



Oracle Database 23ai & Exadata Exascale


O Futuro do Banco de Dados na Nuvem

GUOB Tech Day

Maicon Carneiro

São Paulo, 23 de agosto de 2025





Maicon Carneiro
Oracle ACE Pro

Trajatória com tecnologia Oracle

- *Bacharel em Sistemas de Informação – UniFTC (2015 ~2019)*
- *Trabalha com Oracle Database, SQL & PL/SQL desde 2015*
- *Atua como DBA em consultorias desde 2019*
- *Escreve posts técnicos no Blog do Dibiei desde 2019 (dibiei.blog)*
- *Membro do conselho do Grupo de Usuários Oracle do Brasil (GUOB) desde 2022*
- *Reconhecido como Oracle ACE PRO (Advocate Community Expert) desde 2022*

<https://dibiei.blog>

<https://www.linkedin.com/in/maiconcarneiro>

<https://ace.oracle.com/ords/ace/profile/maiconcarneiro>

Certificações

Oracle Database

ORACLE

Certified Associate

Oracle Database 12c
Administrator



ORACLE

Certified Professional



ORACLE

Certified Expert

Oracle Database 12c
Data Guard Administrator



ORACLE

Certified Expert

Oracle Database 12c:
Oracle RAC and Oracle
Grid Infrastructure
Administrator



ORACLE

Certified Expert

Oracle Database 12c
Maximum Availability
Architecture



Oracle Cloud



Edições Anteriores: GUOB 2023 e 2024

2023 - Como Usar Oracle RAT no Autonomous Database



2024 - Fleet Patching and Provisioning (FPP)





Become an Oracle ACE

Join 500+ technical experts helping peers globally



Members get many benefits:

- Direct **access** to product development
- Global **networking** events
- Exclusive **training** and content
- Cool **swag**, **cloud credits**, and Oracle CloudWorld pass
- **Your profile** listed on oracle.com

And much more! Go to: ace.oracle.com



ACE Member Benefits



Key Benefits

Cool **swag***, Digital awards for social media, Oracle **CloudWorld pass***, & more



Exclusive Content

Exclusive **monthly virtual meetings** with product development teams + engaging guest speakers



Direct Access to Product Management

Multiple **direct communication channels** to product management and fellow ACEs



Networking

In-person & virtual networking opportunities for ACEs to **connect with product development** and each other.




Cloud Account

\$5k USD Cloud account*



Travel Support

ACE Directors are eligible for travel support to give presentations or lead workshops at conferences globally




Agenda

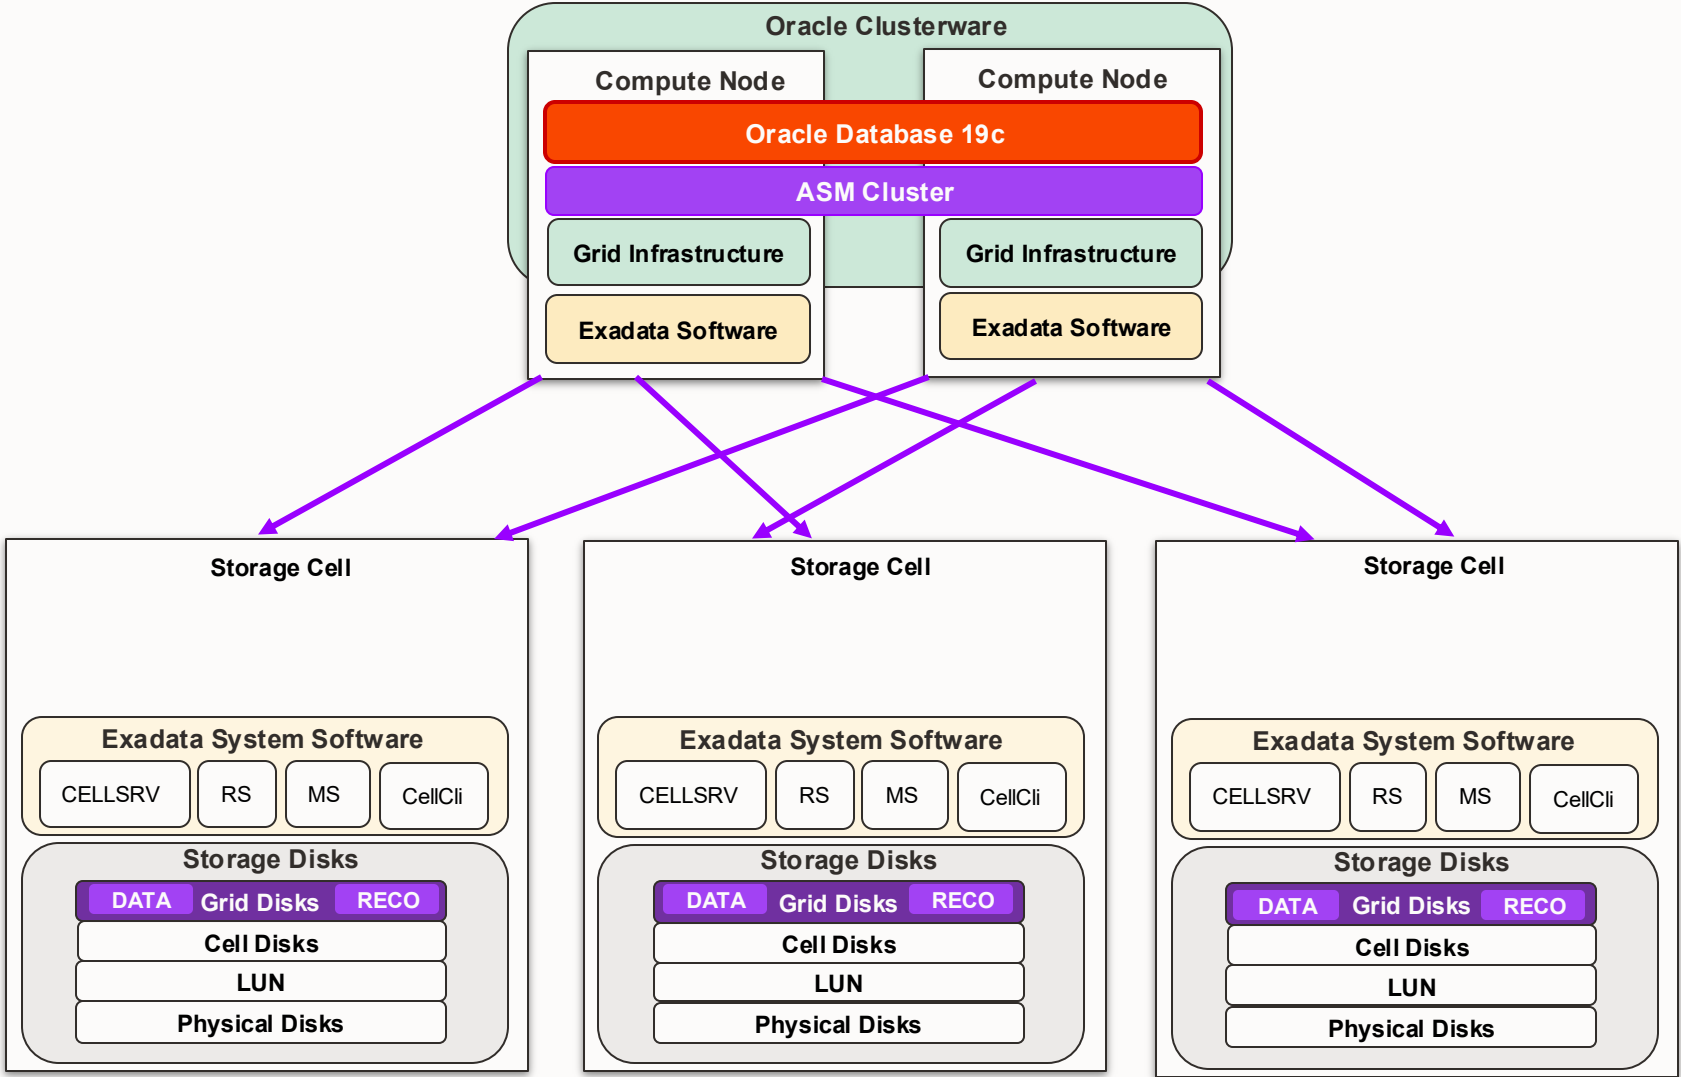

- Exadata Architecture Overview
- Exascale Architecture Overview
- VM Clusters on Exascale
- Databases Files on Exascale
- Exascale Intelligent Snapshot Clones
- Exascale Database Aware PDB Thin Clone
- Exascale CDB Thin Clone with gDBClone
- Monitoring Exascale Storage Vault
- Exascale Deployment Options

Exadata – ASM Architecture

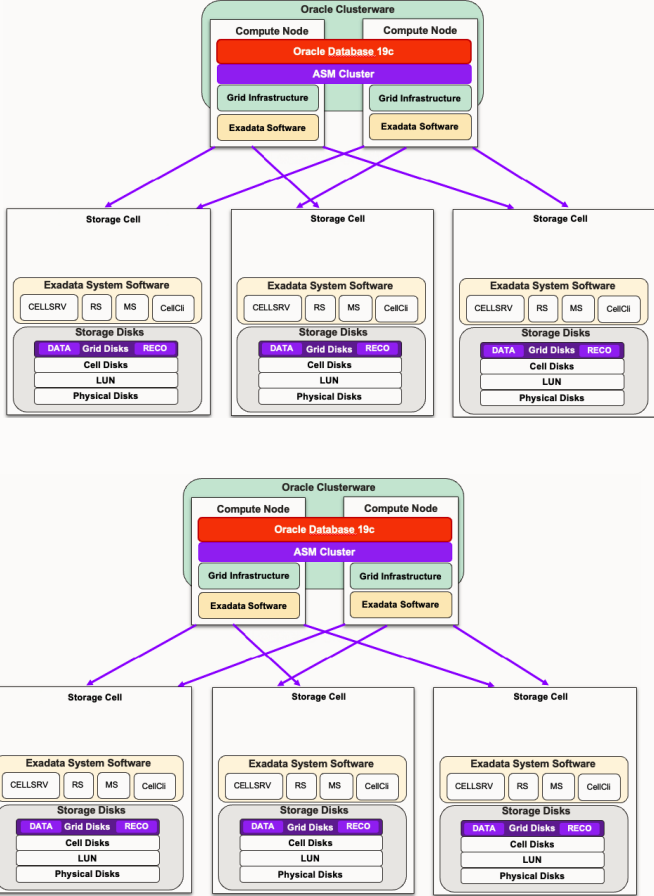
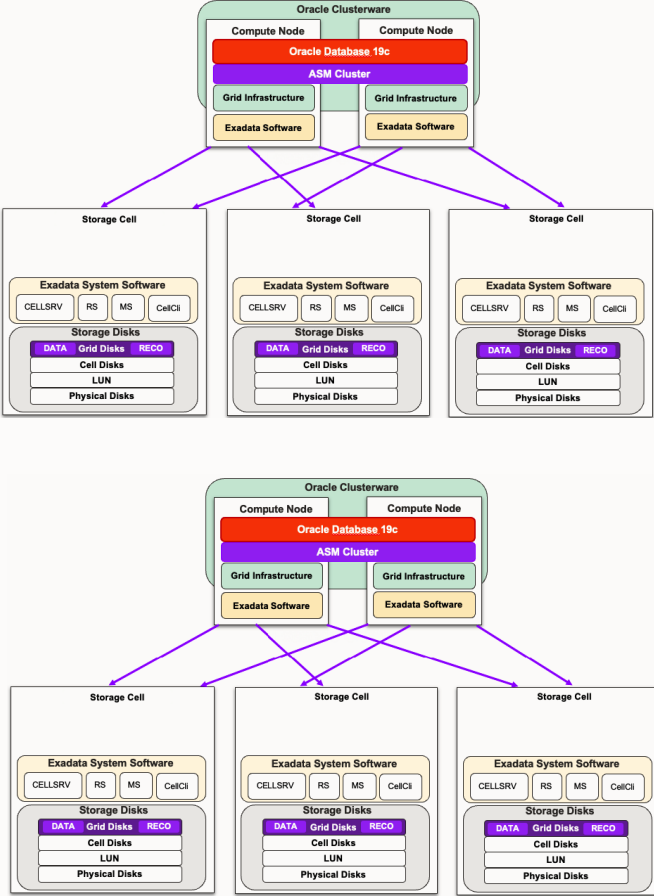
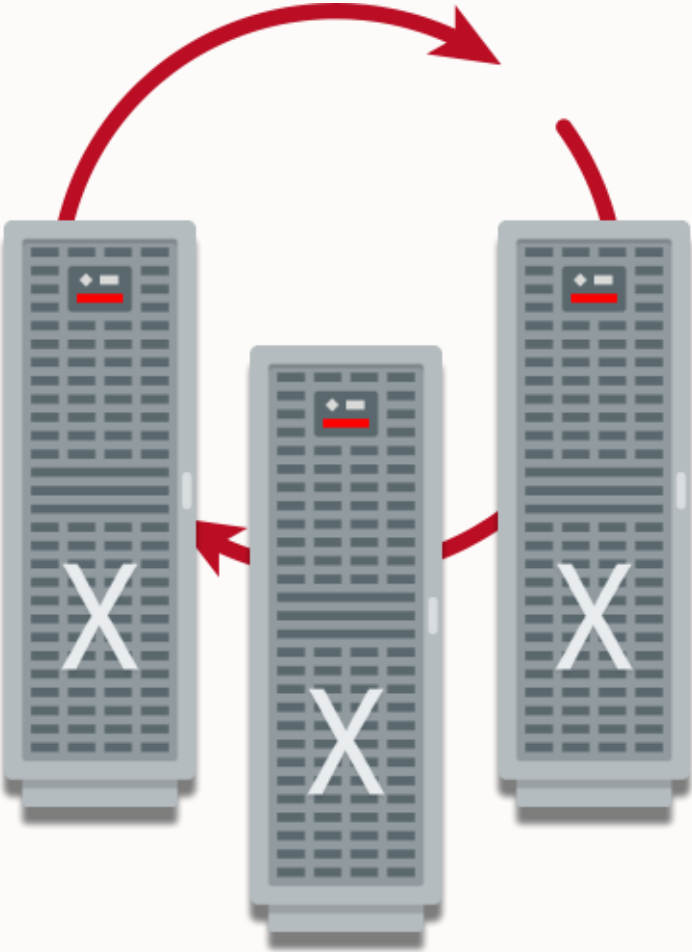
How this worked for 15 Years ...



Exadata Architecture



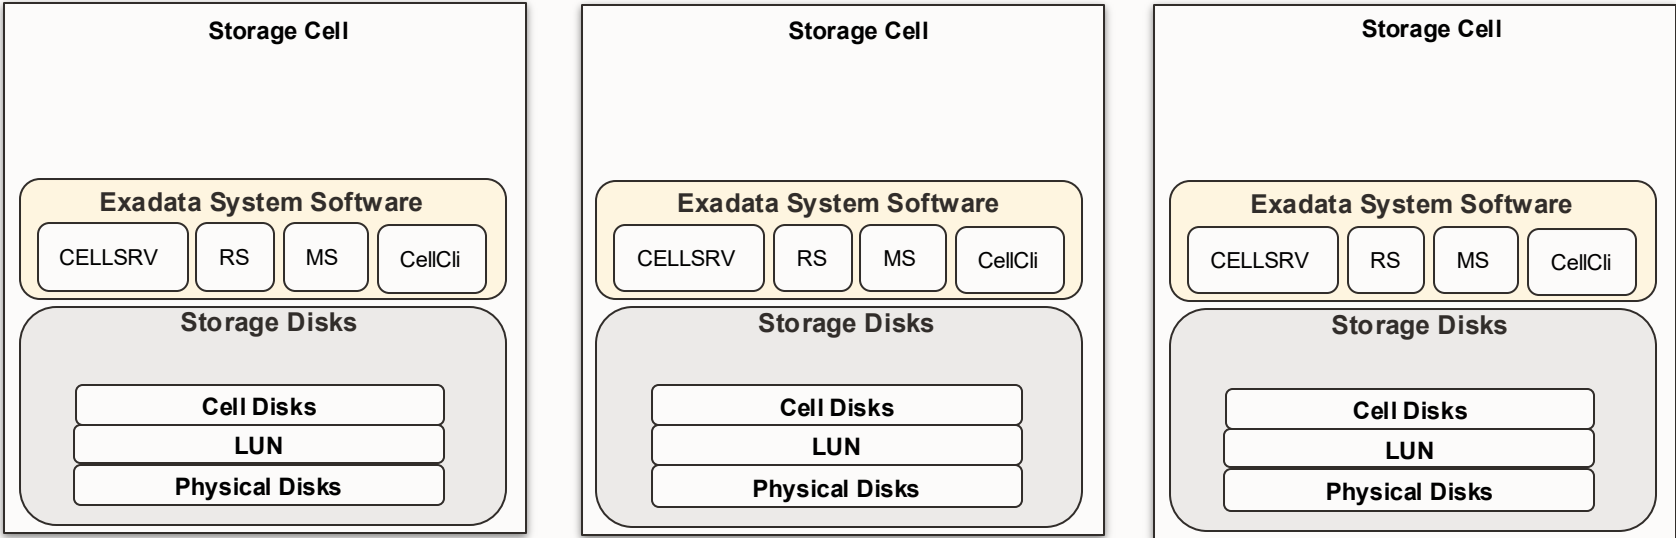
Exadata Architecture – Multiple Clusters



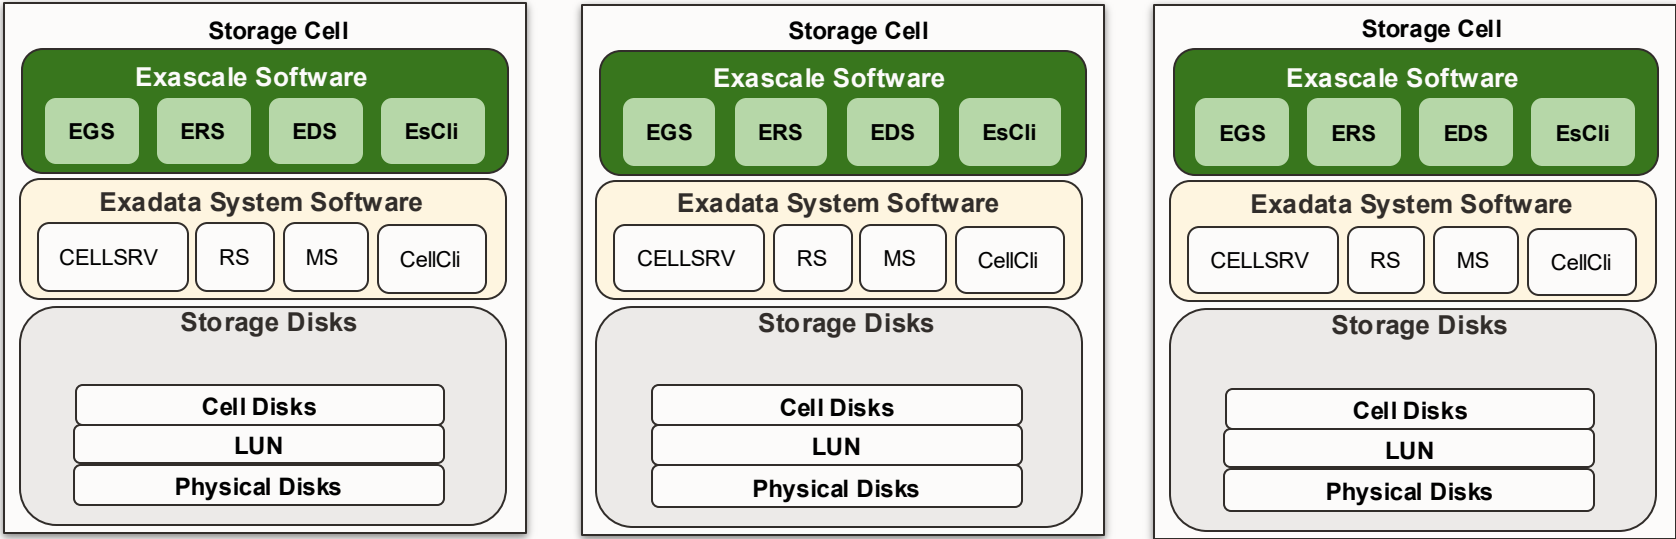
Exascale Architecture

The future of Oracle Database on the Cloud !

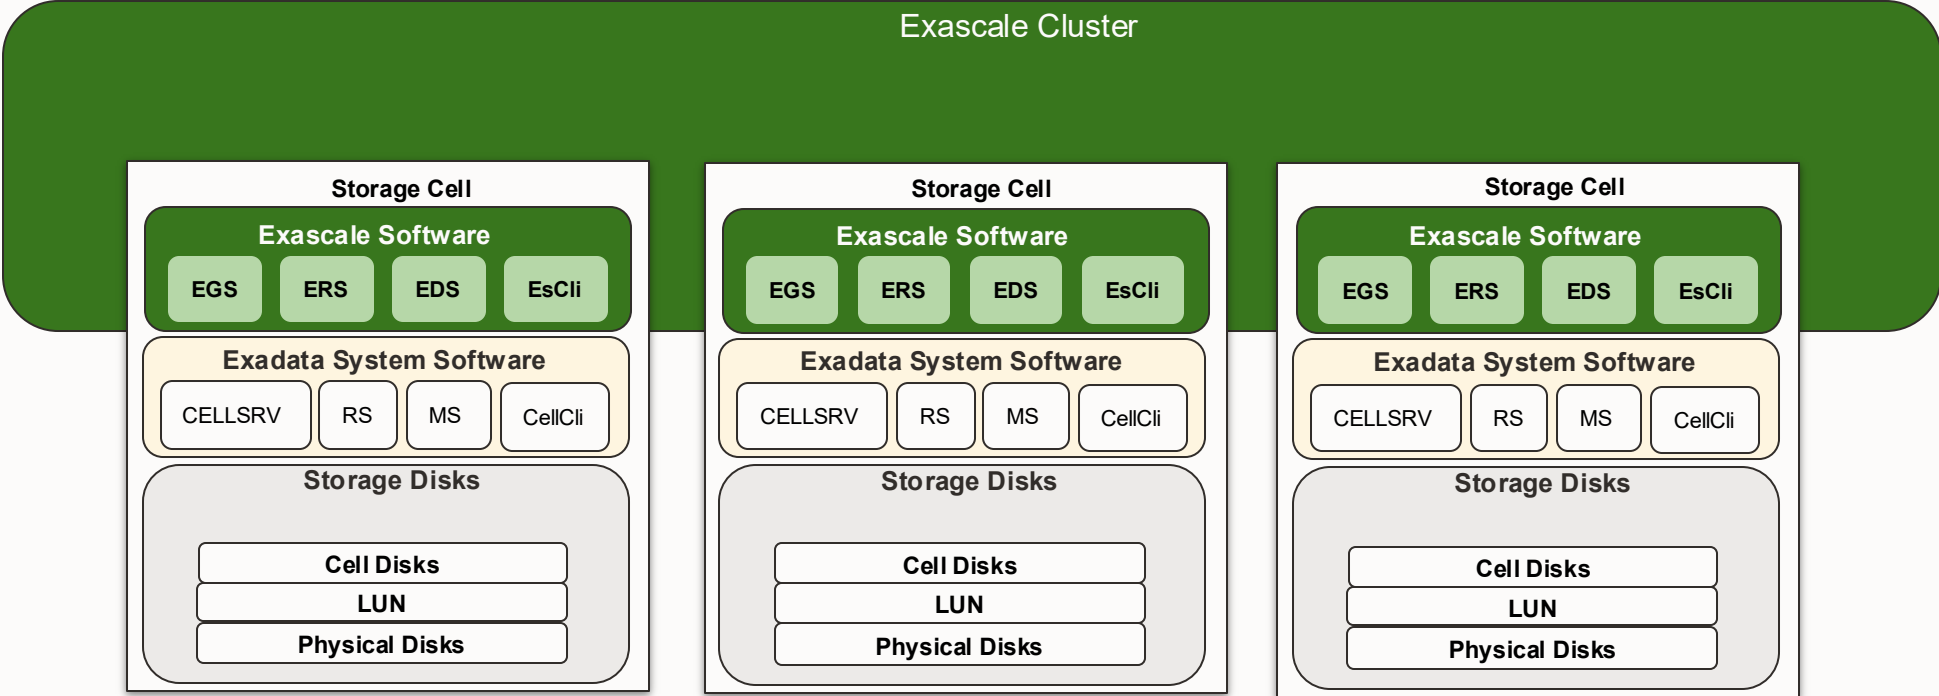
Exascale Architecture



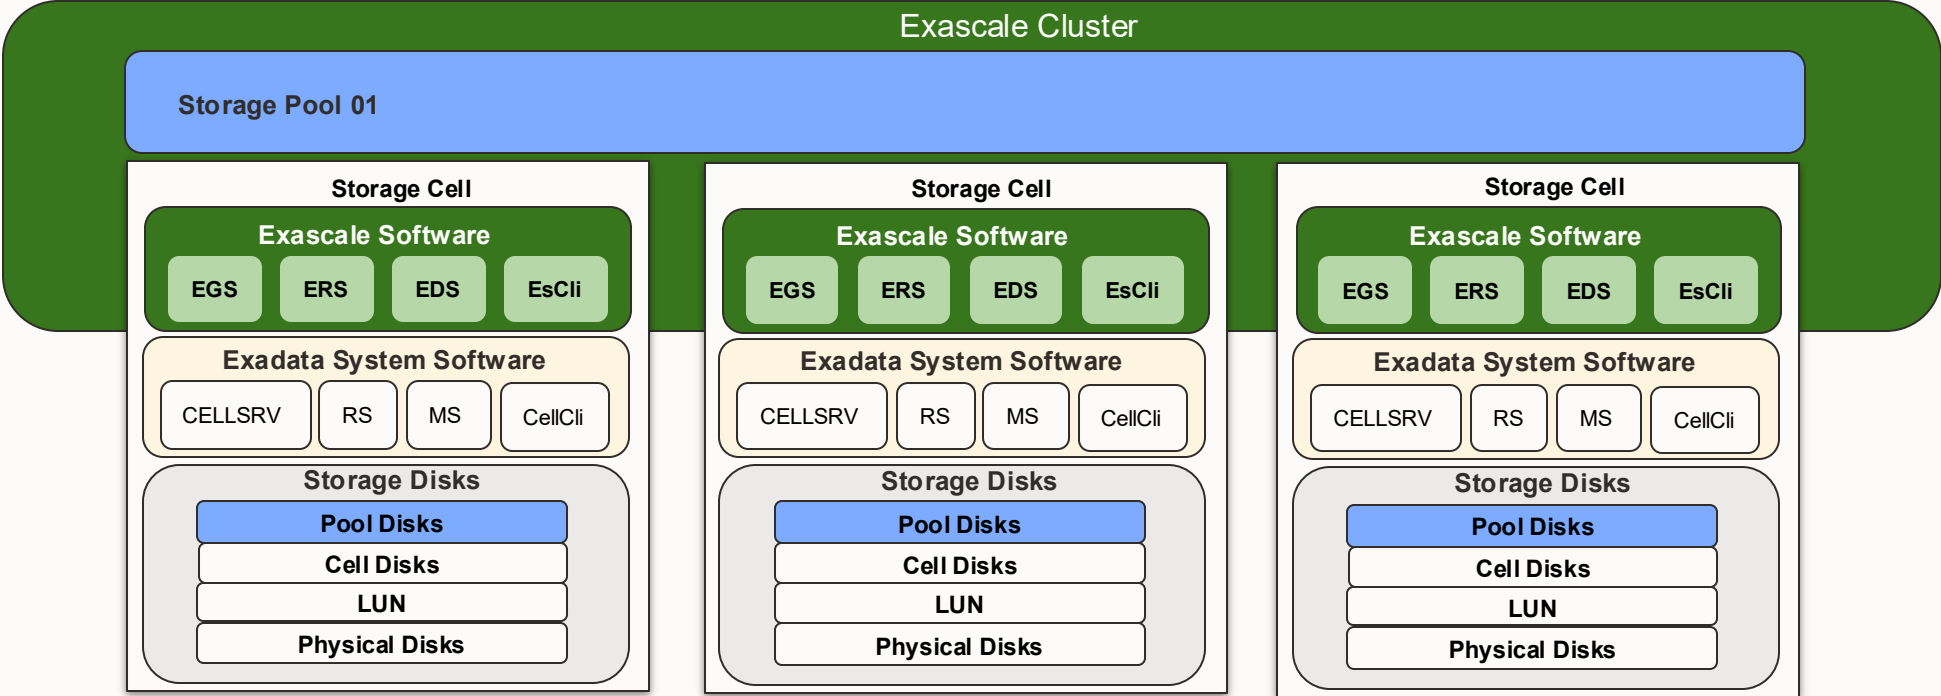
Exascale Architecture



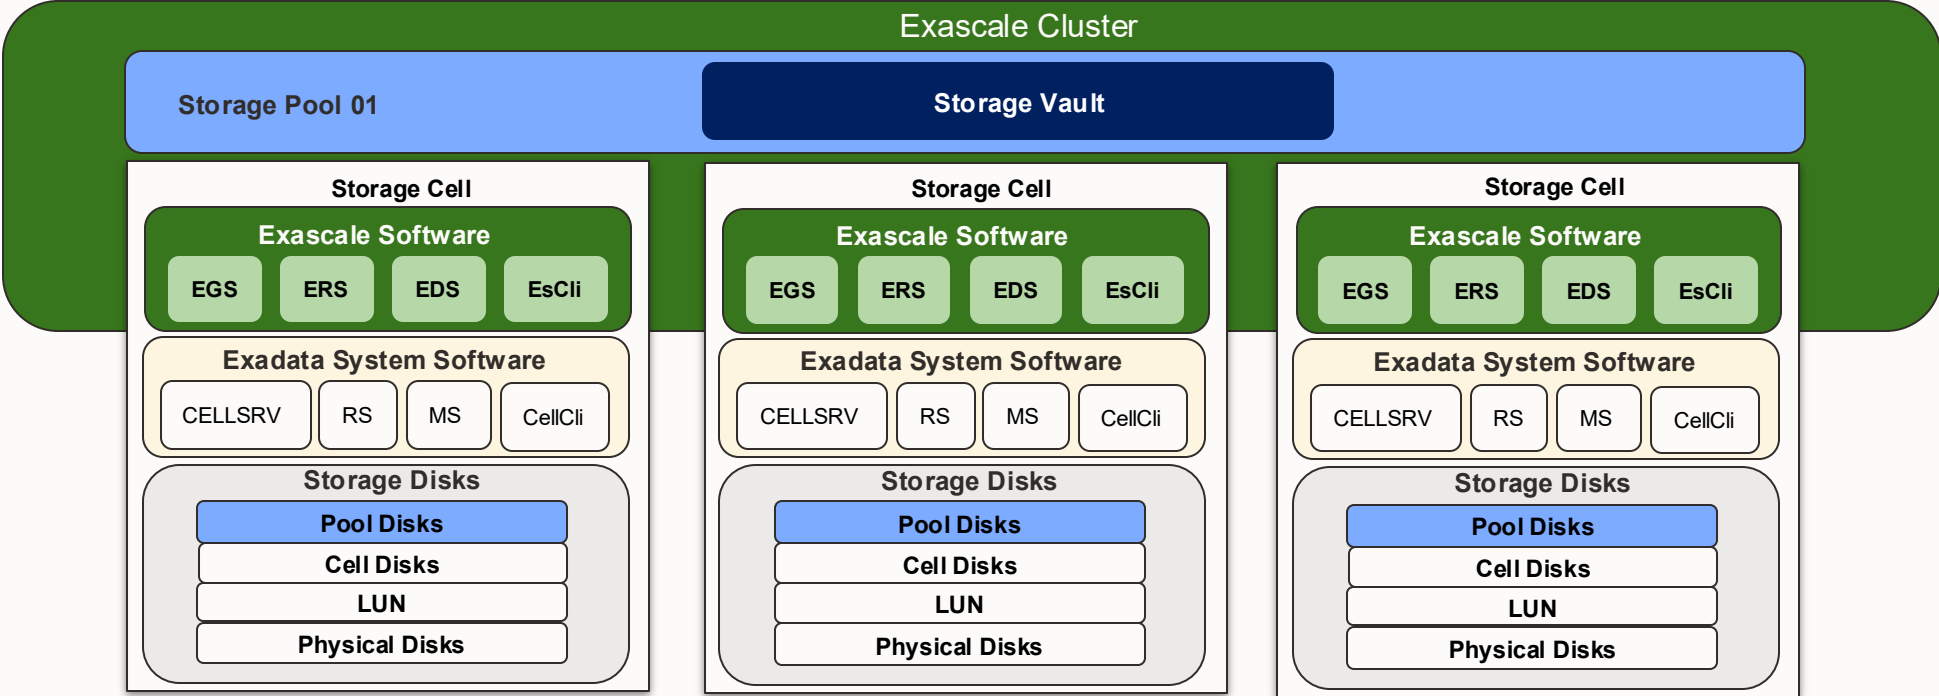
Exascale Architecture



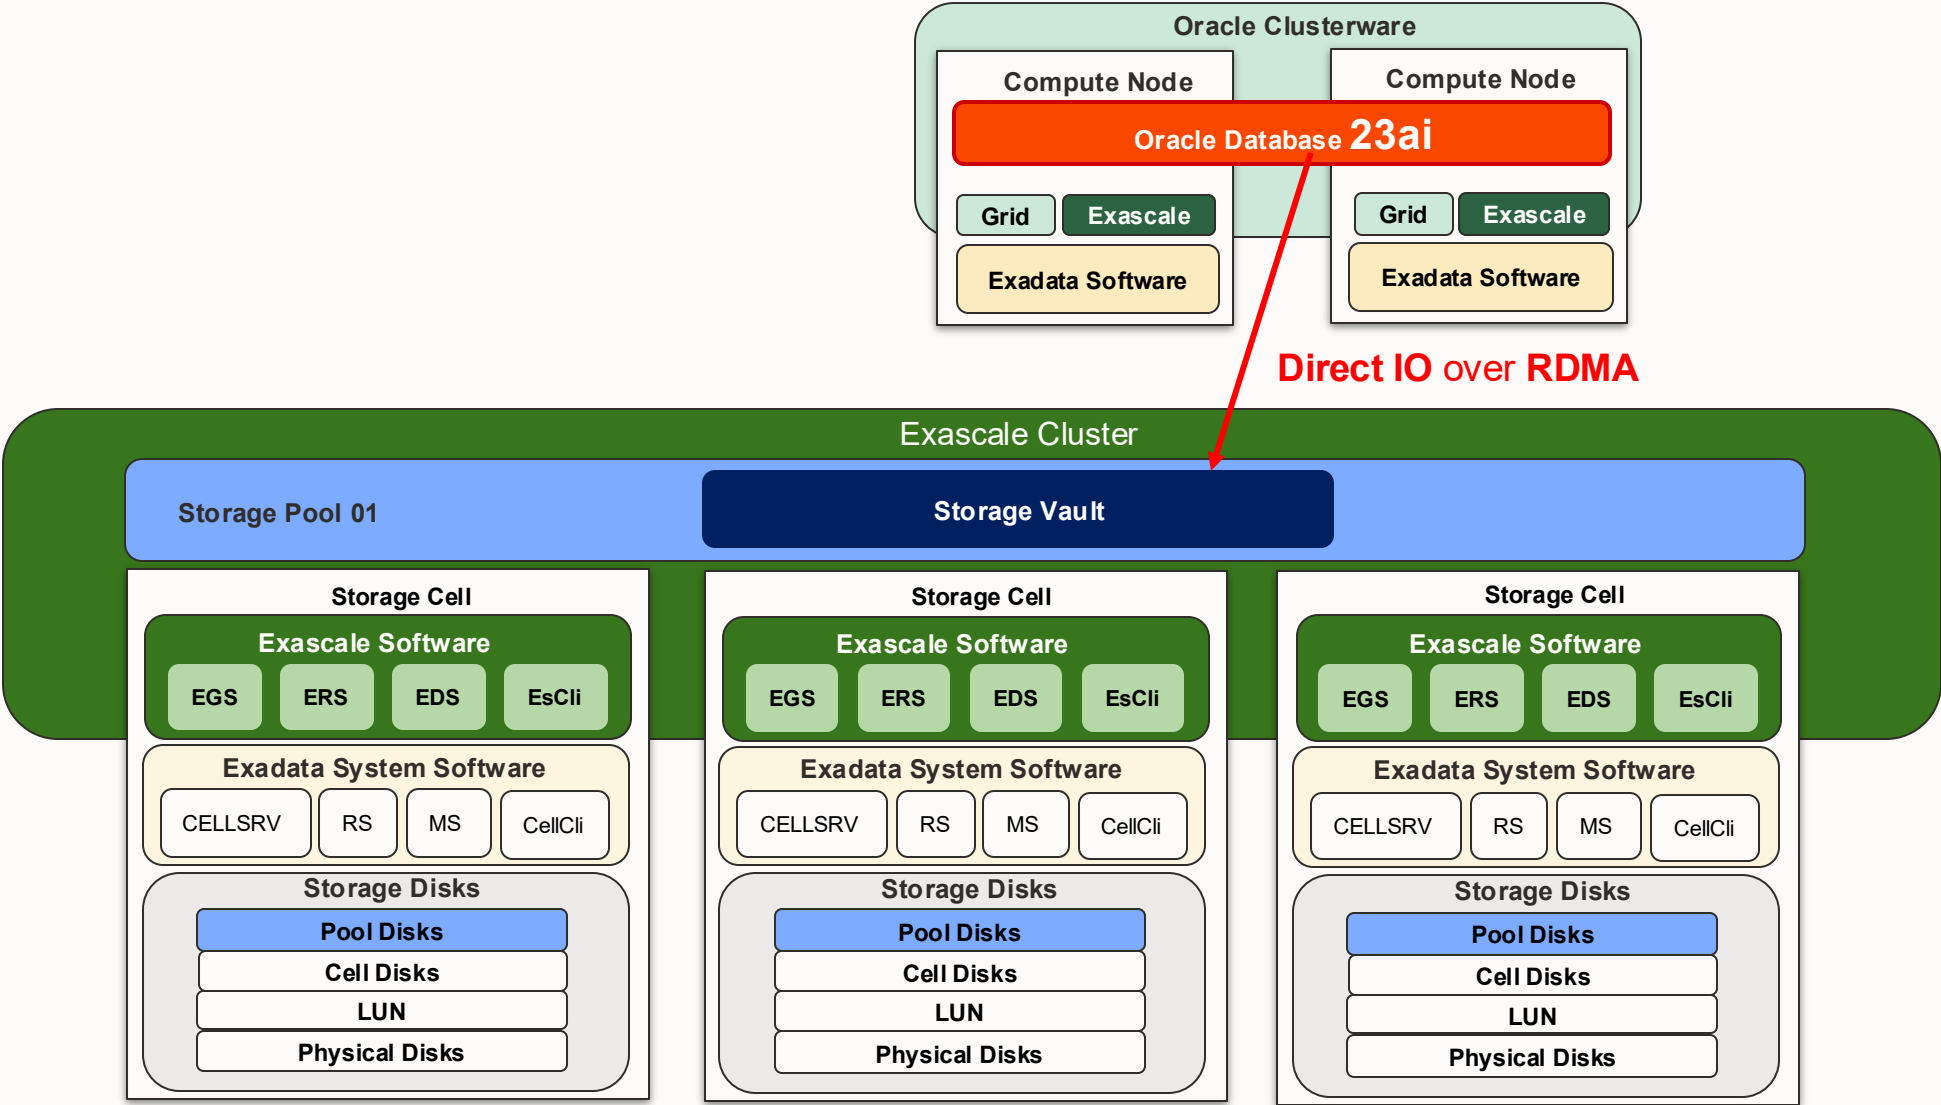
Exascale Architecture



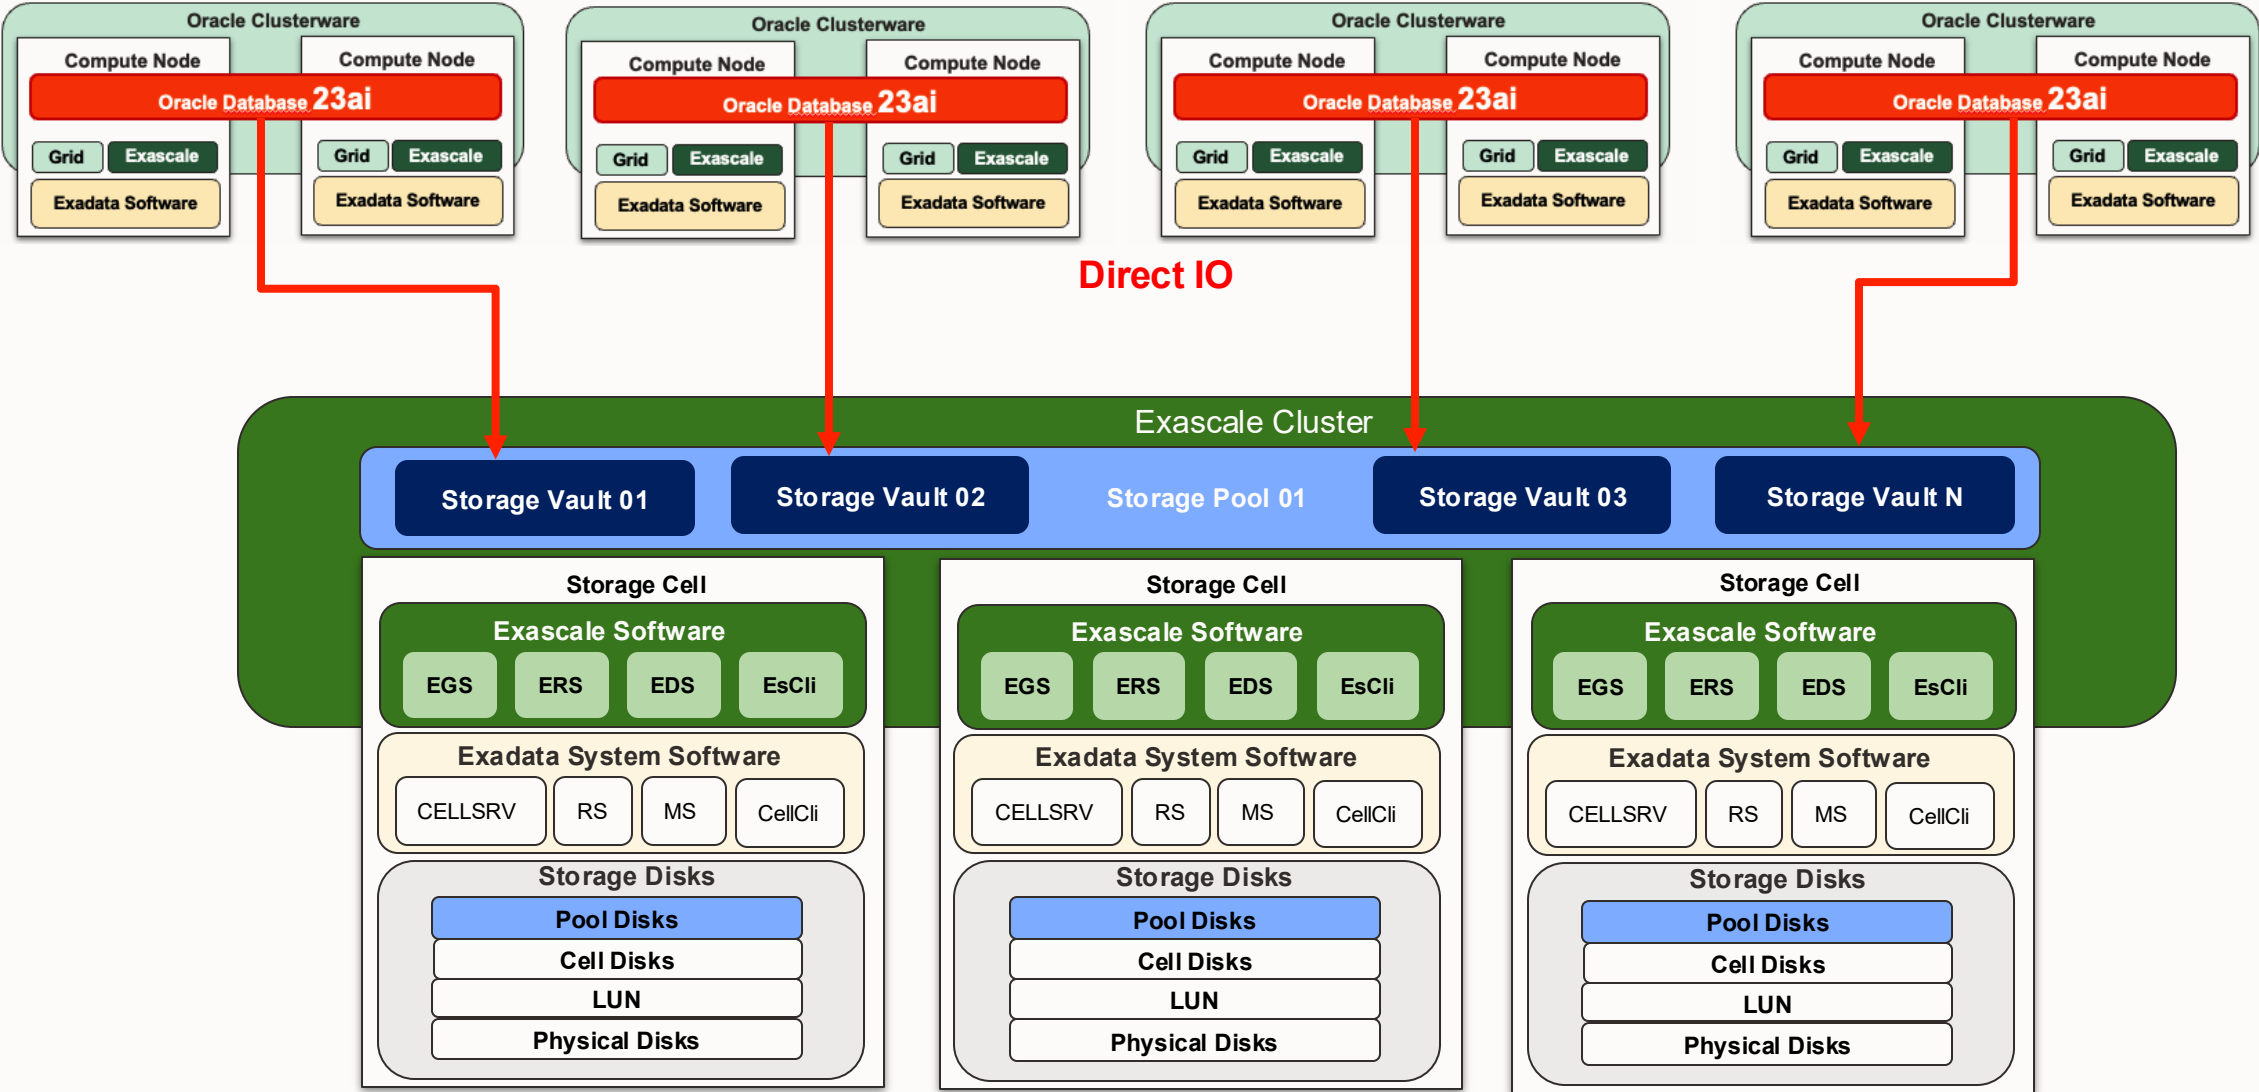
Exascale Architecture



Exascale Architecture



Exascale Architecture – Multiple Clusters

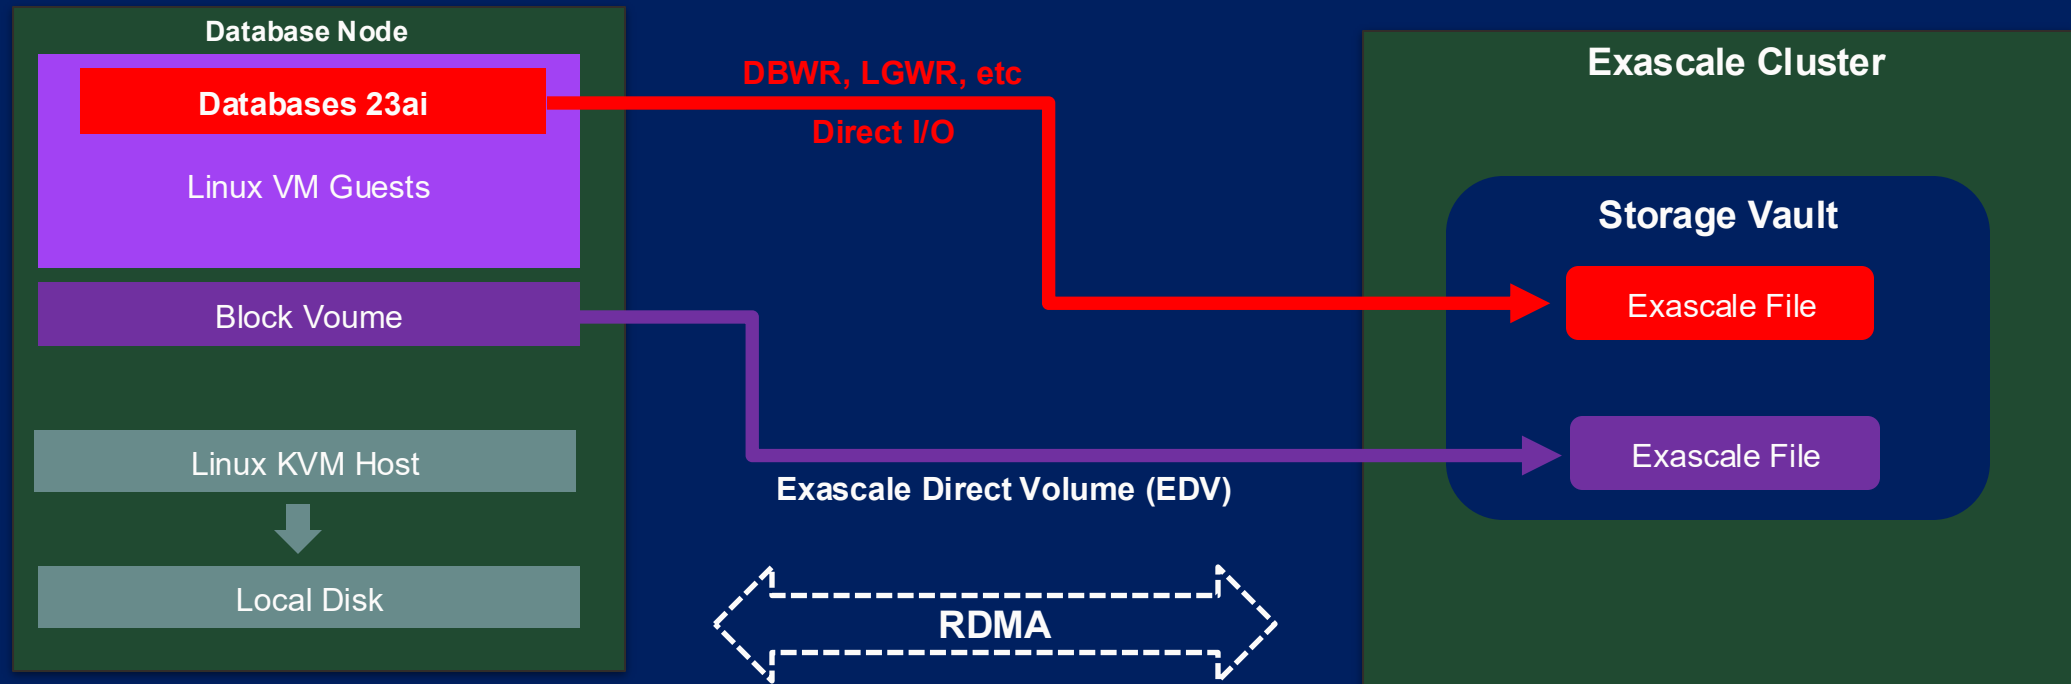


Exascale – VM Cluster

New Limits for VMs per Database Servers

VM Images Architecture – Exascale

VMs are Stored Directly on Exadata Storage using **Exascale Block Volume**



Database Files on Exascale

How database files look like in Exascale storage

Exascale Database Files - Parameters

Using ASM Diskgroup (+)

+ASM_DISKGROUP_NAME

NAME	VALUE
db_create_file_dest	+DATA
db_create_online_log_dest_1	+DATA
db_recovery_file_dest	+RECO

Exascale Database Files - Parameters

Using ASM Diskgroup (+)

+ASM_DISKGROUP_NAME

NAME	VALUE
db_create_file_dest	+DATA
db_create_online_log_dest_1	+DATA
db_recovery_file_dest	+RECO

Using Exascale Storage Vault (@)

@EXASCALE_VAULT_NAME

NAME	VALUE
db_create_file_dest	@eGED0pGD
db_create_online_log_dest_1	@eGED0pGD
db_recovery_file_dest	@eGED0pGD

Database files on Exascale

```
SQL> select name from v$controlfile where rownum = 1;
```

```
NAME
```

```
-----  
+DATA/CDBDG/CONTROLFILE/current.346.1180641287
```

```
SQL>
```

```
SQL>
```

```
SQL> select name from v$datafile where rownum = 1;
```

```
NAME
```

```
-----  
+DATA/CDBDG/DATAFILE/system.298.1180641183
```

```
SQL>
```

```
SQL>
```

```
SQL> select member from v$logfile where rownum = 1;
```

```
MEMBER
```

```
-----  
+DATA/CDBDG_A/ONLINELOG/group_201.338.1193068263
```

Using ASM Diskgroup



Database files on Exascale

```
SQL> select name from v$controlfile where rownum = 1;
NAME
-----
+DATA/CDBDG/CONTROLFILE/current.346.1180641287

SQL>
SQL>
SQL> select name from v$datafile where rownum = 1;
```

Using ASM Diskgroup



```
NAME
-----
+DATA/CDBDG/DATAFILE/system.298

SQL>
SQL>
SQL> select member from v$logfile where rownum = 1;

MEMBER
-----
+DATA/CDBDG_A/ONLINELOG/group_2
```

```
SQL> select name from v$controlfile where rownum = 1;
NAME
-----
+@eGED0pGD/GUOB-C1-E680D9DCFBC24FBDBF0F628A4359374D/CDB1_WG8_IAD/CONTROLFILE/Current.OMF.34A2D8C6

SQL>
SQL>
SQL> select name from v$datafile where rownum = 1;
NAME
-----
+@eGED0pGD/GUOB-C1-E680D9DCFBC24FBDBF0F628A4359374D/CDB1_WG8_IAD/DATAFILE/SYSTEM.OMF.75B52763

SQL>
SQL>
SQL> select member from v$logfile where rownum = 1;

MEMBER
-----
+@eGED0pGD/GUOB-C1-E680D9DCFBC24FBDBF0F628A4359374D/CDB1_WG8_IAD/ONLINELOG/group_101.OMF.4025DA2F
```

Using Exascale Storage Vault



Database files on Exascale

```
[oracle@guob011-rpdrb ~]$ srvctl config database -d CDB1_wg8_iad
Database unique name: CDB1_wg8_iad
Database name: CDB1
Oracle home: /u02/app/oracle/product/23.0.0.0/dbhome_1
Oracle user: oracle
Spfile: @eGED0pGD/GUOB-C1-E680D9DCFBC24FBDBF0F628A4359374D/CDB1_WG8_IAD/PARAMETERFILE/spfile.OMF.216B3DD2
Password file: @eGED0pGD/GUOB-C1-E680D9DCFBC24FBDBF0F628A4359374D/CDB1_WG8_IAD/PASSWORD/pwdCDB1_WG8_IAD.344ED854
Domain: b.vcnus01.oraclevcn.com
Start options: open
Stop options: immediate
Database role: PRIMARY
```

```

RMAN> list archivelog all;
list archivelog all;
using target database control file instead of recovery catalog
List of Archived Log Copies for database with db_unique_name CDB1_WG8_IAD
=====
Key          Thrd Seq      S Low Time
-----
1055        1    1053      A 13/07/2025 11:26:06
Name: @eGED0pGD/GUOB-C1-E680D9DCFBC24FBDBF0F628A4359374D/CDB1_WG8_IAD/ARCHIVELOG/2025_07_13/thread_1_seq_1053.OMF.1E016F70
1056        1    1054      A 13/07/2025 11:36:06
Name: @eGED0pGD/GUOB-C1-E680D9DCFBC24FBDBF0F628A4359374D/CDB1_WG8_IAD/ARCHIVELOG/2025_07_13/thread_1_seq_1054.OMF.7E11A777
1057        1    1055      A 13/07/2025 11:37:05
Name: @eGED0pGD/GUOB-C1-E680D9DCFBC24FBDBF0F628A4359374D/CDB1_WG8_IAD/ARCHIVELOG/2025_07_13/thread_1_seq_1055.OMF.07241CEA
1058        1    1056      A 13/07/2025 11:52:52
Name: @eGED0pGD/GUOB-C1-E680D9DCFBC24FBDBF0F628A4359374D/CDB1_WG8_IAD/ARCHIVELOG/2025_07_13/thread_1_seq_1056.OMF.7009A1E7
1059        1    1057      A 13/07/2025 11:53:07
Name: @eGED0pGD/GUOB-C1-E680D9DCFBC24FBDBF0F628A4359374D/CDB1_WG8_IAD/ARCHIVELOG/2025_07_13/thread_1_seq_1057.OMF.2E77E79E
1060        1    1058      A 13/07/2025 12:35:09
Name: @eGED0pGD/GUOB-C1-E680D9DCFBC24FBDBF0F628A4359374D/CDB1_WG8_IAD/ARCHIVELOG/2025_07_13/thread_1_seq_1058.OMF.50237C5C
1061        1    1059      A 13/07/2025 13:01:04
Name: @eGED0pGD/GUOB-C1-E680D9DCFBC24FBDBF0F628A4359374D/CDB1_WG8_IAD/ARCHIVELOG/2025_07_13/thread_1_seq_1059.OMF.3EF78B0E
```

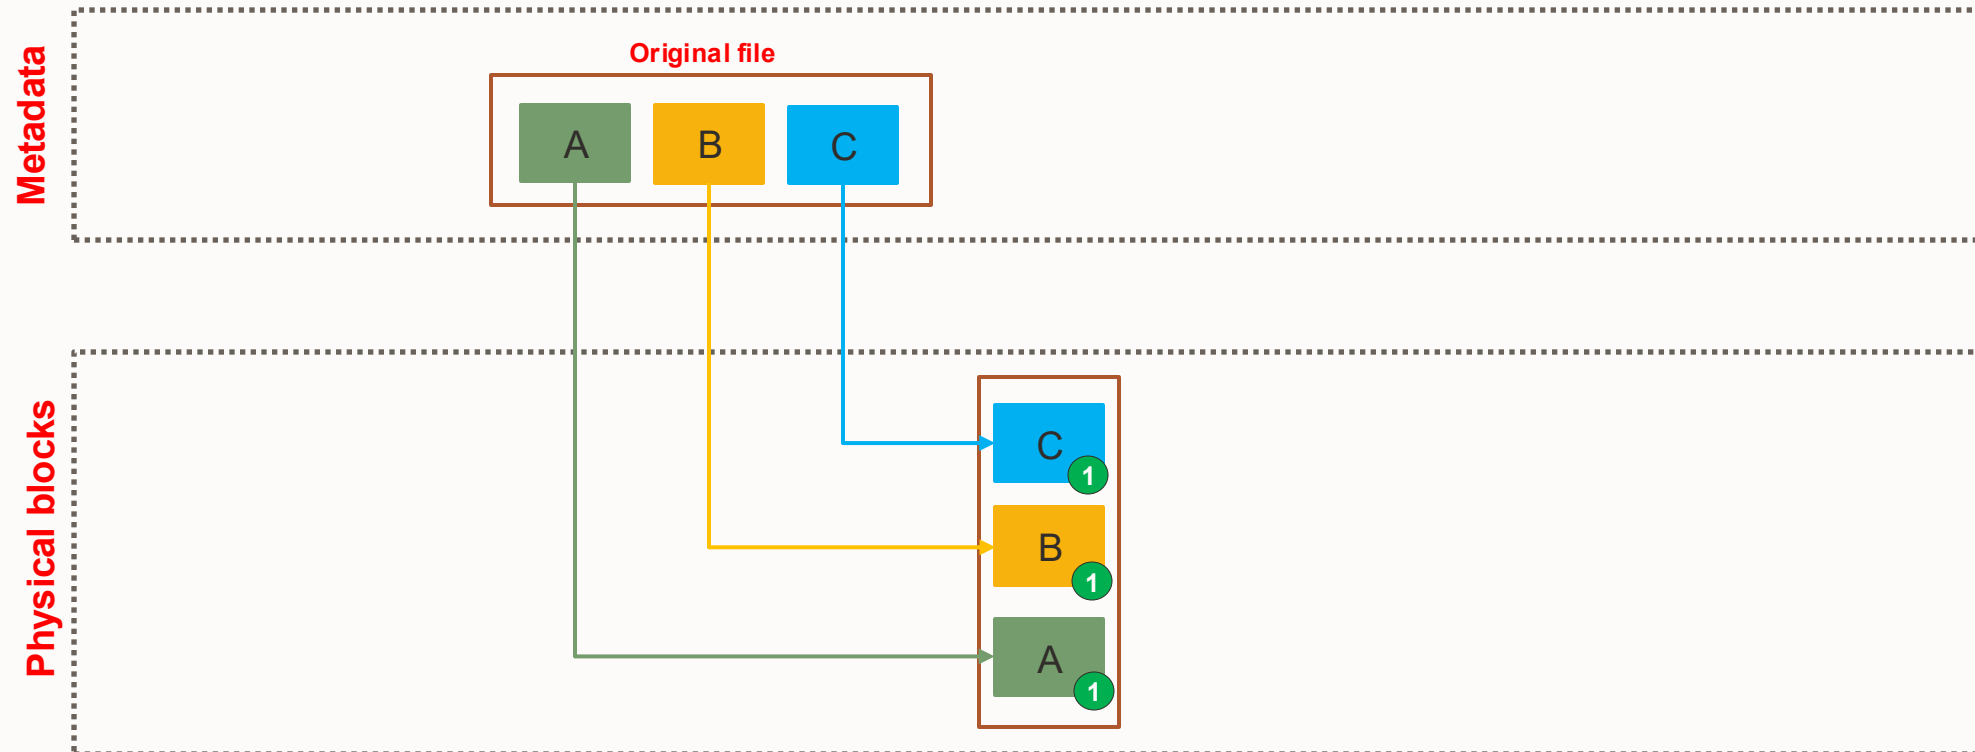
Exascale Intelligent Snapshot Clone

How Exascale uses Redirect-On-Write (ROW) – by Maicon Carneiro

Snapshot Clone Redirect-On-Write Overview

0

Before create a clone

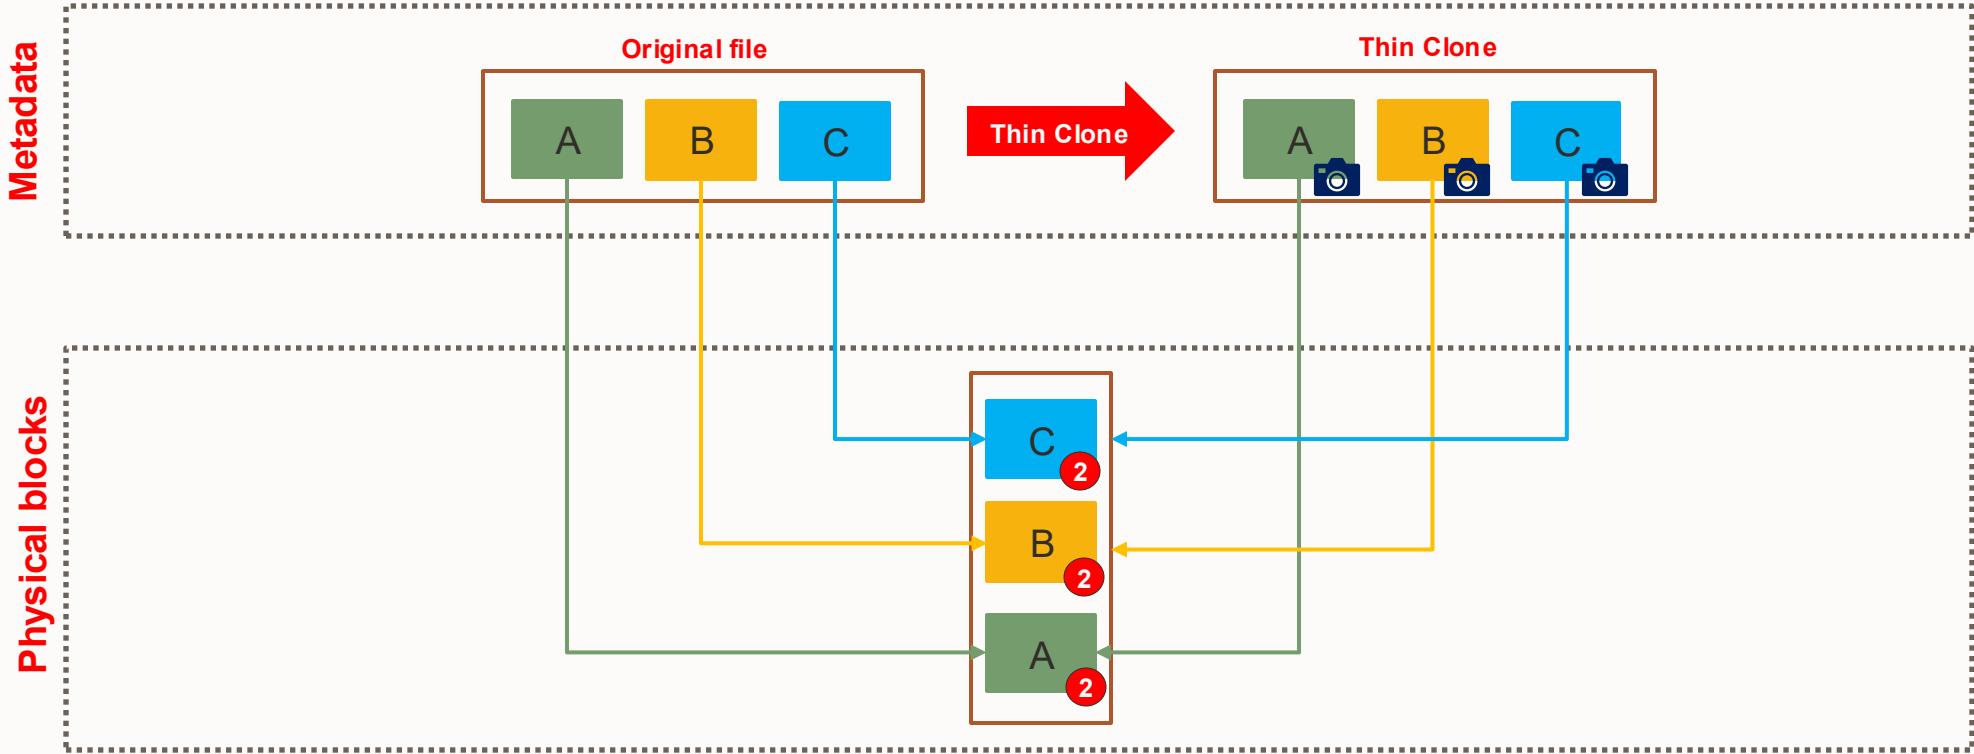


Snapshot Clone Redirect-On-Write Overview

1

Thin clone created

All blocks are shared by 2 references

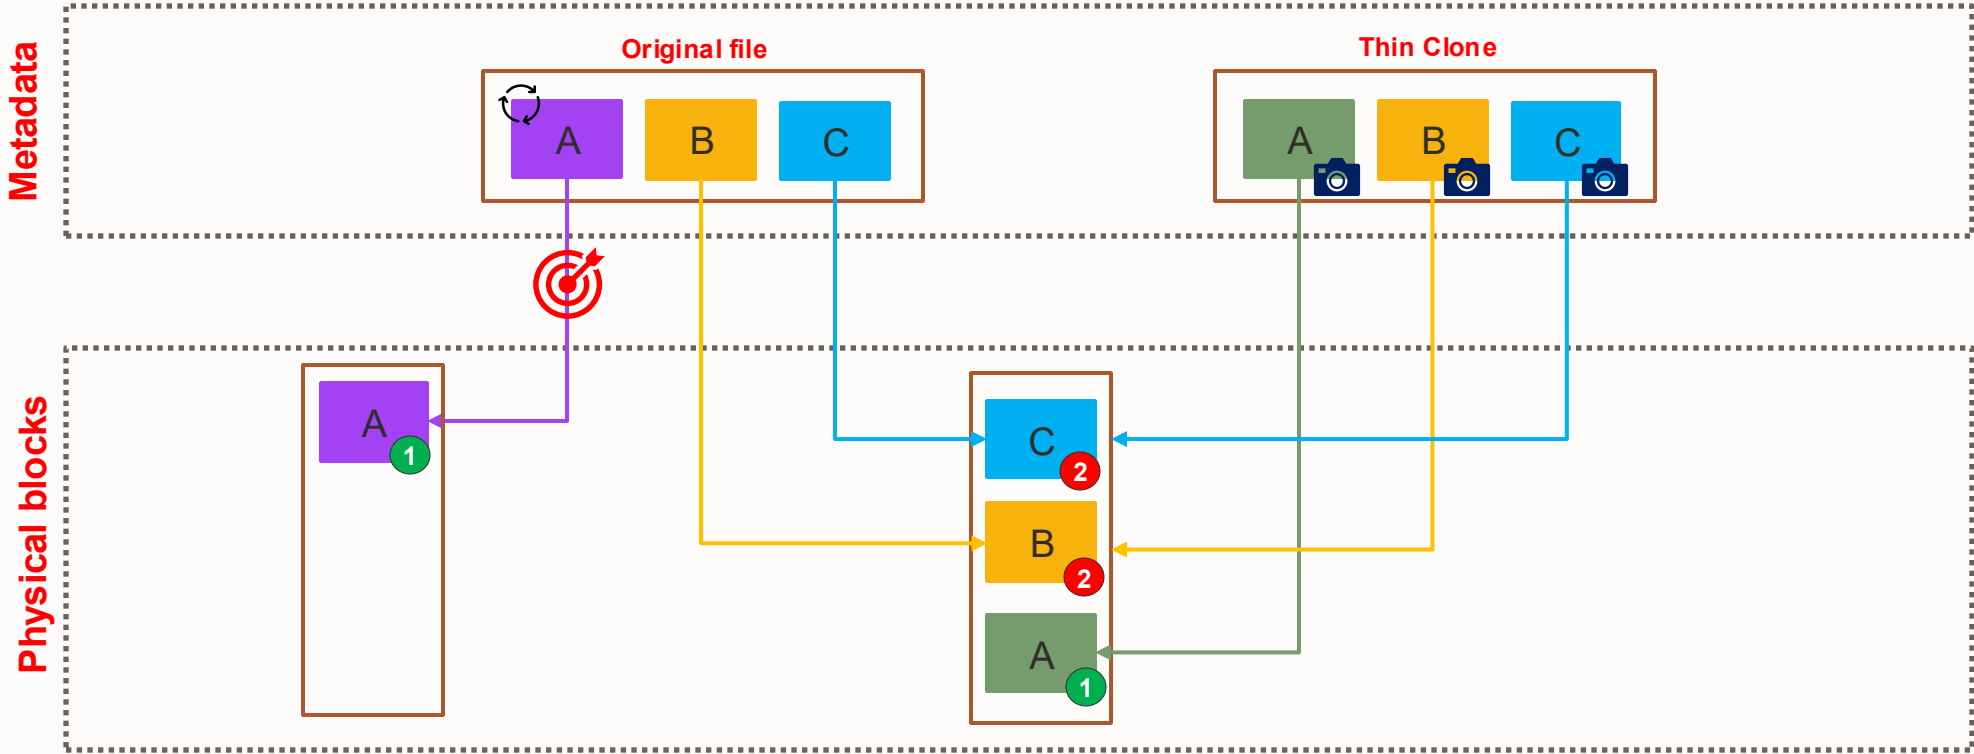


Snapshot Clone Redirect-On-Write Overview

2

Block **A** is modified on **original file**

Write is **redirected** and block **A** is no longer shared

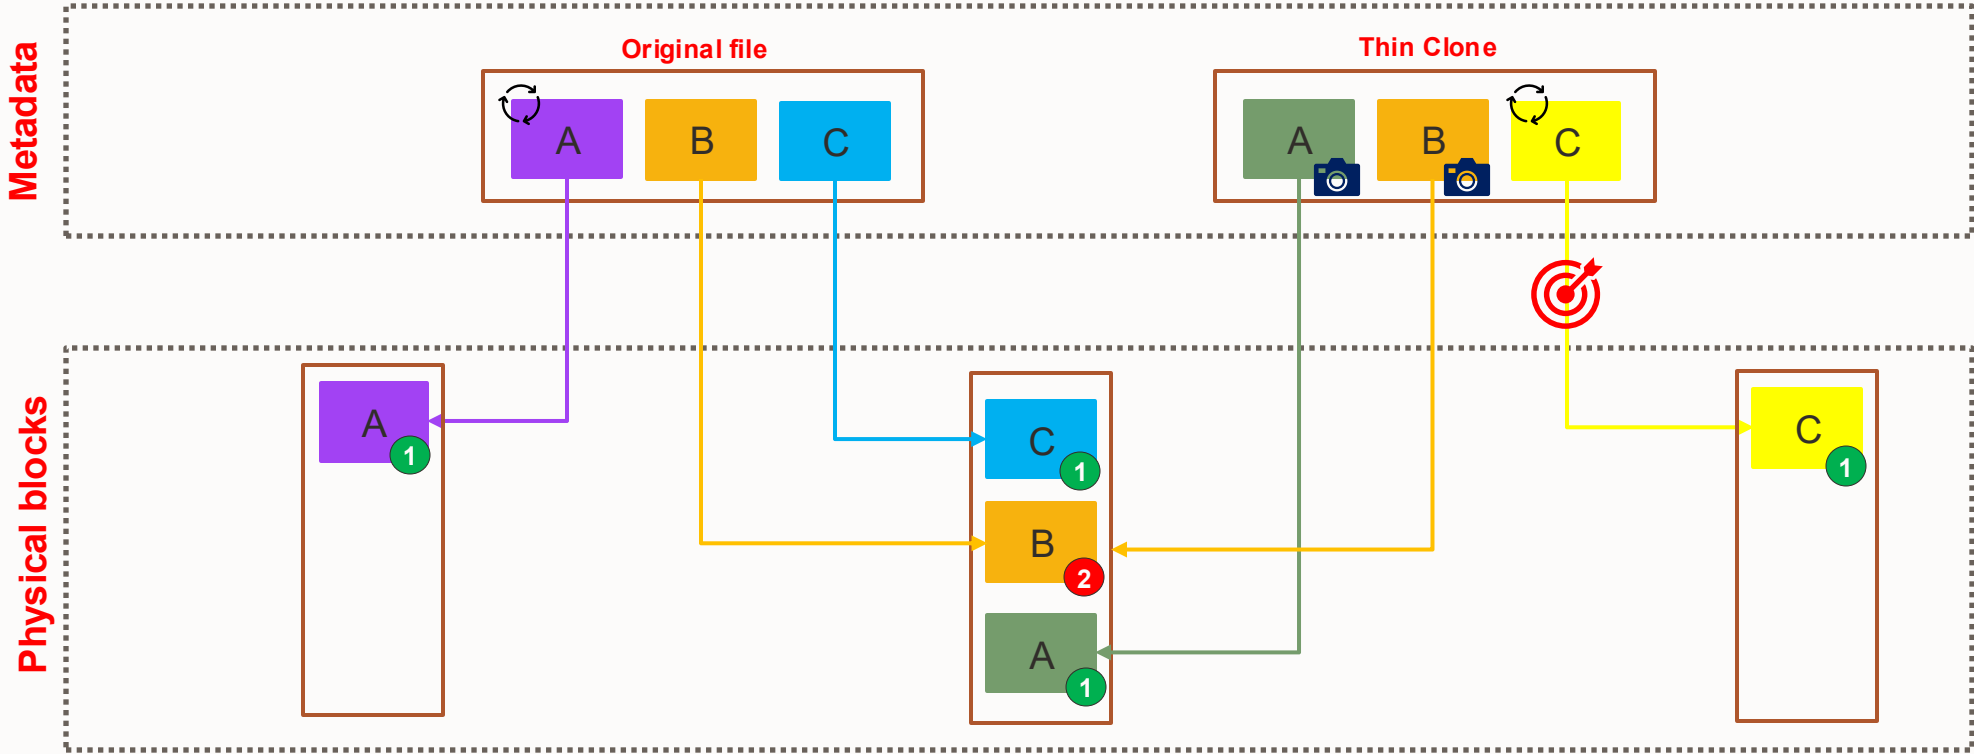


Snapshot Clone Redirect-On-Write Overview

3

Block **C** is modified on **Thin clone**

Write is **redirected** and block **C** is no longer shared

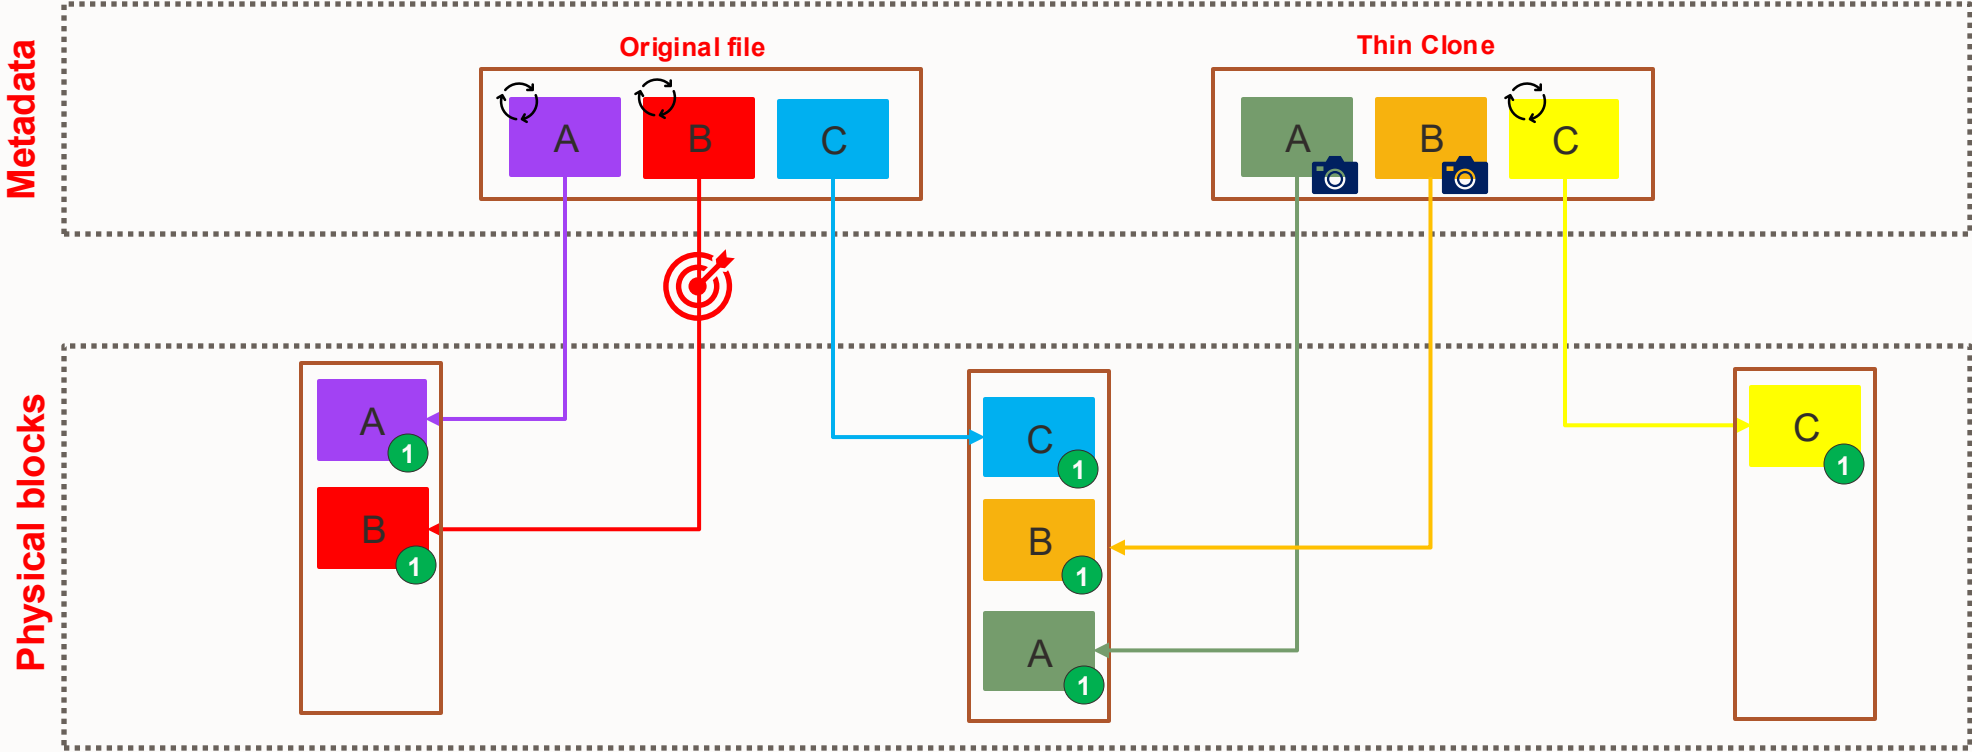


Snapshot Clone Redirect-On-Write Overview

4

Block **B** is modified on **Original file**

Write is **redirected** and block **B** is no longer shared

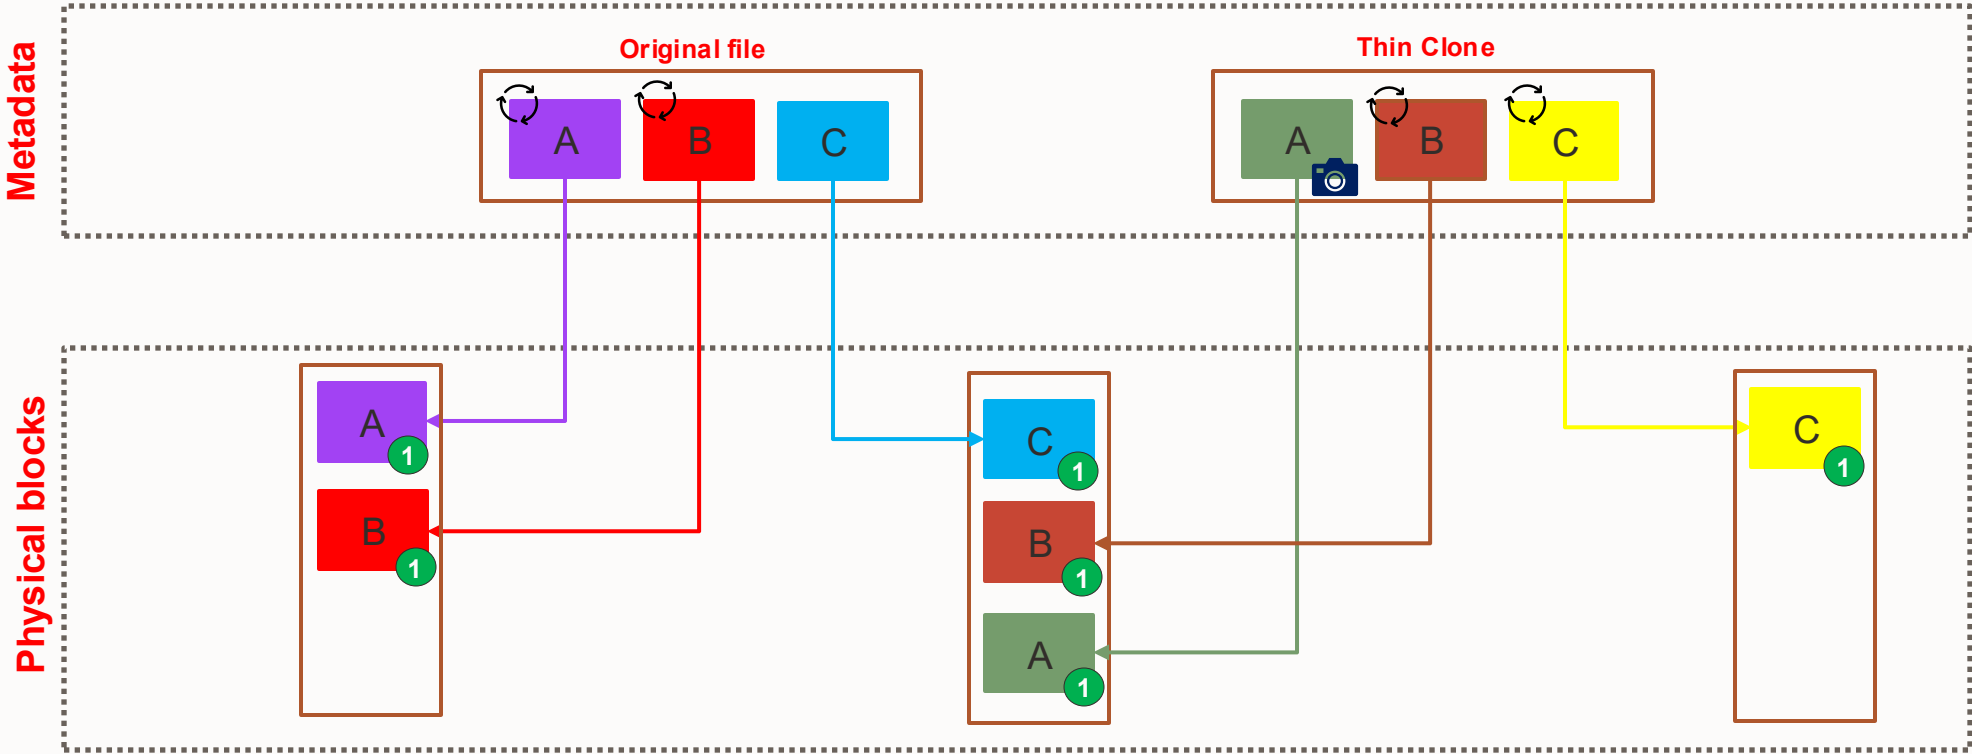


Snapshot Clone Redirect-On-Write Overview

5

Block **B** is modified on **Thin clone**

Block **C** is not shared and can be **overridden**, write is **in-place**



Snapshot Clone Databases on Exascale

Create Thin Clones with one command !

Snapshot clone of Databases on Exascale

PDB Snapshot Clone

- Database Aware Thin Clone
- Just add “**SNAPSHOT COPY**” clause in CREATE PLUGGABLE DATABASE command.
- Oracle Database 23ai uses Exascale snapshot capability transparently.

Snapshot clone of Databases on Exascale

PDB Snapshot Clone

- Database Aware Thin Clone
- Just add “**SNAPSHOT COPY**” clause in CREATE PLUGGABLE DATABASE command.
- Oracle Database 23ai uses Exascale snapshot capability transparently.

CDB Snapshot Clone

- Automated with **gDBClone** tool
- gDBClone creates a snapshot clone of source datafiles.
- gDBClone automation creates a new database instance in the Cluster and mount the new database using the cloned datafiles.

Exascale - Database Aware Thin Clone

Create Thin Clones with one command !

Exascale - Database Aware Thin Clone

```
SQL> show pdbs
```

CON_ID	CON_NAME	OPEN MODE	RESTRICTED
2	PDB\$SEED	READ ONLY	NO
3	PDB1	READ WRITE	NO

```
SQL>
```

```
SQL> @pdbsize
```

PDB_NAME	GUID	GBYTES
PDB1	329A723565C06874E0630E02020A3A53	222.20

Exascale - Database Aware Thin Clone

```
SQL> CREATE PLUGGABLE DATABASE PDB1_CLONE FROM PDB1;
```

Exascale - Database Aware Thin Clone

```
SQL> CREATE PLUGGABLE DATABASE PDB1_CLONE FROM PDB1;
```

```
SQL> create pluggable database PDB1_CLONE from PDB1  
keystore identified by EXTERNAL STORE; 2
```

```
Pluggable database created.
```

```
Elapsed: 00:12:52.46
```

Exascale - Database Aware Thin Clone

```
SQL> CREATE PLUGGABLE DATABASE PDB1_CLONE FROM PDB1;
```

```
SQL> create pluggable database PDB1_CLONE from PDB1  
keystore identified by EXTERNAL STORE; 2  
  
Pluggable database created.  
  
Elapsed: 00:12:52.46
```

SNAPSHOT COPY



```
SQL> CREATE PLUGGABLE DATABASE PDB1_SNAP FROM PDB1 SNAPSHOT COPY;
```

Exascale - Database Aware Thin Clone

```
SQL> CREATE PLUGGABLE DATABASE PDB1_CLONE FROM PDB1;
```

```
SQL> create pluggable database PDB1_CLONE from PDB1  
keystore identified by EXTERNAL STORE; 2
```

```
Pluggable database created.
```

```
Elapsed: (00:12:52.46)
```

+ 12 minutes

```
SQL> CREATE PLUGGABLE DATABASE PDB1_SNAP FROM PDB1 SNAPSHOT COPY;
```

```
SQL> create pluggable database PDB1_SNAP from PDB1 snapshot copy  
keystore identified by EXTERNAL STORE; 2
```

```
Pluggable database created.
```

```
Elapsed: (00:00:09.08)
```

- 10 seconds

SNAPSHOT COPY

CDB Thin Clone with **gDBClone** on Exascale

Create Thin Clones with one command !

gDBClone on Exascale – Example

Basic syntax

```
./gDBClone.bin snap -sdbname CDB1 -tdbname CDB1CLONE
```

Full syntax


```
./gDBClone.bin snap \  
-sdbname CDB1 \  
-tdbname CDB1CLONE \  
-tdomain b.vcnus01.oraclevcn.com \  
-tdbhome OraHome1 \  
-syspwf /opt/gDBClone/cdb1_sys \  
-walletpwf /opt/gDBClone/cdb1_sys \  
-racmod 2 \  
-sga_max_size 2048 -sga_target 2048
```

gDBClone Powerful Database Clone/Snapshot Management Tool (Doc ID 2099214.1)


<https://docs.oracle.com/en/engineered-systems/exadata-database-machine/exscl/cloning-pdb-cdb.html>

CDB Thin Clone with gDBClone on Exascale

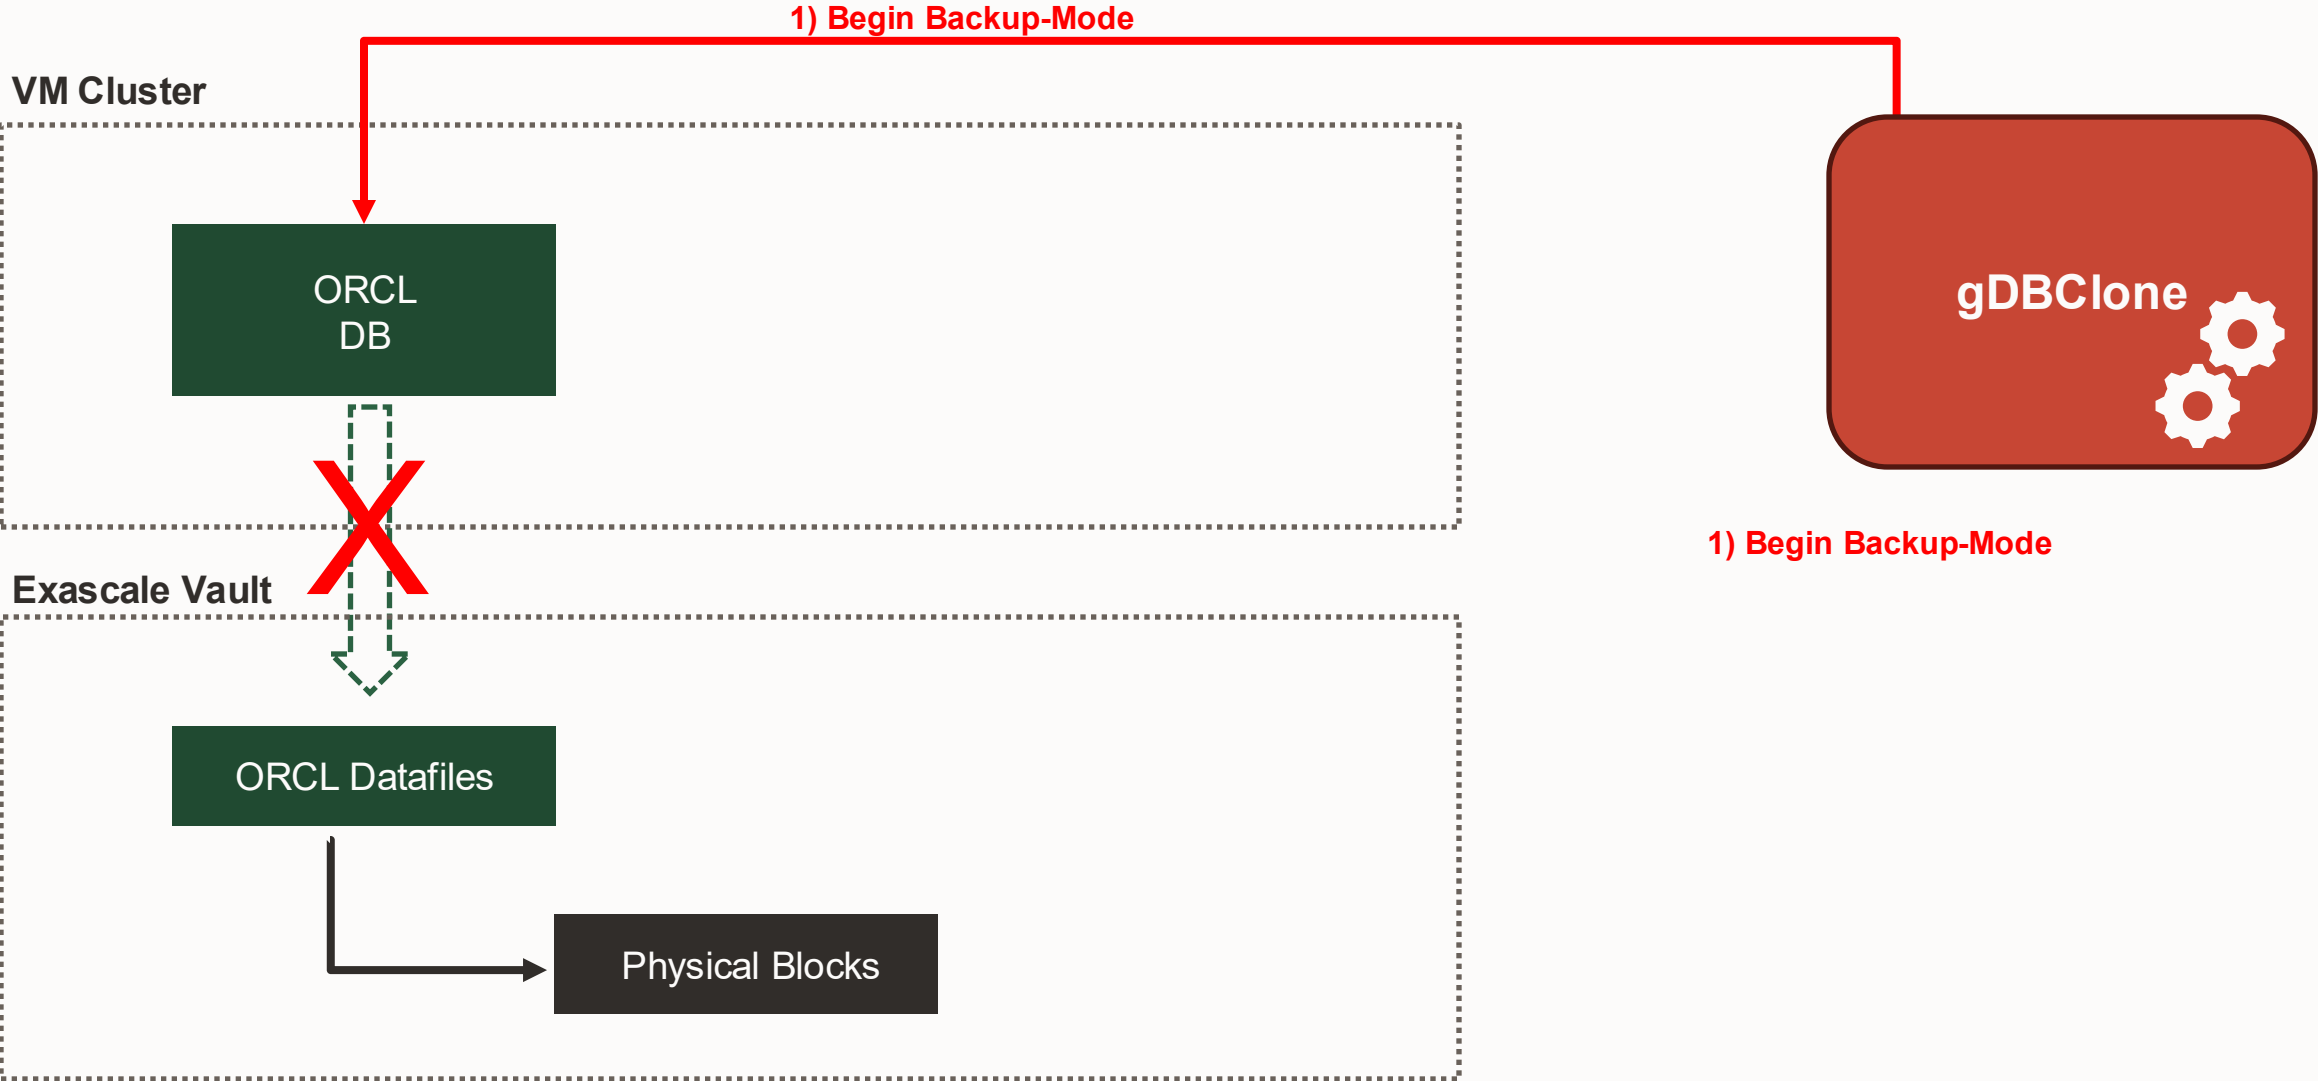
VM Cluster



Exascale Vault

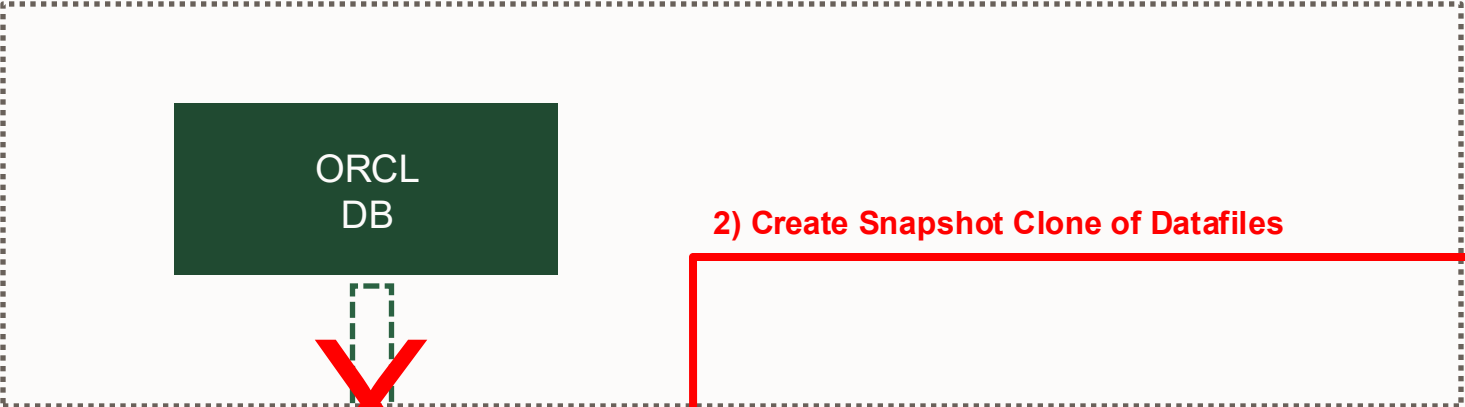


CDB Thin Clone with gDBClone on Exascale

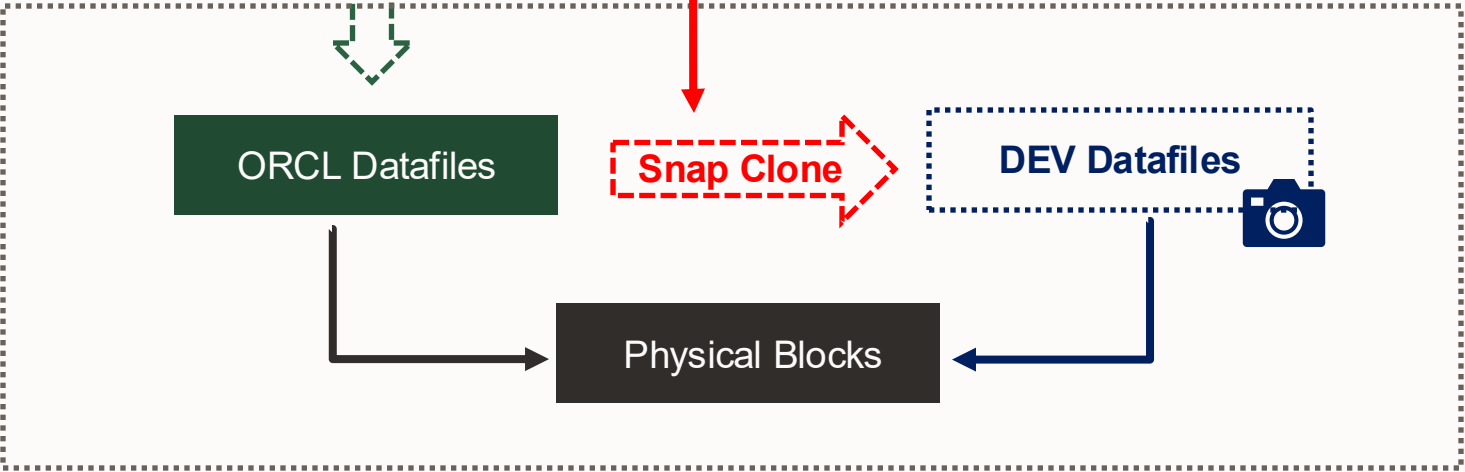


CDB Thin Clone with gDBClone on Exascale

VM Cluster



Exascale Vault



- 1) Begin Backup-Mode
- 2) Create Snapshot Clone of Datafiles

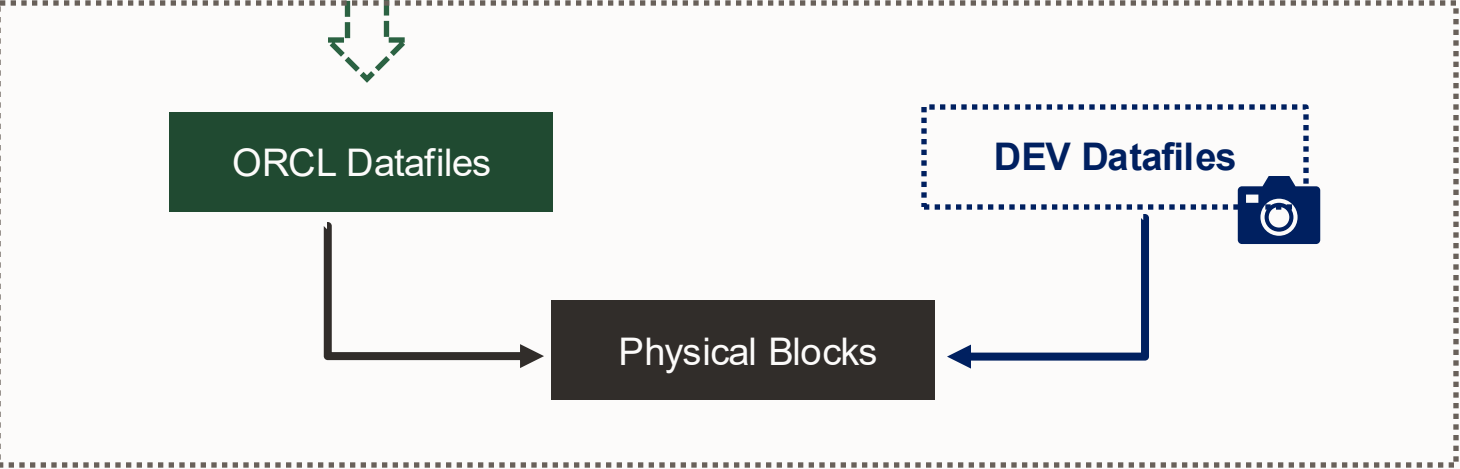
CDB Thin Clone with gDBClone on Exascale

3) End Backup-Mode

VM Cluster



Exascale Vault



- 1) Begin Backup-Mode
- 2) Create Snapshot Clone of Datafiles
- 3) End Backup-Mode

CDB Thin Clone with gDBClone on Exascale

VM Cluster

ORCL
DB

DEV
DB



4) Creates
Auxiliary Instance

Exascale Vault

ORCL Datafiles

DEV Datafiles

Physical Blocks

- 1) Begin Backup-Mode
- 2) Create Snapshot Clone of Datafiles
- 3) End Backup-Mode
- 4) Creates auxiliary instance

CDB Thin Clone with gDBClone on Exascale

VM Cluster

ORCL
DB

DEV
DB

5) MOUNT and OPEN
Cloned Database



Exascale Vault

ORCL Datafiles

DEV Datafiles

Physical Blocks

- 1) Begin Backup-Mode
- 2) Create Snapshot Clone of Datafiles
- 3) End Backup-Mode
- 4) Creates auxiliary instance
- 5) MOUNT and OPEN Cloned Database

gDBClone on Exascale – Demo

```
[root@guob011-rpdrb ~]#  
[root@guob011-rpdrb ~]# /opt/gDBClone/gDBClone.bin snap \  
> -sdbname CDB1_wg8_iad \  
> -tdbname CDB1THIN \  
> -tdomain b.vcnus01.oraclevcn.com \  
> -tdbhome OraHome1 \  
> -syspwf /opt/gDBClone/cdb1_sys \  
> -walletpwf /opt/gDBClone/cdb1_sys \  
> -racmod 2 \  
> -sga_max_size 2048 -sga_target 2048  
INFO: 2025-07-13 12:33:26: Please check the logfile '/opt/gDBClone/out/log/gDBClone_184810.log' for more details  
  
|-----|  
gDBClone - Version: 3.0.6-21  
Copyright (c) 2012-2025 Oracle and/or its affiliates.  
-----  
Author: Ruggero Citton <ruggero.citton@oracle.com>  
RAC Pack, Cloud Innovation and Solution Engineering Team  
|-----|
```

gDBClone on Exascale – Demo: MacroStep1

Validation

```
MacroStep1 - Getting information and validating setup...
INFO: 2025-08-10 19:13:14: Checking whether the database 'CDB1_wg8_iad' is on Exascale 19:13:14
INFO: 2025-08-10 19:13:14: Processing CDB1_wg8_iad;' on Exascale
INFO: 2025-08-10 19:13:14: Source database 'CDB1_WG8_IAD' Exascale vault is '@eGED0pGD/GU0B-C1-E680D9DCFBC24FBDBF0F628A435937
INFO: 2025-08-10 19:13:14: Validating environment...
INFO: 2025-08-10 19:13:14: ...checking superuser usage
INFO: 2025-08-10 19:13:14: ...checking clusterware
INFO: 2025-08-10 19:13:14: ...checking oracle Restart
INFO: 2025-08-10 19:13:14: ...checking minimum crs active version
INFO: 2025-08-10 19:13:14: ...checking source database 'CDB1_wg8_iad' existence
INFO: 2025-08-10 19:13:14: ...checking if source database 'CDB1_wg8_iad' is running
INFO: 2025-08-10 19:13:15: ...checking if source database 'CDB1_wg8_iad' is having all DBFs on same exascale vault '@eGED0pGD
INFO: 2025-08-10 19:13:15: ...checking if source database 'CDB1_wg8_iad' is having all DBFs on same exascale root path
INFO: 2025-08-10 19:13:15: ...checking if source database 'CDB1_wg8_iad' is having invalid PDBs
INFO: 2025-08-10 19:13:16: ...checking if target database name 'CDBSNAP1' is a valid name
INFO: 2025-08-10 19:13:16: ...checking source database 'CDB1_wg8_iad' connectivity
INFO: 2025-08-10 19:13:16: ...checking source database 'CDB1_wg8_iad' and target dbhome version
INFO: 2025-08-10 19:13:21: ...checking minimum source database 'CDB1_wg8_iad' dbhome version
INFO: 2025-08-10 19:13:21: ...checking if target database 'CDBSNAP1' exists
INFO: 2025-08-10 19:13:21: ...checking target database 'CDBSNAP1' registered instance
INFO: 2025-08-10 19:13:24: ...checking if target database 'CDBSNAP1' is having files on exascale
INFO: 2025-08-10 19:13:24: ...checking if source database 'CDB1_wg8_iad' is snapable
NOTE: 2025-08-10 19:13:26: Transparent Data Encryption is enabled on 'CDB1_wg8_iad'
INFO: 2025-08-10 19:13:27: ...checking whether the database 'CDB1_wg8_iad' is a primary/physical standby database.
INFO: 2025-08-10 19:13:28: ...checking whether the database 'CDB1_wg8_iad' is in rw/ro mode
INFO: 2025-08-10 19:13:29: ...checking whether the database 'CDB1_wg8_iad' is running as backup mode
INFO: 2025-08-10 19:13:29: ...checking whether the database 'CDB1_wg8_iad' is running in archivelog mode
INFO: 2025-08-10 19:13:29: ...checking if there are offline datafiles
SUCCESS: 2025-08-10 19:13:29: Environment validation complete
```

gDBClone on Exascale – Demo: MacroStep2

Clone

```
MacroStep2 – Getting database Exascale clone...
INFO: 2025-08-10 19:13:29: Collecting source database 'CDB1_wg8_iad' required information
INFO: 2025-08-10 19:13:29: ...checking 'CDB1_wg8_iad' supplementary logging status
INFO: 2025-08-10 19:13:30: ...checking 'CDB1_wg8_iad' block change tracking
INFO: 2025-08-10 19:13:30: ...checking 'CDB1_wg8_iad' flashback mode
INFO: 2025-08-10 19:13:30: ...getting 'CDB1_wg8_iad' storage information
INFO: 2025-08-10 19:13:31: ...checking 'CDB1_wg8_iad' open mode
INFO: 2025-08-10 19:13:31: ...checking 'CDB1_wg8_iad' role
INFO: 2025-08-10 19:13:31: ...checking 'CDB1_wg8_iad' readonly tablespaces
INFO: 2025-08-10 19:13:32: Entering into Exascale database clone creation phase 1
INFO: 2025-08-10 19:13:32: ...making required action to get consistent database exascale clone
NOTE: 2025-08-10 19:13:34: Setting database 'CDB1_wg8_iad' in backup mode
INFO: 2025-08-10 19:13:35: ...getting control file backup
INFO: 2025-08-10 19:13:38: ...getting source database 'CDB1_wg8_iad' exascale clone
NOTE: 2025-08-10 19:13:39: Ending database 'CDB1_wg8_iad' backup mode
INFO: 2025-08-10 19:13:41: ...getting the exascale archive files clone for database 'CDB1_wg8_iad'
```

Begin Backup

5 seconds

End Backup

gDBClone on Exascale – Demo: MacroStep2

Clone

```
INFO: 2025-08-10 19:13:44: Entering into Exascale database clone creation phase 2
INFO: 2025-08-10 19:13:44: ...startup auxiliary instance
INFO: 2025-08-10 19:13:51: ...duplicating database 'CDB1_wg8_iad' as 'CDBSNAP1', please wait...
INFO: 2025-08-10 19:13:51: ....(rman log file at '/opt/gDBClone/out/log/rman_duplicate4exascale_224780.log')
INFO: 2025-08-10 19:15:41: ...creating spfile for 'CDBSNAP1'
INFO: 2025-08-10 19:15:44: ...creating pfile for 'CDBSNAP1'
INFO: 2025-08-10 19:15:44: ...creating passwordfile
INFO: 2025-08-10 19:15:44: ...copying passwordfile to exascale storage
INFO: 2025-08-10 19:15:45: ...removing temporary control.bkp
INFO: 2025-08-10 19:15:46: Entering into Exascale database clone creation phase 3
INFO: 2025-08-10 19:15:46: ...registering database 'CDBSNAP1' as cluster resource
INFO: 2025-08-10 19:15:55: ...successfully started the database
INFO: 2025-08-10 19:15:55: ...setting database 'CDBSNAP1' express port
INFO: 2025-08-10 19:16:05: ...getting external references information for the database 'CDB1_wg8_iad'
INFO: 2025-08-10 19:16:05: ...disabling the external references in the database 'CDBSNAP1' inherited from 'CDB1_wg8_iad'
-----
Enable these external references running on the database 'CDBSNAP1' the SQL script:
'/u02/app/oracle/product/23.0.0.0/dbhome_1/enable_external_refs_CDBSNAP1_rl38.sql'
Restart the database after running the SQL script is required
-----
INFO: 2025-08-10 19:16:07: ...restarting database 'CDBSNAP1'
INFO: 2025-08-10 19:16:49: ...creating pfile for 'CDBSNAP1'
INFO: 2025-08-10 19:16:49: ...cleanup Archivelog files
SUCCESS: 2025-08-10 19:17:09: Exascale clone database 'CDBSNAP1' created successfully
```

gDBClone on Exascale – Demo: MacroStep3

```
MacroStep3 - Converting clone database 'CDBSNAP1' to cluster mode...
INFO: 2025-08-10 19:17:10: Checking for 'Real Application Clusters' option
INFO: 2025-08-10 19:17:12: Database conversion started...
INFO: 2025-08-10 19:17:12: Making required redolog thread on database 'CDBSNAP1'
INFO: 2025-08-10 19:17:13: ...number of the nodes is less or equal the thread#, no more redolog thread are required
INFO: 2025-08-10 19:17:13: Making required undo tablespaces on database 'CDBSNAP1'
INFO: 2025-08-10 19:17:13: Getting UNDOTBS for 'CDBSNAP1'
INFO: 2025-08-10 19:17:14: Making required instances spfile changes
INFO: 2025-08-10 19:17:16: Stopping database 'CDBSNAP1'
INFO: 2025-08-10 19:17:32: Register 'CDBSNAP1' database as RAC database
INFO: 2025-08-10 19:17:34: Registering additional instance
INFO: 2025-08-10 19:17:37: Register Exascale passwordfile '@eGED0pGD/GUOB-C1-E680D9DCFBC24FBDBF0F628A4359374D/CDBSNAP1/PASSWORD/pwdCDBSNAP1'
INFO: 2025-08-10 19:17:38: Copying password file to the cluster nodes...
INFO: 2025-08-10 19:17:38: ...copying '/u02/app/oracle/product/23.0.0.0/dbhome_1/dbs/orapwCDBSNAP11' to 'guob011-rpdb'
INFO: 2025-08-10 19:17:38: Copying pfile to the cluster nodes...
INFO: 2025-08-10 19:17:38: ...copying '/u02/app/oracle/product/23.0.0.0/dbhome_1/dbs/initCDBSNAP11.ora' to 'guob011-rpdb'
SUCCESS: 2025-08-10 19:17:39: Database 'CDBSNAP1' converted to RAC successfully

INFO: 2025-08-10 19:17:39: Starting database 'CDBSNAP1'
SUCCESS: 2025-08-10 19:18:08: Successfully started clone database 'CDBSNAP1' 19:18:08 (5 minutes)
```

```
SQL>
SQL> @datafiles_used_per_cdb.sql
```

CDB_PATH	SIZE_GBYTES	USED_GBYTES	USED_GBYTES_RAW
@eGED0pGD/GUOB-C1-E680D9DCFBC24FBDBF0F628A4359374D/CDB1_WG8_IAD/	236.49	236.49	709.46
@eGED0pGD/GUOB-C1-E680D9DCFBC24FBDBF0F628A4359374D/CDBSNAP1/	231.96	23.25	69.75

```
SQL>
SQL> □
```

Monitoring Storage Vault usage

How to see the total capacity and space used in my Exascale Vault ?

Oracle Database - Specific views for Exascale

V\$EXA_VAULT

- Capacity and space usage in Storage Vault
- Replacement for **V\$ASM_DISKGROUP**

Oracle Database - Specific views for Exascale

V\$EXA_VAULT

- Capacity and space usage in Storage Vault
- Replacement for **V\$ASM_DISKGROUP**

V\$EXA_FILE

- Space used per file in Storage Vault
- Replacement for **V\$ASM_FILE**

Oracle Database - Specific views for Exascale

V\$EXA_VAULT

- Capacity and space usage in Storage Vault
- Replacement for **V\$ASM_DISKGROUP**

V\$EXA_FILE

- Space used per file in Storage Vault
- Replacement for **V\$ASM_FILE**

V\$EXA_TEMPLATE

- Templates that governs where and how each file type is stored in the Vault
- Replacement for **V\$ASM_TEMPLATE**

Exascale Monitoring Vault

```
SQL> SELECT * FROM V$EXA_VAULT;
```

```
sys@cdb1
sys@CDB1>@exavault2

Vault          High Capacity   High Capacity   High Capacity   High Capacity
Name          Provisioned (GB) Space Used (GB)  Space Free (GB) Space Used (%)
-----
eGED0pGD          600.00          300.58          299.42          50.10

1 row selected.
```

Exascale Monitoring Vault

```
SQL> SELECT * FROM V$EXA_VAULT;
```

```
sys@CDB1>@exavault2

Vault Name          High Capacity Provisioned (GB)  High Capacity Space Used (GB)  High Capacity Space Free (GB)  High Capacity Space Used (%)
-----
eGED0pGD            600.00           300.58                       299.42                       50.10

1 row selected.
```

```
SQL> SELECT * FROM V$EXA_FILE;
```

```
sys@CDB1>@exavault_used_per_cdb.sql

CDB_PATH                                                    SIZE_GBYTES
-----
@eGED0pGD/GU0B-C1-E680D9DCFBC24FBDBF0F628A4359374D/CDB1_WG8_IAD/ 246.95
@eGED0pGD/GU0B-C1-E680D9DCFBC24FBDBF0F628A4359374D/CDB1THIN/      6.02
@eGED0pGD/GU0B-C1-E680D9DCFBC24FBDBF0F628A4359374D/CDBCLONE/      .33

3 rows selected.
```

Exascale Monitoring Vault

```
SQL> SELECT * FROM V$EXA_VAULT;
```

```
sys@CDB1>@exavault2

Vault          High Capacity   High Capacity   High Capacity   High Capacity
Name          Provisioned (GB) Space Used (GB)  Space Free (GB) Space Used (%)
-----
eGED0pGD          600.00         300.58         299.42         50.10

1 row selected.
```

```
SQL> SELECT * FROM V$EXA_FILE;
```

```
sys@CDB1>@exavault_used_per_cdb.sql

CDB_PATH                                             SIZE_GBYTES
-----
@eGED0pGD/GU0B-C1-E680D9DCFBC24FBDBF0F628A4359374D/CDB1_WG8_IAD/ 246.95
@eGED0pGD/GU0B-C1-E680D9DCFBC24FBDBF0F628A4359374D/CDB1THIN/      6.02
@eGED0pGD/GU0B-C1-E680D9DCFBC24FBDBF0F628A4359374D/CDBCLONE/      .33

3 rows selected.
```

```
$ df -t acfs
```

```
sys@CDB1>
sys@CDB1>! df -h -t acfs
Filesystem      Size  Used Avail Use% Mounted on
/dev/exc/vExdb82  50G  7.6G  43G  16% /var/opt/oracle/dbaas_acfs
```

Exascale Monitoring Vault

ACFS space comes from Storage Vault
(Exascale Direct Volume)

```
SQL> SELECT * FROM V$EXA_VAULT;
```

```
sys@cdb1>@exavault2

Vault Name          High Capacity Provisioned (GB)  High Capacity Space Used (GB)  High Capacity Space Free (GB)  High Capacity Space Used (%)
-----
eGED0pGD            600.00           300.58                        299.42                        50.10

1 row selected.
```

```
SQL> SELECT * FROM V$EXA_FILE;
```

```
sys@cdb1>@exavault_used_per_cdb.sql

CDB_PATH                                                    SIZE_GBYTES
-----
@eGED0pGD/GU0B-C1-E680D9DCFBC24FBDBF0F628A4359374D/CDB1_WG8_IAD/ 246.95
@eGED0pGD/GU0B-C1-E680D9DCFBC24FBDBF0F628A4359374D/CDB1THIN/      6.02
@eGED0pGD/GU0B-C1-E680D9DCFBC24FBDBF0F628A4359374D/CDBCLONE/      .33

3 rows selected.
```

```
$ df -t acfs
```

```
sys@cdb1>
sys@cdb1>! df -h -t acfs
Filesystem          Size  Used Avail Use% Mounted on
/dev/exc/vExdb82   50G   7.6G   43G   16% /var/opt/oracle/dbaas_acfs
```

Exascale Tools

Command line tools to manage Exascale

New specific tools for Exascale

Exadata admins

ESCLI

Addition to **CELLCLI**

▼ 6 Using ESCLI

6.1 Start and Use ESCLI

▼ 6.2 ESCLI Command Reference

6.2.1 ESCLI Command Help

6.2.2 Describing Resources and Attributes

▶ 6.2.3 Service and Cluster Management

▶ 6.2.4 Security and User Management

▶ 6.2.5 Storage Pool and Pool Disk Management

Cluster & DB admins

XSH

Replacement for **asmcmd**

▼ 7 Using XSH

7.1 Start and Use XSH

▼ 7.2 XSH Command Reference

7.2.1 cat

7.2.2 chacl

7.2.3 chresourceprofile

7.2.4 chtemplate

7.2.5 clone

7.2.6 cp

Exascale Tools - XSH

```
[oracle@guob011-rpdrb ~]$ xsh ls
@eGED0pGD
[oracle@guob011-rpdrb ~]$
[oracle@guob011-rpdrb ~]$
```

```
[oracle@guob011-rpdrb ~]$ xsh ls @eGED0pGD
@eGED0pGD/GUOB-C1-E680D9DCFBC24FBDBF0F628A4359374D/CDB1_WG8_IAD/24C26E584BA61F76E0633307F40A7940/DATAFILE/SYSAUX.OMF.622E632C
@eGED0pGD/GUOB-C1-E680D9DCFBC24FBDBF0F628A4359374D/CDB1_WG8_IAD/24C26E584BA61F76E0633307F40A7940/DATAFILE/SYSTEM.OMF.20DC0CA5
@eGED0pGD/GUOB-C1-E680D9DCFBC24FBDBF0F628A4359374D/CDB1_WG8_IAD/24C26E584BA61F76E0633307F40A7940/DATAFILE/UNDOTBS1.OMF.01A6FEC3
@eGED0pGD/GUOB-C1-E680D9DCFBC24FBDBF0F628A4359374D/CDB1_WG8_IAD/24C26E584BA61F76E0633307F40A7940/DATAFILE/temp012025-04-12_15-52-32-684-PM.dbf
@eGED0pGD/GUOB-C1-E680D9DCFBC24FBDBF0F628A4359374D/CDB1_WG8_IAD/329A723565C06874E0630E02020A3A53/DATAFILE/IOPS.OMF.18F795FA
@eGED0pGD/GUOB-C1-E680D9DCFBC24FBDBF0F628A4359374D/CDB1_WG8_IAD/329A723565C06874E0630E02020A3A53/DATAFILE/SYSAUX.OMF.369CB05A
@eGED0pGD/GUOB-C1-E680D9DCFBC24FBDBF0F628A4359374D/CDB1_WG8_IAD/329A723565C06874E0630E02020A3A53/DATAFILE/SYSTEM.OMF.208D2075
@eGED0pGD/GUOB-C1-E680D9DCFBC24FBDBF0F628A4359374D/CDB1_WG8_IAD/329A723565C06874E0630E02020A3A53/DATAFILE/UNDOTBS1.OMF.0C3F1EA0
@eGED0pGD/GUOB-C1-E680D9DCFBC24FBDBF0F628A4359374D/CDB1_WG8_IAD/329A723565C06874E0630E02020A3A53/DATAFILE/UNDO_7.OMF.52707605
@eGED0pGD/GUOB-C1-E680D9DCFBC24FBDBF0F628A4359374D/CDB1_WG8_IAD/329A723565C06874E0630E02020A3A53/DATAFILE/USERS.OMF.6B26596A
@eGED0pGD/GUOB-C1-E680D9DCFBC24FBDBF0F628A4359374D/CDB1_WG8_IAD/329A723565C06874E0630E02020A3A53/TEMPFILE/TEMP.OMF.6501DCDD
@eGED0pGD/GUOB-C1-E680D9DCFBC24FBDBF0F628A4359374D/CDB1_WG8_IAD/36384A6A84C0CBF3E0630E02020AEA70/DATAFILE/SH.OMF.770FB600
@eGED0pGD/GUOB-C1-E680D9DCFBC24FBDBF0F628A4359374D/CDB1_WG8_IAD/36384A6A84C0CBF3E0630E02020AEA70/DATAFILE/SYSAUX.OMF.5B5F8397
@eGED0pGD/GUOB-C1-E680D9DCFBC24FBDBF0F628A4359374D/CDB1_WG8_IAD/36384A6A84C0CBF3E0630E02020AEA70/DATAFILE/SYSTEM.OMF.52F72F0E
@eGED0pGD/GUOB-C1-E680D9DCFBC24FBDBF0F628A4359374D/CDB1_WG8_IAD/36384A6A84C0CBF3E0630E02020AEA70/DATAFILE/UNDOTBS1.OMF.6EFC3993
@eGED0pGD/GUOB-C1-E680D9DCFBC24FBDBF0F628A4359374D/CDB1_WG8_IAD/36384A6A84C0CBF3E0630E02020AEA70/DATAFILE/UNDO_5.OMF.749C18DB
@eGED0pGD/GUOB-C1-E680D9DCFBC24FBDBF0F628A4359374D/CDB1_WG8_IAD/ARCHIVELOG/2025_08_10/thread_1_seq_4779.OMF.1B0E6069
@eGED0pGD/GUOB-C1-E680D9DCFBC24FBDBF0F628A4359374D/CDB1_WG8_IAD/ARCHIVELOG/2025_08_10/thread_1_seq_4779.OMF.2919DDD1
@eGED0pGD/GUOB-C1-E680D9DCFBC24FBDBF0F628A4359374D/CDB1_WG8_IAD/ARCHIVELOG/2025_08_10/thread_1_seq_4780.OMF.0985295A
@eGED0pGD/GUOB-C1-E680D9DCFBC24FBDBF0F628A4359374D/CDB1_WG8_IAD/ARCHIVELOG/2025_08_10/thread_1_seq_4780.OMF.37D9FB9D
@eGED0pGD/GUOB-C1-E680D9DCFBC24FBDBF0F628A4359374D/CDB1_WG8_IAD/ARCHIVELOG/2025_08_10/thread_1_seq_4781.OMF.0A283B8F
@eGED0pGD/GUOB-C1-E680D9DCFBC24FBDBF0F628A4359374D/CDB1_WG8_IAD/ARCHIVELOG/2025_08_10/thread_1_seq_4781.OMF.24E55B1B
@eGED0pGD/GUOB-C1-E680D9DCFBC24FBDBF0F628A4359374D/CDB1_WG8_IAD/ARCHIVELOG/2025_08_10/thread_1_seq_4782.OMF.37EF0910
@eGED0pGD/GUOB-C1-E680D9DCFBC24FBDBF0F628A4359374D/CDB1_WG8_IAD/ARCHIVELOG/2025_08_10/thread_1_seq_4782.OMF.649F1347
@eGED0pGD/GUOB-C1-E680D9DCFBC24FBDBF0F628A4359374D/CDB1_WG8_IAD/ARCHIVELOG/2025_08_10/thread_1_seq_4783.OMF.4094C2D4
```

Exascale Tools - XSH

```
[grid@guob011-rpdrb ~]$  
[grid@guob011-rpdrb ~]$ xsh ls @eGED0pGD  
@eGED0pGD/GUOB-C1-E680D9DCFBC24FBDBF0F628A4359374D/backup00.ocr  
@eGED0pGD/GUOB-C1-E680D9DCFBC24FBDBF0F628A4359374D/backup01.ocr  
@eGED0pGD/GUOB-C1-E680D9DCFBC24FBDBF0F628A4359374D/backup02.ocr  
@eGED0pGD/GUOB-C1-E680D9DCFBC24FBDBF0F628A4359374D/backup_20250412_174515.ocr  
@eGED0pGD/GUOB-C1-E680D9DCFBC24FBDBF0F628A4359374D/data.ocr  
@eGED0pGD/GUOB-C1-E680D9DCFBC24FBDBF0F628A4359374D/day.ocr  
@eGED0pGD/GUOB-C1-E680D9DCFBC24FBDBF0F628A4359374D/vfile1  
@eGED0pGD/GUOB-C1-E680D9DCFBC24FBDBF0F628A4359374D/week.ocr  
@eGED0pGD/GUOB-C1-E680D9DCFBC24FBDBF0F628A4359374D/week_.ocr  
@eGED0pGD/vol.55b73f6ef22544a2abd0183b48b8da0c
```

```
[grid@guob011-rpdrb ~]$  
[grid@guob011-rpdrb ~]$ xsh cp @eGED0pGD/GUOB-C1-E680D9DCFBC24FBDBF0F628A4359374D/week.ocr /tmp/week.ocr  
[grid@guob011-rpdrb ~]$  
[grid@guob011-rpdrb ~]$ ls -lrt /tmp/week.ocr  
-rw-rw---- 1 grid oinstall 3059712 Aug 10 18:02 /tmp/week.ocr  
[grid@guob011-rpdrb ~]$  
[grid@guob011-rpdrb ~]$ █
```

Exascale – Deployment Options

Where we can use Exascale

Exascale on Exadata OnPremise

Storage Server Details
Exascale
 Enable Exascale

Exascale

Rack capacity (raw): **737604 GB**
Rack used space (raw): **368802 GB**
Rack available space (raw): **368802 GB**

ExascaleCluster1

Advanced

Cluster Name: ExascaleCluster1
REST IP Address: 192.168.1.17
REST Hostname: dbm0ers0.dibiei.com
ERS Network ID: 194

Select Exascale cluster nodes

Available nodes

- Cell dbm0celadm01
Exadata X11M Cell Node HC 22TB
- Cell dbm0celadm02
Exadata X11M Cell Node HC 22TB
- Cell dbm0celadm03
Exadata X11M Cell Node HC 22TB

Selected nodes

- Cell dbm0celadm01
Exadata X11M Cell Node HC 22TB
- Cell dbm0celadm02
Exadata X11M Cell Node HC 22TB
- Cell dbm0celadm03
Exadata X11M Cell Node HC 22TB

Storage Pool Details

Storage Pool Size(%/GB/TB): 50%
Storage Pool size: [disabled]
Storage Pool Name: hcpool1

Save

Exascale on OCI & Multicloud

Exadata Cloud at Customer
(ExaCC)

Exadata Cloud Service
(ExaCS)

Exadata on Exascale Infrastructure
(ExaXS)

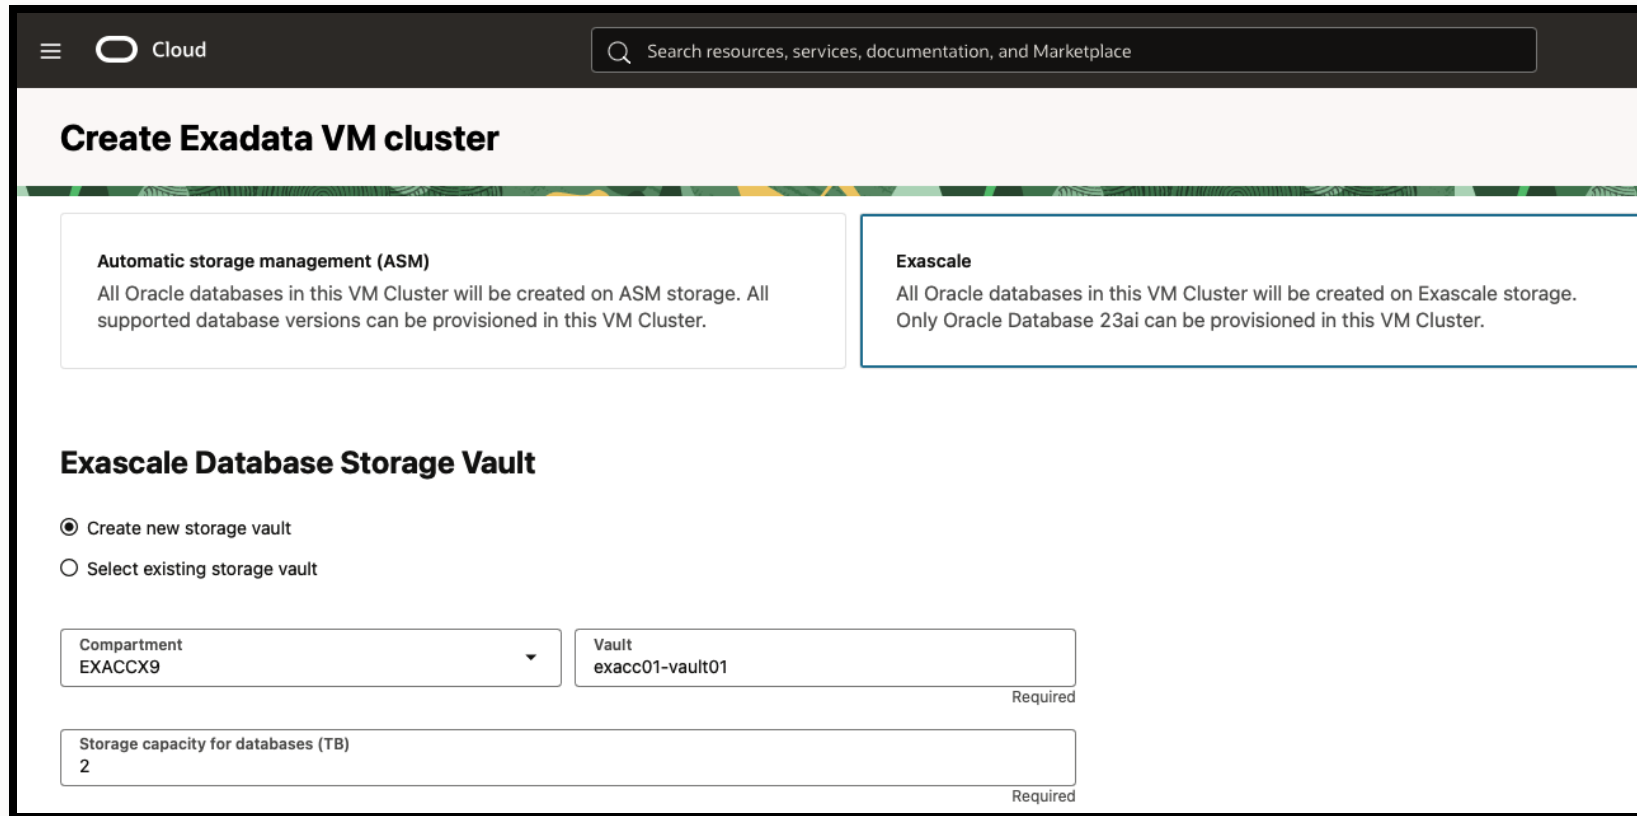


ORACLE®
Cloud Infrastructure

Enable Exascale on ExaCC or ExaCS

The screenshot shows the Oracle Cloud console interface. On the left, a navigation menu is visible with the following items: Oracle Database, Oracle Exadata Database Service on Cloud@Customer, Exadata VM Clusters, Autonomous Databases, Autonomous Container Databases, Autonomous Exadata VM Clusters, and Exadata Infrastructure (highlighted). The main content area shows the 'Exadata Infrastructure' page for 'EXACC01', which is in an 'Active' state. A warning banner is present, stating 'You can now configure Exascale storage on this infrastructure. Learn more'. A modal dialog box titled 'Configure Exascale storage' is overlaid on the screen. It contains the text 'You can now configure Exascale storage on this infrastructure. Learn more' and a text input field for 'Storage capacity for Exascale (TB)' with the value '2'. A 'Required' label is positioned to the right of the input field. Below the input field, another warning banner states: 'Warning: The DB servers will get rebooted to enable Exascale configuration.'

Create VM Cluster with Exascale on ExaCC or ExaCS



Create Exadata VM cluster

Automatic storage management (ASM)
All Oracle databases in this VM Cluster will be created on ASM storage. All supported database versions can be provisioned in this VM Cluster.

Exascale
All Oracle databases in this VM Cluster will be created on Exascale storage. Only Oracle Database 23ai can be provisioned in this VM Cluster.

Exascale Database Storage Vault

Create new storage vault
 Select existing storage vault

Compartment
EXACC9

Vault
exacc01-vault01
Required

Storage capacity for databases (TB)
2
Required

Exadata Database Service on Exascale Infrastructure

Shared Exadata Infrastructure

Exadata Database Service on Exascale Infrastructure (ExaXS)

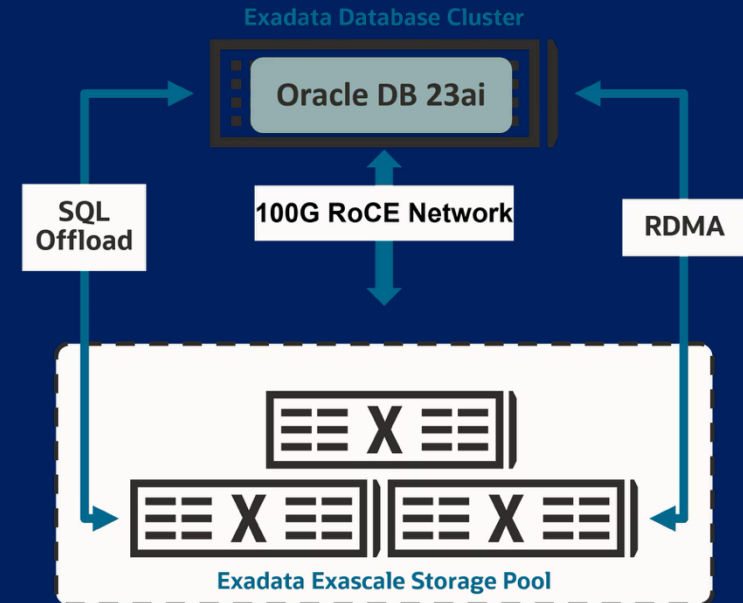
1. Start with a Single VM with **8 ECPU** and **22 GB** of Memory



2. Start with **300 GB** of Storage Vault



3. **Pay per use** (CPU billing per second – Minimum 48h)



ExaXS - Example

The screenshot shows the Oracle Cloud console interface. At the top, there is a search bar with the text "Search resources, services, documentation, and Marketplace". Below the search bar, the "Oracle Database" page is displayed. The left sidebar contains a navigation menu with the following items: Home, Compute, Storage, Networking, Oracle Database (highlighted with a dashed box), Databases, Analytics & AI, Developer Services, and Identity & Security. The main content area is titled "Oracle Database" and contains several sections:

- Overview**
- Autonomous Database**
 - [Autonomous Data Warehouse](#)
 - [Autonomous JSON Database](#)
 - [Autonomous Transaction Processing](#)
- Globally Distributed Autonomous Database**
- Autonomous Dedicated Infrastructure**

On the right side of the page, there are three service options:

- Oracle Base Database Service**
- Oracle Exadata Database Service on Dedicated Infrastructure**
- Oracle Exadata Database Service on Exascale Infrastructure** (highlighted with a red box)
- Oracle Exadata Database Service on Cloud@Customer**
- Exadata Fleet Update**

ExaXS - Example

Cloud

Search resources, services, documentation, and Marketplace

Overview » Exadata Database Service on Exascale Infrastructure » VM Clusters

Exadata Database Service on Exascale Infrastructure

VM Clusters

Exascale Storage Vaults

Resources

Software images

Standalone backups

Interim software updates

VM Clusters *in* dibiei-database *compartment*

The VM Cluster is where Oracle Databases are deployed.

Create VM Cluster


Display name	State	VM count
--------------	-------	----------

ExaXS - Example

☰ Cloud

Create VM Cluster

Provide basic information for Exadata VM Cluster on Exascale Infrastructure

Compartment
dibiei-database 
dibieilab (root)/blog-dibiei/dibiei-database

Display name
exaxs01

Provide the cluster name
exaxs01-c

Select an availability domain

AD-1 <small>rBqO:US-ASHBURN-AD-1</small>	AD-2 <small>rBqO:US-ASHBURN-AD-2</small>	AD-3 <small>rBqO:US-ASHBURN-AD-3</small> ✓
---	---	---

Resources are not presently available in all ADs in this region.

ExaXS - Example

Configure the VM Cluster

Number of VMs in the cluster

Specify a number between 1-10

ECPUs enabled per VM

Max 200. Input needs to be multiple of 4. See [reserve additional ECPU to increase](#).

 [Reserve additional ECPU to increase](#)

Memory per VM (GB) *Read-only*

Memory is calculated based on 2.75 GB per total ECPU.

Total ECPUs enabled across VM Cluster *Read-only*

Total memory across VM Cluster (GB) *Read-only*

VM file system storage

File system storage capacity per VM (GB)

Specify storage capacity (in GB) between 280 and 1100.

Total file system storage capacity across VM Cluster *Read-only*

ExaXS - Example

Exascale Database Storage Vault

Create new vault Select existing vault

Vault in **dibei-database** [\(Change compartment\)](#)

vault01

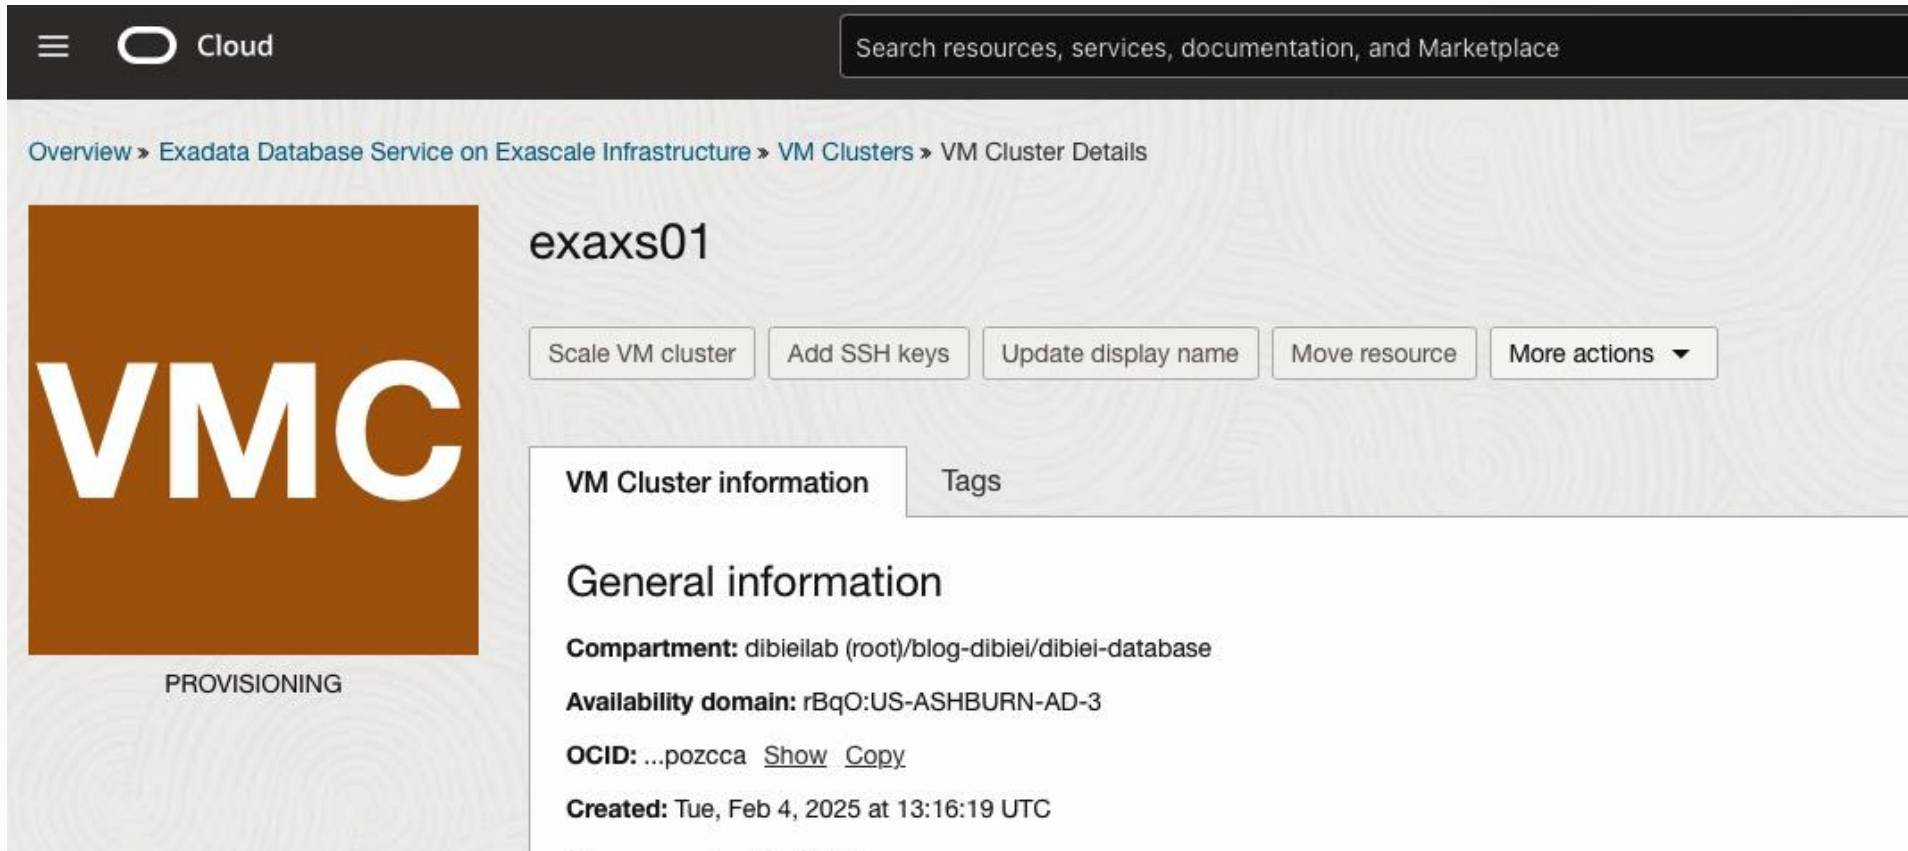
Storage capacity for Databases (GB)

300

Specify a number between 300 to 100000.

 [Show add smart flash and memory cache options](#)

ExaXS - Example



The screenshot displays the Oracle Cloud console interface for a VM Cluster. The breadcrumb navigation path is: Overview > Exadata Database Service on Exascale Infrastructure > VM Clusters > VM Cluster Details. The cluster name is 'exaxs01'. A large orange square with the text 'VMC' is shown in a 'PROVISIONING' state. Action buttons include 'Scale VM cluster', 'Add SSH keys', 'Update display name', 'Move resource', and 'More actions'. The 'VM Cluster information' tab is active, showing the following details:

- General information**
- Compartment:** dibieilab (root)/blog-dibiei/dibiei-database
- Availability domain:** rBqO:US-ASHBURN-AD-3
- OCID:** ...pozcca [Show](#) [Copy](#)
- Created:** Tue, Feb 4, 2025 at 13:16:19 UTC

ExaXS - Service Limits

Limits Name	Description	Limits	Value
exadbxs-vm-instance-base-count	Exadata Database Service on Exascale Infrastructure - Instance Count	Number of VM Instances	4
exadbxs-total-cpu-base-count	Exadata Database Service on Exascale Infrastructure - Total ECPU Count	TotalCpuCores	64
exadbxs-local-storage-base-gb	Exadata Database Service on Exascale Infrastructure - Local Storage (GB)	Local Storage (in GB)	1500
exadbxs-hc-storage-base-gb	Exadata Database Service on Exascale Infrastructure - High Capacity Storage (GB)	High capacity storage (in GB)	2000

<https://docs.oracle.com/en/engineered-systems/exadata-database-exascale/exdxs/overview-exadb-xs-service.html#GUID-E332834C-B356-443D-BF66-F8D8E68229D8>



Thank you

Maicon Carneiro

GUOB Tech Day 2025