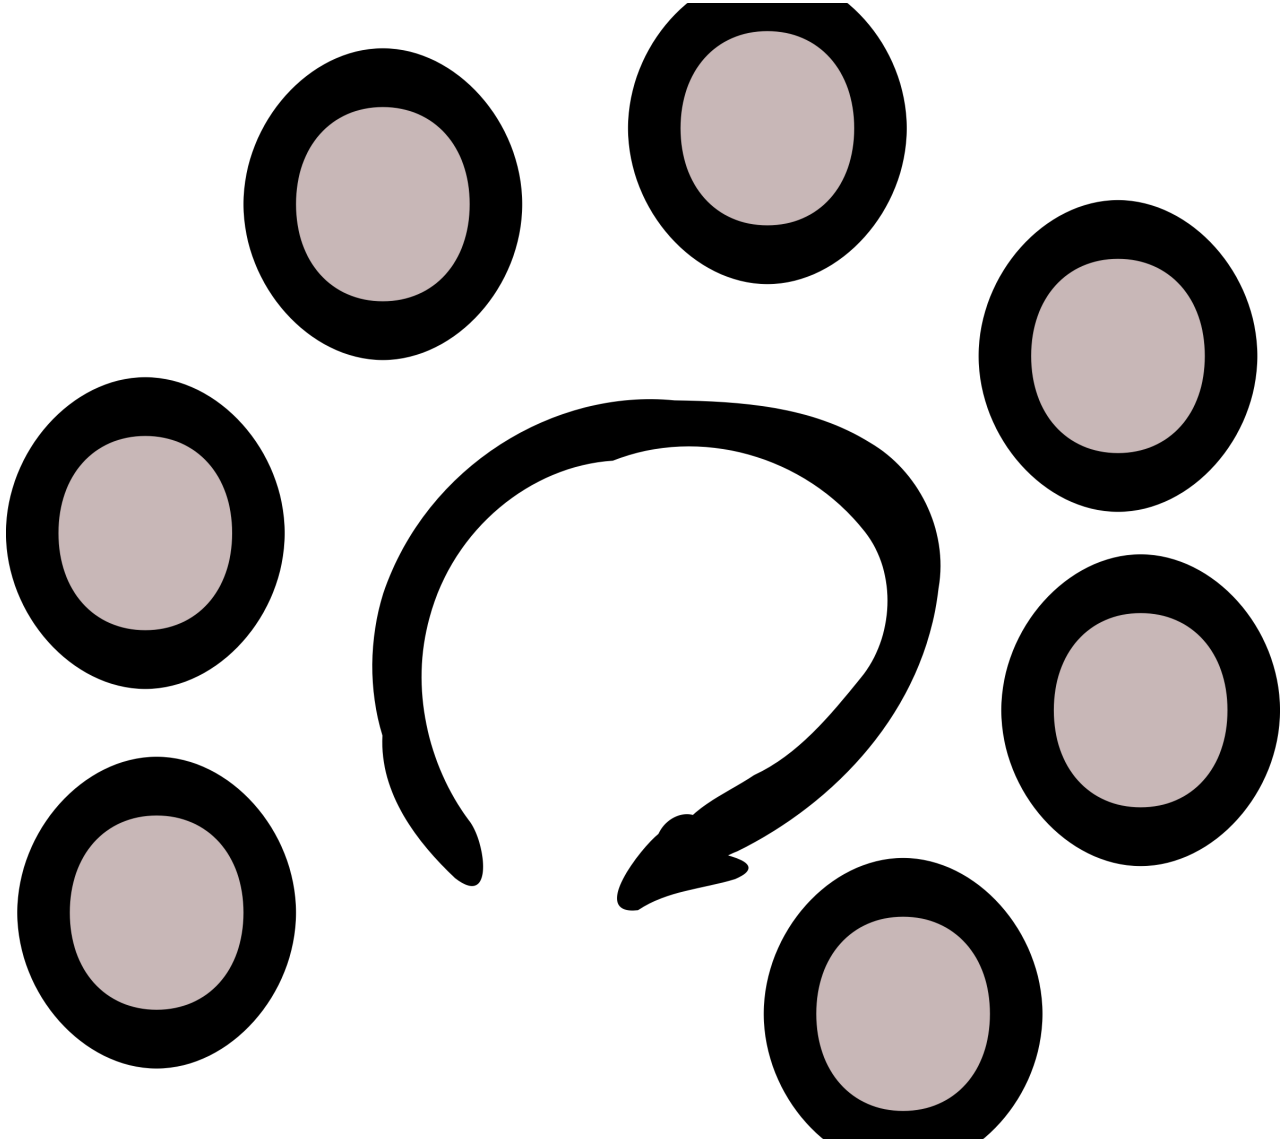

Lessons Learned

- For Continuous Service, use an active/active configuration
- Always use application services
- **Focus on the user experience**
  - Allow application work to complete before progressing to scheduled maintenance
  - Drain the work with **no application changes**
  - Disconnect older style applications transactional at the service level

Set your drain time and stop option at the service level and you are done

Take Away

1. Does your environment support rolling?
2. Are connections properly configured?
3. Is the application well behaved?
4. What is your drain timeout?

NEC Introduction

Company Name: **NEC Corporation**

Established: **July 17, 1899**

Operations: **Japanese multinational provider of information technology (IT) services and products**

Business activities in over

140 Countries and territories

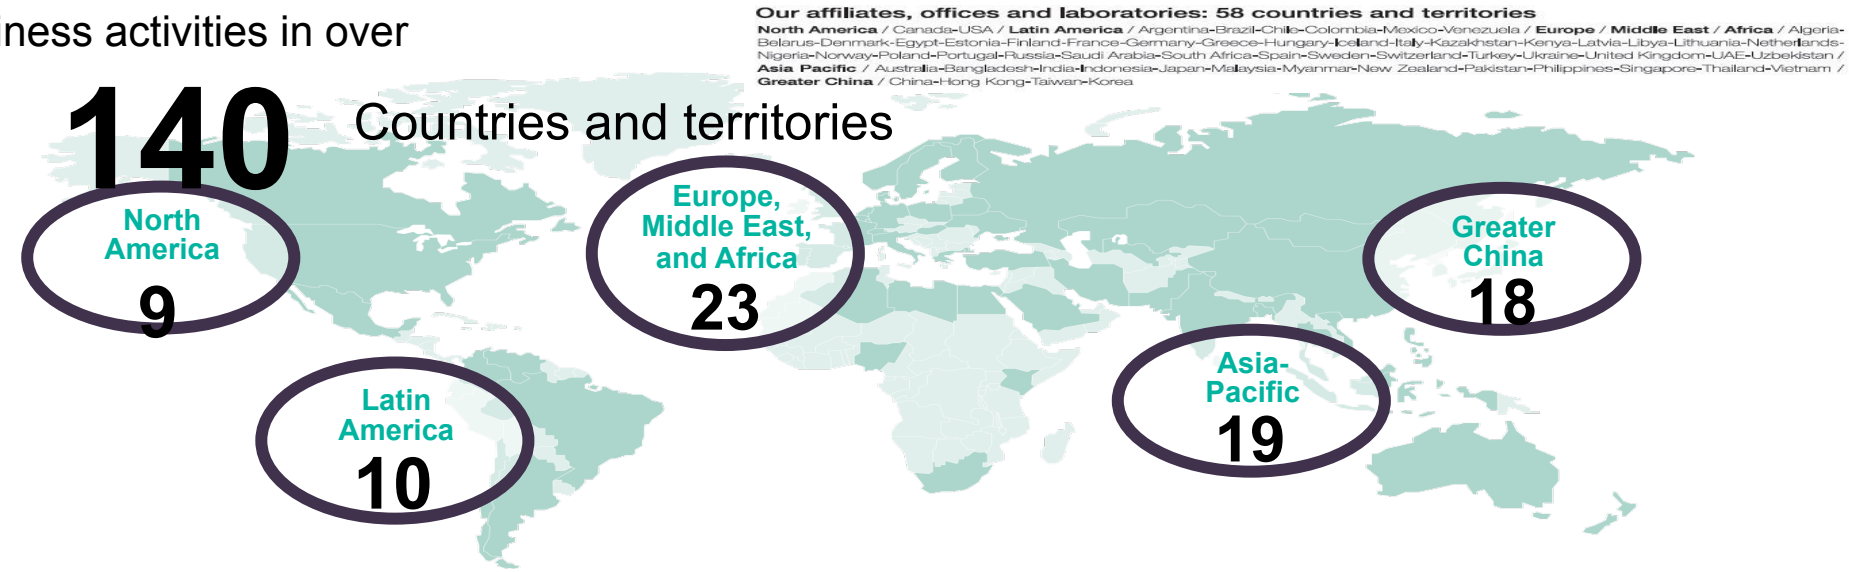

25 Billion dollar in FY2014 sales

History of NEC/Oracle alliance

The NEC and Oracle alliance is continuous over a quarter of a century

Won Global Partner Award: Database 2015,2016

2015 **Zero error solution enhancement for cloud**

2013 **Zero error planned DB maintenance and unplanned DB outage solution**

2012 NEC high available Linux DB platform

2008 NEC's RAC 11gR1 fast failover best practice

2006 NEC's RAC 10gR2 fast failover best practice

2005 STA (Strategic Technology Alliance) started

2000 Development alliance for mission critical systems

1997 NEC and BEA alliance started

1987 NEC and Oracle OEM contract started(first in Japan)

Test Environment

- Oracle DB 12*c* RAC (Services + FAN + Application Continuity)
- Oracle Multitenant
- Oracle WebLogic Server 12*c* + Active GridLink
- Oracle Active Data Guard
- Major hospital application

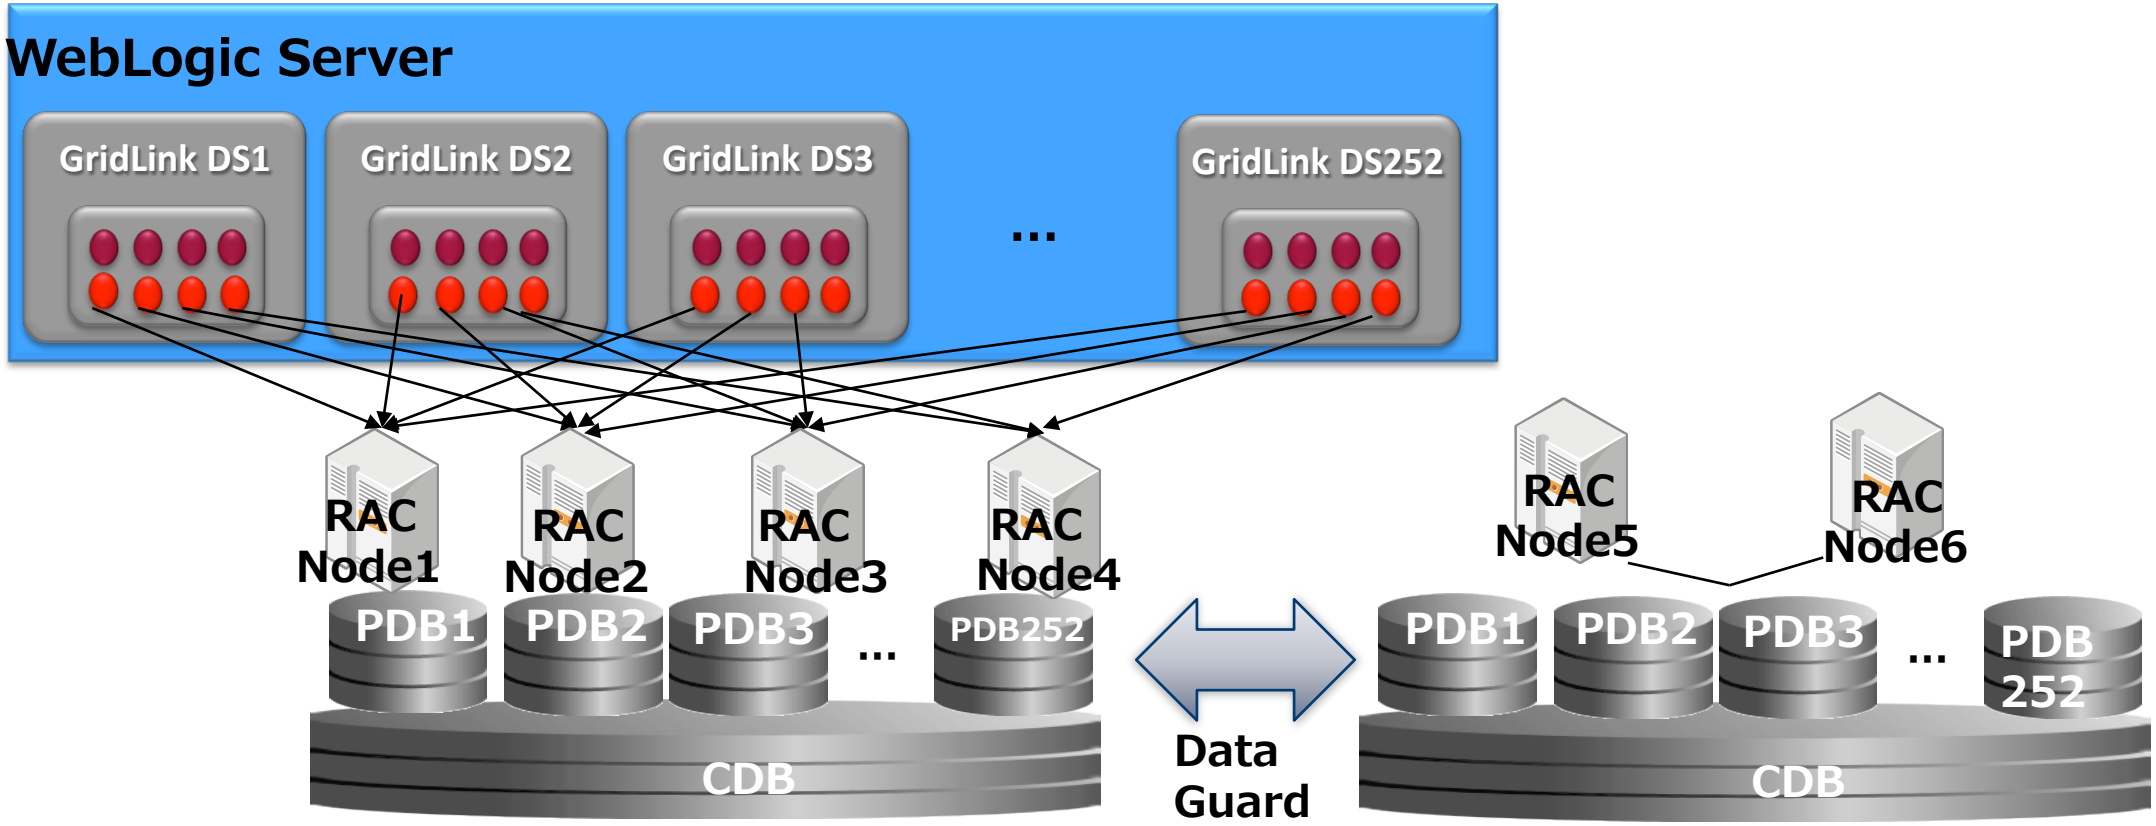

Test cases

1.PDB service stop

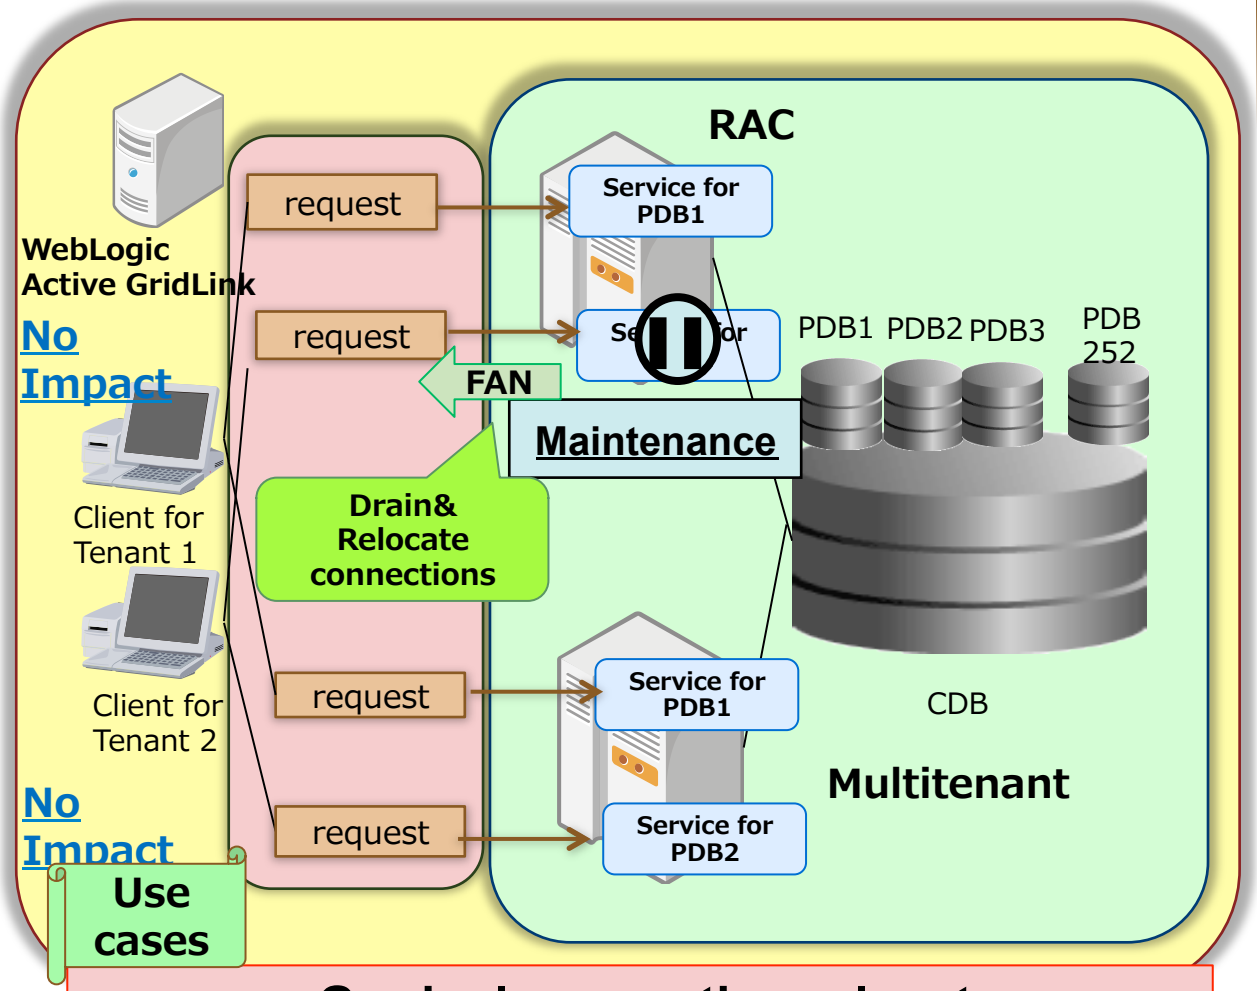

Use cases

- Gradual connection relocate
- Load Balance

2.RAC node maintenance

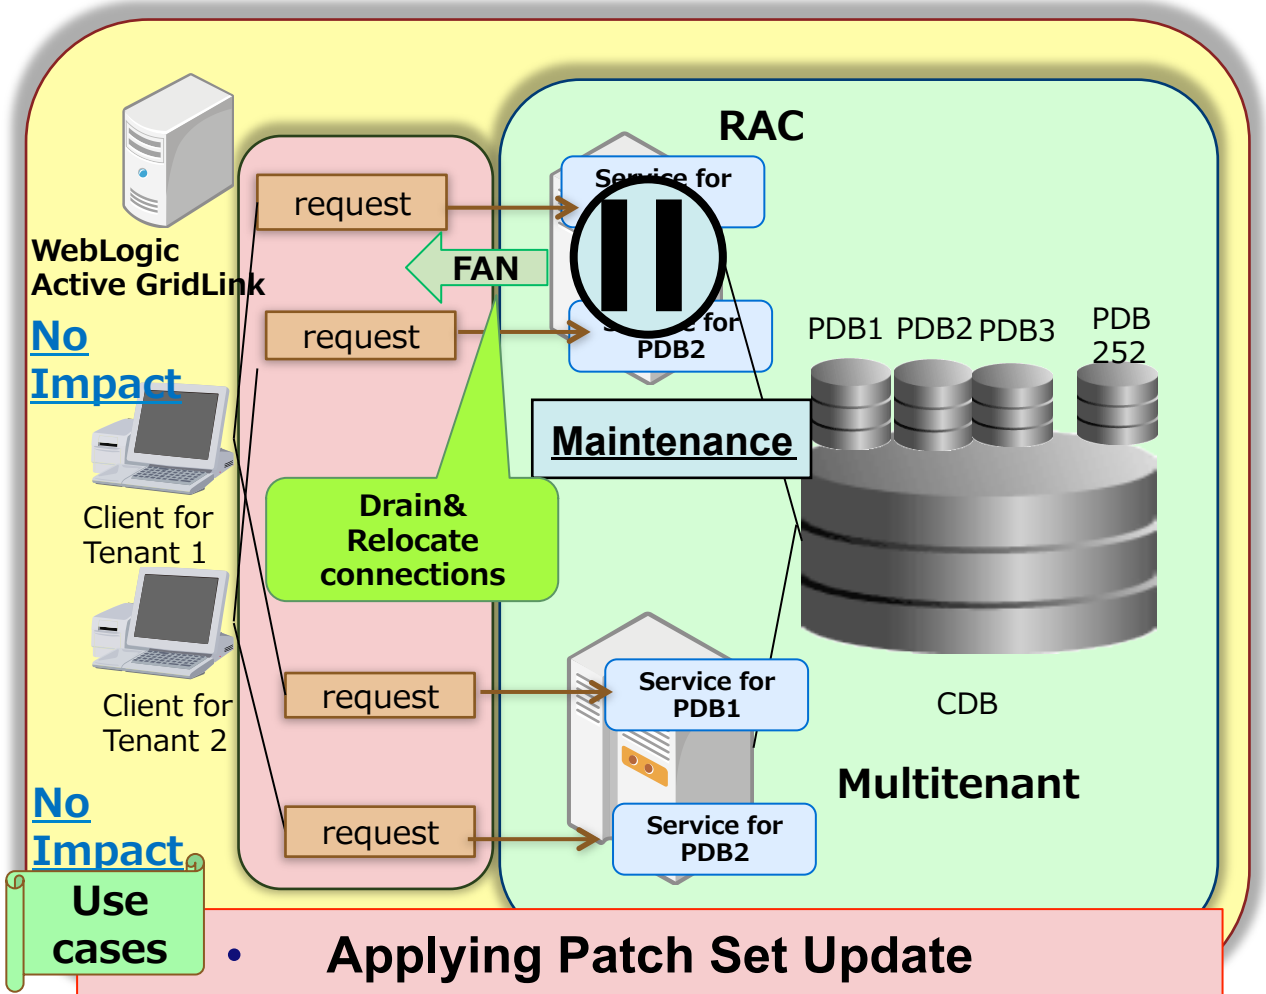

Use cases

- Applying Patch Set Update
- Configuration change
- Hardware maintenance

Test cases

3.Data Center Maintenance

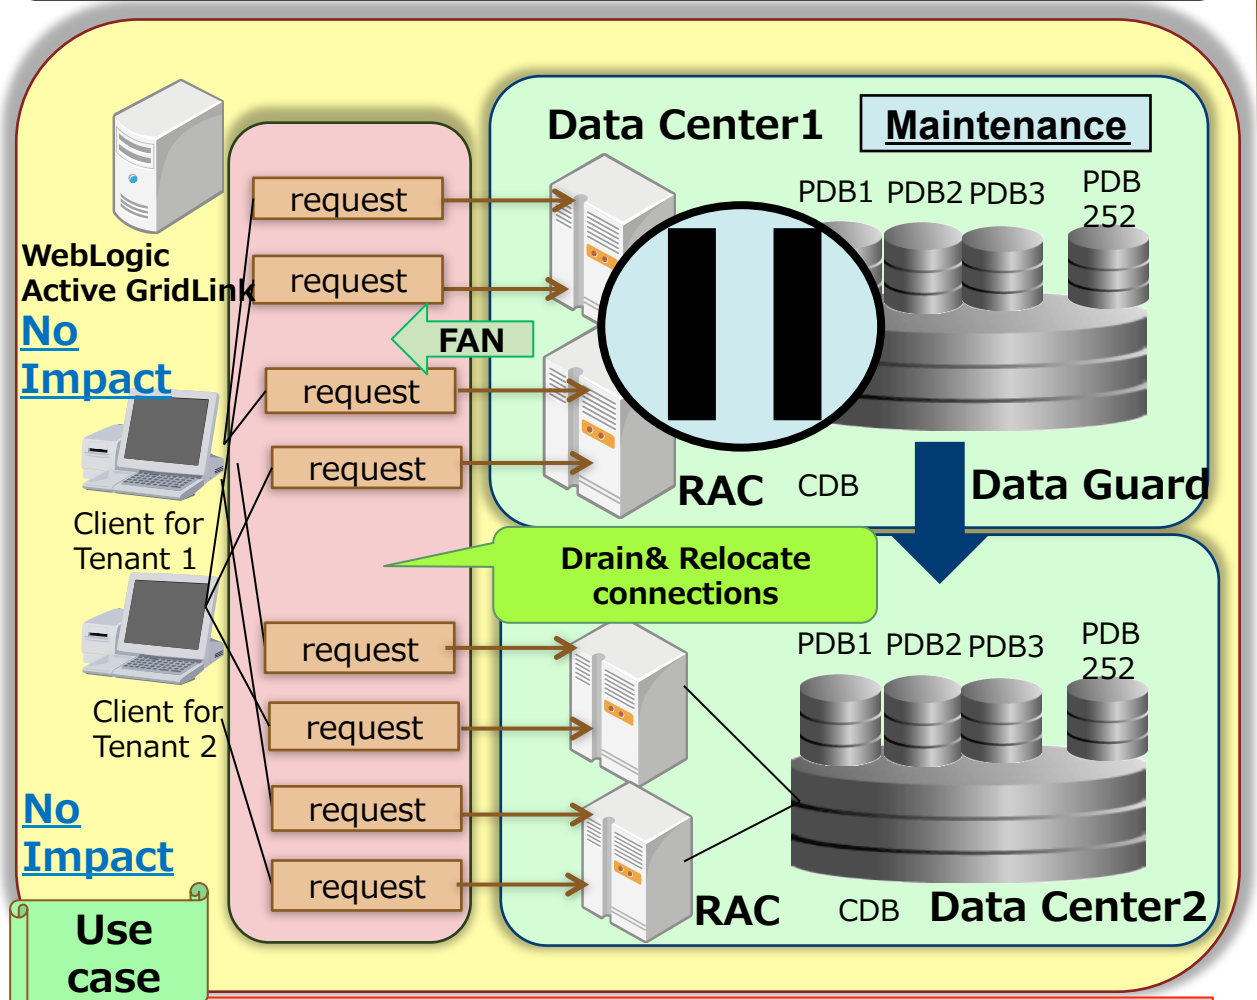

- Data Center maintenance

4. PDB migration by PDB Relocate

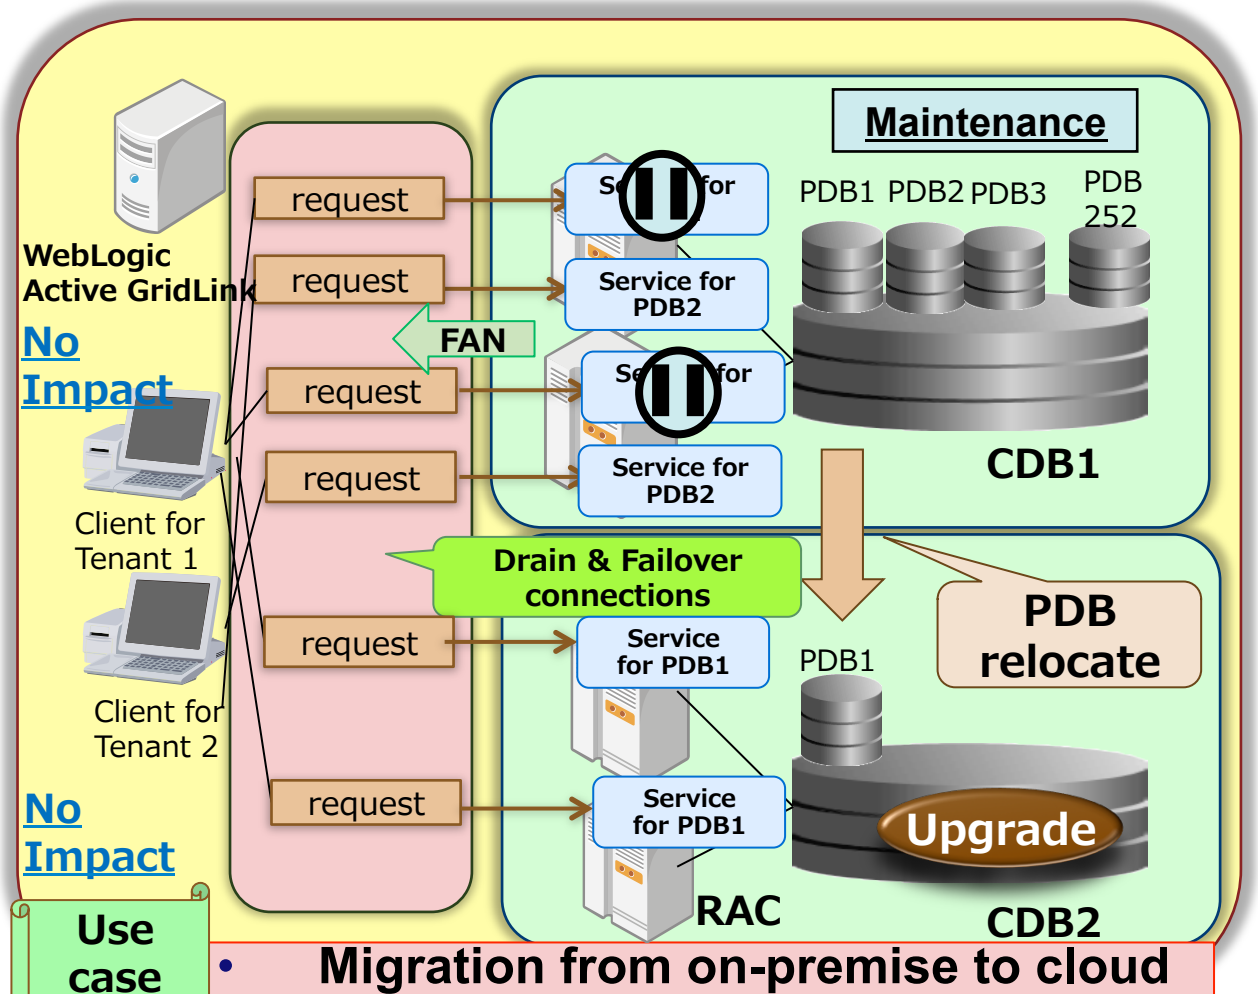

- Migration from on-premise to cloud
- Major version change
- Data center hardware replacement

RAC+Multitenant - Node Maintenance

Confirm no impact to clients

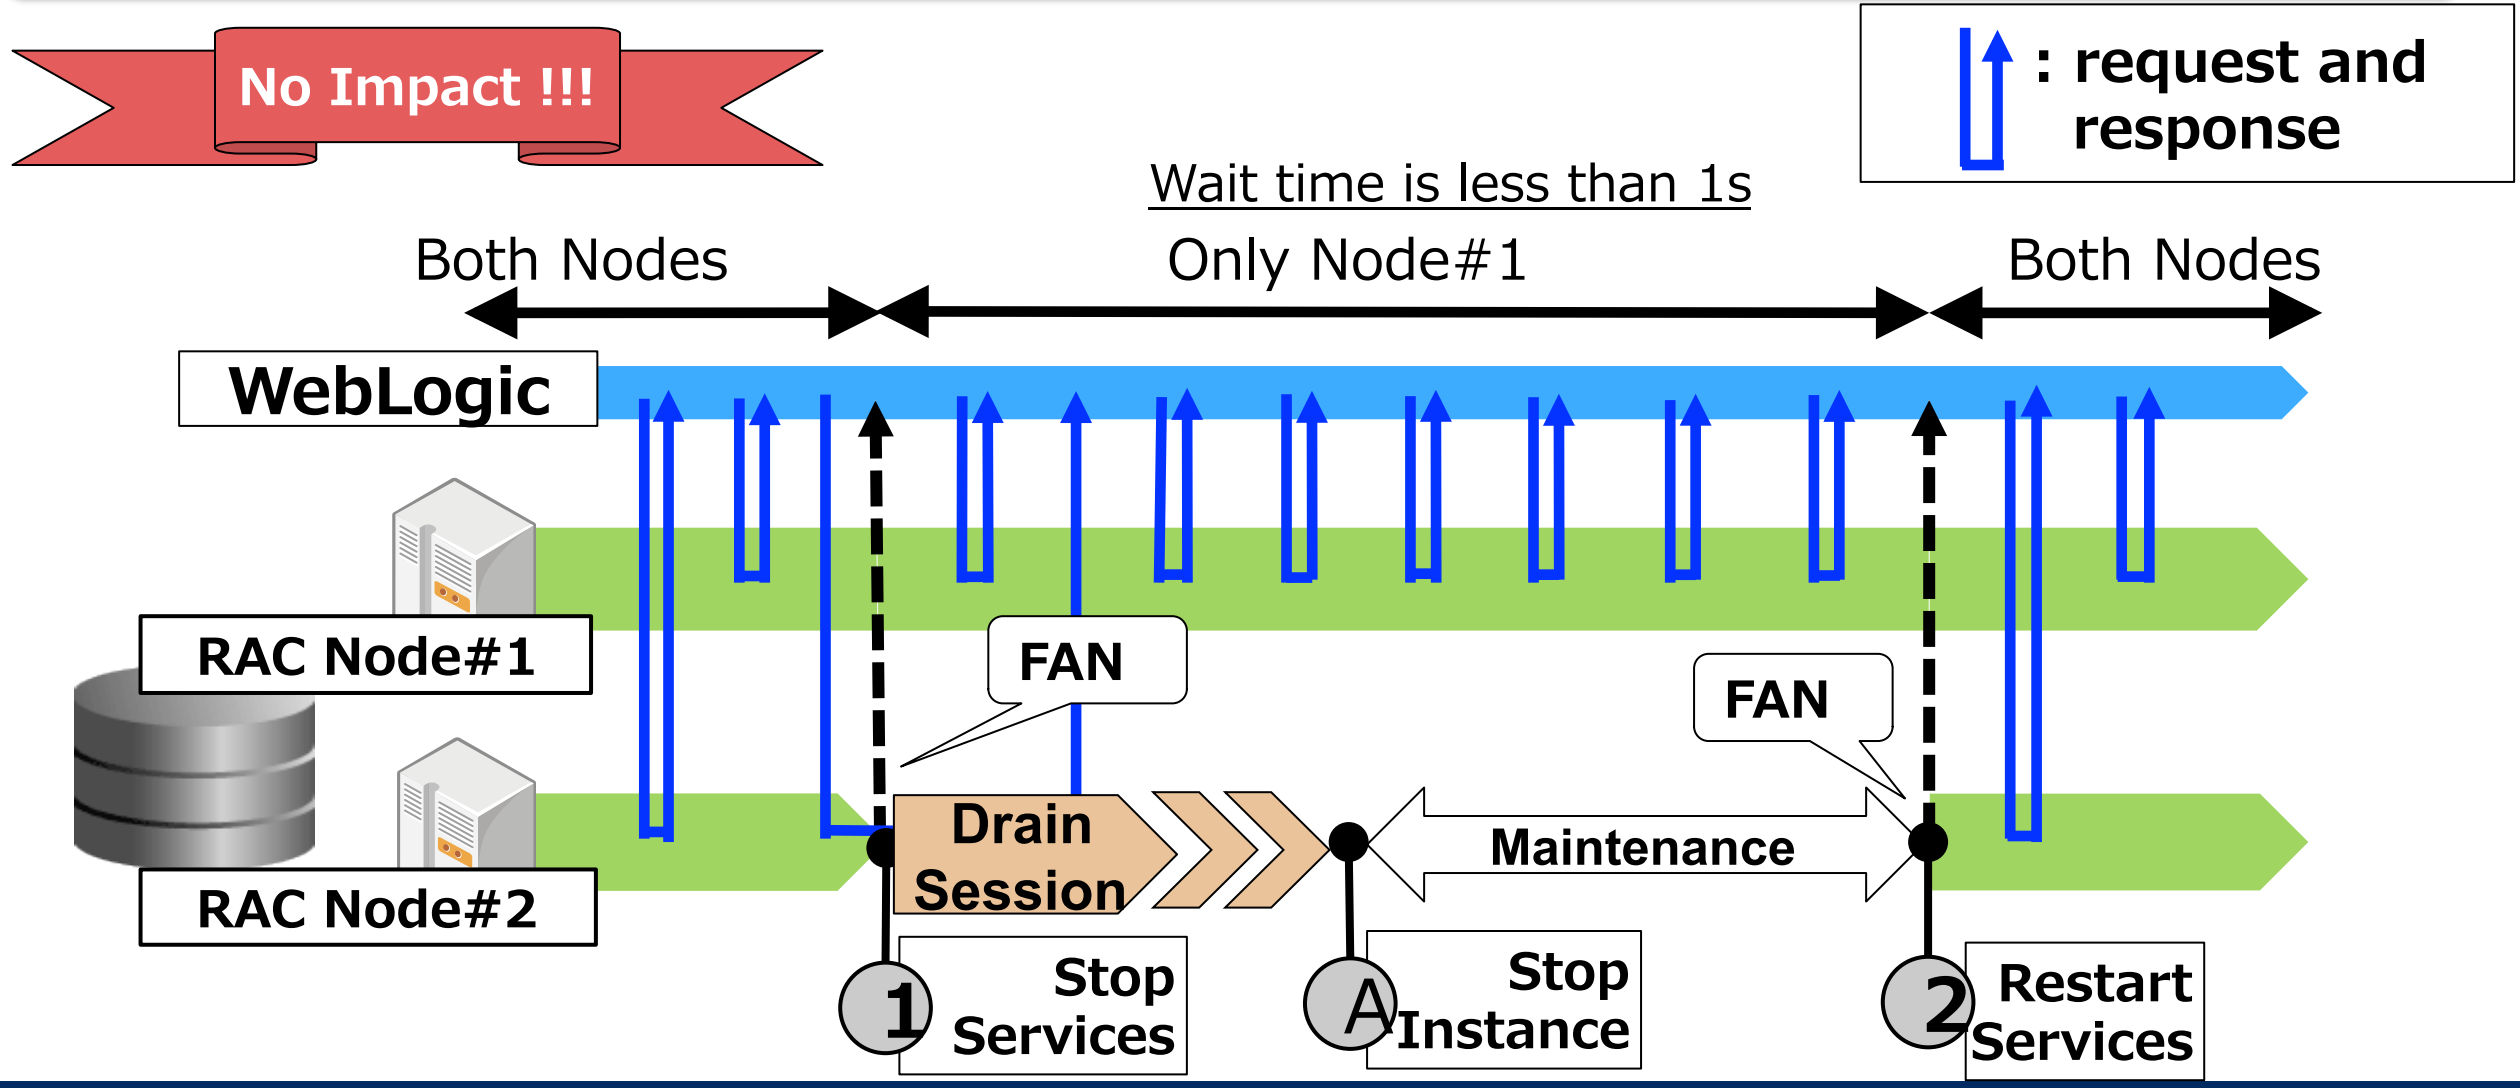

Data Center Maintenance using RAC + Data Guard

Confirm no impact to clients

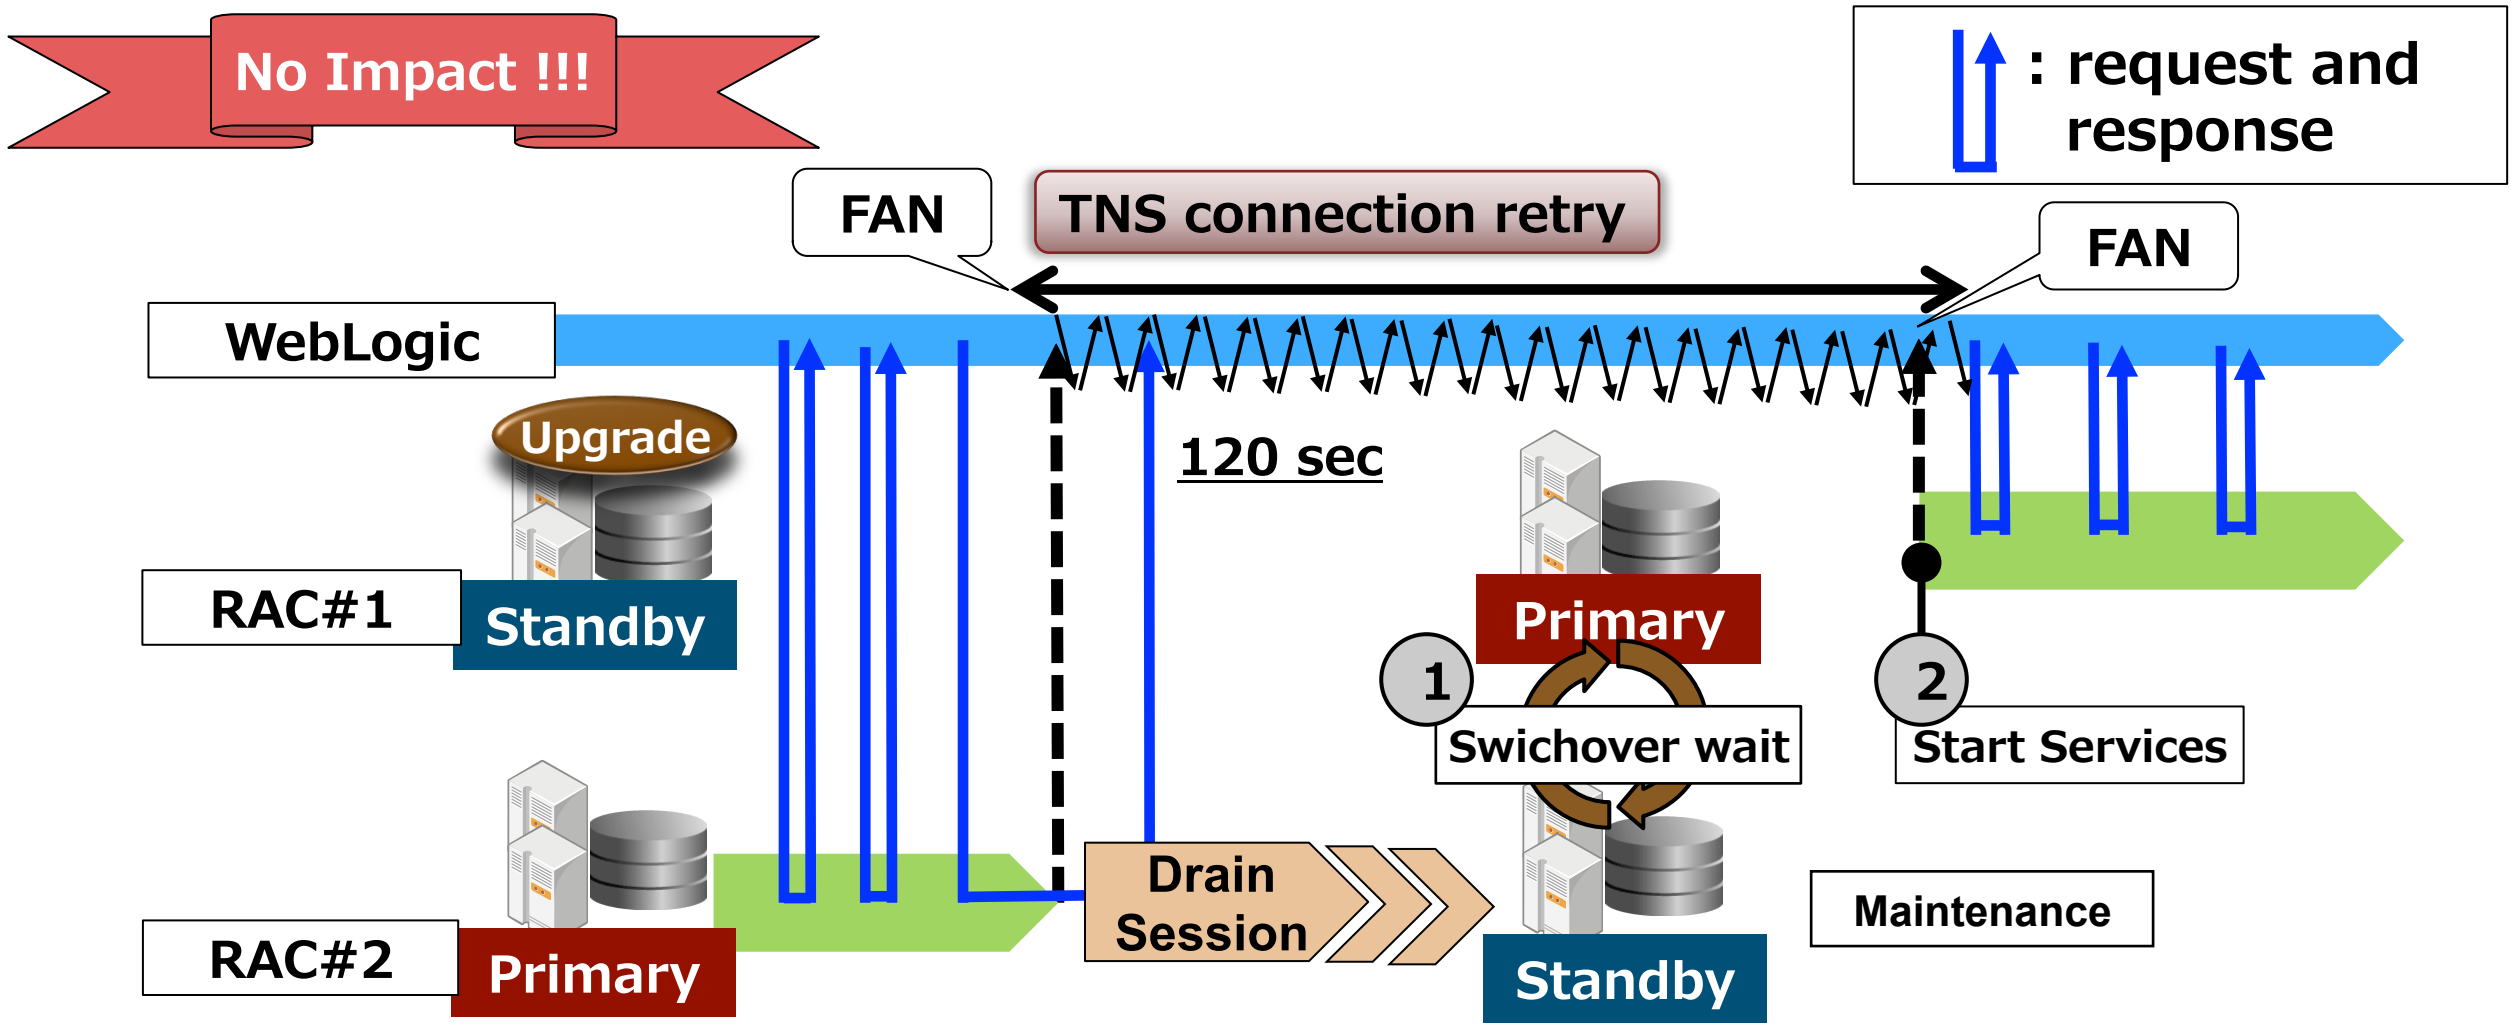

PDB Online Relocate

Confirmed no errors to all tenant's clients

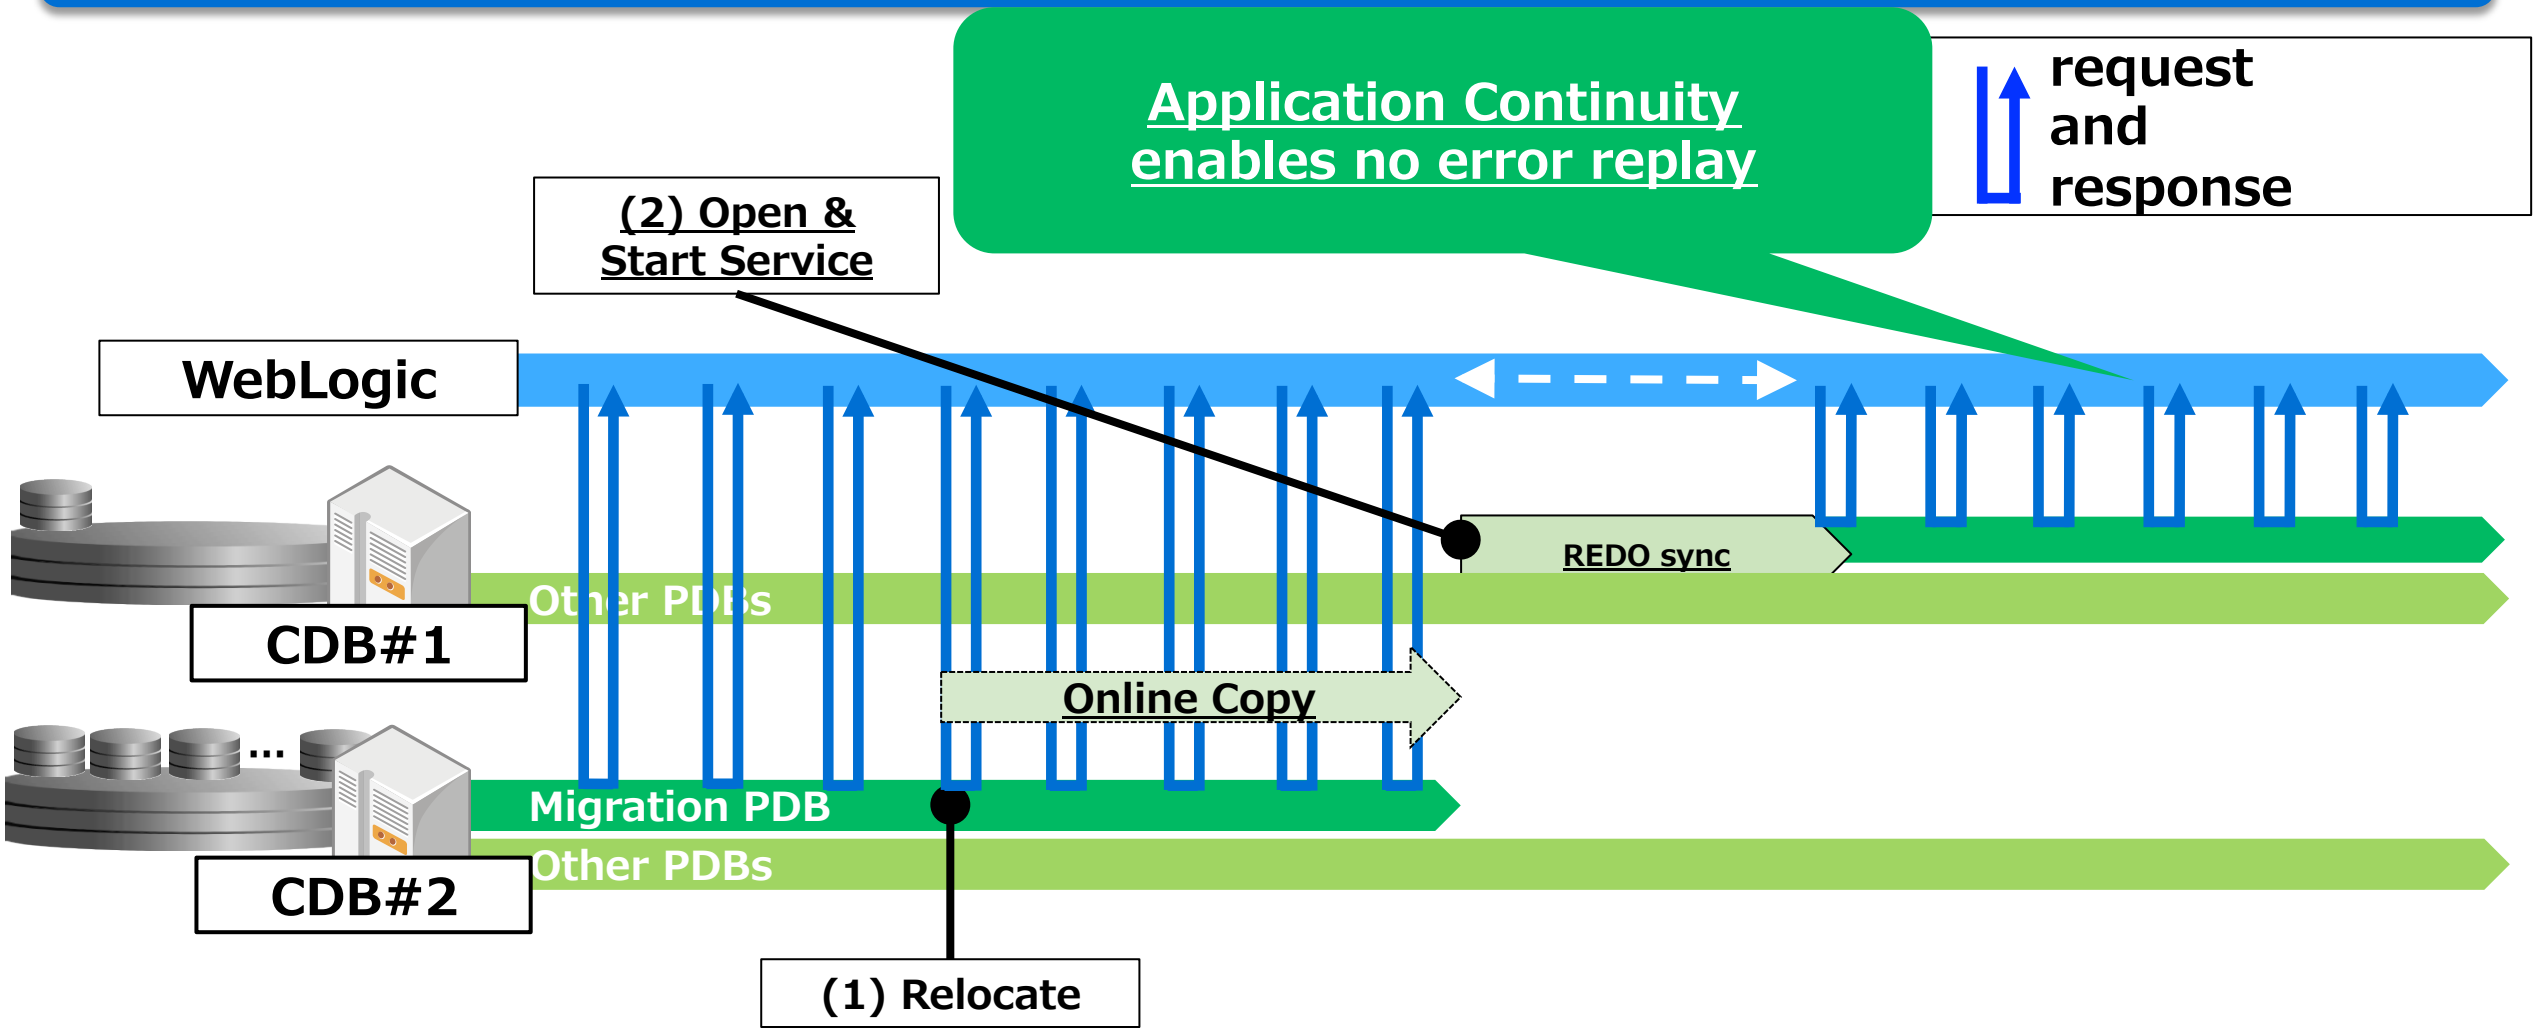

\Orchestrating a brighter world NEC

The combinatorial solution with FAN, Application Continuity, Real Application Clusters, Data Guard, WebLogic Server Active GridLink and NEC hardware and middleware enables us to provide incredibly high available system for our Mission Critical customers. This solution will become our primary solution for cloud and big data areas.

Yuki Moriyama

Senior Manager, NEC Corporation

ORACLE

Safe Harbor Statement

The preceding 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.

Integrated Cloud

Applications & Platform Services


ORACLE®