




There is no MAA without ZDLRA!

HrOUG

Fernando Simon

Senior Database Specialist
Oracle ACE ♠



SAFE HARBOR STATEMENT

- *“The postings on this document are my own and don’t necessarily represent my actual employer positions, strategies or opinions. The information here was edited to be useful for general purpose, specific data and identifications were removed to allow reach the generic audience and to be useful for the community.”*

ORACLE
ACE Program

500+ technical experts helping peers globally

The **Oracle ACE Program** recognizes and rewards community members for their technical contributions in the Oracle community



3 membership tiers



For more details on Oracle ACE Program:
bit.ly/OracleACEProgram



Nominate
yourself or someone you know:

acenomination.oracle.com

Connect:  oracle-ace_ww@oracle.com

 Facebook.com/oracleaces

 [@oracleace](https://twitter.com/oracleace)



AGENDA

- About me.
- What I will (and not) talk about.
- Oracle Maximum Availability Architecture (MAA), have you checked it already?
- ZDLRA.
- Why everything together?
- QA.

ABOUT ME

- Senior Database Specialist at Luxembourg.
 - OCP, OCE RAC, OCI Architect, Autonomous Specialist.
- Oracle ACE ♠, OOW, OOWLA, and User Groups speaker/presentations.



- Contacts:

fernando.simon.br@gmail.com

<https://www.fernandosimon.com/blog/>

<https://twitter.com/FSimonDBA>

<https://www.linkedin.com/in/fernando-simon/>

ABOUT ME

- DBA since 2004:
 - Oracle, PostgreSQL, SQLServer, and DB2.
- Head and DBA Team Manager at Court of Justice – 2010/2017:
 - Exadata since 2010:
 - Exadata V2,X2, X4 (Full), X5 (Full EF), and X6.
 - ZDLRA since 2014/2015:
 - [MAA Project, Multi-Site protection, RAC+RAC, DG, ZDLRA.](#)
- Contributing with Oracle Brazil community since 2010.
- Luxembourg October/2017.
- Consulting at European Commission:
 - LCM (Life Cycle Management) to the Oracle Products.
- Consulting at Bank Institution:
 - DBA Architect: MAA Infrastructure, ExaCC/Exadata and ZDLRA support.
- LUXOUG Co-Founder and Board Member.

WHAT I WILL (AND NOT) TALK ABOUT

- **I will talk about:**

- MAA:
 - RAC, DG, Architectures, Tips&Tricks, New ideas.
- On-Prem:
 - Common mistakes.
- ZDLRA.
- Think different.
- Think big.

- **I will not talk about:**

- Cloud.
- MAA:
 - TAF, FCF, GNS, ONS, Load balancer.
- How to configure:
 - RAC, DG, Broker, Golden Gate, Storage.
- Prices.

MAA - ORACLE MAXIMUM AVAILABILITY ARCHITECTURE

- Principles and goals:
 - Availability of every information.
 - Application Continuity.
 - Low or zero impact over environment.
 - Easy to: operate, control, and verify.
 - Sustain compliances and corporate requirements.
 - 24x7x365.
 - No Data and Access Loss.
- Two words:
 - **RPO** – Recovery Point Objective:
 - Usually, what/how much you can lose.
 - **RTO** – Recovery Time Objective:
 - Usually, time to put everything running again.
- **The goal is zero RPO and zero RTO.**

MAA - ORACLE MAXIMUM AVAILABILITY ARCHITECTURE

- Real Life:
 - Mixed environment.
 - Mix of SLA to cover.
 - Uncovered time, lack of synchronization and data loss.
 - High impact over the environment.
 - A lot of players and huge complexity.
 - Cloud.
 - Endless validation/test/validation/test....

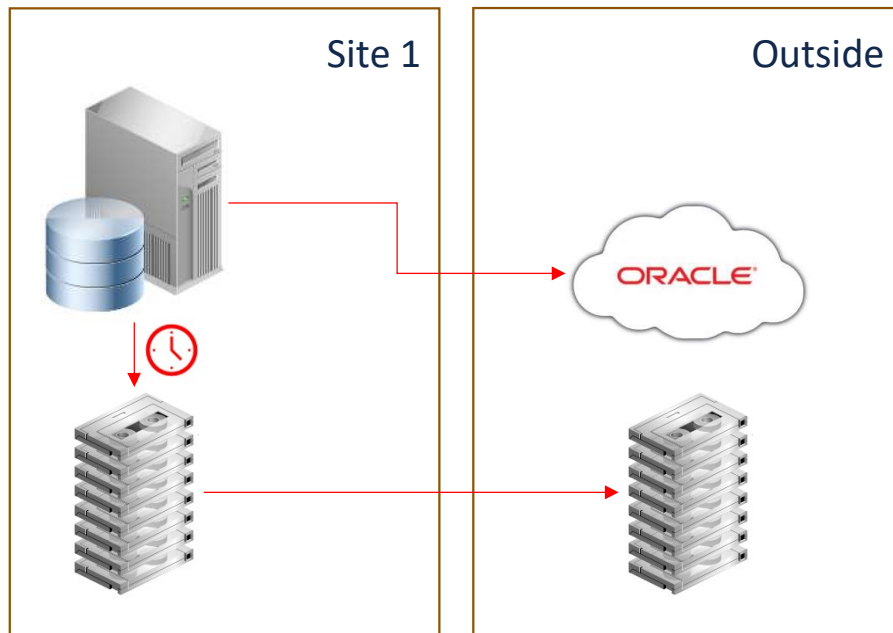
MAA - ORACLE MAXIMUM AVAILABILITY ARCHITECTURE

- To think about:
 - What is your environment today?
 - What do you need to cover?
 - Define Production!!!
 - Production for you is different for what developers think about it.
 - All your database have the same RPO?
 - Can you reach the same RPO in case of failure?
 - Understand what you need to protect, the goals, and the limitations.

MAA - ORACLE MAXIMUM AVAILABILITY ARCHITECTURE

- For Oracle, what really is MAA:
 - *Oracle's best practices blueprint based on proven Oracle high availability technologies, end-to-end validation, expert recommendations and customer experiences.*
 - Guidelines/Reference architecture:
 - Infrastructure.
 - Configuration.
 - Focus:
 - Reduce planned and unplanned downtime for databases.
 - **Application continuity**, data protection, scalable.
 - **ZERO RPO and ZERO RTO.**

MAA – BRONZE ARCHITECTURE



- **Bronze:**

- Single Instance.
- When the restore from the last backup is enough.
- Traditional environment:
 - **DB + Storage.**
 - Susceptible to HW errors. Single point of failure (SPOF).
- RTO and RPO will not be Zero.

Event	RPO	RTO
Hardware error	> Zero	> Zero
Database Error (SW)	Zero	> Zero
Data Corruption	Since last backup	> Zero
Site Outage	Can be > then Zero	> Zero

MAA – SILVER ARCHITECTURE

- **Silver:**

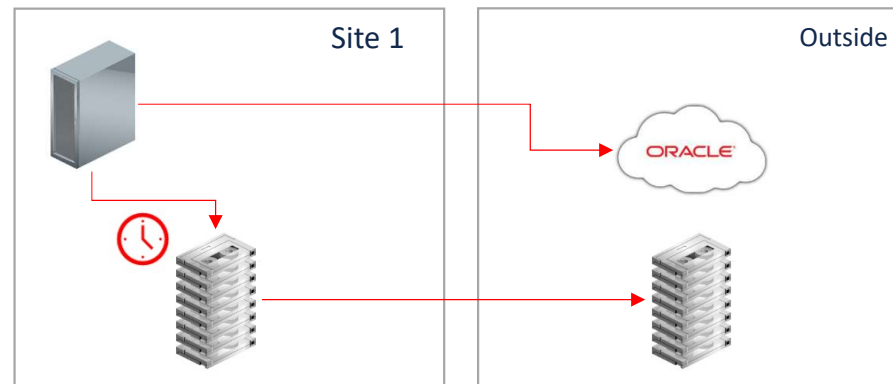
- **Start focus on Application Continuity.**

- Engineered systems (Exa/ODA), RAC:

- Reduces SPOF.
- Susceptible to HW errors.


- RTO can reach zero but not RPO.

- Still relies on backups.



Event	RPO	RTO
Hardware error	Zero*	Can be Zero
Database Error (SW)	Zero*	Can be Zero
Data Corruption	Since last backup	> Zero
Site Outage	Can be > then Zero	> Zero


MAA - GOLD ARCHITECTURE



- **Gold:**
 - **Focus in Application and Data Continuity.**
 - **Multi Site, Data Guard, and Engineered systems.**
 - Golden not protect for simultaneous outages.
 - RTO and RPO zero:
 - **Replication is done by DG.**
 - **RPO zero depends on the DG protection.**

Event	RPO	RTO
Hardware error	Zero*	Zero
Database Error (SW)	Zero	Zero
Data Corruption	Zero	Zero
Site Outage	Zero	Zero

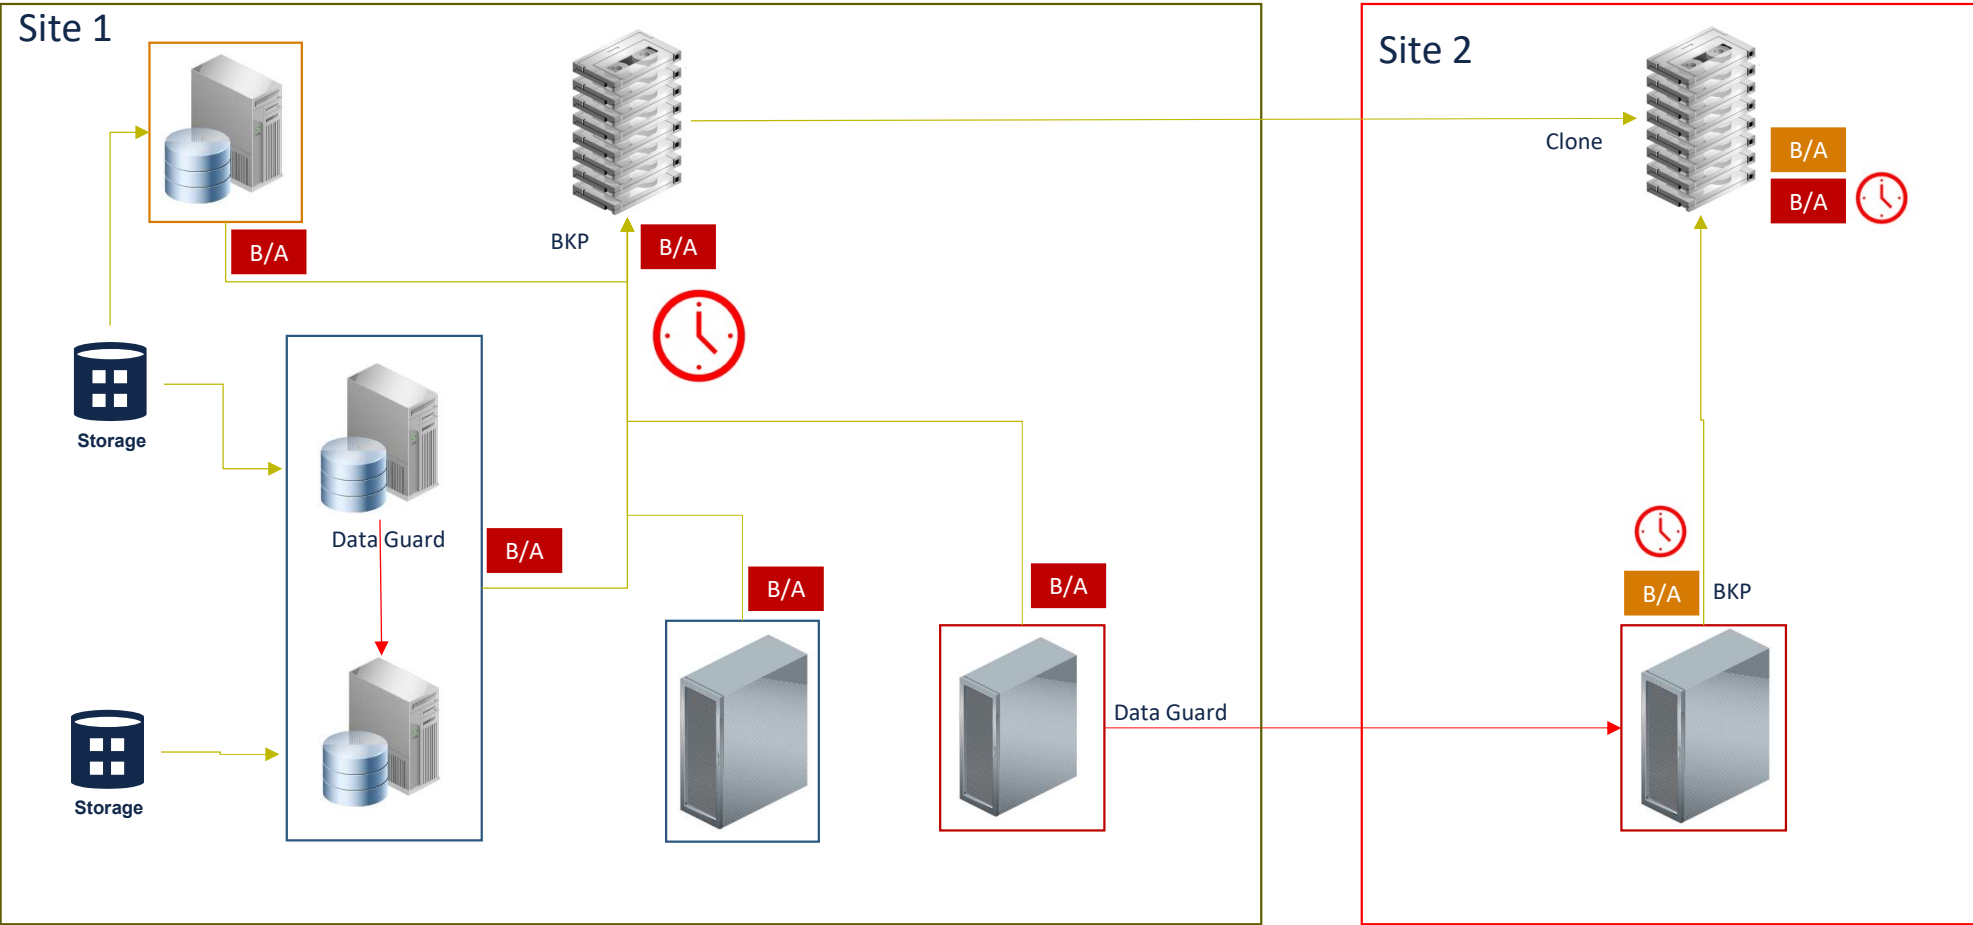
MAA – PLATINUM ARCHITECTURE




- Platinum:
 - More attention in **Application Continuity**.
 - Platinum adds intra and multi-region replication:
 - Golden Gate for zero RTO even for upgrade.

Event	RPO	RTO
Hardware error	Zero	Zero
Database Error (SW)	Zero	Zero
Data Corruption	Zero	Zero
Site Outage	Zero	Zero

MAA – BASE ARCHITECTURE



ZDLRA



WHAT IS ZDLRA

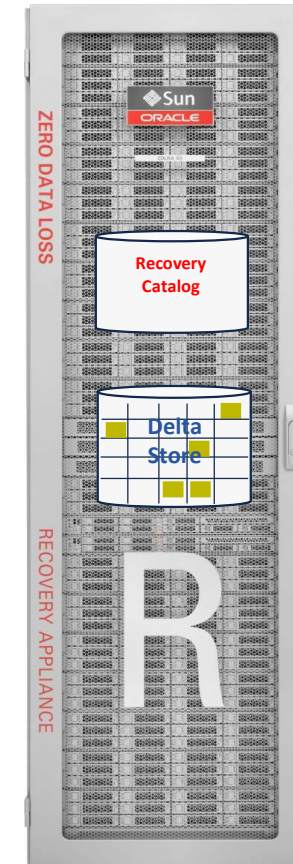
- **ZERO DATA LOSS RECOVERY APPLIANCE – ZDLRA:**

- Engineered Systems.
- Exadata based.
- Hardware + Software:
 - RA Library.
- Media Management Layer (MML) for tape.
- Native replication.
- RMAN Catalog.
- Can be used to protect ExaCC.

- **DOES NOT REDUCES RTO, JUST RPO.**

ZDLRA – INSIDE APPLIANCE

- Oracle Database:
 - Delta Store:
 - Store the backups/archivelogs plus configuration metadata.
 - Backup Policies, Packages, Tasks, and Operations.
 - Rman catalog:
 - Self-driven catalog: **No more crosscheck and validations.**
- EM/CC/CLI.
- Backup library for each client:
 - From Oracle 10g to 19c, RISC and CISC, AIX, Win, Linux, Solaris...



ZDLRA – VIRTUAL FULL BACKUP

- Virtual Full Backup:

- Ingested RMAN datafile backup sets are deconstructed and reconstructed data block per data block.
- Generate one plan/index for every datafile.
- Differs from deduplication:
 - **Content aware.**
 - Opens the backup set and **natively** “see” each Oracle data block.

- **Incremental Forever Strategy:**

- Initial level 0 backup + consequent level 1 backups.
- ZDLRA merge both to generate the virtual full backup.
- Backup and restore commands continue the same.

```
RMAN> list backup of datafile 1;
```

```
List of Backup Sets
```

```
=====
```

```
BS Key   Type LV Size      Device Type Elapsed Time Completion Time
-----
14406    Incr 0   330.29M   SBT_TAPE    00:03:16    05/01/2020 17:40:31
        BP Key: 14407   Status: AVAILABLE Compressed: YES Tag: BKP-DB-INCO
        Handle: VB$_1891149551_14397I Media:
```

```
List of Datafiles in backup set 14406
```

```
File LV Type Ckp SCN    Ckp Time              Abs Fuz SCN Sparse Name
-----
1      0   Incr 1885317  05/01/2020 17:37:15          NO    /u01/app/oracle/oradata/ORCL18C/system01.dbf
```

```
BS Key   Type LV Size      Device Type Elapsed Time Completion Time
-----
14431    Incr 1    56.00K   SBT_TAPE    00:00:02    05/01/2020 17:44:24
        BP Key: 14432   Status: AVAILABLE Compressed: YES Tag: BKP-DB
        Handle: VB$_1891149551_14430I Media:
```

```
List of Datafiles in backup set 14431
```

```
File LV Type Ckp SCN    Ckp Time              Abs Fuz SCN Sparse Name
-----
1      1   Incr 1885774  05/01/2020 17:44:22          NO    /u01/app/oracle/oradata/ORCL18C/system01.dbf
```

```
BS Key   Type LV Size      Device Type Elapsed Time Completion Time
-----
14435    Incr 0   329.18M   SBT_TAPE    00:00:02    05/01/2020 17:44:24
        BP Key: 14436   Status: AVAILABLE Compressed: YES Tag: BKP-DB
        Handle: VB$_1891149551_14430_1 Media:
```

```
List of Datafiles in backup set 14435
```


```
File LV Type Ckp SCN    Ckp Time              Abs Fuz SCN Sparse Name
-----
1      0   Incr 1885774  05/01/2020 17:44:22          NO    /u01/app/oracle/oradata/ORCL18C/system01.dbf
```

```
RMAN>
```

MERGED

ZDLRA - REPLICATION

- Replication:
 - One-Way:
 - One master and one destination.
 - Bi-Directional:
 - Both sides replicate each other.
 - Hub/Spoke:
 - One to many.
 - Executed automatically, as soon as possible.



```
RMAN> list backup of archivelog sequence 49;
```

```
List of Backup Sets  
=====
```

```
BS Key Size  
-----  
26322 256.00K
```

```
List of Archived Logs in backup set 26322
```

Thrd	Seq	Low SCN	Low Time	Next SCN	Next Time
1	49	2418960	27/09/2021 19:07:17	2419284	27/09/2021 19:09:02

```
Backup Set Copy #1 of backup set 26322
```

Device Type	Elapsed Time	Completion Time	Compressed	Tag
SBT_TAPE	00:00:20	27/09/2021 19:09:04	NO	TAG20210927T190903

```
List of Backup Pieces for backup set 26322
```

BP Key	Pc#	Status	Media
26323	1	AVAILABLE	Recovery Appliance

```
Backup Set Copy #2 of backup set 26322
```

Device Type	Elapsed Time	Completion Time
SBT_TAPE	00:00:20	27/09/2021 19:09:20

```
List of Backup Pieces for backup set 26322
```

BP Key	Pc#	Status	Media
26373	1	AVAILABLE	Recovery Appliance

```
RMAN>
```

```
RMAN> list backup of archivelog sequence 48;
```

```
List of Backup Sets  
=====
```

BS Key	Size	Device Type	Elapsed Time	Completion Time
26237	256.00K	SBT_TAPE	00:00:15	27/09/2021 19:07:34

BP Key: 26280 Status: AVAILABLE Compressed: NO Tag: TAG20210927T190719
Handle: RA_SBT_ORCL19C_1487680695_24355_610a4tf7_1_2_26237

```
Media: Recovery Appliance (ZDLRAS2)
```

```
List of Archived Logs in backup set 26237
```

Thrd	Seq	Low SCN	Low Time	Next SCN	Next Time
1	48	2418557	27/09/2021 19:05:16	2418960	27/09/2021 19:07:17

```
RMAN>
```

ZDLRA - REPLICATION

```
[oracle@orcloel7 ~]$ rman target=/ catalog=vpcsrc/vpcsrc@zdlras1-scan:1521/zdlras1
```

```
Recovery Manager: Release 19.0.0.0.0 - Production on Mon Sep 27 19:05:08 2021  
Version 19.3.0.0.0
```

```
Copyright (c) 1982, 2019, Oracle and/or its affiliates. All rights reserved.
```

```
connected to target database: ORCL19C (DBID=1487680695)  
connected to recovery catalog database
```

```
RMAN> restore archivelog sequence 48;
```

```
Starting restore at 27/09/2021 19:15:38  
using channel ORA_SBT_TAPE_1  
using channel ORA_DISK_1
```

```
channel ORA_SBT_TAPE_1: starting archived log restore to default destination  
channel ORA_SBT_TAPE_1: restoring archived log  
archived log thread=1 sequence=48  
channel ORA_SBT_TAPE_1: reading from backup piece 1  
channel ORA_SBT_TAPE_1: piece handle=RA_SBT_ORCL19C_1487680695_24355_630a4tif_1_2_26322 tag=TAG20210927T190903  
channel ORA_SBT_TAPE_1: restored backup piece 1  
channel ORA_SBT_TAPE_1: restore complete, elapsed time: 00:00:25  
Finished restore at 27/09/2021 19:16:24
```

```
RMAN>
```

```
[oracle@orcloel7 ~]$ rman target=/ catalog=vpcprep/vpcprep@zdlras2-scan:1521/zdlras2
```

```
Recovery Manager: Release 19.0.0.0.0 - Production on Mon Sep 27 19:18:19 2021  
Version 19.3.0.0.0
```

```
Copyright (c) 1982, 2019, Oracle and/or its affiliates. All rights reserved.
```

```
connected to target database: ORCL19C (DBID=1487680695)  
connected to recovery catalog database
```

```
RMAN> restore archivelog sequence 49;
```

```
Starting restore at 27/09/2021 19:19:30  
allocated channel: ORA_SBT_TAPE_1  
channel ORA_SBT_TAPE_1: SID=69 device type=SBT_TAPE  
channel ORA_SBT_TAPE_1: RA Library (ZDLRAS2) SID=CCFE4CD091C3196BE053010310ACD738  
allocated channel: ORA_DISK_1  
channel ORA_DISK_1: SID=34 device type=DISK
```

```
channel ORA_SBT_TAPE_1: starting archived log restore to default destination  
channel ORA_SBT_TAPE_1: restoring archived log  
archived log thread=1 sequence=49  
channel ORA_SBT_TAPE_1: reading from backup piece RA_SBT_ORCL19C_1487680695_24355_630a4tif_1_2_26322  
channel ORA_SBT_TAPE_1: piece handle=RA_SBT_ORCL19C_1487680695_24355_630a4tif_1_2_26322 tag=TAG20210927T190903  
channel ORA_SBT_TAPE_1: restored backup piece 1  
channel ORA_SBT_TAPE_1: restore complete, elapsed time: 00:00:25  
Finished restore at 27/09/2021 19:20:00
```

```
RMAN>
```


ZDLRA – CLONE TO TAPE

- Tape and Cloud:
 - Can copy to tapes directly, is MML:
 - It can be by OSB or Third Part (if it is compatible with rman).
 - Natively uses Oracle Secure Backup (OSB) with SAN.
 - Can copy backups to Oracle Cloud, Object Store.
 - **Requires schedule.** Based on dbms_scheduler to call clone job.
 - Can be based on several metadata information:
 - RMAN Tag, Timestamp, Backup Types...

ZDLRA – REAL-TIME REDO TRANSPORT

- Real-Time Redo Transport:
 - Is the “zero data loss” - guarantee the zero RPO:
 - Reduces RPO from the last backup to zero/sub-seconds.
 - Is simple. ZDLRA it is a log_archive_dest destination:
 - Uses same procedure than Data Guard. ZDLRA operates with RFS to receive redo buffers (SYNC/ASYNC).

```
SQL> select value from v$parameter where name = 'log_archive_dest_2';  
  
VALUE  
-----  
SERVICE="zdlras1-scan:1521/zdlras1:VPCSRC" SYNC NOAFFIRM DB_UNIQUE_NAME=zdlras1 VALID_FOR=(ONLINE_LOGFILE,ALL_ROLES)  
  
SQL>
```

- If protected database crash, ZDLRA creates *'partial archived log backup'* with last sent redo info.
- **DOES NOT REQUIRE DG LICENSE! INDEPENDENTLY OF ORACLE EDITION THAT YOU USE!**

ZDLRA – REAL-TIME REDO TRANSPORT

```

RMAN> list copy of archivelog all;

List of Archived Log Copies for database with db_unique_name OR19DG
=====

Key       Thrd Seq      S Low Time
-----
13295    1      90      A 01/01/2020 22:28:18
        Name: +RECO/OR19DG/ARCHIVELOG/2020_01_01/thread_1_seq_90.446.1028586681

RMAN> alter system archive log current;

Statement processed

RMAN> list backup of archivelog sequence 91;

List of Backup Sets
=====

BS Key   Size       Device Type Elapsed Time Completion Time
-----
13314   21.50K     SBT_TAPE    00:00:01     01/01/2020 22:32:09
        BP Key: 13315   Status: AVAILABLE Compressed: YES Tag: TAG20200101T223208
        Handle: $RSCN_1920977_RTIM_1028557385_THRD_1_SEQ_91_CTKEY_13291_BACKUP Media:

List of Archived Logs in backup set 13314
-----
Thrd Seq      Low SCN    Low Time           Next SCN    Next Time
-----
1     91         2519988    01/01/2020 22:31:20 2520099    01/01/2020 22:31:51

RMAN>

```

ZDLRA – REAL-TIME REDO TRANSPORT

```
[oracle@exaclvm01-ORAD18]$ for i in {1..100000}
> do
> echo "Insert Data $i - date +%d-%m-%Y-%H%M%S"
> sqlplus -s / as sysdba<<EOF
> set heading on feedback on;
> insert into testIns(c1, c2, c3) values ($i, sysdate, 'Loop');
> commit;
> EOF
> done
Insert Data 1 - 18-10-2019-230723

1 row created.

Commit complete.
...
```

```
Insert Data 1016 - 18-10-2019-230944
```

```
1 row created.
```

```
Commit complete.
```

```
Insert Data 1017 - 18-10-2019-230944
```

```
1 row created.
```

```
commit
*
```

```
ERROR at line 1:
ORA-03113: end-of-file on communication channel
Process ID: 142277
Session ID: 53 Serial number: 30197
```

```
[oracle@exaclvm01-ORAD18]$ for i in {1..100000}
> do
> echo "Insert Data $i - date +%d-%m-%Y-%H%M%S"
> sqlplus -s / as sysdba<<EOF
> set heading on feedback on;
> insert into testIns(c1, c2, c3) values ($i, sysdate, 'Loop2');
> commit;
> EOF
> done
Insert Data 1 - 18-10-2019-230816

1 row created.

Commit complete.
```

```
Insert Data 646 - 18-10-2019-230944
```

```
1 row created.
```

```
Commit complete.
```

```
Insert Data 647 - 18-10-2019-230944
```

```
1 row created.
```

```
commit
*
```

```
ERROR at line 1:
ORA-03113: end-of-file on communication channel
Process ID: 142274
Session ID: 41 Serial number: 3186
```

Outage

ZDLRA – REAL-TIME REDO TRANSPORT

```
SQL> alter diskgroup data mount;
alter diskgroup data mount
*
ERROR at line 1:
ORA-15032: not all alterations performed
ORA-15017: diskgroup "DATA" cannot be mounted
ORA-15066: offlining disk "1" in group "DATA" may result in a data loss
SQL>
```

```
RMAN> list backup of archivelog all;
...
...
BS Key   Size      Device Type Elapsed Time Completion Time
-----
50958809 2.05M      SBT_TAPE    00:00:00     2019-10-18_23-10-12
          BP Key: 50958810   Status: AVAILABLE Compressed: YES Tag: TAG20191018T231012
          Handle: $RSCN_1129803_RTIM_1022011106_THRD_1_SEQ_5_CTKEY_50958757_BACKUP Media:
```

List of Archived Logs in backup set 50958809

Thrd	Seq	Low SCN	Low Time	Next SCN	Next Time
1	5	1131667	2019-10-18_23-05-09	1135762	2019-10-18_23-09-44

```
RMAN>
```

```
RMAN> run{
2> set until scn 1135762;
3> restore database;
4> recover database;
5> }
```

```
SQL> select count(*) from testIns group by c3;
```

```
      COUNT (*)
-----
          646
          1016
```

```
SQL>
```

WHY NOT EVERYTHING TOGETHER?

- From now what we need to handle:
 - Single instances Databases (Bronze):
 - Relies purely on backups.
 - RAC instances (Silver):
 - Storage failures (Single point of failure).
 - DG Databases (Gold):
 - In case of failures, drop protection to the Bronze level.
 - Shared hardware between different architectures:
 - Same storage to Single instances and DG.

- Why ZDLRA:
 - Real-Time Redo drops everything to zero RPO.
 - External protection even for unexpected failures and outages.
 - Multiple copies for protection.
 - First and Second line of defense.
 - Reduces load and impact for backups.
 - **All databases will have the same RPO.**

WHY EVERYTHING TOGETHER?

- MAA + ZDLRA:
 - Think BIG.

 - Today there is no MAA without ZDLRA.
 - Don't need fancy solutions.

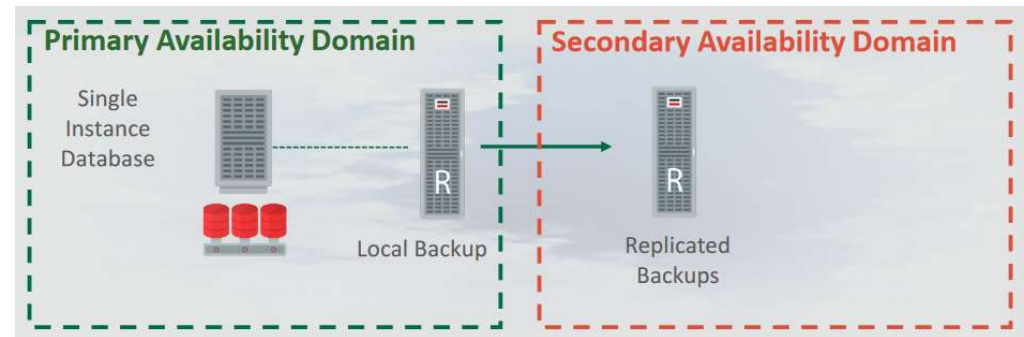
- Reference architecture for On-Prem:
 - [Oracle Maximum Availability Architecture \(MAA\)](#).
 - [Oracle MAA Reference Architectures](#).

MAA – BRONZE ARCHITECTURE + ZDLRA

- BRONZE:
 - When the restore from the last backup is enough.

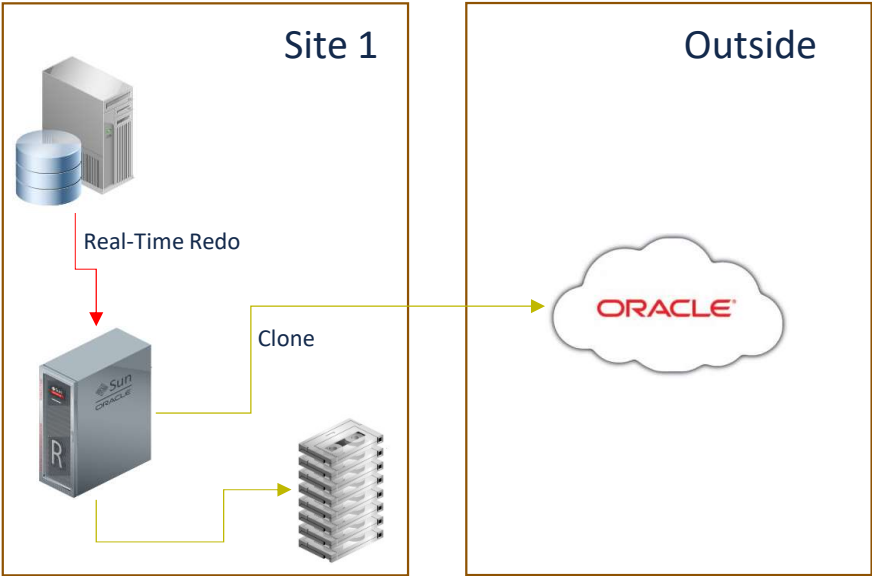
- Single instances + Traditional environment:
 - Not Engineering systems. **DB + Storage.**
 - Susceptible to HW errors. Single point of failure (SPOF).

– RTO will not be Zero.

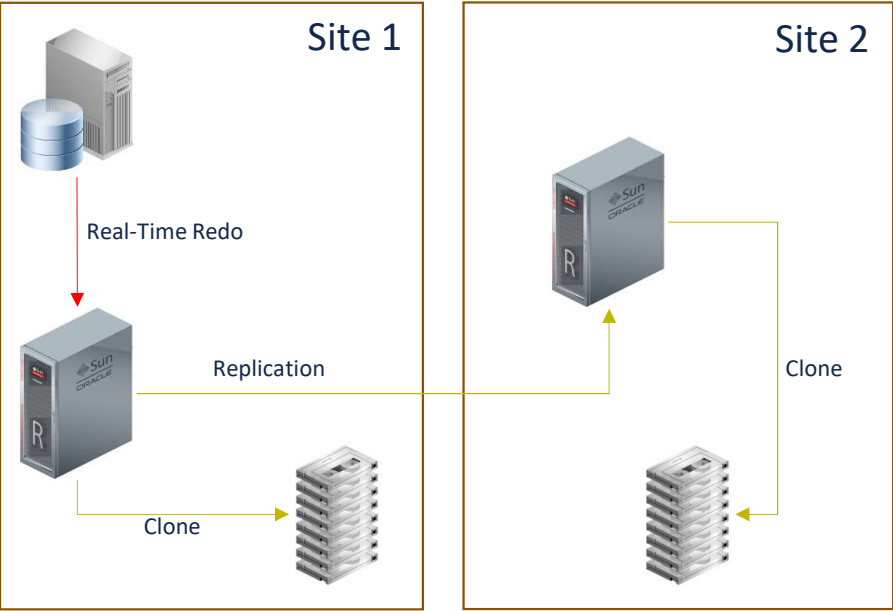


– With ZDLRA you can have ZERO RPO and multi-site protection for a Standard Edition.

MAA – BRONZE ARCHITECTURE + ZDLRA

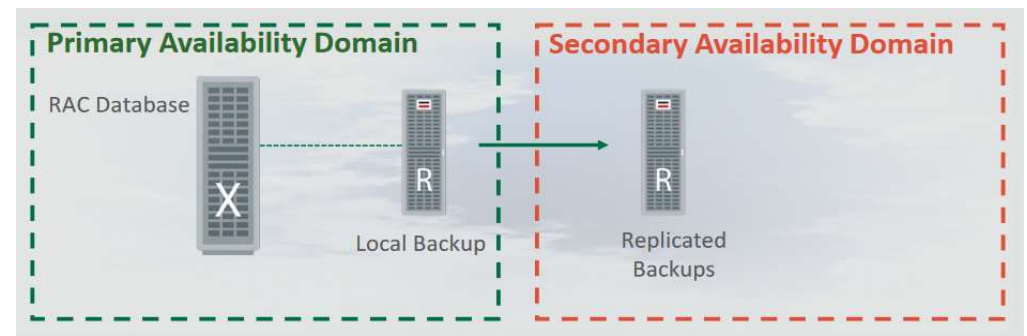


Event	RPO	RTO
Hardware error	Zero	> Zero
Database Error (SW)	Zero	> Zero
Data Corruption	Zero	> Zero
Site Outage	Can be Zero	> Zero




MAA – SILVER ARCHITECTURE + ZDLRA

- SILVER:
 - RAC or Single instance (With DG):
 - Reduces SPOF.
 - Two “ways” to do:
 - Traditional environment:
 - Not Engineered systems, more susceptible to HW errors.
 - Engineered systems (Exa/ODA), RAC.
 - **RTO and RPO can be zero.**



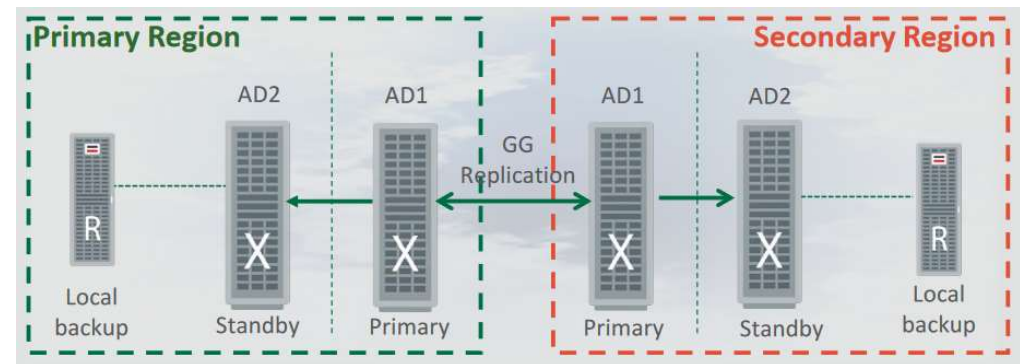

MAA – SILVER ARCHITECTURE + ZDLRA



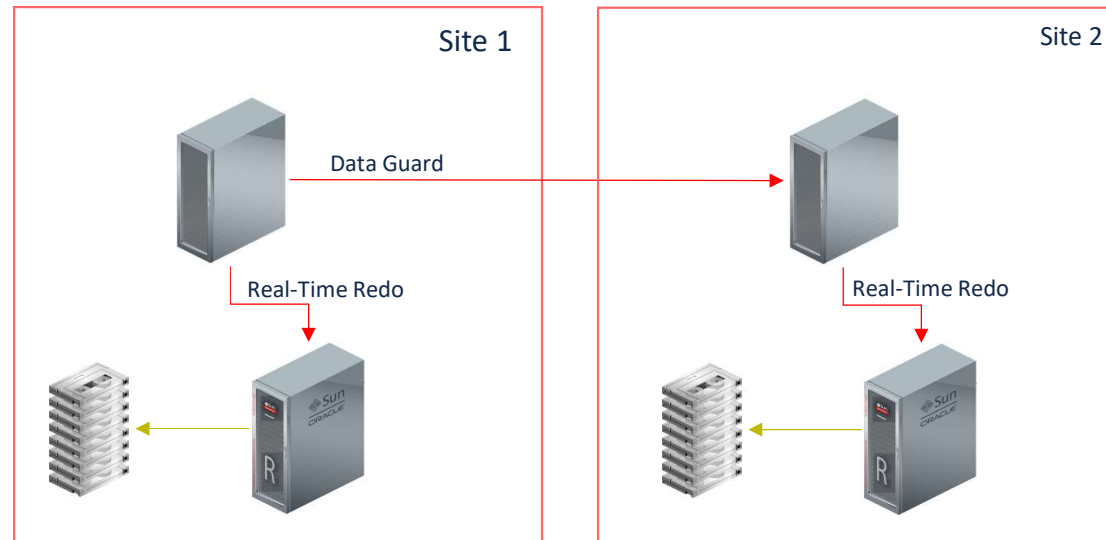
Event	RPO	RTO
Hardware error	Zero	Zero
Database Error (SW)	Zero	Can be Zero
Data Corruption	Zero	> Zero
Site Outage	Can be Zero*	> Zero

MAA – GOLD AND PLATINUM + ZDLRA

- GOLD and PLATINUM:
 - **Focus on Continuity.**
 - **Multi-Site**, Data Guard, and maybe Engineered systems.
 - **Golden not protect for simultaneous outages.**
 - Platinum adds intra and multi-region replication.
 - RTO and RPO zero:
 - Replication is done by DG, not by ZDLRA.
 - RPO zero between sites depends on the DG protection.

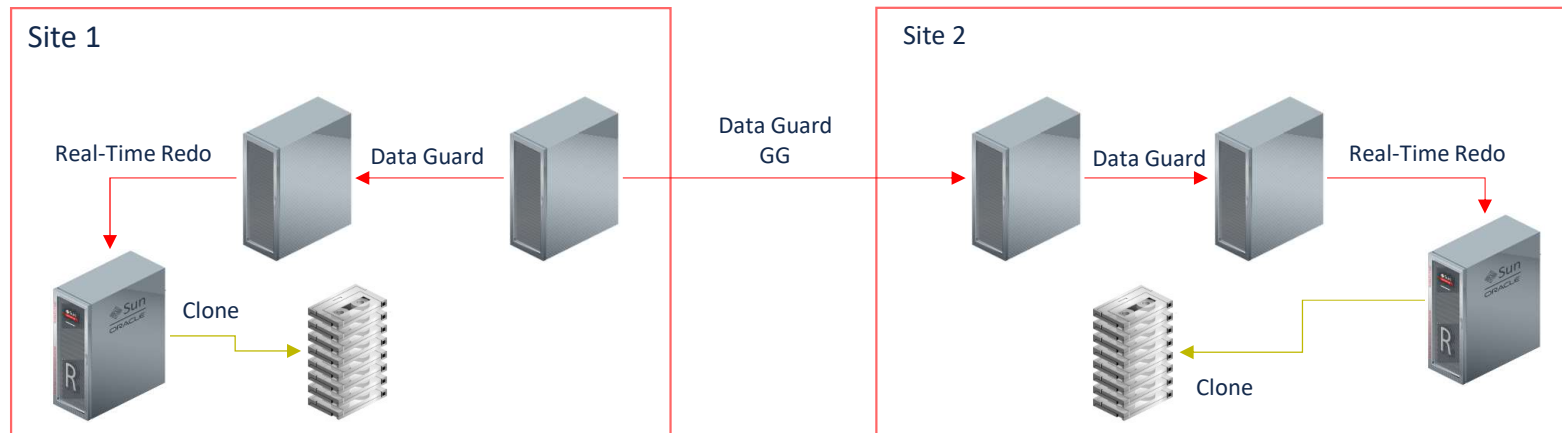


MAA – GOLD ARCHITECTURE + ZDLRA



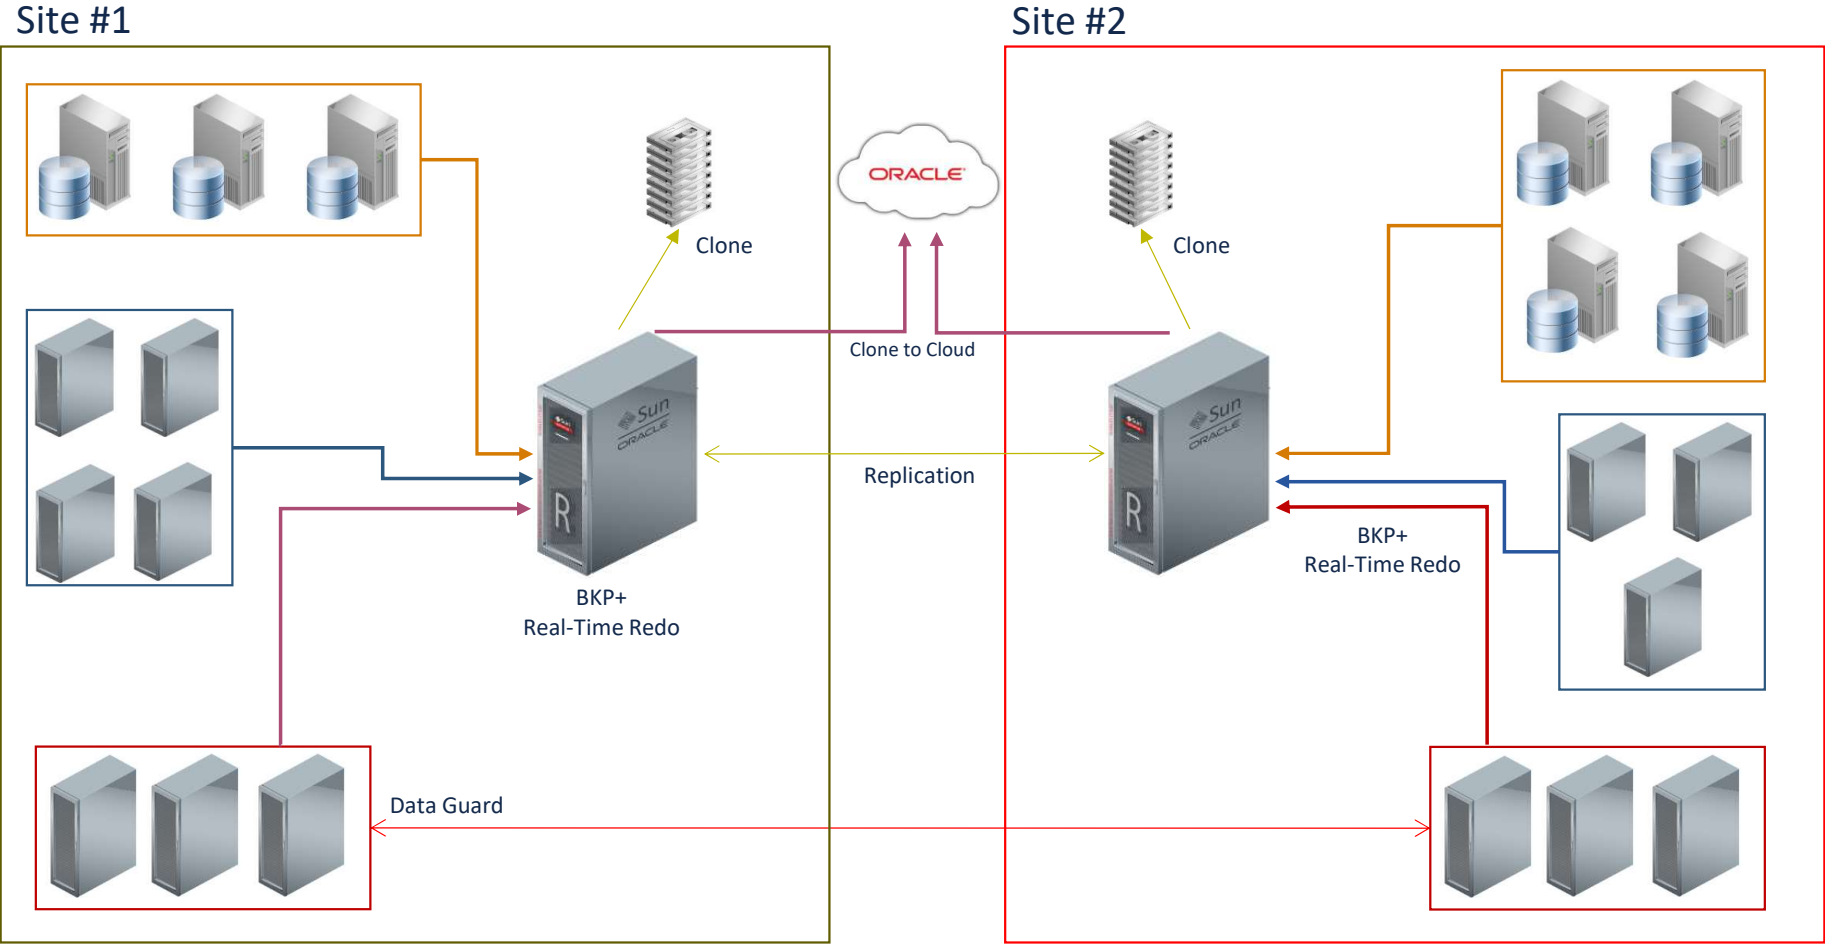
Event	RPO	RTO
Hardware error	Zero	Zero
Database Error (SW)	Zero	Zero
Data Corruption	Zero	Zero
Site Outage	Zero	Zero

MAA – PLATINUM ARCHITECTURE + ZDLRA



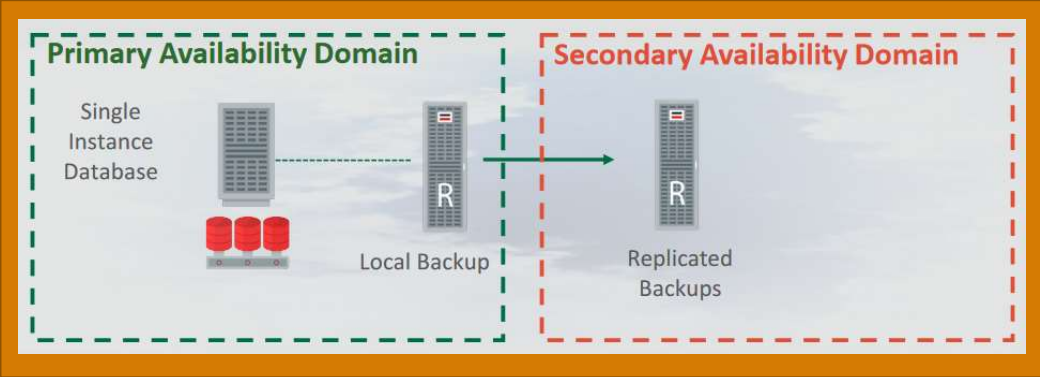
Event	RPO	RTO
Hardware error	Zero	Zero
Database Error (SW)	Zero	Zero
Data Corruption	Zero	Zero
Site Outage	Zero	Zero

MAA – EVERYTHING INTEGRATED

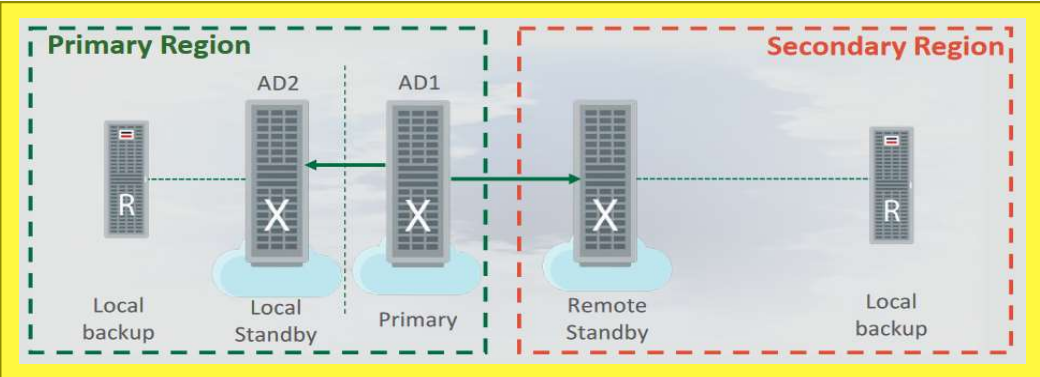
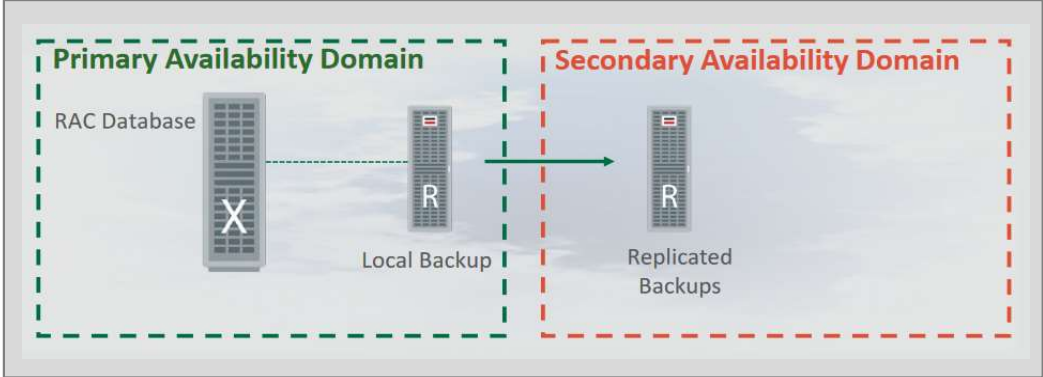


MAA - RESUME

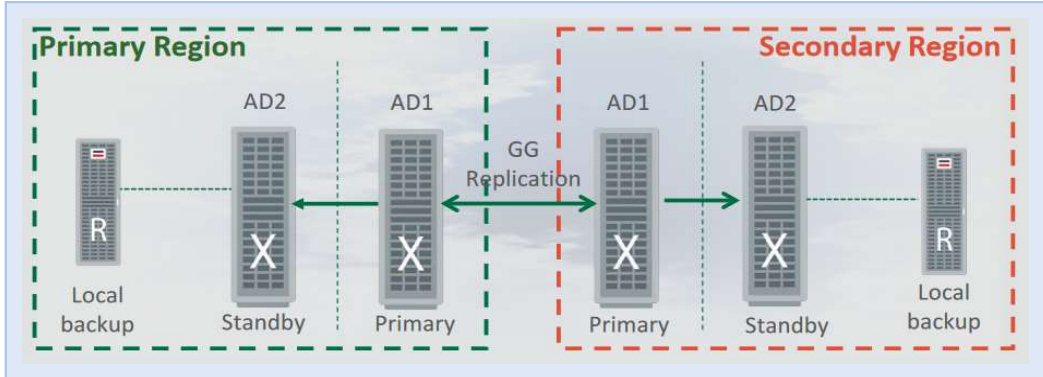
Bronze



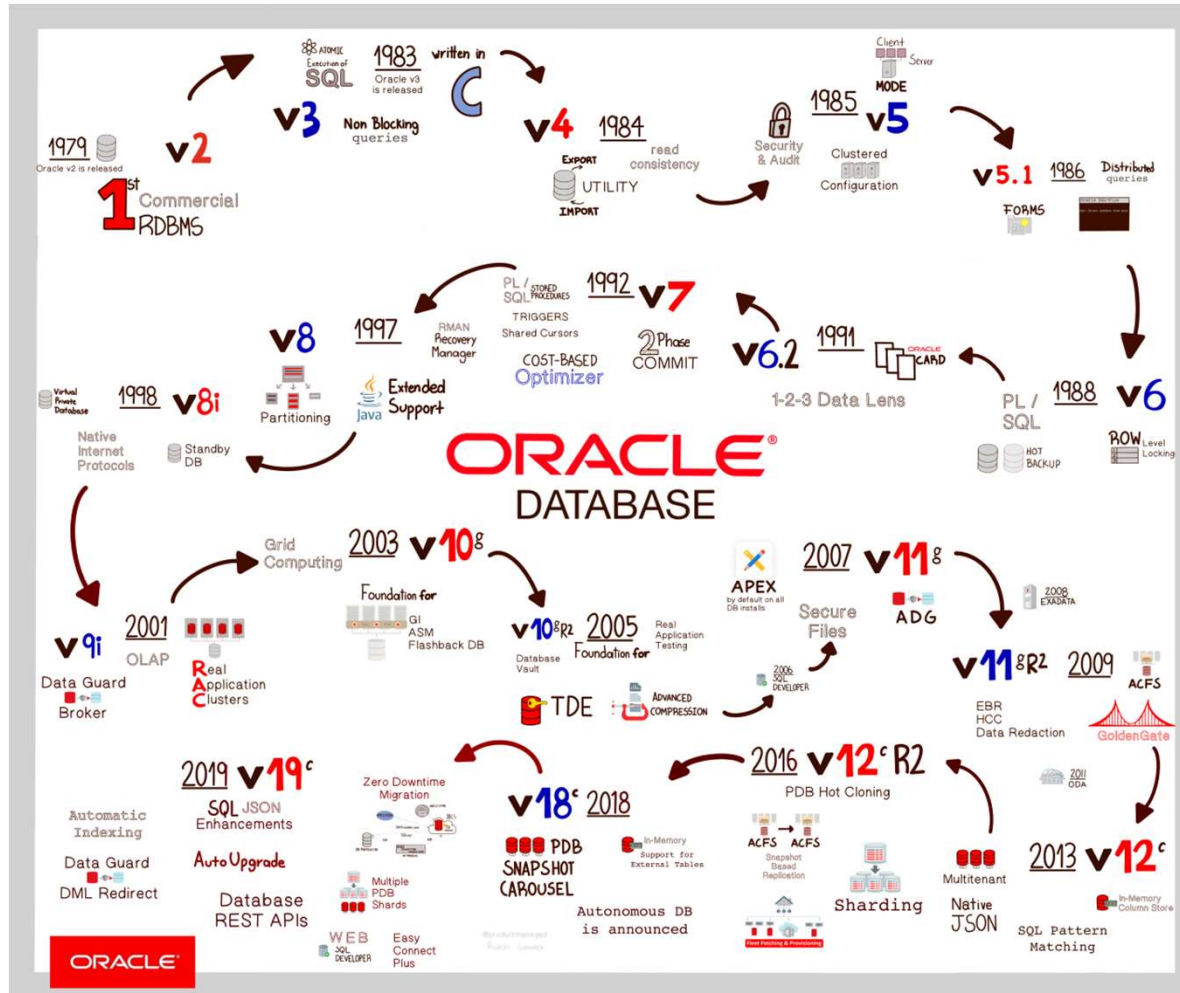
Silver



Golden



Platinum



Credit: Ricardo Gonzalez - Senior Principal Product Manager at Oracle - <https://www.linkedin.com/posts/activity-6719433973749166080-VGoR/>