

High Availability for Oracle Database SE2

Markus Flechtner

 @markusdba

 www.markusdba.net|.de

BASEL | BERN | BRUGG | BUCHAREST | DÜSSELDORF | FRANKFURT A.M. | FREIBURG I.BR. | GENEVA
HAMBURG | COPENHAGEN | LAUSANNE | MANNHEIM | MUNICH | STUTTGART | VIENNA | ZURICH

trivadis

Markus Flechtner

- Principal Consultant / Trivadis Germany GmbH
- Studied Mathematics a long time ago
- Focus
 - Oracle High Availability
 - Database Upgrade + Migration
- Teacher:
RAC, New Features, Multitenant, PostgreSQL



@markusdba



www.markusdba.net|.de



ORACLE®
ACE

DOAG

BASEL | BERN | BRUGG | BUKAREST | DÜSSELDORF | FRANKFURT A.M. | FREIBURG I.B.R. | GENÈVE
HAMBURG | KOPENHAGEN | LAUSANNE | MANNHEIM | MÜNCHEN | STUTTGART | WIEN | ZÜRICH

trivadis

trivadis

FOUNDED IN
1994

15 TRIVADIS
WORKSPACES

SWITZERLAND, GERMANY,
AUSTRIA, ROMANIA


118 MILLION
CHF
TURNOVER 

300 SLA's
(SERVICE LEVEL AGREEMENTS)

4000 
TRAINING PARTICIPANTS PER YEAR

800 
CUSTOMERS

 **700**
EMPLOYEES

5 MILLION
CHF 
BUDGET FOR SCIENCE
AND DEVELOPMENT PER YEAR

EXPERIENCE FROM
1900 PROJECTS
PER YEAR 

Now part of Accenture

trivadis

Agenda



- Introduction
- Standard Edition High-Availability (SEHA)
- High-Availability by using Refreshable PDBs
- "SEMAA"
- Summary & Further information

Introduction

HA for the Oracle Database = ?

trivadis

- Enterprise Edition
- Real Application Cluster
- DataGuard



HA for the DB Standard Edition 2

trivadis


- Real Application Cluster (until Oracle Database 18c)
 - Max. 2 nodes, 1 socket each
 - Max. 8 threads per instance
- Failover Cluster
 - Own scripts
- "DataGuard—like" third-Party-Tools
 - DBVisit Standby
 - Trivadis db*Standby
 - ... And more



From the 19c licensing information

Table 1-9 Scalability

Feature / Option / Pack	SE2	EE	EE-ES	DBCS SE	DBCS EE	DBCS EE-HP	DBCS EE-EP	ExaCS	Notes
Oracle Real Application Clusters (Oracle RAC)	N	Y	Y	N	N	N	Y	Y	EE and EE-ES: Extra cost option




 **Johannes Ahrends** @carajandb 5 Apr 2019
 I'm worried! Oracle 19c Documentation says that SE2 doesn't include RAC!
[#Documentationbug?](#)

DBCS ES	DBCS SE	DBCS EE	DBCS EE-HP	DBCS EE-EP
N	N	N	Y	Y


About 11 months later ...


MAXIMUM AVAILABILITY ARCHITECTURE

March 2, 2020

Standard Edition 2 – We Heard You! Announcing: Standard Edition High Availability

 **Markus Michalewicz**
SENIOR DIRECTOR OF PRODUCT MANAGEMENT



The diagram illustrates the SE2 High Availability architecture. It features two server racks, each labeled 'SE'. A red arrow points from the left server to the right server, and a green arrow points from the right server back to the left server, indicating bidirectional data replication. Below the servers, a central database instance is shown, represented by a stack of three red cylinders. To the left of the database is a white box labeled 'ORACLE DATABASE STANDARD EDITION 2 LICENSE'. The text 'SE2 High Availability' is positioned to the right of the database cylinders.

**What are the current
"out-of-the-box"
High Availability Solutions
for the
Oracle Database 19c Standard Edition 2?**

Standard Edition High Availability "SEHA"

Standard Edition High Availability



- **SEHA is a failover cluster database configuration for SE2 databases, based on Oracle Grid Infrastructure**
 - 2 node cluster
 - must use ASM or ACFS as database storage
 - Currently supported on Linux x86-64, Oracle Solaris on SPARC (64-bit) and Microsoft Windows
- **SEHA** was introduced with the Release Update in April 2020 (RU 19.7)
 - Due to some issues in RU 19.7, using RU 19.8 or higher is recommended
 - OCW RU must have the same version in GI home and RDBMS home
- On Windows it may serve as a replacement for Oracle Fail Safe
 - Oracle Fail Safe is desupported with Oracle Database 21c

SEHA – Licensing



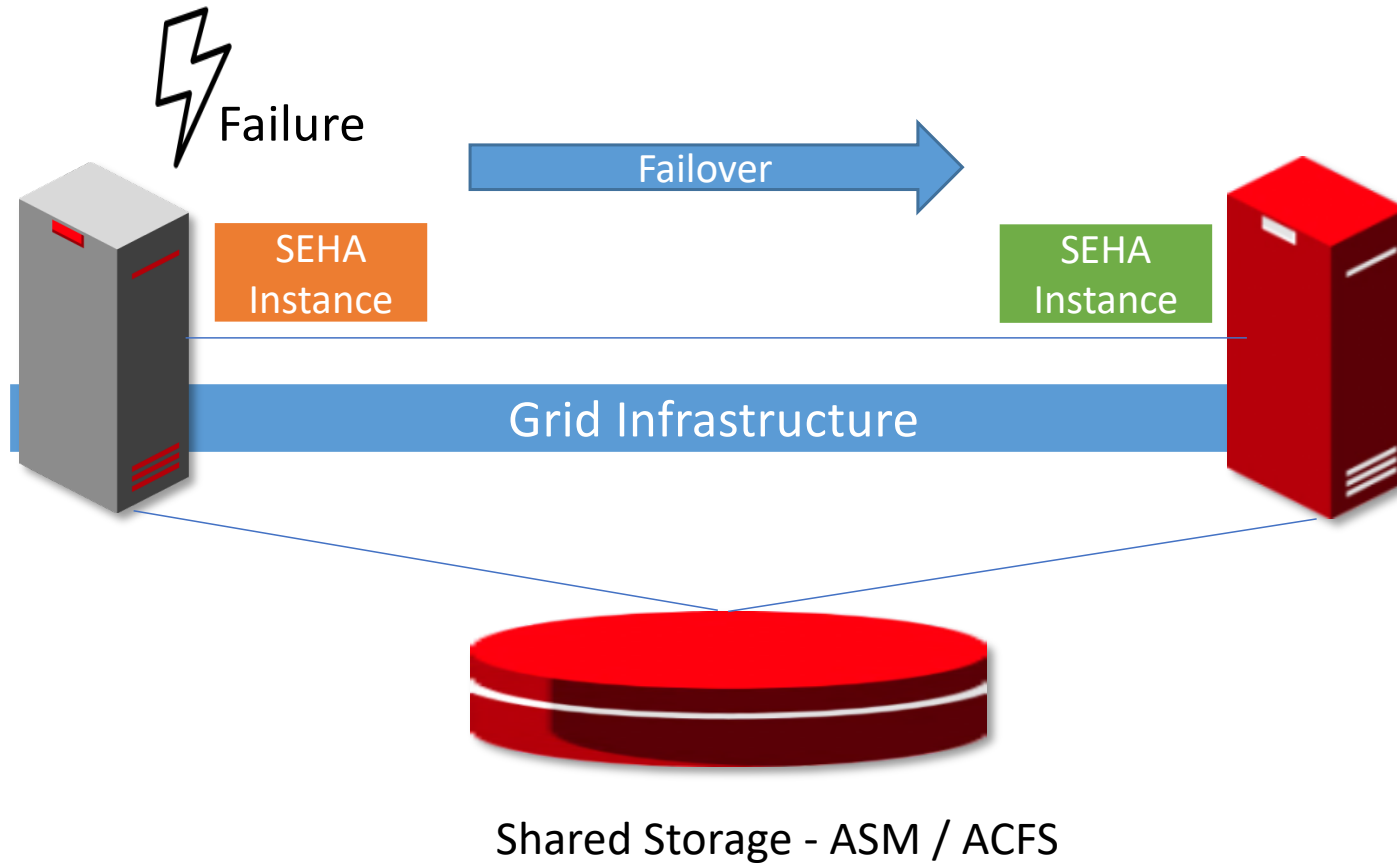
- 2 sockets per server
- 16 threads per server
- 10-days-rule applies (10 x 24 h)

Data Recovery using Clustered Environments (Failover)

The failover data recovery method is an example of a clustered deployment, where multiple nodes/servers have access to one Single Storage/SAN. In such cases your license for the programs listed on the US Oracle Technology Price List (<http://www.oracle.com/corporate/pricing/pricelists.html>) is eligible for the 10-day rule, which includes the right to run the licensed program(s) on an unlicensed spare computer in a failover environment for up to a total of ten separate 24-hour periods in any given calendar year (for example, if a failover node is down for two hours on Tuesday and three hours on Friday, it counts as two 24-hour periods). The above right only applies when a number of physical or logical machines as defined in Oracle's Partitioning Policy (detailed in <https://www.oracle.com/assets/partitioning-070609.pdf>) are arranged in a cluster and share one logical disk array located in a single data center. When the primary node fails, the failover node acts as the primary node. Once the primary node is repaired, you must either switch back or designate that repaired server as the failover node. Once the failover period has exceeded ten 24-hour periods, the failover node must be licensed. In addition, only one failover node per clustered environment is at no charge even if

- Source: <http://www.oracle.com/us/corporate/pricing/data-recovery-licensing-070587.pdf>
- More information on the "10 days rule / 10 x 24 h rule": <https://mpaage.wordpress.com/2020/08/24/failover-und-die-10-tage-regel-aeh-10x24-stunden-regel/>

SEHA – Overview



SEHA – Setup

1. Create the SE2 database on one node
 - there's nothing special in this step, e.g. you cannot specify "SEHA" in DBCA
2. Move Spfile and password file to shared storage and modify the OCR accordingly

```
$ asmcmd pcopy /u00/app/oracle/dbs/orapwSEHADB +DATA/SEHADB/orapwSEHADB  
$ srvctl modify database -db SEHADB -pwfile +DATA/SEHADB/orapwSEHADB
```

3. specify the list of available nodes for your database

```
$ srvctl modify database -db SEHADB -node node1,node2
```

- The required directories on the second node are created when the database is started for the first time

SEHA – Database Relocate



```
$ srvctl status database -db SEHADB
Instance SEHADB is running on node node2

$ srvctl relocate database -db SEHADB -node node1

$ srvctl status database -db SEHADB
Instance SEHADB is running on node node1
```

- Database is closed on node1 (SHUTDOWN IMMEDIATE) and opened on node2
- Duration depends on the number and size of open transactions (Rollback due to SHUTDOWN IMMEDIATE)
- The parameter "-timeout" and "-drain_timeout" do not have any effect on a SEHA relocate

SEHA – Database Failover



- Very simple:

In case of a node failure of the node on which the SEHA database instance is running, the Oracle Clusterware detects the failure and starts the SEHA database instance on the other node

SEHA – Client Connect



- There's nothing special
 - Use an application service for client connects (do not use the default service of the PDB)
 - Use the cluster name (SCAN name) for client connects

```
SEHA_APP.MARKUSDBA.LOCAL =  
  (DESCRIPTION =  
    (ADDRESS = (PROTOCOL = TCP) (HOST = sehacluster) (PORT=1521))  
    (CONNECT_DATA =  
      (SERVER = DEDICATED)  
      (SERVICE_NAME = SEHA_APP_SVC.markusdba.local)  
    )  
  )  
)
```

- Fast-Application-Notification (FAN) und Fast-Connection-Failover (FCF) work fine with SEHA

SEHA – Miscellaneous

- When starting a SEHA database manually, you can specify the node

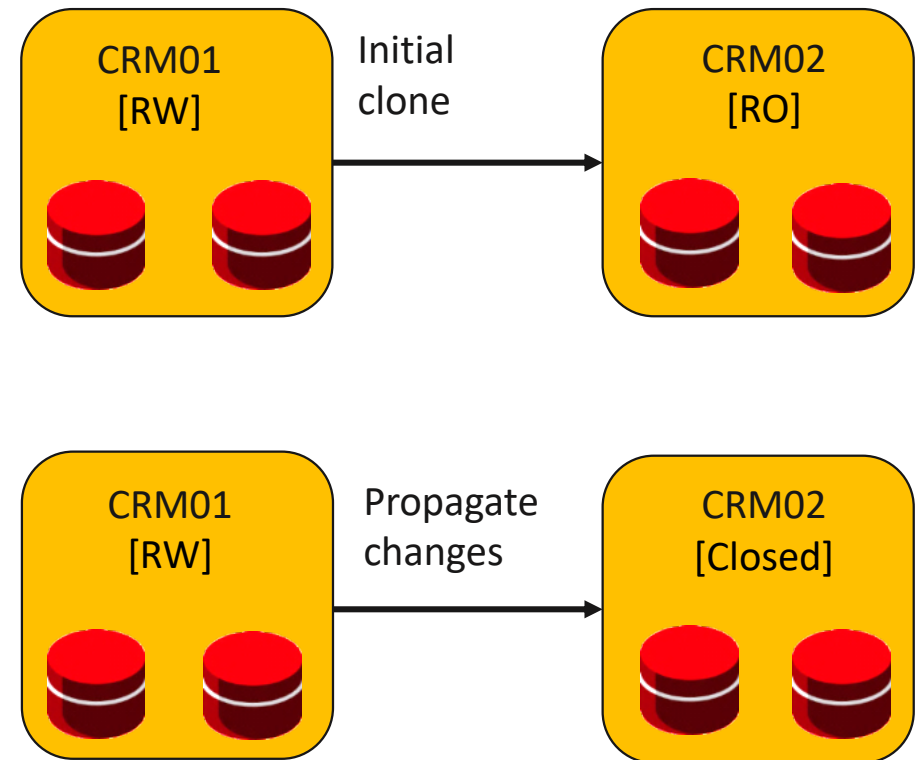
```
$ srvctl start database -db SEHADB -node node2
```

- SEHA is not "RAC One Node"! There is **always only one instance** running!
- SEHA does not work with the Enterprise Edition of the database ("PRCD-2088 : failed to configure the single instance database EEHACDB with multiple nodes because it is not a Standard Edition database")
- Application Continuity is not available with SEHA
 - AC requires Enterprise Edition + Oracle Active Data Guard option or RAC One Node or RAC option

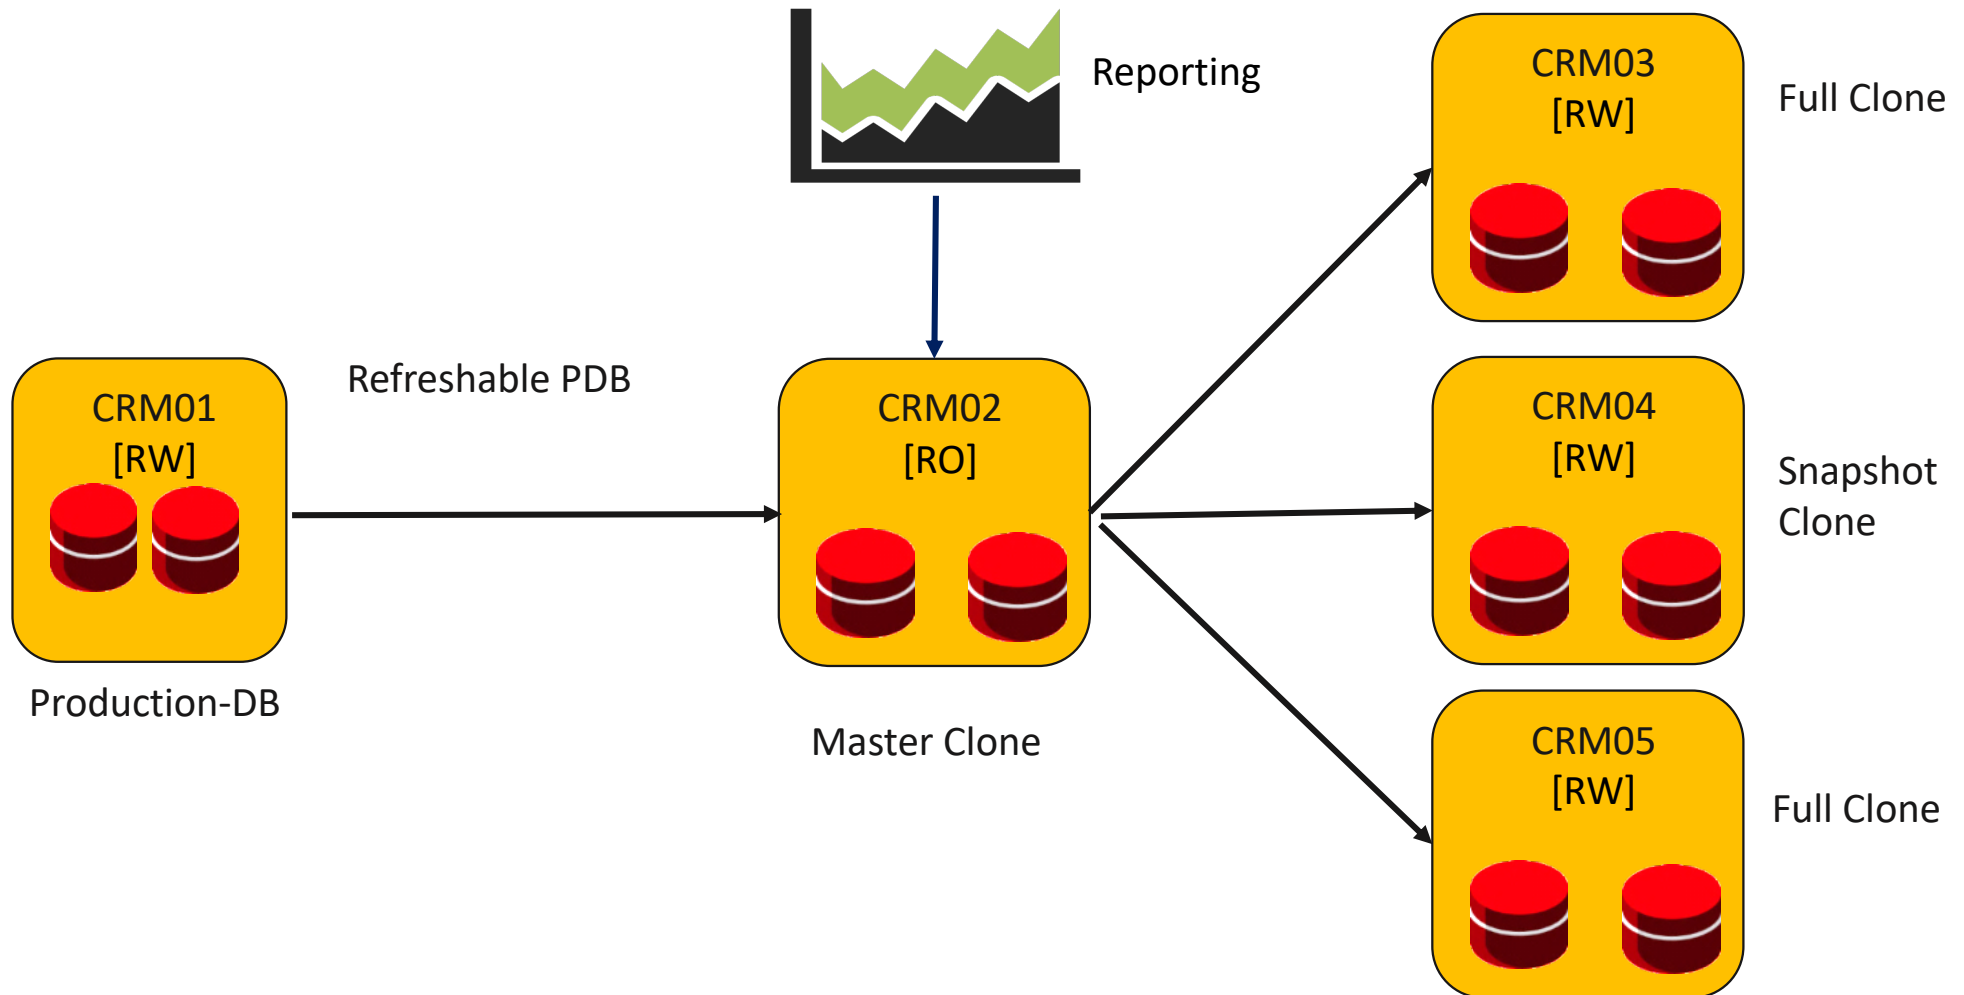
High-Availability by using Refreshable PDBs

Refreshable PDB - Basics

- **Idea: refresh a remote PDB in regular intervals**
- Use cases:
 - Use cloned PDB as source for further PDB cloning
 - Use cloned PDB as read-only-PDB for reporting
 - "High-Availability" for SE2-PDBs
- Cloned PDB can be opened in read-only mode only
- Refresh can be done manually or automatically
- Source and target PDB must be in different CDBs
- Target PDB must be closed for refresh
- Refreshable PDB can be based on a Non-CDB
- Available since Oracle Database 12c Release 2



Refreshable PDB – "normal" Use Cases



Refreshable PDB - Setup



- Source CDB must use local Undo

1. Create a common user in the source CDB

```
SQL> CREATE USER C##REFRESH identified by <password>;  
SQL> GRANT CREATE SESSION,RESTRICTED SESSION,SYSOPER to C##REFRESH  
2 CONTAINER=ALL;  
SQL> GRANT CREATE PLUGGABLE DATABASE TO C##REFRESH;
```

2. Create a public database link from the target CDB to the source CDB

- Instead of using a database to the source CDB you can use a database link which points to the source PDB

```
SQL> CREATE PUBLIC DATABASE LINK <SOURCE_CDB>  
2 CONNECT TO C##REFRESH IDENTIFIED BY '<PASSWORD>' USING '<TNS>';
```

3. Create a refreshable PDB and specify the refresh mode

```
CREATE PLUGGABLE DATABASE <PDB> FROM <PDB>@<SOURCE_CDB>  
file_name_convert= ... REFRESH MODE EVERY 240 MINUTES;
```

Refreshable PDB – Refresh Mode

- Possible refresh Modes
 - REFRESH MODE EVERY <n> MINUTES
 - REFRESH MODE MANUAL
 - REFRESH NONE
- If the PDB is not closed for a scheduled refresh, the refresh will be skipped

```
SQL> ALTER PLUGGABLE DATABASE <PDB> CLOSE IMMEDIATE;  
SQL> ALTER PLUGGABLE DATABASE <PDB> REFRESH;
```

- You cannot open a refreshable PDB "read-write":
"ORA-65341: cannot open pluggable database in read/write mode"

Refreshable PDB – Data Dictionary

```
SQL> select PDB_NAME, REFRESH_MODE, REFRESH_INTERVAL, LAST_REFRESH_SCN  
2      from cdb_pdbs;
```

PDB_NAME	REFRES	REFRESH_INTERVAL	LAST_REFRESH_SCN
PDB\$SEED	NONE		
SE2PDB	AUTO	1	1672347

From Source DB

```
SQL> select scn_to_timestamp(1672347) from dual@<source_db>;
```

```
SCN_TO_TIMESTAMP(1672347)
```

```
-----  
16-AUG-20 01.07.16.000000000 PM
```

Refreshable PDB – alert.log



```
SE2PDB(4):alter pluggable database refresh
2020-08-16T11:58:12.076936+02:00
Applying media recovery for pdb-4099 from SCN 1672311 to SCN 1672347
Remote log information: count-1
thr-1,seq-20,logfile-
/u02/fast_recovery_area/SE2CDBA/foreign_archivelog/SE2PDB/2020_08_16/o1_mf_1_20_hml0smw3_
.arc,los-1592646,nxs-18446744073709551615,maxblks-366859
SE2PDB(4):Media Recovery Start
2020-08-16T11:58:12.077338+02:00
SE2PDB(4):Serial Media Recovery started
SE2PDB(4):max_pdb is 5
2020-08-16T11:58:12.154566+02:00
SE2PDB(4):Media Recovery Log
/u02/fast_recovery_area/SE2CDBA/foreign_archivelog/SE2PDB/2020_08_16/o1_mf_1_20_hml0smw3_
.arc
2020-08-16T11:58:12.665203+02:00
SE2PDB(4):Incomplete Recovery applied until change 1672347 time 08/16/2020 11:58:10
2020-08-16T11:58:12.668710+02:00
SE2PDB(4):Media Recovery Complete (SE2CDBB)
SE2PDB(4):Completed: alter pluggable database refresh
```

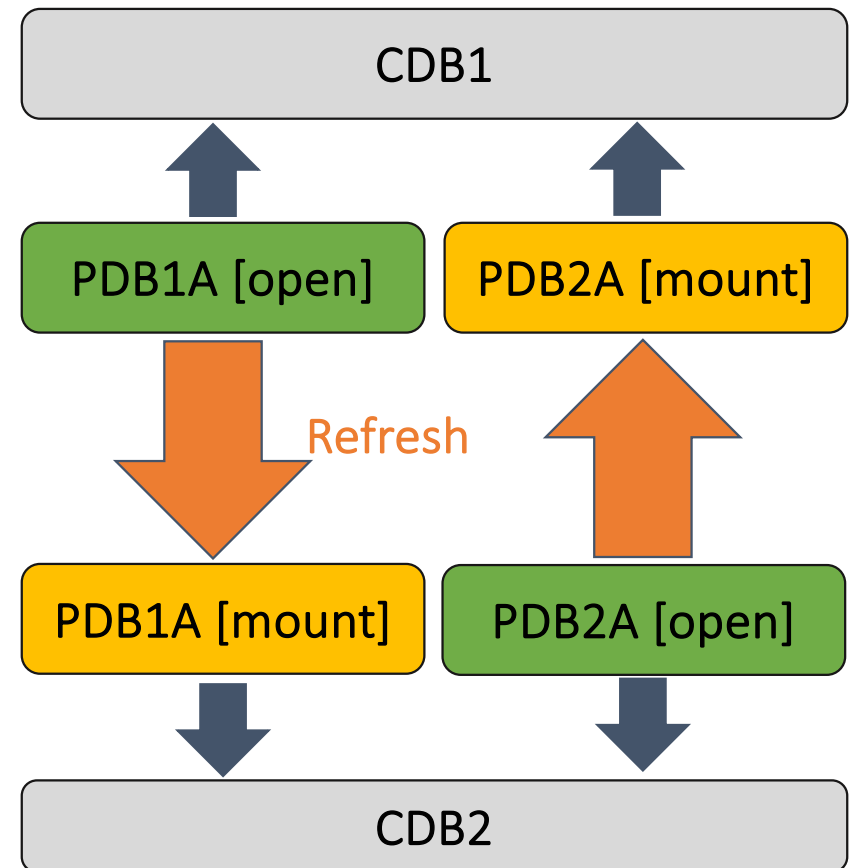
Refresh Mechanism

- For a refresh, the redolog information is used
 - Redo information is accessed via the database link
- If you are using longer refresh intervals, the archivelog files must be available on the source CDB
 - Not an issue for "HA-configurations" with 1-minute-interval 😊
 - Keep this in mind if you open a refreshable PDB temporarily read only so that it cannot be refreshed
- You can specify "REMOTE_RECOVERY_FILE_DEST" in the refreshable PDB as a location from where the target CDB can read the redolog files

- RMAN does not backup a refreshable PDB

Refreshable PDB for HA

- Create a refreshable "Standby PDB" in another CDB
- "Standby PDB" is refreshed at a high frequency (minimum refresh interval is 1 minute)
- You can open the "Standby-PDB" read-only e.g. for reporting
 - !! This will cause a gap
- Disable the refresh and open the "Standby-PDB" in case of a failover
- **Idea:**
mutual protection of 2 CDBs resp. 2 PDBs
- Unfortunately, there is no "observer"



Refreshable PDB for HA -Preparation



1. Create a "Standby-CDB"
2. Create a refreshable PDB on the "Standby-CDB" – refresh interval 1 minute

```
SQL> CREATE PLUGGABLE DATABASE PDB04
2 FROM PDB04@CDB1.markusdba.local
3 file_name_convert=
4 ('/u01/oradata/CDB1/PDB04', '/u01/oradata/CDB2/PDB04')
5 REFRESH MODE EVERY 1 MINUTES;
```

3. Keep this PDB closed (mounted) for automatic refresh

Refreshable PDB for HA – Failover (1)



- **(in case of a failover) you cannot use the refreshable PDB as a regular PDB**
 - There's still a "connection" to the source PDB, but the source PDB is not available anymore ☹️
 - The refreshable PDB cannot be opened in read/write-mode, but in read-only-mode only
- **You have to clone the refreshable PDB first**
 - Regular "full" clone (takes longer, needs more disk space)
 - Snapshot clone (almost immediately available, later on it can be merged to a regular PDB) (!! Does not work with ASM)
- "Reinstantiation":
 - As soon as the other CDB (= original source CDB) is up and running again, you have to drop the PDB there and create a refreshable PDB in the other direction
 - **!! When the "original source PDB" comes up again, the application service will start there, too.**
- There is no "Switchover" for SE2 ("Refreshable PDB Switchover" is available for Exadata and Oracle Cloud only)

Refreshable PDB for HA – Failover (2)

- Cloning a PDB

```
SQL> ALTER PLUGGABLE DATABASE <pdb_name> OPEN READ ONLY;  
SQL> CREATE PLUGGABLE DATABASE <work_pdb> from <pdb_name>  
2 FILE_NAME_CONVERT=(...);
```

- Create a “snapshot clone”

```
SQL> ALTER PLUGGABLE DATABASE <pdb_name> OPEN READ ONLY;  
SQL> CREATE PLUGGABLE DATABASE <work_pdb> from <pdb_name>  
2 FILE_NAME_CONVERT=(...) SNAPSHOT COPY;
```

- Later on you can make the snapshot clone a “regular PDB”

```
SQL> ALTER PLUGGABLE DATABASE <work_pdb> MATERIALIZED  
2 FILE_NAME_CONVERT=(...);
```

Client Connect

- Use an application service for client connects (do not use the default service of the PDB)
- Place both hosts in the client connect

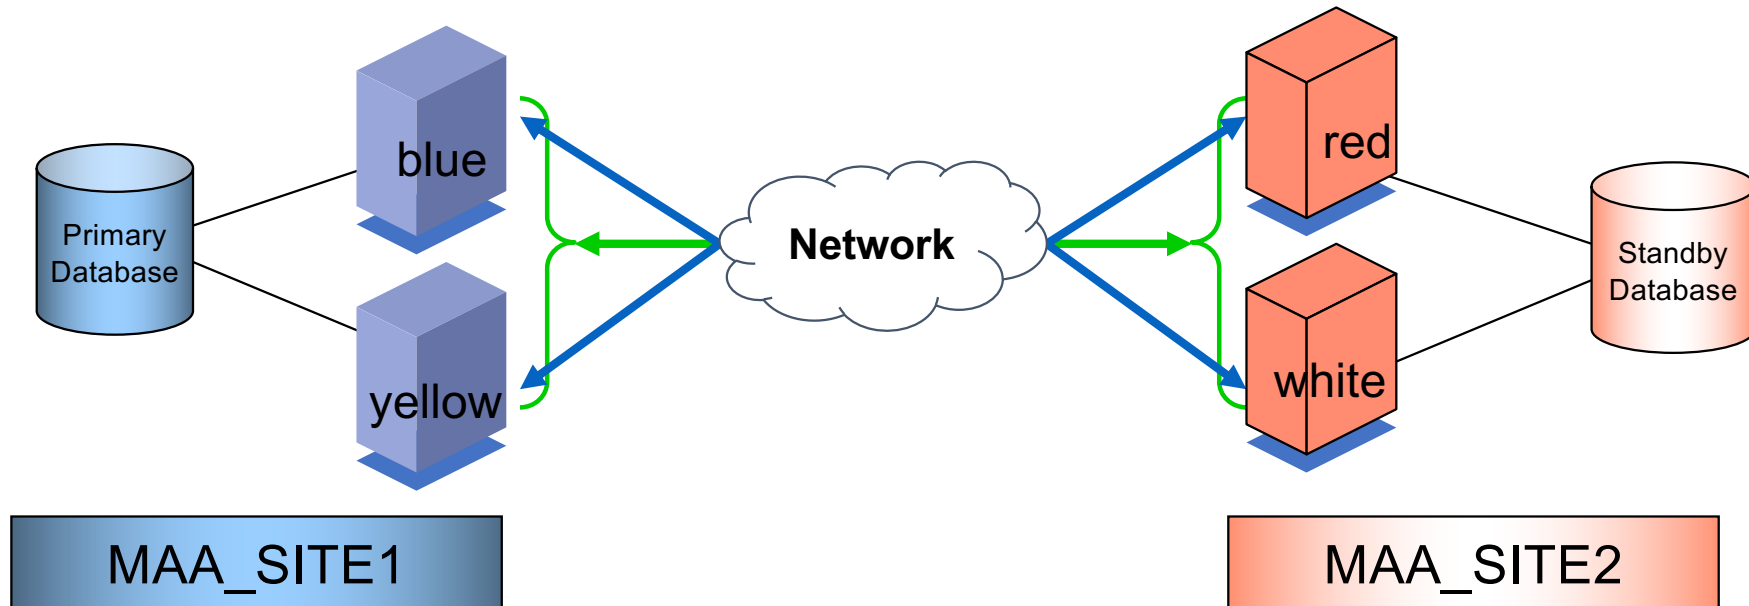
```
PDB_APP.MARKUSDBA.LOCAL =  
  (DESCRIPTION =  
    (RETRY_COUNT=20) (RETRY_DELAY=3)  
    (ADDRESS_LIST=  
      (LOAD_BALANCE = OFF )  
      (ADDRESS = (PROTOCOL = TCP) (HOST = node1) (PORT=1521))  
      (ADDRESS = (PROTOCOL = TCP) (HOST = node2) (PORT=1521))  
    )  
    (CONNECT_DATA =  
      (SERVER = DEDICATED)  
      (SERVICE_NAME = PDB_APP_SVC.markusdba.local)  
    )  
  )
```


"SEMAA"

Maximum Availability Architecture

trivadis

- Combination of RAC and DataGuard
- Provides service high availability (RAC) and data high availability (DataGuard)

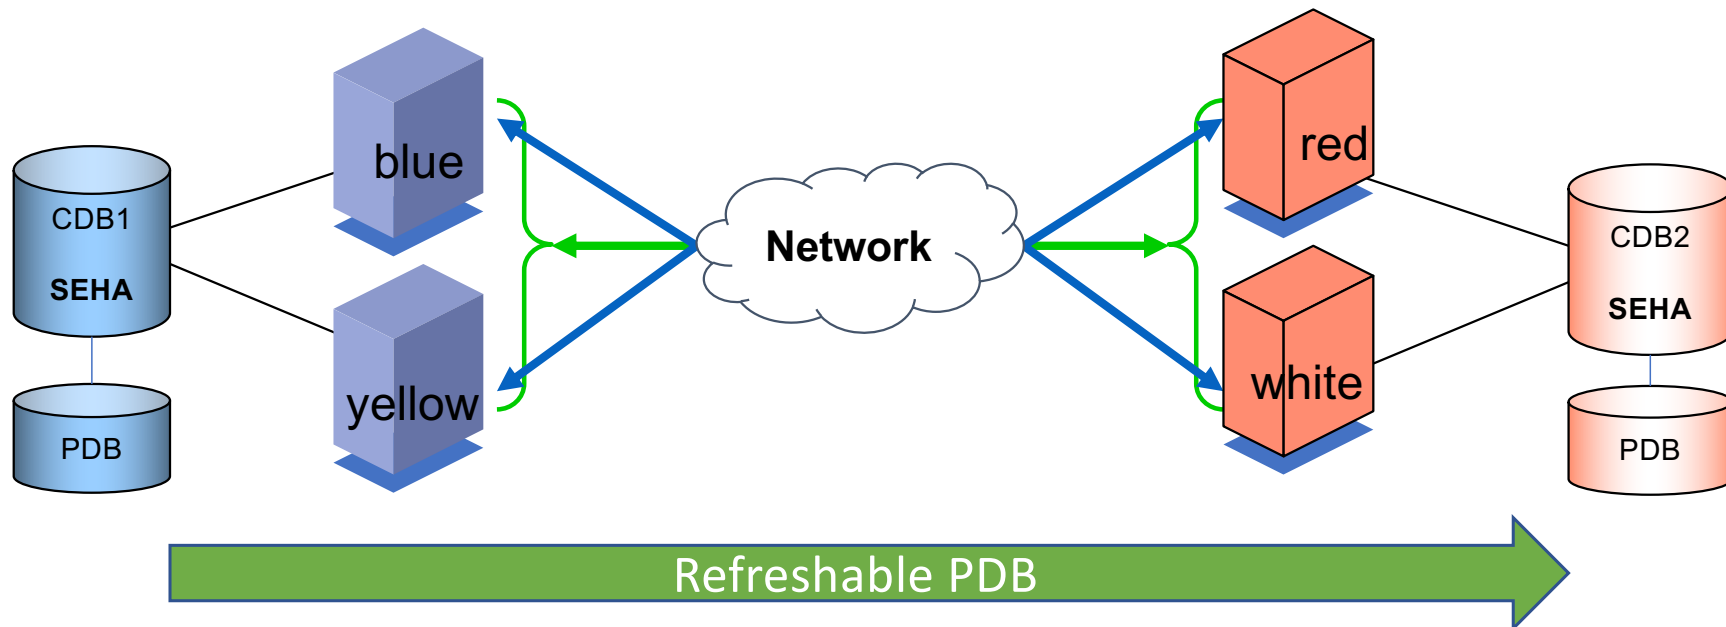


"SEMAA"

trivadis

- Combination of SEHA and Refreshable PDB

\$\$ €€



"SEMAA" - Setup



- The same rules apply as with SEHA and Refreshable PDB
- Procedure
 1. Create "Primary CDB"
 2. Make "Primary CDB" a "SEHA-CDB"
 3. Create "Standby CDB"
 4. Make "Standby CDB" a "SEHA-CDB"
 5. Create "Primary PDB"
 6. Create the "Standby-PDB" as a refreshable clone of the "Primary-PDB"
- Client Connect
 - Use an application service for client connects (not the default service of the PDB)
 - Use both cluster names (SCAN names) in TNS-Connect

Summary & Further Information

Summary

- "SEHA" and "Refreshable PDBs" offer High Availability solutions for the Oracle Database 19c SE2 **without additional costs**
- SEHA
 - Failover Cluster "out of the box"
 - Easy setup
- "Refreshable PDBs"
 - Max. one minute Data Loss (in theory)
 - No Switchover ☹️
 - No immediate use of the refreshable PDB in case of a failover (cloning required) ☹️
- **It's not Enterprise Edition like High Availability, but it's sufficient for a lot of customers**

Further Information (1)



- **Oracle Documentation**

- Documentation - Oracle Database 19c - Database Installation Guide - Chapter 10 "Installing Standard Edition High Availability"
- Documentation – Oracle Database 19c "SEHA – Known issues"
<https://docs.oracle.com/en/database/oracle/oracle-database/19/rnrdb/linux-platform-issues.html#GUID-E462F73F-E37B-4543-B55F-F5CD9543D417>
- Documentation – Oracle Database 19c – Multitenant Administrators Guide – Section 7.5 "About Refreshable Clone PDBs"

- **MOS-Notes**

- Some Data File Missing From DBA_DATA_FILES or CDB_DATA_FILES in Refreshable PDB (Doc ID 2610743.1)
- How to create refreshable PDB (Doc ID 2441949.1)
- RMAN Skipping Backup of Refreshable PDB (Doc ID 2567138.1)

Further Information (2)



- **Other Oracle Sources**

- Standard Edition 2 – We Heard You! Announcing: Standard Edition High Availability - <https://blogs.oracle.com/maa/standard-edition-2-announcing-standard-edition-high-availability>
- Licensing Data Recovery Environments: <http://www.oracle.com/us/corporate/pricing/data-recovery-licensing-070587.pdf>

- **Blog Posts**

- <https://www.carajandb.com/blog/2020/standard-edition-high-availability/>
- <https://blog.ordix.de/neue-hochverfuegbarkeitsloesung-standard-edition-high-availability-seha-fuer-die-standard-edition-2-se2>
- <https://www.markusdba.net/2020/05/11/standard-edition-high-availability-first-impressions/>
- <https://gotodba.com/2020/09/02/seha-standard-edition-high-availability-in-oracle-db/>

Questions & Answers

Markus Flechtner

markus.flechtner@trivadis.com

Phone +49 211 5866 64725

*Credits to Markus Michalewicz and Johannes Ahrends
for their support when preparing this presentation*



@markusdba



www.markusdba.net|.de

BASEL | BERN | BRUGG | BUKAREST | DÜSSELDORF | FRANKFURT A.M. | FREIBURG I.B.R. | GENÈVE
HAMBURG | KOPENHAGEN | LAUSANNE | MANNHEIM | MÜNCHEN | STUTTGART | WIEN | ZÜRICH

trivadis



Making a **WORLD** possible
in which **intelligent IT**
facilitates **LIFE and WORK** as a
matter of course.