



Photo by [Sinisa Lekovic](#) on [Unsplash](#)

Upgrade, Migration and Patching to Oracle AI Database 26ai

Advanced Edition

Berlin, February 2026



O

Oracle
DBAs
run the world





Mike Dietrich

Vice President

 mikedietrich

 @mikedietrichde.com

 <https://mikedietrichde.com>



Daniel Overby Hansen

Distinguished Product Manager

 dohdatabase

 @dohdatabase.com

 <https://dohdatabase.com>

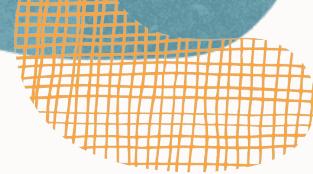
Get the Slides

<https://dohdatabase.com/slides>





Find Slides and Much More on Our Blogs



MikeDietrichDE.com

Mike.Dietrich@oracle.com



dohdatabase.com

Daniel.Overby.Hansen@oracle.com



DBArj.com.br

Rodrigo.R.Jorge@oracle.com



AlexZaballa.com

Alex.Zaballa@oracle.com



| Web Seminar | Slides |
|--|--|
| <p>Episode 16</p> <p>(replaces Episode 1 from Feb 2021)</p> <p>Oracle Database Release and Patching Strategy for 19c and 23c</p> <p>115 minutes – May 10, 2023</p> |  |
| <p>Episode 17</p> <p>From SR to Patch – Insights into the Oracle Database Development process</p> <p>55 minutes – June 22, 2023</p> |  |
| <p>*NEW* Episode 18</p> <p>Cross Platform Migration – Transportable Tablespaces to the Extreme</p> <p>145 min – February 22, 2024</p> |  |
| <p>Episode 2</p> <p>AutoUpgrade to Oracle Database 19c</p> <p>115 minutes – Feb 20, 2021</p> |  |
| <p>Episode 3</p> <p>Performance Stability, Tips and Tricks and Underscores</p> <p>120 minutes – Mar 4, 2021</p> |  |
| <p>Episode 4</p> <p>Migration to Oracle Multitenant</p> |  |

Recorded Web Seminars

<https://MikeDietrichDE.com/videos>

More than 40 hours of technical content, on-demand, anytime, anywhere



Scan me to sign up

Database Patching for DBAs

– Patch smarter, not harder

February 12, 14:00 CET

[Sign up](#)

AGENDA

09:30

Welcome
Release Strategy
Patching

11:15

AutoUpgrade
Multitenant

13:30

Data Pump
Autonomous AI Database

15:15

Migrating
"The Beast"

11:00

Coffee break

12:45

Lunch

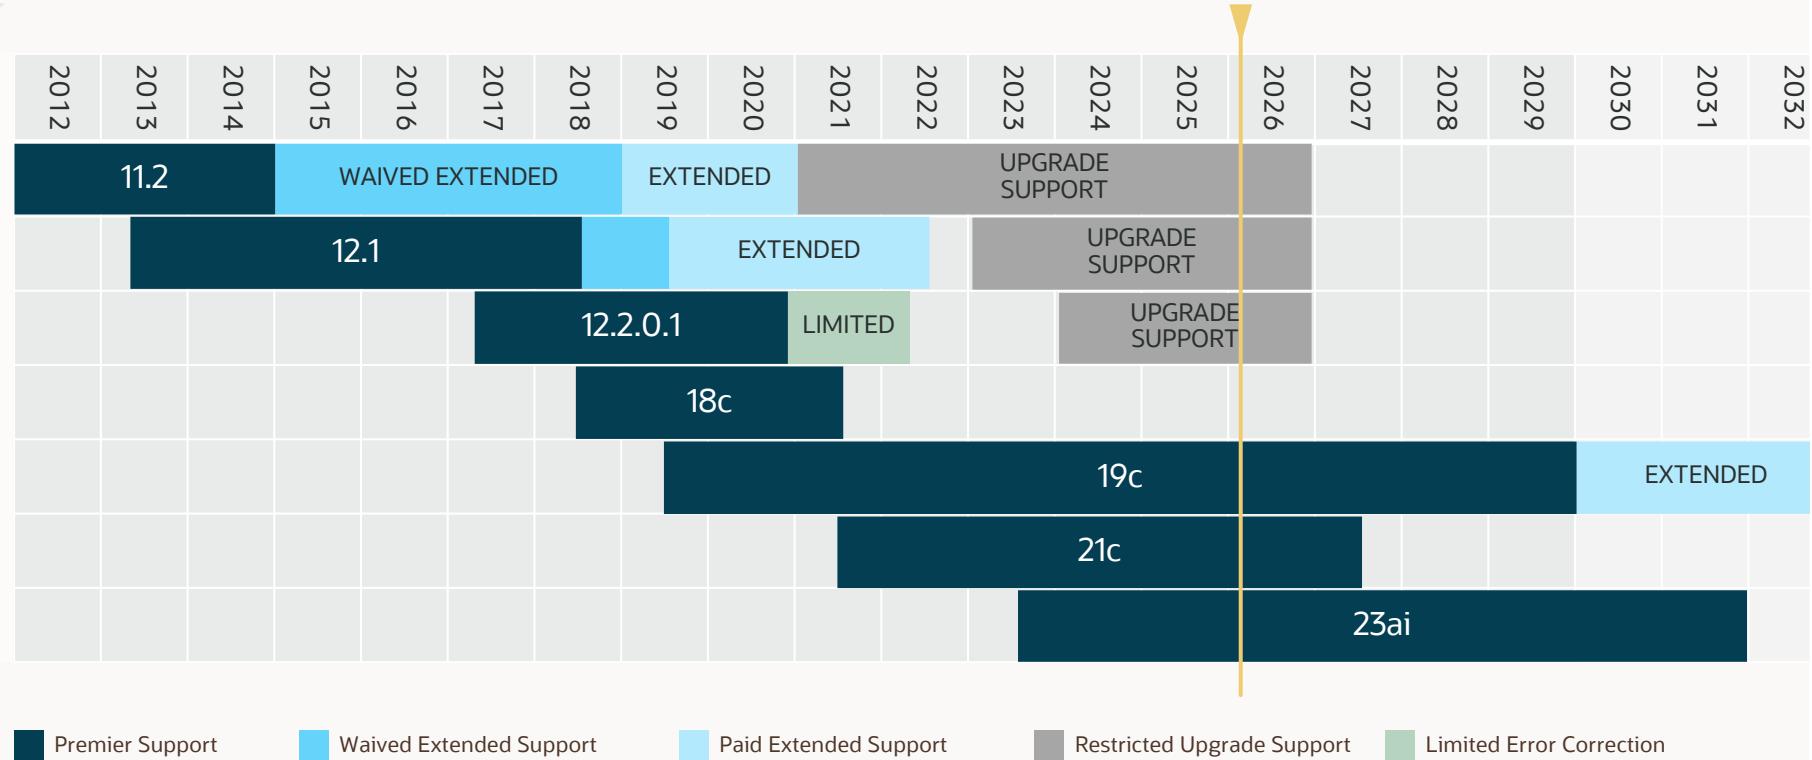
15:00

Coffee break

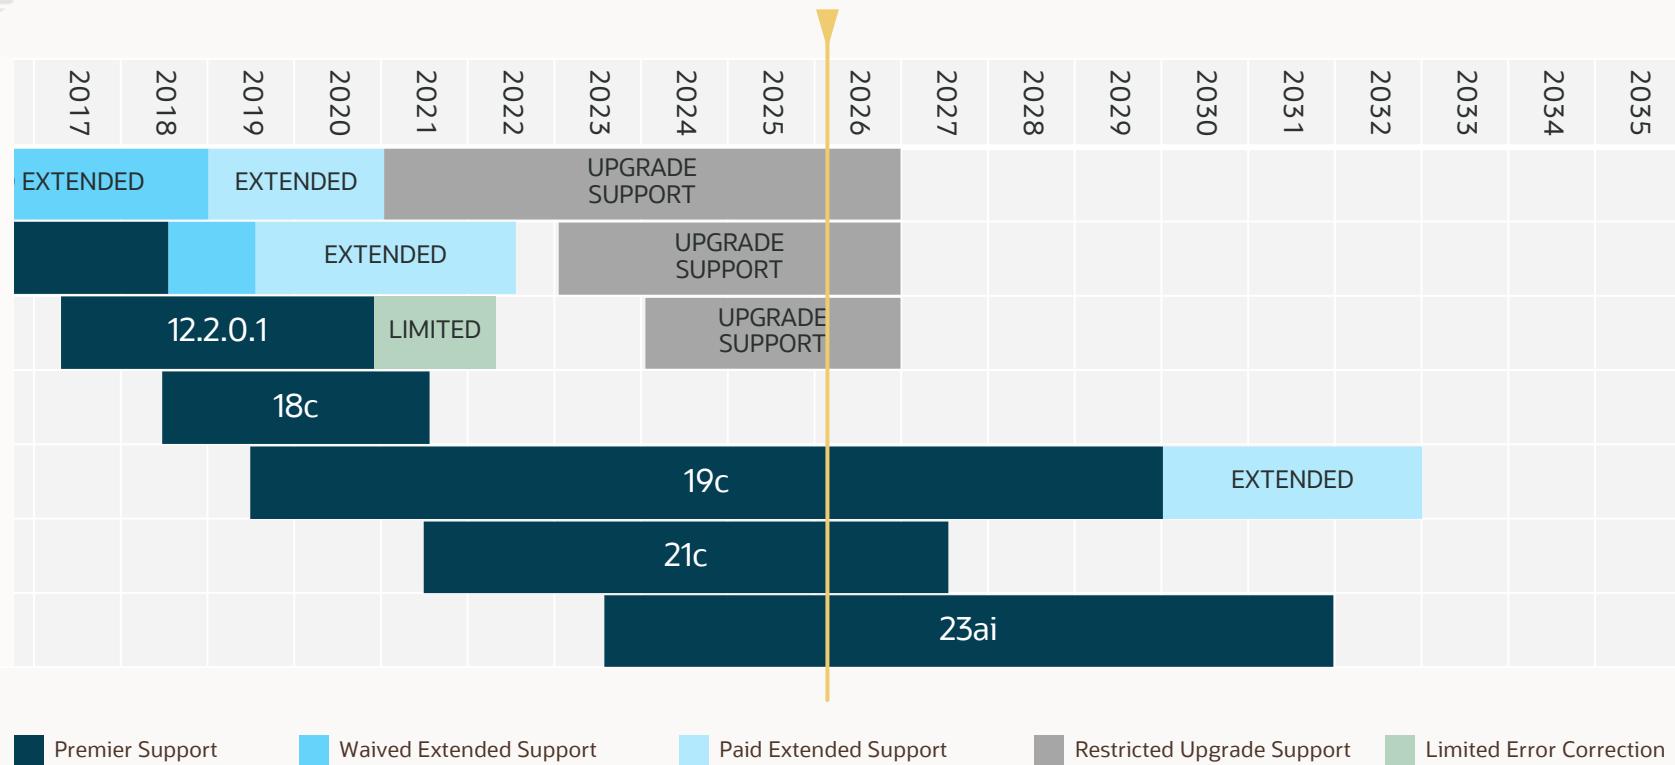


Release Strategy

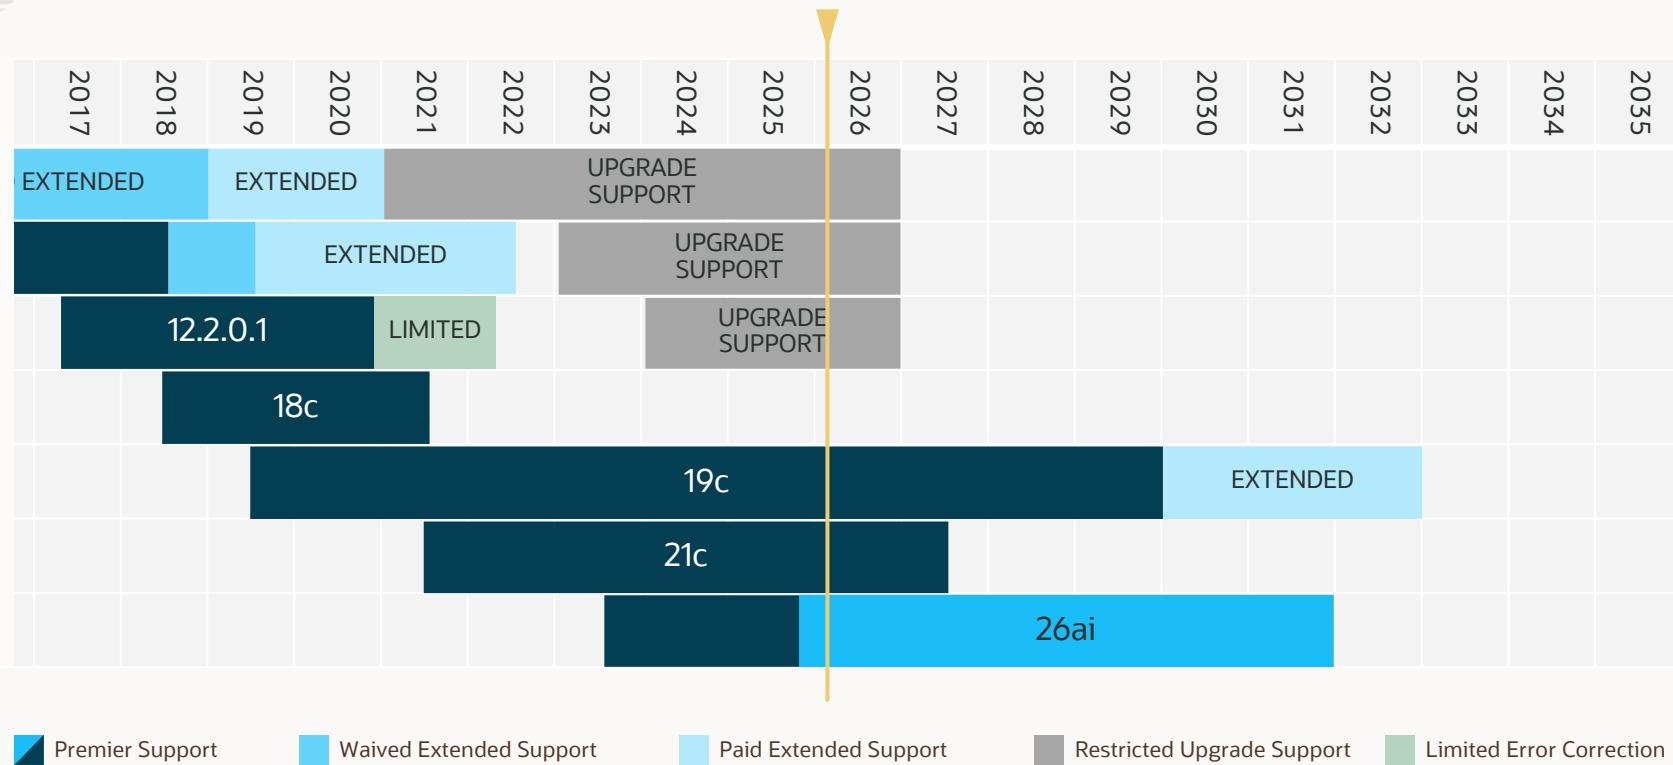
Lifetime Support Policy



Lifetime Support Policy



Lifetime Support Policy



Do you want to upgrade?

Oracle Database 11.2.0.4

Oracle Database 12.1.0.2

Oracle Database 12.2.0.1

Oracle Database 18c

Oracle Database 11.2.0.4
Oracle Database 12.1.0.2
Oracle Database 12.2.0.1
Oracle Database 18c



Oracle Database 19c



Oracle Database 23ai

Oracle Database 11.2.0.4
Oracle Database 12.1.0.2
Oracle Database 12.2.0.1
Oracle Database 18c



Oracle Database 19c



Oracle Database 23ai
Oracle AI Database 26ai

26^{ai}

When is a database upgrade
required?

Oracle Database 19c \Rightarrow Oracle Database 23ai \Rightarrow Oracle AI Database 26ai

UPGRADE

Oracle Database 19c



Oracle AI Database 26ai

UPGRADE

Oracle Database 19c \Rightarrow Oracle Database 23ai \Rightarrow Oracle AI Database 26ai

UPDATE

July 2025

Oracle Database 23ai

23.9.0.25.07

October 2025

Oracle AI Database 26ai

23.26.0.0.0

January 2026

Oracle AI Database 26ai

23.26.1.0.0

April 2026

Oracle AI Database 26ai

23.26.2.0.0

| | | |
|--------------|-------------------------|--------------|
| July 2025 | Oracle Database 23ai | 23.9.0.25.07 |
| October 2025 | Oracle AI Database 26ai | 23.26.0.0.0 |
| January 2026 | Oracle AI Database 26ai | 23.26.1.0.0 |
| April 2026 | Oracle AI Database 26ai | 23.26.2.0.0 |

July 2025

Oracle Database 23ai

23.9.0.25.07

October 2025

Oracle AI Database 26ai

23.26.0.0.0

January 2026

Oracle AI Database 26ai

23.26.1.0.0

April 2026

Oracle AI Database 26ai

23.26.2.0.0

| | | |
|--------------|-------------------------|--------------|
| July 2025 | Oracle Database 23ai | 23.9.0.25.07 |
| October 2025 | Oracle AI Database 26ai | 23.26.0.0.0 |
| January 2026 | Oracle AI Database 26ai | 23.26.1.0.0 |
| April 2026 | Oracle AI Database 26ai | 23.26.2.0.0 |

| | | |
|--------------|-------------------------|--------------|
| July 2025 | Oracle Database 23ai | 23.9.0.25.07 |
| October 2025 | Oracle AI Database 26ai | 23.26.0.0.0 |
| January 2026 | Oracle AI Database 26ai | 23.26.1.0.0 |
| April 2026 | Oracle AI Database 26ai | 23.26.2.0.0 |

| | | |
|--------------|-------------------------|--------------|
| July 2025 | Oracle Database 23ai | 23.9.0.25.07 |
| October 2025 | Oracle AI Database 26ai | 23.26.0.0.0 |
| January 2026 | Oracle AI Database 26ai | 23.26.1.0.0 |
| April 2026 | Oracle AI Database 26ai | 23.26.2.0.0 |

26^{ai}

When will it be available for
non-Oracle hardware on-premises?

- See [MOS Note: PNEWS1360](#)

On-Premises Server Releases (includes client)**(Download here)**

| | | | | | | | | | |
|--------------------------------------|---|---|--|------------------------|------------------------|------------------------|------------------------|------------------------|------------------------|
| Linux x86 | <i>Not Planned</i> | <i>Not Planned</i> | <i>Not planned</i> | <i>Not planned</i> | <i>Not planned</i> | <i>Not planned</i> | <i>Not planned</i> | <i>Not planned</i> | 28-Aug-2013 |
| Linux x86-64 | January 2026 Release Update (23.26.1) (Instant Client and Full Client are available today) | 13-Aug-2021 | 25-Apr-2019 | 23-Jul-2018 | 1-Mar-2017 | 22-Jul-2014 | 25-Jun-2013 | 27-Aug-2013 | |
| Linux on Arm | TBA | <i>Not Planned</i> | 28-June-2023 Client: May-2021 download | <i>Not Planned</i> | <i>Not Planned</i> | <i>Not Planned</i> | <i>Not Planned</i> | <i>Not Planned</i> | <i>Not Planned</i> |
| Oracle Solaris SPARC (64- | TBA | See KR126264 | 26-Apr-2019 | 30-Jul-2018 | 1-Mar-2017 | 22-Jul-2014 | 25-Jun-2013 | 29-Aug-2013 | |

Patching

Question

How often do you patch your Oracle environments?

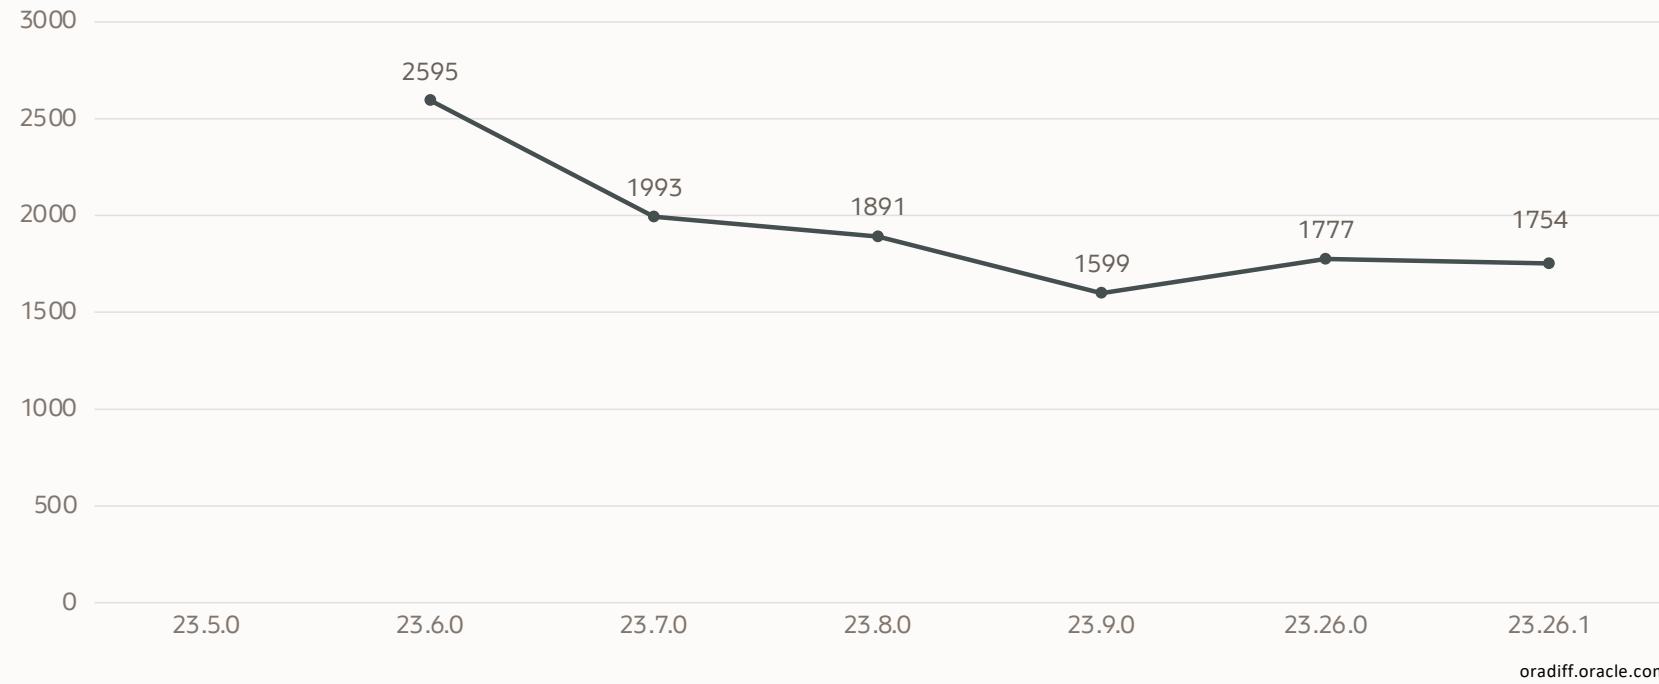
- Every quarter
- Twice per year
- Once per year
- Never

Release Update Contents – Oracle Database 19c



[Database 19 Release Updates and Revisions Bugs Fixed Lists \(Doc ID 2523220.1\)](#)

Release Update Contents – Oracle AI Database 26ai



oradiff.oracle.com

If you don't apply a recent Release Update,
you will miss **thousands** of fixes

- More than 14k fixes with 19.29.0
- More than 750 security fixes
- Almost 10k fixes with 23.26.0
- More than 850 security fixes

Just as easy as patching your smart phone

AutoUpgrade's mission for patching Oracle Database

AutoUpgrade and Patching



1 Download

- Find the right patch numbers
- For the right platform
- Get latest OPatch



2 Install

- Install brand-new Oracle home
- Update OPatch
- Apply all patches



3 Patch

- Datapatch Sanity Check
- Move instances and files
- Datapatch
- Recompilation
- Post-tasks

```
$ cat just_patch.cfg
```

```
patch1.source_home=/u01/app/oracle/product/19/dbhome_19_29_0
patch1.target_home=/u01/app/oracle/product/19/dbhome_19_30_0
patch1.sid=DB19
```



AutoUpgrade and Patching



1 Download

- Find the right patch numbers
- For the right platform
- Get latest OPatch



2 Install

- Install brand-new Oracle home
- Update OPatch
- Apply all patches



3 Patch

- Datapatch Sanity Check
- Move instances and files
- Datapatch
- Recompilation
- Post-tasks

```
$ cat do_it_all.cfg
```

```
global.keystore=/home/oracle/autougrade-patching/keystore
patch1.source_home=/u01/app/oracle/product/19/dbhome_19_29_0
patch1.target_home=/u01/app/oracle/product/19/dbhome_19_30_0
patch1.sid=DB19
patch1.folder=/home/oracle/autougrade-patching/patch
patch1.patch=RU,OPATCH,OJVM,DPBP
```

```
$ cat do_it_all.cfg
```

```
global.keystore=/home/oracle/autoupgrade-patching/keystore
patch1.source_home=/u01/app/oracle/product/19/dbhome_19_29_0
patch1.target_home=/u01/app/oracle/product/19/dbhome_19_30_0
patch1.sid=DB19
patch1.folder=/home/oracle/autoupgrade-patching/patch
patch1.patch=RU,OPATCH,OJVM,DPBP
```

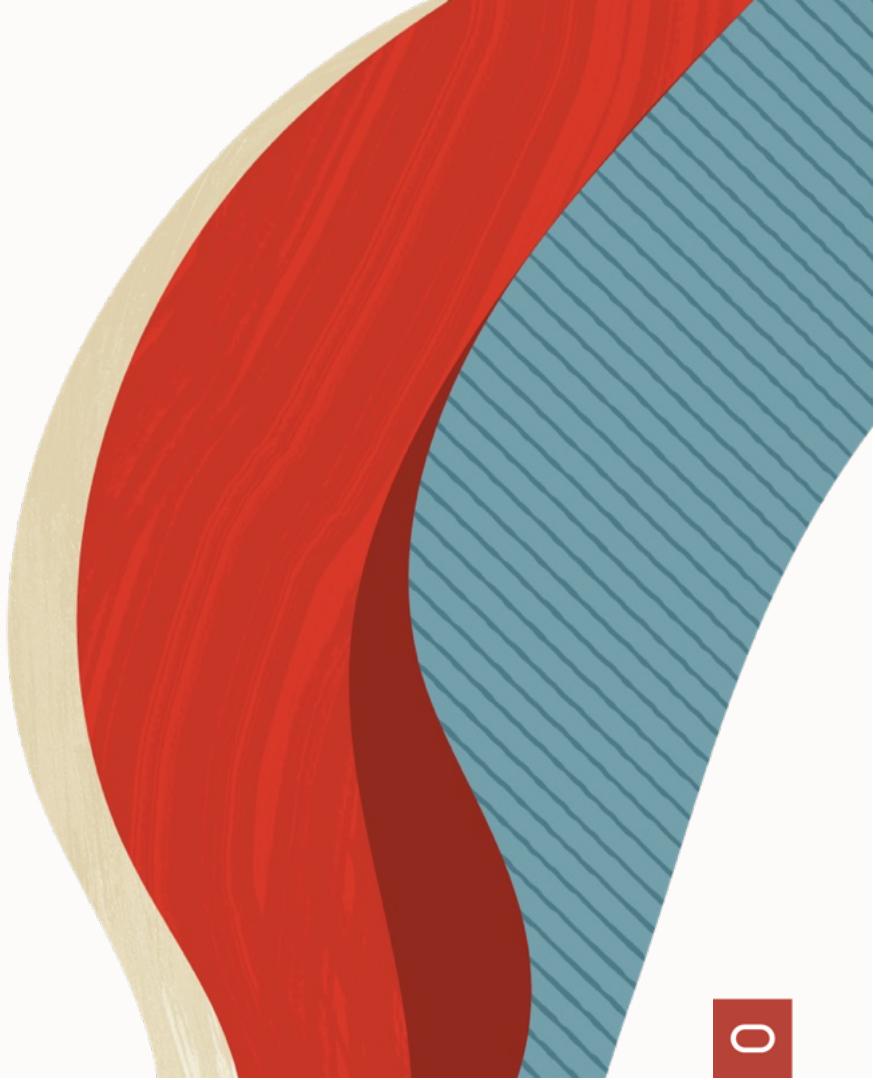
```
$ java -jar autoupgrade.jar -config do_it_all.cfg -patch -mode deploy
```

Demo

One-Button Patching

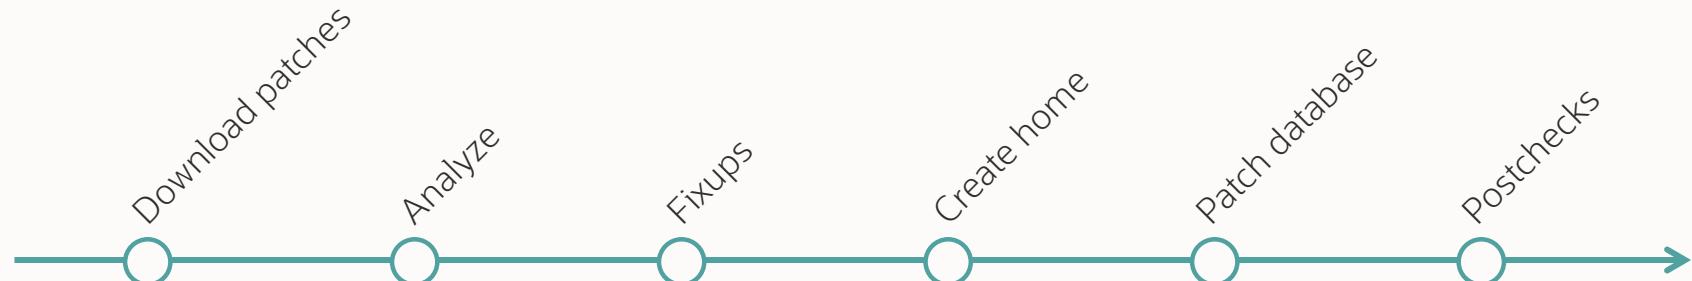
- Download patches
- Install Oracle home
- Patch database

Watch on [YouTube](#)





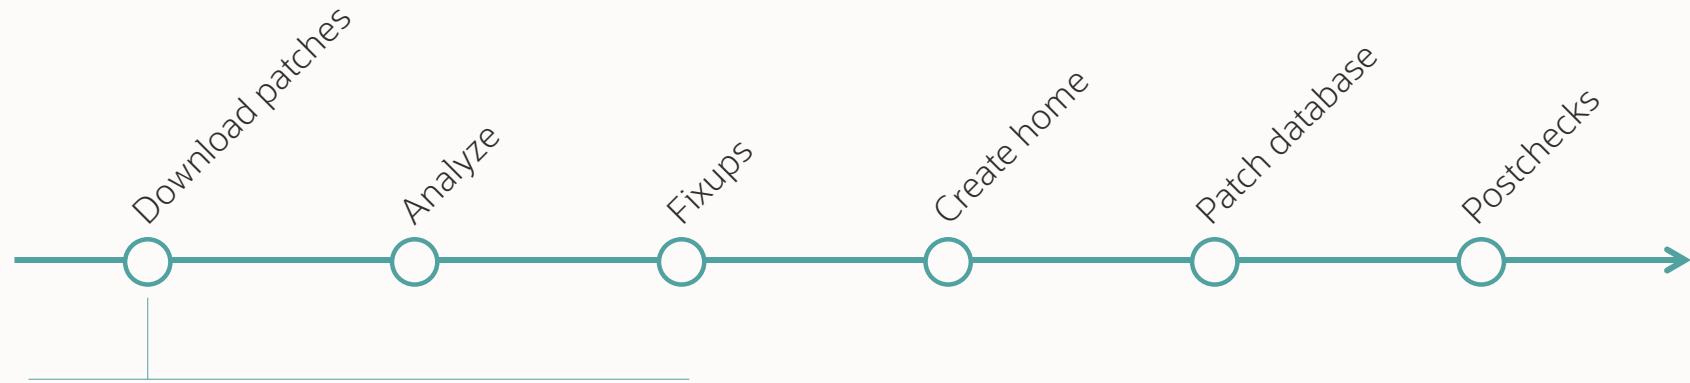
AutoUpgrade Patching



-mode deploy



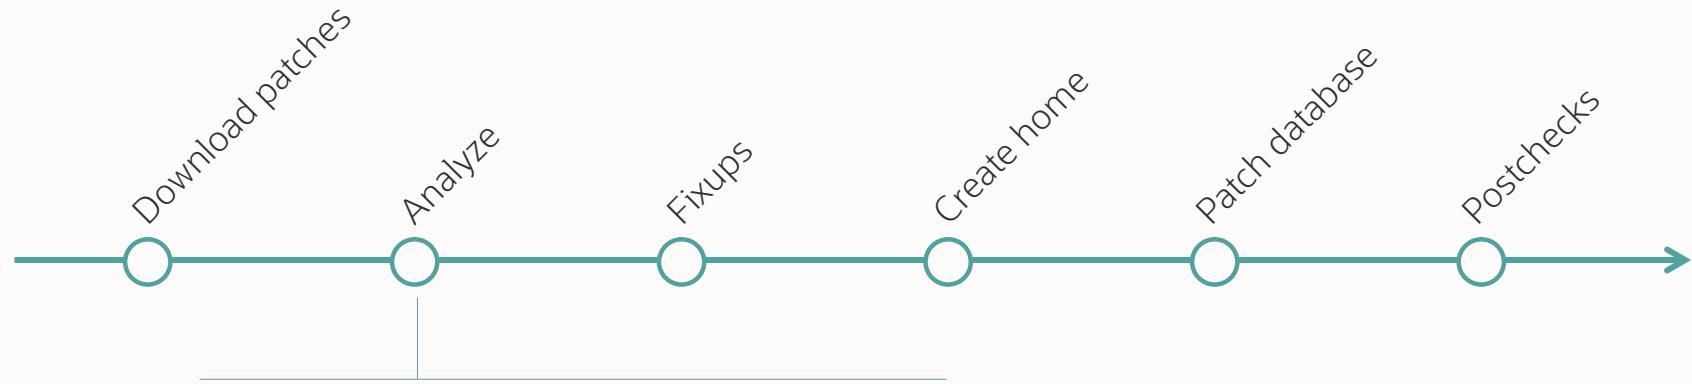
AutoUpgrade Patching



- Finds and downloads patches
- Uses your MOS credentials
- Chooses the right platform
- Stores patches in local repository



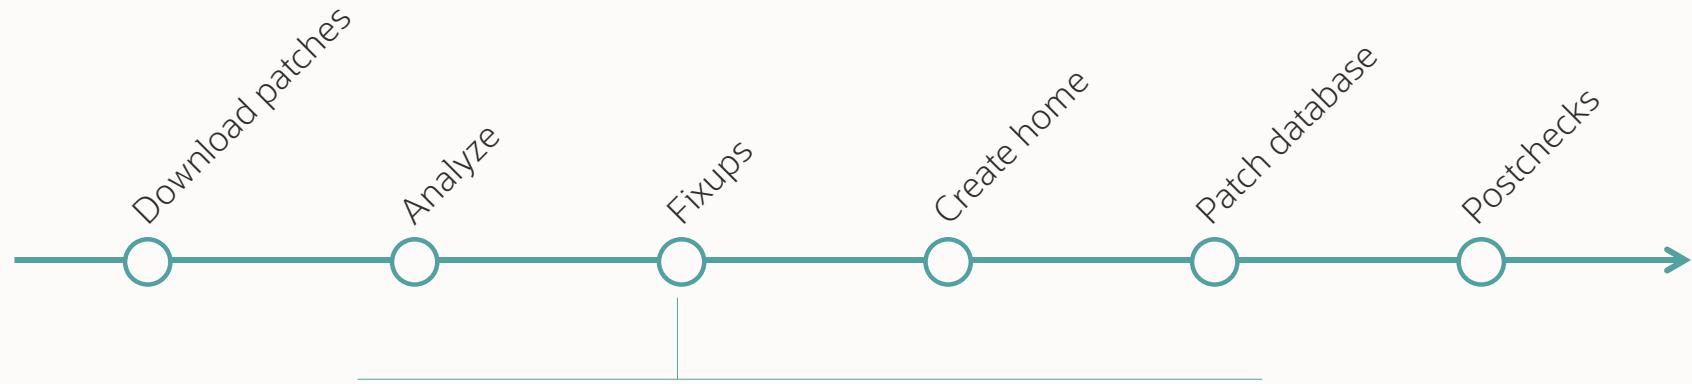
AutoUpgrade Patching



- Analyzes database for patch readiness
- Datapatch Sanity Checks
- Lightweight
- Non-intrusive
- Recommended, not required



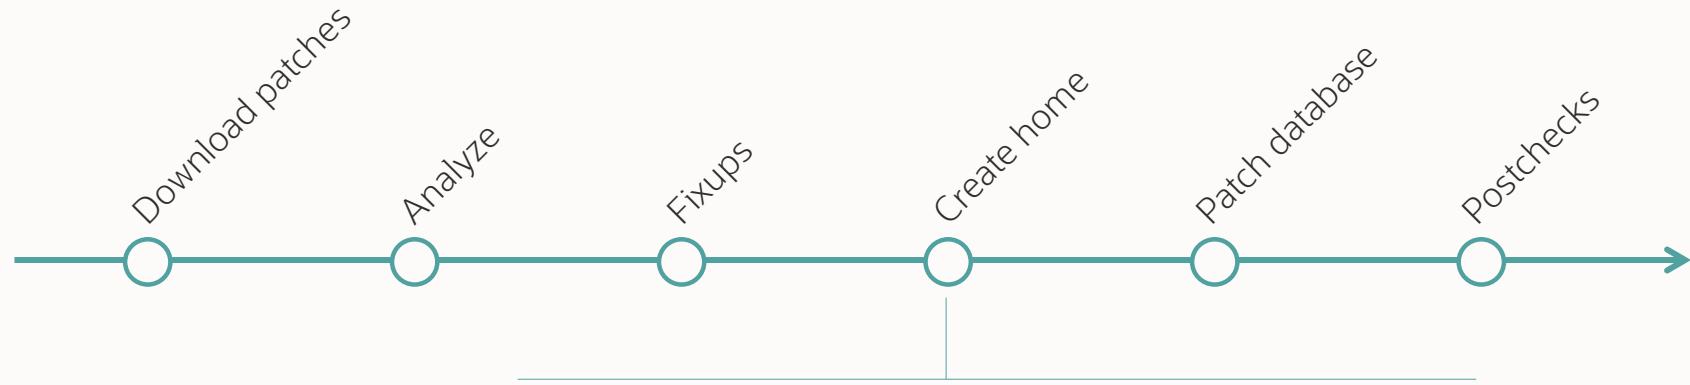
AutoUpgrade Patching



- Gathers dictionary statistics if needed
- Recompiles Oracle-maintained objects if needed
- Executes checks - see Doc ID [2380601.1](#)



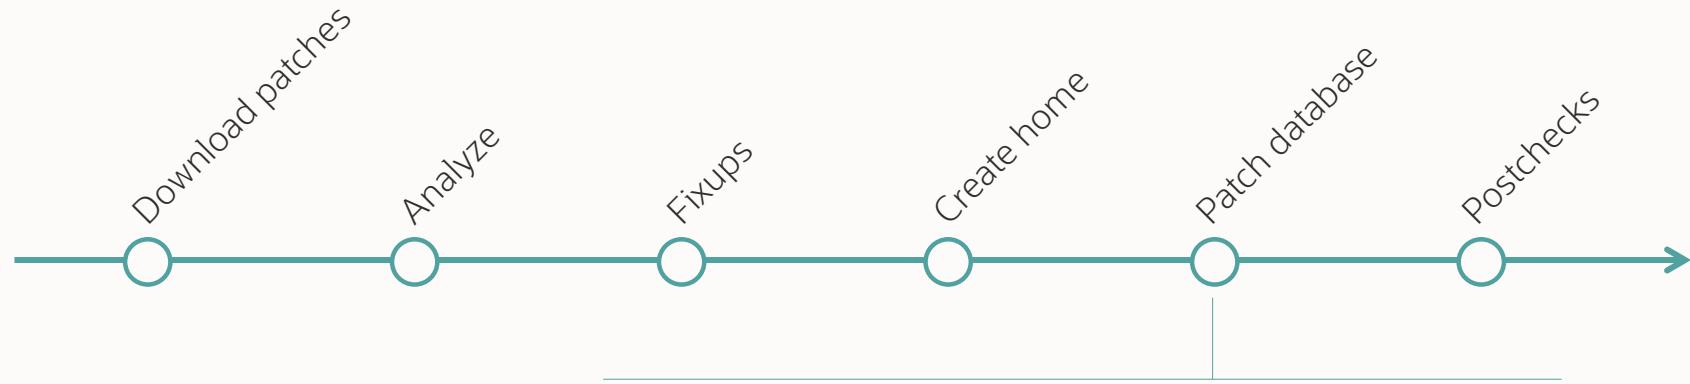
AutoUpgrade Patching



- Out-of-place patching
- Creates a brand-new Oracle home
- Uses *runInstaller* settings and binary options from source Oracle home
- Execute **root.sh** via sudo



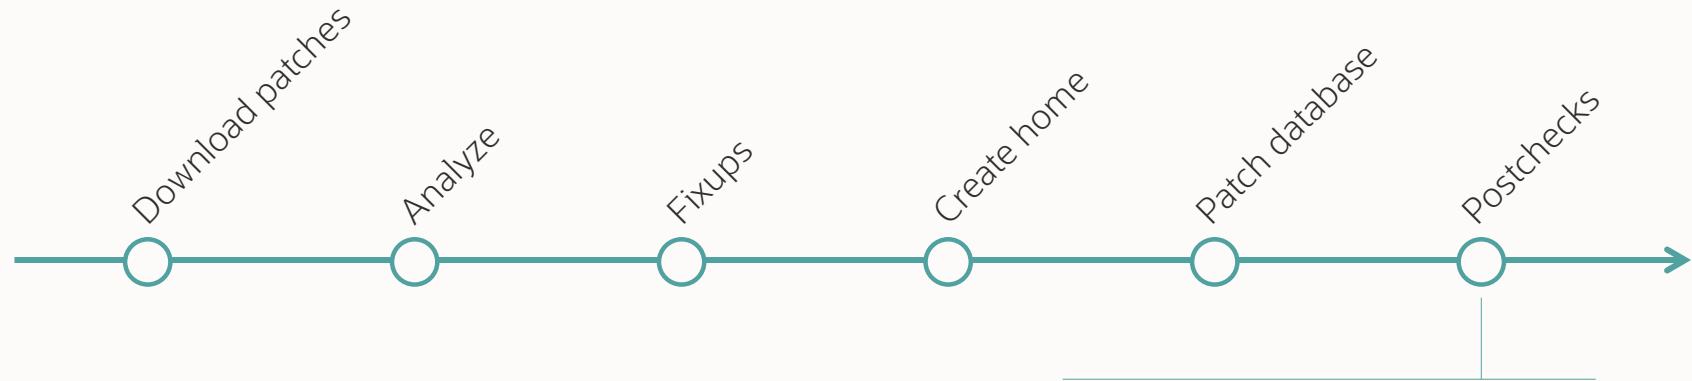
AutoUpgrade Patching



- Moves database instance to new Oracle home
- Moves configuration files
- Executes Datapatch
- Updates system directories (`utlfixdirs.sql`)
- Updates `/etc/oratab`
- Supports read-only Oracle home



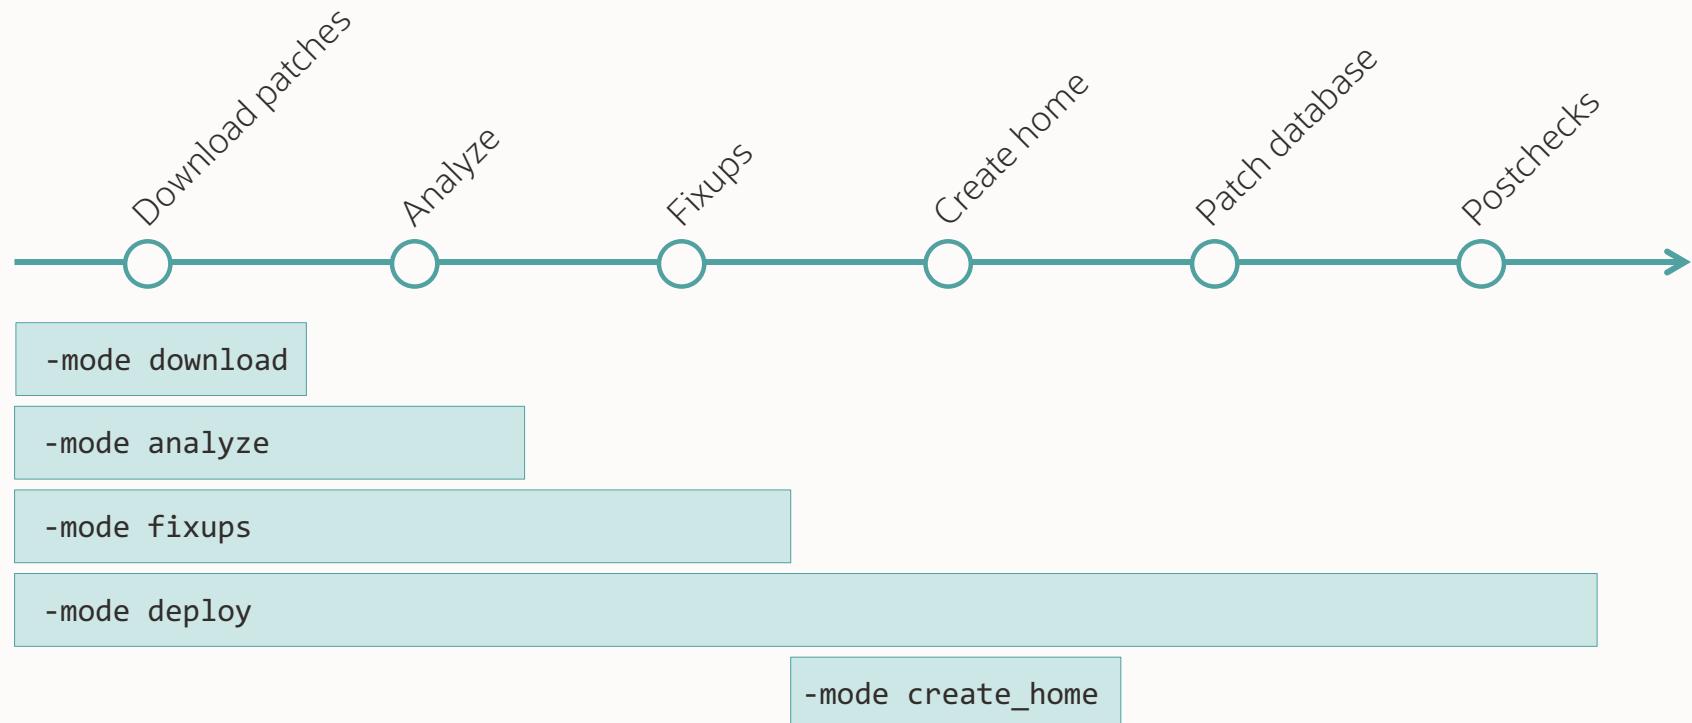
AutoUpgrade Patching



- Post-patching fixups
- Recompiles if needed
- Cleans up



Patching Modes





Requirements

Current Requirements



Single instance



Oracle Database 19c

Current Requirements



Single instance



Oracle Database 19c [or newer](#)



Current Requirements



All databases, including RAC and Data Guard



Oracle Database 19c or newer



Always download
the latest version of AutoUpgrade

- My Oracle Support [KB123450](#)

--Download the latest version of AutoUpgrade directly from oracle.com
--No authentication needed

```
wget https://download.oracle.com/otn-pub/otn_software/autoupgrade.jar
```

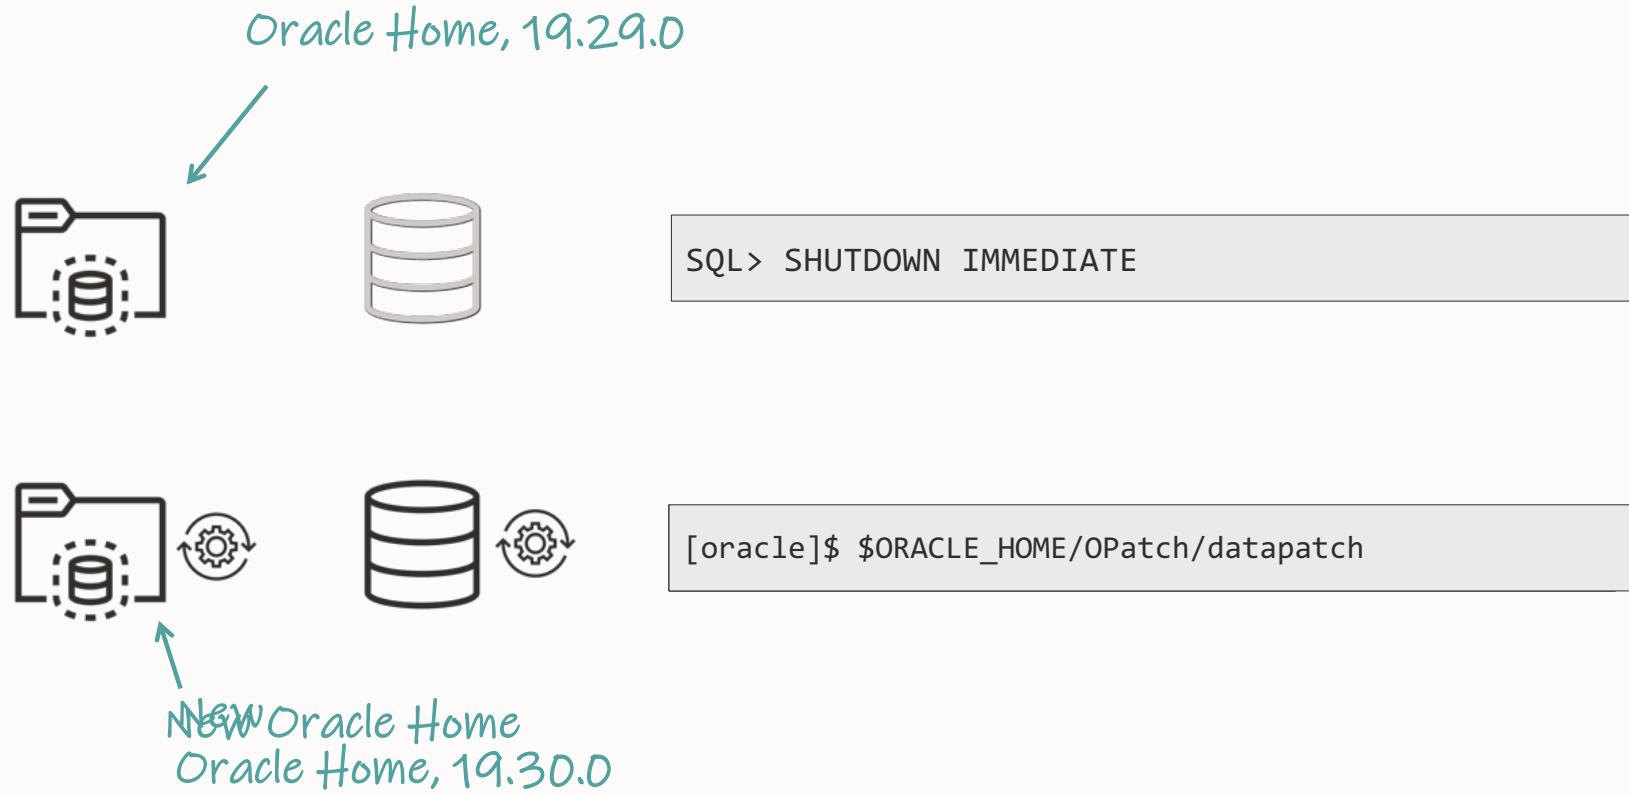


Patching

In-Place Patching



Out-of-Place Patching





Tim Hall  +∞ 
@oraclebase

...

When patching your production Oracle GI/DB installations, which method do you use?

In-Place = Current ORACLE_HOME

Out-Of-Place = New ORACLE_HOME

If you don't look after have production kit, then don't answer.

In-Place

55.4%

Out-Of-Place

44.6%



AutoUpgrade patches out-of-place



Which patches should you install?

```
global.keystore=/home/oracle/autoupgrade-patching/keystore
patch1.source_home=/u01/app/oracle/product/19/dbhome_19_29_0
patch1.target_home=/u01/app/oracle/product/19/dbhome_19_30_0
patch1.sid=DB19
patch1.folder=/home/oracle/autoupgrade-patching/patch
patch1.patch=RECOMMENDED
```

Recommended Patches

patch1.patch=RECOMMENDED

OPATCH

The latest OPatch

RU

The latest Release Update

OJVM

OJVM bundle matching Release Update

DPBP

Data Pump bundle patch matching RU



OJVM is embedded in Release Updates

- No separate download
- Complete RAC Rolling patching support



You can also add MRPs

- Get patch number from Primary Note for Database Quarterly Release Updates ([KB106822](#))

Primary Note for Database Quarterly Release Updates

KB106822

Last Updated

Dec 17, 2025

Service

Exadata

Service

Gen 2 Exadata Cloud at

Gen 2 Exadata Cloud at

Customer, Generation 1 -

Exadata Cloud at Customer

(First Generation Cloud

Machine), Oracle Cloud

Infrastructure - Exadata Cloud

Service, Oracle Database -

Enterprise Edition, Oracle

Database - Standard Edition



1.7

Authoring Instructions

This is a crossover article and must be edited only in Legacy MOS for Document ID : 888.1. Any changes made here will be lost in the next update.

Applies To

All Users

Summary

The purpose of this document is to list Database patches for both proactive and reactive maintenance.

This was [MOS Document ID: 888.1 in Legacy MOS](#). Post migration, we will be moving to KB888 (Date TBA)

Solution

1.0 Overview

2.0 Oracle Recommendation

3.0 Latest Release Updates

- 3.1 Oracle Database Long Term Releases

- 3.1.1 Database 19c

```
$ cat DB19.cfg
```

```
global.keystore=/home/oracle/autougrade-patching/keystore
patch1.source_home=/u01/app/oracle/product/19/dbhome_19_29_0
patch1.target_home=/u01/app/oracle/product/19/dbhome_19_30_0
patch1.sid=DB19
patch1.folder=/home/oracle/autougrade-patching/patch
patch1.patch=RECOMMENDED,38716952
```

```
$ cat DB19.cfg
```

```
global.keystore=/home/oracle/autougrade-patching/keystore
patch1.source_home=/u01/app/oracle/product/19/dbhome_19_29_0
patch1.target_home=/u01/app/oracle/product/19/dbhome_19_30_0
patch1.sid=DB19
patch1.folder=/home/oracle/autougrade-patching/patch
patch1.patch=RECOMMENDED,38716952
patch1.patch=RECOMMENDED,MRP
```



You can also add specific one-off fixes

- Oracle Database 19c and Oracle AI Database 26ai Important Recommended One-off Patches ([KB188772](#))

Oracle Database 19c and Oracle AI Database 26ai Important Recommended One- off Patches

KB188772

Last Updated
Yesterday 11:23 PM

Service
Oracle Database - Enterprise
Edition



1.9

Authoring Instructions

MOS Document ID : 555.1. Any changes made here will be lost in the next update.

Applies To

All Issues

Summary



Solution

NOTE: This document will be updated once a week on Tuesdays at 9AM EST if any new patch is identified.

Through our review of service requests, we often find that issues encountered are the result of customers being on an older Release Update (RU) or Monthly Recommended Patch (MRP). Many issues will have been fixed in the latest updates, which we always recommend.

As noted in the RU-specific tables below, fixes for known issues are targeted for inclusion in the next available MRP and RU. You can always find the latest RUs, MRPs, other patches, lists of fixed bugs and known issues in [KB106822](#) Identifying and installing the latest updates (patches) helps ensure you are using the most current content for security, functional, regression and bug fixes, as well as minor enhancements and any emergency one-offs.

Note:

Oracle encourages customers to install the latest MRP whenever possible. MRPs deliver the safest, fully-tested path to receive the fixes in this document. Refer to [FAQ2283](#) for further details.

In addition to the relevant patches listed below, you should apply patches based on the specific RU after reviewing the following My Oracle Support knowledge documents:

```
$ cat DB19.cfg
```

```
global.keystore=/home/oracle/autougrade-patching/keystore
patch1.source_home=/u01/app/oracle/product/19/dbhome_19_29_0
patch1.target_home=/u01/app/oracle/product/19/dbhome_19_30_0
patch1.sid=DB19
patch1.folder=/home/oracle/autougrade-patching/patch
patch1.patch=RECOMMENDED,37690446,34672698,34774667,29213893
```

```
$ cat DB19.cfg
```

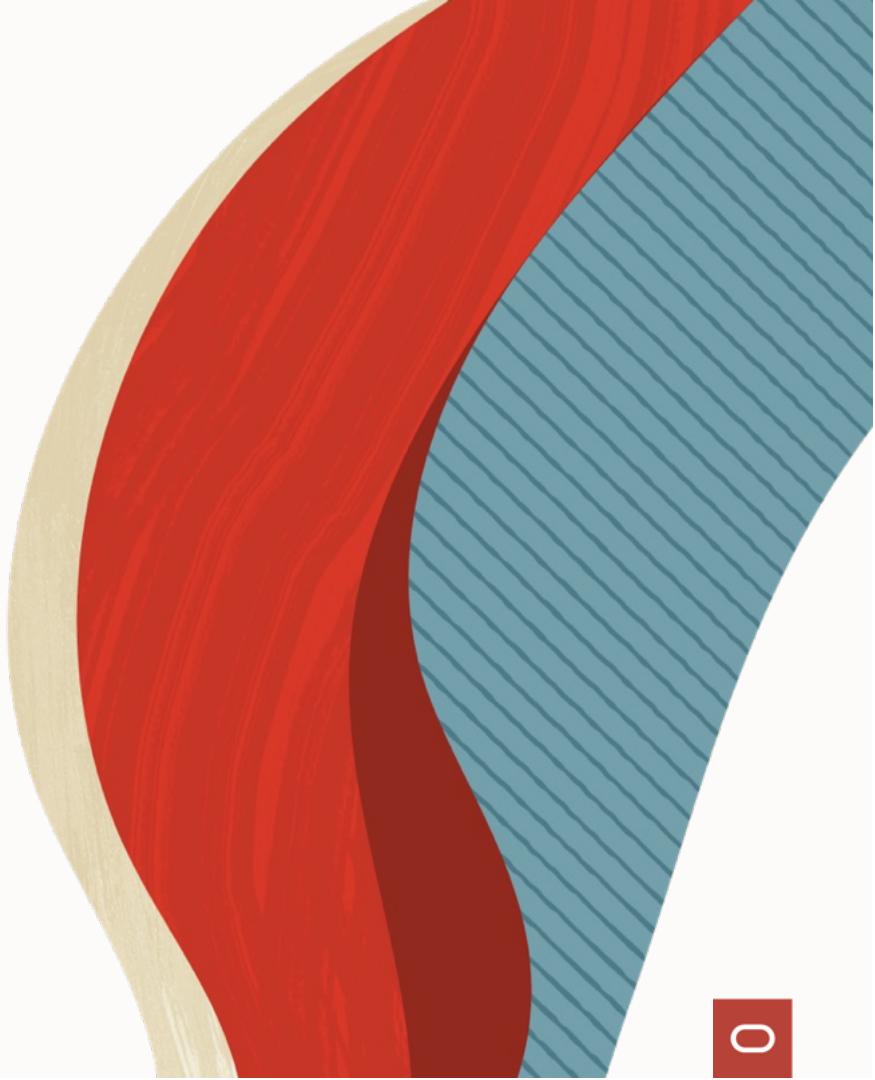
```
global.keystore=/home/oracle/autougrade-patching/keystore
patch1.source_home=/u01/app/oracle/product/19/dbhome_19_29_0
patch1.target_home=/u01/app/oracle/product/19/dbhome_19_30_0
patch1.sid=DB19
patch1.folder=/home/oracle/autougrade-patching/patch
patch1.patch=RECOMMENDED,OCW
```

Demo

Applying recommended patches

- Specifying patches
- Patch database

Watch on [YouTube](#)





How to download patches

Network Connectivity

The download stage requires:

1. Internet access
2. My Oracle Support credentials

Internet Access

AutoUpgrade connects to:

- `https://updates.oracle.com`
- `https://login-ext.identity.oraclecloud.com`
- `https://aru-akam.oracle.com`

URLs are part of a CDN, so expect changing IP addresses

- Use DNS names instead of IP addresses in your firewall

For connections via proxy use environment variables

- `https_proxy`

MOS Credentials

You must have:

- A valid My Oracle Support credential
- Connected to a Customer Support Identifier (CSI)
- Privilege to download patches using that CSI

Failure to meet the requirements:

- *Connection Failed - You entered an incorrect user name or password.*

MOS Credentials

For *download patches* privilege contact an MOS administrator in your organization.

- In *User Details*
- Select the appropriate *Support Identifier*
- Set the *Patches* drop-down to *Download*

Keystore

AutoUpgrade stores MOS credentials in a keystore

- Config file parameter: `global.keystore`
- Governs directory of AutoUpgrade keystore
- Password protected software keystore
- Optionally, an auto-open keystore
- No additional license needed

```
$ java -jar autoupgrade.jar ... -patch -load_password
```

```
$ java -jar autoupgrade.jar ... -patch -load_password
```

...

```
MOS> add -user <MOS username>
```

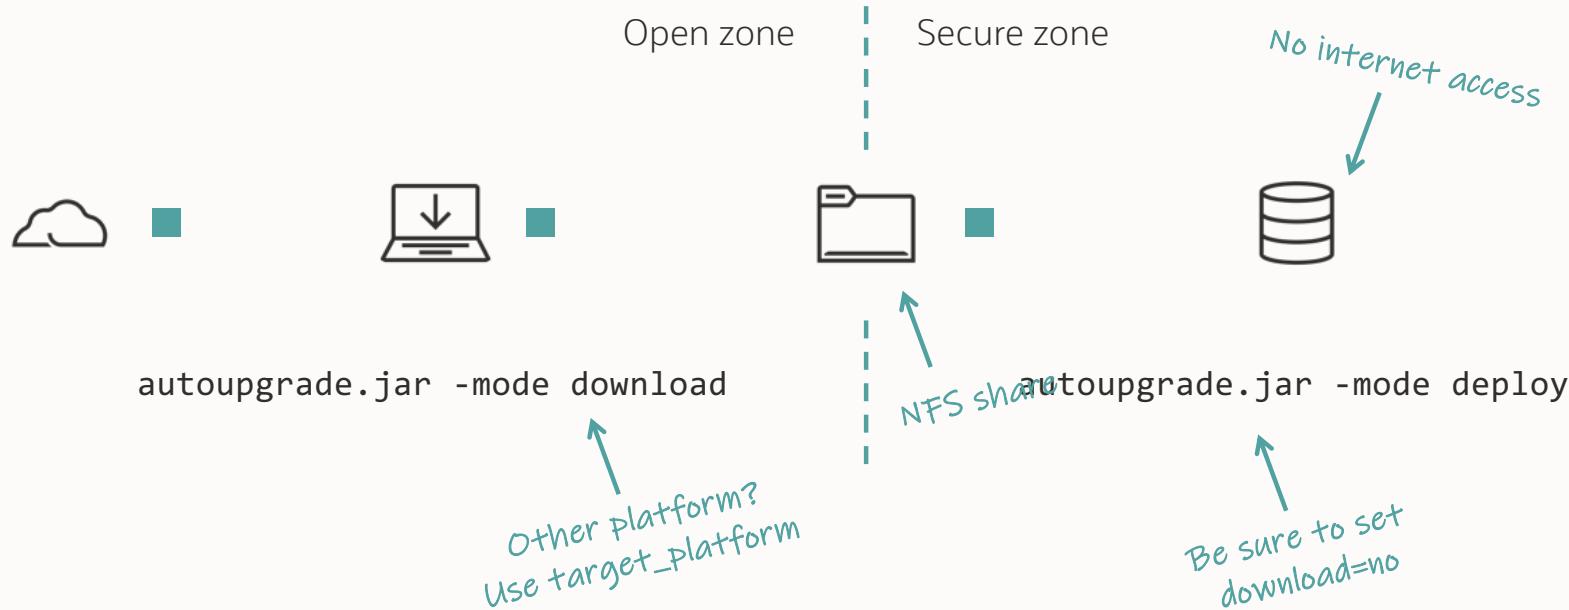
```
Enter your secret/Password:
```

```
Re-enter your secret/Password:
```



Your database host
doesn't have internet access?

Using Download Mode





You can download patches manually and place them in the *patch* folder

- Use a shared folder accessible to all hosts



In The Plans

We're working to support

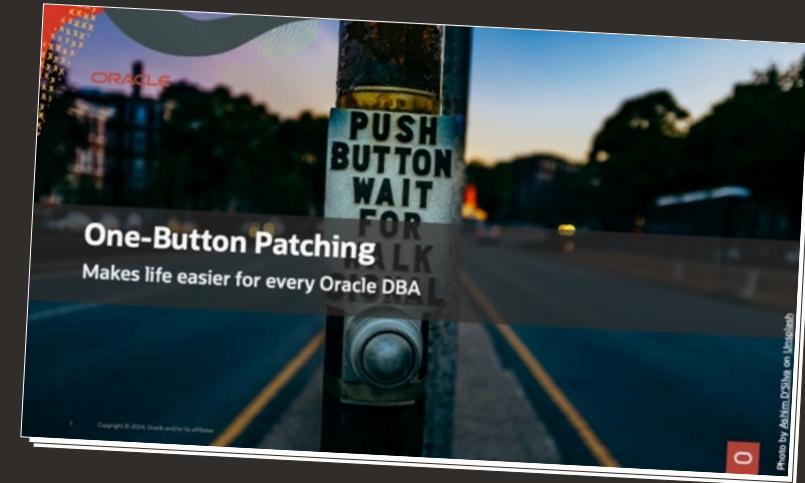


- Cloning Oracle homes
- Gold images

Webinar

One-Button Patching

Makes life easier for every Oracle DBA



Recording on [YouTube](#)
Get the [slides](#)

Hands-on Lab

Patch Me If You Can

It's better to fail in our lab, than in production



[Access lab on Oracle Live Labs](#)

Datapatch

What is Datapatch?

`$ORACLE_HOME/OPatch/datapatch`



`$ORACLE_HOME/sqlpatch/sqlpatch`



`$ORACLE_HOME/sqlpatch/sqlpatch.pl`

```
use strict;
use Getopt::Long;

use sqlpatch.pm
```



Datapatch uses **OPatch** to determine patches installed in Oracle home

- From within the database via **DBMS_QOPATCH**

ORA-20001: Latest xml inventory is not loaded into table

- Avoid the use of Queryable Inventory by using OPatch directly
- Safe for single instance databases. In RAC, it's vital you ensure
- all nodes are patched and the inventories on each node is similar.

./datapatch -noqi

./datapatch -local_inventory

Patch Apply Sequence

datapatch



1



Java patches

2



Release Update

3



One-off patches

Datapatch | Patch Rollback and Apply Queue

Binary Registry after opatch:

Patch 444 – Java Patch

Patch 555 – Release Update

Patch 666 – One-off Patch

Oracle Home

Rollback:

Apply:

Rollback:

SQL Registry before datapatch:

Patch 111 – Java Patch

Patch 222 – Release Update

Patch 333 – One-off Patch

Database

\$./datapatch

Cumulative:

Patch 222 to 555 – Release Update

Apply:



Datapatch uses
REGISTRY\$SQLPATCH_RU_INFO to
control the patching operations

Rollback Scripts



Apply/rollback scripts

```
$ORACLE_HOME/sqlpatch/.../nnn_apply.sql  
$ORACLE_HOME/sqlpatch/.../nnn_rollback.sql
```



Rollback scripts (zipped as BLOB):

One-off patches:

```
SELECT PATCH_DIRECTORY  
FROM   REGISTRY$SQLPATCH
```

Release Updates:

```
SELECT PATCH_DIRECTORY  
FROM   REGISTRY$SQLPATCH_RU_INFO
```

| PATCH_ID | RU_VERSION | RU_BUILD_TS | LOB_SIZE_MB |
|----------|-------------|---------------------|-------------|
| 29517242 | 19.3.0.0.0 | 04/10/2019 12:27:20 | 4 |
| 30125133 | 19.5.0.0.0 | 09/09/2019 18:05:49 | 18 |
| 30557433 | 19.6.0.0.0 | 12/17/2019 15:50:04 | 24 |
| 30869156 | 19.7.0.0.0 | 04/04/2020 03:50:18 | 30 |
| 31281355 | 19.8.0.0.0 | 07/03/2020 03:15:01 | 37 |
| 31771877 | 19.9.0.0.0 | 09/30/2020 18:32:49 | 43 |
| 32218454 | 19.10.0.0.0 | 01/08/2021 18:50:17 | 51 |
| 32545013 | 19.11.0.0.0 | 04/13/2021 00:40:09 | 59 |
| 32904851 | 19.12.0.0.0 | 07/16/2021 14:18:10 | 67 |
| 33192793 | 19.13.0.0.0 | 10/04/2021 16:50:50 | 75 |
| 33515361 | 19.14.0.0.0 | 12/25/2021 12:21:23 | 83 |
| 34133642 | 19.16.0.0.0 | 07/03/2022 02:22:23 | 99 |
| 34419443 | 19.17.0.0.0 | 09/24/2022 22:40:51 | 108 |
| 34765931 | 19.18.0.0.0 | 01/11/2023 17:17:38 | 116 |
| 35643107 | 19.21.0.0.0 | 09/30/2023 15:19:51 | 141 |

15 rows selected.

| CON_ID | LOB_SIZE_MB |
|--------|-------------|
|--------|-------------|

| CON_ID | LOB_SIZE_MB |
|--------|-------------|
|--------|-------------|

| | |
|---|-----|
| 1 | 953 |
| 2 | 953 |
| 3 | 953 |

| | |
|-----|-----|
| ... | ... |
|-----|-----|

| | |
|------|-----|
| 4098 | 953 |
|------|-----|

- Available since 19.28.0. Purges no longer needed zip files.
- Run in off-peak period after patching

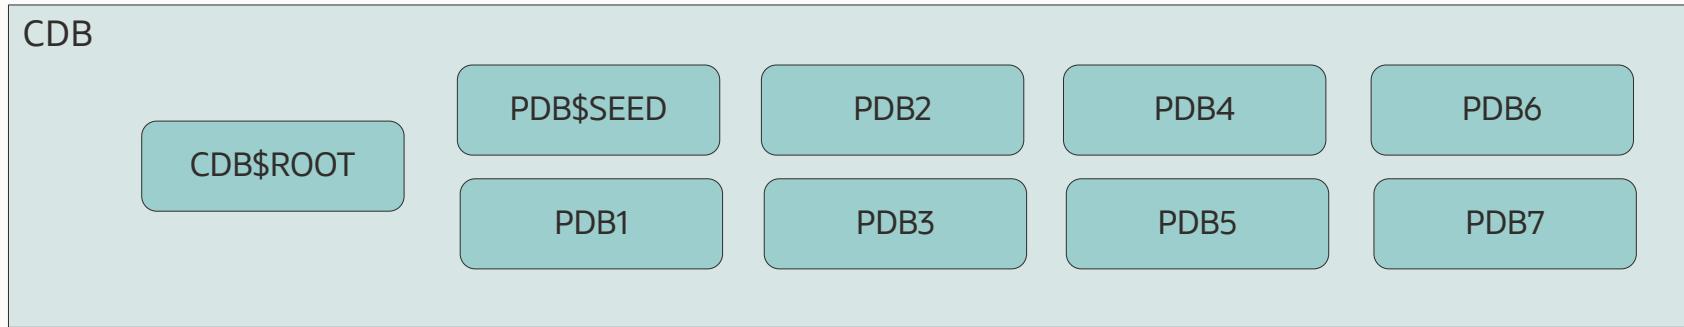
```
./datapatch -purge_old_metadata
```

Major retailer in
the US

Saved **20 TB/year** in total:

- *5000 DBs*
- *5 PDBs (3 PDBs + Root + Seed)*
- *200 Mb in each per patch*
- *4 patches annualy*

Multitenant



- Datapatch starts with CDB\$ROOT
- Datapatch sorts PDBs by *priority* and *con_id*
 - Set priority using `ALTER PLUGGABLE DATABASE ... PRIORITY`
- Datapatch determines parallel degree based on `CPU_COUNT`



Datapatch patches *PDB\$SEED* automatically

- New PDBs are ready to go
- No need to execute Datapatch on new PDBs
- Always set: `_exclude_seed_cdb_view=FALSE`



Datapatch only patches open PDBs

- READ WRITE, READ ONLY, or UPGRADE



Unpatched PDBs will open
in **RESTRICTED** mode

- Applies to failed or incomplete Datapatch runs
- Reported as plug-in violation

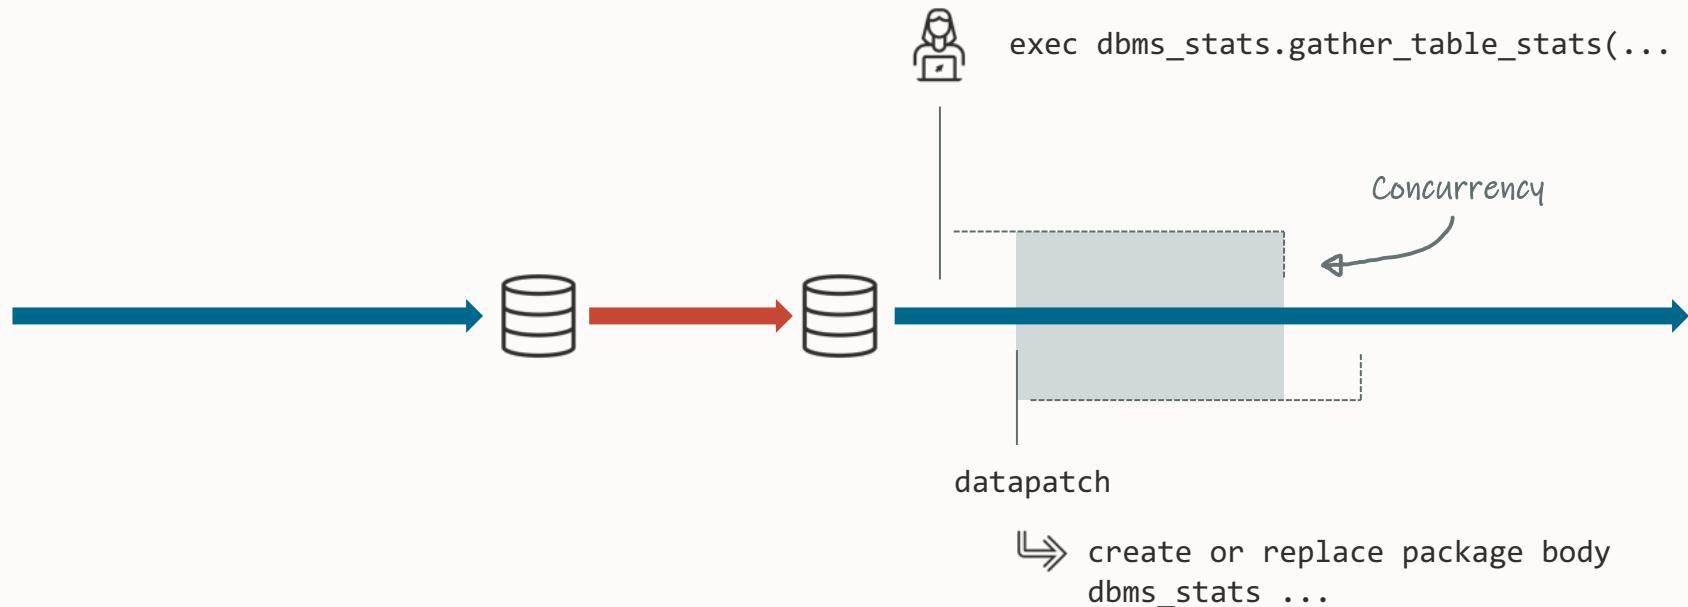
--Prevent PDBs from opening in restricted mode after a failed/incomplete
--Datapatch run. Quickly resolve the patching issue and re-run Datapatch.
--Use with caution.

```
alter system set "_pdb_datapatchViolation_restricted"=false;
```

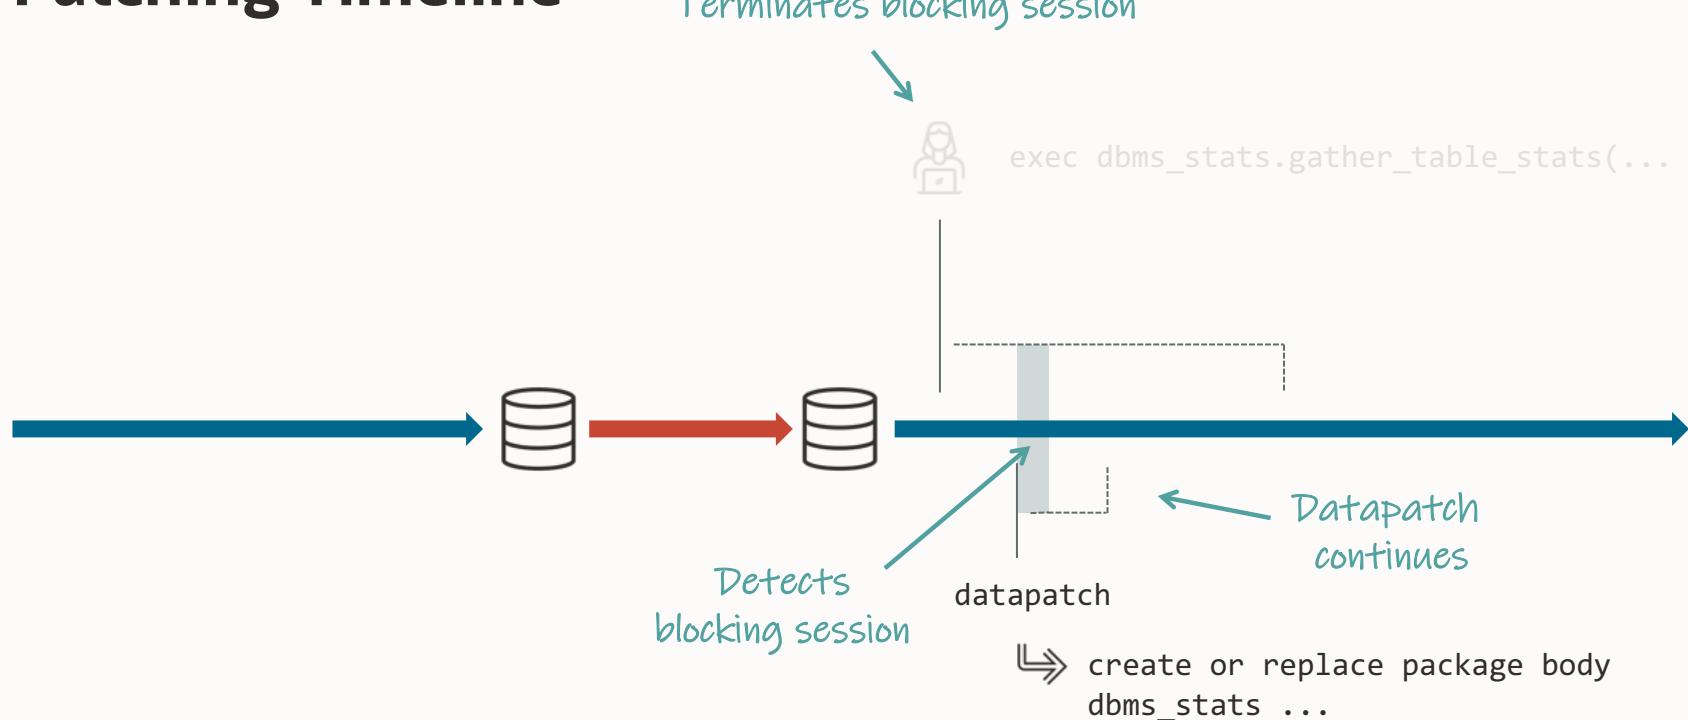


Datapatch now kills blocking sessions

Patching Timeline



Patching Timeline



```
SQL> exec dbms_stats.gather_table_stats(..
```

ERROR at line 1:

ORA-03113: end-of-file on communication channel

Process ID: 94

Session ID: 424 Serial number: 30685

Help: <https://docs.oracle.com/error-help/db/ora-03113/>

```
$ cd u01/app/oracle/diag/rdbms/cdb26/CDB26/trace
```

```
$ grep -i "Kill Final Blocker" *
```

```
CDB26_ora_50498.trc: Attempting to Kill Final Blocker blkr_sid:  
424 blkr_ser: 30685 session_S0: 0xa19df200 sess_ser: 30685
```



Datapatch kills **first blocker** within two minutes;
kills subsequent blockers immediately



Default in Oracle AI Database 26ai
and Oracle Database 19c (19.30)

```
-- Disable the automatic kill switch
```

```
./datapatch -force_terminate_blocking_sessions false
```

Concurrency Without Kill Switch

- Datapatch waits 15 min to acquire a lock
 - On timeout, ORA-04021 `timeout` occurred while waiting to lock object
- Optionally, [find blocking session](#) and kill it
- Increase timeout using `-ddl_lock_timeout <time-in-seconds>`

Datapatch Kill Switch



- Blog post: [A kill-switch terminates blocking sessions during patching](#)

Advanced Patching



Reduce downtime to the time it takes
to perform a switchover

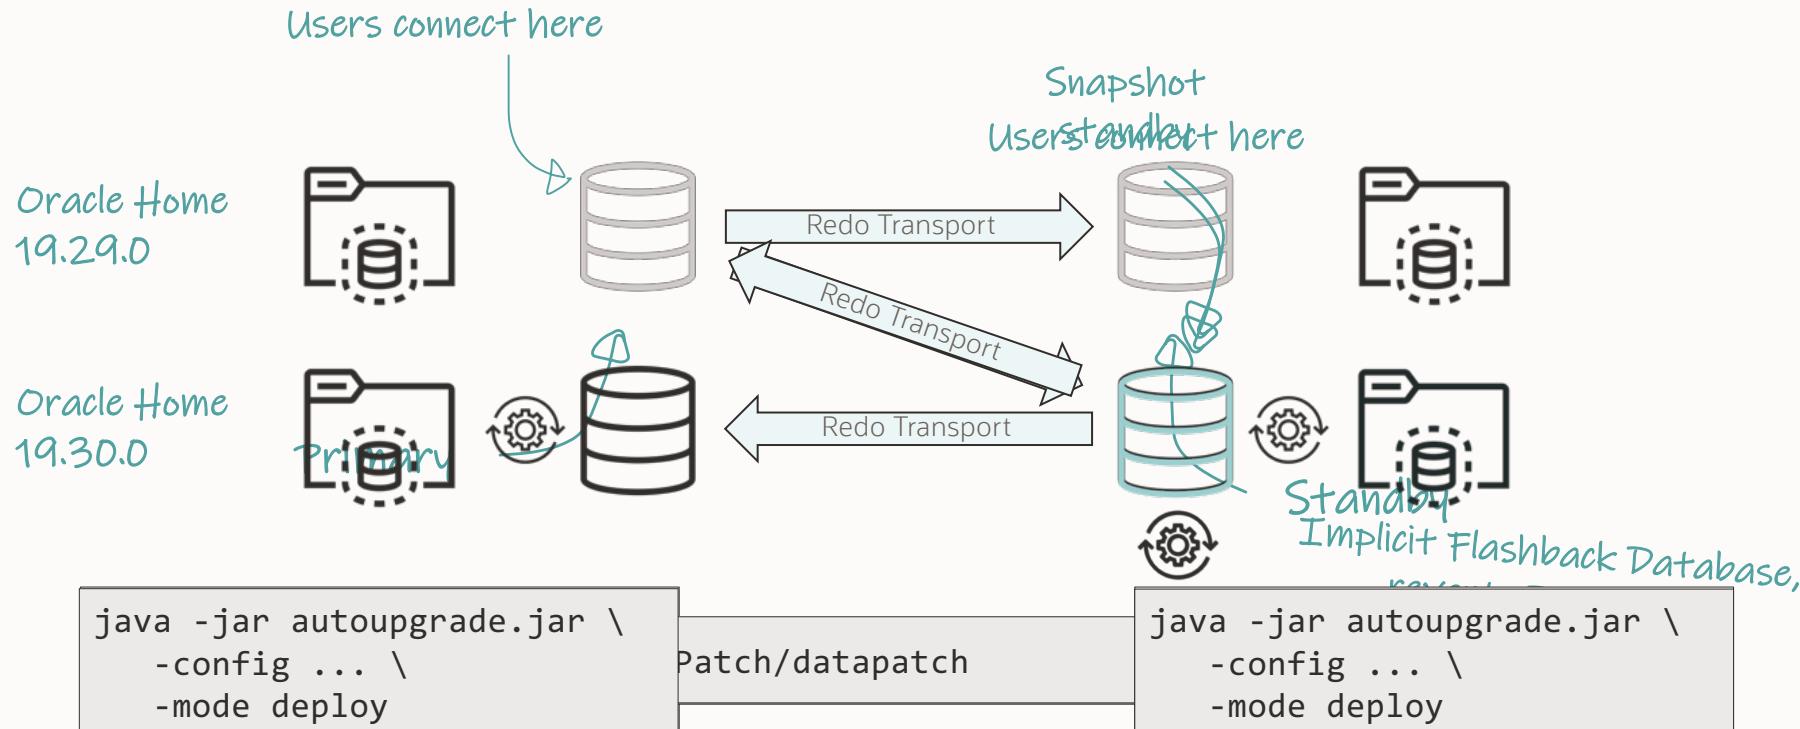
- [KB137118](#)
[Oracle Patch Assurance - Data Guard Standby-First Patch Apply](#)



Safely test and verify patches with Standby-First Patch Apply

- [KB137118](#)
[Oracle Patch Assurance - Data Guard Standby-First Patch Apply](#)

Standby-First Patching





Patch must Standby-First installable

- Release Updates (RU) are always standby-first installable
- Always check the patches' readme



Execute Datapatch on the primary database

- Only execute Datapatch when all homes are on the new patch



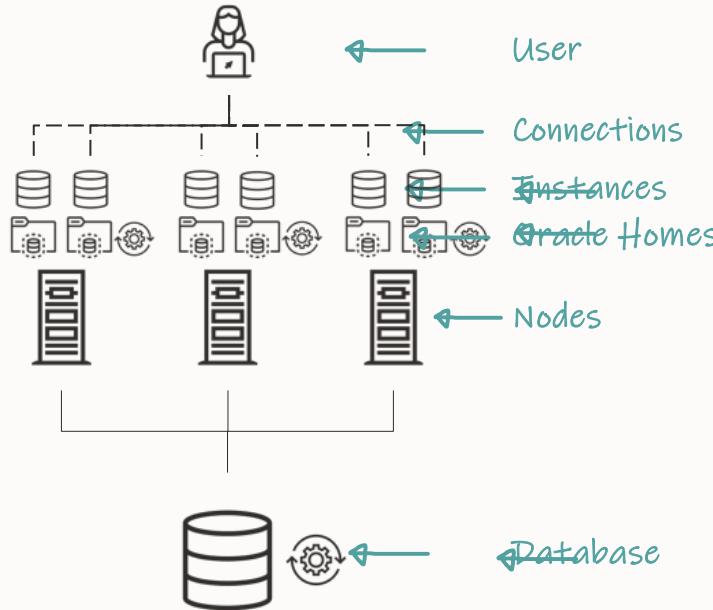
Use AutoUpgrade to patch
your standby database

- Check [blog posts](#) for details



Avoid database downtime with
RAC Rolling Patch Apply

RAC Rolling Patching



- New Oracle Home
- Patch Oracle Home
- Move to new Oracle Home
- Execute Datapatch



Keep GI and DB patch levels in sync

- This is what we test and run in our Cloud

GI and DB Patch Levels

| Method | Status | Example |
|---|-------------------------------------|--|
| In sync | Recommended, best practice | GI 19.30 + DB 19.30 |
| Within two Release Updates | Supported, good practice | GI 19.30 + DB 19.28 GI 19.28 + DB 19.30 |
| Within three or more Release Updates | Supported, not recommended | GI 19.30 + DB 19.27 GI 19.27 + DB 19.30 |
| Different releases | Supported, use only for upgrades | GI 26.1 + DB 19.30 |



Complete a rolling patching operation
always as **quickly as possible**

- [FAQ7365](#)
[RAC: Frequently Asked Questions](#)



Unusual combinations are supported, but
we **strongly** advice against it

- Node 1 with GI 19.30.0, node 2 with GI 19.29.0
- Patching node 1 on Monday, node 2 on Tuesday ...
- Mix of GI and DB versions on various nodes

Release updates are **always**:



Standby-First installable



RAC Rolling installable



Key Learnings



- 1 Patch your Oracle Database regularly
- 2 Patch out-of-place
- 3 Use AutoUpgrade Patching

Break

We start again at 11:15

AutoUpgrade

Overview and new features





AutoUpgrade is **the only supported tool to upgrade** your Oracle Database

- DBUA is **desupported**

```
$ java -jar autoupgrade.jar -version
```

```
build.version 26.1.260109
```

```
build.date 2026/01/09 12:19:19 -0500
```

```
build.supported_target_versions 12.2,18,19,21,23
```

```
build.type production
```

```
build.MOS_NOTE 2485457.1
```

```
build.MOS_LINK https://support.oracle.com/.../?id=2485457.1
```

- Download the latest version of AutoUpgrade directly from oracle.com
- No authentication needed

```
wget https://download.oracle.com/otn-pub/otn_software/autoupgrade.jar
```

New Features

--Upgrade RMAN catalog after upgrade

--<https://dohdatabase.com/autoupgrade-new-features-upgrade-rman-catalog-schema>

```
$ cat CDB1.cfg
```

```
upg1.source_home=/u01/app/oracle/product/19/dbhome_19_30
```

```
upg1.target_home=/u01/app/oracle/product/26/dbhome_26_1
```

```
upg1.sid=CDB1
```

```
upg1.rman_catalog_connect_string=catalogdb
```

--Update OEM configuration after upgrade

-- <https://dohdatabase.com/autougrade-new-features-update-enterprise-manager-configuration>

```
$ cat CDB1.cfg
```

```
upg1.source_home=/u01/app/oracle/product/19/dbhome_19_30
```

```
upg1.target_home=/u01/app/oracle/product/26/dbhome_26_1
```

```
upg1.sid=CDB1
```

```
upg1.upg1.emcli_path=/u01/app/oracle/oem
```

```
upg1.em_target_name=ORCL_myhost.domain.int
```



Upgrading and patching on Windows

Services

E... Act... Vi... H...

| Services (Local) | Name | Description | Status | Startup Type | Log On As |
|------------------|--|--------------------|----------|---------------|-----------------|
| | OpenSSH Authentication Agent | Agent to handle... | Disabled | Local System | |
| | OpenSSH SSH Server | SSH protocol... | Disabled | Local System | |
| | Optimize drives | Helps the c... | Manual | Local System | |
| | Oracle Balloon Service | | Running | Automatic | Local System |
| | Oracle Cloud Agent | Oracle Clo... | Running | Automatic | NT SERVICE\O... |
| | Oracle Cloud Agent Cloud Guard Work... | Oracle Clo... | Running | Manual | NT SERVICE\O... |
| | Oracle Cloud Agent Run Command Serv... | Oracle Clo... | Running | Manual | NT SERVICE\O... |
| | Oracle Cloud Agent Updater | Oracle Clo... | Running | Automatic | NT SERVICE\O... |
| | Oracle Cloud Block Storage Manageme... | Oracle Clo... | | Manual | NT SERVICE\O... |
| | Oracle Cloud Operating System Manage... | Oracle Clo... | | Manual | NT SERVICE\O... |
| | Oracle Cloud Operating System Manage... | Oracle Clo... | | Manual | NT SERVICE\O... |
| | Oracle Cloud Unified Monitoring Install... | Oracle Clo... | Running | Manual | NT SERVICE\O... |
| | Oracle Cloud Vulnerability Scanning Ser... | Oracle Clo... | | Manual | NT SERVICE\O... |
| | Oracle Unified Agent | Oracle Unif... | Running | Manual | Local System |
| | Oracle VioGpu Resolution Service | | Running | Automatic | Local System |
| | Oracle VirtIO Service | Oracle Virtl... | Running | Automatic | Local System |
| | OracleJobSchedulerDB19 | | | Disabled | .\OraDb |
| | OracleRemExecServiceV2 | | | Manual | Local System |
| | OracleServiceDB19 | | | Automatic | .\OraDb |
| | OracleVssWriterDB19 | | | Automatic | .\OraDb |
| | Payments and NFC/SE Manager | Manages p... | Disabled | Local Service | |
| | Performance Counter DLL Host | Enables re... | Manual | Local Service | |
| | Performance Logs & Alerts | Performan... | Manual | Local Service | |
| | Phone Service | Manages t... | Disabled | Local Service | |

Extended Standard

```
$ cat patch_on_win.cfg
```

```
global.keystore=c:\oracle\autoupgrade\keystore
patch1.source_home=c:\oracle\product\dbhome_19_29_0
patch1.target_home=c:\oracle\product\dbhome_19_30_0
patch1.sid=DB19
patch1.folder=c:\oracle\patches
patch1.patch=RECOMMENDED
patch1.wincredential=c:\oracle\autoupgrade\credential
```

```
C:\> java -jar autoupgrade.jar  
-config patch_on_win.cfg  
-patch  
-loadwincredential "DB19"
```

```
C:\> java -jar autoupgrade.jar  
-config patch_on_win.cfg  
-patch  
-loadwincredential "patch1"
```

```
C:\> java -jar autoupgrade.jar  
-config patch_on_win.cfg  
-patch  
-mode deploy
```

Windows



--AutoUpgrade deletes the credential file after use for security reasons
--For automation it is desirable to keep the credential file

patch1.delete_credential_file=false



AutoUpgrade fully supports upgrading
Oracle RAC Database on Windows

Refreshable Clone PDBs

Non-CDB to PDB conversion is irreversible

What are your rollback options?

ROLLBACK



Backup / restore

Ensure you have a recent backup and requires time to restore and recover

Copy data files

Requires time and disk space to hold a copy of the data files

ROLLBACK



Backup / restore

Ensure you have a recent backup and requires time to restore and recover

Copy data files

Requires time and disk space to hold a copy of the data files

Refreshable clone PDB

Requires **time and** disk space to hold a copy of the data files

Requires Oracle Database 12.2 or newer

Refreshable Clone PDB



CREATE

Create PDB from non-CDB over a database link



REFRESH

Apply redo from non-CDB to keep PDB up-to-date



OUTAGE

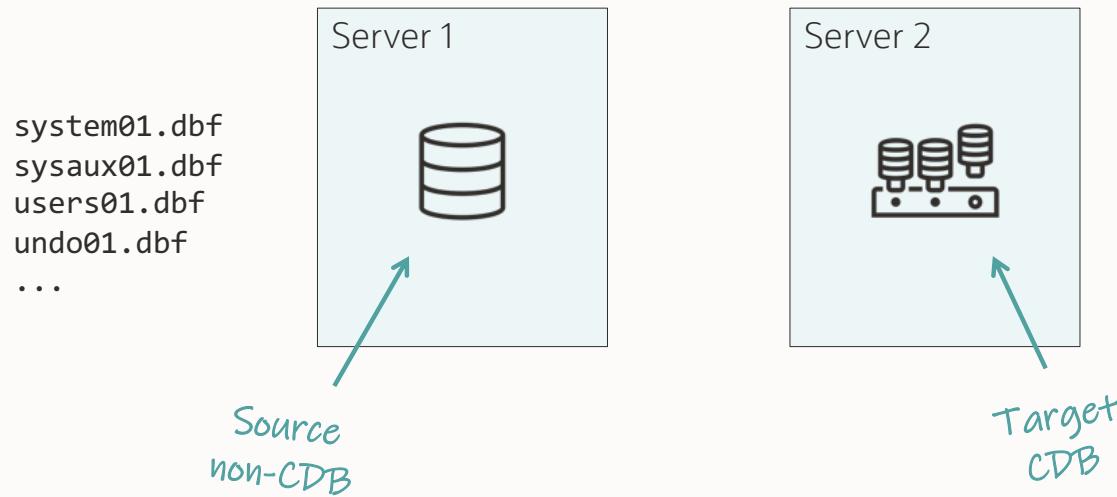
Disconnect users and refresh PDB for the last time



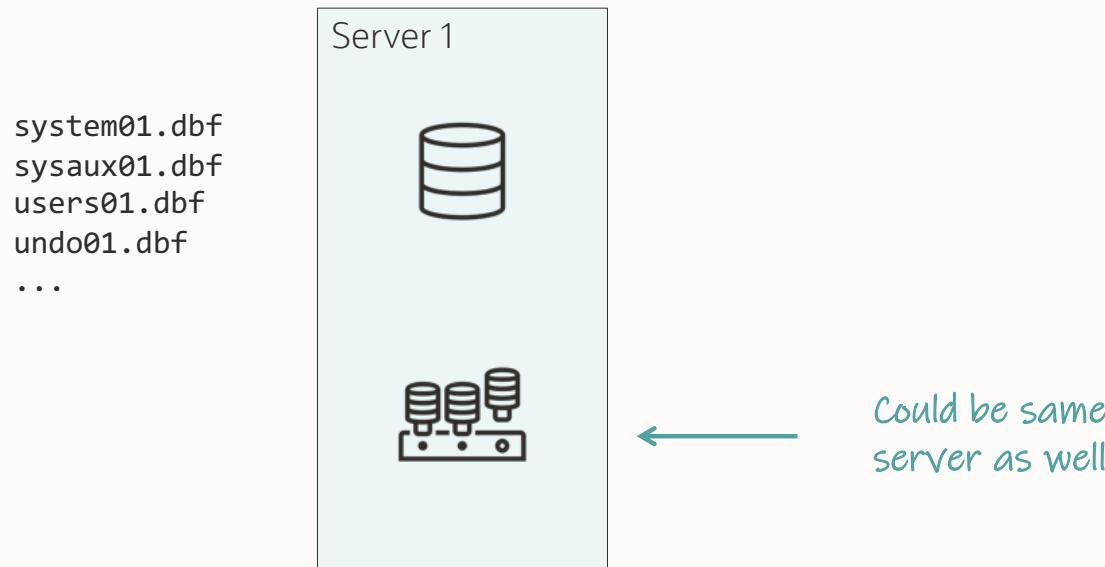
CONVERT

To become a proper PDB, it must be converted

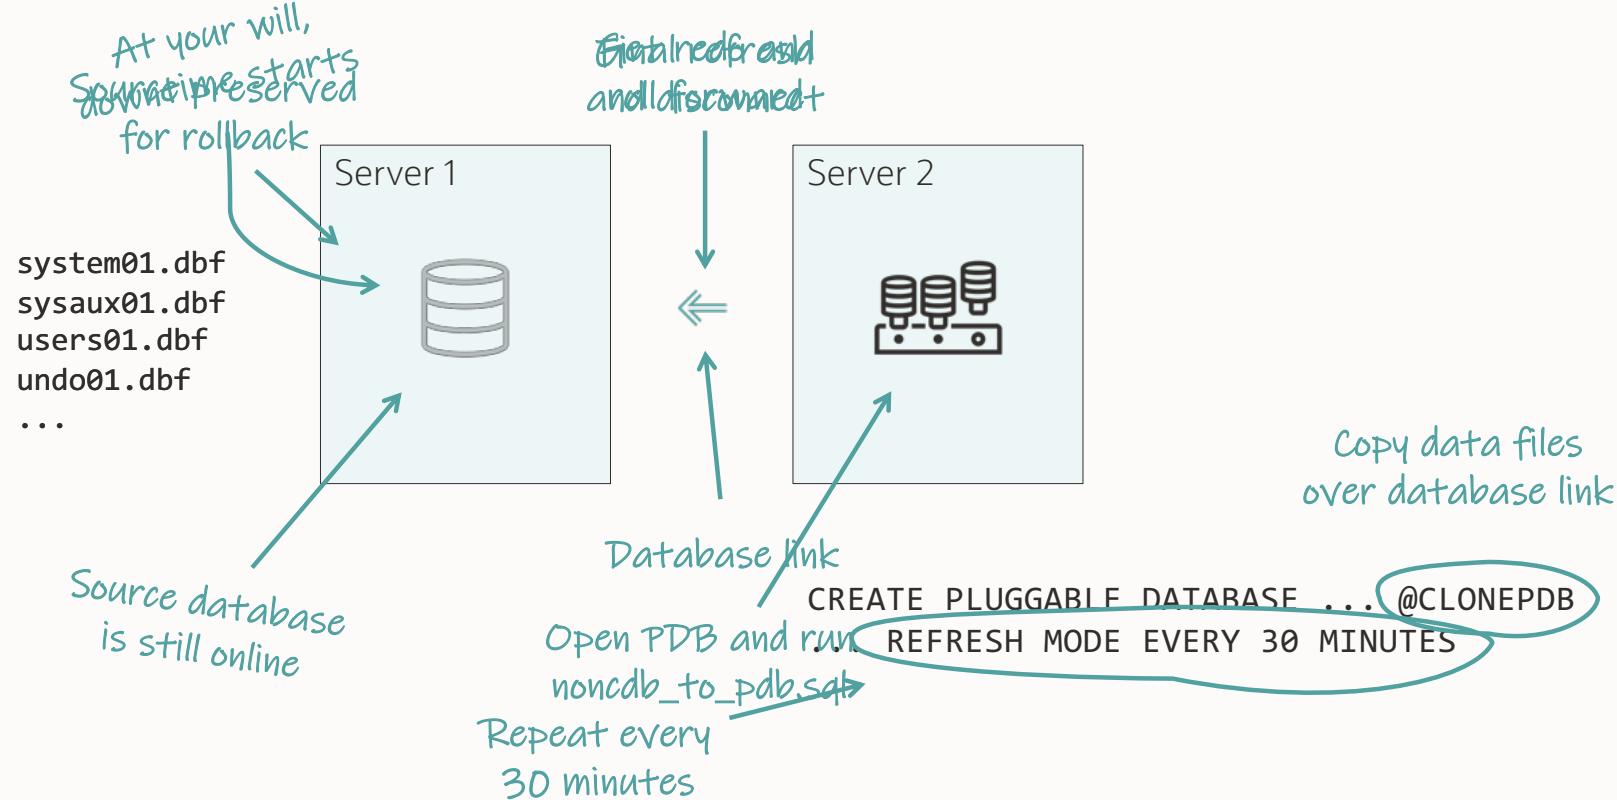
Refreshable Clone PDB



Refreshable Clone PDB



Refreshable Clone



Refreshable Clone PDB

Source non-CDB



```
CREATE USER dblinkuser  
    IDENTIFIED BY ... ;  
  
GRANT CREATE SESSION,  
    CREATE PLUGGABLE DATABASE,  
    SELECT_CATALOG_ROLE TO dblinkuser;  
  
GRANT READ ON sys.ENC$ TO dblinkuser;
```

Target CDB



```
CREATE DATABASE LINK CLONEPDB  
    CONNECT TO dblinkuser  
    IDENTIFIED BY ...  
    USING 'noncdb-alias';
```

Refreshable Clone

Source non-CDB

Target CDB



```
upg1.source_home=/u01/app/oracle/product/dbhome_19_30
upg1.target_home=/u01/app/oracle/product/dbhome_26_1
upg1.sid=NONCDB1
upg1.target_cdb=CDB1
upg1.target_version=23
upg1.source_dblink.NONCDB1=CLONEPDB
upg1.target_pdb_name.NONCDB1=PDB1
```



You can drop user and database link
after migration

- Config file parameter `drop_dblink`

Refreshable Clone PDB

Source non-CDB

Target CDB



```
upg1.source_home=/u01/app/oracle/product/dbhome_19_30
upg1.target_home=/u01/app/oracle/product/dbhome_26_1
upg1.sid=NONCDB1
upg1.target_cdb=CDB1
upg1.target_version=23
upg1.source_dblink.NONCDB1=CLONEPDB 300
upg1.target_pdb_name.NONCDB1=PDB1
```

Refreshable Clone PDB

Source non-CDB

Target CDB



```
upg1.source_home=/u01/app/oracle/product/dbhome_19_30
upg1.target_home=/u01/app/oracle/product/dbhome_26_1
upg1.sid=NONCDB1
upg1.target_cdb=CDB1
upg1.target_version=23
upg1.source_dblink.NONCDB1=CLONEPDB 300
upg1.target_pdb_name.NONCDB1=PDB1
```



Rename your PDB to avoid name collision

- If CDB is on same host,
it also registers for the default service

Refreshable Clone

Source non-CDB

Target CDB



```
upg1.source_home=/u01/app/oracle/product/dbhome_19_30
upg1.target_home=/u01/app/oracle/product/dbhome_26_1
upg1.sid=NONCDB1
upg1.target_cdb=CDB1
upg1.target_version=23
upg1.source_dblink.NONCDB1=CLONEPDB 300
upg1.target_pdb_name.NONCDB1=PDB1
upg1.start_time=19/01/2038 03:14:07
```

Refreshable Clone PDB

Source non-CDB

Target CDB



```
upg1.source_home=/u01/app/oracle/product/dbhome_19_30
upg1.target_home=/u01/app/oracle/product/dbhome_26_1
upg1.sid=NONCDB1
upg1.target_cdb=CDB1
upg1.target_version=23
upg1.source_dblink.NONCDB1=CLONEPDB 300
upg1.target_pdb_name.NONCDB1=PDB1
upg1.start_time=19/01/2038 03:14:07
upg1.parallel_pdb_creation_clause=4
```



Refreshable Clone

1

Run on source

```
autoupgrade.jar ... -mode analyze  
autoupgrade.jar ... -mode fixups
```

2

Run on target

```
autoupgrade.jar ... -mode deploy
```



Refreshable Clone PDB



autoupgrade.jar ... -mode deploy

upg1.start_time=19/01/2038 03:14:07

--When a job is in REFRESHPDB stage,
--you can force it to start immediately
upg> **proceed -job 101**

--When a job is in REFRESHPDB stage,
--you can force it to start immediately
upg> proceed -job 101

--Or postpone it
upg> proceed -job 101 -newstarttime +2h30m

--When a job is in REFRESHPDB stage,
--you can force it to start immediately
upg> proceed -job 101

--Or postpone it
upg> proceed -job 101 -newstarttime +30m

--Or reschedule it
upg> proceed -job 101 -newstarttime 07/02/2026 02:30:00



Works for unplug-plug upgrades as well



The source non-CDB stays intact
to allow rollback



**Techniker
Krankenkasse**

Move hundreds of Oracle Database 19c PDBs to
Oracle Database 23ai with ExaScale-only

Customer Case | Techniker Krankenkasse

| Customer |
|-------------|
| Project |
| Constraints |
| Preparation |
| Upgrade |
| Success? |
| Remarks |

Large German Public Health Insurance

- **Founded:** 1884
- **Legal Form:** Public Corporation
- **Headquarters:** Hamburg, Germany
- **Members:** 11.9 million
- **Employees:** aprx. 19,000
- **Customer Service:** 174 service centers
- **Budget 2024:** €62.5 billion

Customer Case | Techniker Krankenkasse

| |
|----------------|
| Customer |
| Project |
| Constraints |
| Preparation |
| Upgrade |
| Success? |
| Remarks |

Move to Oracle Database 23ai with ExaScale

- No ASM, no Oracle Database 19c anymore



Customer Case | Techniker Krankenkasse

Customer

Project

Constraints

Preparation

Upgrade

Success?

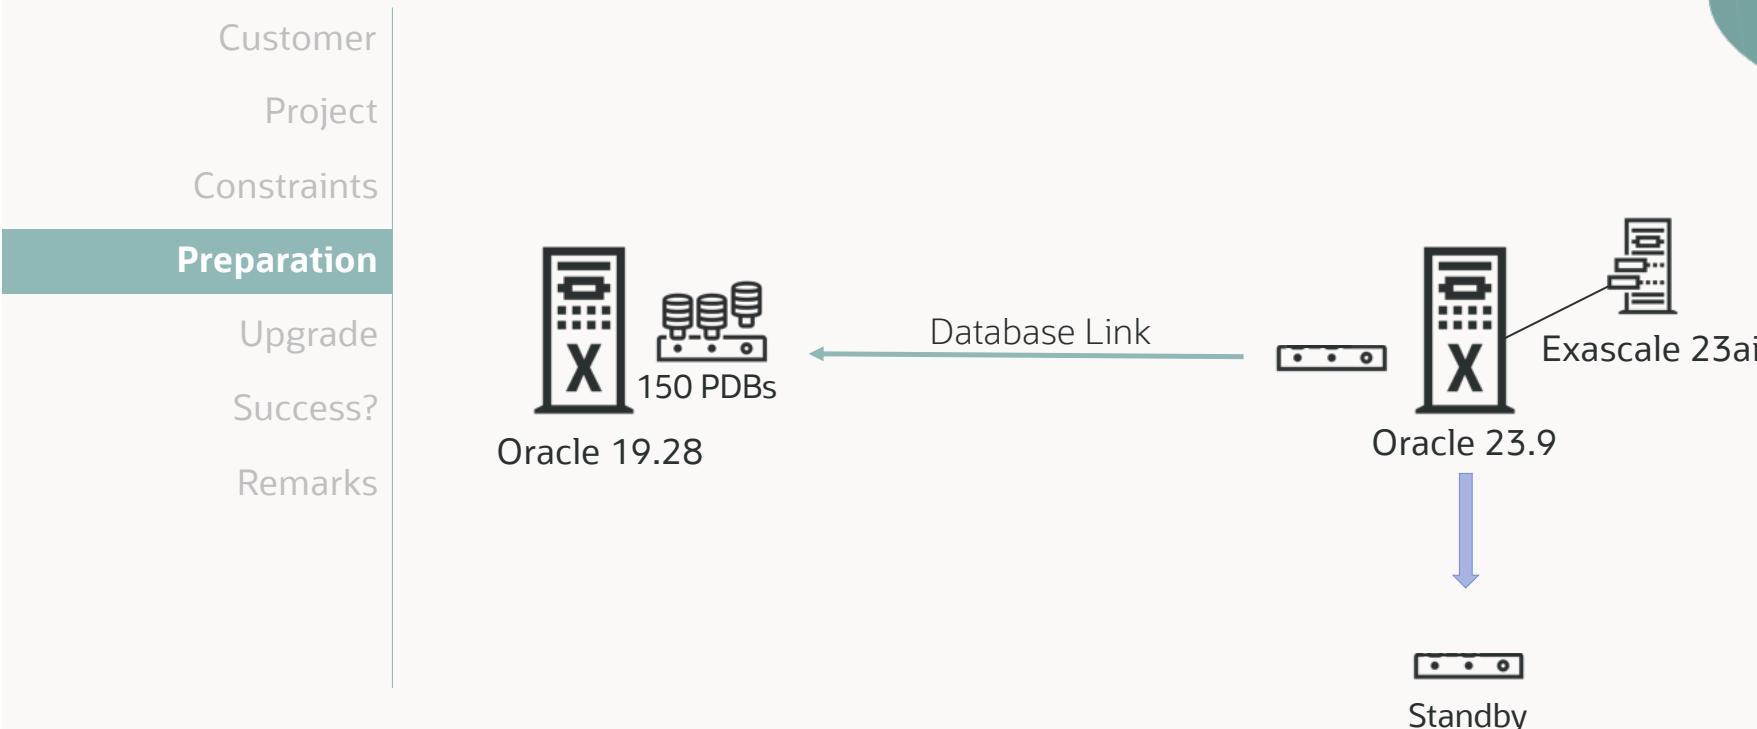
Remarks

COMPATIBLE=23.0.0

No ASM, no Oracle Database 19c anymore

Data Pump as rollback

Customer Case | Techniker Krankenkasse



Customer Case | Techniker Krankenkasse

Customer

Project

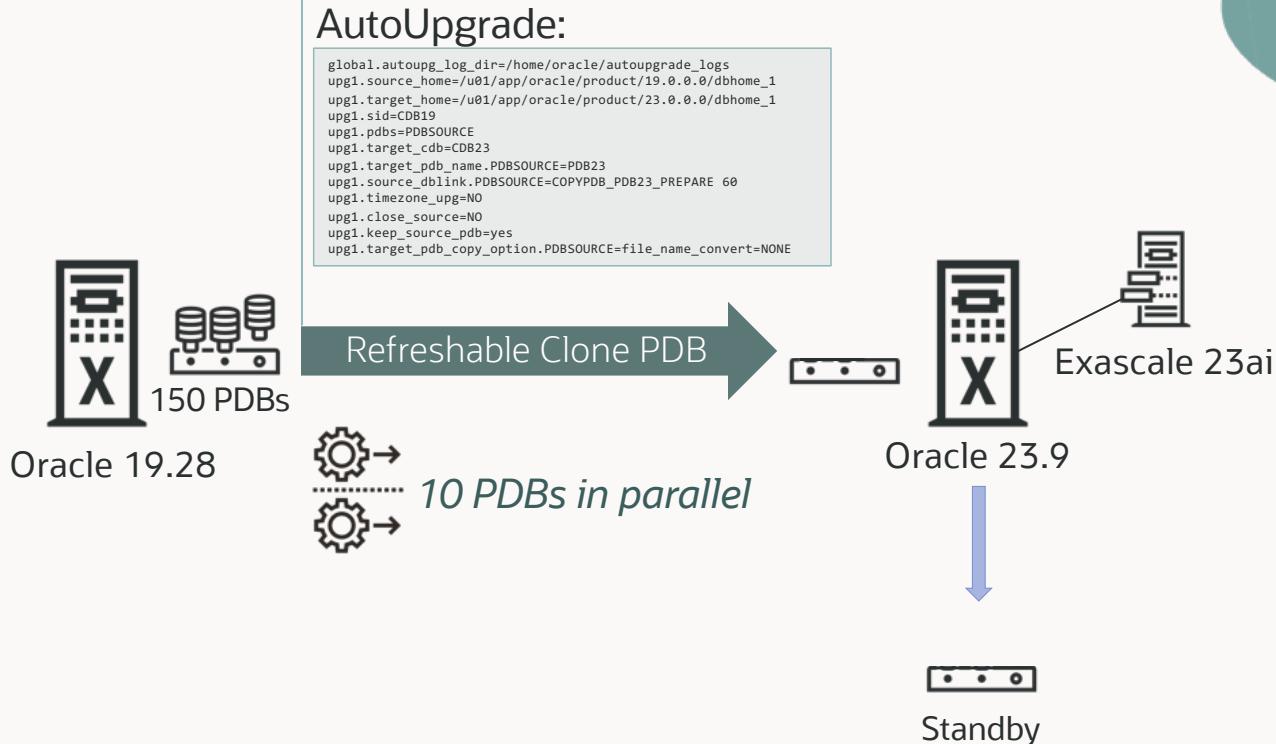
Constraints

Preparation

Upgrade

Success?

Remarks



Customer Case | Techniker Krankenkasse



Customer Case | Techniker Krankenkasse

| | |
|-----------------|-------|
| Customer | Yes!! |
| Project | |
| Constraints | |
| Preparation | |
| Upgrade | |
| Success? | |
| Remarks | |

Customer Case | Techniker Krankenkasse

| | |
|----------------|---|
| Customer | Standby building |
| Project | <ul style="list-style-type: none">• There were some issues |
| Constraints | <ul style="list-style-type: none">• Refreshable Clones don't propagate |
| Preparation | <ul style="list-style-type: none">• PDB recovery isolation doesn't work as expected |
| Upgrade | |
| Success? | A number of SRs had to be opened for various areas |
| Remarks | Optimizer works very well |



In the cloud, upgrading Oracle databases
is as **easy** as pressing a button



The easy approach



Upgrade entire CDB, including all PDBs



The fast approach



Upgrade single PDB

Upgrading in the cloud

Blog post:

[Upgrade Encrypted PDB in Cloud to Oracle AI Database 26ai](#)

Blog post:

[Upgrade Oracle Base Database Service to Oracle AI Database 26ai](#)

Blog post:

[Upgrade from 19c to 26ai using AutoUpgrade –ExaDB-D or ExaC@C – Part 1](#)



Blog post:

[When A Refreshable Clone Takes Over The Service](#)

Blog post:

[Upgrade to Oracle AI Database 26ai](#)

Documentation: [Proceed command](#)



Refreshable clone works only with deferred recovery on standby database

- You must restore the PDB on standby database after disconnect from non-CDB

Multitenant

Create Container Database



1

Character set

2

Components

3

COMPATIBLE

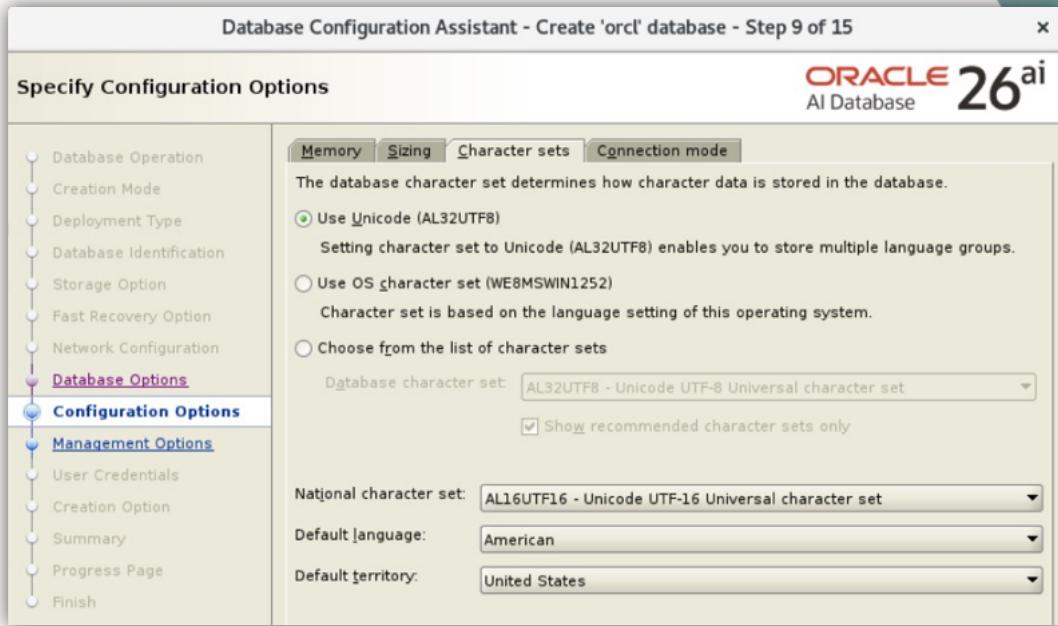
Create Container Database

1 Character set

- Always choose AL32UTF8
- Allows PDBs with any character set

2 Components

3 COMPATIBLE



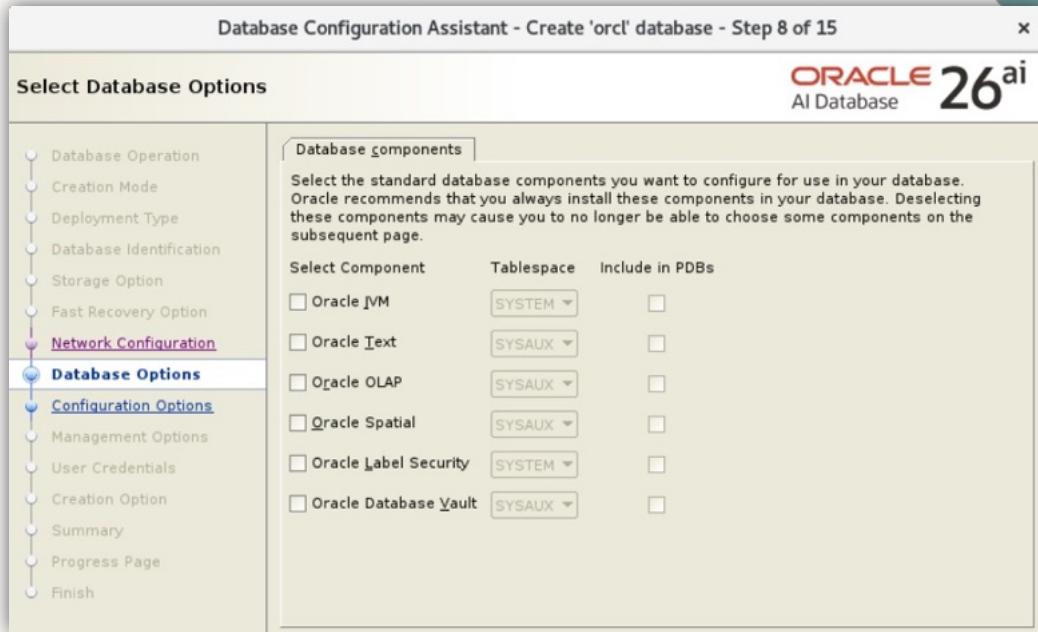
Create Container Database

1 Character set

2 Components

- Install as many as you need
- No more than that

3 COMPATIBLE



Create Container Database

1 Character set

2 Components

3 COMPATIBLE

- Keep at **23.6.0** to benefit from AI Vector Search features
- Unless you want the option of downgrade

All initialization parameters

⚠ Update the initialization parameters only when it is required. Refer to the Oracle documentation to learn more about each initialization parameter and its valid set of values.

(Storage related parameter(s) value is shown in MB)

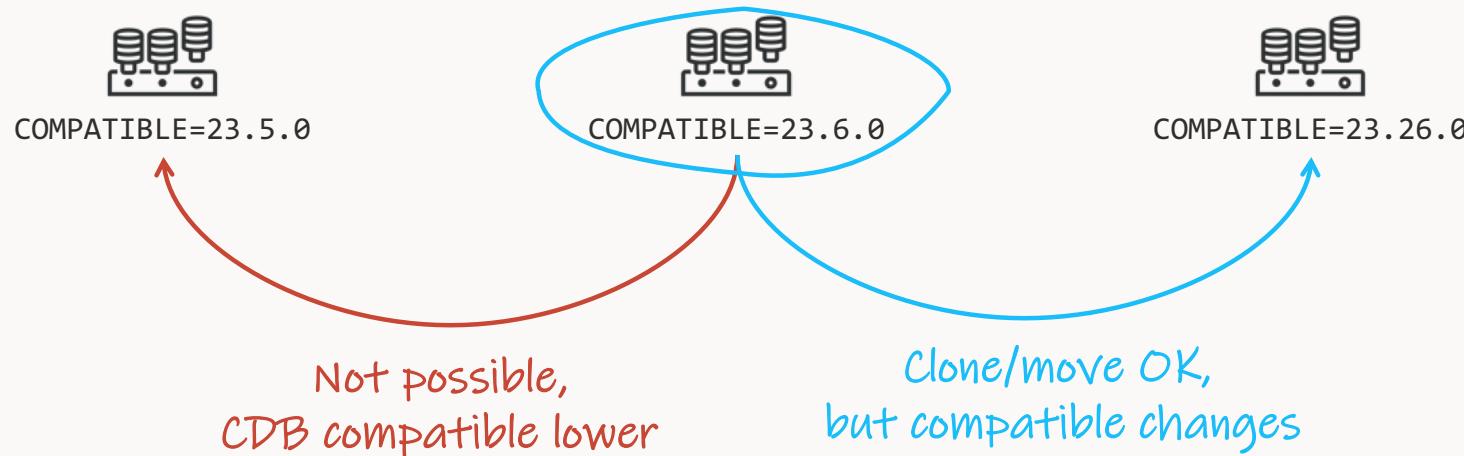
Show advanced parameters

| Name | Value | Include in spfile | Category |
|----------------------------------|------------------------------------|-------------------------------------|---|
| undo_tablespace | UNDOTBS1 | <input checked="" type="checkbox"/> | Cluster Database |
| sga_target | 19202 | <input checked="" type="checkbox"/> | SGA Memory |
| db_block_size (bytes) | 8192 | <input checked="" type="checkbox"/> | Cache and I/O |
| nls_language | AMERICAN | <input checked="" type="checkbox"/> | NLS |
| control_files | ('{ORACLE_BASE}/oradata/... | <input checked="" type="checkbox"/> | File Configuration |
| remote_login_passwordfile | EXCLUSIVE | <input checked="" type="checkbox"/> | Security and Auditing |
| processes | 640 | <input checked="" type="checkbox"/> | Processes and Sessions |
| pga_aggregate_target | 6401 | <input checked="" type="checkbox"/> | Sort, Hash Joins, Bitmap Indexes |
| nls_territory | AMERICA | <input checked="" type="checkbox"/> | NLS |
| open_cursors | 300 | <input checked="" type="checkbox"/> | Cursors and Library Cache |
| db_domain | livelabs.oraclecn.com | <input checked="" type="checkbox"/> | Database Identification |
| compatible | 23.6.0 | <input checked="" type="checkbox"/> | Miscellaneous |
| db_name | orcl | <input checked="" type="checkbox"/> | Database Identification |
| cluster_database | FALSE | <input type="checkbox"/> | Cluster Database |

Description:

compatible: Allows you to use a new release, while at the same time guaranteeing backward compatibility with an earlier release.
Range of Values: Default to current release. Default Value: Release dependent

Compatible



- You can now update the COMPATIBLE parameter without a restart
- Works within a release, e.g., from 23.6.0 to 23.26.0
- Cross-release changes, e.g., 19.0.0 to 23.6.0, still requires a restart

```
alter database set downgrade compatibility to '23.26.0';
```

Further Information

Container Database

YouTube:
[Create CDB in Oracle AI Database](#)

Blog post:
[How I Think You Should Create Your Container Databases](#)

[COMPATIBLE Parameter - Documentation](#)

Blog post:
[Online COMPATIBLE raise starting with 23.9](#)

```
SQL> select con$name, comp_id, status from cdb_registry;
```

| CON\$NAME | COMP_ID | STATUS |
|-----------|---------|--------|
| CDB\$ROOT | CATALOG | VALID |
| CDB\$ROOT | CATPROC | VALID |
| CDB\$ROOT | XDB | VALID |
| RED | CATALOG | VALID |
| RED | CATPROC | VALID |
| RED | XDB | VALID |

No information on
PDB\$SEED

```
alter system set "_exclude_seed_cdb_view"=false;
```

```
alter system set "_exclude_seed_cdb_view"=false;
```

```
SQL> select con$name, comp_id, status from cdb_registry;
```

| CON\$NAME | COMP_ID | STATUS |
|-----------|---------|--------|
|-----------|---------|--------|

| | | |
|-----------|---------|-------|
| CDB\$ROOT | CATALOG | VALID |
| CDB\$ROOT | CATPROC | VALID |
| CDB\$ROOT | XDB | VALID |
| PDB\$SEED | CATALOG | VALID |
| PDB\$SEED | CATPROC | VALID |
| PDB\$SEED | XDB | VALID |
| RED | CATALOG | VALID |
| RED | CATPROC | VALID |
| RED | XDB | VALID |



How does a standby database handles a plug-in operation

- Take care of your Data Guard during migration

Data Guard



1

Enabled recovery

2

Deferred recovery

*Plug-in on primary propagates
to standby database via redo*



Enabled Recovery

1

Enabled recovery

```
create pluggable database ... standbys=all
```

Standby records PDB creation

Standby locates data files

MRP applies redo to PDB

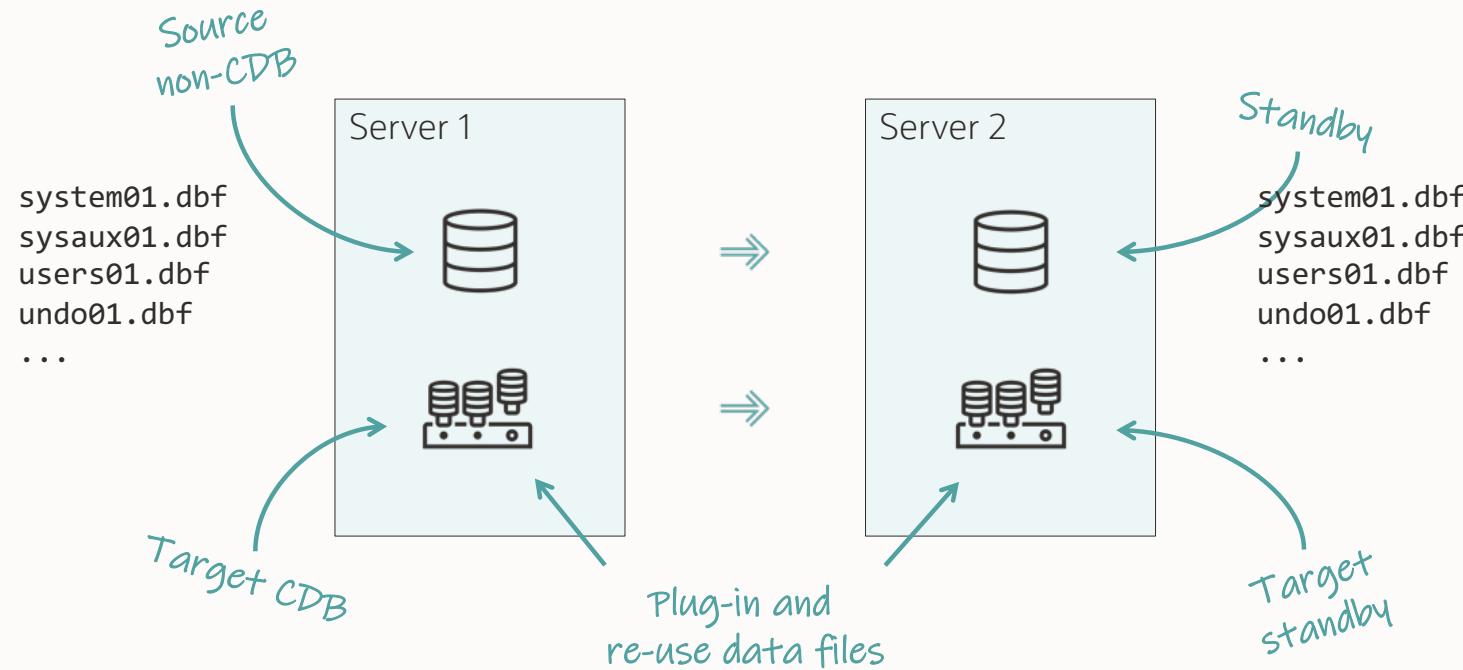
PDB is immediately protected

Default

2

Deferred recovery

Enabled Recovery





All data files on primary and standby
must be at the same SCN

Enabled Recovery

- The plug-in happens on the primary database
- The plug-in uses the *manifest file*
- The manifest file contains information on data files from the primary database only

```
<?xml version="1.0" encoding="UTF-8"?>
<PDB>
  <xmlversion>1</xmlversion>
  <pdbsname>DB12</pdbsname>
  <cid>0</cid>
  <byteorder>1</byteorder>
  <vsn>203424000</vsn>
  <vsn>
    <vsnnum>12.2.0.1.0</vsnnum>
    <cdbscompt>12.2.0.0.0</cdbscompt>
    <pdbscompt>12.2.0.0.0</pdbscompt>
    <vsnlibnum>0.0.0.0.24</vsnlibnum>
    <vsnsql>24</vsnsql>
    <vsnbsv>8.0.0.0.0</vsnbsv>
  </vsn>
  <dbid>1852833295</dbid>
  <ncdb2pdbs>1</ncdb2pdbs>
  <cdbid>1852833295</cdbid>
  <guid>8605DC2587337002E0532AB2A8C0A57C</guid>
  <uscnbas>4437941</uscnbas>
  <uscnwrc>0</uscnwrc>
  <undoscn>8</undoscn>
  <rdba>4194824</rdba>
  <tablespace>
    <name>SYSTEM</name>
    <type>0</type>
    <tsn>0</tsn>
    <status>1</status>
    <issft>0</issft>
    <isnft>0</isnft>
    <encts>0</encts>
    <flags>0</flags>
    <bmunitsize>8</bmunitsize>
    <file>
      <path>/u02/oradata/DB12/system01.dbf</path>
      <afn>1</afn>
      <rfn>1</rfn>
```

Enabled Recovery

- The plug-in happens on the primary database
- The plug-in uses the *manifest file*
- The manifest file contains information on data files from the primary database only

**How does the standby database
know which files to plug in?**

Enabled Recovery

How does the standby database know which files to plug in?

- 1 Regular files
- 2 OMF in regular file system
- 3 ASM

Enabled Recovery

1 Regular files

- Standby search for data files at the same location as the primary
- Override with `DB_FILE_NAME_CONVERT`
- Or, override with `STANDBY_PDB_SOURCE_FILE_DIRECTORY`

Enabled Recovery

2 OMF in regular file system

- Standby search for data files at the OMF location (`DB_CREATE_FILE_DEST`)
- Move data files from non-CDB location into OMF location
- Or, create soft links in OMF location pointing to data file location

Enabled Recovery

3 ASM

- Standby search for data files at the OMF location (`DB_CREATE_FILE_DEST`)
- Use ASM aliases to find data files at non-CDB OMF location

Enabled Recovery | ASM

Primary

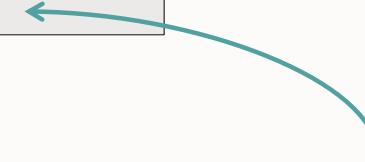


Standby



```
SQL> select name from v$datafile;

NAME
-----
+DATA/DB_BOSTON/DATAFILE/system.269.1103046537
+DATA/DB_BOSTON/DATAFILE/sysaux.270.1103046537
+DATA/DB_BOSTON/DATAFILE/users.273.1103046827
```



Same file,
but different name

```
SQL> select name from v$datafile;

NAME
-----
+DATA/DB_CHICAGO/DATAFILE/system.265.1103050007
+DATA/DB_CHICAGO/DATAFILE/sysaux.266.1103050007
+DATA/DB_CHICAGO/DATAFILE/users.269.1103050009
```

Enabled Recovery | ASM

19c
Non-CDB
Primary



19c
Non-CDB
Standby



The manifest file contains

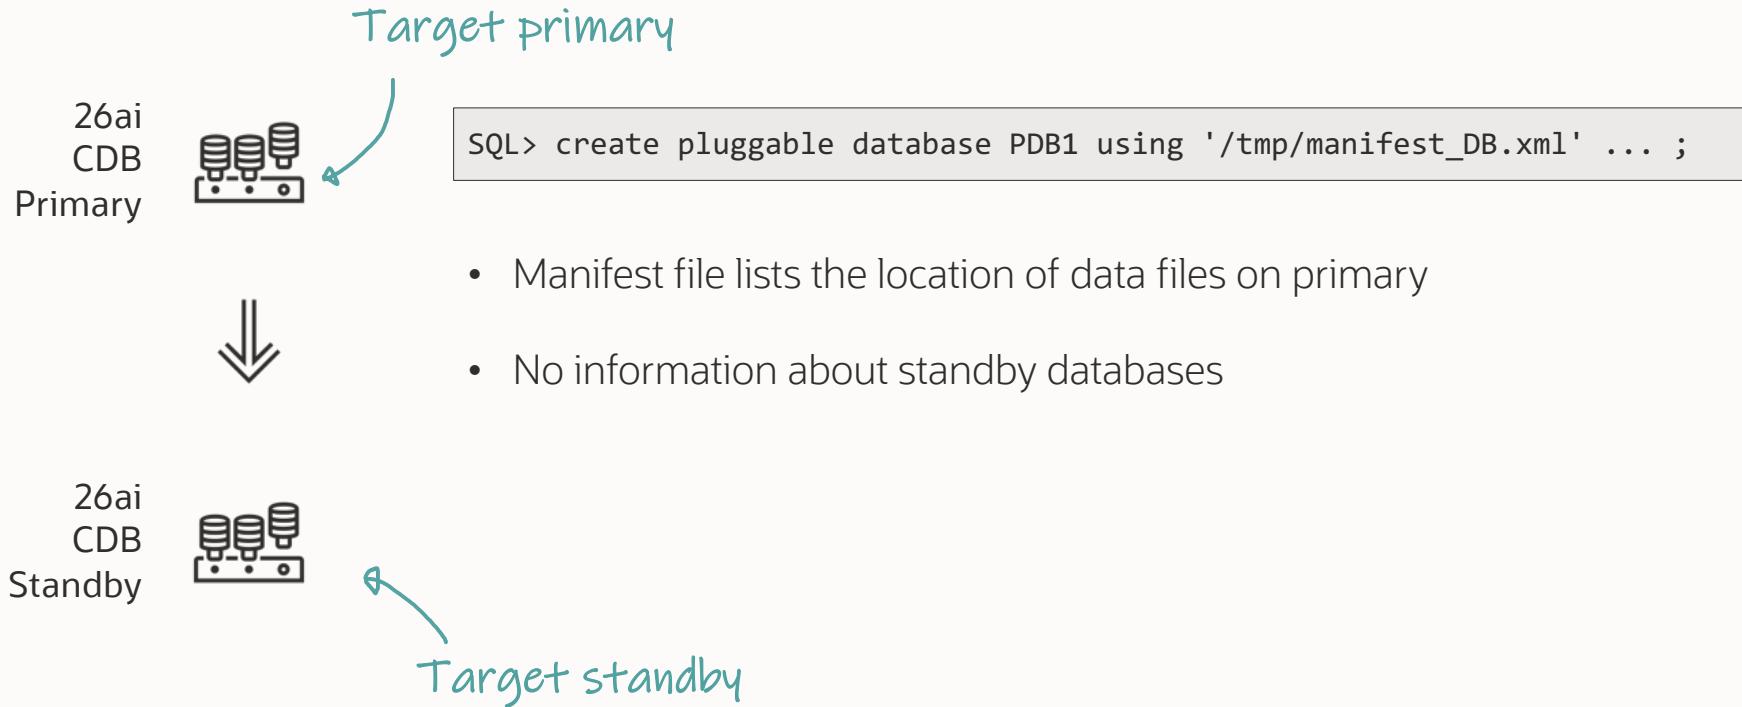
```
SQL>exec dbms_pdb.describe('tmp/manifest_DB.xml');
```

- Not standby database

⋮

```
<?xml version="1.0" encoding="UTF-8"?>
<PDB>
  <xmlversion>1</xmlversion>
  <pdbname>PDB1</pdbname>
  ...
  <guid>DDB49CFEFD8ED4FCE053E801000A078C</guid>
  ...
  <tablespace>
    <name>USERS</name>
    ...
    <file>
      <path>+DATA/DB_BOSTON/DATAFILE/users.273.1103046827</path>
```

Enabled Recovery | ASM



Enabled Recovery | ASM

26ai
CDB
Primary



+DATA/DB_BOSTON/DATAFILE/users.273.1103046827



Redo record says:
Plug in this data file

26ai
CDB
Standby



+DATA/DB_CHICAGO/DATAFILE/users.269.1103050009



No good, data file
has a different name



Enabled Recovery | ASM

26ai
CDB
Primary



+DATA/DB_BOSTON/DATAFILE/users.273.1103046827



26ai
CDB
Standby



+DATA/DB_CHICAGO/DATAFILE/users.269.1103050009

+DATA/CDB1_CHICAGO/<PDB_GUID>/DATAFILE

OK, let's check the OMF directory



It's empty

Enabled Recovery | ASM

26ai
CDB
Primary



+DATA/DB_BOSTON/DATAFILE/users.273.1103046827



26ai
CDB
Standby



OK, let's check the OMF directory
+DATA/DB_CHICAGO/DATAFILE/users.269.1103050009
+DATA/CDB1_CHICAGO/<PDB_GUID>/DATAFILE

It's empty



I'll just move the file in ASM



There's no **move** command in ASM.
But you can create *aliases*

- Similar to file system soft links

```
SQL> alter diskgroup data add alias
  '+DATA/DB_CHICAGO/DATAFILE/users.269.1103050009'
  for
  '+DATA/CDB1_CHICAGO/<PDB_GUID>/DATAFILE/users.269.1103050009' :
```

Non-CDB standby data file



```
SQL> alter diskgroup data add alias
  '+DATA/DB_CHICAGO/DATAFILE/users.269.1103050009'
for
  '+DATA/CDB1_CHICAGO/<PDB_GUID>/DATAFILE/users.269.1103050009' :
```

```
SQL> alter diskgroup data add alias  
'+DATA/DB_CHICAGO/DATAFILE/users.269.1103050009'  
for  
'+DATA/CDB1_CHICAGO/<PDB_GUID>/DATAFILE/users.269.1103050009':
```



Standby PDB OMF location

```
SQL> alter diskgroup data add alias  
'+DATA/DB_CHICAGO/DATAFILE/users.269.1103050009'  
for  
'+DATA/CDB1_CHICAGO/<PDB_GUID>/DATAFILE/users.269.1103050009':
```



Name does not matter.
Standby scans all files in OMF directory



Data Guard | Re-use Data Files

26ai
CDB
Primary



26ai
CDB
Standby



- Standby database scans its own OMF directory for data files
- Standby ignores file names and look at file headers
- Standby will find aliases and find the real file locations

Data Guard | Re-use Data Files

Primary



Standby



Looking for file like on primary

```
Recovery scanning directory +DATA/DB_BOSTON/... for any matching files
Deleted Oracle managed file +DATA/DB_BOSTON/...
Successfully added datafile 37 to media recovery
Datafile #37: +DATA/DB_CHICAGO/DATAFILE/users.269.1103050009
```

Follows alias and finds the real file



What happens with enabled recovery
if the standby fails to find the data files?

Enabled Recovery | Missing Data Files

What if a standby database fails to find data files?

- If Active Data Guard and PDB Recovery Isolation is turned on
 - New feature in Oracle Database 21c
 - Recovery disabled for PDB
 - Recovery proceeds in the entire CDB, except in specific PDB
 - Standby automatically restores data files from primary and re-enables recovery afterward
 - PDB protected after auto-restore
- If not, recovery halts in the **entire** CDB
 - **This is a critical situation**



What about AutoUpgrade
and enabled recovery?

Enabled Recovery | AutoUpgrade

AutoUpgrade does not support plugging in with enabled recovery

- Enabled recovery requires work on both primary and standby hosts
- You must execute commands at specific times
- It's complicated - but we're working on it



Deferred Recovery

1

Enabled recovery

`create pluggable database ... standbys=all`

Standby records PDB creation

Standby locates data files

MRP applies redo to PDB

PDB is immediately protected

2

Deferred recovery

`create pluggable database ... standbys=none`

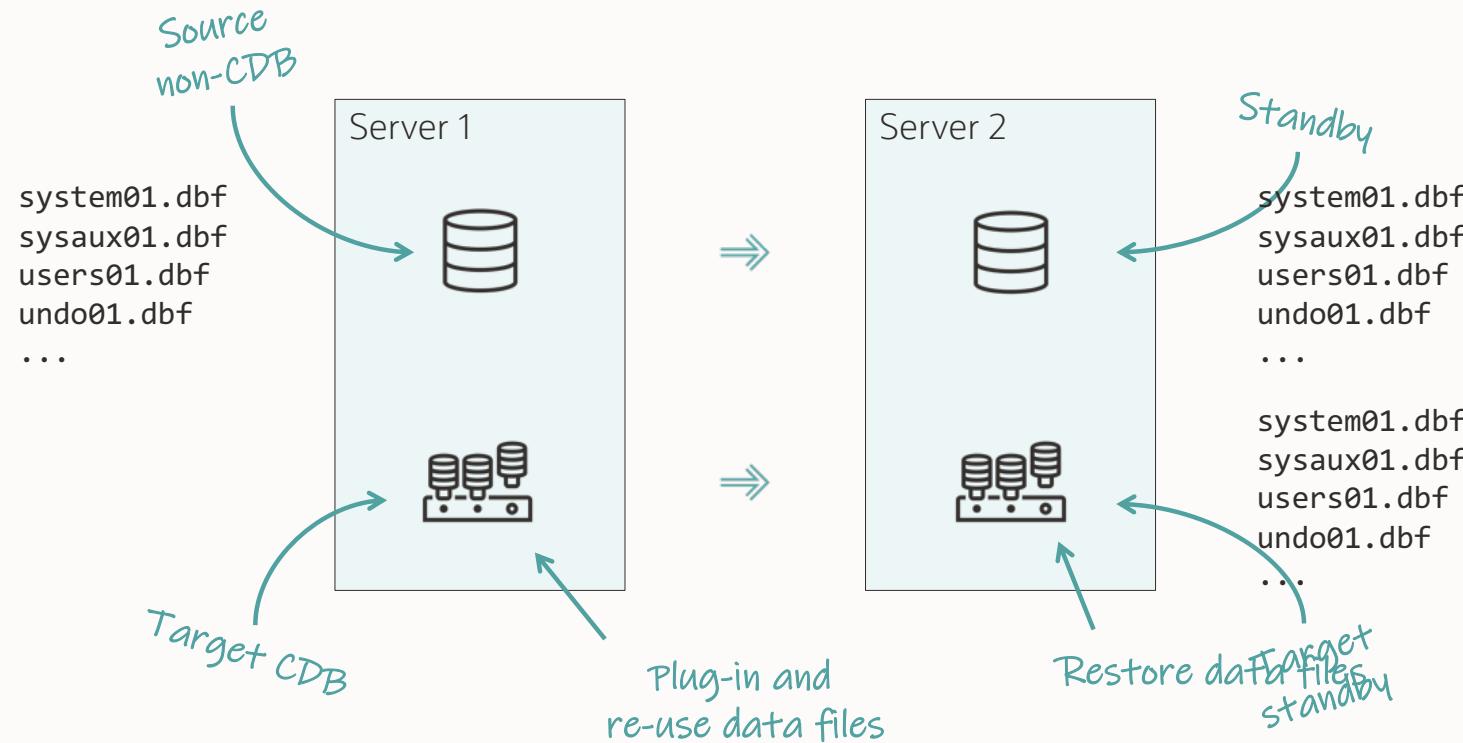
Standby records PDB creation

Standby ignores data files

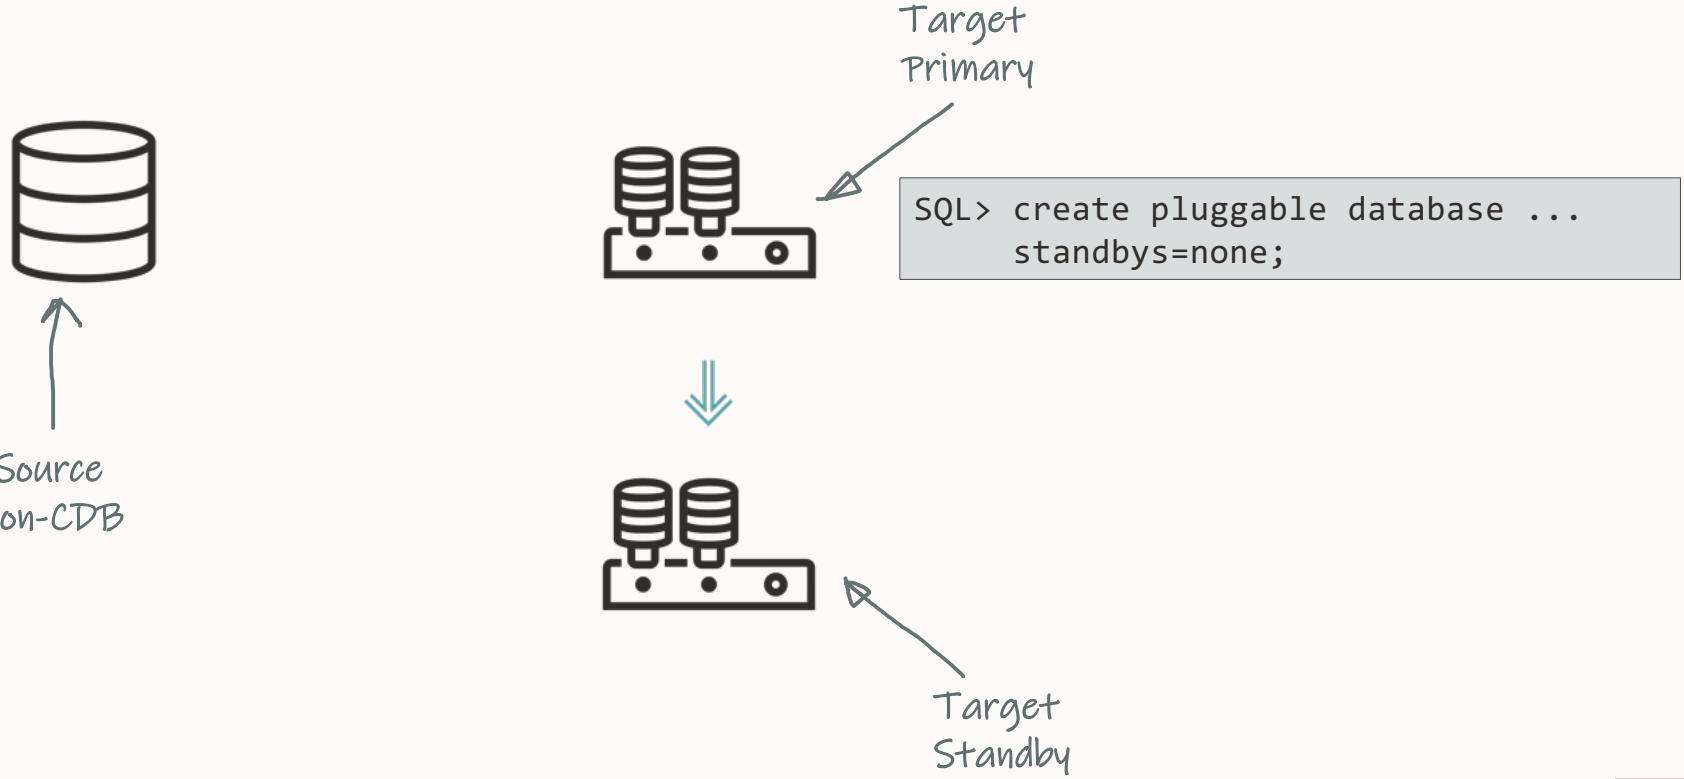
MRP skips redo

PDB protected after restore

Deferred Recovery



Deferred Recovery



Deferred Recovery



```
SQL> create pluggable database ...  
standbys=none;
```



PDB created
Data files missing

Deferred Recovery



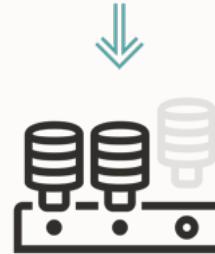
```
SQL> show pdbs
```

| CON_NAME | OPEN MODE |
|----------|------------|
| PDB1 | READ WRITE |

```
SQL> show pdbs
```

| CON_NAME | OPEN MODE |
|----------|-----------|
| PDB1 | MOUNTED |

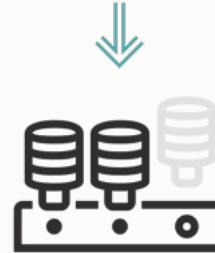
Deferred Recovery



```
SQL> select name, recovery_status  
      from v$pdbs;
```

| NAME | RECOVERY_STATUS |
|------|-----------------|
| PDB1 | DISABLED |

Deferred Recovery



```
RMAN> restore pluggable database  
... from service ... ;  
  
SQL> alter pluggable database  
enable recovery;  
SQL> alter database datafile  
... online;
```



Deferred Recovery



```
RMAN> restore pluggable database
      ... from service ... ;

SQL> alter pluggable database
      enable recovery;
SQL> alter database datafile
      ... online;
```

- Automated process in Oracle Database 21c
- PDB Recovery Isolation
- Requires Active Data Guard

Further Details

Webinar:

[Move to Oracle Database 23ai –
Everything you need to know about Oracle Multitenant – Part 1](#)



Don't jeopardize your Data Guard

- Test the procedure and verify your environment



Operations



Share resources between PDBs

Resource Consolidation

Non-CDB
database



Memory



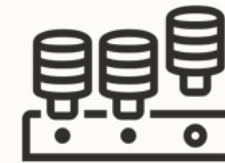
Background processes



Files

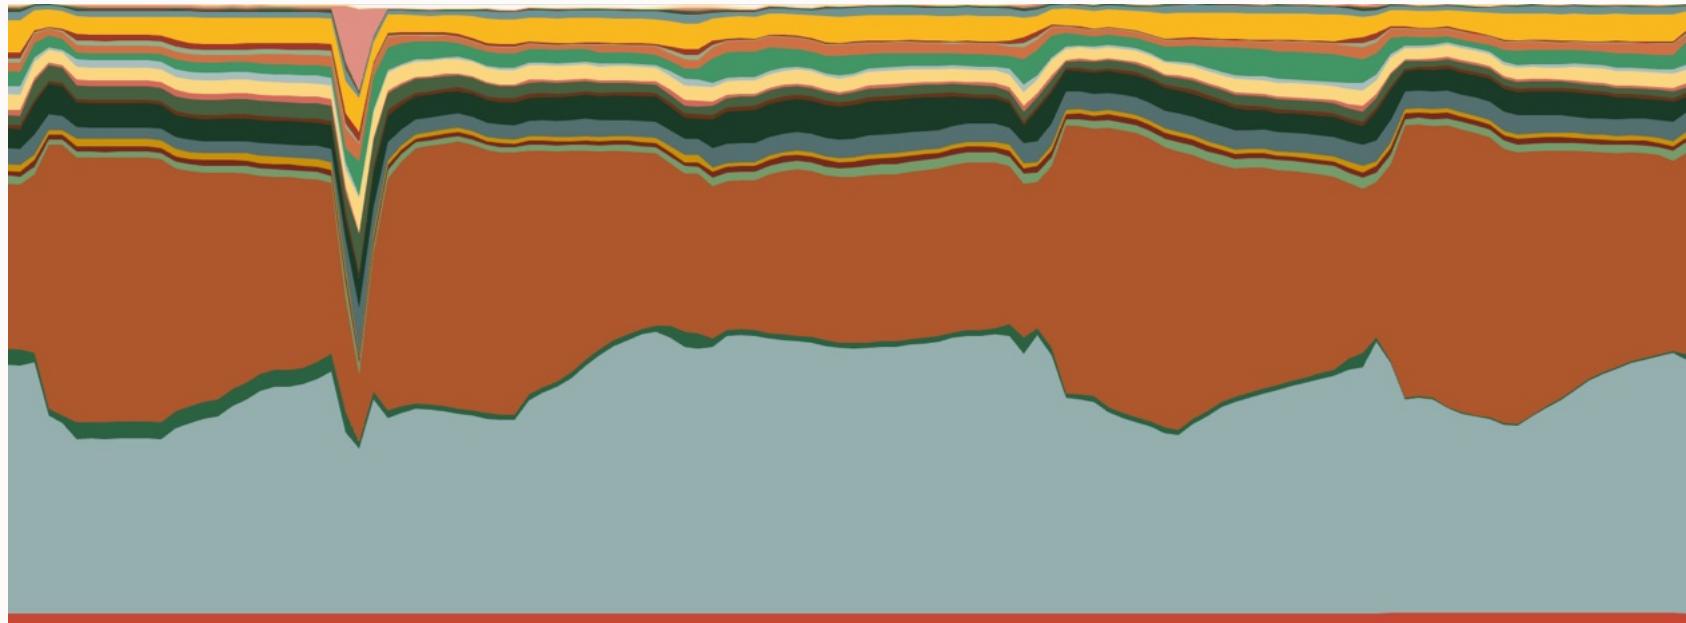


Resource Consolidation





Resource Consolidation



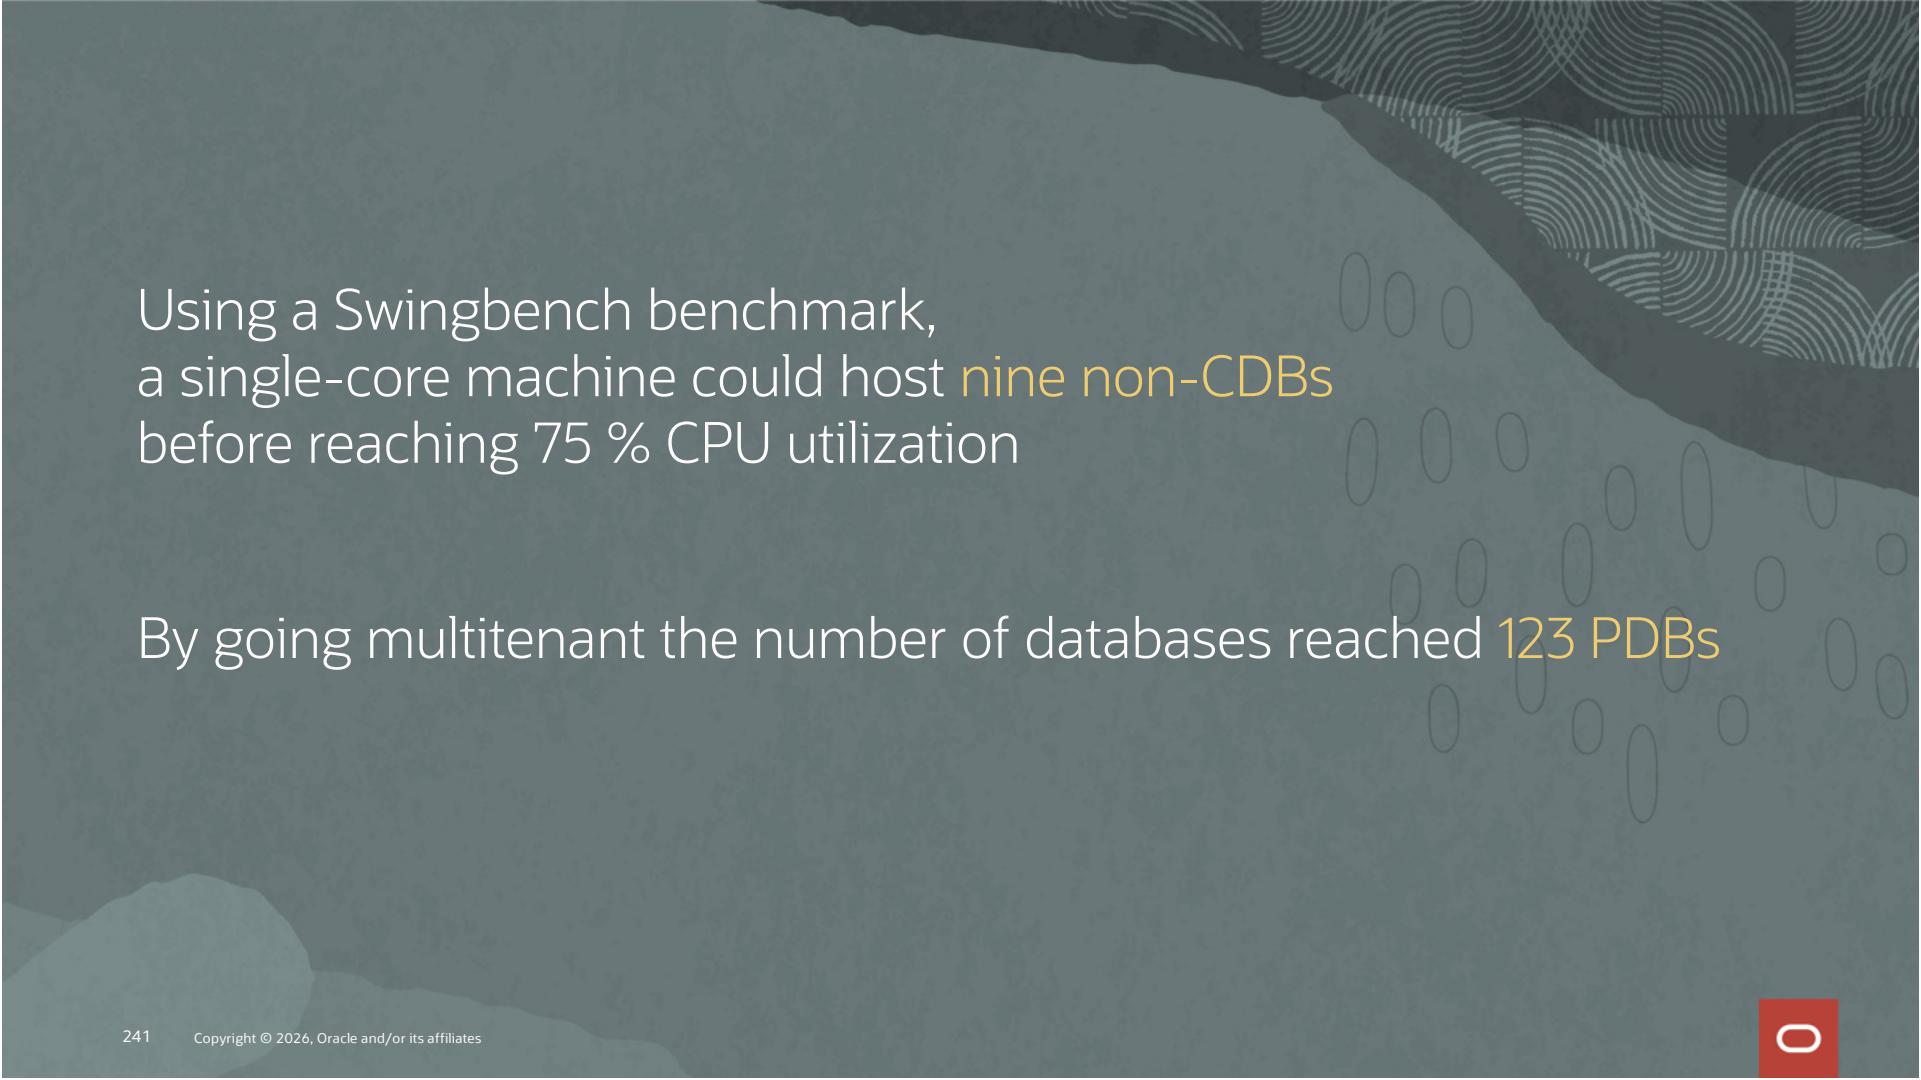
Consolidation Strategies?

There is no "*best*" strategy



- Don't mix PDBs with different SLAs
- Total memory consumption minus 20%-30%
- Increase consolidation factor slowly
- Constrain PDBs with instance caging (CPU_COUNT)
- Be open to start with a completely new naming schema



The background of the slide features a dark grey gradient with a subtle, abstract pattern. It includes faint, light-grey binary digits (0s and 1s) scattered across the surface, and several wavy, horizontal lines of the same binary pattern that create a sense of depth and motion.

Using a Swingbench benchmark,
a single-core machine could host **nine** non-CDBs
before reaching 75 % CPU utilization

By going multitenant the number of databases reached **123** PDBs

A US Health Care provider managed to

- Reduce the number of database instances by 7x
- Reduce the number of physical servers by 50 %

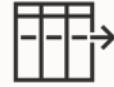


You can run multiple CDBs on the same host and out of the same Oracle home

Consolidation



Schema consolidation



Virtual Private Database



PDB consolidation

- Less complexity
- Better isolation
- Operational benefits
- Easier cloning

A global provider of financial services states

The multitenant architecture gives us complete client separation out of the box, without having to maintain a Virtual Private Database setup.

We went away from Virtual Private Database and consolidated our different clients in individual PDBs.

This reduced the complexity of our database implementation and made operations much easier.



Avoid noisy neighbors

- Allow sharing resources but everyone must get a fair share



Method 1

Instance caging



- Most simple
- Define **CPU_COUNT** for each PDB
- Hard limit

Method 1



8 CPUs



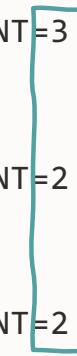
CPU_COUNT=3



CPU_COUNT=2

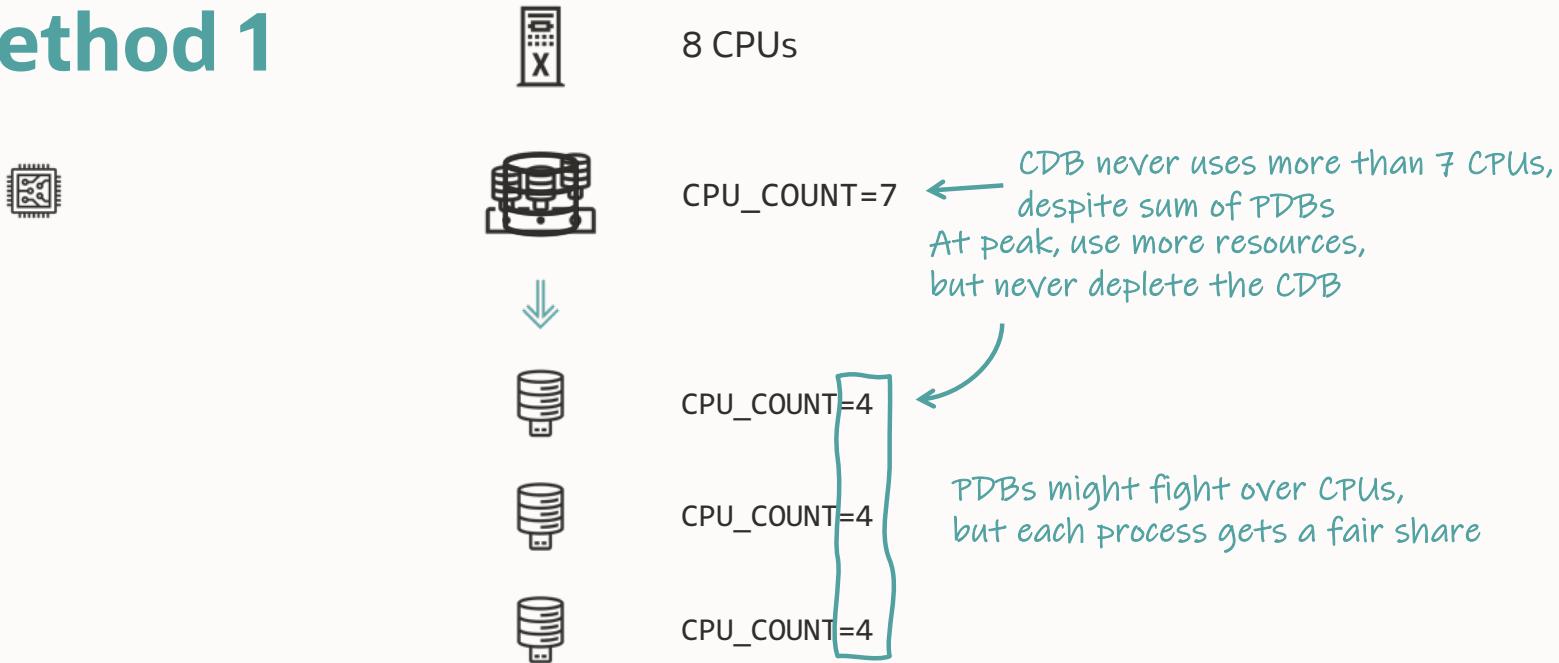


CPU_COUNT=2



All non-CDBs
share 7 CPU

Method 1



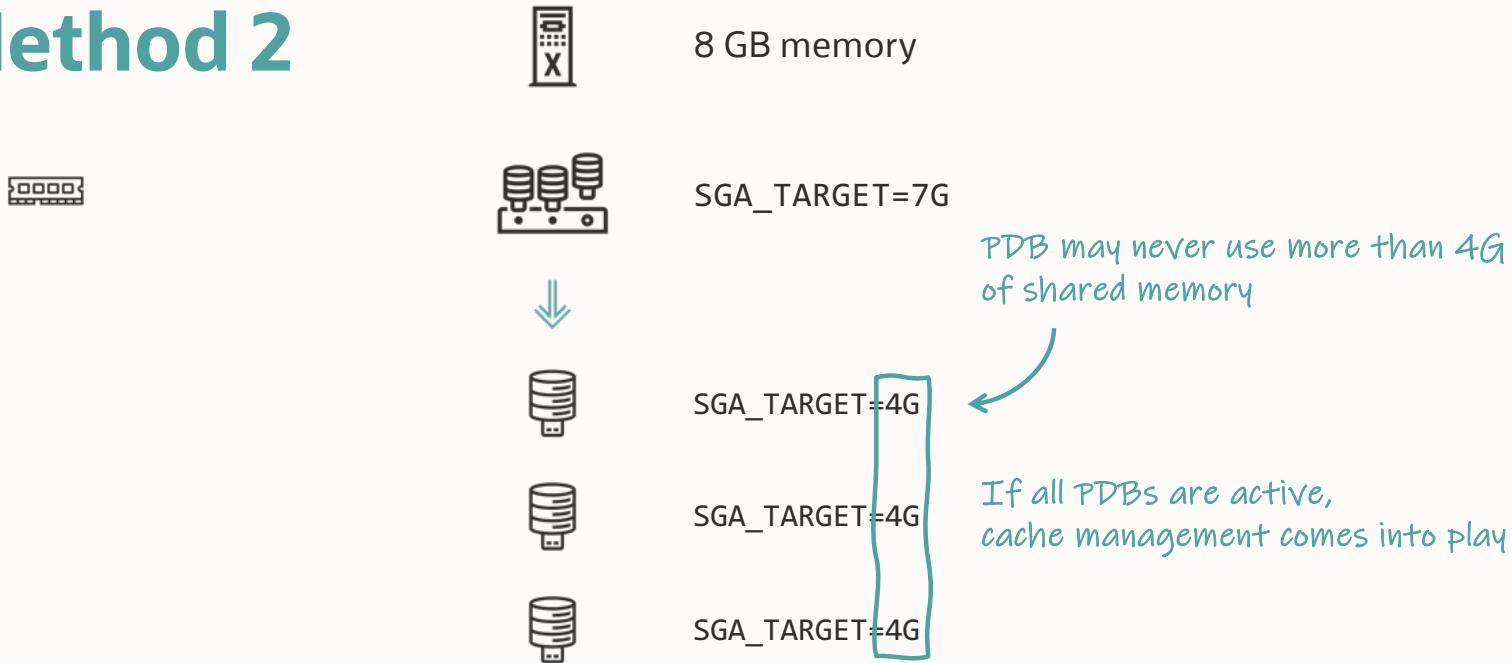
Method 2

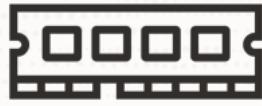
Memory allocation



- Simple
- Define **SGA_TARGET** for each PDB
- Hard limit

Method 2





Requires use of
Automatic Shared Memory Management

- Both in CDB and PDB



Optionally, allocate minimum shared pool and buffer cache for a PDB

- Use `SHARED_POOL_SIZE` and `DB_CACHE_SIZE`



You can combine method 1 and 2

- Instance caging and memory allocation

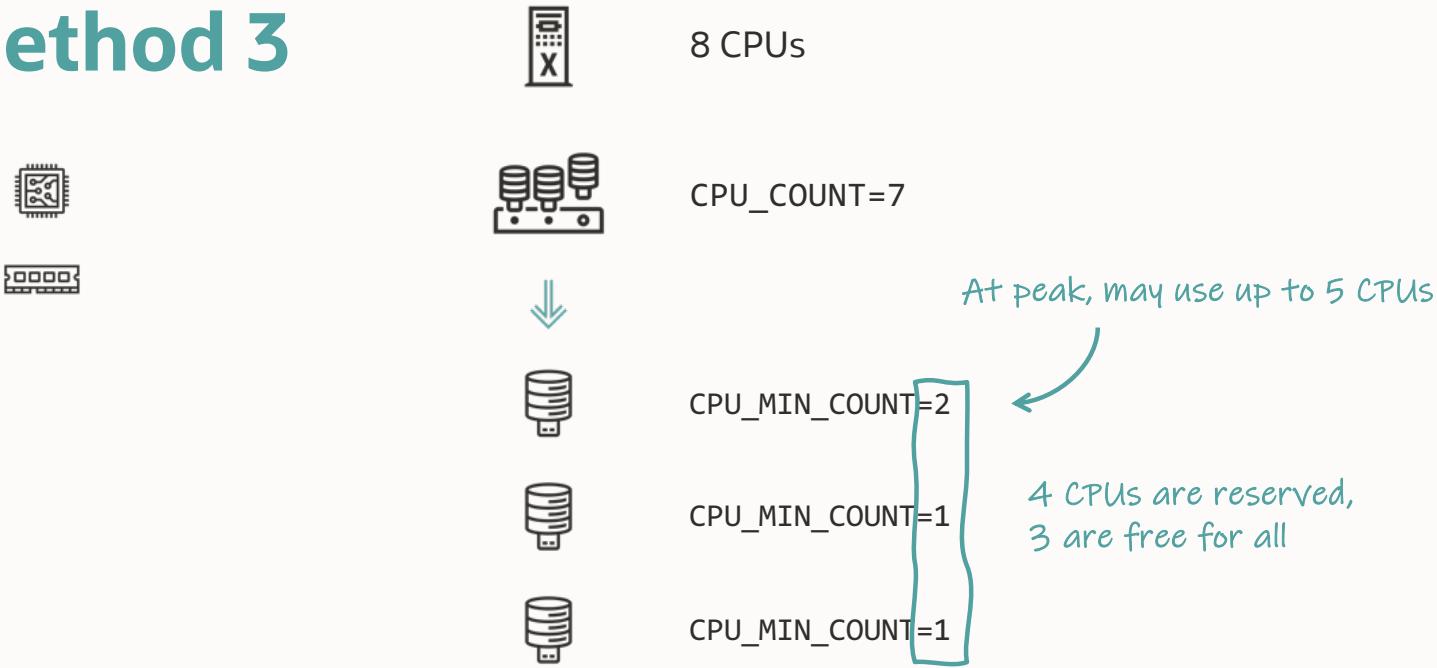
Method 3

Simple Resource Manager

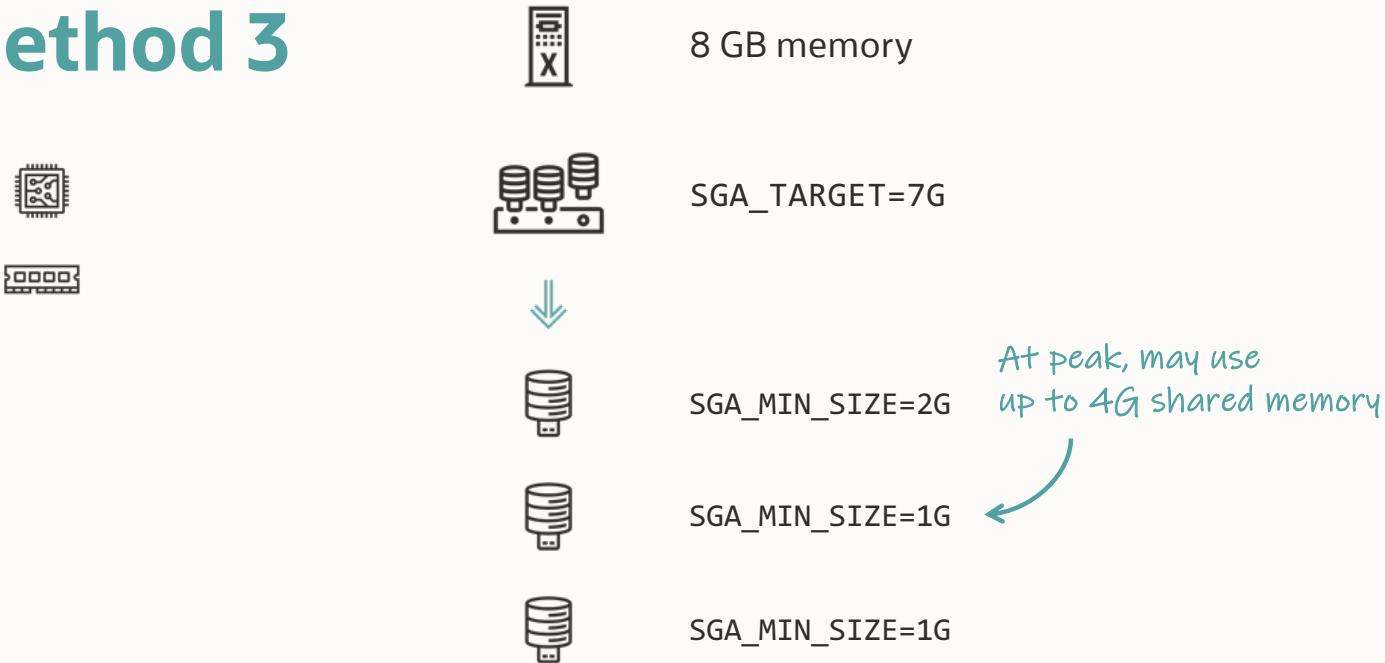


- Elaborate, yet simple to implement
- Enable CDB resource manager
- Allocate minimum shares instead of hard limits
- For advanced use cases

Method 3



Method 3





Requires Resource Manager at root level

```
alter session set container=cdb$root;

-- Create an empty resource manager plan with no directives
exec dbms_resource_manager.clear_pending_area;
exec dbms_resource_manager.create_pending_area;
exec dbms_resource_manager.create_cdb_plan('CDB_PLAN');
exec dbms_resource_manager.validate_pending_area;
exec dbms_resource_manager.submit_pending_area;

-- Make plan active in root to enable CDB resource manager
alter system set resource_manager_plan=CDB_PLAN;
```

Method 4

Advanced Resource Manager



- Requires additional configuration, but much greater control
- Use directives instead of shares



You can still control resources inside a PDB
with Resource Manager



What about I/O?

- Exadata I/O Resource Management
- Or, **MAX_MBPS** and **MAX_IOPS**



You can run multiple CDBs on the same host and out of the same Oracle home

Inter-instance Resource Management

Shares resources like with non-CDBs:

- CPU_COUNT
- SGA_MAX_SIZE

Inter-instance CPU resource manager:

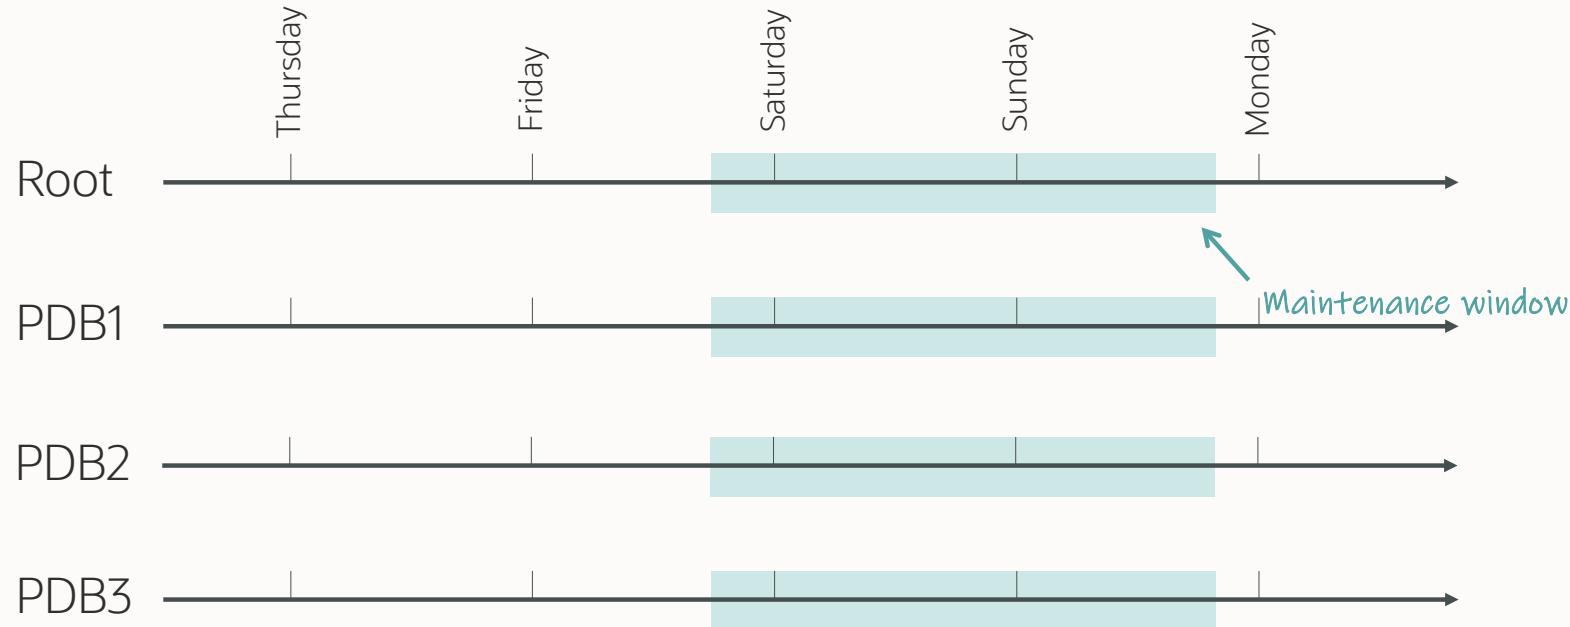
- Controls resource sharing using Linux c-groups
- Check [RESOURCE_MANAGER_CPU_SCOPE](#)
- Exadata Database Machine and Autonomous AI Database



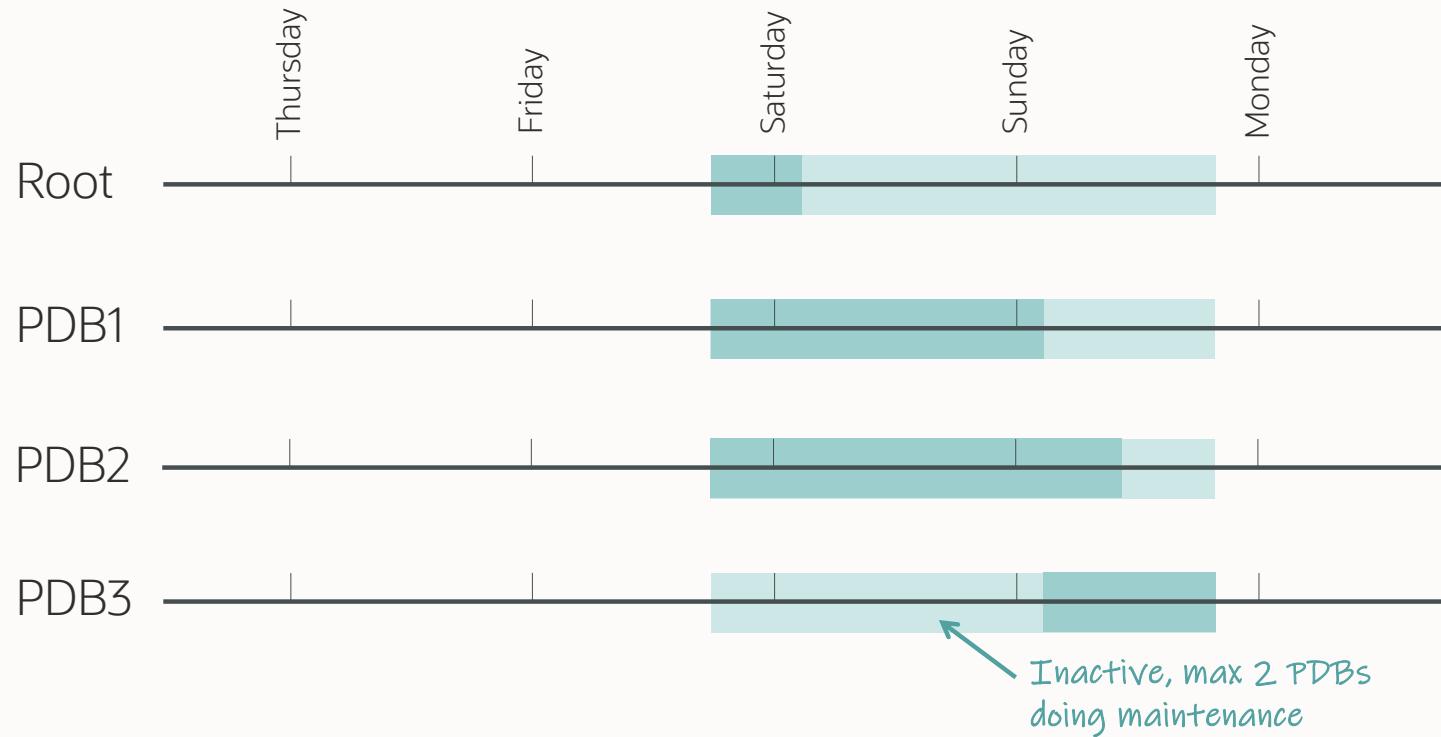


A word about
automated maintenance tasks

Automated Maintenance Tasks



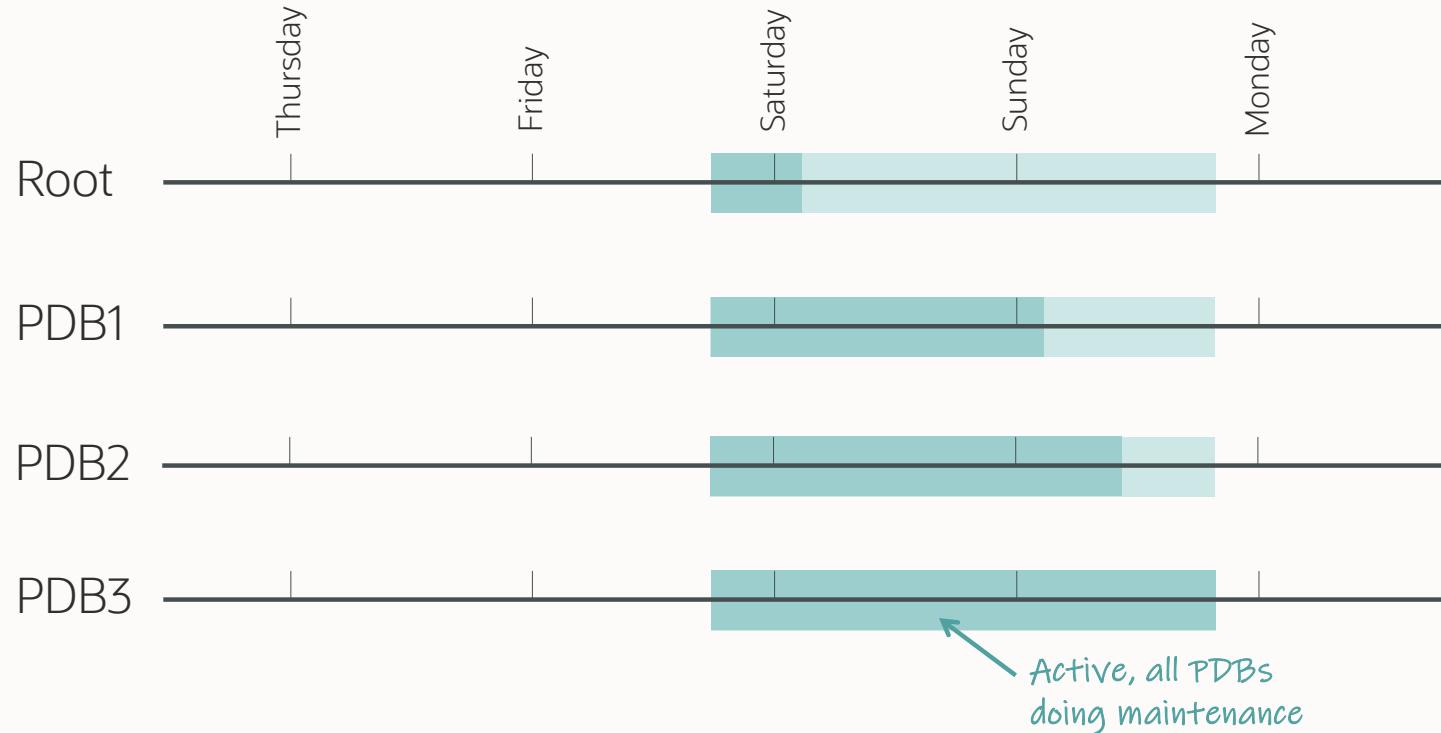
Automated Maintenance Tasks



--Change the amount of PDBs that can run maintenance tasks at the same time
--Default value 2

```
alter system set autotask_max_active_pdb=3;
```

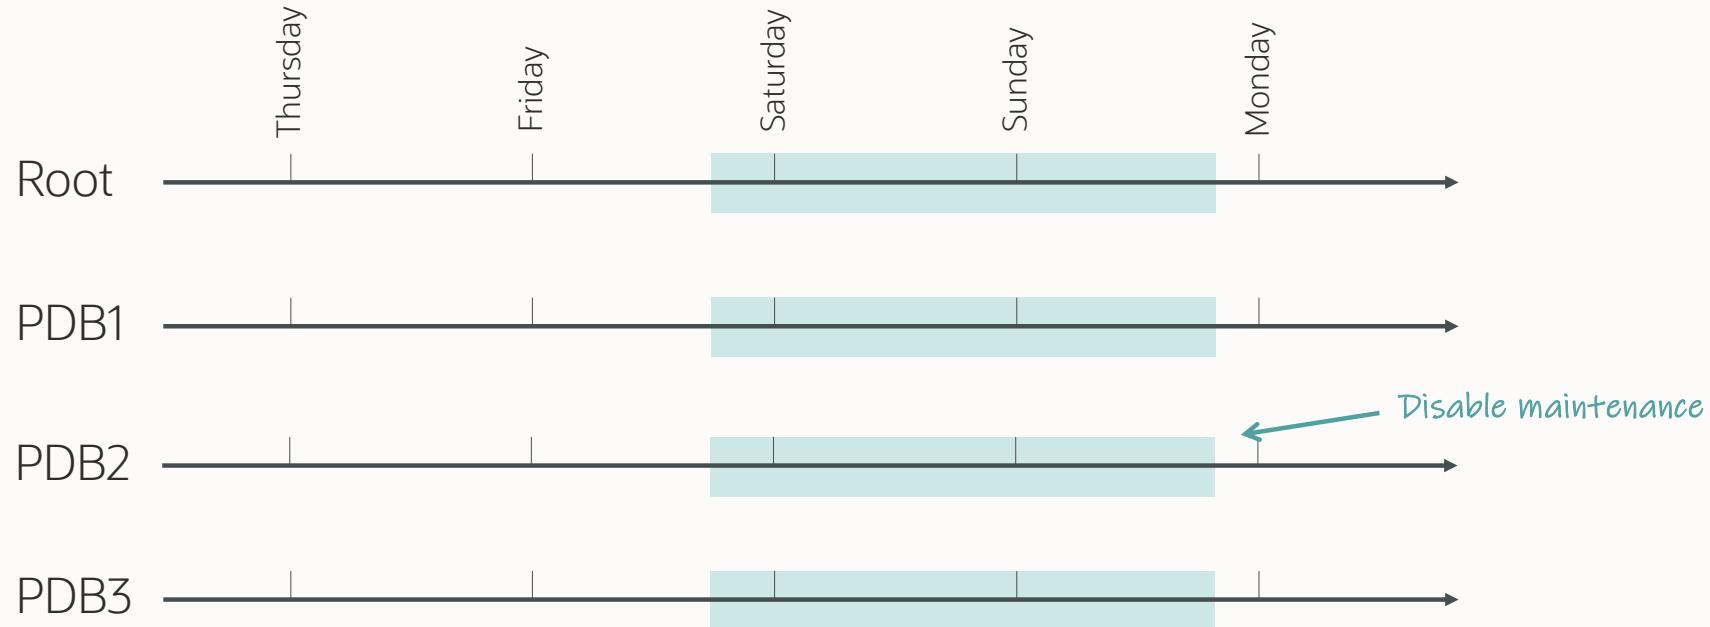
Automated Maintenance Tasks



--Selectively disable maintenance tasks in a PDB
--For instance, test databases or databases that are rebuilt frequently

```
alter session set container=PDB2;  
alter system set enable_automatic_maintenance_pdb=false;
```

Automated Maintenance Tasks

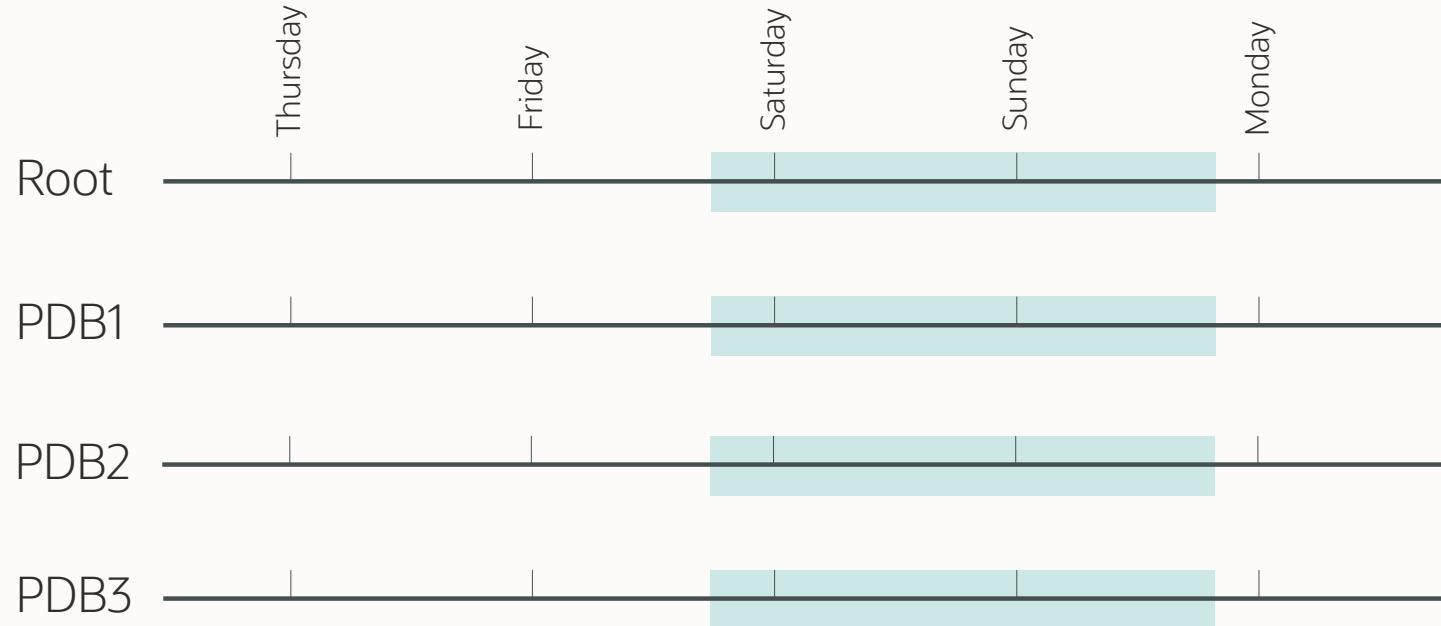




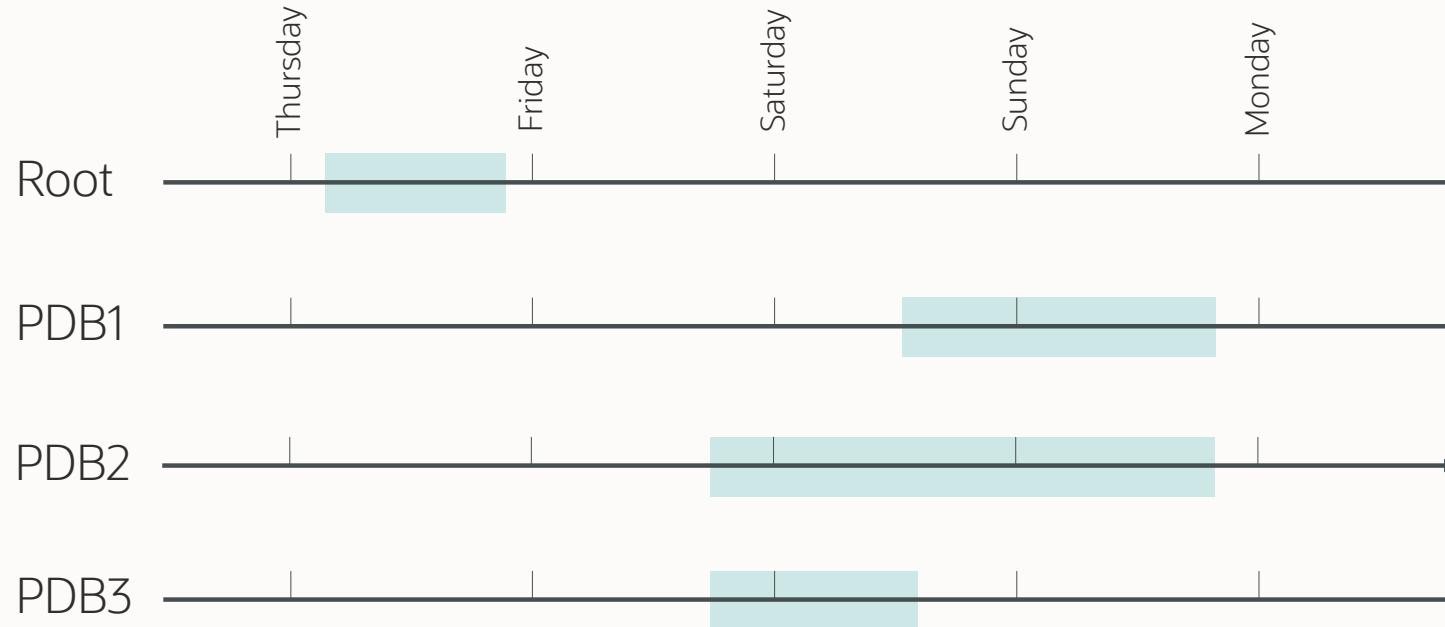
Shift maintenance windows

- Optionally, shorten maintenance windows

Automated Maintenance Tasks



Automated Maintenance Tasks





Selectively disable individual maintenance tasks using `DBMS_AUTO_TASK_ADMIN`

- Does a test database need Automatic Segment Advisor?
- Or Evolve Advisor?



Resource Manager prevents maintenance tasks from *stealing* resources from users

- Consumer group **ORA\$AUTOTASK**



Key Learnings



- 1 Get started with Multitenant
- 2 Understand plug-in on standby databases
- 3 Finetune your container database

Lunch Break

We start again at 13:30

Data Pump Top Tips

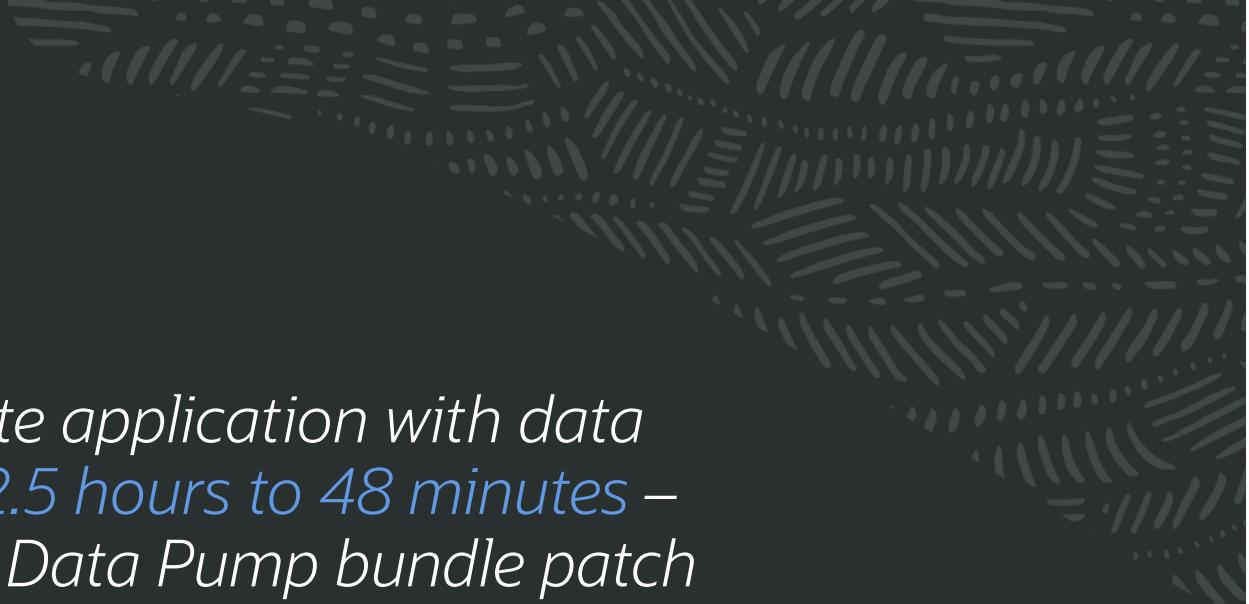
Supercharge data loading/unloading



Always use the Data Pump Bundle Patch



More than 230 functional
and performance fixes



Importing a complete application with data drops from almost 2.5 hours to 48 minutes – by just applying the Data Pump bundle patch

A global provider of financial services



Apply the Data Pump Bundle Patch
without downtime

DEMO

Apply Data Pump Bundle Patch



[Watch on YouTube](#)

Data Pump Bundle Patch

- [Data Pump Recommended Proactive Patches For 19.10 and Above \(Doc ID 2819284.1\)](#)
- The patch is not RAC rolling installable
 - The patch is non-binary online installable
 - Apply while the database instance is running
 - Don't use Data Pump or `DBMS_METADATA`



Use parallel and multiple dump files

--Apply parallelism by simply specifying a degree

```
expdp ... parallel=8
```

--Use different parallel degree on import

```
impdp ... parallel=32
```



Oracle Cloud Infrastructure

Number of ECPUs / 4

Number of OCPUs



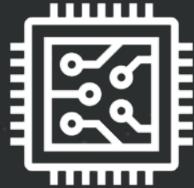
On-prem (x86-64)

2 x physical cores



On-prem (other)

Depends



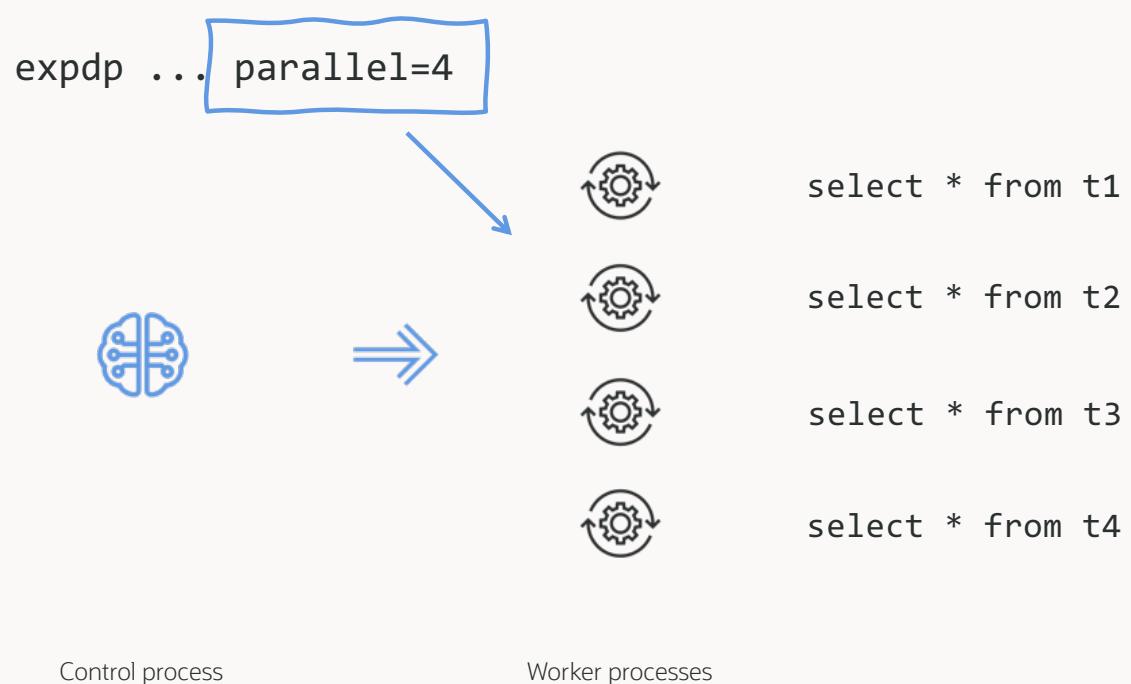
Use at least 32 ECPUs when importing into Autonomous AI Database

- Use the elasticity in Autonomous AI Database and scale even higher

Parallel Architecture

```
expdp ... parallel=4
```

Parallel Architecture



Parallel Architecture

expdp ... parallel=4



Control process



Worker processes

select /*+ parallel(2) */ * from t1



select * from t2



select * from t3



select * from t4

Worker 4 goes idle

--Use %L to allow multiple dump files

```
expdp ... parallel=8 dumpfile=exp%L.dmp
```

--Split dump files into minor files for easier transport
expdp ... parallel=8 dumpfile=exp%L.dmp filesize=10000M

- After export, store a checksum in the dump file.
- Detects in-flight corruption or alteration.
- Specify other algorithms using checksum_algorithm parameter.

`expdp ... checksum=yes`

`impdp ... verify_checksum=yes`
`verify_only=yes`



For best protection against dump file tampering, use encrypted dump files

- Requires Advanced Security Option

- Protect your dump files from alteration by using encryption
- Creating an encrypted dump file requires Advanced Security Option

```
expdp ... encryption=all encryption_algorithm=AES256
```

- Protect your dump files from alteration by using encryption
- Creating an encrypted dump file requires Advanced Security Option

```
expdp ... encryption=all encryption_algorithm=AES256
```

DEMO

Encrypted exports

 [Watch on YouTube](#)



Transportable jobs can use parallel
in Oracle Database 21c

- Any transportable jobs can now run in parallel
- Parallel unload/load of metadata provide a significant performance boost

```
expdp ... full=y transportable=always parallel=16
```

```
expdp ... tablespace=<list> parallel=16
```

```
impdp ... parallel=16
```

Parallel Transportable Benchmark

Oracle E-Business Suite database
600.000+ objects

| | |
|-------------------|---------------|
| Export parallel 1 | 2h 2m |
| Import parallel 1 | 6h 44m |
| Total | 8h 46m |

| | |
|--------------------|---------------|
| Export parallel 16 | 1h 8m |
| Import parallel 16 | 1h 23m |
| Total | 2h 31m |



Speed up imports by using
NOVALIDATE constraints

A Constraint Can Be

VALIDATED

All data in the table obeys the constraint.
The database guarantees that data is good.

NOT VALIDATED

All data in the table *may* obey the constraint.
The database *does not know* if data is good.

ii

Most constraints are **VALIDATED**



On import, Data Pump creates constraints
in the same state as in the source

--Example of which commands Data Pump import might execute as part of an import

```
create table sales ( .... );
```

Recursive full table scan

```
insert into sales as select ... ;
```

Recursive full table scan

```
alter table sales add constraint c_sales_1 check (c1 in (0,1)) enable validate;
```

```
alter table sales add constraint c_sales_2 check (c2 in ('A','B')) enable validate;
```

```
alter table sales add constraint c_sales_3 check (c3 > 0) enable validate;
```

Recursive full table scan

-- Add constraints with NOVALIDATE keyword regardless of state in source database
-- Significantly speeds up add constraints for larger tables

`impdp ... transform=constraint_novalidate:y`

```
--Transforming constraints to NOVALIDATE to speed up import
```

```
alter table sales add constraint c_sales_1 check (c1 in (0,1)) enable novalidate;  
alter table sales add constraint c_sales_2 check (c2 in ('A','B')) enable novalidate;  
alter table sales add constraint c_sales_3 check (c3 > 0) enable novalidate;
```

No full table scan

```
--Transforming constraints to NOVALIDATE to speed up import
```

```
alter table sales add constraint c_sales_1 check (c1 in (0,1)) enable novalidate;  
alter table sales add constraint c_sales_2 check (c2 in ('A','B')) enable novalidate;  
alter table sales add constraint c_sales_3 check (c3 > 0) enable novalidate;
```

Database validates new rows

Benchmark, 1 billion rows

Importing VALIDATE constraints

```
10-AUG-24 00:32:28.716: W-1 Processing object type TABLE_EXPORT/TABLE/TABLE_DATA
10-AUG-24 00:36:42.762: W-1 . . imported "FUSION"."hwr_topic_t1" 151.2 GB 1044625000 rows in 254 seconds using external_table
10-AUG-24 00:45:41.226: W-1 Processing object type TABLE_EXPORT/TABLE/CONSTRAINT/CONSTRAINT
10-AUG-24 00:55:35.787: W-1      Completed 7 CONSTRAINT objects in 594 seconds
```

Importing NOVALIDATE constraints

```
10-AUG-24 00:14:56.050: W-1 Processing object type TABLE_EXPORT/TABLE/TABLE_DATA
10-AUG-24 00:19:10.311: W-1 . . imported "FUSION"."hwr_topic_t1" 151.2 GB 1044625000 rows in 254 seconds using external_table
10-AUG-24 00:29:20.841: W-1 Processing object type TABLE_EXPORT/TABLE/CONSTRAINT/CONSTRAINT
10-AUG-24 00:29:21.101: W-1      Completed 7 CONSTRAINT objects in 1 seconds
```

Real-World Example, Importing 180G Schema

| | |
|-------------------------------|---------|
| Starting point | 44m 50s |
| Adding Data Pump Bundle Patch | 36m 53s |
| Using NOVALIDATE constraints | 7m 25s |

o
i

NOVALIDATE constraints prevent
the optimizer from certain query rewrites

Validate constraints after import, or even after go-live

- Still requires a full scan of the table
- But can use parallel query
- And no table lock!

```
alter table sales add constraint c_sales_1 check (c1 in (0,1)) enable novalidate;
```

```
-----  
----- GO LIVE -----  
-----
```

```
-- Validate constraints  
-- Optionally, use parallel query
```

```
alter session force parallel query;  
alter table sales modify constraint c_sales_1 enable validate;
```



Also available in Oracle Database 19c
via 19.27 Data Pump Bundle Patch



Also available in
Oracle Autonomous AI Database 19c

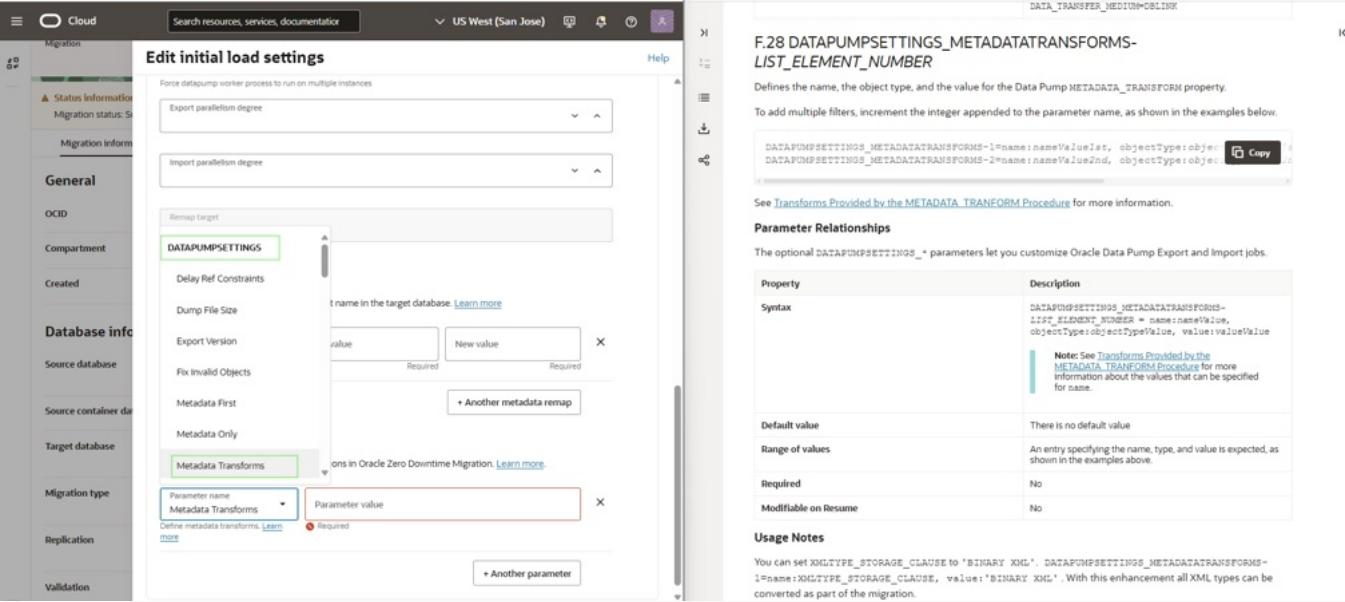
```
--Add the following to your ZDM response file to set the parameter during
--ADB migrations using ZDM
DATAPUMPSETTINGS_METADATATRANSFORMS-1=name:CONSTR_NOVALIDATE,value:1
```

NOVALIDATE Constraints

- 19.26 and earlier also requires patch 37280692
- Allow optimizer to perform query rewrites even with NOVALIDATE constraints
 - using QUERY REWRITE INTEGRITY
- Data Pump always validates certain constraints:
 - On DEFAULT ON NULL columns
 - Used by a reference partitioned table
 - Used by a reference partitioned child table
 - Table with Primary key OID
 - Used as clustering key on a clustered table

NOVALIDATE Constraints

- Available in Database Migration Service (DMS)
- In the OCI console, Advanced parameters> DATAPUMPSETTINGS: Metadata Transforms



The image shows two screenshots of the Oracle Cloud Infrastructure (OCI) console. The left screenshot is titled 'Edit initial load settings' and shows the 'DATAPUMPSETTINGS' section. It includes fields for 'Export parallelism degree', 'Import parallelism degree', 'Remap target', and 'Metadata Transforms'. The 'Metadata Transforms' section is expanded, showing a 'Parameter name' dropdown set to 'Metadata Transforms' and a 'Parameter value' input field. The right screenshot is a detailed view of the 'F.28 DATAPUMPSETTINGS_METADATATRANSFORMS-LIST_ELEMENT_NUMBER' parameter. It provides syntax examples: `DATAPUMPSETTINGS_METADATATRANSFORMS-1=name:nodeValue1, objectType:objectType1` and `DATAPUMPSETTINGS_METADATATRANSFORMS-2=name:nodeValue2nd, objectType:objectType2`. It also includes a 'Copy' button, a note about using the `METADATA_TRANSFORM` procedure, and a table of properties and their descriptions. The table includes columns for 'Property', 'Syntax', 'Description', 'Default value', 'Range of values', 'Required', 'Modifiable on Resume', and 'Usage Notes'.

| Property | Description |
|----------------------|--|
| Syntax | <code>DATAPUMPSETTINGS_METADATATRANSFORMS-LIST_ELEMENT_NUMBER = name:nodeValue, objectType:objectTypeValue</code> |
| Note | See <code>Transforms Provided by the METADATA_TRANSFORM Procedure</code> for more information about the values that can be specified for name. |
| Default value | There is no default value. |
| Range of values | An entry specifying the name, type, and value is expected, as shown in the examples above. |
| Required | No |
| Modifiable on Resume | No |
| Usage Notes | You can set <code>XMLTYPE_STORAGE_CLAUSE</code> to 'BINARY XML'. <code>DATAPUMPSETTINGS_METADATATRANSFORMS-1=name:XMLTYPE_STORAGE_CLAUSE, value:'BINARY XML'</code> . With this enhancement all XML types can be converted as part of the migration. |



Use index size to determine parallel degree on index creation

Index Creation

```
impdp ... parallel=16
```

Before 12.1

Worker 1

```
CREATE INDEX .... PARALLEL 16
```

Really good for few big indexes

Index Creation

```
impdp ... parallel=16
```

From 12.1

| | |
|-----------|------------------------------|
| Worker 1 | CREATE INDEX PARALLEL 1 |
| Worker 2 | CREATE INDEX PARALLEL 1 |
| ... | CREATE INDEX PARALLEL 1 |
| Worker 16 | CREATE INDEX PARALLEL 1 |

Really good for many small indexes



Index Creation

```
impdp ... parallel=16
```

From 26

| | |
|----------|------------------------------|
| Worker 1 | CREATE INDEX PARALLEL 1 |
| Worker 2 | CREATE INDEX PARALLEL 8 |
| Worker 3 | CREATE INDEX PARALLEL 4 |
| Worker 4 | CREATE INDEX PARALLEL 3 |

The best of both worlds



How Data Pump Create Indexes

- 1 Calculate the optimal parallel degree
- 2 Create indexes

How Data Pump Create Indexes

1 Calculate the optimal parallel degree

- Always parallel 1 when a table is less than 150 MB
- Customizable via `INDEX_THRESHOLD`
- Get optimal parallel degree using `EXPLAIN PLAN`

```
SQL> explain plan for create index i1 on t1(c1) parallel;
```

Explained.

```
SQL> explain plan for create index i1 on t1(c1) parallel;
```

Explained.

```
SQL> select * from table(dbms_xplan.display(format => 'ALL'));
```

...

Note

- automatic DOP: Computed Degree of Parallelism is 4 because of degree limit
- estimated index size: 655K bytes

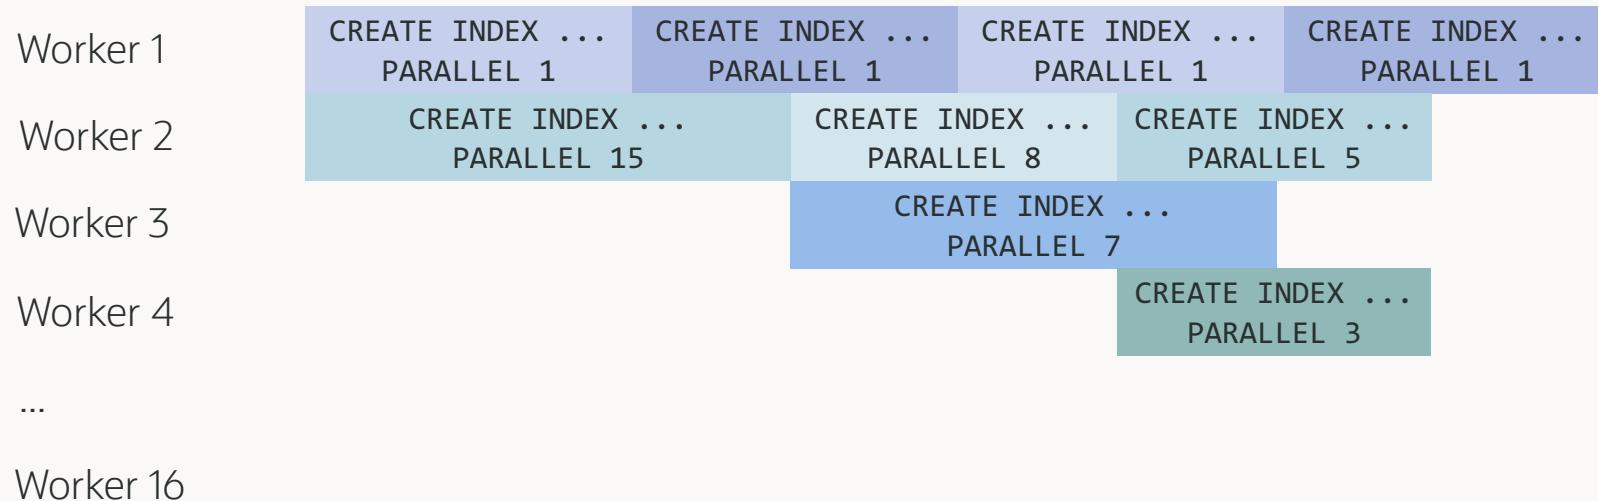
How Data Pump Create Indexes

2 Create indexes

- One worker creates small indexes (parallel 1) in large batches
- The next worker starts with the biggest index (measured by optimal parallel degree)

How Data Pump Create Indexes

impdp ... parallel=16



Importing with former index method

```
10-MAY-25 16:18:55.130: W-12 Processing object type SCHEMA_EXPORT/TABLE/INDEX/INDEX
10-MAY-25 16:36:46.902: W-30      Completed 480 INDEX objects in 1071 seconds
```

Importing with new index method

```
10-MAY-25 16:47:50.267: W-4 Processing object type SCHEMA_EXPORT/TABLE/INDEX/INDEX
10-MAY-25 16:59:17.006: W-3      Completed 480 INDEX objects in 686 seconds
```

Index Creation



Requires:

- Oracle AI Database 26ai
- Oracle Database 23.8 plus Data Pump Bundle Patch
- Oracle Database 19.26 plus Data Pump Bundle Patch



Also available in Oracle Database 19c
via 19.26 Data Pump Bundle Patch



Convert time zone data on import

Time Zone File Version Check



Source

Version 43

Target

Version 42

```
create table t1 (
  ...
  c1 timestamp with timezone
  ...
)
```

Import: Release 19.0.0.0.0 - Production on Sun Sep 30 06:17:06 2025

Version 19.27.0.0.0

ORA-39002: invalid operation

ORA-39405: Oracle Data Pump does not support importing from a source database with TSTZ version 43 into a target database with TSTZ version 42.

DEMO

Convert time zone data on import

 [Watch on YouTube](#)

Time Zone Data Convert



- If needed, Data Pump converts **TIMESTAMP WITH TIMEZONE** to the target time zone file version.
- This works if the target time zone file is higher or lower than the source.
- Expect a small overhead for the conversion.
- Requires 19.27 including Data Pump Bundle Patch.



Also available in Oracle Database 19c
via 19.27 Data Pump Bundle Patch



Copy the dump files
before the export completes

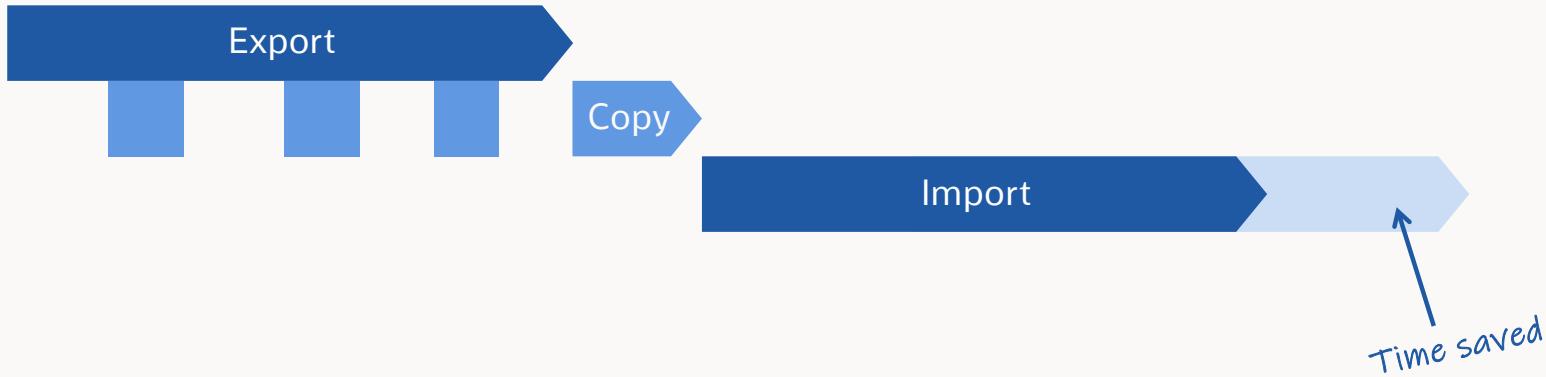
- Oracle AI Database 26ai



Export

Copy

Import



```
dbms_datapump.get_status(l_handle, l_stsmask, NULL, l_job_state, l_status);
```

```
dbms_datapump.get_status(l_handle, l_stsmask, NULL, l_job_state, l_status);

for i in l_status.job_status.files.first..l_status.job_status.files.last() loop

end loop;
```

```
dbms_datapump.get_status(l_handle, l_stsmask, NULL, l_job_state, l_status);

for i in l_status.job_status.files.first..l_status.job_status.files.last() loop

  l_dump_file := l_status.job_status.files(i);

  if (l_dump_file.file_bytes_written = l_dump_file.file_size) then
    dbms_output.put_line('DONE: ' || l_dump_file.file_name);
  end if;

end loop;
```

Copy Dump Files



- Blog post: [Copy Data Pump Files Before the End of the Export](#)



Use diagnostics views to assist in troubleshooting

- Oracle AI Database 26ai

```
SQL> select waiting_session, event, dp_state_in_wait
  from v$datapump_sessionwait_info;
```

| WAITING_SESSION | EVENT | DP_STATE_IN_WAIT |
|-----------------|----------------------|------------------|
| 10 | direct path sync | WAITING |
| 77 | log buffer space | WAITING |
| 191 | log buffer space | WAITING |
| 428 | enq: TT - contention | WAITING |

Diagnostic Views



- New views available in Oracle AI Database 26ai:
 - `v$datapump_process_info`
 - `v$datapump_sessionwait_info`
 - `v$datapump_processwait_info`



Key Learnings



- 1 Use Data Pump Bundle Patch
- 2 Use NOVALIDATE constraints
- 3 Create indexes faster

Autonomous AI Database

A migration approach

What's the story?

Flashback to October 2017



The idea?

A mostly self-managed database environment,
taking care on many tasks

KEYNOTE PRESENTATION

Oracle OpenWorld San Francisco 2017



ry E sdN

OOOW 2017

ORACLE

A photograph of a man from a side-on perspective, sitting at a desk. He is wearing a grey t-shirt and has his hand on his head in a gesture of stress or despair. In front of him is a computer keyboard and a monitor that is mostly out of frame. The lighting is low, creating a moody atmosphere.

Will the DBA's be fired?

<https://www.complexsql.com/oracle-18c-impact-on-dbas/>

A silhouette of a Tyrannosaurus Rex is positioned against a dark, star-filled background. The silhouette is partially visible, with its head and front legs on the left and its body and tail extending towards the bottom right. The background is filled with numerous small stars of varying brightness.

"It's that sort of attitude that has turned some
DBAs into inflexible dinosaurs.
You've got to evolve or die, people!"

—
Tim Hall

<https://oracle-base.com/blog/2017/10/02/oracle-autonomous-database-and-the-death-of-the-dba/>

Autonomous AI Database – Where?

Public cloud

Autonomous AI Database

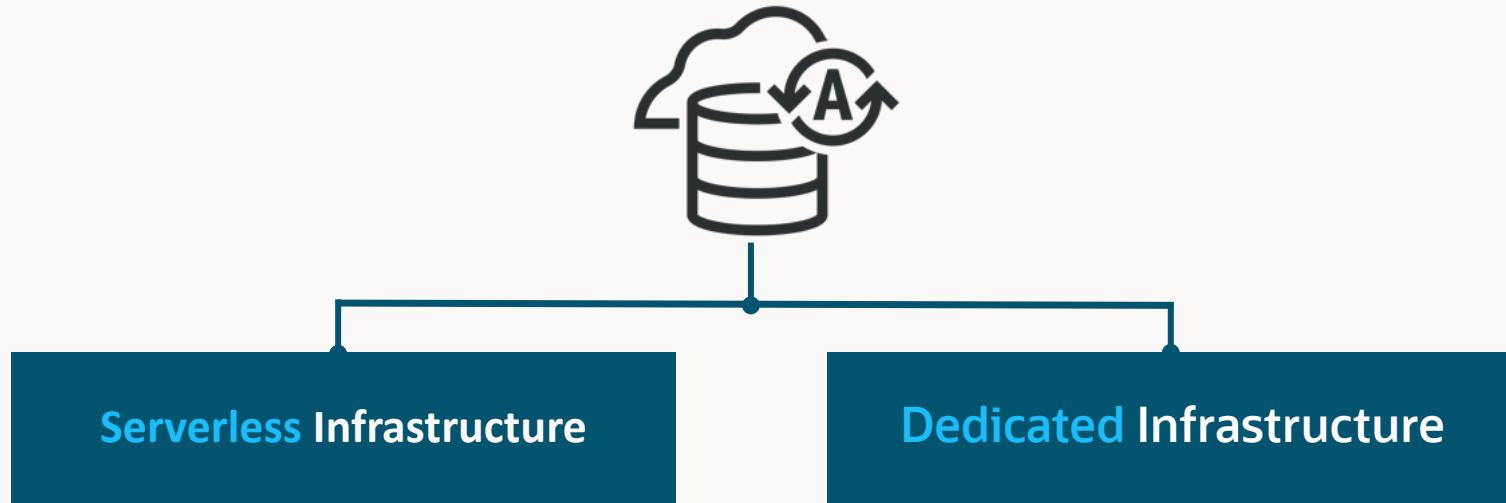


Cloud@Customer

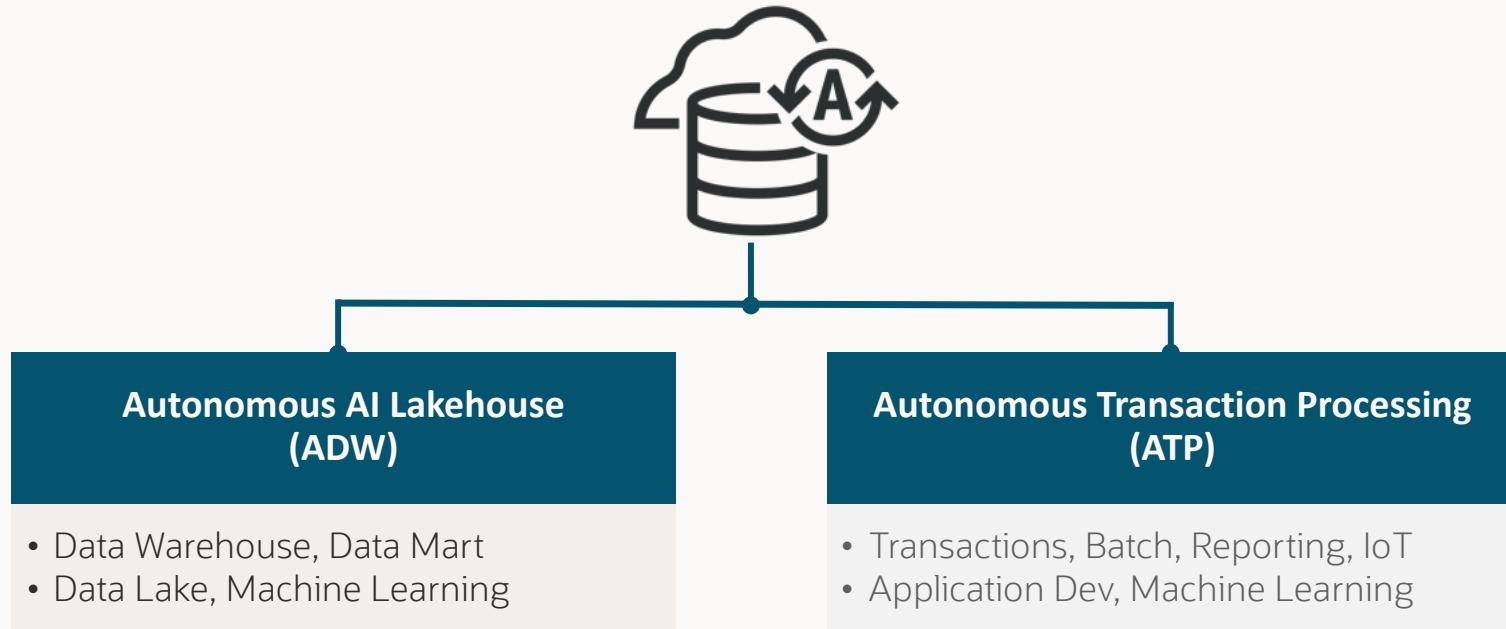
Autonomous AI Database in a VM environment



One Autonomous AI Database – Two Deployment Choices



One Autonomous AI Database – Workload Choices



Migration Planning

No migration without a proper runbook



Photo by [bert b](#) on [Unsplash](#)

Estate Modernization

But not every database is a great candidate for ADB



Migration to Autonomous AI Database is always a **logical** migration

- Move the data, not the database

Tools out-of-the-box



SQL Developer Web

Web-based Function rich,
low code development env
No client software needed



Oracle REST Data Services

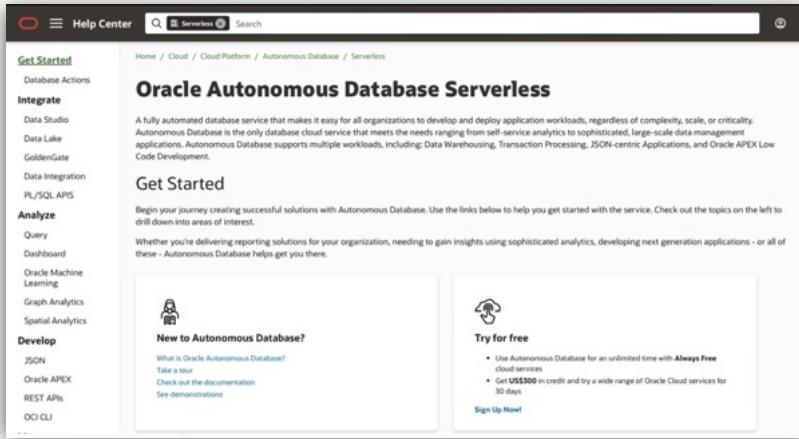
Ability to REST enable a
schema and autogenerate
REST endpoints for tables,
views, and procedures



APEX

Execute SQL and PL/SQL
Build Data Models,
generate DDL statements
Monitor and manage the DB

Essentials

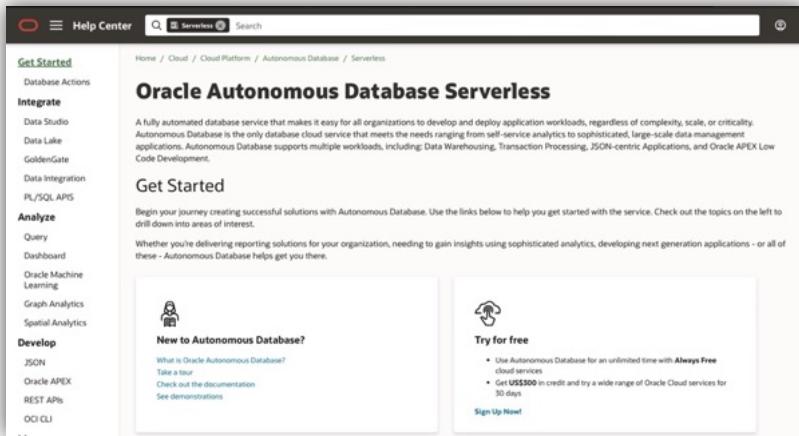


The screenshot shows the Oracle Autonomous Database Serverless Help Center. The page title is "Oracle Autonomous Database Serverless". The main content area is titled "Get Started" and contains a brief introduction to the service. Below this, there are two sections: "New to Autonomous Database?" and "Try for free". The "New to Autonomous Database?" section includes links to "What's Oracle Autonomous Database?", "Take a tour", "Check out the documentation", and "See demonstrations". The "Try for free" section includes a list of benefits: "Use Autonomous Database for an unlimited time with **Always Free** cloud services", "Get **US\$300** in credit and try a wide range of Oracle Cloud services for 30 days", and a "Sign Up Now!" button. The left sidebar lists various database actions and services, including Get Started, Database Actions, Integrate, Analyze, Develop, and Oracle Machine Learning.

Outside the database,
check [OCI Documentation](#)

Example: Deploy, start, stop, scale

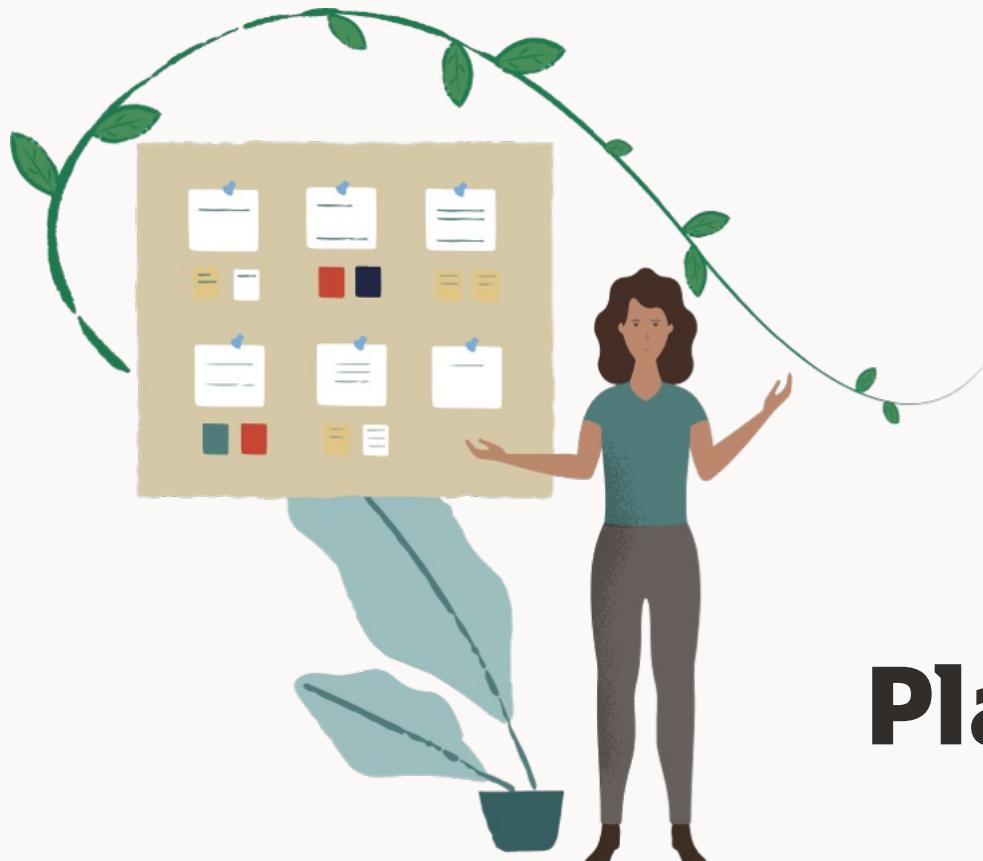
Essentials



The screenshot shows the Oracle Autonomous Database Serverless documentation page. The top navigation bar includes 'Help Center', a search bar, and a 'Serverless' link. The main content area is titled 'Oracle Autonomous Database Serverless' and describes it as a fully automated database service. It highlights that Autonomous Database is the only database cloud service that meets the needs ranging from self-service analytics to sophisticated, large-scale data management applications. The page includes sections for 'Get Started', 'Integrate', 'Analyze', and 'Develop', each with sub-links. A 'Try for free' button is prominently displayed on the right.

Inside the database,
check [Database Documentation](#)

Example: Schema, capabilities, connecting



Planning



How do we migrate our 500 databases to Oracle Autonomous AI Database?

- And which ones are good candidates?

Getting an Overview

1

Estate Explorer



2

Cloud Premigration
Advisor Tool



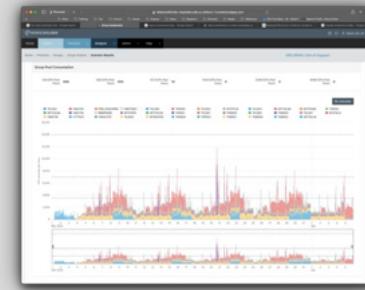
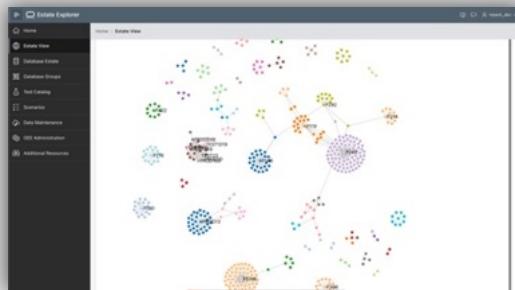
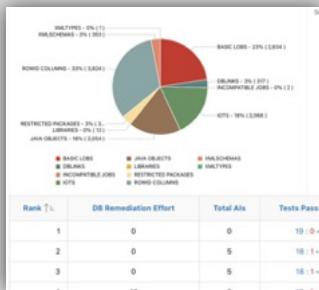
3

Cloud Migration
Advisor

Getting an Overview

1

Estate Explorer



Getting an Overview

1

Estate Explorer



2

Cloud Premigration
Advisor Tool



3

Cloud Migration
Advisor



Analyze 1000's of databases in
just a few hours



Provide a detailed TCO to
compare on-premises and cloud



View innovative visualizations
and detailed reports



Optimize your Autonomous AI
Databases using Elastic Pools

Getting an Overview

1

Estate Explorer



2

Cloud Premigration
Advisor Tool



3

Cloud Migration
Advisor



Getting an Overview

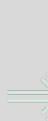
1

Estate Explorer



2

Cloud Premigration
Advisor Tool



3

Cloud Migration
Advisor

▼ Premigration Advisor Check Details List

Source Database

Expand All Close All

▼ Action Required (2 checks)

▼ OGG Minimal Supplemental Logging Not Enabled

Description: Minimal supplemental logging is not enabled on the Database.

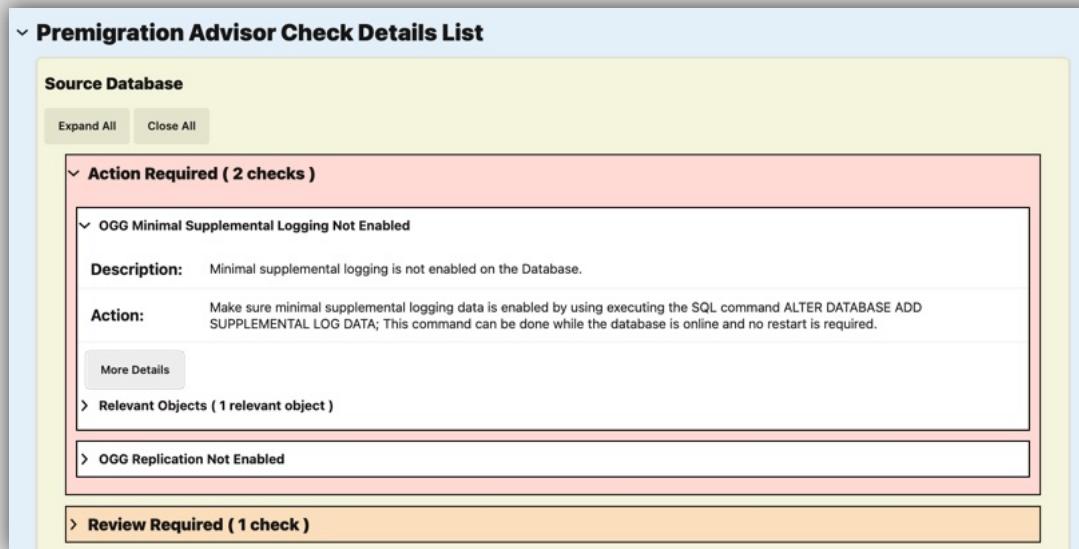
Action: Make sure minimal supplemental logging data is enabled by using executing the SQL command ALTER DATABASE ADD SUPPLEMENTAL LOG DATA; This command can be done while the database is online and no restart is required.

More Details

➤ Relevant Objects (1 relevant object)

➤ OGG Replication Not Enabled

➤ Review Required (1 check)



Getting an Overview

1

Estate Explorer



2

Cloud Premigration
Advisor Tool



3

Cloud Migration
Advisor



Getting an Overview

1

Estate Explorer



2

Cloud Premigration
Advisor Tool



3

Cloud Migration
Advisor



The screenshot shows the Oracle Cloud Migration Advisor (CMA) interface. The top navigation bar includes a user icon, a search bar with the text "mike.dietrich@oracle.com", and a dropdown menu. The main content area features a green rocket icon and the title "Oracle Cloud Migration Advisor". Below the title, a "Welcome to the Cloud Migration Advisor (CMA)" message is displayed, followed by a paragraph about the tool's purpose. A "Guided Mode" section lists two bullet points: "Which databases can be **most easily migrated** to Oracle Autonomous Database, or" and "What is the **best migration method** to move chosen databases to a desired Oracle Cloud platform?". A "For more options, you can" section lists two bullet points: "Create and configure migration scenarios using **Create Project**, or" and "Access your existing migration projects and scenarios with **My Available Projects**". At the bottom, a message says "Now it is time to start - let's move to the Oracle Cloud!" and a footer bar contains three buttons: "Guided Mode" (blue), "My Available Projects" (light blue), and "Create New Project" (green).

Getting an Overview

Estate Explorer ⇒ Cloud Premigration Advisor Tool ⇒ Cloud Migration Advisor

- OCI Database Migration Service
- Zero Downtime Migration
- Autonomous Migration Automation
- Data Pump
- GoldenGate
- O2O / 000



Cloud Premigration Advisor Tool



Evaluate an Oracle AI Database for compatibility with Oracle Autonomous AI Database

- Use Cloud Premigration Advisor Tool (CPAT)

Overview



Download CPAT from MOS Note: 2758371.1

 **Patch 32613591: Cloud Premigration Advisor Tool (CPAT) for version 11.2.0.4 and Higher**

Last Updated 11-Feb-2025 17:31 (12 days ago)

| | | | |
|----------|-----------------------------------|-----------------|----------|
| Product | Oracle Database Upgrade Assistant | Size | 8.6 MB |
| Release | Oracle 11.2.0.4.8 | Download Access | Software |
| Platform | Generic Platform | Classification | General |
| | | Patch Tag | |

Bugs Resolved by This Patch
List of bugs fixed is not available. Consult the Readme.

[View Related Knowledge to this Patch](#)

Release Oracle 11.2.0.4.8
Platform Generic Platform
Language American English

[Read Me](#) [Download](#)  [Add to Plan](#) 
[Analyze with OPatch...](#)

All-time Downloads 50 [View Trends](#)

 [Discuss this patch in the community](#)

```
# One or more report formats separated by spaces
# json html text

./premigration.sh \
  --connectstring jdbc:oracle:thin:@<host>:<port>/<service> \
  --username CPAT_CHECK \
  --pdbname PDB_COMPLEX \
  --schemas appuser,reportuser \
  --outdir /home/oracle/cpat-db \
  --targetcloud atps \
  --migrationmethod goldengate \
  --reportformat html
```



That's a lot of options.
Help me out, please!

CPAT COMPOSER

<https://macsdata.com/oracle/cpat-composer>

- Free to use
- Available online
- Not an official Oracle tool
- Created by Marcus Doeringer
Migration Specialist @Oracle





CLOUD PREMIGRATION ADVISOR TOOL (CPAT) COMPOSER

Version: 25.2.0 (Default)

Template: SA: Recommended Settings



Input Form

Operating System

Select the operating system to run CPAT

Linux
Use for Linux operating system

Windows
Use for Windows operating system

CPAT Mode

Select the mode you want to run CPAT

Source Analysis
Run CPAT on the source database for analysis

Target Properties
Generate a properties file from the target database

Copyright © 2026, Oracle and/or its affiliates

Output

Status Info

All required options set

Recommended Options have been applied

Recommended Options

Analysis Mode: SCHEMA File Prefix

Command

Parameter File

Command Line

```
./premigration.sh --connectstring 'jdbc:oracle:thin:@myhost:1521:ORCL' --username sys --targetcloud ATPS --migrationmethod DATAPUMP --reportformat HTML JSON TEXT --full --zip
```

Cloud Premigration Advisor Tool (CPAT) Report

CPAT Version: 25.2.1-1

Version Date: Feb 17, 2025

Days Since Last CPAT Update: 38 days

Table of Contents

- [Premigration Advisor Report Summary](#)
- [Report Details](#)
- [Report Analysis Notes](#)
- [Source Database Details](#)
- [Source Database Version Information](#)
- [Source Database Patch Information](#)
- [Source Database Redo Information](#)
- [Source Database Supplemental Information](#)
- [Source Database Schema Summary Information](#)
- [Premigration Advisor Check Details List](#)
- [Report Legend](#)

[Expand All](#)[Close All](#)

▼ Premigration Advisor Report Summary

Report Result

Action Required

Number of schemas analyzed:

4

- Premigration Advisor Check Details List
- Report Legend

[Expand All](#)[Close All](#)

▼ Premigration Advisor Report Summary

Report Result

Action Required

Number of schemas analyzed:

4

List of schemas analyzed:

[PDBUSER, CO, HR, MYAPP]

▼ Report Results Overview

| Source Database | | Target Database | | Migration Method | | Additional Tasks | |
|------------------|-----------|------------------|-----------|------------------|----------|------------------|-----------|
| Action Required | 2 | Action Required | 10 | Action Required | 1 | Action Required | 0 |
| Review Required | 1 | Review Required | 7 | Review Required | 2 | Review Required | 0 |
| Review Suggested | 2 | Review Suggested | 4 | Review Suggested | 1 | Review Suggested | 4 |
| Passed | 16 | Passed | 17 | Passed | 4 | Passed | 16 |

[Return to Table of Contents](#)

▼ Report Details

Report Details

| | |
|-------------------------|---|
| CPAT | |
| Application Version: | 25.2.1-1 |
| Report Generated On: | Sat Feb 22 20:59:44 UTC 2025 |
| Analysis Property File: | premigration_advisor_analysis.properties |
| Analysis Mode: | FULL |
| Target Cloud Type: | ALL |
| Migration Method(s): | [DATAPUMP, DATAPUMP_DBLINK, GOLDENGATE] |
| Command Line Options: | --connectstring jdbc:oracle:thin:@dbsystemaz:1521/pdb_complex.sub07021512520.upgradeteam.oraclevcn.com --targetcloud ALL --username SYS --sysdba --analysisprops premigration_advisor_analysis.properties --outdir /home/oracle/cpat_22_feb_2025 --logginglevel FINE --migrationmethod ALL --reportformat JSON HTML TEXT --resultlevel R0 --zip --gatherdetails ALL |

[More Details](#)

More Details

[Return to Table of Contents](#)

› **Report Analysis Notes**

› **Source Database Details**

› **Source Database Version Information**

› **Source Database Patch Information**

› **Source Database Redo Information**

› **Source Database Supplemental Information**

› **Source Database Schema Summary Information**

› **Premigration Advisor Check Details List**

Source Database

Source Database Details

Source Cloud Vendor: Oracle Cloud Infrastructure (Database)

Source Database Host Name: dbsystemaz

Source Oracle SID: ORCL

Source Database Created Date: Fri Jan 24 22:23:51 UTC 2025

Source Database DBID: 1719058167

Source Database Unique Name: ORCL_5tr_iad

Source Instance Name: ORCL

Source Database Name: ORCL

Source Database Username: SYS

Source Database Port String: x86_64/Linux 2.4.xx

Source Database Platform ID: 13

Source Database Container Name: PDB_COMPLEX

Source DB Block Size in KB: 8

Source DB Combined Size of DATA, TEMP, LOG, and CONTROL File Usage in GB: 5.044

Source DB Size of DATA File Usage in GB: 1.856

Source DB Size of TEMP File Usage in GB: 0.17

More Details

[Return to Table of Contents](#)

› **Report Analysis Notes**

› **Source Database Details**

› **Source Database Version Information**

› **Source Database Patch Information**

› **Source Database Redo Information**

› **Source Database Supplemental Information**

› **Source Database Schema Summary Information**

› **Premigration Advisor Check Details List**

Source Database

▼ Action Required (2 checks)

▼ OGG Minimal Supplemental Logging Not Enabled

Description: Minimal supplemental logging is not enabled on the Database.

Action: Make sure minimal supplemental logging data is enabled by using executing the SQL command ALTER DATABASE ADD SUPPLEMENTAL LOG DATA; This command can be done while the database is online and no restart is required.

More Details

➤ Relevant Objects (1 relevant object)

▼ OGG Replication Not Enabled

Description: ENABLE_GOLDENGATE_REPLICATION init.ora parameter is not set.

Action: Make sure ENABLE_GOLDENGATE_REPLICATION is set to TRUE by using executing the SQL command: ALTER SYSTEM SET ENABLE_GOLDENGATE_REPLICATION=TRUE SCOPE=BOTH; This command can be done while the database is online and no restart is required.

More Details

➤ Relevant Objects (1 relevant object)



The documentation has additional information on each CPAT check

[Utilities Guide, Oracle AI Database 26ai](#)



--Generates fixup scripts whenever possible
--Stores the scripts on disk for review

```
./premigration.sh ... --genfixups
```



You can run CPAT on any live database.
It is completely non-intrusive.

CPAT integration



Generate CPAT report

ZERO DOWNTIME MIGRATION

Run as part of its migration assessment

OCI DATABASE MIGRATION SERVICE

Run as part of its migration assessment

ENTERPRISE MANAGER MIGRATION WORKBENCH

Run as part of its migration assessment

SQL DEVELOPER / SQLcl

Through the MIGRATEADVISOR command



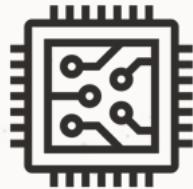
Data Pump

The simple approach



Data Pump Bundle Patch aren't yet applied in ADB Serverless (October 2025)

- You may request one-off fixes via an SR



Allocate a sufficient number of ECPUs

- 32 should be the minimum when you import



Export: PARALLEL 2x of physical cores



Import: PARALLEL=ECPU/4, or higher

- Scale up to the maximum for migrations

Most simple method: Data Pump



Datapump **with Files**



Datapump **with DB Links**



Automation

How AMA scripts ease migrations

What is AMA?

Autonomous AI Database Migration Automation (AMA)

- Simple migration solution for ADB Serverless
- Script based
- Single configuration file
- Migrates in phases
- Can act fully automated
- Not a new product, just a solution to ease migrations

An ADB-S migration is a bit like making a movie

You won't start with filming right away

You need a script book

You need to cast actors

You need a film set

You need ...

Now you can start filming your scenes

And then there's plenty of work on editing and cutting the movie



AMA Workflow



- Examination of source database (CPAT)
- Create migration directories
- Configure AMA parameter file

```
--Create migration directories  
--Copy parameter file into INPUT  
--Edit parameter file and make adjust with your values
```

```
mkdir -p /home/oracle/CPAT_MIG_SCRIPTS/INPUT  
mkdir -p /home/oracle/CPAT_MIG_SCRIPTS/OUTPUT
```

```
cp CPAT_MIGRATION_PARAMETERS.txt /home/oracle/CPAT_MIG_SCRIPTS/INPUT
```

```
vi /home/oracle/CPAT_MIG_SCRIPTS/INPUT/CPAT_MIGRATION_PARAMETERS.txt
```

Parameter File

Adjust:

- Connect strings source and target
- Data Pump encryption
- Storage (FSS or Object Store)
- Format: TAB or SCRIPT

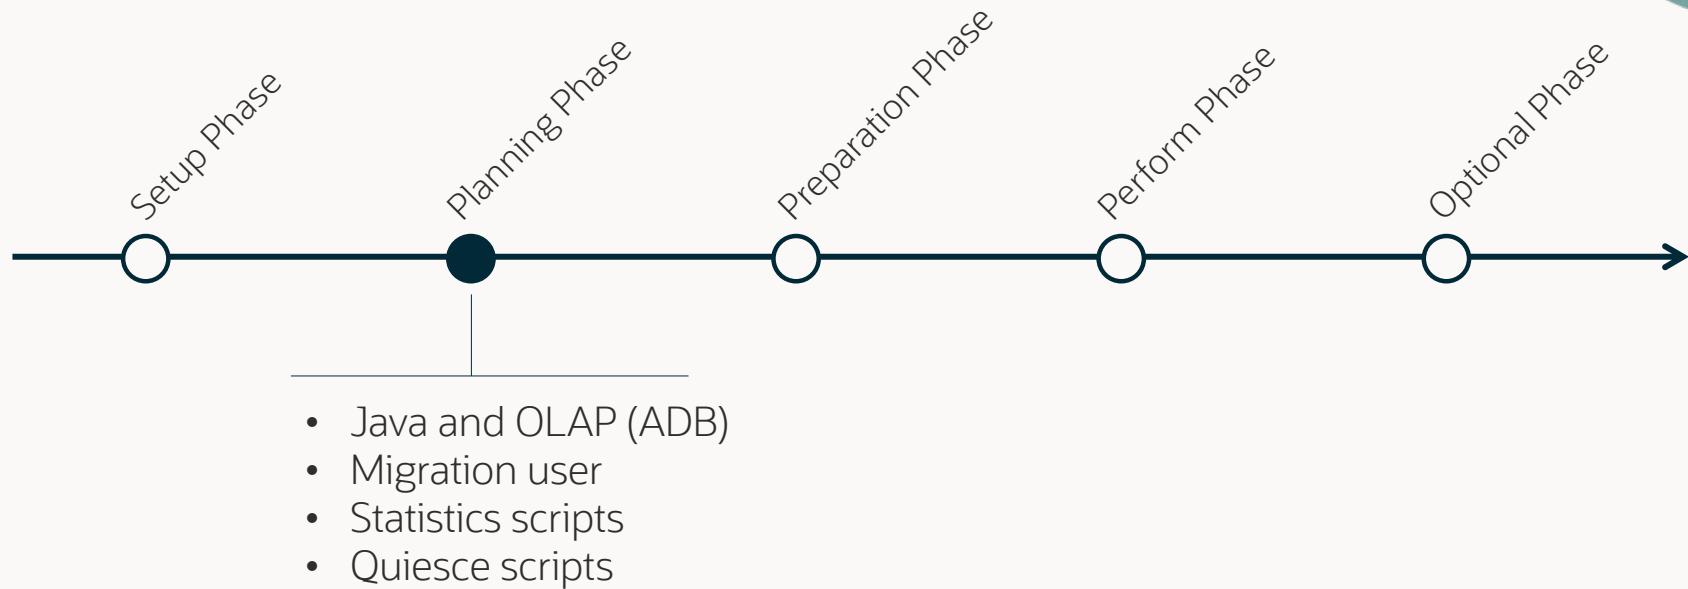
Documentation

AMA Documentation is available at request

- Documents the entire flow and all options and parameters



AMA Workflow



AMA | Planning Phase

On-Prem - Source

ADB-S - Target

Gather stats for SYS / SYSTEM

Enable OLAP / JAVA in ADB-S

Create Migration user

Enable restricted session

Set JOB_QUEUE_PROCESSES=0

```
[oracle@ephx31vm1-jlosd1 OUTPUT]$ cat __US3BLDW_MIGRATION_CONTROL_FILE.ctl
```

```
---
```

```
--- PLAN PHASE ---
```

```
--- All steps in this phase affect the source database ---
```

```
--- * You can collect the statistics or create the migration user in advance ---
```

```
--- * Get familiar with the restricted session privilege and how to prepare it ---
```

```
--- * shortly before the migration starts make sure no unwanted user is connected ---
```

```
--- to the source database, turn on restricted session and disable the scheduler ---
```

```
---
```

```
---
```

```
#### SOURCE ####
```

```
#### TARGET ####
```

```
PLAN TARGET 00001 01 ..... 00001_US3BLDW_SQL_ENABLE OLAP JAVA.sh
```

```
PLAN SOURCE 00002 01 00002_US3BLDW_SQL_OPTIONAL_SOURCE_STATS.sh
```

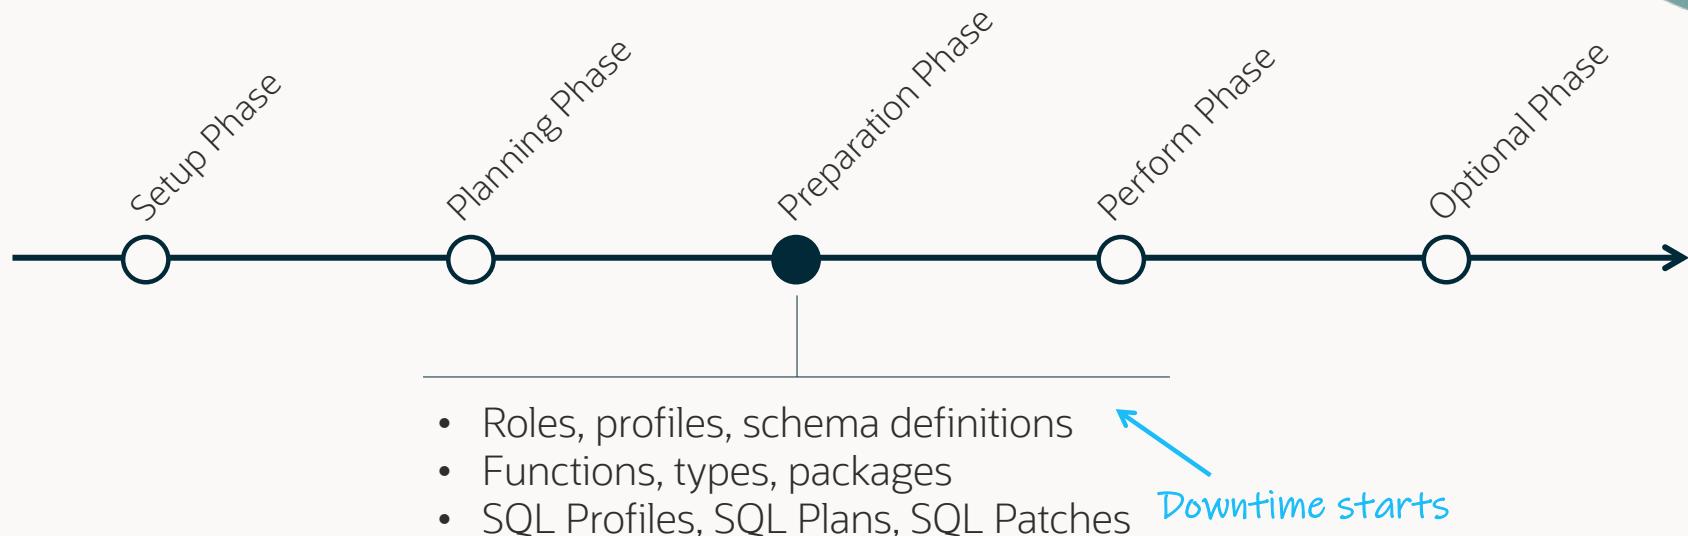
```
PLAN SOURCE 00003 01 00003_US3BLDW_SQL_CREATE_MIG_USER_SRC.sh
```

```
PLAN SOURCE 00004 01 00004_US3BLDW_SQL_SET_JOB_QUEUE_PROCESSES.sh
```

```
PLAN SOURCE 00005 01 00005_US3BLDW_SQL_ENABLE_RESTRICTED_SESSION.sh
```

```
---
```

AMA Workflow



AMA | Preparation Phase

On-Prem - Source

ADB-S - Target

| | |
|--|---|
| Collect allowed ROLES | Create ROLES |
| Collect PROFILES | Create PROFILES |
| Export schema definition | Create storage credential (NFS, Object Store) |
| Export FUNCTIONS, TYPES, PACKAGES | Import schema definition |
| Collect SQL Profiles, SQL Plans, SQL Patches | Import FUNCTIONS, TYPES, PACKAGES |
| | Granting migration privileges |
| | Alter user profiles |
| | Create SQL Profiles, SQL Plans, SQL Patches |

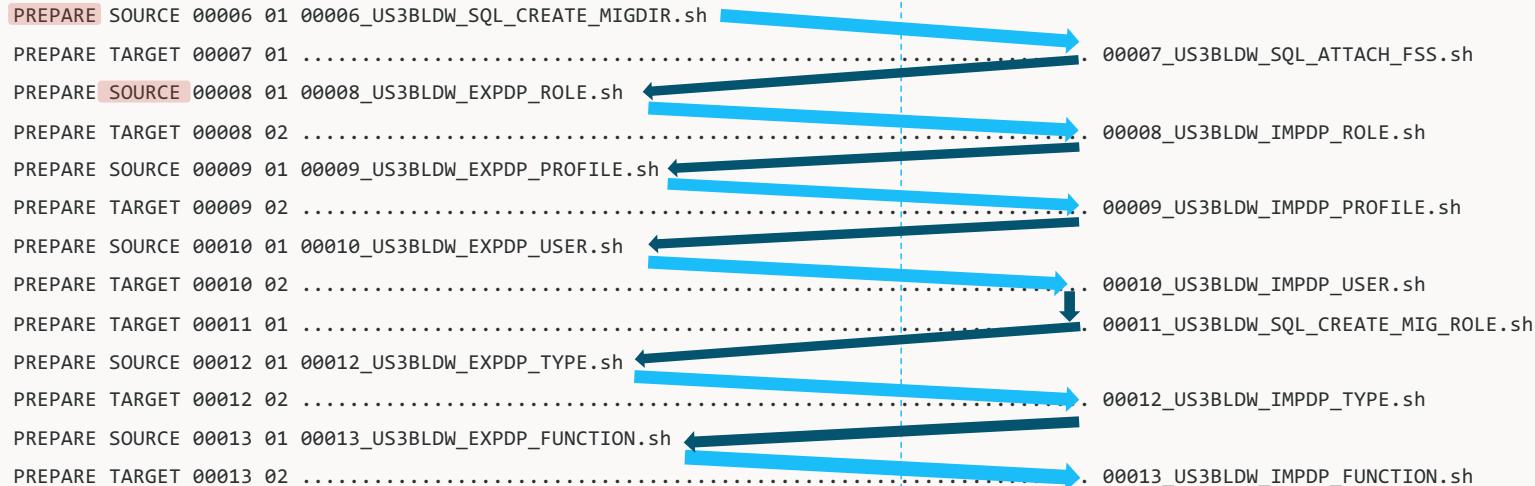
--- PREPARATION PHASE ---

--- All steps in this phase will prepare the source and target database ---

--- The scripts depend on each other, so execute in this phase one script after the other ---

SOURCE

TARGET



AMA Workflow

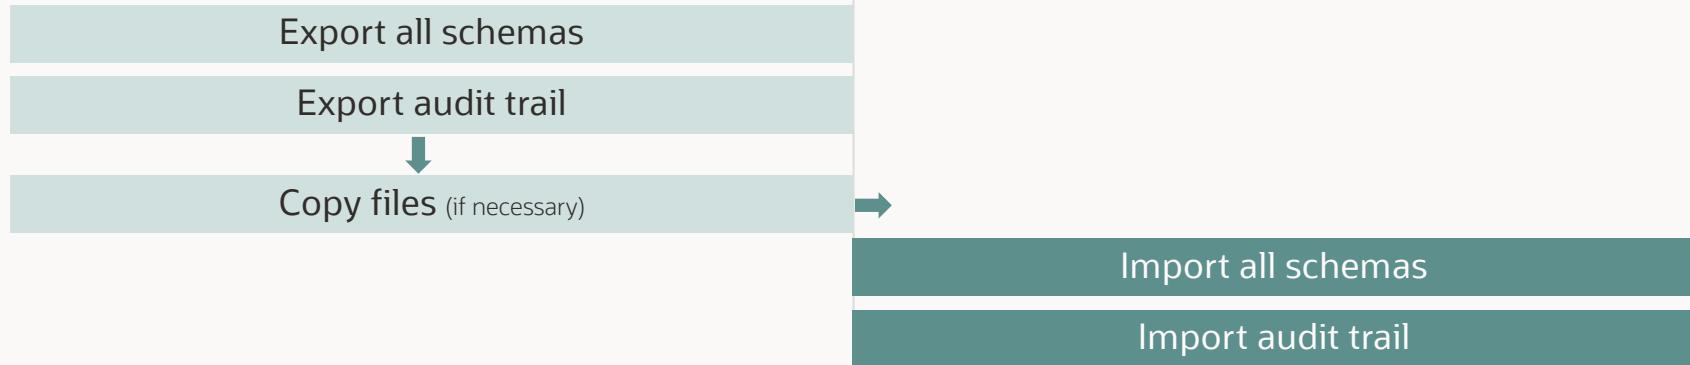


- Export schemas and audit trail
- Copy files (if necessary)
- Import schemas and audit trail

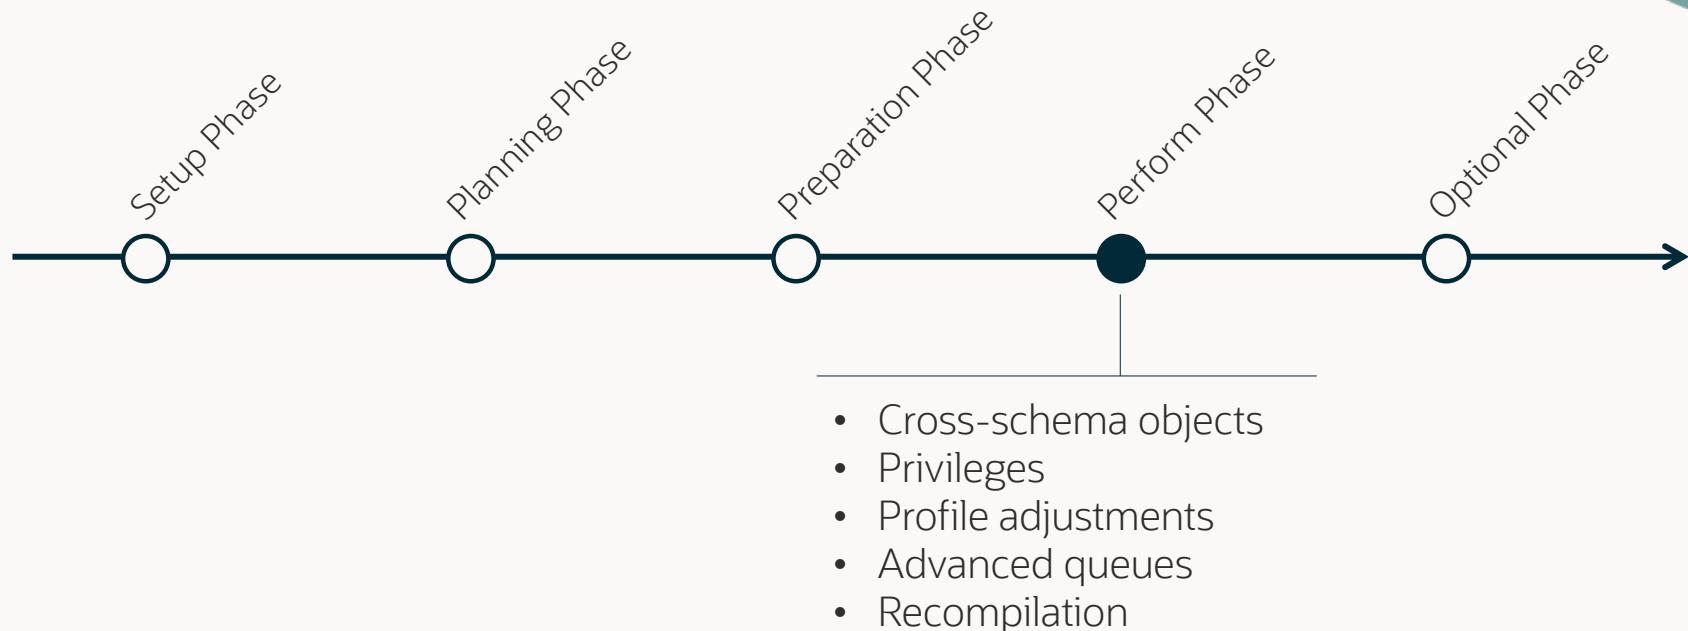
AMA | Perform Phase

On-Prem - Source

ADB-S - Target



AMA Workflow



--- **PERFORM PHASE** ---

--- Commonly in this phase nothing depends on each other (except you for example have objects that depend on objects stored in other schema) ---
--- So export jobs can be started in parallel and imports once the export finished ---

SOURCE

TARGET

| | |
|--|--|
| PERFORM SOURCE 00014 01 00014_US3BLDW_EXPDP_AUDIT_TRAILS.sh | |
| PERFORM TARGET 00014 02 | 00014_US3BLDW_IMPDP_AUDIT_TRAILS.sh |
| PERFORM SOURCE 00015 01 00015_US3BLDW_SQL_GEN_SQL_PROFILE_STAGE_TAB.sh | |
| PERFORM SOURCE 00015 02 00015_US3BLDW_EXPDP_SQL_PROFILES.sh | |
| PERFORM TARGET 00015 03 | 00015_US3BLDW_IMPDP_SQL_PROFILES.sh |
| PERFORM TARGET 00015 04 | 00015_US3BLDW_SQL_APPL_SQL_PROFILE_STAGE_TAB.sh |
| PERFORM SOURCE 00016 01 00016_US3BLDW_SQL_GEN_SQL_PATCHES_STAGE_TAB.sh | |
| PERFORM SOURCE 00016 02 00016_US3BLDW_EXPDP_SQL_PATCHES.sh | |
| PERFORM TARGET 00016 03 | 00016_US3BLDW_IMPDP_SQL_PATCHES.sh |
| PERFORM TARGET 00016 04 | 00016_US3BLDW_SQL_APPL_SQL_PATCHES_STAGE_TAB.sh |
| PERFORM SOURCE 00017 01 00017_US3BLDW_EXPDP_SCHEMA_FUSION.sh | |
| PERFORM TARGET 00017 02 | 00017_US3BLDW_IMPDP_SCHEMA_FUSION.sh |
| PERFORM SOURCE 00018 01 00018_US3BLDW_EXPDP_SCHEMA_FUSION_OCSERVER11G.sh | |
| PERFORM TARGET 00018 02 | 00018_US3BLDW_IMPDP_SCHEMA_FUSION_OCSERVER11G.sh |
| ... | |

AMA | Perform Phase

On-Prem - Source

Export network ACLs

ADB-S - Target

FOREIGN KEYS cross-schemas

INDEXES cross-schemas

FUNCTIONAL INDEXES enabling

REVOKE transition privileges

GRANT privs SYS, SYSTEM, CTXSYS, objects

Restore final profiles

Set tablespace quotas

Import network ACLS

Enable Advanced Queues

Recompilation

--- POST PHASE ---

--- Here execute again all scripts one after the other as they might have dependencies again ---

SOURCE

TARGET

| | |
|---|---|
| POST TARGET 00082 01 | 00082_US3BLDW_SQL_REMOVE_MIG_ROLE.sh |
| POST TARGET 00083 01 | 00083_US3BLDW_SQL_SYS_PRIVS.sh |
| POST TARGET 00084 01 | 00084_US3BLDW_SQL_CTXSYS_PRIVS.sh |
| POST TARGET 00085 01 | 00085_US3BLDW_SQL_DATAMINING_PRIVS.sh |
| POST TARGET 00086 01 | 00086_US3BLDW_SQL_OBJECT_PRIVS.sh |
| POST TARGET 00087 01 | 00087_US3BLDW_SQL_ROLE_PRIVS.sh |
| POST TARGET 00088 01 | 00088_US3BLDW_SQL_TBS_QUOTES.sh |
| POST TARGET 00089 01 | 00089_US3BLDW_SQL_DETACH_FSS.sh |
| POST SOURCE 00090 01 00090_US3BLDW_EXPDP_NETWORK_ACL.sh | |
| POST TARGET 00090 02 | 00090_US3BLDW_IMPDP_NETWORK_ACL.sh |
| POST TARGET 00091 01 | 00091_US3BLDW_SQL_SET_AQ_STATUS.sh |
| POST TARGET 00092 01 | 00092_US3BLDW_SQL_RECOMPILE.sh |

--- END OF MIGRATION ---

AMA Workflow



- Object comparison
- Row export/import comparison
- OLAP Analytic Workspace



Done!!



AMA can run a migration fully automated and completely unattended



Works with Windows as source database



Database links, external tables, APEX applications

- Work-in-progress

Key Learnings



- 1 Find the right candidates for ADB
- 2 Follow our migration approach
- 3 Ask us about your ADB migration project

Virtual Classroom Seminar Series #22 – #25



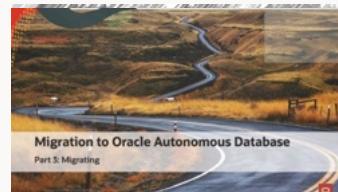
1 PLANNING

Watch [recording](#)
Get [slides](#)



2 PREPARING

Watch [recording](#)
Get [slides](#)



3 MIGRATING

Watch [recording](#)
Get [slides](#)



4 OPERATING

Watch [recording](#)
Get [slides](#)



Try it out, please!!

- We are looking for reference customers
- Get in touch with us when you tested it

Break

We start again at 15:15

Cross-Platform Migrations

Pushing the limits

The Beast



Introduction

Who is who?



ANDREAS GROETZ

Oracle DBA Tech Lead
Entain Services Austria GmbH



Entain is one of the world's largest sports betting and gaming groups. Leveraging the power of the Entain Platform, they bring moments of excitement into their customers' lives through more than 30 iconic brands such as bwin, Coral, Ladbrokes and many more.

Entain operates on over 140 licenses across 40+ territories and employs over 29,000 talented workforce. Entain is listed on the London Stock Exchange and is a constituent of the FTSE 100 Index.

Entain

Ladbrokes



sportingbet

CORAL
GET CLOSER TO THE ACTION

bwin

EUROBET

SuperSport

partypoker

Foxy BINGO

Challenges

What is special, what makes it so complex?

Migration Challenges



SPARC SuperCluster



Exadata X9M Extreme Flash

Migration Challenges

180TB
size



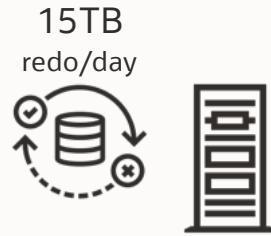
SPARC SuperCluster



Exadata X9M Extreme Flash



Migration Challenges



Migration Challenges



SPARC SuperCluster



Exadata X9M Extreme Flash



5 Physical Standby DBs

Local, and in different region, 2500km away

Constraints

Limiting factors, and other things to know



Up to 15TB redo/day is beyond what Oracle GoldenGate will be able to synch

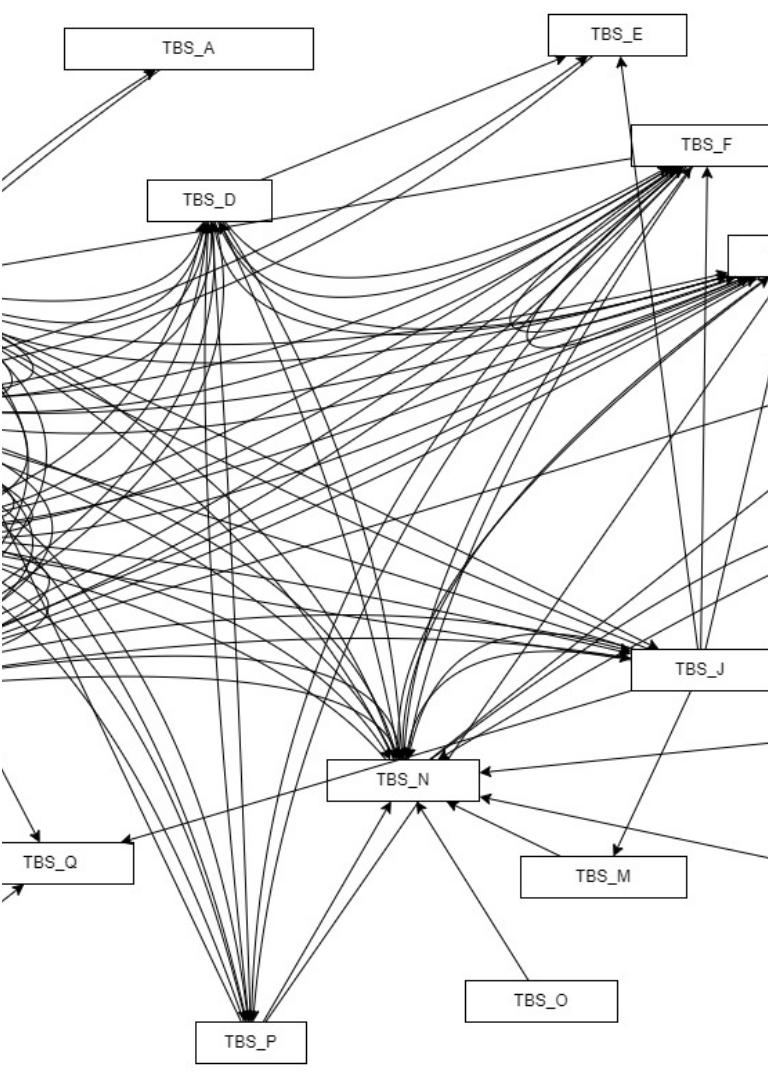
The system is highly active 24 x 7 x 365



Photo by [Dave Hoeffler](#) on [Unsplash](#)

Very large database, very constrained downtime

- 180+ TB database size
- 5-6 TB growth/month
- Every minute of downtime costs \$\$\$



Migrating tablespaces upfront or separately **definitely** not an option

- Way too many cross-dependencies
- Tablespaces aren't isolated

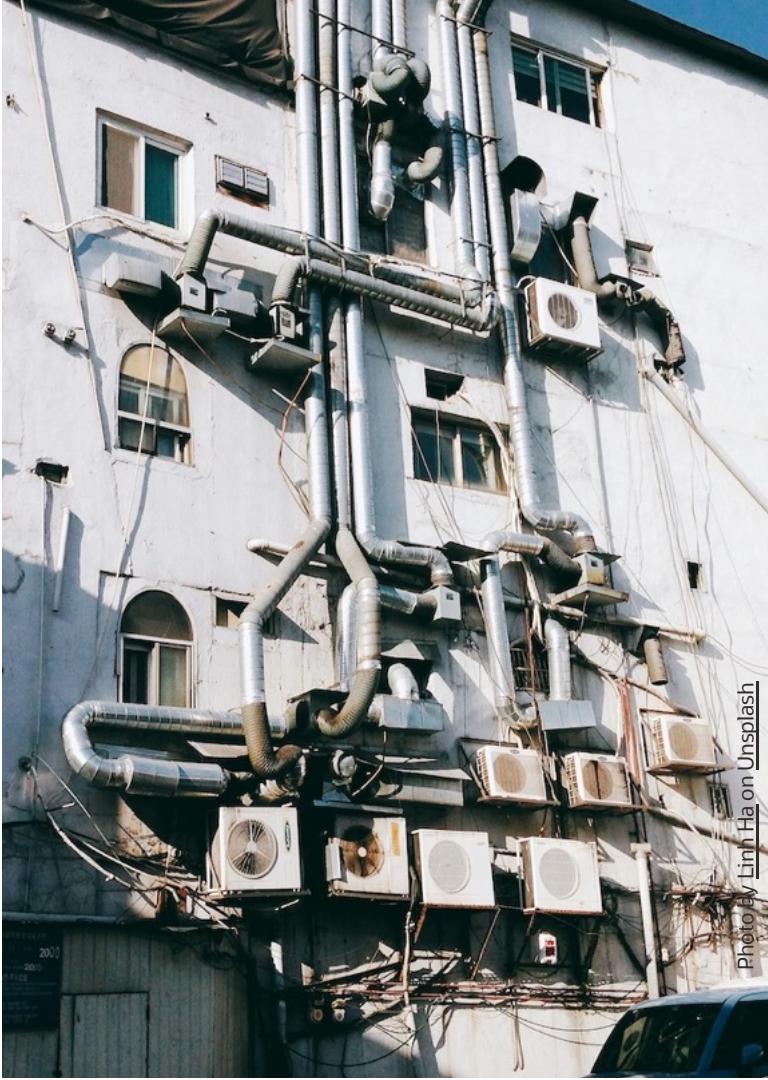


Photo by Linh Ha on Unsplash

Every complex Oracle data type you can imagine is used

- XML binary types
- Nested partitioned tables
- Evolved object types



Tight downtime window

- Dry run: 2 hours outage approved
 - Tablespace read-only
 - Full Transportable Export
- Live migration: 13 hours approved



The available **Oracle V4 PERL** migration scripts would have been worked technically, but ...

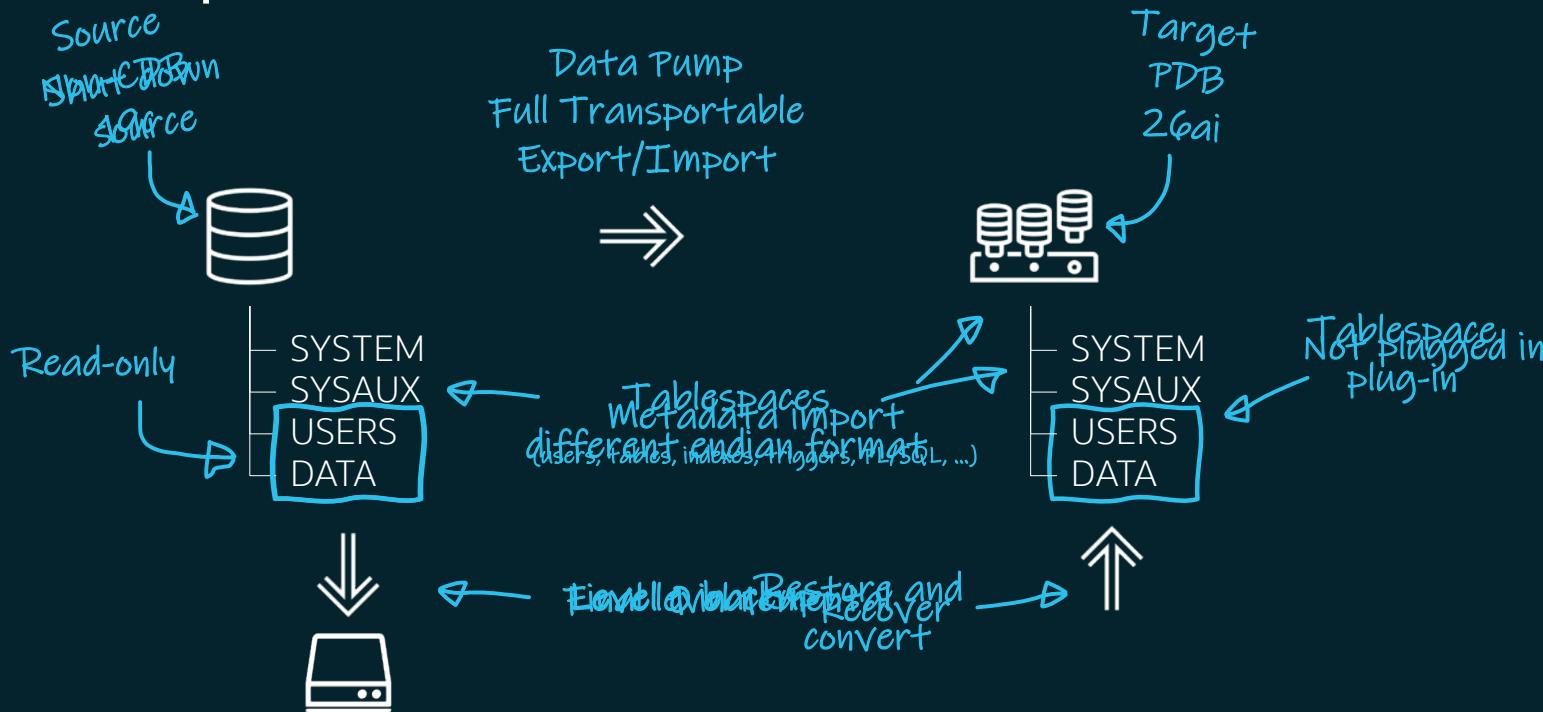
- No section-size backup support
- No standby backup support
- No selective PDB migration support

Migration



Typically, we use **Full Transportable Export/Import** for large cross-endian migrations

Concept



Scripts for Incremental Backup Automation

Backup / restore / recover



- No multisection backups
- No encrypted tablespaces
- Inefficient parallelism
- Incomplete multitenant support



M5 is the [next-generation](#) cross-platform
transportable tablespace procedure

- New RMAN functionality combined with
Full Transportable Export/Import
- Doc ID [2999157.1](#)

M5 Migration Script

The new migrations scripts superseeding the V4 PERL scripts

```
# source database
RUN
{
  ALLOCATE CHANNEL d1 DEVICE TYPE DISK FORMAT '...';
  ALLOCATE CHANNEL d2 DEVICE TYPE DISK FORMAT '...';
  BACKUP
    FILESPERSET 1
    SECTION SIZE 64G
    TAG UP19_L0_240206101548
    TABLESPACE <list-of-tablespace>;
}
```

```
# source database
RUN
{
  ALLOCATE CHANNEL d1 DEVICE TYPE DISK FORMAT '...';
  ALLOCATE CHANNEL d2 DEVICE TYPE DISK FORMAT '...';
  BACKUP
    FILESPERSET 1
    SECTION SIZE 64G
    TAG UP19_L0_240206101548
    TABLESPACE <list-of-tablespace>;
}

# target database
RUN
{
  ALLOCATE CHANNEL DISK1 DEVICE TYPE DISK FORMAT '...';
  ALLOCATE CHANNEL DISK2 DEVICE TYPE DISK FORMAT '...';
  RESTORE ALL FOREIGN DATAFILES TO NEW FROM BACKUPSET
  '<backup-set-1>',
  '<backup-set-2>',
  ...
  '<backup-set-n>'
};
```

Benefits

M5 procedure supports:

- Encrypted tablespaces
- Multisection backups
- Migrating multiple databases into the same CDB simultaneously
- Compressed backup sets
- Better parallelism

Requirements

- Source and target database must
 - be 19.18.0 or higher
 - use Data Pump Bundle Patch



Always use the latest version of M5 script

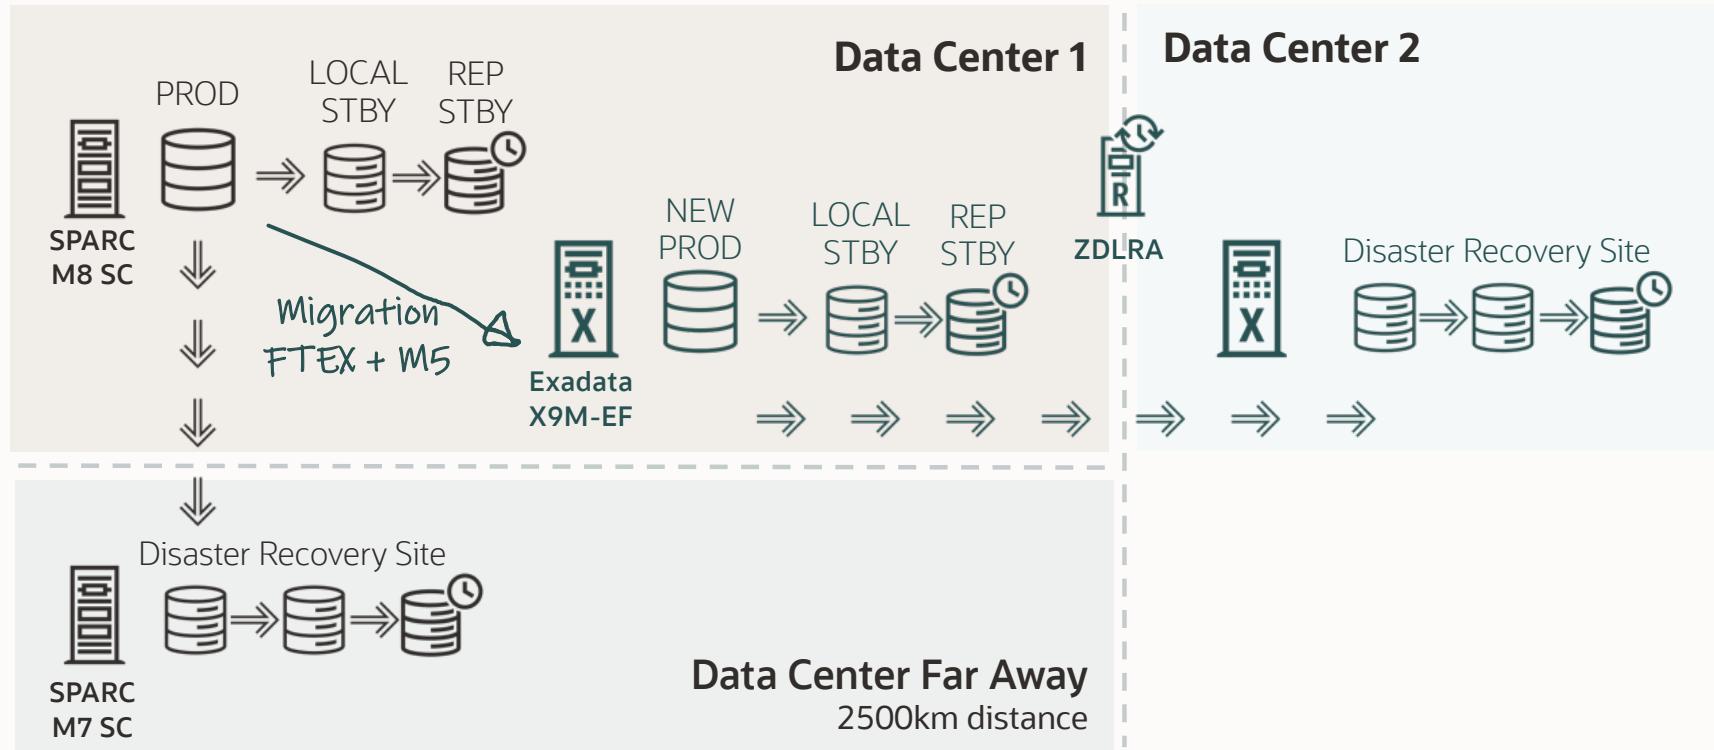
- Download from Doc ID [2999157.1](#)



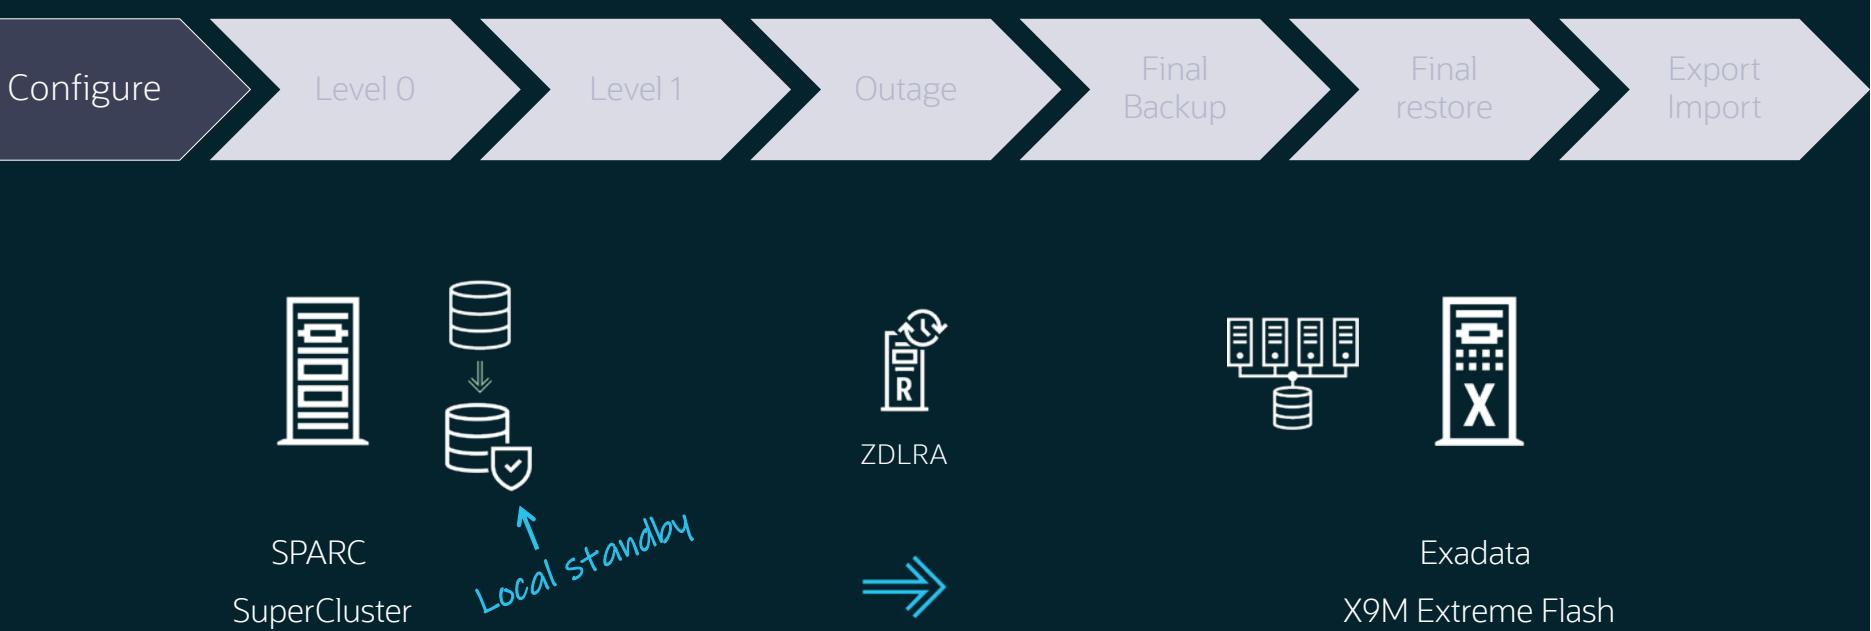
Use Block Change Tracking for faster incremental backups

- Check the License Guide for details

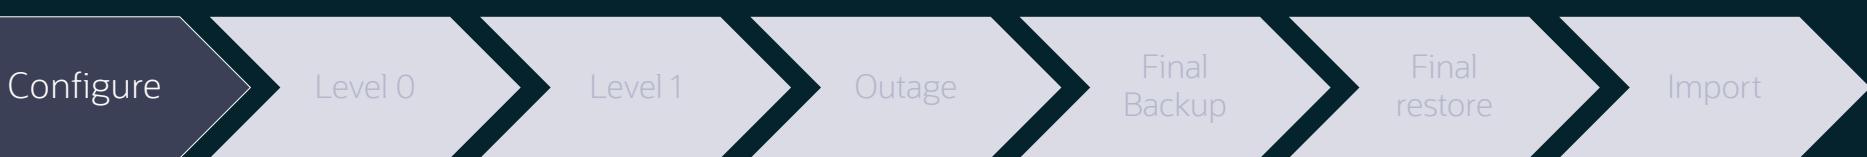
Migration Plan



M5 Workflow

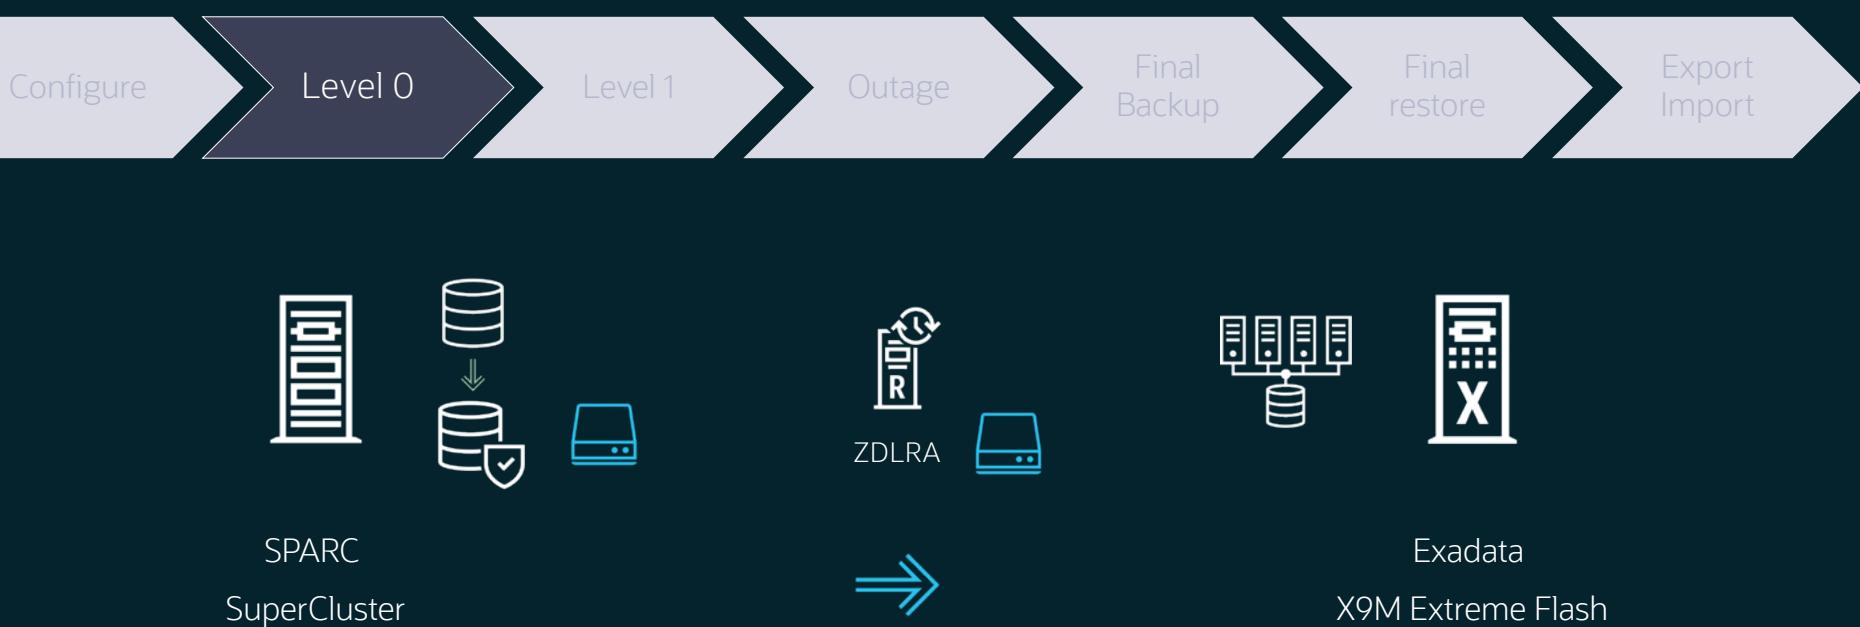


M5 Workflow



- Download M5 script from Doc ID [2999157.1](#)
- Configure shared NFS
- Edit **dbmig_ts_list.txt**
- Edit **dbmig_driver.properties**
- Create new, empty target database

M5 Workflow

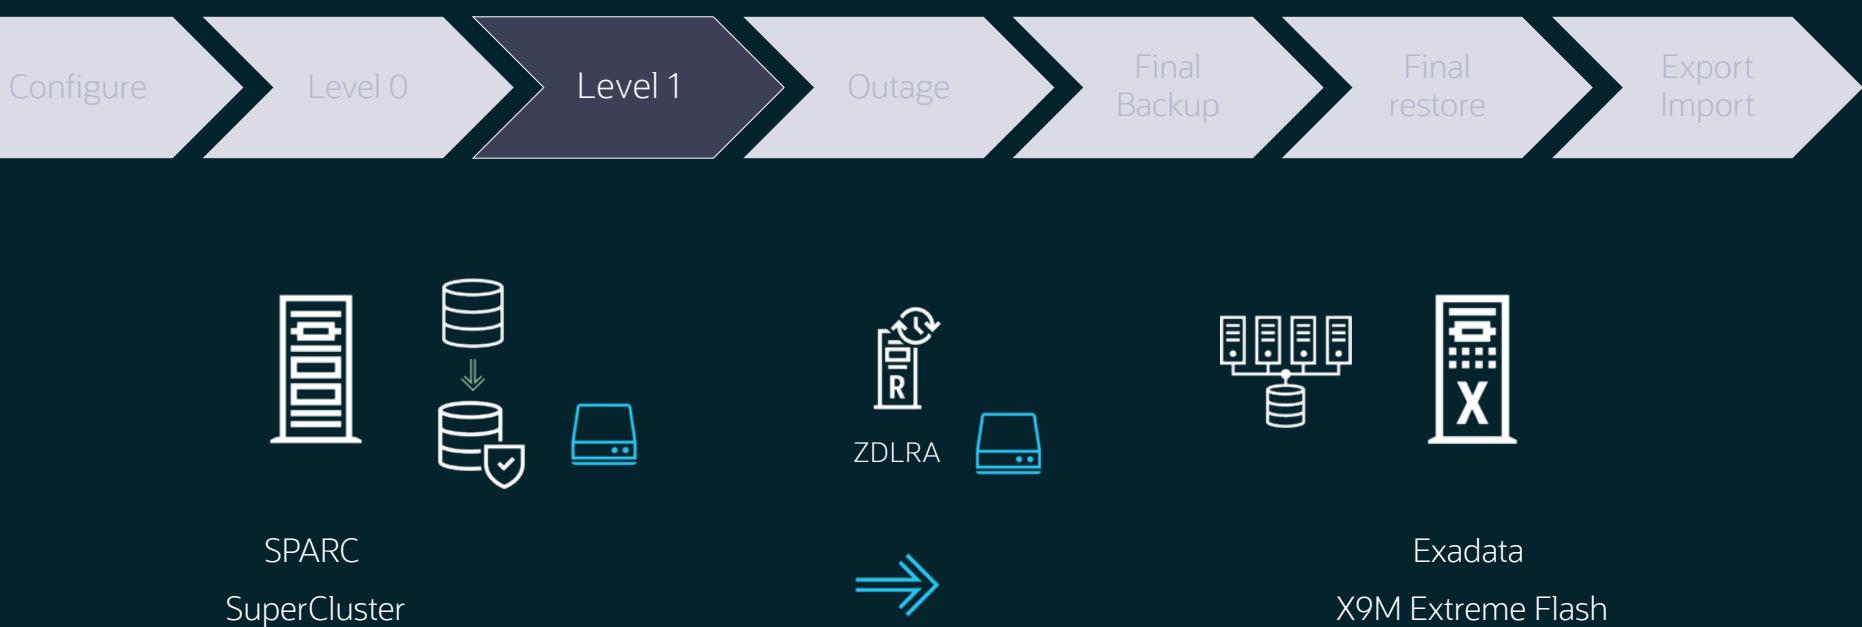


M5 Workflow



- Start initial level 0 backup
 - Use driver script `dbmig_driver_m5.sh L0`
- Driver script creates a restore script
 - Restore using `restore_L0_<source_sid>_<timestamp>.cmd`
- Check logs

M5 Workflow

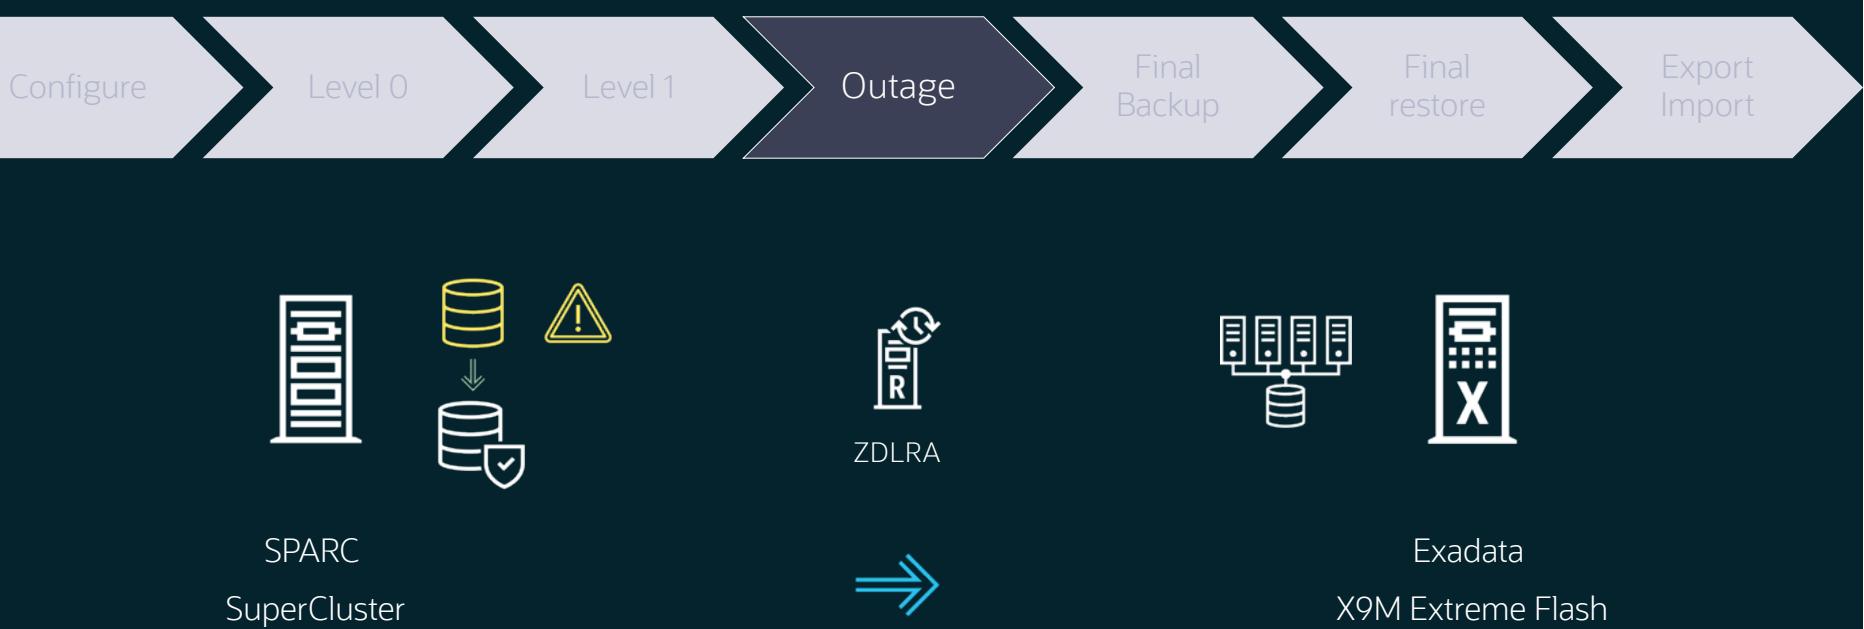


M5 Workflow



- Start level 1 incremental backup
 - Use driver script `dbmig_driver_m5.sh L1`
- Driver script creates a restore script
 - Restore using `restore_L1_<source_sid>_<timestamp>.cmd`
- Check logs
- Repeat as often as desired

M5 Workflow

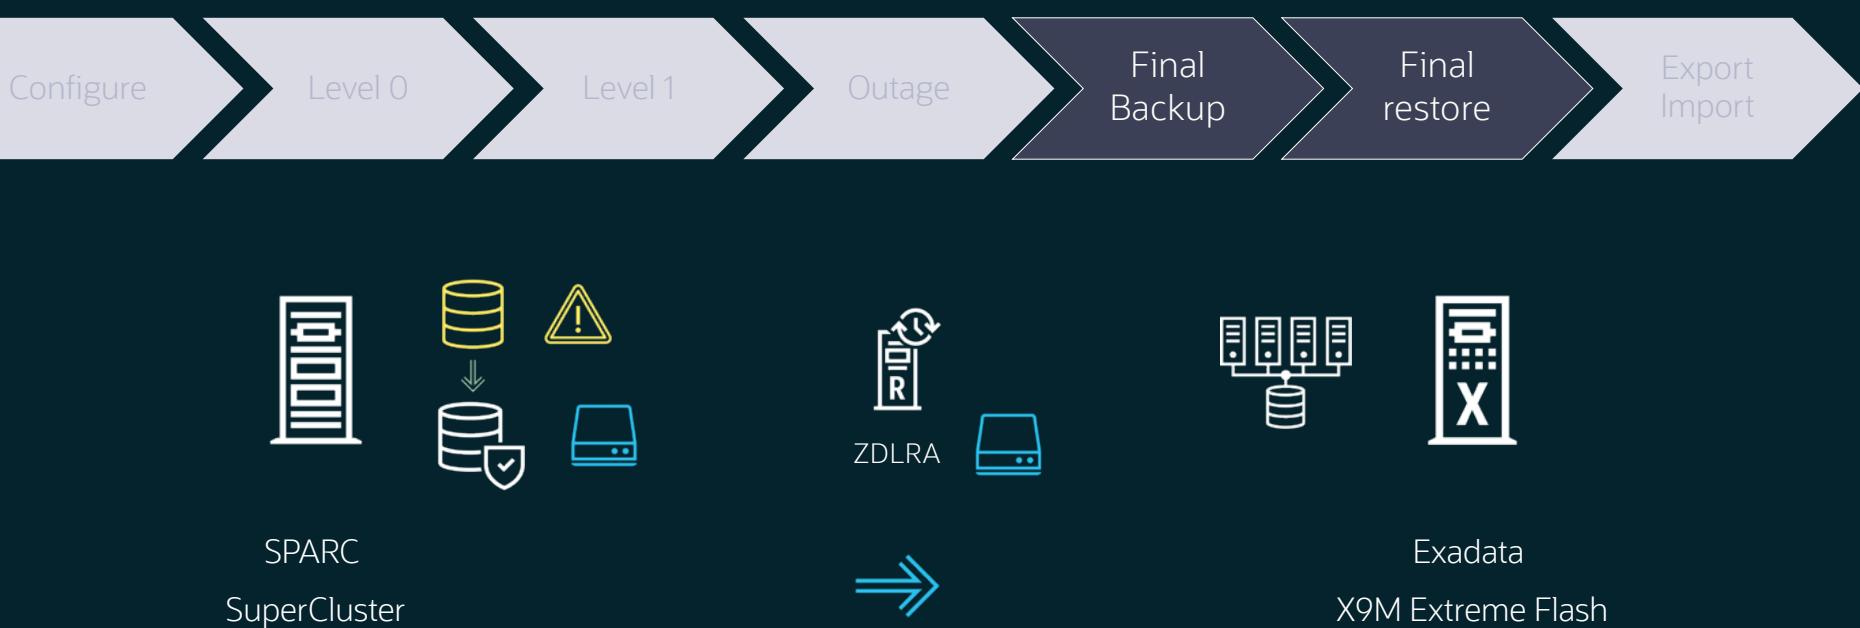


M5 Workflow



- Maintenance window begins
- Read-only sessions can still use the database

M5 Workflow



M5 Workflow



- Start final level 1 incremental backup
 - Use driver script `dbmg_driver_m5.sh L1F`
 - Sets tablespaces read-only
 - Performs level 1 incremental backup
 - Starts Data Pump full transportable export
- Optionally, shuts down source database

M5 Workflow



- Driver script created a restore script
 - Restore using `restore_L1F_<source_sid>_<timestamp>.cmd`
- Check logs

M5 Workflow



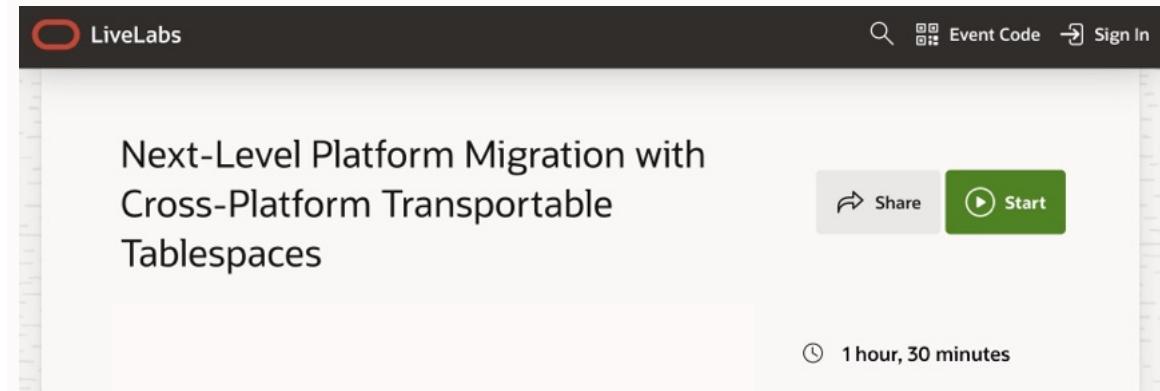
M5 Workflow



- Copy Data Pump dump file to *DATA_PUMP_DIR*
- Use import driver script in test mode
 - Start `impdp.sh <dump_file> <restore_Log> test`
- Check generated parameter file
 - Use `impdp.sh <dump_file> <restore_Log> run`
- Check Data Pump log file



Wanna try it out?



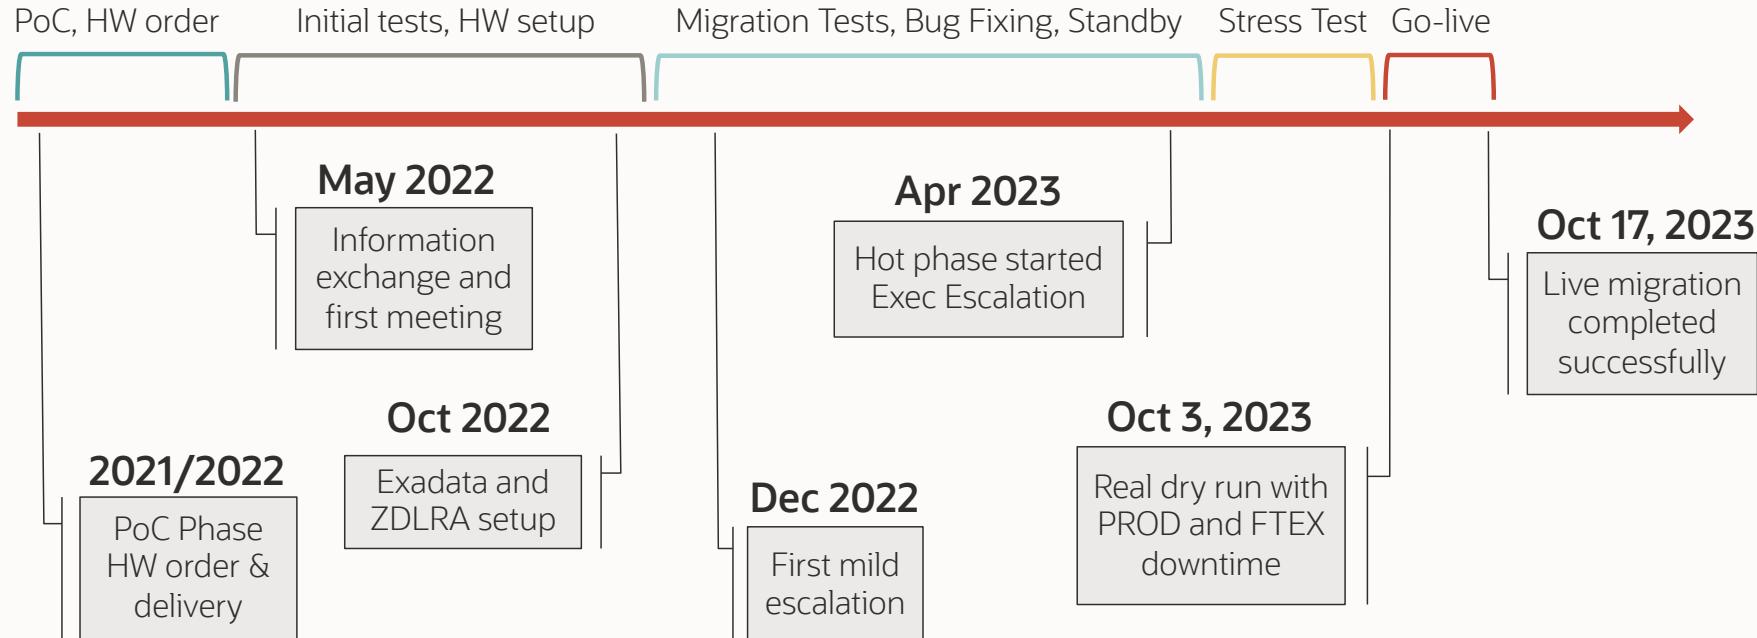
The screenshot shows the Oracle LiveLabs interface. At the top, there is a navigation bar with the "LiveLabs" logo, a search icon, an "Event Code" icon, and a "Sign In" button. The main content area features a title: "Next-Level Platform Migration with Cross-Platform Transportable Tablespaces". Below the title are two buttons: "Share" (with a share icon) and "Start" (a green button with a play icon). At the bottom of the content area, there is a timer icon followed by the text "1 hour, 30 minutes".

Oracle LiveLabs – Run the lab just inside your browser!

The Project Plan

Timelines and the Run Book

Overall Project Timeline



Key to Success: Runbook

Complex projects absolutely require a **runbook**

| ID | Task | Status | Primary Person | Secondary Person | Predecessor | Start Time (CEST) | Duration (hh:mm) | End Time (CEST) | Start Time (IST) | End Time (IST) | Actual Start Time (CEST) | Actual Duration | Actual End Time (CEST) | Comments - Blocker |
|----|------|--------|----------------|------------------|-------------|-------------------|------------------|-----------------|------------------|----------------|--------------------------|-----------------|------------------------|--------------------|
|----|------|--------|----------------|------------------|-------------|-------------------|------------------|-----------------|------------------|----------------|--------------------------|-----------------|------------------------|--------------------|

- This run book covered over 200 individual tasks



The screenshot shows a Microsoft Excel spreadsheet titled 'Runbook' with a tab labeled 'Runbook'. The spreadsheet contains a large grid of data representing 200+ tasks. The columns are labeled as follows: ID, Task, Status, Primary Person, Secondary Person, Predecessor, Start Time (CEST), Duration (hh:mm), End Time (CEST), Start Time (IST), End Time (IST), Actual Start Time (CEST), Actual Duration, Actual End Time (CEST), and Comments - Blocker. The data is organized into several sections, each with a title and a list of tasks. The tasks are described in detail, including specific actions, dependencies, and timelines. The 'Actual Start Time (CEST)' and 'Actual Duration' columns are filled with specific times and durations, indicating the execution history of the tasks.



Timeline Live Migration

How many people were involved?



2 data center engineers



4 product engineering



Migration Issues

Some of them...

Where we started ...



PoC first FTEX export:

```
01-OCT-21 05:32:36.275: Job "SYSTEM"."SYS_EXPORT_FULL_01" successfully completed at  
Fri Oct 1 05:32:36 2021 elapsed 0 04:25:22
```

PoC first FTEX import:

```
05-OCT-21 01:48:59.534: Job "SYSTEM"."SYS_IMPORT_FULL_01" completed with 103000  
error(s) at Tue Oct 5 01:48:59 2021 elapsed 3 18:34:09
```



The way forward ...

80 SRs opened and solved in various areas

Migration
ZDLRA
Standby
Infrastructure

18 one-off patches

5 merges

Daily calls with Oracle, countless evening / night / weekend hours

Many areas required special attention....

- Optimizer Statistics
- Scheduler jobs
- Resource Manager
- Cross Schema objects
- AQ
- Evolved Types/partitioned nested tables
- Binary XML
- Standby DBs

...



Photo by [Justin Chinn](#) on [Unsplash](#)

Issue 1 | Long Running Metadata Import

Fix applied to remedy export errors

- BUG 34201281 - MERGE ON DATABASE RU 19.12.0.0.0 OF 33963454 34052641

Result:

- Now Full Transportable import alone took over **6 days** (!!)

```
08-JUN-22 16:21:17.887: W-1 Processing object type DATABASE_EXPORT/.../PROCACT_INSTANCE
14-JUN-22 18:56:58.813: W-1          Completed 108 PROCACT_INSTANCE objects in 527737 seconds
...
14-JUN-22 19:15:50.016: Job "SYSTEM"."SYS_IMPORT_TRANSPORTABLE_01" completed with 316 error(s)
at Tue Jun 14 19:15:49 2022 elapsed 6 06:45:38
```

Issue 1 | Long Running Meta Import

Long running action identified via tracing:

```
UPDATE "POSTMAN"."T_MAIL_LOG"  
      SET "C_CVAR"=SYS_REMAP_XMLTYPE("C_CVAR")
```

- 300+ million rows

Issue in internal package DBMS_CSX_INT

- Fast merge of XMLTYPE is not happening as expected
 - *Reason:* Incorrect internal check query
- Tokens between source and target are not identical
 - *Reason:* Different Endianness

Issue 1 | Long Running Meta Import

Solution:

- Use workaround from MOS Note: 2309649.1 in UPGRADE mode
 - [MOS Note: 2309649.1 - How to Migrate Large Amount of Binary XML Data between Databases](#)

```
25-JUL-22 12:28:40.813: Job "SYSTEM"."SYS_IMPORT_TRANSPORTABLE_01" completed
with 317 error(s) at Mon Jul 25 12:28:40 2022 elapsed 0 04:33:03
```

Issue 2 | Metadata API and Nested Tables

Full transportable import errors out for a **nested partitioned table**

```
PLS-00172: string literal too long
```

```
ORA-39151: Table "DBA_XY"."X_GAMES" exists.
```

```
    All dependent metadata and data will be skipped due to table_exists_action
```

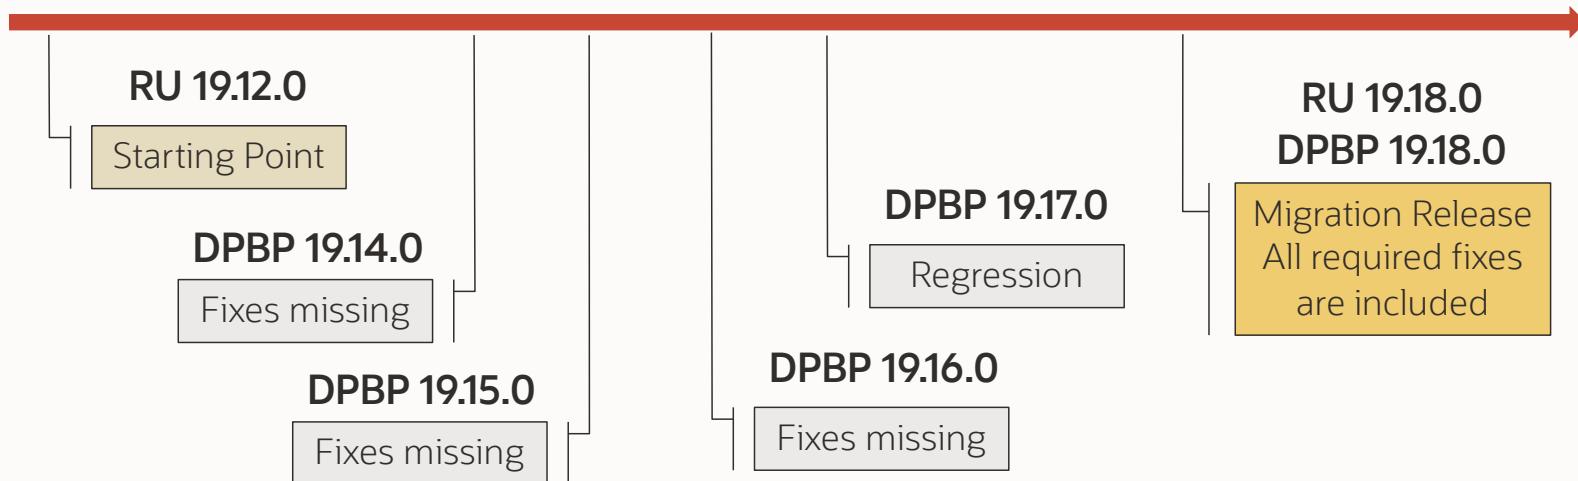
Root cause was a string overflow in the Metadata API

- Data Pump creates the index, and then alters it – here the overflow happened
- Side effect was a misleading error message

Issue 3 | Data Pump Bundle Patch

Many of the TTS and Metadata fixes got included into Data Pump Bundle

- [Data Pump Recommended Proactive Patches For 19.10 and Above \(Doc ID 2819284.1\)](#)



Issue 4 | Evolved Object Types

Evolved TYPEs can lead to Data Pump errors during transportable import:

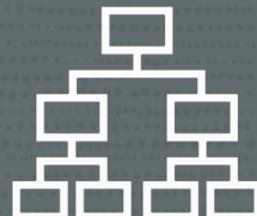
ORA-39083: Object type TABLE:"APPUSER"."CARS" failed to create with error:

ORA-39218: type check on object type "APPUSER"."CAR_TYPE" failed

ORA-39216: object type "APPUSER"."CAR_TYPE" hashcode or version number mismatch

Further Information:

- [Blog Post: Understand why Data Pump errors with evolved types](#)



Using evolved types in table definitions

```
--Create a new type. The type is now version 1
```

```
--Use the type in a table
```

```
CREATE TYPE CAR_INFO_TYPE IS OBJECT (model VARCHAR2(40));
```

```
CREATE TABLE CARS (id number, car_info car_info_type);
```

```
INSERT INTO CARS VALUES (1, car_info_type('Volvo V90'));
```

The type is now evolving

```
--Make a change to the type. The type is now version 2
```

```
ALTER TYPE CAR_INFO_TYPE ADD ATTRIBUTE horsepower NUMBER CASCADE NOT INCLUDING TABLE DATA;
```

```
INSERT INTO CARS VALUES (2, car_info_type('BMW 530i', 250));
```

Existing data is not updated

```
--Make another change to the type. The type is now version 3
```

```
ALTER TYPE CAR_INFO_TYPE ADD ATTRIBUTE color VARCHAR2(20) CASCADE NOT INCLUDING TABLE DATA;
```

```
INSERT INTO CARS VALUES (3, car_info_type('Hyundai Sonata', 160, 'Black'));
```

Evolved Types

```
SELECT * FROM CARS
```



| CARS | |
|------|--|
| 1 | car_info_type v1: Volvo V90 |
| 2 | car_info_type v2: BMW 530i, 250 |
| 3 | car_info_type v3: Hyundai Sonata, 160, Black |



| DICTIONARY | |
|------------------|--------------------------|
| car_info_type v1 | model |
| car_info_type v2 | model, horsepower |
| car_info_type v3 | model, horsepower, color |



Data Pump recreates types during
Full Transportable Export/Import

Evolved Types

- To avoid **data corruption**,
Data Pump must recreate the exact same type evolution in target database
- Due to **implementation restrictions**,
it is not always possible to recreate the exact same type evolution
- In such situations, to avoid corruption,
Data Pump reports ORA-39218 or ORA-39216 on **import**

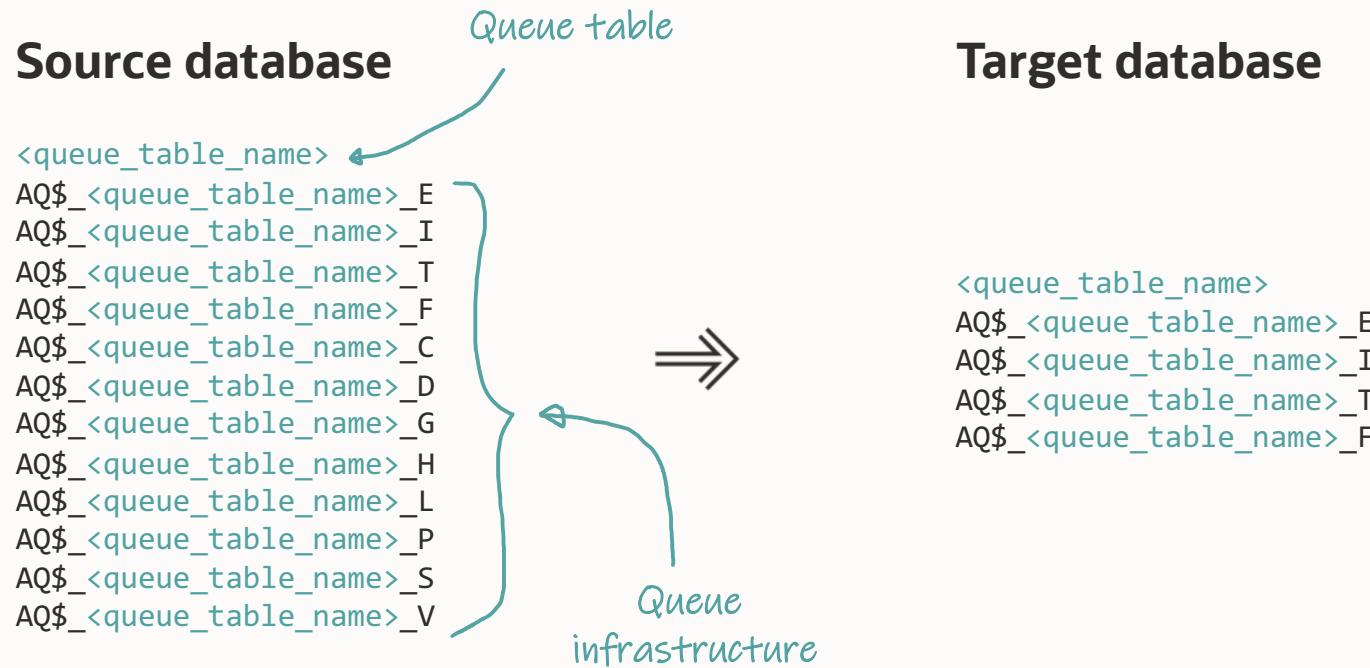


Evolved Types | Possible Solutions

- 1 Conventional Data Pump export
- 2 Manually recreate `type` in target database with matching evolution
- 3 Recreate `type` without evolution before export

[Blog post](#) with details

Issue 5 | Advanced Queueing



Issue 5 | Advanced Queueing

Queue tables and underlying objects may change during import

- COMPATIBLE **during creation** of queue tables matters
- COMPATIBLE **during import** matters as well
- [MOS Note: 2291530.1 - Understanding How AQ Objects Are Exported And Imported](#)
- [Blog post: Changing data types in queue tables during import](#)

Options:

- Recreate the queue tables with "old" COMPATIBLE setting
- Benefit from new COMPATIBLE setting and test the application



Take into account when comparing source and target databases' object count

- Understanding How Advanced Queueing (AQ) Objects Are Exported And Imported (Doc ID [2291530.1](#))



Data Pump does not start queues

- Manually start queues after migration
- Use `DBMS_AQADM.START_QUEUE`

Issue 6 | Default Tablespaces

Due to a security fix export wants to write into the user's default tablespace

- Bug 27692190
- But default tablespaces are read-only while full transportable export runs

Workaround:

- Change default tablespace for all users to SYSTEM
- Full Transportable Export
- Full Transportable Import
- Revert default tablespaces back to original in target (and source)

Issue 7 | Exporting Statistics

Exporting statistics is slow using
DBMS_STATS.EXPORT_SCHEMA_STATS

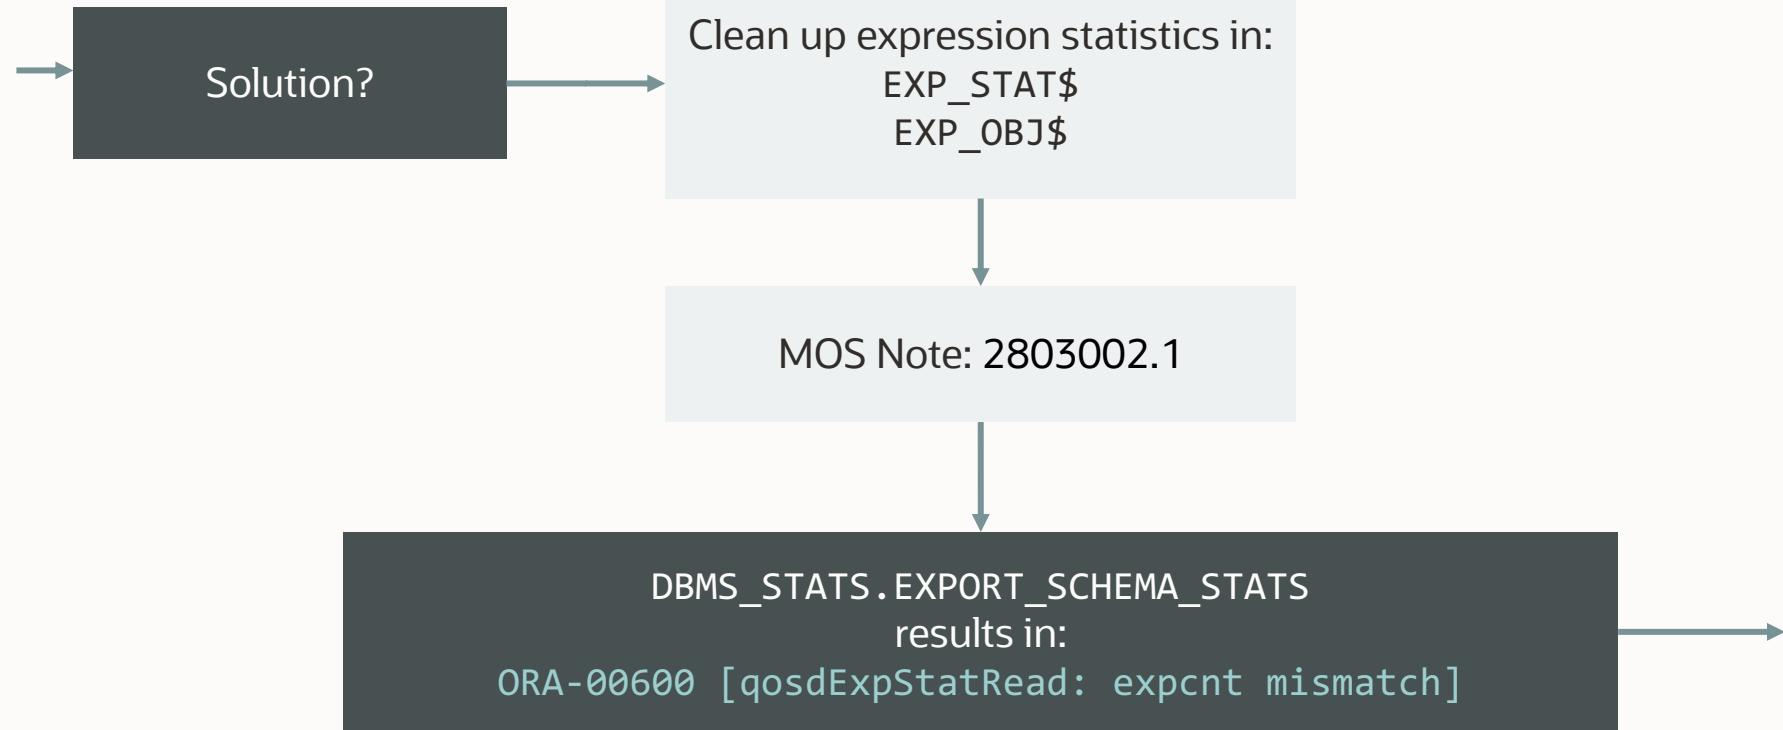


10046 trace reveals long runtime on:
EXP_STAT\$
EXP_OBJ\$

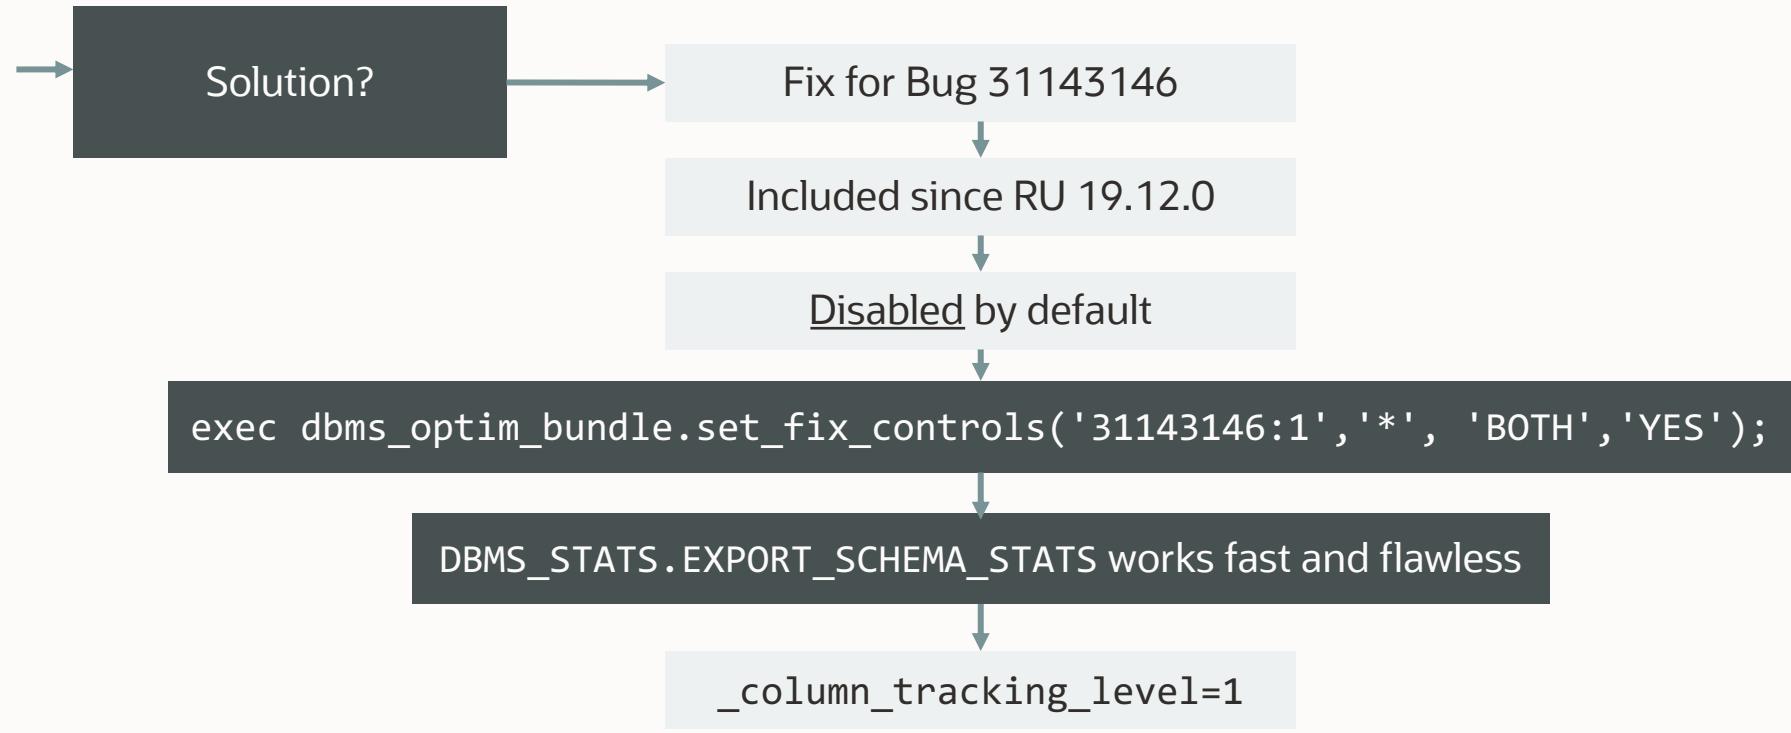
Expression Statistics for Auto-Indexing

Default: `_column_tracking_level=53;`

Issue 7 | Exporting Statistics



Issue 7 | Exporting Statistics



Issue 8 | Auditing

Tablespaces containing auditing tables **can't be set read-only**

Data Pump always unloads the audit records into the dump file

- Huge audit trail will lead to a huge dump file and longer outage

Options:

- Export audit records, and eventually import them afterwards
- Archive audit records, purge the audit trail



Key Learnings



- 1 Practice, practice, practice
- 2 Create a detailed runbook
- 3 Work together

AI changes everything



Harvard Business Review

***“AI won’t replace humans -
but humans with AI will
replace humans without AI”***



Karim Lakhani



AI Database



Using AI



Filter



Year



Genre



Nominations

Awards

Question

Message

Motivational movies about athletics

35

Search Type

 Text Search Vector Search Search

Results

1 - 10 of 10

| | Title | Awards | Year | Nominations | Genre |
|--|------------------|--------|------|--|------------------------------------|
|  | Goodbye Columbus | NONE | 1969 | Academy Award for Best Writing, Adapted Screenplay | Romance, Comedy, Family |
|  | The Formula | NONE | 2002 | NONE | Unknown |
|  | Batman Returns | NONE | 1992 | Academy Award for Best Makeup and Hairstyling, Academy Award for Best Visual Effects, MTV Movie Award for Best Villain | Film-Noir, Action, Family, Fantasy |

Why Batman Returns?

“Motivational Movies about Athletics”

NJQ15253, *Batman Returns*, 1992, https://upload.wikimedia.org/wikipedia/en/8/83/Batman_returns_poster2.jpg, 'Film-Noir, Action, Family, Fantasy', 'Academy Award for Best Makeup and Hairstyling, Academy Award for Best Visual Effects, MTV Movie Award for Best Villain', , *Batman Returns* is a 1992 American superhero film directed by Tim Burton and produced by Denise Di Novi and Burton, based on the DC Comics character Batman. The sequel to the 1989 film *Batman* , it is the second installment of Warner Bros. initial Batman film series, and stars Michael Keaton as Bruce Wayne / Batman, alongside Danny DeVito, Michelle Pfeiffer, Christopher Walken, Michael Gough, Pat Hingle and Michael Murphy. In *Batman Returns* , Batman faces the Penguin, who plots to kill all of Gotham City's firstborn sons, while dealing with Catwoman, who seeks vengeance against Max Shreck, a corrupt tycoon who allies with the Penguin to bring Gotham City under his control. Burton originally did not want to direct another Batman film. Warner Bros. developed a script with Sam Hamm which had the Penguin and Catwoman going after hidden treasure. Burton agreed to return after they granted him more creative control and replaced Hamm with Daniel Waters. Wesley Strick was later chosen to do an uncredited rewrite shortly before filming. This included normalizing dialogue, fleshing out the Penguin's

motivations and master plan, and removing scenes due to budget concerns. Strick continued working as the on-set writer through filming. Annette Bening was originally cast as Catwoman, but became pregnant and was replaced with Pfeiffer. *Batman Returns* was released on June 19, 1992. It grossed \$266.8 million worldwide on a total budget of \$80 million and received positive reviews. Critics praised its action sequences, performances, Danny Elfman's score, effects and villains, although its dark tone and high level of violence for a PG-13 film, was criticized. The film was nominated for two Academy Awards: Best Visual Effects and Best Makeup, as well as two BAFTA awards. A stand- alone sequel, *Batman Forever* , was released in 1995, with Val Kilmer replacing Keaton as Batman. An alternate comic book continuation, which ignores the events of the subsequent films, will be published by DC Comics starting in July 2021. Keaton is also set to reprise the role of Batman in the DC Extended Universe beginning with *The Flash* (2022).



Filter



Year



Genre



Nominations

Awards

Question

Message

Motivational movies about athletics

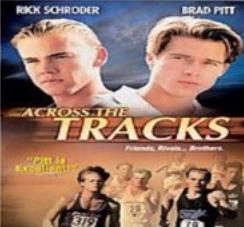
35

Search Type

 Text Search Vector Search Search

Results

1 - 10 of 10

| | Title | Awards | Year | Nominations | Genre |
|--|------------------------|--------|------|-------------|--------------|
|  | The Do-Deca-Pentathlon | NONE | 2012 | NONE | Comedy |
|  | Champions | NONE | 2023 | NONE | Comedy,Sport |
|  | Across the Tracks | NONE | 1991 | NONE | Drama,Action |
|  | Hustle | NONE | 2022 | NONE | Sport |



Filter



Year
2023

Genre
Sport

Nominations

Awards

Question

Message

Motivational movies about athletics

35

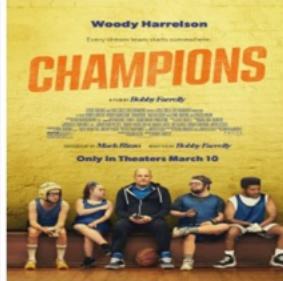
Search Type

Text Search Vector Search

Search

Results

1 - 3 of 3

| | Title | Awards | Year | Nominations | Genre |
|--|--------------|--------|------|-------------|------------------|
|  | Champions | NONE | 2023 | NONE | Comedy, Sport |
|  | Sweetwater | NONE | 2023 | NONE | Biography, Sport |
|  | 80 for Brady | NONE | 2023 | NONE | Sport, Comedy |

1 - 3 of 3

Key new AI data type: AI Vector

Stores semantic content
of a complex object

Source
Object

AI
Model

AI
Vector



Document



33
42
16
21
50



Image



94
31
16
21
12

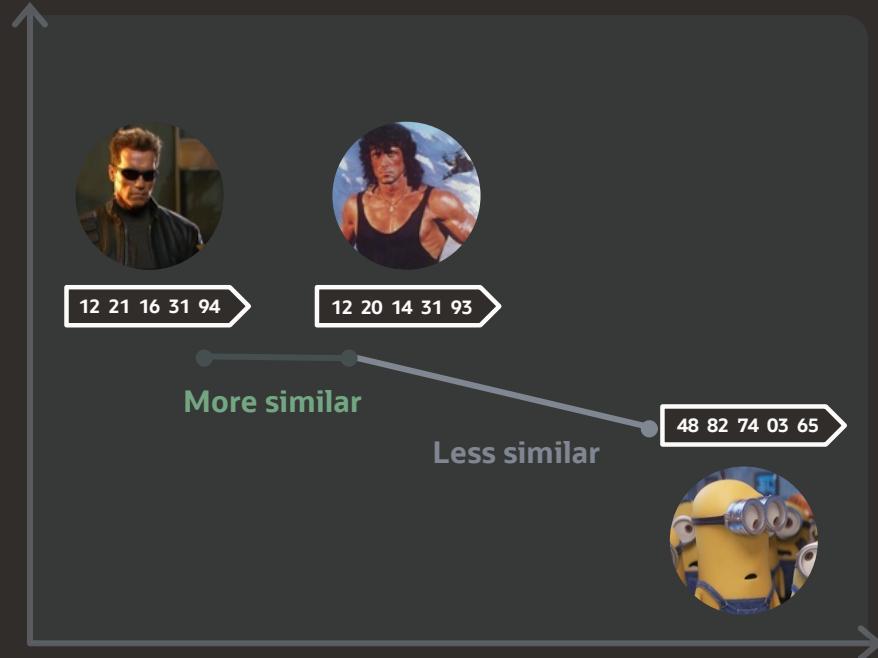


Video



65
03
74
82
48

Oracle AI Database
can **store** AI vectors and
use them to **find**
similar objects



AI needs **business data** to answer your questions

Finding relevant business data requires searching **private data** with **traditional data search** and **AI Vector Search**



AI Database **unifies** search across traditional data and AI vectors



Does my dental plan cover braces for my 19-year-old?

Find semantically similar benefit documents that are relevant to the employee



**Bring AI search to
both operational and
analytical systems**

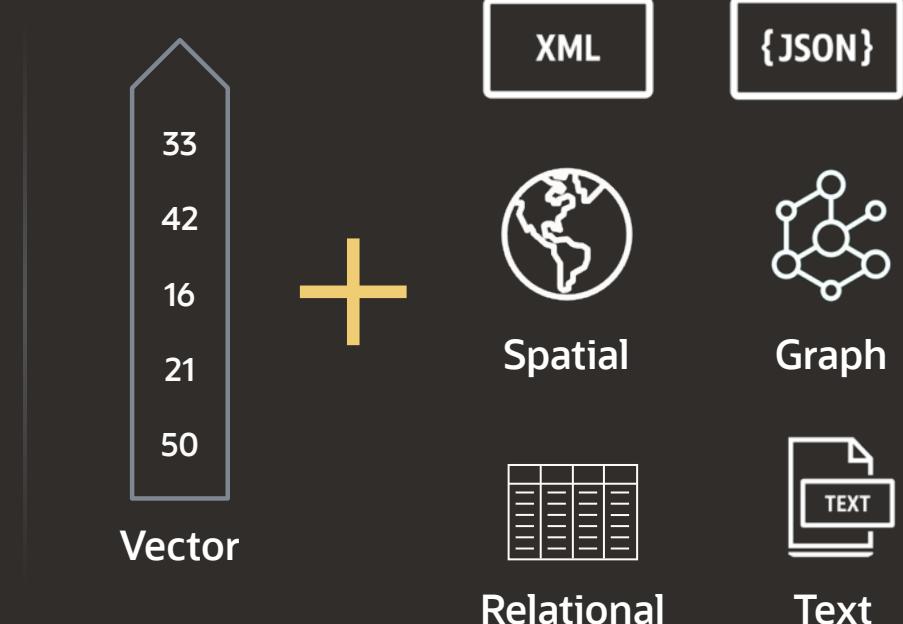
**with just one extra
line of SQL**



Does my dental plan cover
braces for my 19-year-old?

```
SELECT document
  FROM benefits_documents
 WHERE doc_plan_id IN (SELECT plan_id from emp_plans...)
 ORDER BY VECTOR_DISTANCE(:question_vec,
 document_vec);
```

**Unify AI Vector Search
with **all** types of data,
not just relational**



Easily pass results of search to LLM to answer the question



Does my dental plan cover
braces for my 19-year-old?



AI Vector Search



**Retrieval
Augmented
Generation**

Yes, your “Dental Plus” plan covers braces
for your children until they are 21 years old.

Vector Database



- New data type: **VECTOR**
- A vector is
*a numerical representation of data
that captures key features and relationships of the data*

Vector Database



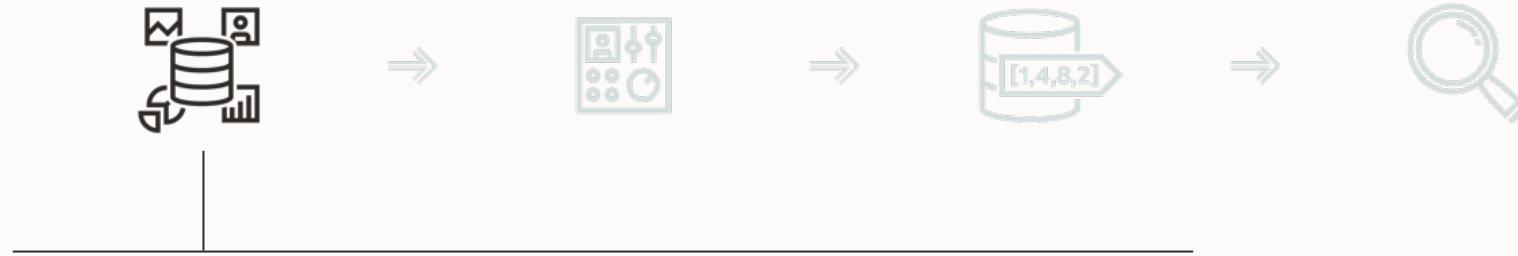
```
CREATE TABLE movies (
    id          NUMBER,
    description CLOB,
    photo       BLOB,
    my_vector   VECTOR(768, FLOAT32));
```

It's just an array ...

50 21 16 42 33

The dimension is
"how many numbers"

Vector Database



- Your data is already in the database
- How do you perform a semantic search instead of a literal search?

Comparing numbers is trivial

$100 > 50 = \text{true}$

How to compare complex data types?



How to search for search complex data types?



Vector Database



- Search vector data using vector indexes
- Calculate the vector distance



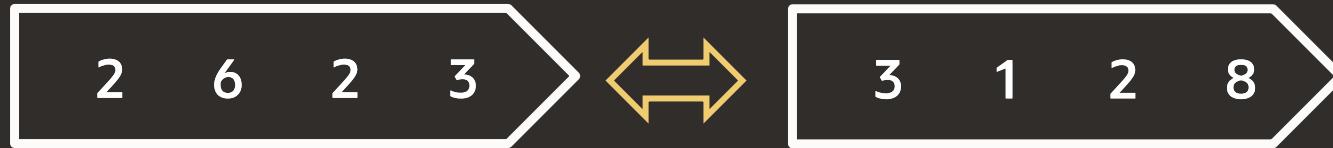
Type of movie



Genre



Mood



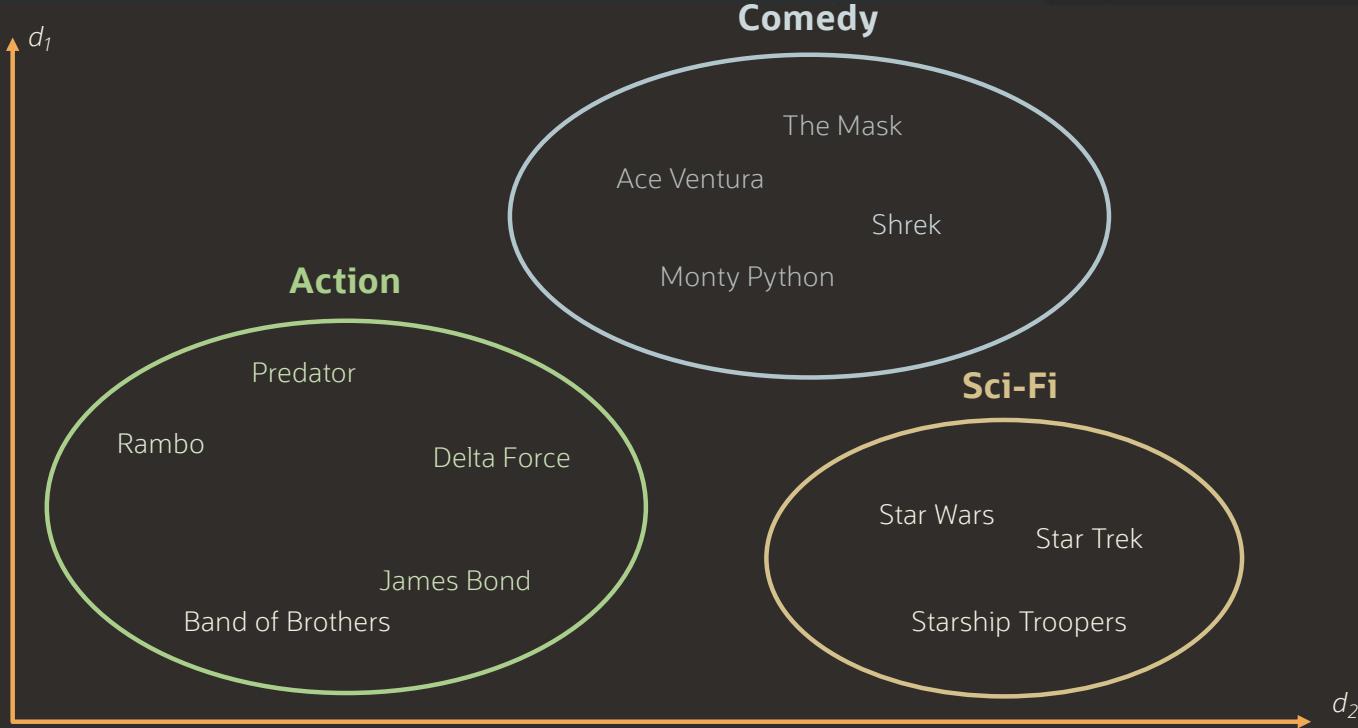
Distance (Euclidean Squared)

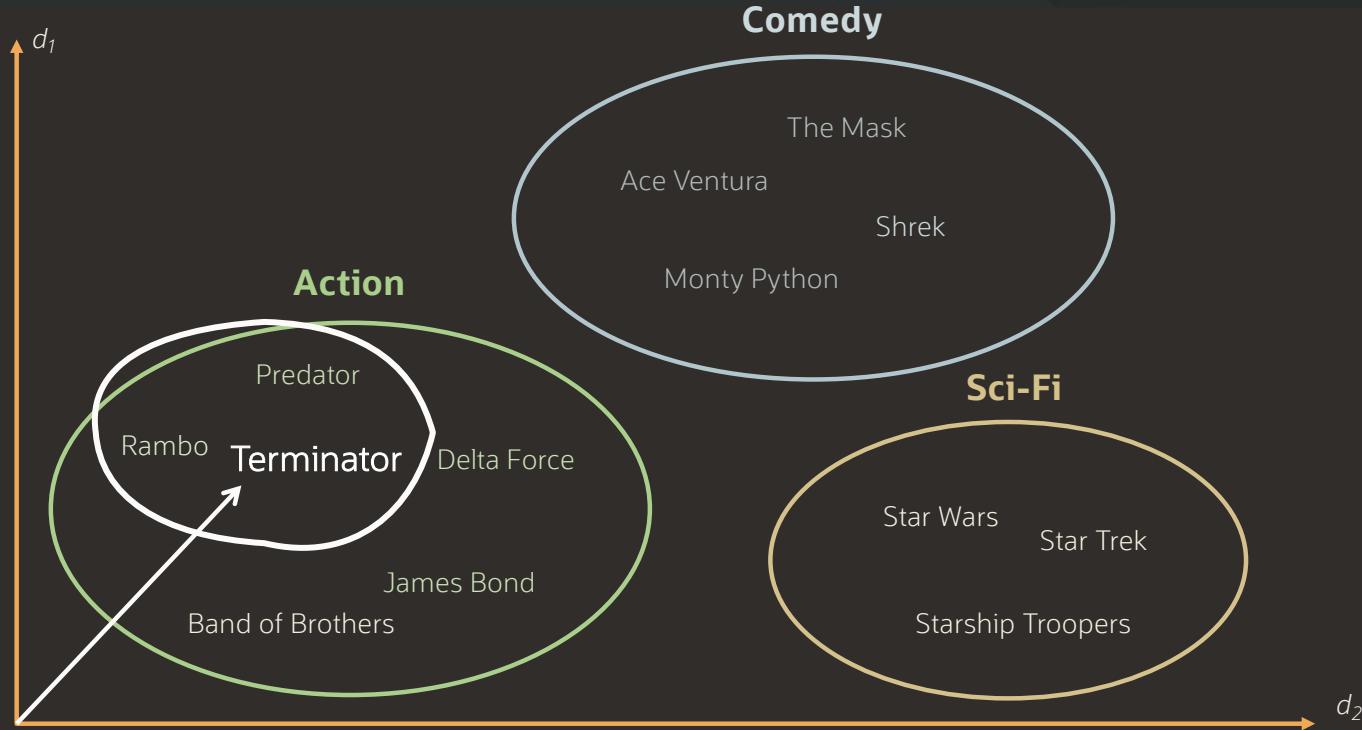
$$= (3-2)^2 + (1-6)^2 + (2-2)^2 + (8-3)^2$$

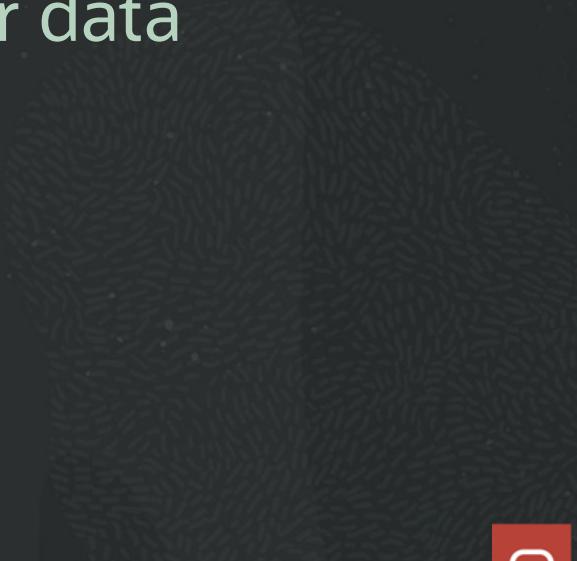
Vector Database



```
SELECT ...  
FROM movies  
ORDER BY vector_distance(movie1, movie2, EUCLIDEAN_SQUARED);
```







How do you convert your data into vectors?

Vector Database



- Embedding models transform your data into a vector

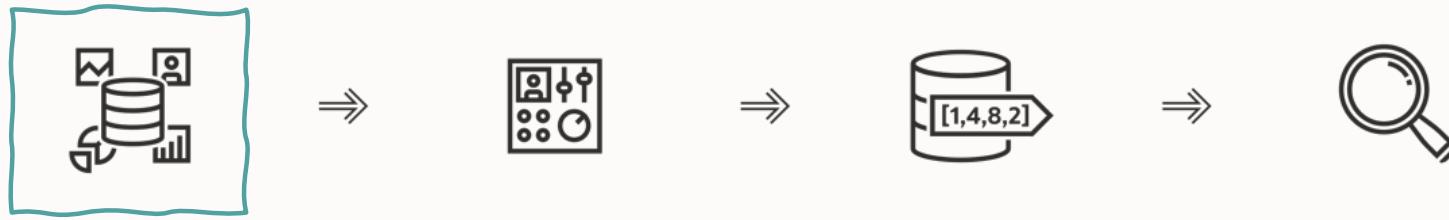


10,000 GPU ?

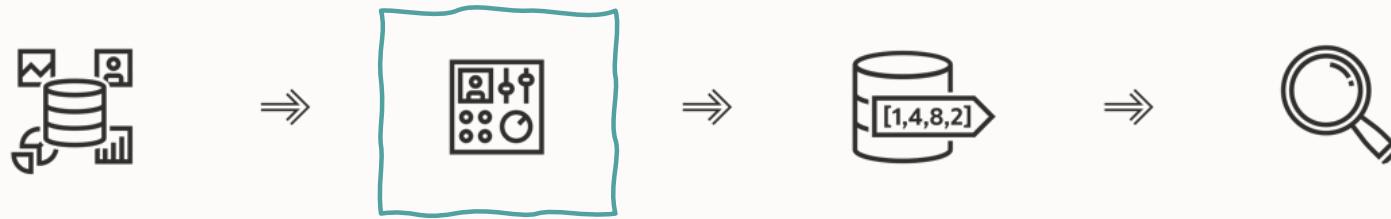
```
DBMS_DATA_MINING.import_onnx_model(  
    model_name => 'All-MiniLM-L6-v2',  
    model_data => 'All-MiniLM-L6-v2.onnx'  
    ...  
)
```

```
INSERT INTO movies
VALUES (:id,
       :description,
       :photo,
       VECTOR_EMBEDDING(All-MiniLM-L6-v2, :description));
```

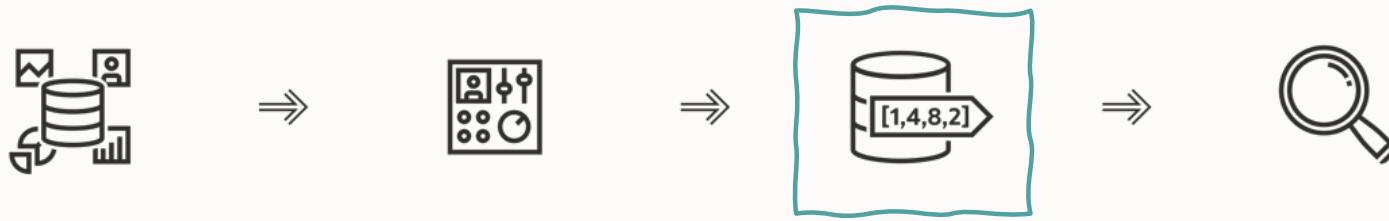
Vector Database



Vector Database



Vector Database



Vector Database





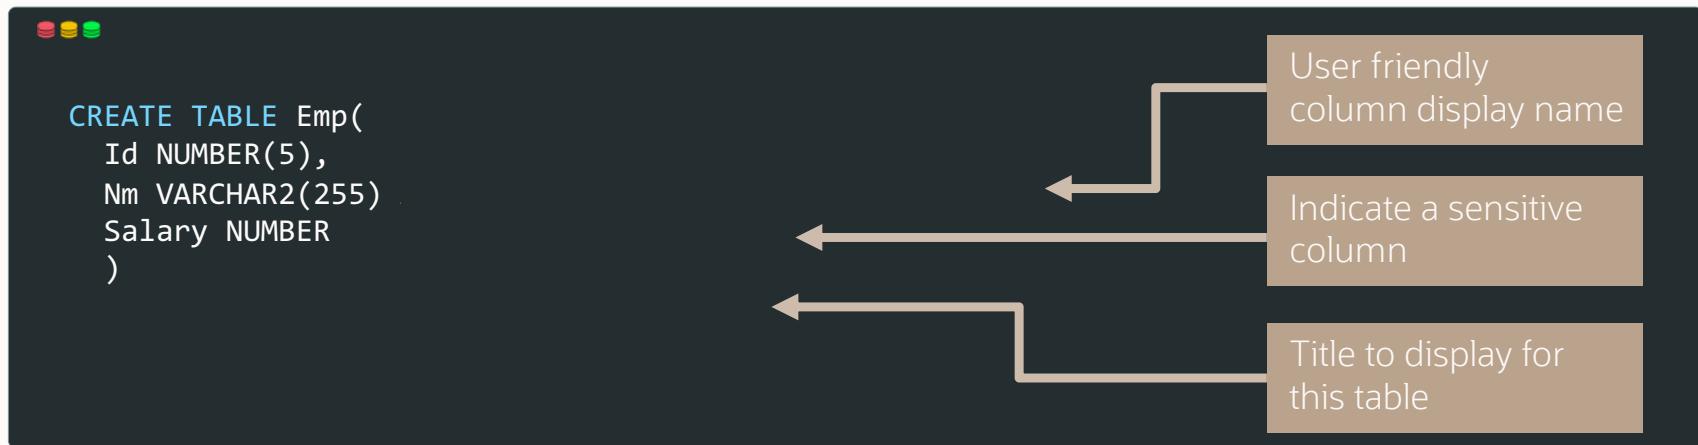
Included in any edition
of Oracle AI Database 26ai

- No extra license required

Oracle AI Database 26ai Annotations

Annotations add **intended usage** to schema elements such as tables, views, columns

User-friendly display names, identify sensitive data, etc.



Annotations can be used by low code frameworks to **generate user interfaces**



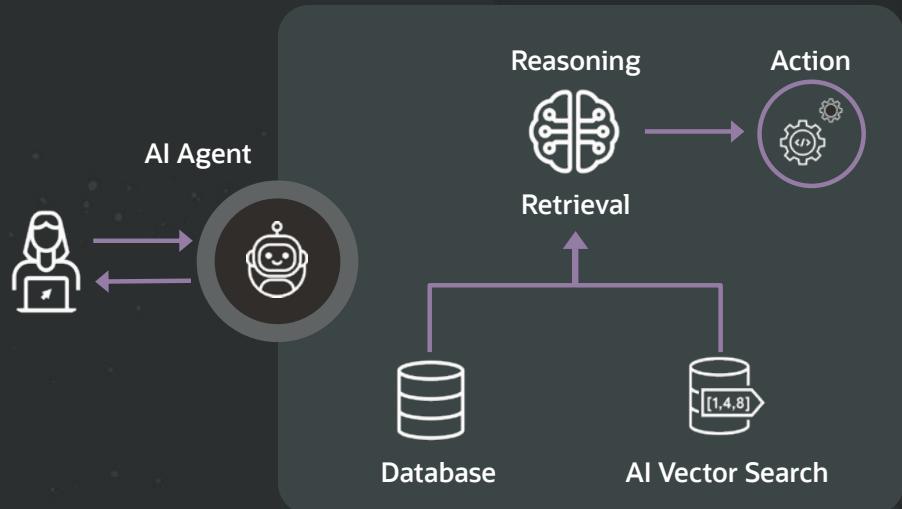
AI Database



Using AI

Architecting Agentic AI into the database

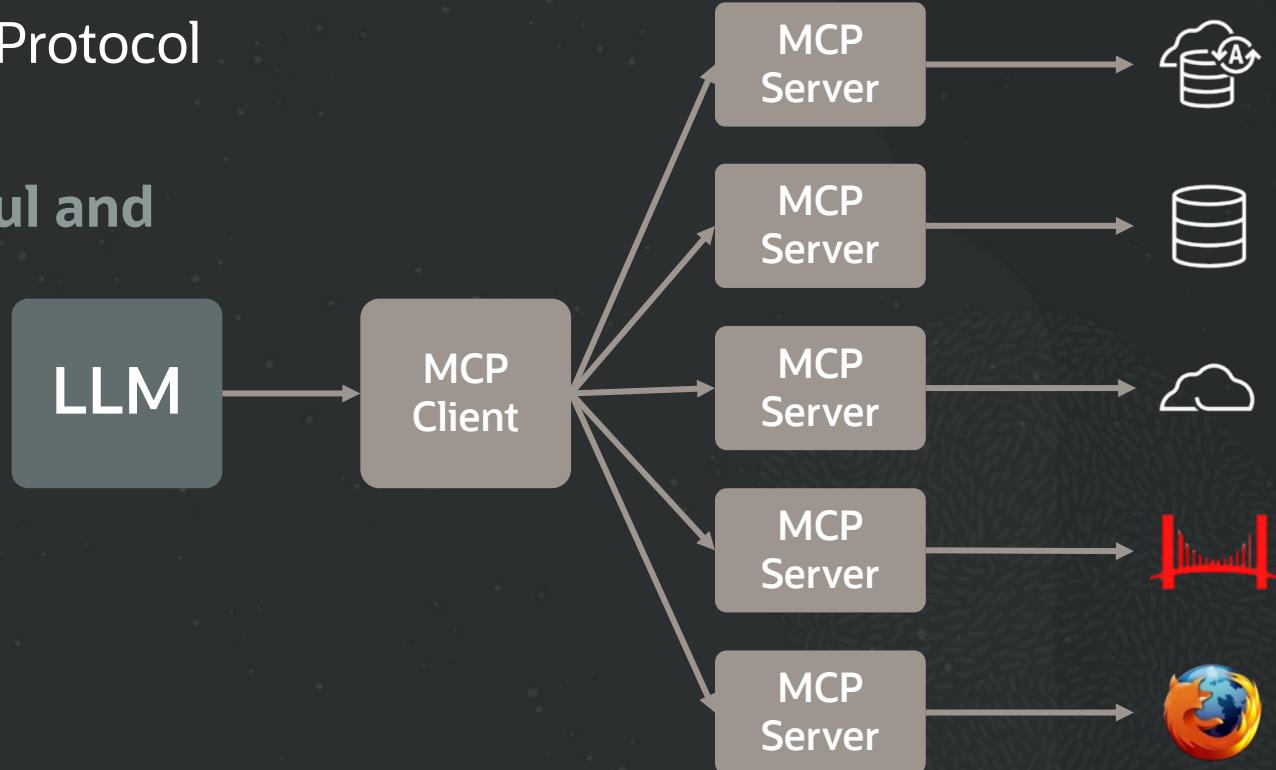
- › Select AI Agent
- › Private Agent Factory
- › SQLcl MCP Server



MCP Server

Model Context Protocol

Enable powerful and
simplistic AI



Untitled-1 X CHAT

Generate code (⌘I), or select a language (⌘K M). Start typing to dismiss or don't show this again.

EXTENSIONS

Search Extensions in Marketplace

INSTALLED

- Cline** Autonomous coding agent right in your IDE, capable of crea... 5 504ms Cline
- GitHub Copilot Chat** AI chat features powered by Copilot 404ms GitHub
- Oracle Code Assist** Oracle Code Assist Oracle 91ms Oracle
- Oracle SQL Developer Extension for VSCode** End-to-end development tools for your SQL and PL/SQL ap... 91ms Oracle Corporation
- GitHub Copilot** Your AI pair programmer GitHub

RECOMMENDED

- SQLTools** Connecting users to many of the most commonly used data... 6M 3.5 Matheus Teixeira Install
- Vim** Vim emulation for Visual Studio Code 8.3M 4 vscodevim Install

MCP SERVERS

MCP Servers

Browse and Install Model Context Protocol (MCP)

Build with Agent

AI responses may be inaccurate.

Generate Agent Instructions to onboard AI onto your codebase.

Untitled-1

Describe what to build next

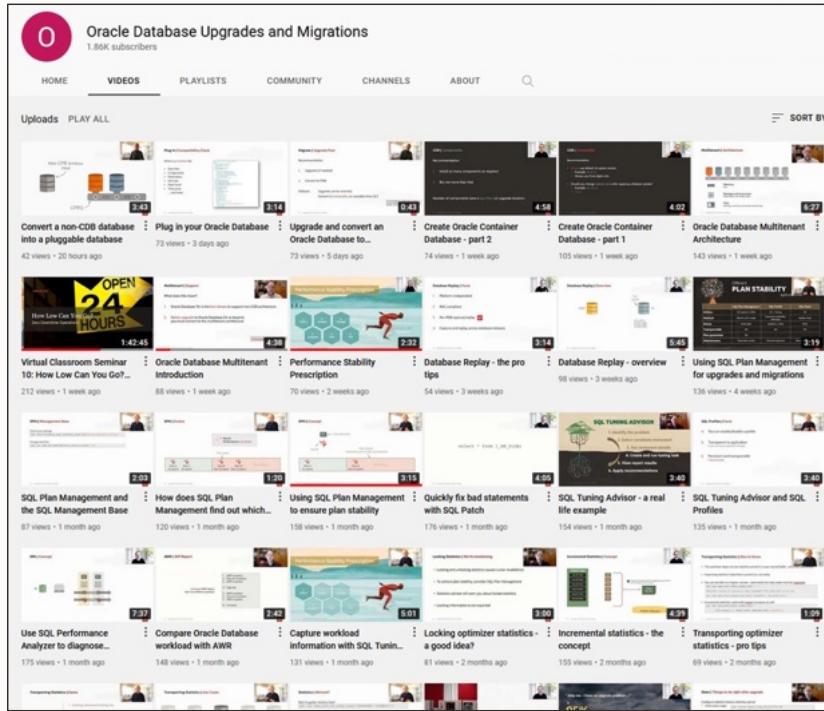
Agent v GPT-5 mini v 0

Ln 1, Col 1 Spaces: 4 UTF-8 LF {} PL/SQL (Dedicated) ADB26ai

Oracle
DBAs
run the world



YouTube | Oracle Database Upgrades and Migrations



<https://www.youtube.com/@upgradenow>

- 300+ videos
- New videos every week
- No marketing
- No buzzword
- All tech



Thank You

