



Real World Oracle Database Upgrade and Migration to Oracle AI Database 26ai

Helsinki, March 2026



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>



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



Get the Slides

<https://dohdatabase.com/slides>



Virtual Classroom Seminars

Slides

Episode 16

(replaces Episode 1 from Feb 2021)

[Oracle Database Release and Patching Strategy for 19c and 23c](#)

115 minutes – May 10, 2023



Episode 17

[From SR to Patch – Insights into the Oracle Database Development](#)

[process](#)

55 minutes – June 22, 2023



Episode 18

[Cross Platform Migration – Transportable Tablespaces to the Extreme](#)

145 min – February 22, 2024



Episode 19

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

[Multitenant PART 1](#)

145 min – May 16, 2024



Episode 20

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

[Multitenant PART 2](#)

100 min – June 28, 2024



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

March 12, 14:00 CET

[Sign up](#)



AGENDA

09:30

Welcome
Release Strategy
Patching

11:00

Upgrade
Multitenant

13:30

Data Pump
Performance Stability

15:15

Insights into development
Oracle Database 26ai
Best Features for DBAs

10:45

Coffee break

12:45

Lunch

15:00

Coffee break



Release Strategy

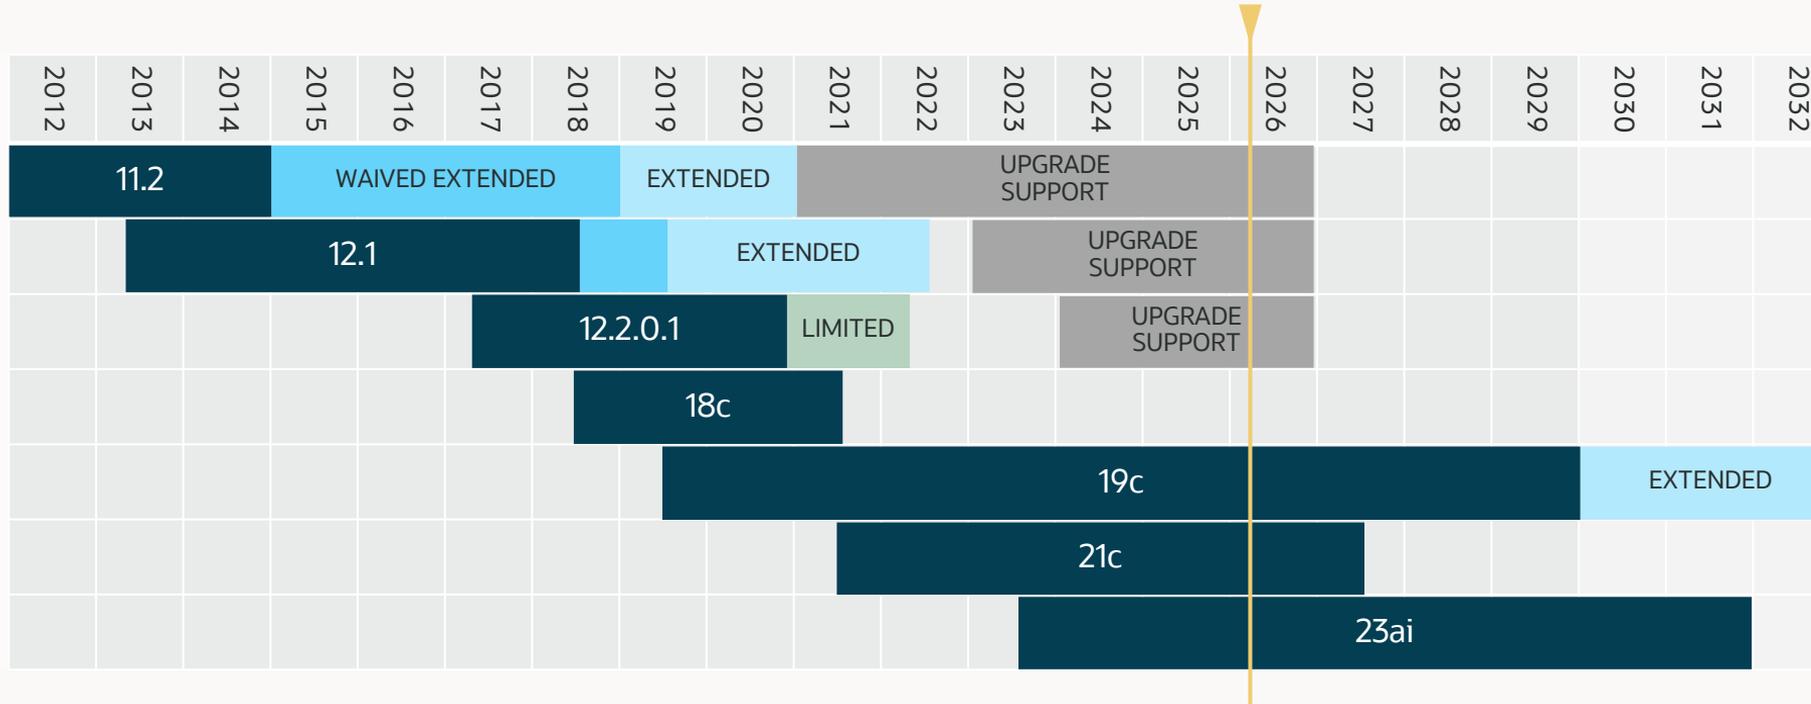


Move production databases from one Long Term Support release to the next

- Release Schedule of Current Database Releases (Doc ID [742060.1](#))



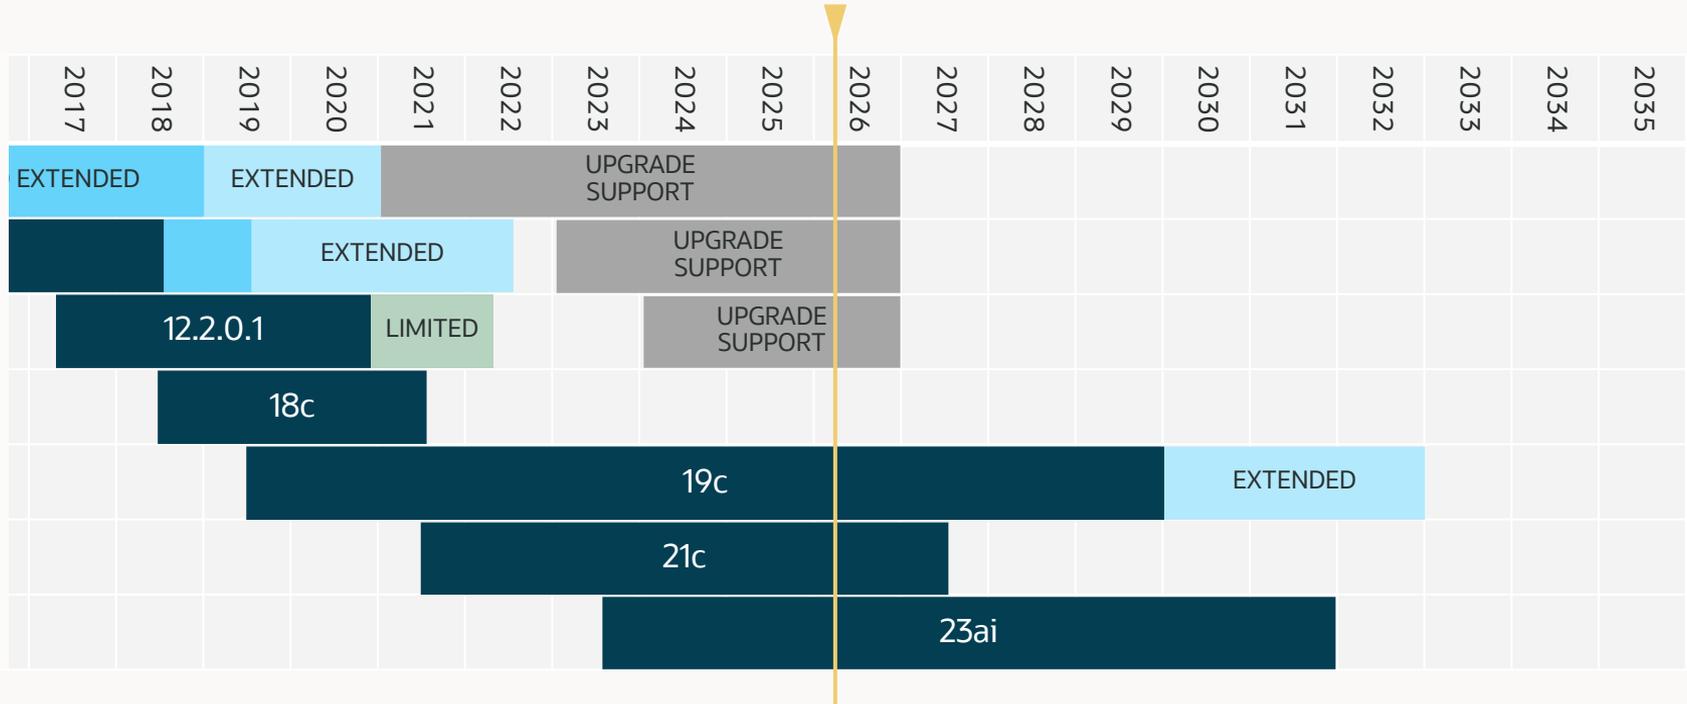
Lifetime Support Policy



Premier Support
 Waived Extended Support
 Paid Extended Support
 Restricted Upgrade Support
 Limited Error Correction



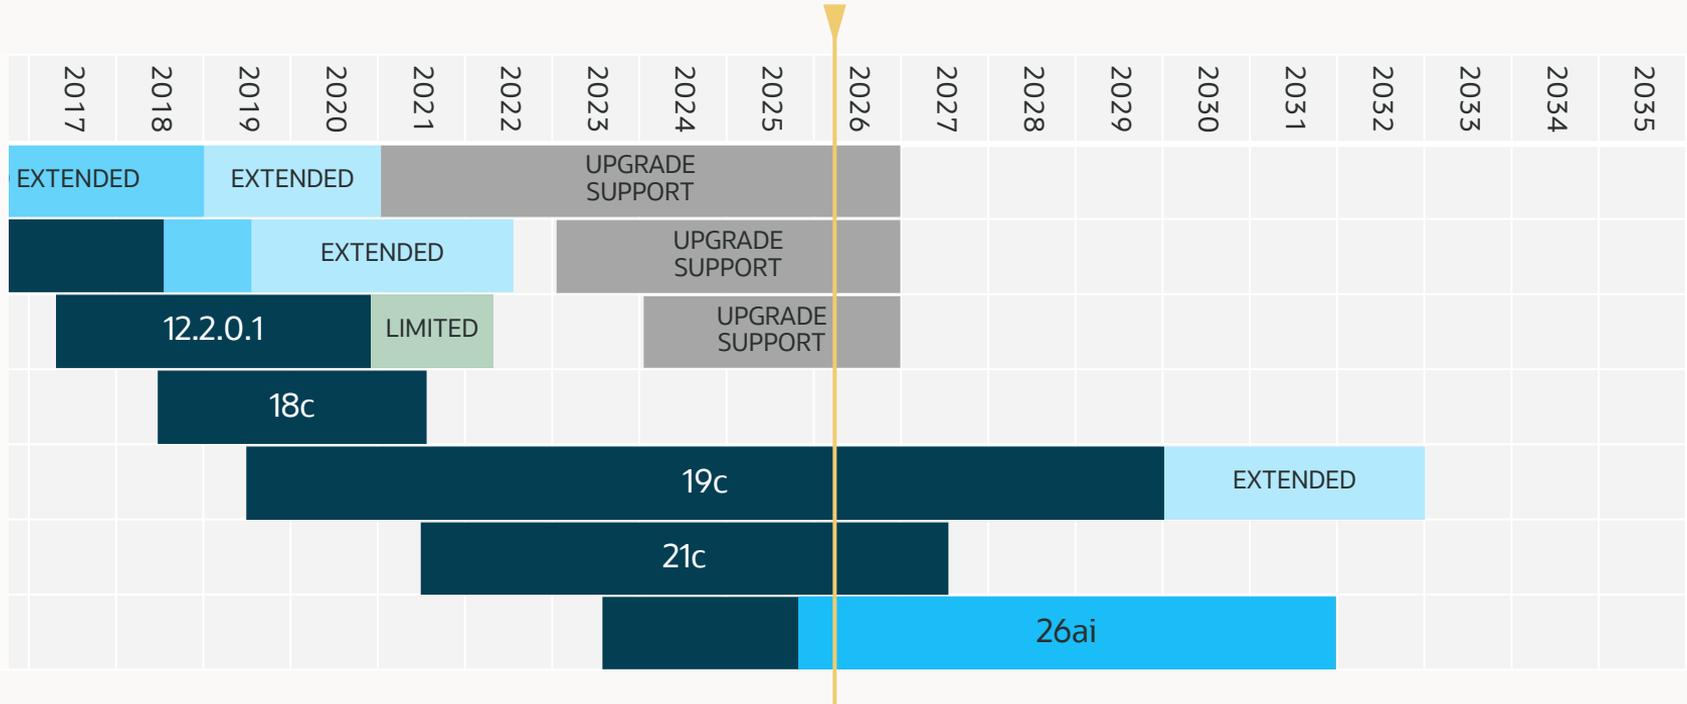
Lifetime Support Policy



Premier Support
 Waived Extended Support
 Paid Extended Support
 Restricted Upgrade Support
 Limited Error Correction



Lifetime Support Policy



■ Premier Support
 ■ Waived Extended Support
 ■ Paid Extended Support
 ■ Restricted Upgrade Support
 ■ Limited Error Correction



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 ⇒ Oracle Database 23ai ⇒ Oracle AI Database 26ai

UPGRADE

Oracle Database 19c



Oracle AI Database 26ai

UPGRADE

Oracle Database 19c ⇒ Oracle Database 23ai ⇒ 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](#)

Oracle Database Enterprise Edition

Use these downloads for developing, prototyping, and testing in a non-production environment. If you have a commercial license for Oracle AI Database, you can download all supported versions from the [Oracle Software Delivery Cloud](#).

Database Version and Platform	Download	Related Resources
Oracle AI Database 26ai for Linux x86-64	🔒 ZIP (2.2 GB) 🔒 RPM (OL9) (2.0 GB)	Installation Guide Individual Component Downloads
Oracle Database 19c for Microsoft Windows x64 (64-bit)	🔒 ZIP (2.9 GB)	Installation Guide Individual Component Downloads
Oracle Database 19c for Linux x86-64	🔒 ZIP (2.8 GB) 🔒 RPM (2.5 GB)	Installation Guide Individual Component Downloads
Oracle Database 19c for Linux ARM (aarch64)	🔒 ZIP (2.2 GB)	Installation Guide

[Link](#)

Database and Grid Infrastructure Patching

Patching Best Practices

Installation

Basics

Methods

Grid Infrastructure

Datapatch

AutoUpgrade



Until Oracle Database 19c, you always start with the base release

- Oracle Database 19.3.0





From Oracle Database 23ai onwards, you download a **Gold Image**

- Updated with latest Release Update
- You have the choice, RUs are available as well



Download the most recent RU or Gold Image

Use the Patch Download Assistant [KA958](#)

Details

Oracle Version: **26ai** Patch Type: **Oracle Database Release Updates (RUs)**

Oracle Database Release Updates (RUs)

Release	Database Update	GI Update
OCT2025 (23.26.0.0.0)	▶ 38404116	▶ 38402446
JUL2025 (23.9.0.25.07)	▶ 38048302	▶ 38041895
APR2025 (23.8.0.25.04)	37701421	37689703
JAN2025 (23.7.0.25.01)	37366180	37353687
OCT2024 (23.6.0.24.10)	37037086	37031054
JUL2024 (23.5.0.24.07)	36741532	36739878

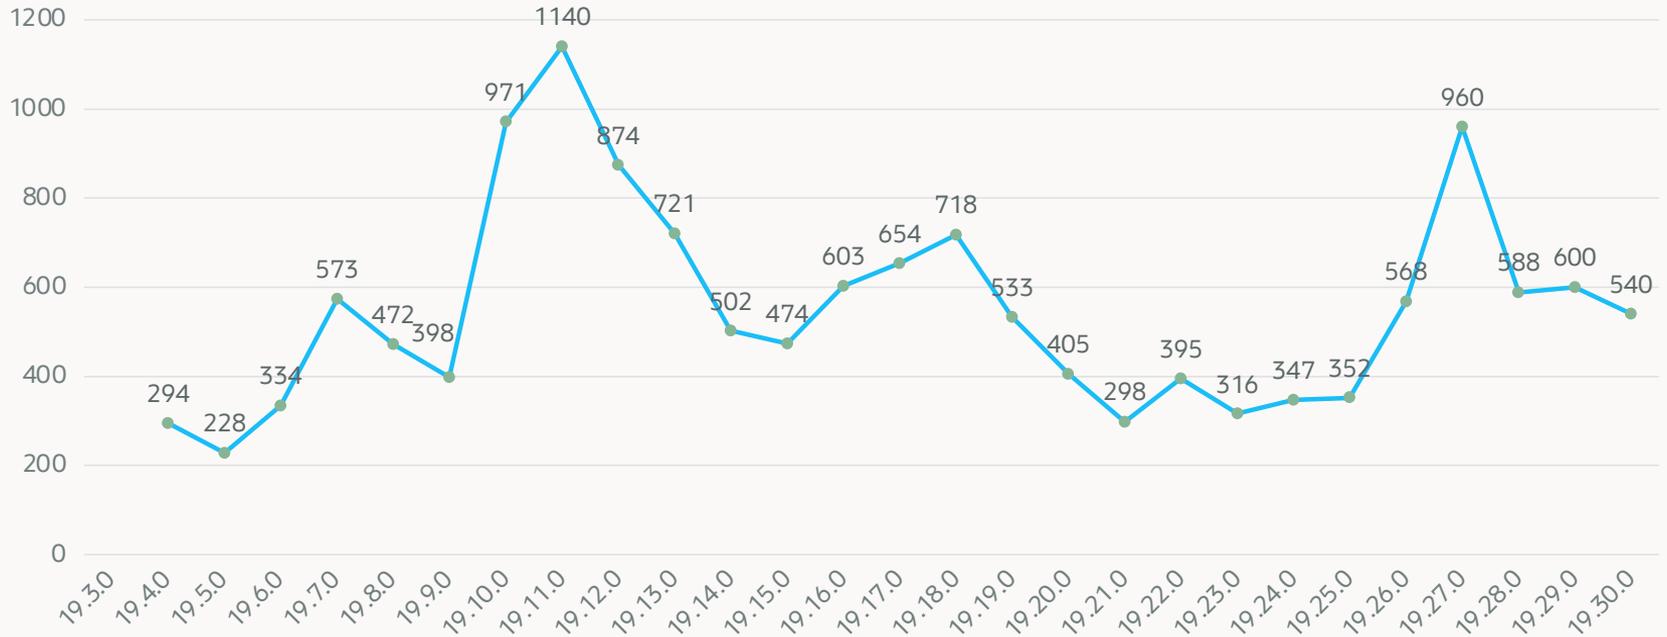
Oracle Version: **26ai** Patch Type: **Oracle Database Release Update as Gold Image**

Oracle Database Release Update as Gold Image

Release	Database Update as Gold Image	GI Update as Gold Image
OCT2025 (23.26.0.0.0)	38411960	38435666
JUL2025 (23.9.0.25.07)	38053646	38053676
APR2025 (23.8.0.25.04)	37703495	37703569
JAN2025 (23.7.0.25.01)	37370465	37370503
OCT2024 (23.6.0.24.10)	37037908	37037934



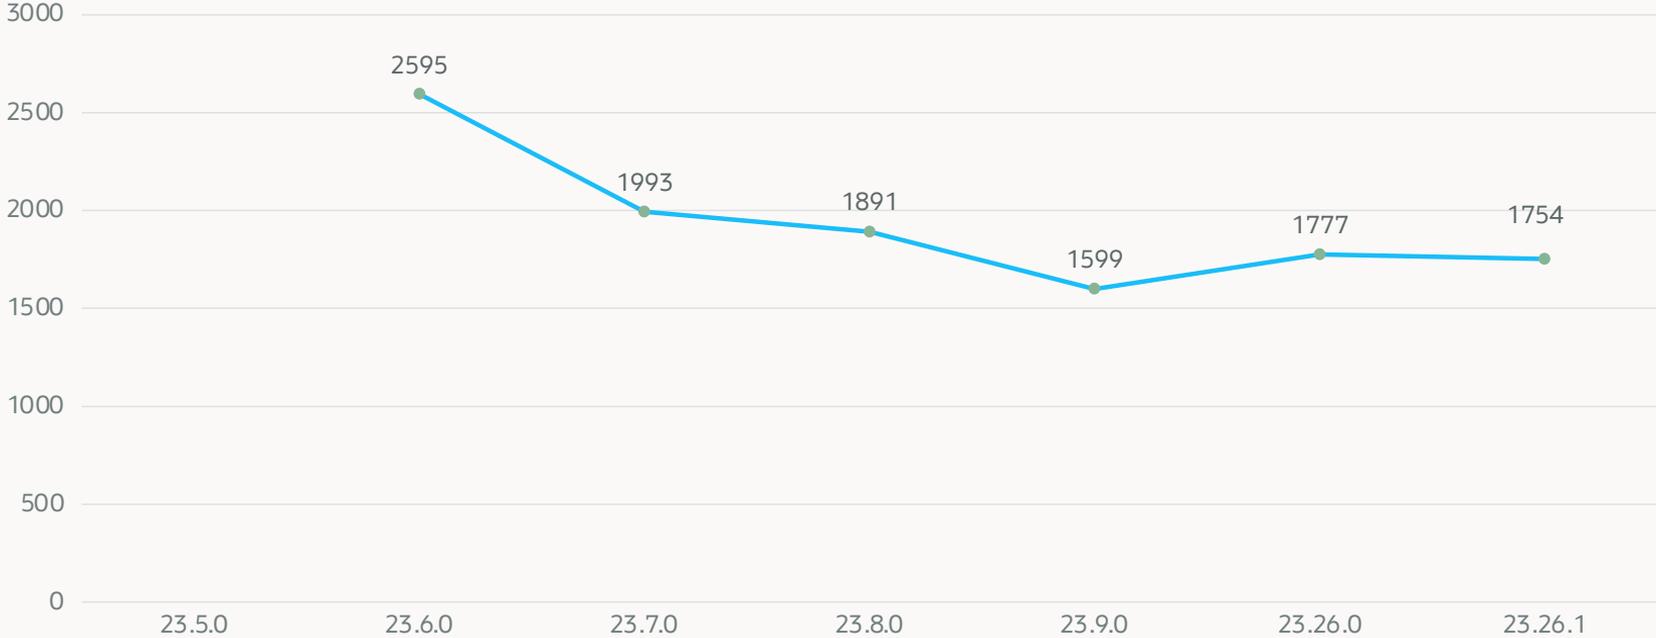
Release Update Contents – Oracle Database 19c



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



Release Update Contents – Oracle Database 26ai



Source: 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



Apply the Most Important Patches - [KB188772](#)

Recommended Patches for 23.26.0 DB Home

Below is the list of important patches to consider applying on top of 23.26.0.

Bug	KM Article	Fixed in RU	Fixed in MRP	Description	Patches	RAC Rolling Installable	Database Online Installable	Added
38765450 (replaces 37019735)	KI54683 (replaces KI39650)			[RECOVERY] LGWR 'ges inquiry response' --> 'ges remote message' hang with fix for bug 37019735 applied	PATCH 38765450	YES	YES	6- JAN-2026
38450494	KI54680			[RAC] ORA-600: internal error code, arguments: [99999]	PATCH 38450494	YES	NO	6- JAN-2026
37393792	KI40133			[PAR EXECUTION] ORA-12850 Reported after applying DBRU on First Node	PATCH 37393792	YES	YES	04- APR-2025

Version DB 23.26.0_555.1: 38765450,38450494,37393792





Oracle delivers security fixes via Release Updates

- Occasionally, also via Monthly Recommended Patches



MRP might contain some, but not all Security Fixes

- Important one-off fixes
- Monthly intervals
- Single downloadable patch

Quarterly Release Updates

	2023				2024				2025				2026		
	January	April	July	October	January	April	July	October	January	April	July	October	January	April	July
19c	19.18.0	19.19.0	19.20.0	19.21.0	19.22.0	19.23.0	19.24.0	19.25.0	19.26.0	19.27.0	19.28.0	19.29.0	19.30.0	19.31.0	19.32.0
21c	21.9.0	21.10.0	21.11.0	21.12.0	21.13.0	21.14.0	21.15.0	21.16.0	21.17.0	21.18.0	21.19.0	21.20.0	21.21.0	21.22.0	21.23.0
23ai							23.5.0	23.6.0	23.7.0	23.8.0	23.9.0	23.26.0	23.26.1	23.26.2	23.26.3



Monthly Recommended Patches

	2025			2026										
	October	November	December	January	February	March	April	May	June	July	August	September	October	November
19.29.0	19.29.0	MRP1	MRP2	MRP3	MRP4	MRP5	MRP6							
19.30.0				19.30.0	MRP1	MRP2	MRP3	MRP4	MRP5	MRP6				
19.31.0							19.31.0	MRP1	MRP2	MRP3	MRP4	MRP5	MRP6	
19.32.0										19.32.0	MRP1	MRP2	MRP3	MRP4
19.33.0													19.33.0	MRP1



Sorry, but there is more to talk about ...

Apply Additional Important Fixes and Bundles – [KA912](#)

Getting Started

Performance

Golden Gate

Oracle Text

Platform Specific

Data Guard High Availability

DN FS

Data Pump

Partitioning

Multitenant

General

Oracle Spatial

Search This Document

Print

Getting Started

When applying Database patches, Oracle recommends that you take a 3-tiered step-by-step approach.

LEVEL 1: Apply latest quarterly patches:

- Apply latest quarterly updates using [Master Note for Database Proactive Patch Program \(Doc ID KB106822\)](#)

LEVEL 2: Apply Critical/Recommended patches:

- - For Exadata environments: [Exadata Critical Issues \(Doc ID KB623062\)](#)
 - For Database environments:
 - Customers on Linux x86-64 - Apply the latest [Monthly Recommended Patches - MRP \(Doc ID FAQ2283\)](#) for the specific RU
 - For customers on other platforms, apply critical patches using [Oracle Database 19c Important Recommended One-off Patches \(Doc ID KB858385\)](#)

LEVEL 3: Apply additional patches based on features or focus areas:

- Use the tabs in this document for quick access to additional feature based patches



Always use the latest OPatch

- [Patch 6880880](#)





Patching an Oracle home is faster when you use a brand-new home

- Avoid cloned Oracle Homes and In-Place Patching
- Use `./opatch util deleteinactivepatches`

Patching Best Practices

Installation

Basics

Methods

Grid Infrastructure

Datapatch

AutoUpgrade



What Can Be in a Patch?

FILES

New or changed executables, libs or files

```
bin/oracle  
bin/srvctl  
oracore/zoneinfo/timezone_42.dat
```

Apply and rollback scripts

```
sqlpatch/.../nnn_apply.sql  
sqlpatch/.../nnn_rollback.sql
```

SQL PL/SQL

New or changed objects

```
alter table sys.tab$ ...  
create index sys.i_tab1 ...  
create or replace package sys.dbms_scheduler ...
```



How to Apply a Patch?

`opatch`



Applies binaries to an Oracle Home



All instances using this Oracle Home are down

`datapatch`



Applies SQL and PL/SQL changes to a database



Database is up



What Is Installed?

In the Oracle Home?

```
$ opatch lsinventory
```

```
$ opatch lspatches
```

```
SQL> select  
xmltransform(dbms_qopatch.get_opatch_lsinventory,  
dbms_qopatch.get_opatch_xslt) from dual;
```

- [Oracle Database 12.1 : FAQ on Queryable Patch Inventory \(FAQ2083\)](#)

In the database / PDB?

```
SQL> select * from  
cdb_registry_sqlpatch;
```



Patching Best Practices

Installation

Basics

Methods

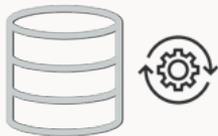
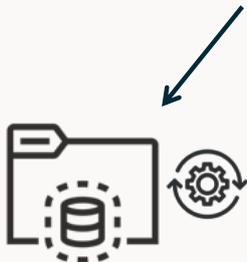
Grid Infrastructure

Datapatch

AutoUpgrade

In-Place Patching

Oracle Home, 19.30.0

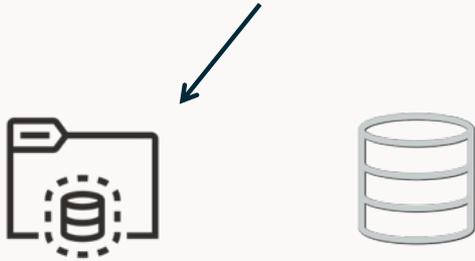


```
$ORACLE_HOME/OPatch/opatch rollback -id ...
```



Out-of-Place Patching

Oracle Home, 19.28.0



```
SQL> SHUTDOWN IMMEDIATE
```



New Oracle Home
Oracle Home, 19.30.0

```
[oracle]$ $ORACLE_HOME/OPatch/datapatch
```





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%



Always patch Out-of-Place

- Don't argue with us 😊



Reduce downtime to the time it takes to perform a switchover

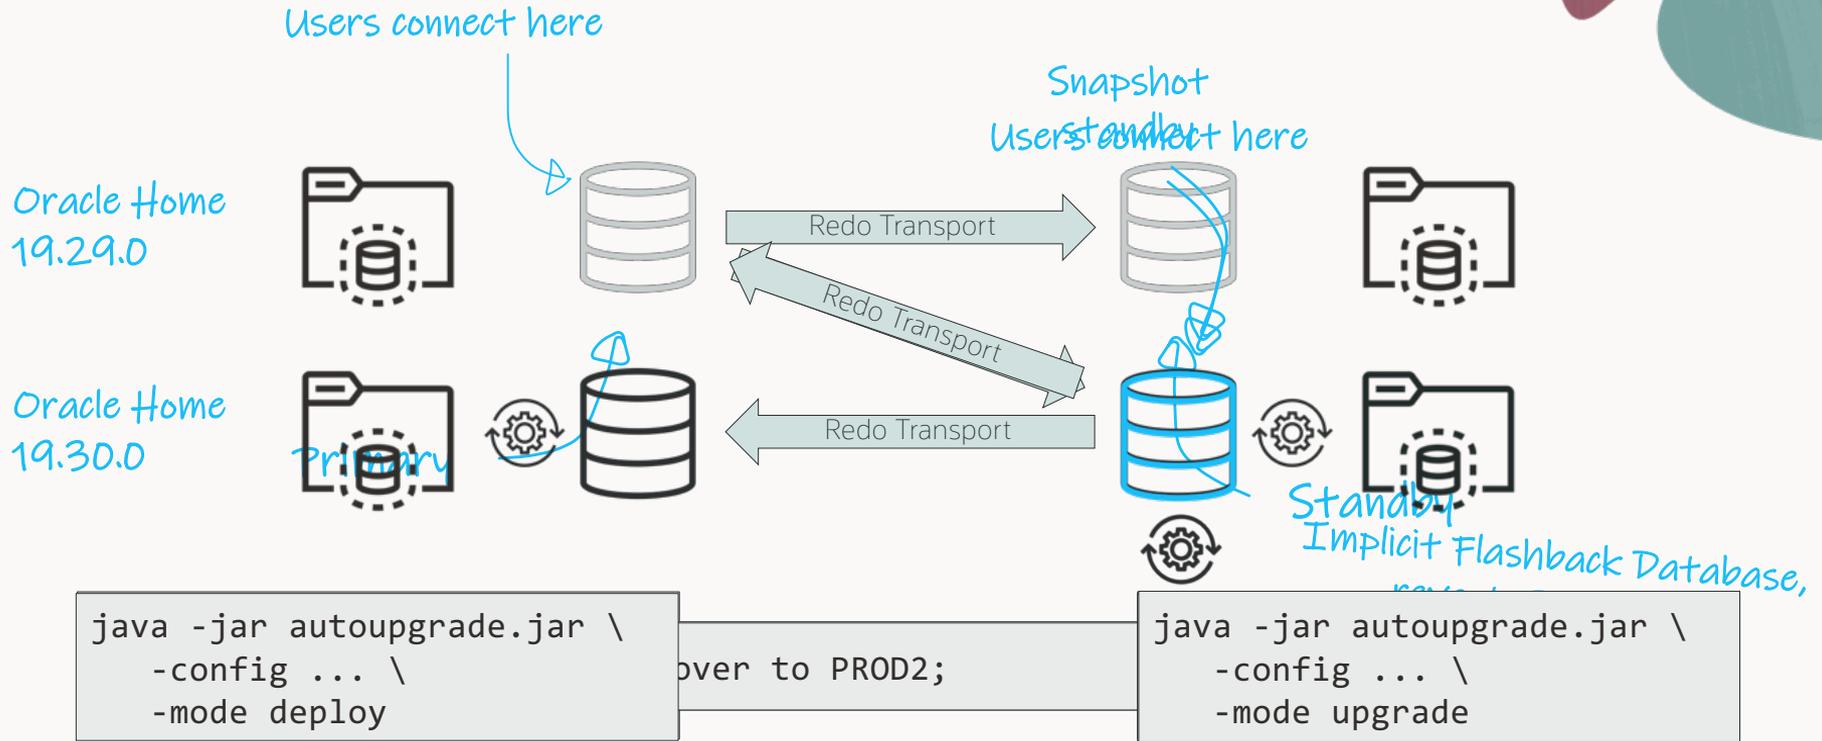
- [KB137118](#)
[Oracle Patch Assurance - Data Guard Standby-First Patch Apply](#)



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

- Check the patch readme



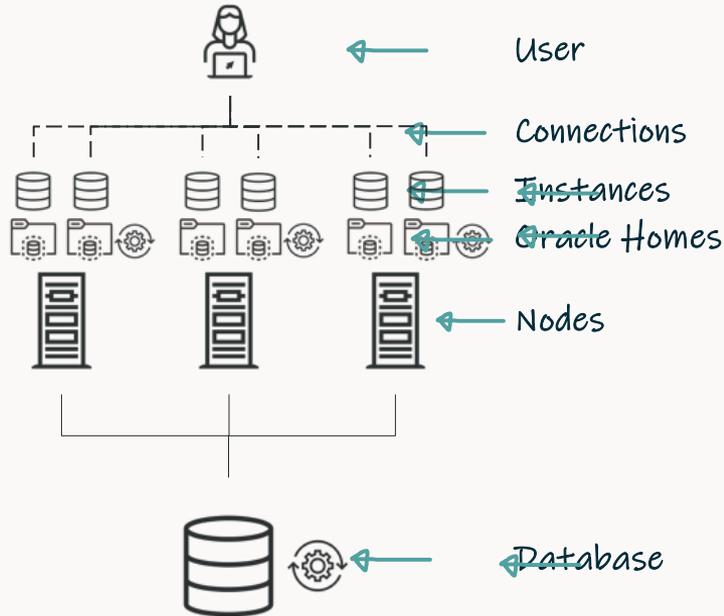
Use AutoUpgrade to patch your standby database

- 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



Release updates are **always**:



Standby-First installable



RAC Rolling installable

Patching Best Practices

Installation

Basics

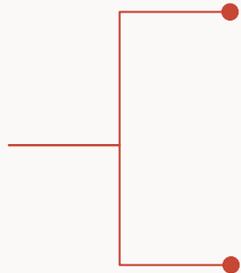
Methods

Grid Infrastructure

Datapatch

AutoUpgrade

Grid Infrastructure Patching Methods



1

In-place

Replaces existing Oracle Home

Uses opatchauto

2

Out-of-place

Creates a new Oracle Home

Uses opatchauto or SwitchGridHome



Grid Infrastructure Patching Methods



2

Out-of-place

Creates a new Oracle Home

Uses opatchauto or SwitchGridHome



--Ahead of maintenance window, create and patch new GI homes

--Optionally, create and use gold images

```
./gridSetup.sh -applyRU ... \  
                -applyOneOffs ... \  
                -clusterNodes ... \  
                ...
```



--Ahead of maintenance window, create and patch new GI homes

--Optionally, create and use gold images

```
./gridSetup.sh -applyRU ... \  
                -applyOneOffs ... \  
                -clusterNodes ... \  
                ...
```

--During maintenance window, switch to the new GI home by stopping

--the old home and start the new home. Repeat on all nodes.

```
$ORACLE_HOME/gridSetup.sh -switchGridHome ...
```

```
$ORACLE_HOME/root.sh
```



Demo

Patch a 2-node RAC system
GI and database

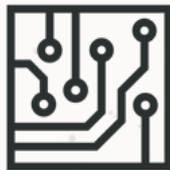
Watch on [YouTube](#)





26ai GI home disk space greatly reduced to 3 GB

- 12 GB in 19c



Use golden images

- [Blog post](#)
- Also for database homes

Should you patch Grid Infrastructure and Database Homes **together**, or **separately**?

Patching GI and DB together?

Option 1

TOGETHER

One maintenance window

Longer, single patching window

Several changes

When draining is a problem

Option 2

SEPARATELY

Two maintenance windows

Shorter window, but longer overall patching

One change at a time

For well-behaving applications





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 19.28 + DB 19.28
Within two Release Updates	Supported, good practice	GI 19.28 + DB 19.26 GI 19.26 + DB 19.28
Within three or more Release Updates	Supported, not recommended	GI 19.28 + DB 19.25 GI 19.25 + DB 19.28
Different releases	Supported, use only for upgrades	GI 23.9 + DB 19.28





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

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





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

- Node 1 with GI 19.25.0, node 2 with GI 19.28.0
- Patching node 1 on Monday, node 2 on Tuesday ...
- Mix of GI and DB versions on various nodes

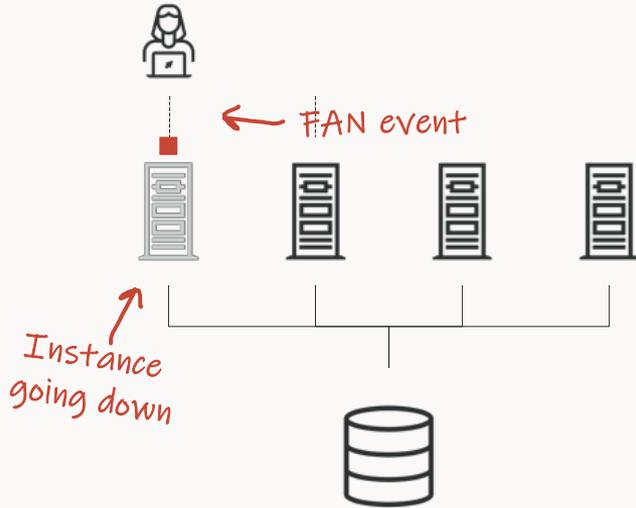


Rolling patching requires efficient draining

- Optionally, consider a *batched* approach



Draining Connections



DRAINING

- Allows users to finish their work and reconnect to another instance
- New sessions connect to other instances
- Sessions that don't drain in time are forcefully terminated
- Controlled by `drain_timeout` parameter in [`srvctl`](#) and [`DBMS_SERVICE`](#)



Drain Timeout



Setting `drain_timeout` very low?

- This may cause login storms
 - Be cautious on databases with many connections



Setting `drain_timeout` very high?

- Load is spread on fewer instances
 - Cluster is in rolling patch mode for an extended period of time



Patching Best Practices

Installation

Basics

Methods

Grid Infrastructure

Datapatch

AutoUpgrade

Patching a Database



1

Start database in new Oracle Home

Start in normal open

Open all PDBs



2

Complete patching with datapatch

Found in \$ORACLE_HOME/OPatch

One database per invocation

Multiple Datapatch sessions in parallel

[Datapatch User Guide](#)





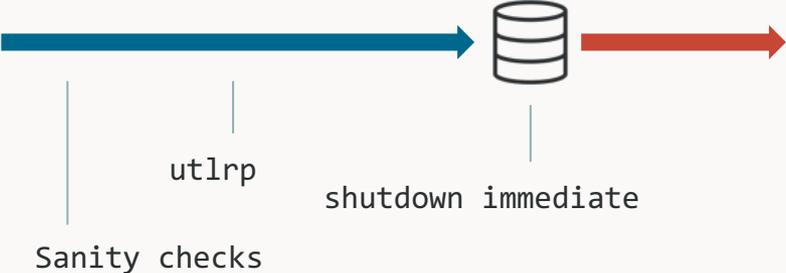
Analyze the database for patching readiness using Datapatch Sanity Checks

- [KB123801: Datapatch User Guide](#)
- Executed by AutoUpgrade in analyze mode



Recompile invalid objects
before invoking datapatch

Patching Timeline

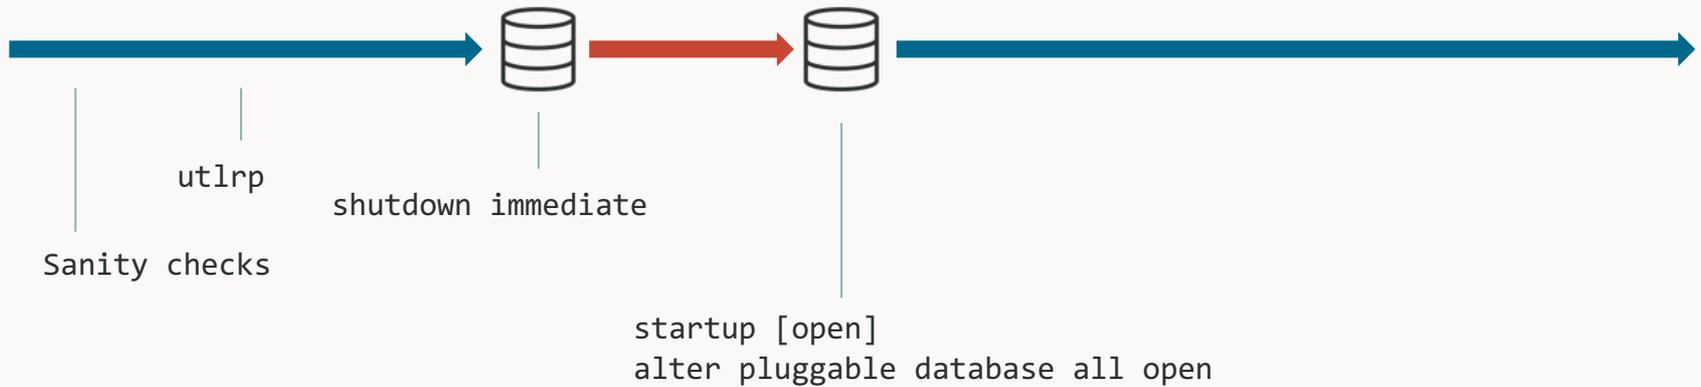




The database must be open Only open PDBs are patched

- Upgrade mode or restricted session is **not** needed

Patching Timeline

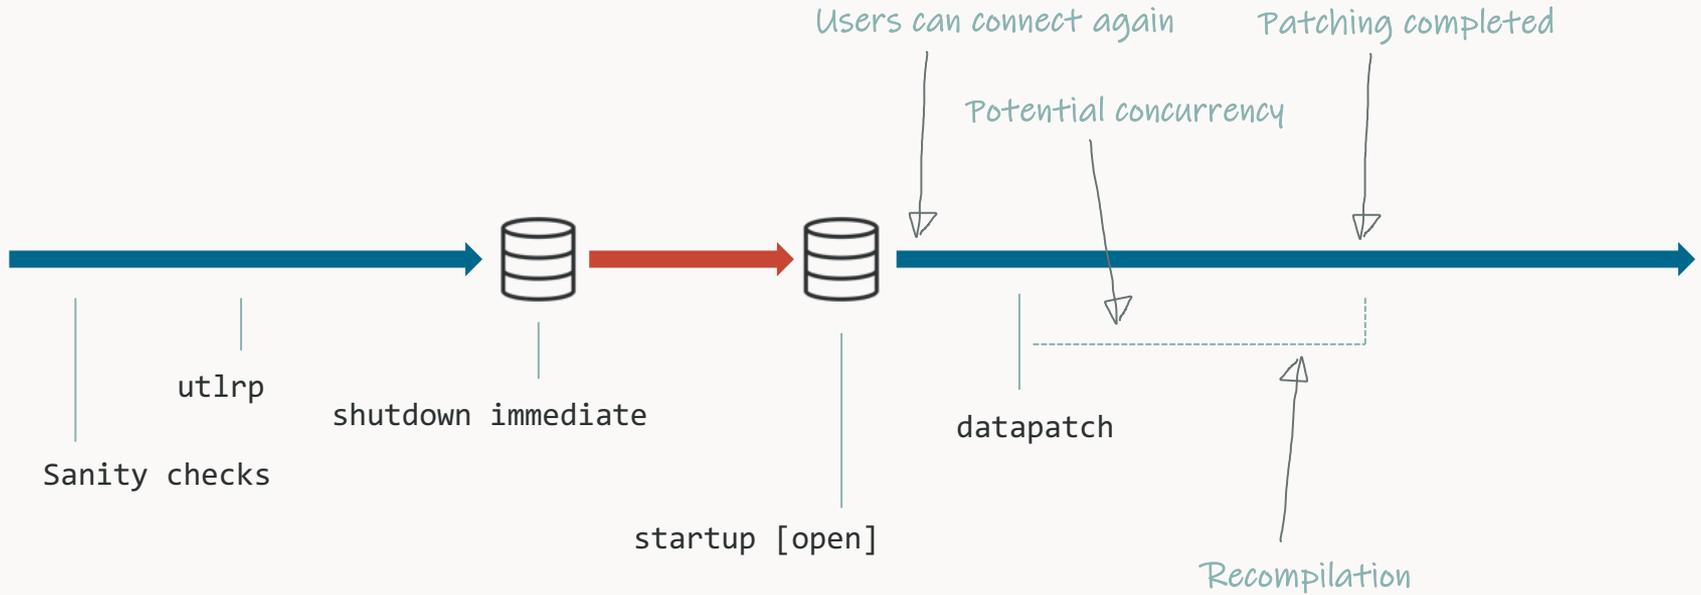




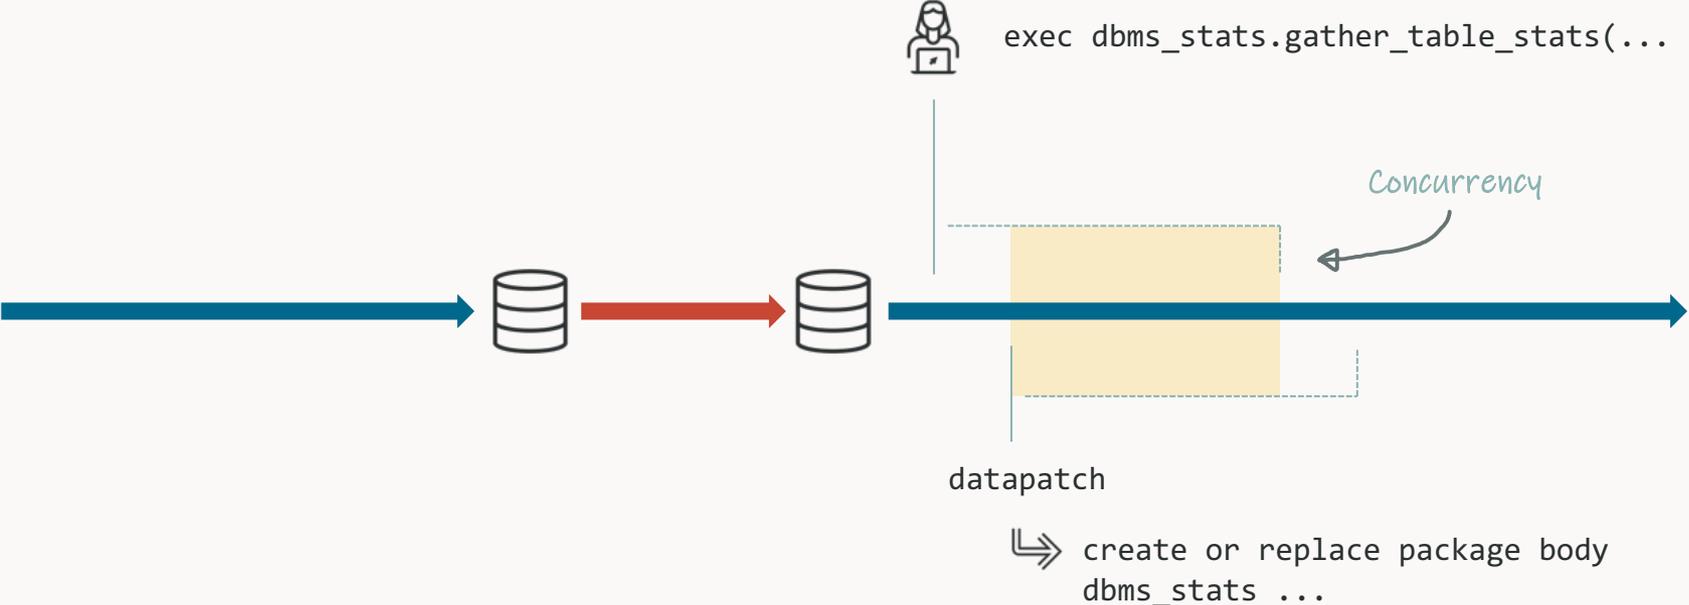
You can run datapatch while users are connected to the database

- Details in [blog post](#)

Patching Timeline



Patching Timeline

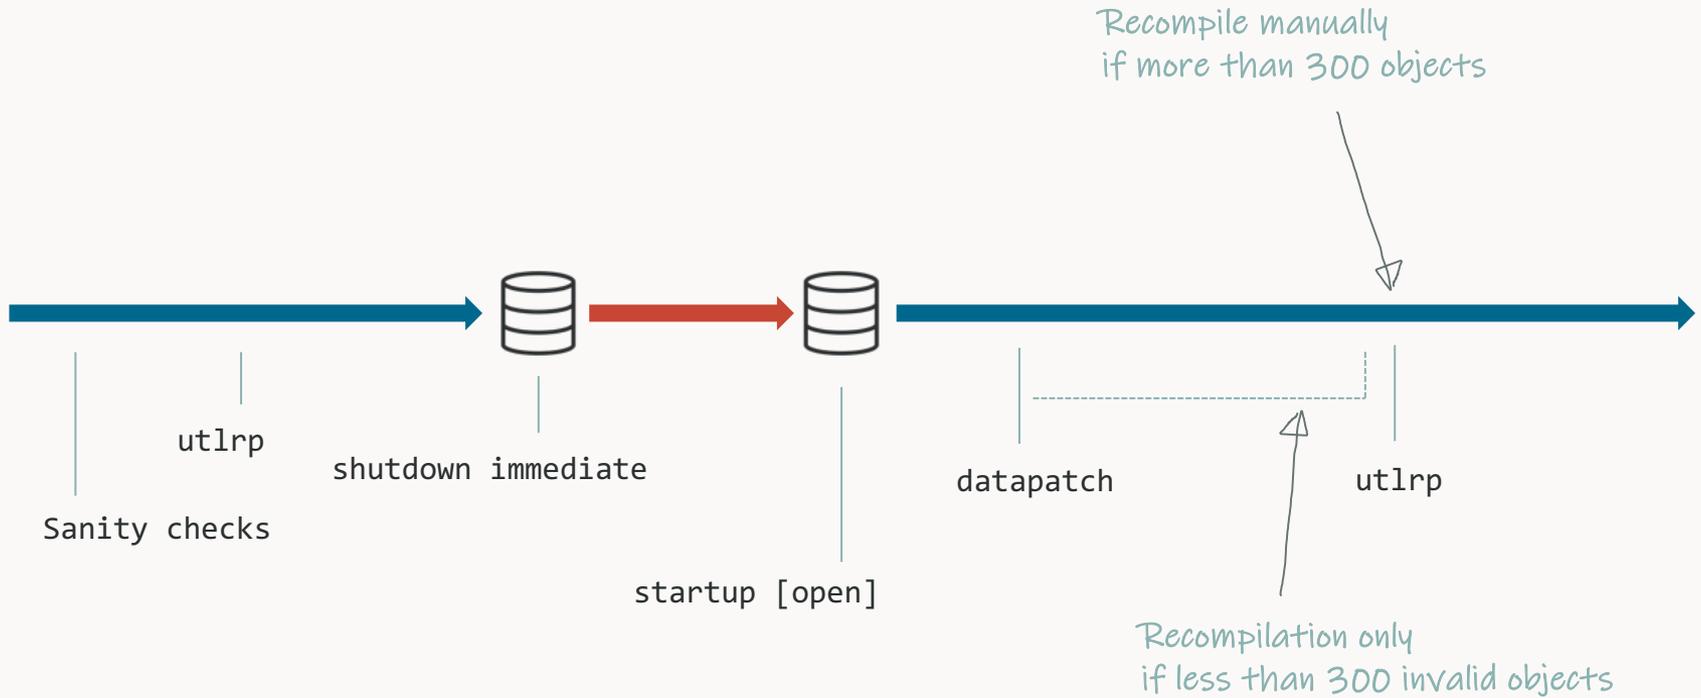


Concurrency

- 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
 - From 19.30 and 23.26.0 onwards, the [kill-switch is active](#)
- Increase timeout using `-ddl_lock_timeout <time-in-seconds>`



Patching Timeline



```
$ ./datapatch
SQL Patching tool version 19.19.0.0.0 Production on Sun Jun 25 07:12:19 2023
```

-
-
-
-

```
Automatic recompilation incomplete; run utlrp.sql to revalidate.
PDBs: PDB1 PDB$SEED
```

```
SQL Patching tool complete on Sun Jun 25 07:12:19 2023
```





Datapatch uses
`REGISTRY$SQLPATCH_RU_INFO` to
control the patching operations



If in doubt run datapatch again

- Datapatch only does what is needed
- You can run **datapatch** as many times as you like

Datapatch Rollback Scripts



Apply/rollback scripts:

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

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



Rollback scripts (zipped as BLOB):

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





Update database directories using
`rdbms/admin/utlfixdirs.sql`

- AutoUpgrade executes the script

Patching Best Practices

Installation

Basics

Methods

Grid Infrastructure

Datapatch

AutoUpgrade

AutoUpgrade



- 1 Upgrading
- 2 Non-CDB to PDB
- 3 Patching



AutoUpgrade

Upgrading

Non-CDB to PDB

Patching

Oracle
Database 19c



Oracle AI
Database 26ai



AutoUpgrade

Upgrading

Non-CDB to PDB

Patching

Non-CDB



Multitenant



AutoUpgrade

Upgrading

Non-CDB to PDB

Patching

23.9.0



23.26.1





One single tool for everything
- on all platforms

-- Always use the latest version of AutoUpgrade

-- Available as a direct download from oracle.com or MOS (Doc ID 2485457.1)

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

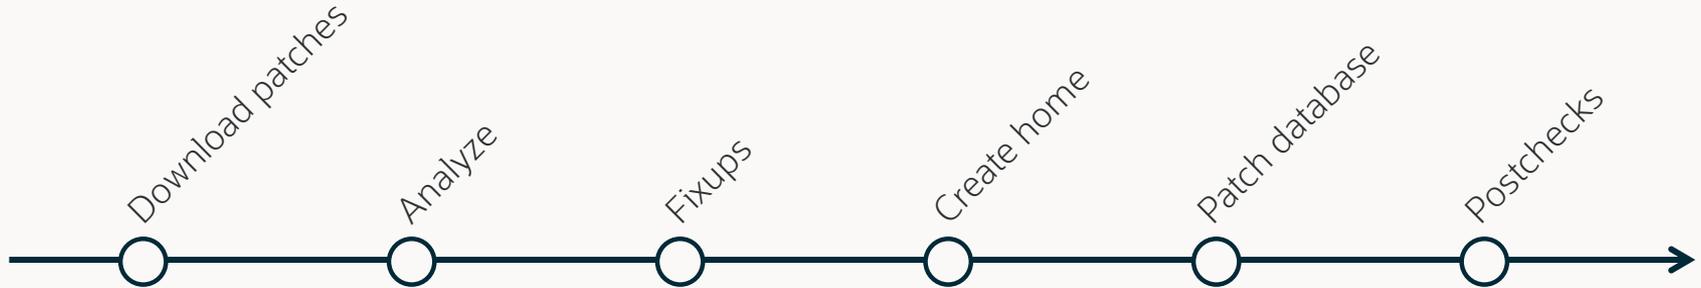




Simplify downloading and installing Oracle software and patches

- Oracle AI Database 26ai comes as fully updated gold images

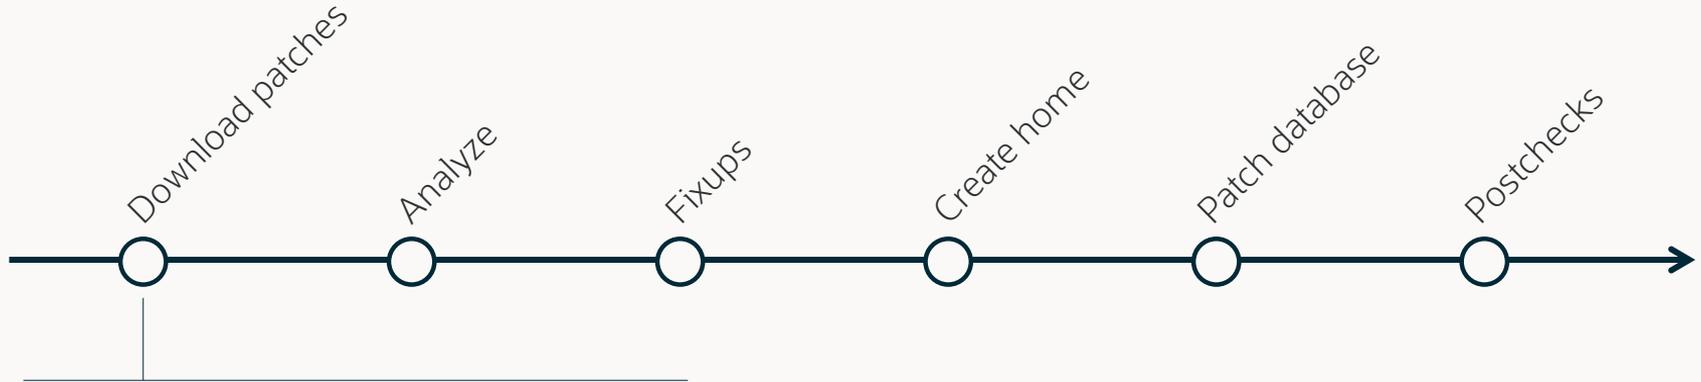
AutoUpgrade Patching



`-mode deploy`



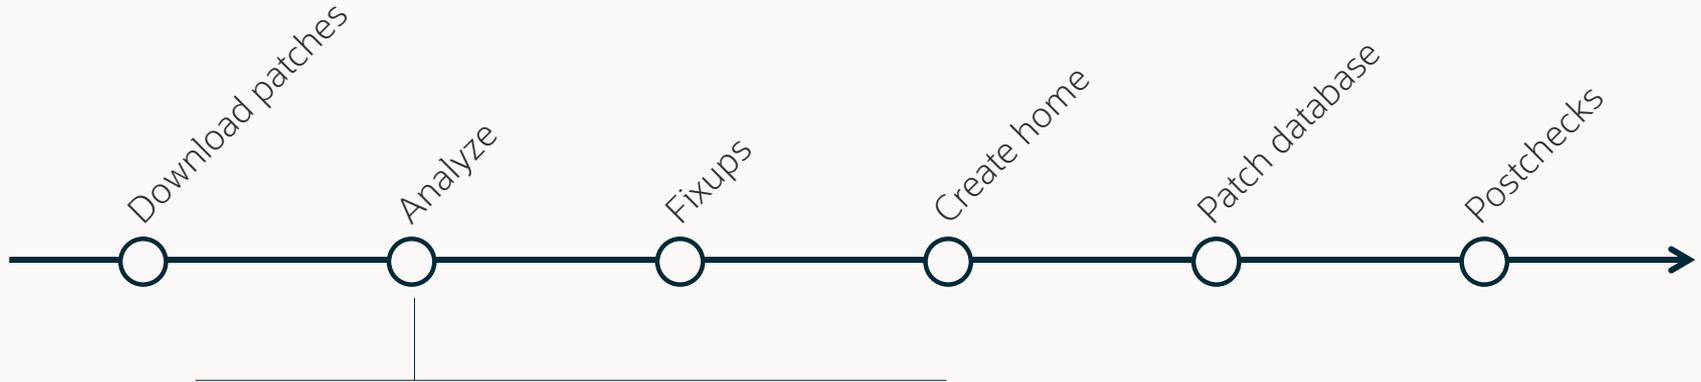
AutoUpgrade Patching



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



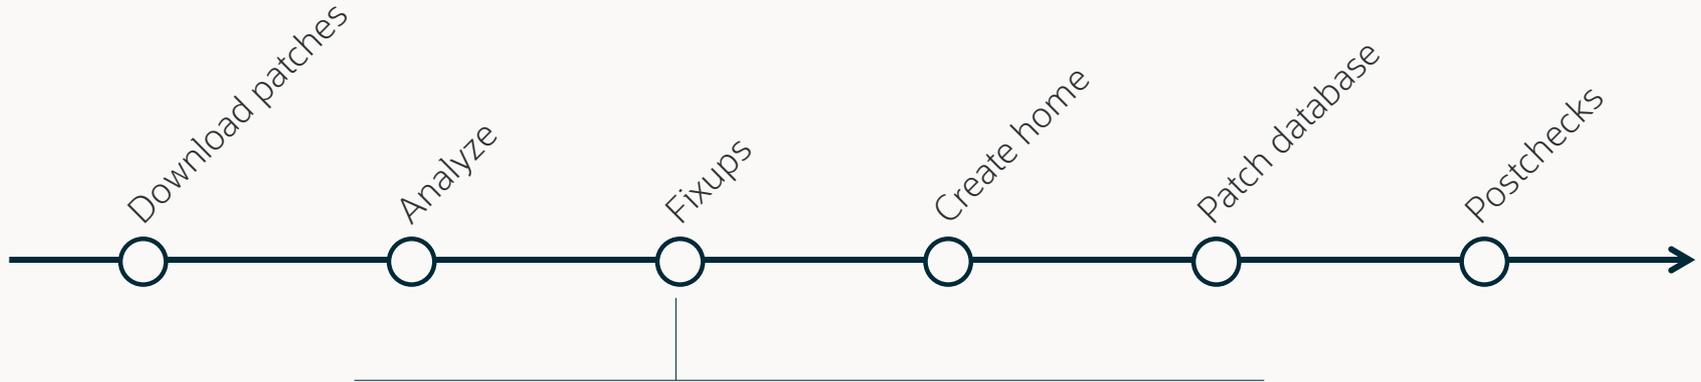
AutoUpgrade Patching



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



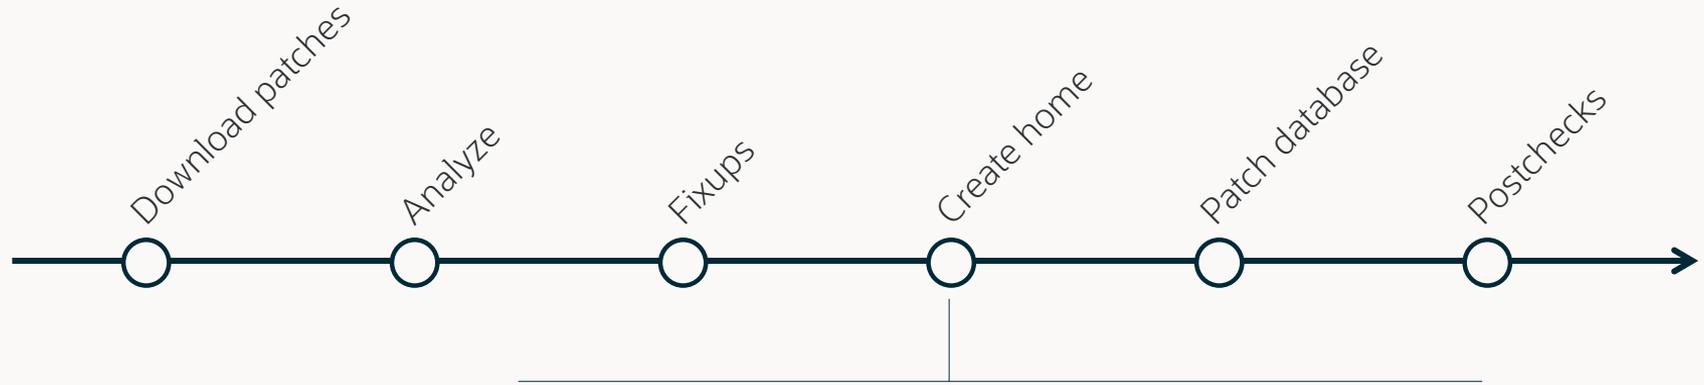
AutoUpgrade Patching



- Gathers dictionary statistics if needed
- Recompiles Oracle-maintained objects if needed
- Executes checks - see [KB142594](#)



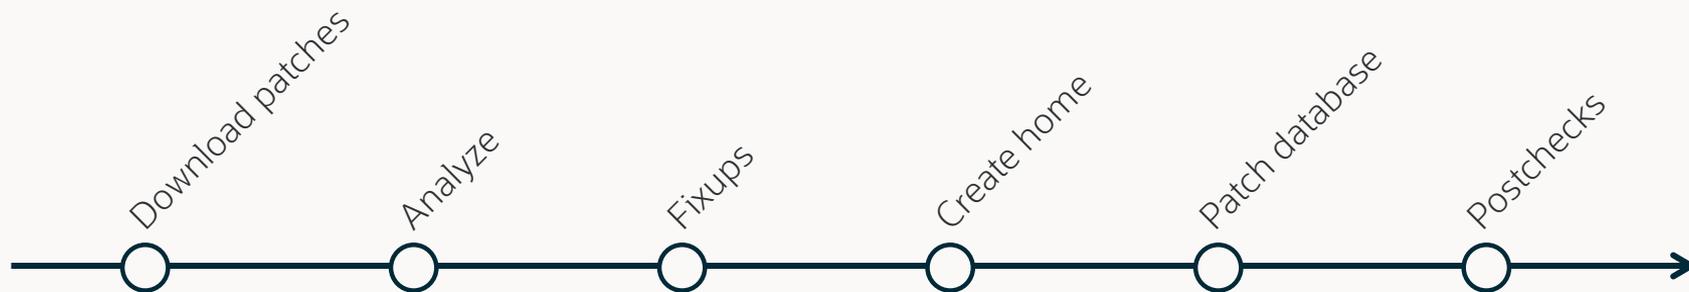
AutoUpgrade Patching



- 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`



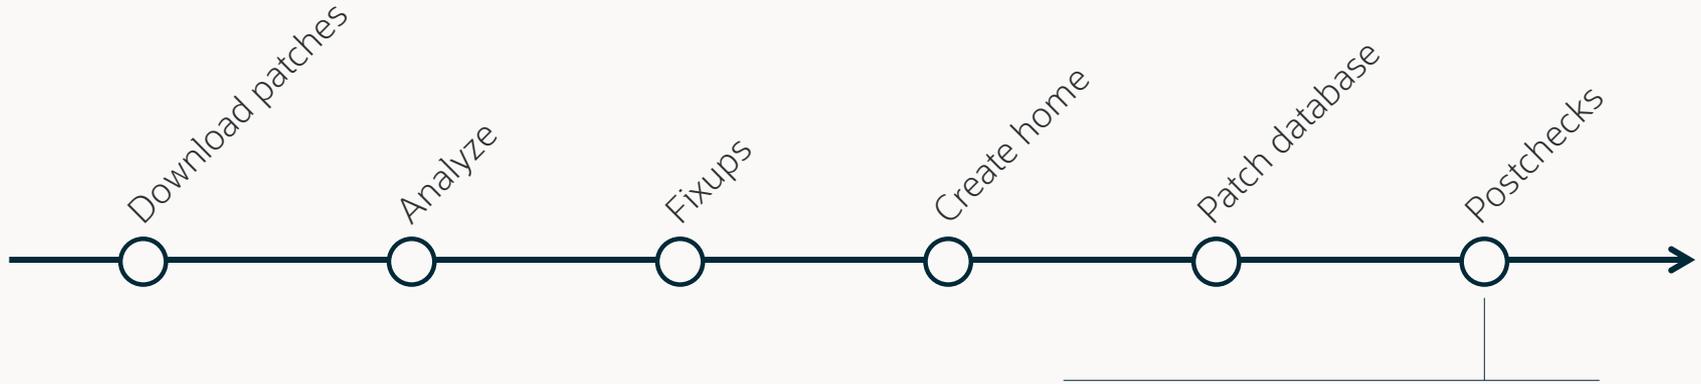
AutoUpgrade Patching



- 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



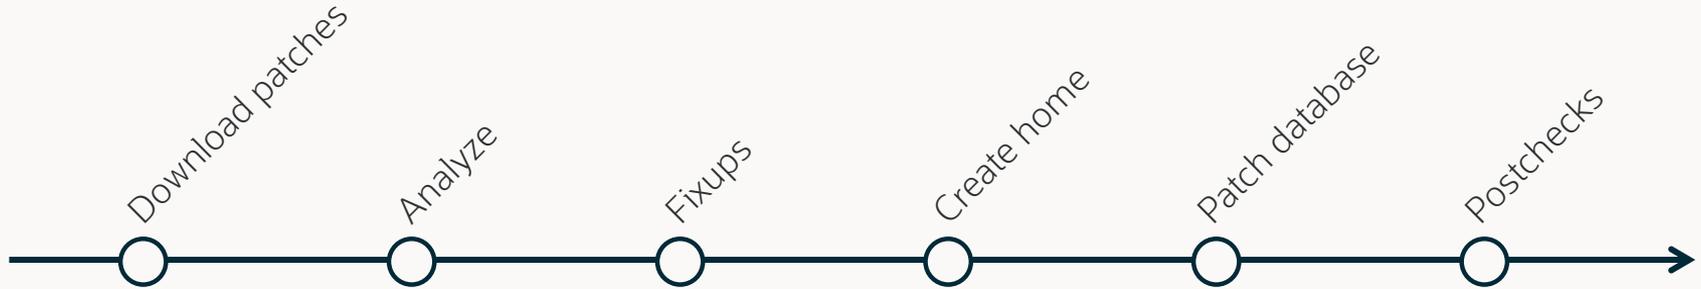
AutoUpgrade Patching



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



Patching Modes



`-mode download`

`-mode analyze`

`-mode fixups`

`-mode deploy`

`-mode create_home`





Which patches should you install?

```
global.keystore=/home/oracle/autoupgrade-patching/keystore
patch1.source_home=/u01/app/oracle/product/19/dbhome_19_27_0
patch1.target_home=/u01/app/oracle/product/19/dbhome_19_28_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



Important MOS Notes

Assistant: Download Reference for Oracle Database Patchsets and Base Releases (Doc ID 2118136.2)
Availability: EXTERNAL

Selection(s)

What would you like to download?

- Oracle Database Base Releases
- Oracle Database Patchsets
- Oracle Database Release Updates (RUs)
- Oracle Database Release Update Revisions (RURs)
- Oracle Database PSU, SPUJ(CPU), Bundle Patch
- OJVM Update/PSU/Bundle Patches
- Latest Available Microsoft Windows Patches
- Monthly Recommended Patches (MRPs)

Download Assistant

[KA958](#)

Download Release Update, OJVM, Bundles

Recommended Patches for 19.25 DB Home
In the list of important patches to consider applying on top of 19.25. In addition to the table, also review patches in Database Patchset (Doc ID 2118136.2), External Security Patches (Doc ID 2118136.2) and Patches to Consider for 19c (Doc ID 2118136.2), which contains patches to consider for Golden Gate etc.

Bug	Fixed First in RU MRP	Description	SQL Patch					
34272167	Yes	[748] NIGHT (SQL) (ORA-1970) oracle data block corrupted on tempfile (seen after 200A262), 20170916	Doc	Doc	Doc	Doc	Doc	Doc
34622088	Yes	[VCS] 20058: OSA-800: soft network error, arguments: Application [see priority related], 240801	Doc	Doc	Doc	Doc	Doc	Doc
34728627	Yes	[A2] OSA-7445 in Plug Oracle Net8 High-CPU usage in TNS Server	Doc	Doc	Doc	Doc	Doc	Doc
34113933	Yes	[CPU] OPMON2 (ORA-51475) failing with error ORA-1422 when Gathering Stats for Users Table	Doc	Doc	Doc	Doc	Doc	Doc

Important One-Offs

[KB188772](#)

Add critical fixes on top for known issues

Database Patches to Consider for 19c (Doc ID 2118136.2)

Doc ID	Performance	GoldenGate	Oracle Text	Platform Specifics	Data Guard/High Availability	Other
2118136.2	Yes	Yes	Yes	Yes	Yes	Yes

Applying Database patches, Oracle recommends that you take a 5-step step-by-step approach:

1. Apply latest quarterly patches:
Apply latest quarterly updates using Master Note for Database Patchset Program (Doc ID 2118136.2)
2. Apply Critical/Recommended patches:
For Exadata environments: Exadata Critical Issues (Doc ID 1755994.6)
For Database environments:
Customers on Linux x86_64 - Apply the Latest Monthly Recommended Patch
For customers on other platforms, apply critical patches using Oracle Database Patchset Program (Doc ID 2118136.2)
3. Apply additional patches based on features or needs across:
Use the table in this document for quick access to additional feature-based patches.

Additional Fixes

[KA912](#)

Important fixes such as GoldenGate, Spatial, Data Pump Bundle, Text and more

Sign to Consider to Avoid SQL Performance Problems on 19c (Doc ID 2118136.2)

Document

Doc ID: 2118136.2

Doc Title: Sign to Consider to Avoid SQL Performance Problems on 19c (Doc ID 2118136.2)

Doc Type: Knowledge Base Article

Doc Status: Published

Doc Version: 1.0

Doc Language: English

Doc Category: Database

Doc Subcategory: Performance

Doc Keywords: SQL Performance, 19c, Avoid Problems

Doc Description: This document provides information on how to avoid SQL performance problems on Oracle Database 19c. It includes a list of patches to consider for 19c, a list of patches to consider for 19c, and a list of patches to consider for 19c.

Doc Author: Oracle Support

Doc Last Modified: 2019-09-16

Doc URL: [https://www.oracle.com/technetwork/database/enterprise-edition/19c-sign-to-consider-to-avoid-sql-performance-problems-on-19c-2118136-2.pdf](#)

SQL Performance

[KB138000](#)

Avoid SQL performance problems. Has links to "wrong results", "SPM" and more included



DEMO

Download patches
Install Oracle home

 [Watch on YouTube](#)

JDK kit in \$ORACLE_HOME is always 3 months old

How to solve this with AutoUpgrade?

```
$ cat get-software.cfg
```

```
global.global_log_dir=/home/oracle/autoupgrade/logs
```

```
global.keystore=/home/oracle/autoupgrade/keystore
```

```
install1.target_version=23
```

```
install1.patch=RECOMMENDED,37393792,JDK
```

```
install1.folder=/u01/app/oracle/software
```



```
$ cat DB19.cfg
```

```
global.keystore=/home/oracle/autoupgrade-patching/keystore  
patch1.source_home=/u01/app/oracle/product/19/dbhome_19_27_0  
patch1.target_home=/u01/app/oracle/product/19/dbhome_19_28_0  
patch1.sid=DB19  
patch1.folder=/home/oracle/autoupgrade-patching/patch  
patch1.patch=RECOMMENDED,MRP
```



KB188772

Recommended Patches for 19.25 DB Home

Below is the list of important patches to consider applying on top of 19.25. In addition to the relevant patches listed below, you should also review patches in Database PSU/BP/Update/Revision - [Known Issues Primary Note \(Doc ID 1227443.1\)](#) and [Oracle Database Patches to Consider for 19c \(Doc ID 2781612.2\)](#) which contains patches to consider for specific areas such as Data Pump, Golden gate etc.

Bug	Fixed in RU	Fixed in MRP	Description	Patches	RAC Rolling Installable	Database Online Installable	Added
36760879	19.26		[OSS SRV LAYER] Exadata: Warm up flash cache for disks that are being synced	[list-patches]	YES		11-JAN-2025
36006239	19.26		[VOS] Exadata: Flash Cache Space Pressure After Cell Goes Down	[list-patches]	YES		11-JAN-2025
37166484	19.26	DBMRP 19.24.0.0.241217 , DBMRP 19.25.0.0.241217	[SPACE LOB] To defer foreground materialize(chunks movement ufs->cfs)	[list-patches]	YES	YES	28-NOV-2024
35197819	19.26	DBMRP 19.24.0.0.241119 , DBMRP 19.25.0.0.241119 , GIMRP 19.24.0.0.241119 , GIMRP 19.25.0.0.241119	[ASM] CSSD Leaking Memory When There Are Voting File Add/Drops	[list-patches]	YES	YES	19-OCT-2024
34672698 (replaces 34286265)		Not Applicable	[VOS] DB50: ORA-800: soft external error, arguments: [set priority failed], [vktm] , dism(16)	[list-patches]	YES		22-FEB-2024
34774667			[AQ] ORA-7445 in Purge Queue Table / High CPU usage in SVCB Service	[list-patches]			11-JUL-2023
29213893		Not Applicable	[QRY OPTIMIZER] DBMS_STATS Failing With Error Ora-1422 When Gathering Stats for User\$ Table	[list-patches]			01-SEP-2021

Version DB 19.25_555.1: [36760879](#),[36006239](#),[37166484](#),[35197819](#),[34672698](#),[34774667](#),[29213893](#)

This list may not be conflict-free!

```
$ cat DB19.cfg
```

```
global.keystore=/home/oracle/autoupgrade-patching/keystore  
patch1.source_home=/u01/app/oracle/product/19/dbhome_19_27_0  
patch1.target_home=/u01/app/oracle/product/19/dbhome_19_28_0  
patch1.sid=DB19  
patch1.folder=/home/oracle/autoupgrade-patching/patch  
patch1.patch=RECOMMENDED,36006910,36908826,35398148,36916250,  
36273767,34672698,34774667,29213893
```



```
$ cat DB19.cfg
```

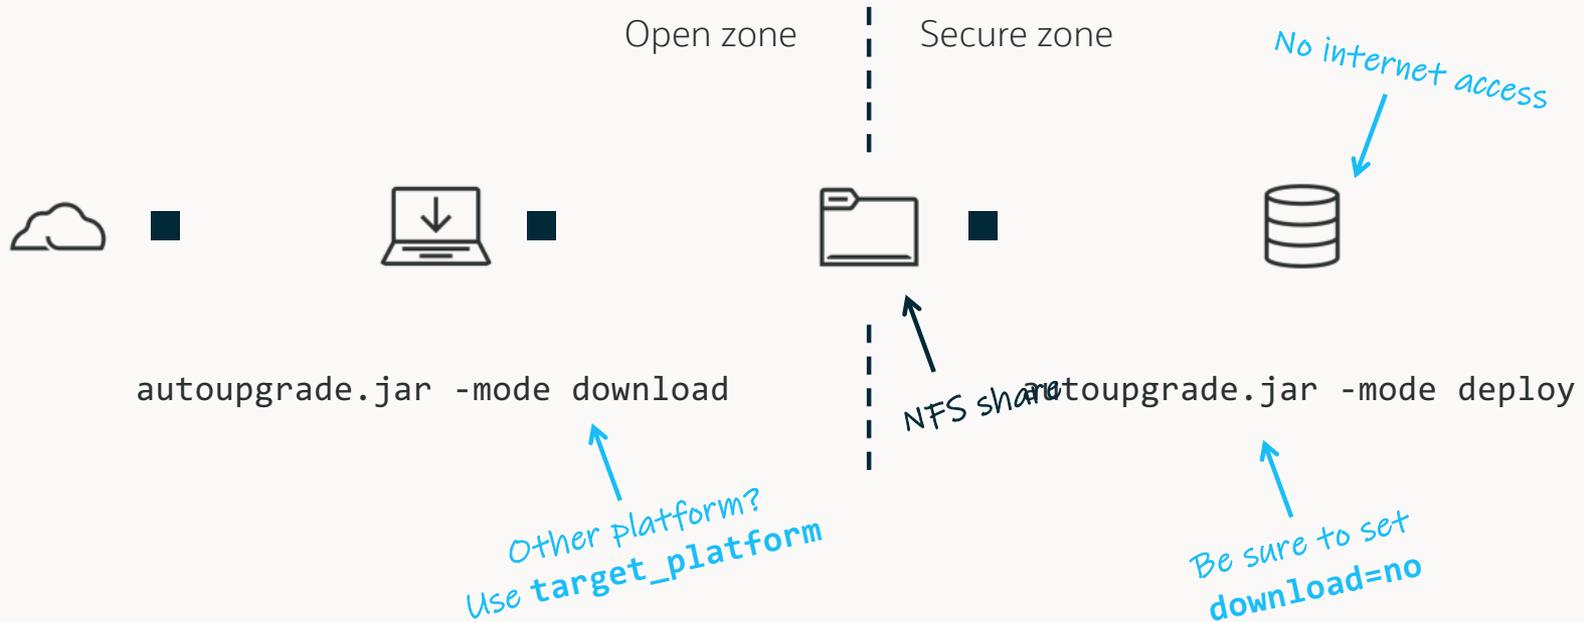
```
global.keystore=/home/oracle/autoupgrade-patching/keystore  
patch1.source_home=/u01/app/oracle/product/19/dbhome_19_27_0  
patch1.target_home=/u01/app/oracle/product/19/dbhome_19_28_0  
patch1.sid=DB19  
patch1.folder=/home/oracle/autoupgrade-patching/patch  
patch1.patch=RECOMMENDED,OCW
```





Your database host
doesn't have internet access?

Using Download Mode





Without internet, AutoUpgrade applies the recent-most RU found in *patch* folder

- Unless you explicit specify a different RU



We made upgrading easy. Now we make patching just as easy.

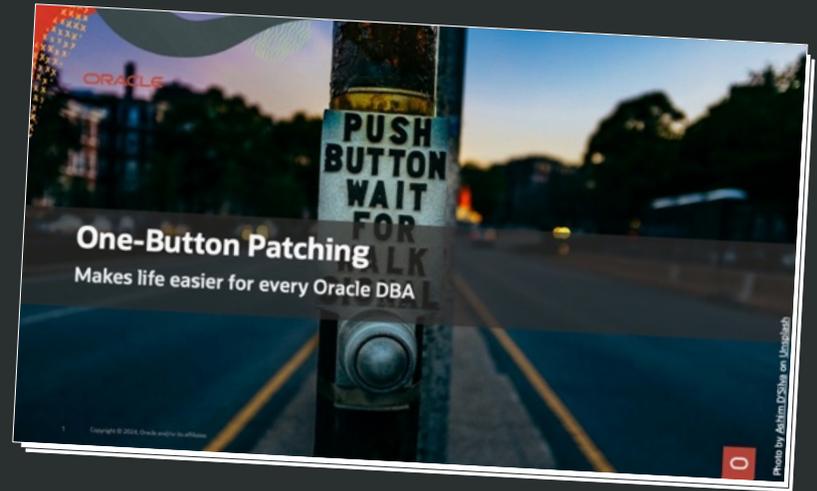
AutoUpgrade functionality extended to patching

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](#)

Key Learnings



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



Break

We start again at 11:15





Before upgrade

How to upgrade and convert

After upgrade



Your Road to Oracle AI Database 26ai

ANALYZE ⇒ **UPGRADE** ⇒ **GO-LIVE**



Oracle AI Database 26ai supports the multitenant architecture only

- You must convert your database to a PDB

--Use up to 3 user-created PDBs
--without a license for Multitenant option.
--Applies to Oracle Database 19c and newer, including SE2

```
alter system set max_pdb=3;
```





Generally, you don't need to change your application to use a pluggable database



Single vs. Multitenant



Single Tenant

One PDB
No extra license



Multitenant

Multiple PDBs
Extra license if more than 3 PDBs





Update Exadata System Software to the latest version

- Exadata Database Machine and Exadata Storage Server Supported Versions ([KB153930](#))



Upgrade to Oracle Grid Infrastructure 26ai

- Manages Oracle Database 19c and newer
- Upgrade older databases in advance

--Grid Infrastructure 26ai requires a minimum compatible setting

--Databases older than 19c may not use this disk group anymore

```
alter diskgroup DATA set attribute 'compatible.rdbms' = '19.0.0.0.0';
```





Confirm your clients can connect to Oracle AI Database 26ai

- Upgrade your clients well in advance of the upgrade

Client / Server Interoperability

Client Version	Server Version						
	26ai	21c	19c	18c	12.2.0	12.1.0	11.2.0
26ai#11	Yes	Yes	Yes	No	No	No	No
21c	Yes	Yes	Yes	Was	Was	Yes#12	No
19c	Yes	Yes	Yes	Was	Was	Yes#12	Yes#9
18c	No	Was	Was	Was	Was	Was	Was#9
12.2.0	No	Was	Was	Was	Was	Was	Was
12.1.0	No	Yes#12	Yes#12	Was	Was	Yes#12	Yes#12
11.2.0	No	No	Yes#9	Was#9	Was	Yes#12	Yes#9

[KB141043 - Client / Server Interoperability Support Matrix](#)



Client / Server Interoperability

Client Version	Server Version						
	26ai	21c	19c	18c	12.2.0	12.1.0	11.2.0
26ai#11	Yes	Yes	Yes	No	No	No	No
21c	Yes	Yes	Yes	Was	Was	Yes#12	No
19c	Yes	Yes	Yes	Was	Was	Yes#12	Yes#9
18c	No	Was	Was	Was	Was	Was	Was#9
12.2.0	No	Was	Was	Was	Was	Was	Was
12.1.0	No	Yes#12	Yes#12	Was	Was	Yes#12	Yes#12
11.2.0	No	No	Yes#9	Was#9	Was	Yes#12	Yes#9

[KB141043 - Client / Server Interoperability Support Matrix](#)



```
--List current connections and their driver details  
--Join to gv$session for more details.
```

```
select * from gv$session_connect_info;
```





The same certification matrix applies to **database links** as well



Traditional Auditing is desupported in Oracle AI Database 26ai

- Migrate your policies to Unified Auditing



TRADITIONAL AUDITING

- Database generates audit records
- You can't create new polices or change existing ones

UNIFIED AUDITING

- Use syntax converter script ([MOS Note: KB120309](#))
- Do it before the upgrade





Familiarize yourself with the changes in Oracle AI Database

- Check the [Upgrade Guide](#)

- Home
- Parameters
- Users, Roles, etc
- Privileges
- Included Fixes
- Oracle Database Home
- Objects
- Fixed Objects
- Audit
- System
- Scheduler
- Release & Updates Timeline
- Report
- Custom Patch Sets



Oracle Release Analyzer Diff Utility

Find the differences between two Oracle Database releases

Oracle Release Analyzer Diff Utility (ORAdiff) allows you to compare two database releases to each other - with or without patch bundles on top.

Get Started

1. Click the menu icon on the upper left corner of the page
2. Select an object type from the left-hand navigation menu
3. Choose the Source and Target versions and patch levels
4. View the report. You may optionally choose a filter



The tool will display the differences such as "new tables", "added parameters", "changed columns", "removed privileges" and much more. It search can tell you when a parameter was added and which files changed in your Oracle Home.



On important databases,
execute a dictionary check before
upgrade

- Formerly known as *Health Check*
- [KB150043](#)

```
upg1.sid=DB19
upg1.source_home=/opt/oracle/product/19c
upg1.target_home=/opt/oracle/product/23ai
upg1.target_cdb=CDB1
upg1.run_dictionary_health=full
#To run only the critical checks
#upg1.run_dictionary_health=critical
```



Ensure dictionary and fixed objects statistics are accurate

- Save downtime by gathering in advance



```
begin
```

```
--dbms_stats.gather_dictionary_stats;
```

```
dbms_stats.gather_schema_stats('SYS');
```

```
dbms_stats.gather_schema_stats('SYSTEM');
```

```
dbms_stats.gather_fixed_objects_stats;
```

```
end;
```

```
/
```





Before upgrade

How to upgrade and convert

After upgrade

How Do You Start?



Installation

Download and install
Oracle AI Database 26ai



Container Database



AutoUpgrade



Installation of Oracle Home is simpler

- Gold images with recent Release Update
- Available for Oracle AI Database 26ai

Simplified Installation

- 1** Download software
- 2** Download patches
- 3** Unzip
- 4** Update OPatch
- 5** Install
- 6** Apply patches



Simplified Installation

1 Download software

2 ~~Download patches~~

3 Unzip

4 ~~Update OPatch~~

5 Install

6 ~~Apply patches~~



Simplified Installation

1 Download software

2 Unzip

3 Install



Fully updated
Oracle Home



How Do You Start?



Installation



Container Database

Create a new CDB in
Oracle AI Database 26ai



AutoUpgrade



Create Container Database in advance

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

Database Configuration Assistant - Create 'orcl' database - Step 9 of 15

Specify Configuration Options

ORACLE 26^{ai}
AI Database

Memory | Sizing | **Character sets** | Connection mode

The database character set determines how character data is stored in the database.

Use **Unicode (AL32UTF8)**
Setting character set to Unicode (AL32UTF8) enables you to store multiple language groups.

Use **OS character set (WE8MSWIN1252)**
Character set is based on the language setting of this operating system.

Choose from the list of character sets

Database character set: AL32UTF8 - Unicode UTF-8 Universal character set

Show recommended character sets only

National character set: AL16UTF16 - Unicode UTF-16 Universal character set

Default language: American

Default territory: United States

Navigation: Database Operation, Creation Mode, Deployment Type, Database Identification, Storage Option, Fast Recovery Option, Network Configuration, Database Options, Configuration Options, Management Options, User Credentials, Creation Option, Summary, Progress Page, Finish



Create Container Database

1 Character set

2 Components

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

3 COMPATIBLE

Select Component	Tablespace	Include in PDBs
<input type="checkbox"/> Oracle JVM	SYSTEM	<input type="checkbox"/>
<input type="checkbox"/> Oracle Text	SYSAUX	<input type="checkbox"/>
<input type="checkbox"/> Oracle OLAP	SYSAUX	<input type="checkbox"/>
<input type="checkbox"/> Oracle Spatial	SYSAUX	<input type="checkbox"/>
<input type="checkbox"/> Oracle Label Security	SYSTEM	<input type="checkbox"/>
<input type="checkbox"/> Oracle Database Vault	SYSAUX	<input type="checkbox"/>



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

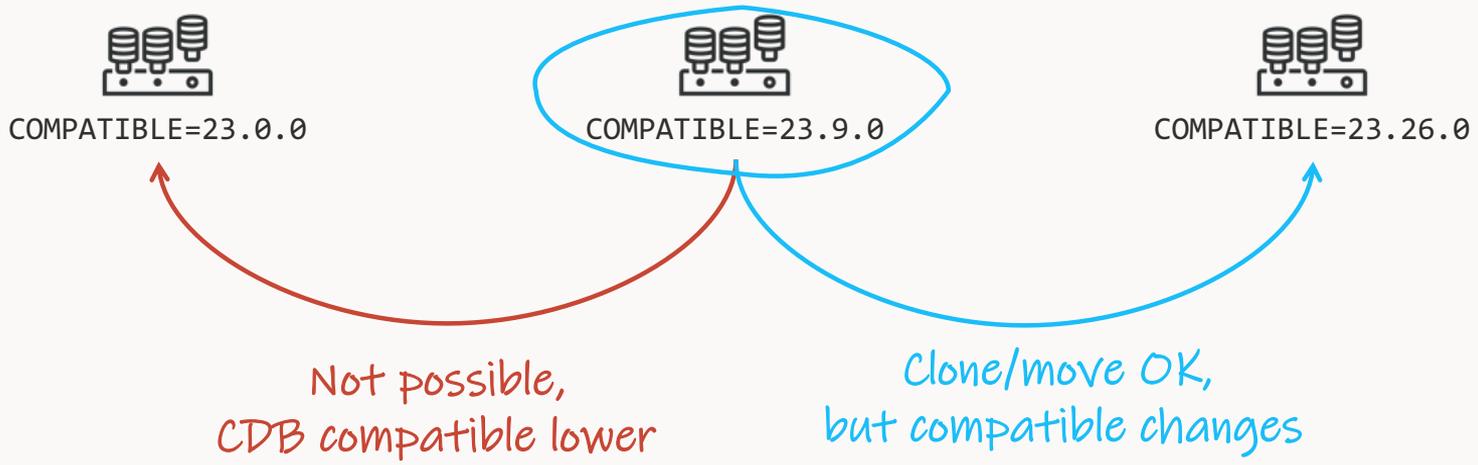
Name	Value	Include in spfile	Category
undo_tablespace	UNDOTBS1	<input checked="" type="checkbox"/>	Cluster Database
sga_target	19202	<input checked="" type="checkbox"/>	SGA Memory
db_block_size (bytes)	8192	<input checked="" type="checkbox"/>	Cache and I/O
nls_language	AMERICAN	<input checked="" type="checkbox"/>	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.oraclevcn.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 23.9.0 to 23.26.0
- Cross-release changes, e.g., 19.0.0 to 23.9.0, still requires a restart

```
alter database set downgrade compatibility to '23.9.0';
```



```
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 Do You Start?



Installation



Container Database



AutoUpgrade

Download latest version,
create your config file
and start the process

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

```
build.version 26.2.260205
```

```
build.date 2026/02/05 04:36:27 +0000
```

```
build.hash 226a8557b
```

```
build.hash_date 2026/02/04 13:52:42 +0000
```

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

```
build.type production
```

```
build.label (HEAD, tag: v26.2, origin/rdbms_19.31)
```

```
build.MOS_NOTE 2485457.1
```

```
build.MOS_LINK
```

```
https://support.oracle.com/epmos/faces/DocumentDisplay?id=2485457.1
```

1

Plug in

2

Upgrade

3

Convert



*Irreversible!
Flashback no good*



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



Upgrade via Refreshable Clone PDB



CREATE

In CDB on Oracle AI Database 26ai, create a copy of PDB over a database link



REFRESH

Apply redo from source to keep PDB up-to-date



OUTAGE

Disconnect users and refresh PDB for the last time



UPGRADE

Upgrade PDB to Oracle AI Database 26ai

Upgrade via Refreshable Clone PDB

system01.dbf
sysaux01.dbf
users01.dbf
undo01.dbf
...



*19c
database*



*21ai
CDB*

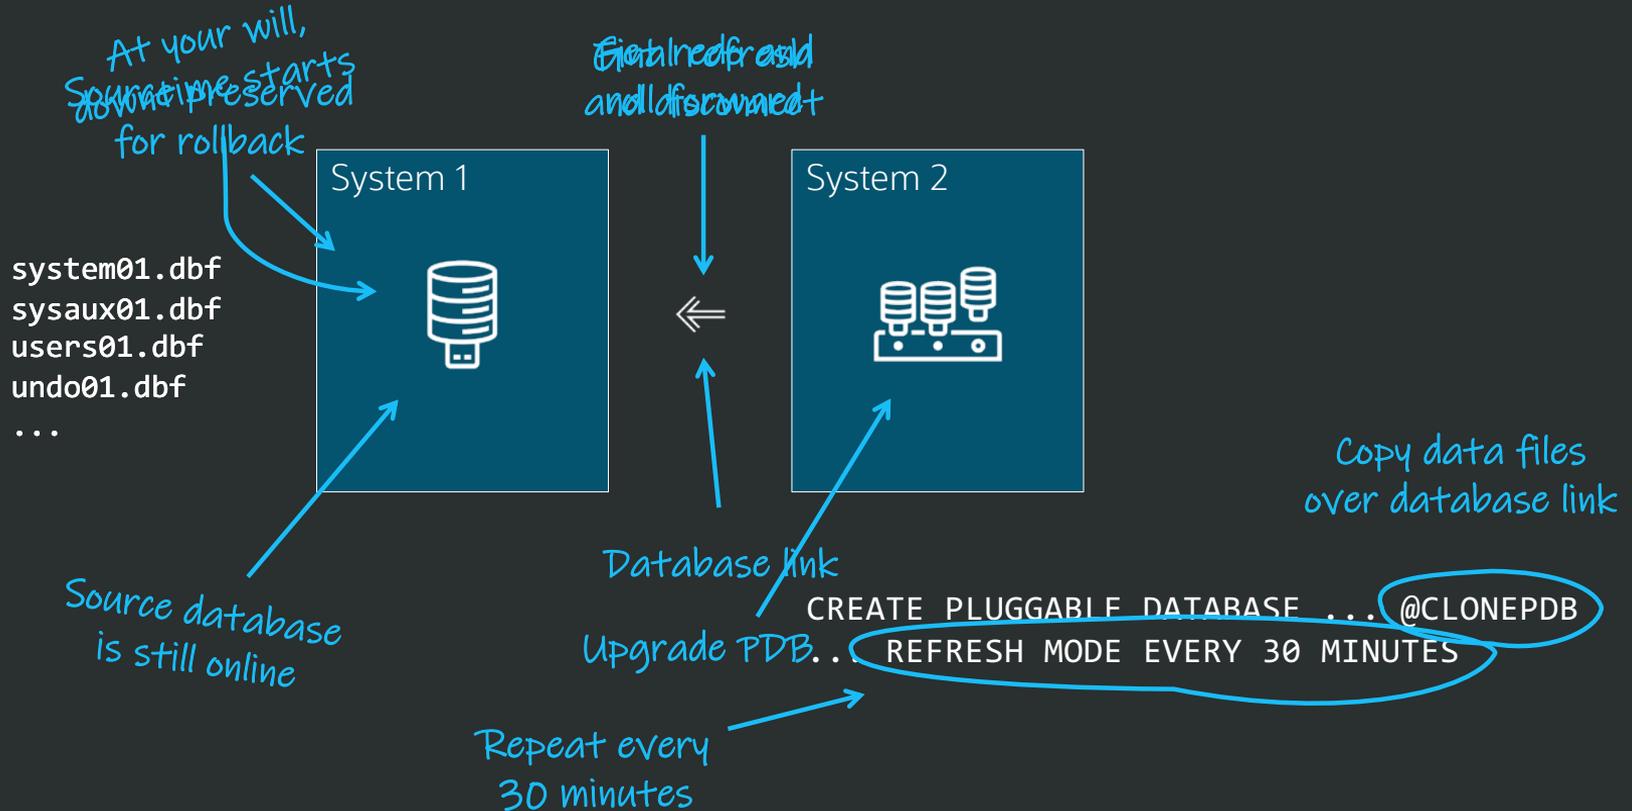
Upgrade via Refreshable Clone PDB

system01.dbf
sysaux01.dbf
users01.dbf
undo01.dbf
...



*Could be same
system as well*

Upgrade via Refreshable Clone PDB



Refreshable Clone PDB

Source non-CDB

Target CDB



```
CREATE USER dblinkuser
  IDENTIFIED BY ... ;

GRANT CREATE SESSION,
  CREATE PLUGGABLE DATABASE,
  SELECT_CATALOG_ROLE TO dblinkuser;

GRANT READ ON sys.enc$ TO dblinkuser;
```

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





You can drop user and database link
after migration



Refreshable Clone

Source non-CDB

Target CDB



```
upg1.source_home=/u01/app/oracle/product/12.2.0.1  
upg1.target_home=/u01/app/oracle/product/19  
upg1.sid=NONCDB1  
upg1.target_cdb=CDB1  
upg1.target_version=23  
upg1.source_dblink.NONCDB1=CLONEPDB  
upg1.target_pdb_name.NONCDB1=PDB1
```

Refreshable Clone PDB

Source non-CDB

Target CDB



```
upg1.source_home=/u01/app/oracle/product/12.2.0.1
upg1.target_home=/u01/app/oracle/product/19
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/12.2.0.1
upg1.target_home=/u01/app/oracle/product/19
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/12.2.0.1
upg1.target_home=/u01/app/oracle/product/19
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=21/09/2025 02:00:00
```

Refreshable Clone PDB

Source non-CDB

Target CDB



```
upg1.source_home=/u01/app/oracle/product/12.2.0.1
upg1.target_home=/u01/app/oracle/product/19
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=21/09/2025 02:00:00
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

1.

PDB
is created

2.

Data files
are copied

3.

Redo is
applied

4.

Final refresh

5.

Disconnect
and convert

`autoupgrade.jar ... -mode deploy`

`upg1.start_time=21/09/2025 02:00:00`



--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 +2h30m
```

--Or reschedule it

```
upg> proceed -job 101 -newstarttime 21/09/2025 06:30:00
```





Works for unplug-plug upgrades as well





The source non-CDB stays intact
to allow rollback



DEMO

Upgrade to Oracle AI Database 26ai





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

- Refreshable clone works only with deferred recovery on standby database

Key Benefits of Upgrade via Refreshable Clone PDB



- 1 Less downtime
- 2 Excellent rollback option
- 3 For PDBs and non-CDBs



Works everywhere

Refreshable Clone PDB

- After creating the refreshable clone PDB, don't restart the source database
- In the source database, refreshable clone PDB supports:
 - Creating new tablespaces
 - Extending existing data files
 - Adding new data files





In OCI, upgrade individual PDBs using AutoUpgrade and refreshable clone PDB

- See [blog post](#) for details



In OCI, upgrade entire CDBs using cloud tooling

- Be sure to update cloud tooling beforehand
- See [blog post](#) for details



You can also migrate with Data Pump or Transportable Tablespaces

- Suitable when direct upgrade is not possible
- Smaller databases
- Reorganizing data



Before upgrade

How to upgrade and convert

After upgrade



Fallback Options | After Go-Live



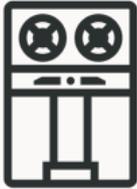
1 Back to 19c non-CDB

- Data Pump
- GoldenGate

2 Back to 19c, stay multitenant

- Downgrade
- **COMPATIBLE** must be 19.0.0 in 26ai CDB





Backup your database after migration

- Level 0
- Practice restore with pre-plugin backups





Check your standby databases

- Special attention is needed for standby databases



Data Guard



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

1 Enabled recovery

2 Deferred recovery



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

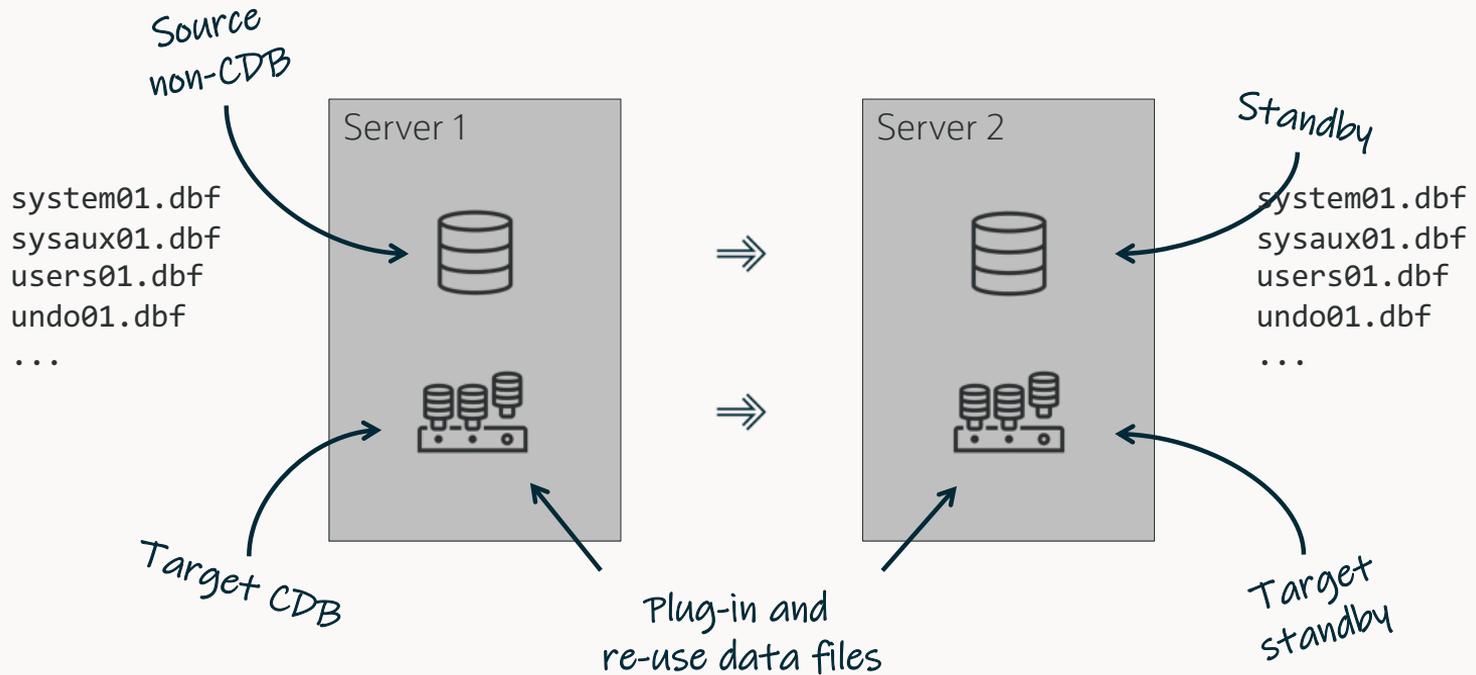
Default

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>
  <pdbname>DB12</pdbname>
  <cid>0</cid>
  <byteorder>1</byteorder>
  <vsn>203424000</vsn>
  <vsns>
    <vsnum>12.2.0.1.0</vsnum>
    <cdbcompt>12.2.0.0.0</cdbcompt>
    <pdbcompt>12.2.0.0.0</pdbcompt>
    <vsnlibnum>0.0.0.0.24</vsnlibnum>
    <vsnsql>24</vsnsql>
    <vsnbsv>8.0.0.0.0</vsnbsv>
  </vsns>
  <dbid>1852833295</dbid>
  <ncdb2pdb>1</ncdb2pdb>
  <cdbid>1852833295</cdbid>
  <guid>86D5DC2587337002E0532AB2A8C0A57C</guid>
  <uscnbas>4437941</uscnbas>
  <uscnrwp>0</uscnrwp>
  <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>
    </file>
  </tablespace>
</PDB>
```

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



```
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
```



Standby



```
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
```

Same file,
but different name

Enabled Recovery | ASM

19c
Non-CDB
Primary



19c
Non-CDB
Standby



The manifest file contains

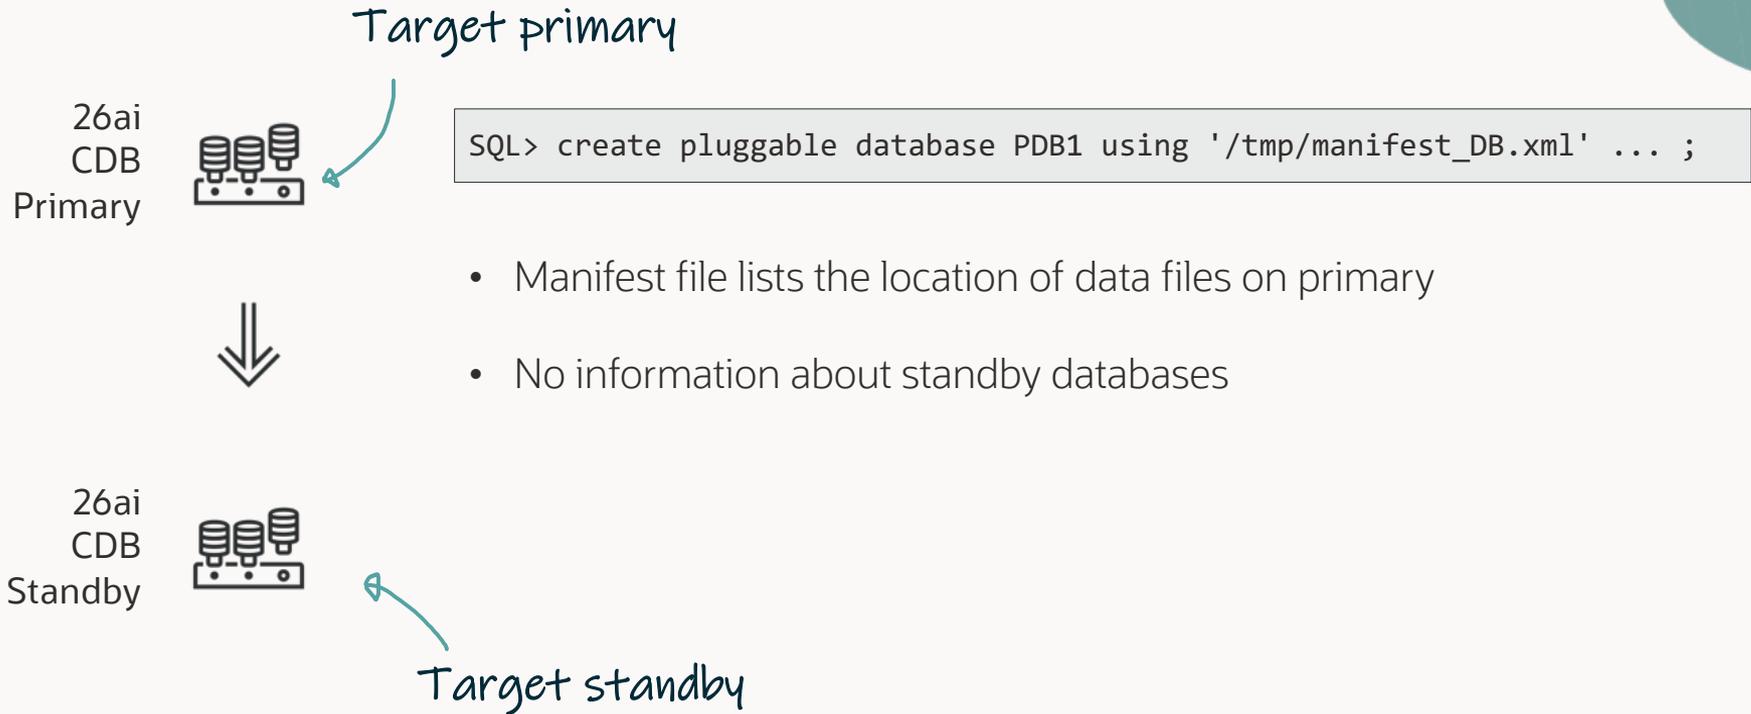
```
SQL> exec dbms_pdb.describe('/tmp/manifest_DB.xml');
```

File path on primary database only

- 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



26ai
CDB
Standby



+DATA/DB_BOSTON/DATAFILE/users.273.1103046827

Redo record says:
Plug in this data file

No good, data file
has a different name

+DATA/DB_CHICAGO/DATAFILE/users.269.1103050009



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



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



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



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



Enabled Recovery | AutoUpgrade

The current version (25.6) 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

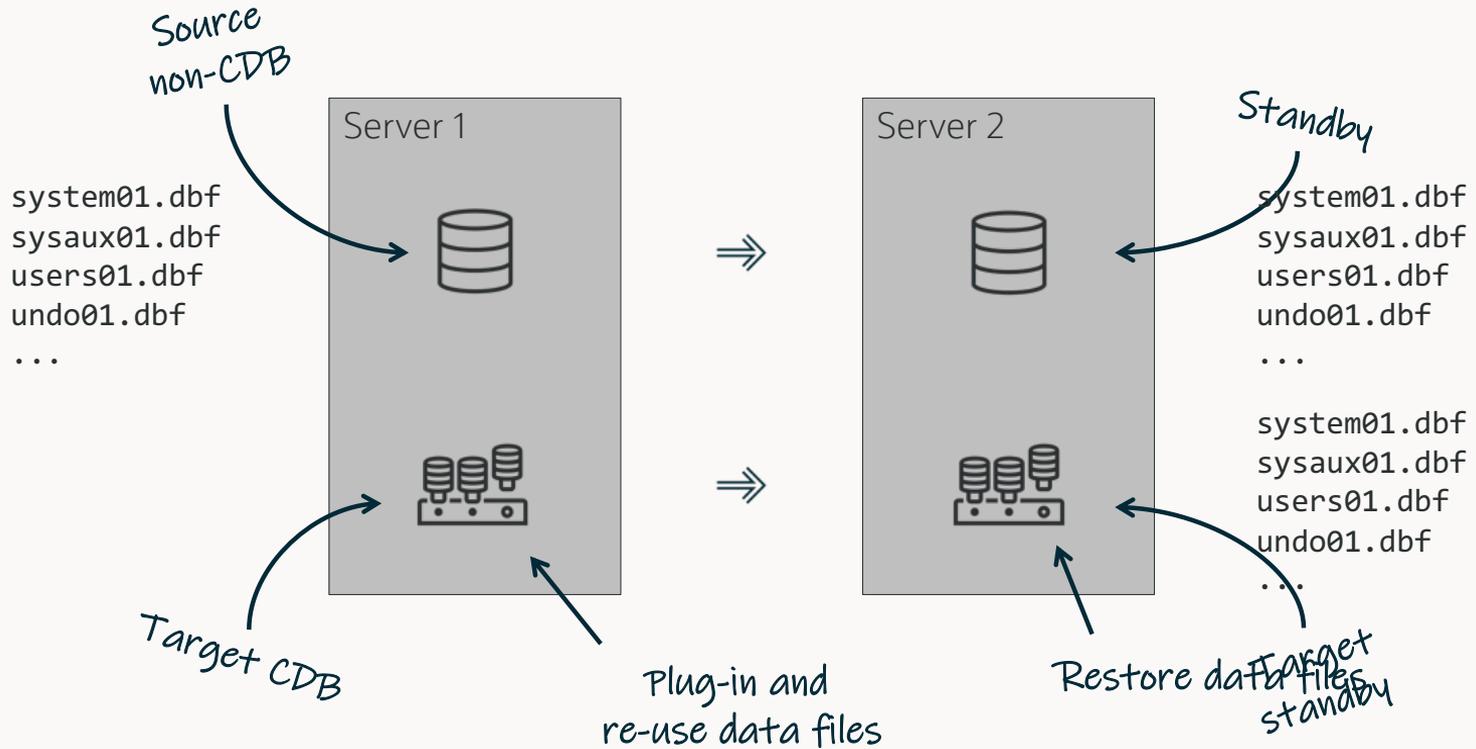
Standby ignores data files

MRP skips redo

PDB protected after restore



Deferred Recovery



Deferred Recovery



Source
Non-CDB



Target
Primary

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



Target
Standby



Deferred Recovery



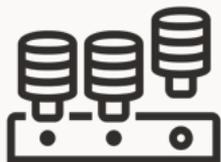
```
SQL> create pluggable database ...  
standbys=none;
```



PDB created
Data files missing



Deferred Recovery



```
SQL> show pdbs
```

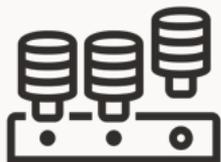
CON_NAME	OPEN MODE
PDB1	READ WRITE

```
SQL> show pdbs
```

CON_NAME	OPEN MODE
PDB1	MOUNTED



Deferred Recovery

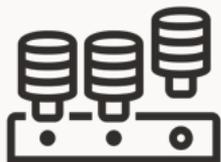


```
SQL> select name, recovery_status  
       from v$pdb;
```

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

--Default value is for CDBs with lots of PDBs
--In other cases it leads to concurrency issues
--Reset back to 12.1 default as described in MOS 2431353.1

```
alter system set "_cursor_obsolete_threshold"=1024;
```



--Default value is for CDBs with lots of PDBs
--In other cases it leads to concurrency issues
--Reset back to 12.1 default as described in MOS 2431353.1

```
alter system set "_cursor_obsolete_threshold"=1024  
comment="Added 2024-03-25 - Advice from MOS Note: 2431353.1";
```



--Database collects expression statistics for
--SQL Plan Directives and Auto-Indexing.
--If you don't use any of these, then turn it to the 11.2 value of "1".
--Tables being used in dictionary: EXP_HEAD\$, EXP_OBJ\$, EXP_STAT\$
--Doc ID 2674400.1, Value of "0" fully disables tracking.

```
alter system set "_column_tracking_level"=1;
```



Key Learnings



- 1 Start preparing today
- 2 Use AutoUpgrade
- 3 Caution with Data Guard and Multitenant conversion



Real World Customer Case

*We used AutoUpgrade and refreshable clone to move a mission-critical 76 TB database to Oracle Database 23ai in **less than 45 minutes** - while preserving the source as rollback option*

Portuguese Government Agency



**Techniker
Krankenkasse**

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



Customer Case | Techniker Krankenkasse

Customer

Large German Public Health Insurance

Project

Constraints

Preparation

Upgrade

Success?

Remarks

- **Founded:** 1884
- **Legal Form:** Public Corporation
- **Headquarters:** Hamburg, Germany
- **Members:** 11.9 million
- **Employees:** approx. 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

COMPATIBLE=23.0.0

Project

Constraints

No ASM, no Oracle Database 19c anymore

Preparation

Upgrade

Success?

Data Pump as rollback

Remarks

Customer Case | Techniker Krankenkasse

Customer

Project

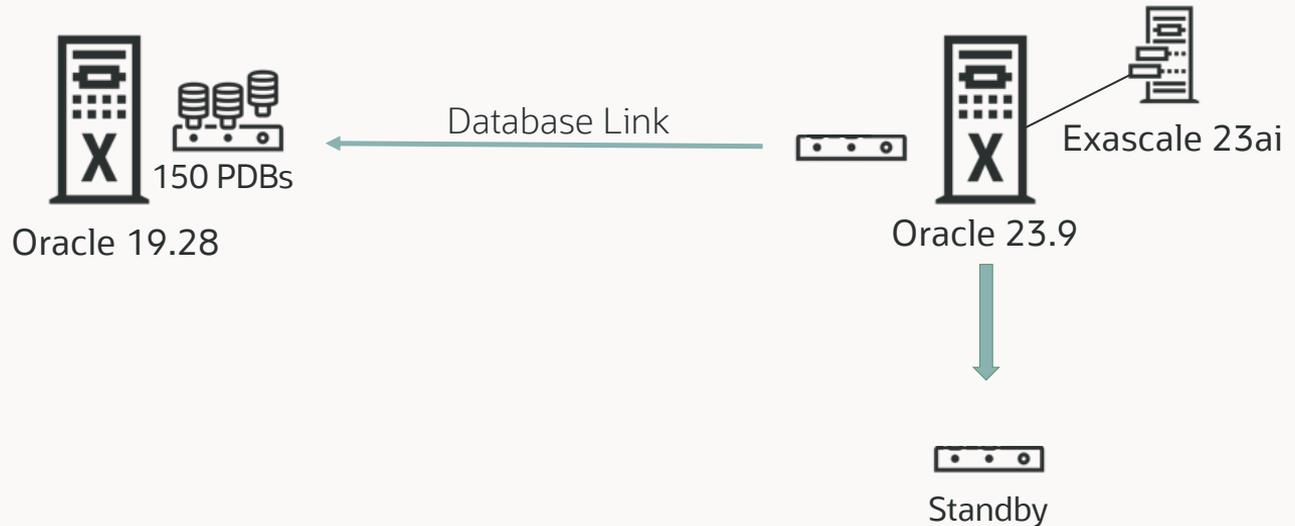
Constraints

Preparation

Upgrade

Success?

Remarks



Customer Case | Techniker Krankenkasse

Customer

Project

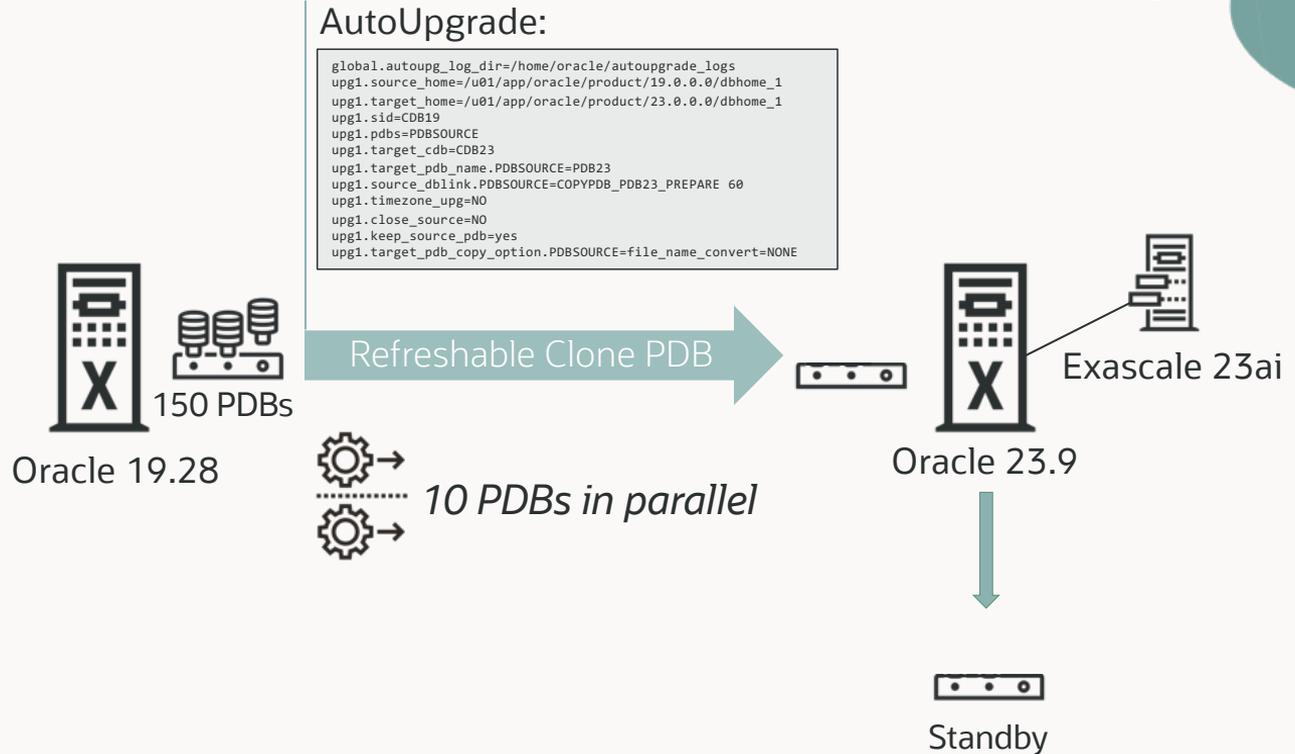
Constraints

Preparation

Upgrade

Success?

Remarks



Customer Case | Techniker Krankenkasse

Customer

Project

Constraints

Preparation

Upgrade

Success?

Remarks



Oracle 19.28



Oracle 23.9



Standby



Customer Case | Techniker Krankenkasse

Customer

Project

Constraints

Preparation

Upgrade

Success?

Remarks

Yes!!

- Migrated 150 PDBs from TEST & DEV within less than 4 weeks
- Migration of all environments completed by December 2025
- Exascale snapshots are super-fast

Customer Case | Techniker Krankenkasse

Customer

Project

Constraints

Preparation

Upgrade

Success?

Remarks

Standby building

- There were some issues
- Refreshable Clones don't propagate
- PDB recovery isolation doesn't work as expected

A number of SRs had to be opened for various areas

Optimizer works very well

How to Upgrade Refreshable PDB Clones from 19c to 23ai

Panasonic Information Systems Co., Ltd.



Overview

How we upgraded the database of **Japan's largest used car sales website**, operated by PROTO CORPORATION managed by Panasonic, on Exadata to Oracle Database 23ai

<https://www.goo-net.com/>



かんたん中古車検索

登録台数 **538,715**台

中古車

キーワードで検索



メーカーを選択する

地域を選択する

他条件を選択する

検索する

メーカーから探す

車のタイプから探す



レクサス



トヨタ



日産



ホンダ



スバル



ダイハツ



スズキ



マツダ



三菱



三菱ふそう



いすゞ



国産車
すべて



メルセデス・ベンツ



BMW



フォルクスワーゲン



アウディ



ポルシェ



MINI



ポルボ



プジョー



ランドローバー



ジープ



アルファロメオ



ジャガー



輸入車
すべて



Mike Dietrich
Oracle Vice President of Product Management and Development for Database Upgrade, Cloud Migrations and Patching writing just another blog.
Opinions certainly are my own as is this page.

Q To search type and hit enter

Archives
Select Month

Categories
Select Category

Upgrade @YouTube

*** Click on the pictures get to the download page for each slide deck ***

- [Oracle AI World 2025](#)
- [Seminars 2025](#)
- [Oracle Cloud World 2024](#)
- [Seminars 2024](#)
- [Oracle Cloud World 2023](#)
- [Seminars 2023](#)
- [Oracle Cloud World 2022](#)
- [Seminars 2022](#)
- [Seminars 2021](#)
- [Seminars 2020](#)
- [DOAG / UKOUG 2019](#)
- [OOW 2019](#)
- [AOUG 2019, OUGF 2019, POUG 2019](#)
- [OOW 2018](#)
- [Current and Comprehensive Slide Decks](#)

Oracle AI World 2025
AI World 2025



Try it out

Hitchhiker's Guide for Upgrading to Oracle AI Database

Share

Start



6 hours

Outline

- Upgrade Oracle AI Database using AutoUpgrade
- Convert to multitenant architecture using AutoUpgrade
- Convert to multitenant architecture using Refreshable Clone PDBs
- Convert to multitenant architecture using Data Pump
- Ensure performance stability
- Capture workload information
- Detect regressing statements using SQL Performance Analyzer
- Fix bad plans using SQL Tuning Advisor
- Avoid plan regressions with SQL Plan Management
- Restore failed upgrade
- Downgrade a pluggable database

Prerequisites

PLEASE NOTE: This workshop is currently only available on the LiveLabs Sandbox!

[Link: Upgrade 26ai Lab](#)



Lunch



We start again at 13:30





Data Pump Top Tips

Supercharge data loading/unloading

Always use Data Pump Bundle Patch



More than 230 functional and performance fixes

- Data Pump Recommended Proactive Patches For 19.10 and Above ([KB107134](#))



*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

Ensure dictionary and fixed objects statistics
are accurate



Ensure dictionary and fixed objects statistics are accurate

- Before export
- Before import
- Immediately after import



```
begin
```

```
    --dbms_stats.gather_dictionary_stats;
```

```
    dbms_stats.gather_schema_stats('SYS');
```

```
    dbms_stats.gather_schema_stats('SYSTEM');
```

```
    dbms_stats.gather_fixed_objects_stats;
```

```
end;
```

```
/
```



```
begin
  --dbms_stats.gather_dictionary_stats;
  dbms_stats.gather_schema_stats('SYS');
  dbms_stats.gather_schema_stats('SYSTEM');
  dbms_stats.gather_fixed_objects_stats;
end;
/
```

"After gathering dictionary stats, our Data Pump export went from 46 to 8 minutes"



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 OCPUs

Number of ECPUs / 4



On-prem (x86-64)

2 x physical cores



On-prem (other)

Depends

--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**



```
-- New in Oracle Database 21c
--
-- 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

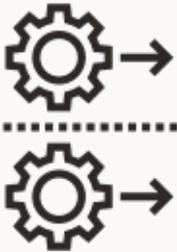
- Checksum is a weaker protection
- Requires Advanced Security Option



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

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





Transportable jobs can use parallel
from 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



Always convert to SecureFile LOBs



If exporting SecureFile LOBs is slow,
apply 19.23.0 Data Pump Bundle Patch

- Alternatively, trick Data Pump with [fake stats](#)





Do you still have BasicFile LOBs?

- Use [DIY parallelism](#) during export



--Converting a BasicFile LOB to SecureFile during import,
--is faster than not converting it.
--Overview of Oracle LOBs (Doc ID: 1490228.1)

`impdp ... transform=lob_storage:securefile`



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.



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 ( .... );
```

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

```
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

Recursive full table scan

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
```



NOVALIDATE constraints prevent the optimizer from certain **query rewrites**

- Check QUERY REWRITE INTEGRITY

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!

Exceptions

Data Pump always validates certain constraints:

1. On DEFAULT ON NULL columns
2. Used by a reference partitioned table
3. Used by a reference partitioned child table
4. Table with Primary key OID
5. Used as clustering key on a clustered table





Use with care if
you are transforming data on import



Available in Oracle Database 19c
via 19.27.0 Data Pump Bundle Patch

Even faster index imports



Use index size to determine parallel degree on index creation

- Requires 23.8 and Data Pump Bundle Patch



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 23

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 Creates 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 Creates Indexes

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



Benchmark

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 15:47:50.267: W-4 Processing object type SCHEMA_EXPORT/TABLE/INDEX/INDEX  
10-MAY-25 15:59:17.006: W-3      Completed 480 INDEX objects in 686 seconds
```



Also available in Oracle Database 19c
via 19.26.0 Data Pump Bundle Patch

Performance Stability Prescription

A person is standing in the ocean, their body partially submerged. The sky is filled with large, dramatic clouds, and the water shows some ripples. The overall mood is contemplative and somewhat somber.

what's your

biggest fear

when making changes?



A hammock is strung between two palm trees on a sandy beach. The background shows a sunset over the ocean with waves crashing on the shore. The sky is filled with orange and yellow light, and there are some clouds. The hammock is made of a dark, woven material and has some tassels hanging from it. The palm trees are silhouetted against the bright sky.

Performance Stability Prescription

Performance Stability Prescription



Collect workload information

- Sample from cursor cache
- Gather from AWR



Performance Stability Prescription



```
SQL> --Load statements from cursor cache
SQL> exec dbms_sqlset.capture_cursor_cache_sqlset( ...
SQL>
SQL> --Load statements from AWR
SQL> open ... table(dbms_sqltune.select_workload_repository) ..
SQL> dbms_sqlset.load_sqlset(...', cur);
```



SQL Tuning Set | Definition

SQL statement



Context



Statistics



Plans





Gather at least a full month of workload data

- Assist in testing your database
- Useful in solving post-upgrade performance problems



Workload Information



AWR – Automatic Workload Repository

Change the retention to a minimum of 40 days

```
exec dbms_workload_repository.modify_snapshot_settings(  
    retention=>57600,  
    interval=>30);
```



Collect SQL statements and plans

Use AWR as main source

Capture from Cursor Cache for OLTP

Collect statements, plans and stats in SQL Tuning Sets

Performance Stability Prescription



Upgrade test database

Load workload data
(SQL Tuning Set)



Performance Stability Prescription



AWR Diff Report

SQL Performance Analyzer tests your workload

Report with all regressing statements



AWR | Diff Report

Use script `awrddrpt.sql`

Top Timed Events

- Events with a "-" did not make the Top list in this set of snapshots, but are displayed for comparison purposes

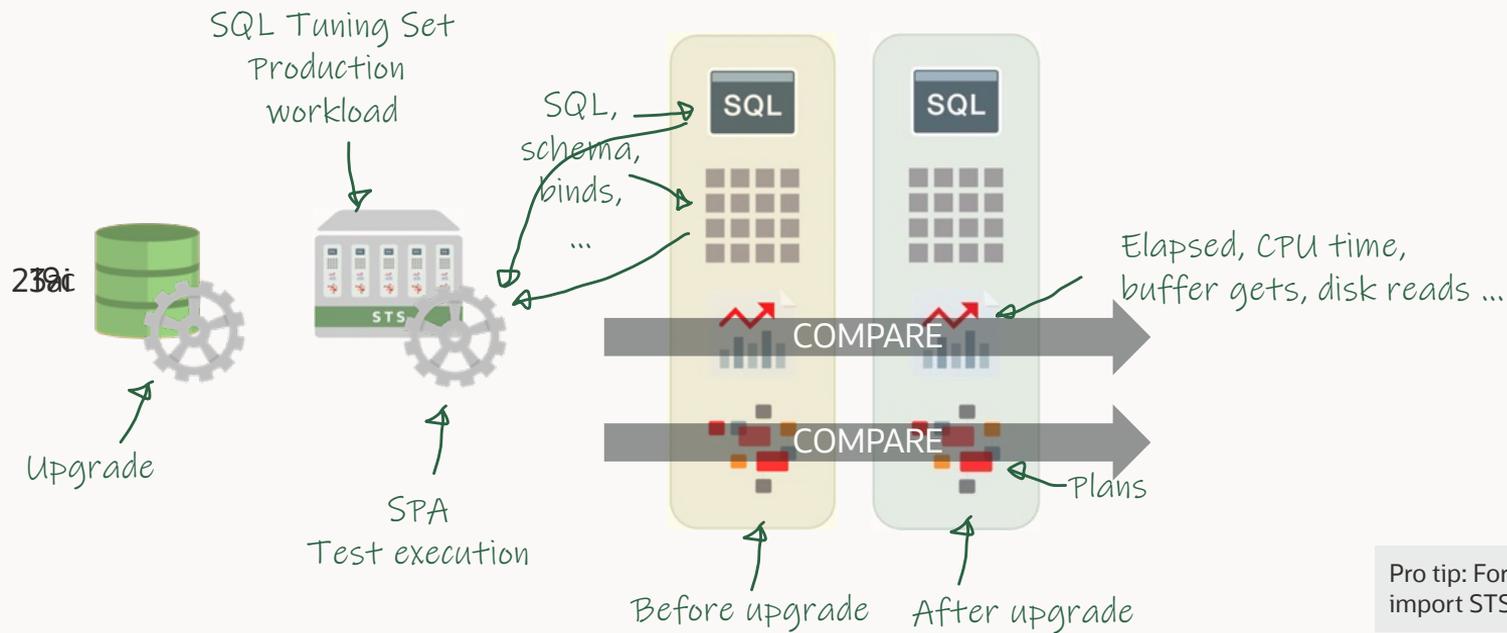
1st						2nd					
Event	Wait Class	Waits	Time(s)	Avg Time(ms)	%DB time	Event	Wait Class	Waits	Time(s)	Avg Time(ms)	%DB time
CPU time			68,289.05		43.73	db file sequential read	User I/O	22,193,998	114,919.21	5.18	23.17
db file sequential read	User I/O	6,686,953	37,737.81	5.64	24.17	enq: SS - contention	Configuration	3,913	98,997.90	25,299.74	19.96
gc buffer busy	Cluster	12,508,244	23,886.55	1.91	15.30	CPU time			73,786.55		14.88
TCP Socket (KGAS)	Network	680,629	12,514.65	18.39	8.01	row cache lock	Concurrency	73,940	48,472.30	655.56	9.77
db file scattered read	User I/O	1,572,296	4,271.68	2.72	2.74	reliable message	Other	41,148	47,600.87	1,156.82	9.60

Requires Enterprise Edition + Diagnostic pack

Pro tip: For migrations, you can [transport AWR data](#)



SQL Performance Analyzer | Concept



Pro tip: For migrations, import STS into target database



SQL Performance Analyzer | Report

Top 21 SQL Sorted by Absolute Value of Change Impact on the Workload

object_id	sql_id	Impact on Workload	Execution Frequency	Metric Before	Metric After	Impact on SQL	Plan Change
52	csv0xdm9c394t	4.02%	3262	6149.0885959534	4208	31.57%	n
41	7m5h0wf6stq0q	2.79%	21694	692.311883470084	490	29.22%	y
34	4wg725nwpbx1z	2.3%	19715	692.202079634796	509	26.47%	y
40	7jyw5gy3d1t1b	-1.43%	31816	12.0617299471964	83	-588.13%	n
36	5ps73nuy5f2vj	1.06%	31819	61.1872151858952	9	85.29%	n
44	88fgqncchy6wg	-.41%	325424	6.00316202861498	8	-33.26%	n
57	g5u7xuchhfu62	.39%	32790	26.6833180847819	8	70.02%	n



SQL Performance Analyzer | Report



Execution Plan Before Change:

Plan Hash Value : 3642382161

Id	Operation	Name	Rows	Bytes	Cost	Time
0	SELECT STATEMENT				245	
1	SORT AGGREGATE		1	24		
2	TABLE ACCESS BY INDEX ROWID	CUSTOMER	23	552	245	00:00:03
3	INDEX RANGE SCAN	CUSTOMER_I1	2888		10	00:00:01



Execution Plan After Change:

Plan Id : 138

Plan Hash Value : 1075826057

Id	Operation	Name	Rows	Bytes	Cost	Time
0	SELECT STATEMENT		1	24	245	00:00:01
1	SORT AGGREGATE		1	24		
* 2	TABLE ACCESS BY INDEX ROWID BATCHED	CUSTOMER	23	552	245	00:00:01
* 3	INDEX RANGE SCAN	CUSTOMER_I1	2888		10	00:00:01



SQL Performance Analyzer | Report

Regressed SQL Statements						
	SQL ID	Net Impact on Workload (%)	Buffer Gets		Net Impact on SQL (%)	New Plan
			SQL Trial 1	SQL Trial 2		
↓	3fv28gfu9y0aq	-0.050	26,504	29,573	-11.580	Y
↓	czzzubf8fjz96	-0.030	1,410	1,981	-40.500	Y

From production workload

From test execution



SQL Performance Analyzer | Report

Regressed SQL Statements						
	SQL ID	Net Impact on Workload (%)	Buffer Gets		Net Impact on SQL (%)	New Plan
			SQL Trial 1	SQL Trial 2		
↓	3fv28gfu9y0aq	-0.050	26,504	29,573	-11.580	Y
↓	czzzubf8fjz96	-0.030	1,410	1,981	-40.500	Y



SQL Performance Analyzer | Report

Regressed SQL Statements						
	SQL ID	Net Impact on Workload (%)	Buffer Gets		Net Impact on SQL (%)	New Plan
			SQL Trial 1	SQL Trial 2		
↓	3fy28qfu9y0aq	-0.050	26,504	29,573	-11.580	Y
↓	czzubf8fjz96	-0.030	1,410	1,981	-40.500	Y

SQL Details: czzubf8fjz96

Parsing Schema APPS

Execution Frequency 3

SQL Text 

```
SELECT /* my_query_21 */ /*+ ORDERED INDEX(t1) USE_HASH(t1) */ 'B' || t2.take_02 take_02, 'B' || t2.take_15
take_15, 'B' || t2.take_08 take_08, 'r' || t3.record_nr price_eur_id,...
```

Single Execution Statistics

	Execution Statistic Name	Net Impact on Workload (%)	Execution Statistic Collected		Net Impact on SQL (%)
			SQL Trial 1	SQL Trial 2	
↓	Elapsed Time (sec)	-0.240	0.112	0.164	-46.170
↑	Parse Time (sec)	0.220	0.001	0.001	14.490
↓	CPU Time (sec)	-0.030	0.108	0.114	-5.040
↔	User I/O Time (sec)	0.000	0.000	0.000	0.000
↓	Buffer Gets	-0.030	1,410	1,981	-40.500



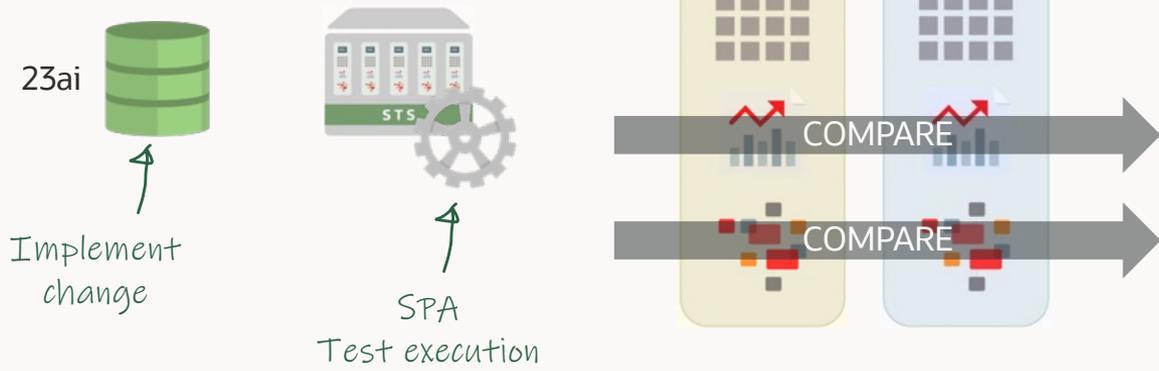
SQL Performance Analyzer | Report

Regressed SQL Statements						
	SQL ID	Net Impact on Workload (%)	Buffer Gets		Net Impact on SQL (%)	New Plan
			SQL Trial 1	SQL Trial 2		
↓	3fy28qfu9y0aq	-0.050	26,504	29,573	-11.580	Y
↓	czzzubf8fjz96	-0.030	1,410	1,981	-40.500	Y

Plan Comparison						
SQL_TRIAL_1353942463446						
Plan Hash Value 1165613724						
Expand All Collapse All						
Operation	Line ID	Object	Rows	Cost	Predicate	
▽ SELECT STATEMENT	0		1	9,830		
▽ HASH GROUP BY	1		1	9,830		
▽ MERGE JOIN	2		1	9,829		
▽ SORT JOIN	3		8	9,795		
▽ HASH JOIN	4		8	9,794	"T1"."PERIOD_CODE"="T4"."FLYER...	
INDEX RANGE SCAN	5	APPS.IDX\$\$_080F0004	1	2	"T4"."EXPORT_LIC_NR"=14659	
▽ HASH JOIN	6		14,210	9,792	"T1"."SKU_NR"="T2"."SKU_NR" AN...	



SPA | Continuous Improvement



```
--If your queries have a lot of binds, you may tweak the below
--underscore. It defines how much bind data will be kept in memory
--for each query. Default at 400 is often too low.
--Find procedures and more on:
--https://blogs.oracle.com/coretec/post/spa-in-autonomous-database
```

```
alter system set "_cursor_bind_capture_area_size"=3999;
```



Performance Stability Prescription



Tune SQLs with regressed plans

Create SQL Plan Baselines

Transport to production database



```
declare
    l_task varchar2(64);
    l_report clob;
begin
    l_task := dbms_sqltune.create_tuning_task(sql_id=> ... );

    dbms_sqltune.execute_tuning_task(l_task);

    l_report := dbms_sqltune.report_tuning_task(l_task);

    dbms_output.put_line(l_report);
end;
/
```



SQL Tuning Advisor | Example

FINDINGS SECTION (8 findings)

1- Statistics Finding

Optimizer statistics for table "SYS"."TABPART\$" and its indices are stale.

Recommendation

- Consider collecting optimizer statistics for this table.

```
execute dbms_stats.gather_table_stats(ownname => 'SYS', tabname =>
'TABPART$', estimate_percent => DBMS_STATS.AUTO_SAMPLE_SIZE,
method_opt => 'FOR ALL COLUMNS SIZE AUTO');
```

Rationale

The optimizer requires up-to-date statistics for the table in order to select a good execution plan.

SQL Tuning Advisor | Example

FINDINGS SECTION (8 findings)
1- Statistics Finding
Optimizer statistics for table "SY
Recommendation
- Consider collecting optimizer st execute dbms_stats.gather_table_ 'TABPARTS', estimate_per method_opt => 'FOR ALL C
Rationale
The optimizer requires up-to-dat select a good execution plan.

```
6- SQL Profile Finding (see explain plans section below)
-----
A potentially better execution plan was found for this statement.

Recommendation (estimated benefit: 67.2%)
-----
- Consider accepting the recommended SQL profile.
  execute dbms_sqltune.accept_sql_profile(task_name => 'TASK_21944',
    task_owner => 'SYS', replace => TRUE);

Validation results
-----
The SQL profile was tested by executing both its plan and the original plan
and measuring their respective execution statistics. A plan may have been
only partially executed if the other could be run to completion in less time.
```



SQL Tuning Advisor | Example

7- Index Finding (see explain plans section below)

The execution plan of this statement can be improved by creating one or more indices.

DBMS_SQLTUNE.REPORT_TUNING_TASK(:STMT_TASK)

Recommendation (estimated benefit: 88.23%)

- Consider running the Access Advisor to improve the physical schema design or creating the recommended index.
create index SYSTEM.IDX\$\$_55B80001 on SYSTEM.STATS("N13");

Rationale

Creating the recommended indices significantly improves the execution plan of this statement. However, it might be preferable to run "Access Advisor" using a representative SQL workload as opposed to a single statement. This will allow to get comprehensive index recommendations which takes into account index maintenance overhead and additional space consumption.

FINDINGS SECTION IS FINDINGS
6- SQL Profile Finding
6- Statistics Finding
Optimizer statistics for s
Recommendation
- Consider collecting stats
execute dbms_sqlt
'(spare)', est
method_opt = 'S
Rationale
The optimizer requires us
to select a good execution

A potentially better
Recommendation (est
- Consider acceptin
execute dbms_sqlt
task_owne
Validation results
The SQL profile was
and measuring their
only partially exec



SQL Tuning Advisor | Findings

Types of findings:

- Rewriting SQL statements
- Creation of SQL plan baselines
- Gathering object statistics
- Creation of indexes
- Creation of SQL profiles
- and more

Pro tip: SQL Developer has a good [interface](#) to SQL Tuning Advisor



SQL Plan Management | Concept



← Restorable SQL

No plans in baseline

Filter
Plan A

Plan A is used

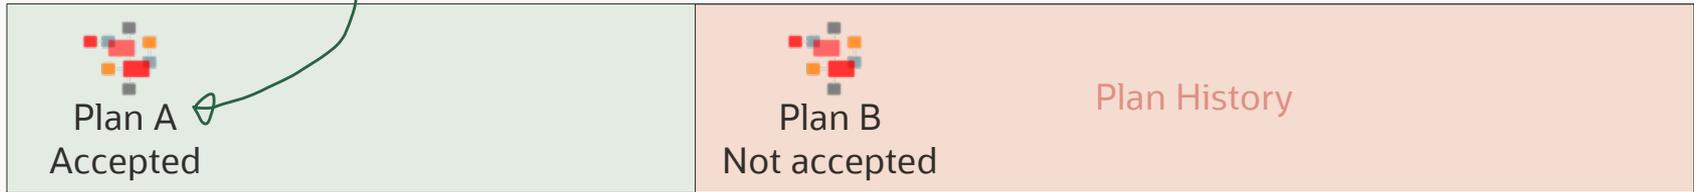
Filtering changed

- ~~SQL~~ statistics
- ~~SQL~~ text parameters
- ~~SQL~~ profile
- Action

```
OPTIMIZER_USE_SQL_PLAN_BASELINES=TRUE
```

Plan history

Dedicated part of SQL plan baseline



SPM | Evolve

 Result:
Performance better

Test execute Test execute
Plan stays



SPM | Load from STS

SQL Tuning set



Plan C



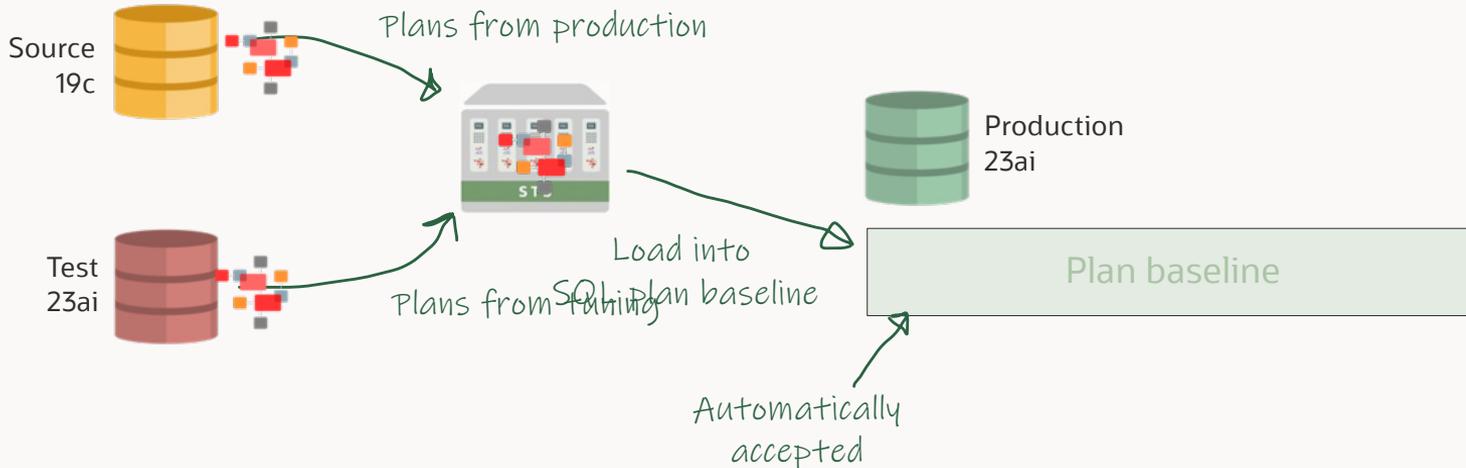
```
SQL> DECLARE
      cnt number;
BEGIN
      cnt := DBMS_SPM.LOAD_PLANS_FROM_SQLSET('UPG_STS_1');
END;
/
```



Automatically
accepted



SPM | Use Case



SPM | What If ... Literals

SQL Plan Management in a system with literals is not a good fit

- Many distinct statements
- `CURSOR_SHARING = FORCE?` [No!](#)

Optimal solution: Change your application to use bind variables





Use SQL Profiles for statements with literals

- Part of Tuning Pack

```
dbms_sqltune.accept_sql_profile(..., category=>'TEST_ENV');
```

```
alter session set sqltune_category='TEST_ENV';
```

```
dbms_sqlltune.accept_sql_profile(..., category=>'TEST_ENV');  
  
alter session set sqlltune_category='TEST_ENV';  
  
--After testing, fully enabled profile  
dbms_sqlltune.alter_sql_profile(..., attribute_name=>'CATEGORY', value=>'DEFAULT');
```

SQL Profiles | Facts

- Stores a set of hints that causes the optimizer to select a plan
- Affects one statement only
- Transparent to application
- Useful with literals using **FORCE_MATCH=TRUE**
- Persistent and transportable

There is only one tool to ensure plan stability:

SQL Plan Management

Don't use

- OPTIMIZER_FEATURES_ENABLE
- COMPATIBLE

COMPATIBLE vs. OPTIMIZER_FEATURES_ENABLE

COMPATIBLE

- Enables features
- Always use the default value of a release (e.g. 19.0.0)

OPTIMIZER_FEATURES_ENABLE

- Just reverts to the parameters used in a previous release
- Avoid using it if possible
- This is not a Swiss Army knife!
- You will turn off a lot of great features





Hitchhiker's Guide to Upgrade to Oracle AI Database 26ai

try it out for free

IT'S EASY



If you don't have the right license, use a database system in OCI

- Relevant options and packs included in most cloud offerings

Use OCI for Performance Testing



Export SQL Tuning Set



Export SQL Plan Baselines

Generate SPA report
Create SQL Plan Baselines



Performance Tips & Tricks



Use as few initialization parameters as possible

- Stick to the defaults
- Stick to vendor recommendations



Only use underscores and events to solve specific situations

- Only under guidance of Oracle Support

Patches For Optimal Performance

- 1 Install the latest Release Update
- 2 Install the latest Monthly Recommended Patches
- 3 Check for important recommended one-off patches (Doc ID [555.1](#))
- 4 Check for other SQL performance bug fixes (Doc ID [2773715.1](#))



Enable Optimizer Fixes

5 Selectively enable optimizer fixes using `DBMS_OPTIM_BUNDLE`

```
begin
  dbms_optim_bundle.enable_optim_fixes(
    action          => 'ON',
    scope           => 'BOTH',
    current_setting_precedence => 'YES');
end;
/
```

Find available bug fixes in ORAdiff or `dbms_optim_bundle.GetBugsForBundle`



Should You Enable Optimizer Fixes?

oradiff.oracle.com

The screenshot shows the Oracle Release Analyzer Diff Utility interface. The top navigation bar includes 'Oracle Release Analyzer Diff Utility', 'Administrator', 'MOS Account', 'Oracle Account', and a user profile 'mike.dietrich@oracle.com'. A left sidebar contains navigation options: Home, Parameters, Users, Roles, etc., Privileges, and Included Fixes (highlighted in green). Under 'Included Fixes', there are links for List of Fixes, Total Fixes, Search Fixes, List of CVEs, CPU Explore, and Fix Control.

The main content area is titled 'Version Selector' and contains four dropdown menus: Source Release (19c), Source Patch Level (19.3.0 (Base Release)), Target Release (19c), and Target Patch Level (19.24.0). A 'List Mode' toggle is visible on the right.

Below the version selector is a section titled 'V_\$SYSTEM_FIX_CONTROL'. It has tabs for 'Added', 'Removed', and 'Modified'. A search bar with a 'Go' button and an 'Actions' dropdown is present. Below the search bar is a table with the following columns: BUGNO, VALUE, SQL_FEATURE, DESCRIPTION, OPTIMIZER_FEATURE_ENABLE, EVENT, and IS FAULT.

BUGNO	VALUE	SQL_FEATURE	DESCRIPTION	OPTIMIZER_FEATURE_ENABLE	EVENT	IS FAULT
33636280	0	QKSFM_EXTTAB_33636280	fix for serializable transaction		0	1
33649782	0	QKSFM_OR_EXPAND_33649782	Try LORE after ORE rejected due to LOB structure		0	1
33667505	0	QKSFM_OR_EXPAND_33667505	enhance index check in presence of NLSORT function based index		0	1

At the bottom right of the table, there is a pagination control showing '1 - 50 of 235' and a right arrow button. A red arrow points to this pagination control.



Should You Enable Optimizer Fixes?

Upgrade
New database

Enable optimizer fixes using `DBMS_OPTIM_BUNDLE`

Patching

Do proper testing before enabling
optimizer fixes using `DBMS_OPTIM_BUNDLE`





Don't gather new optimizer statistics after upgrade

- Upgrades from 11.2.0.4 might be different



Don't gather system statistics

- In most cases, the defaults are fine
- Might be a good idea on Exadata Database Machine
- [Optimizer blog](#)

Autonomous AI Database

A migration approach



The idea?

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

What's the story?

Flashback to October 2017

KEYNOTE PRESENTATION

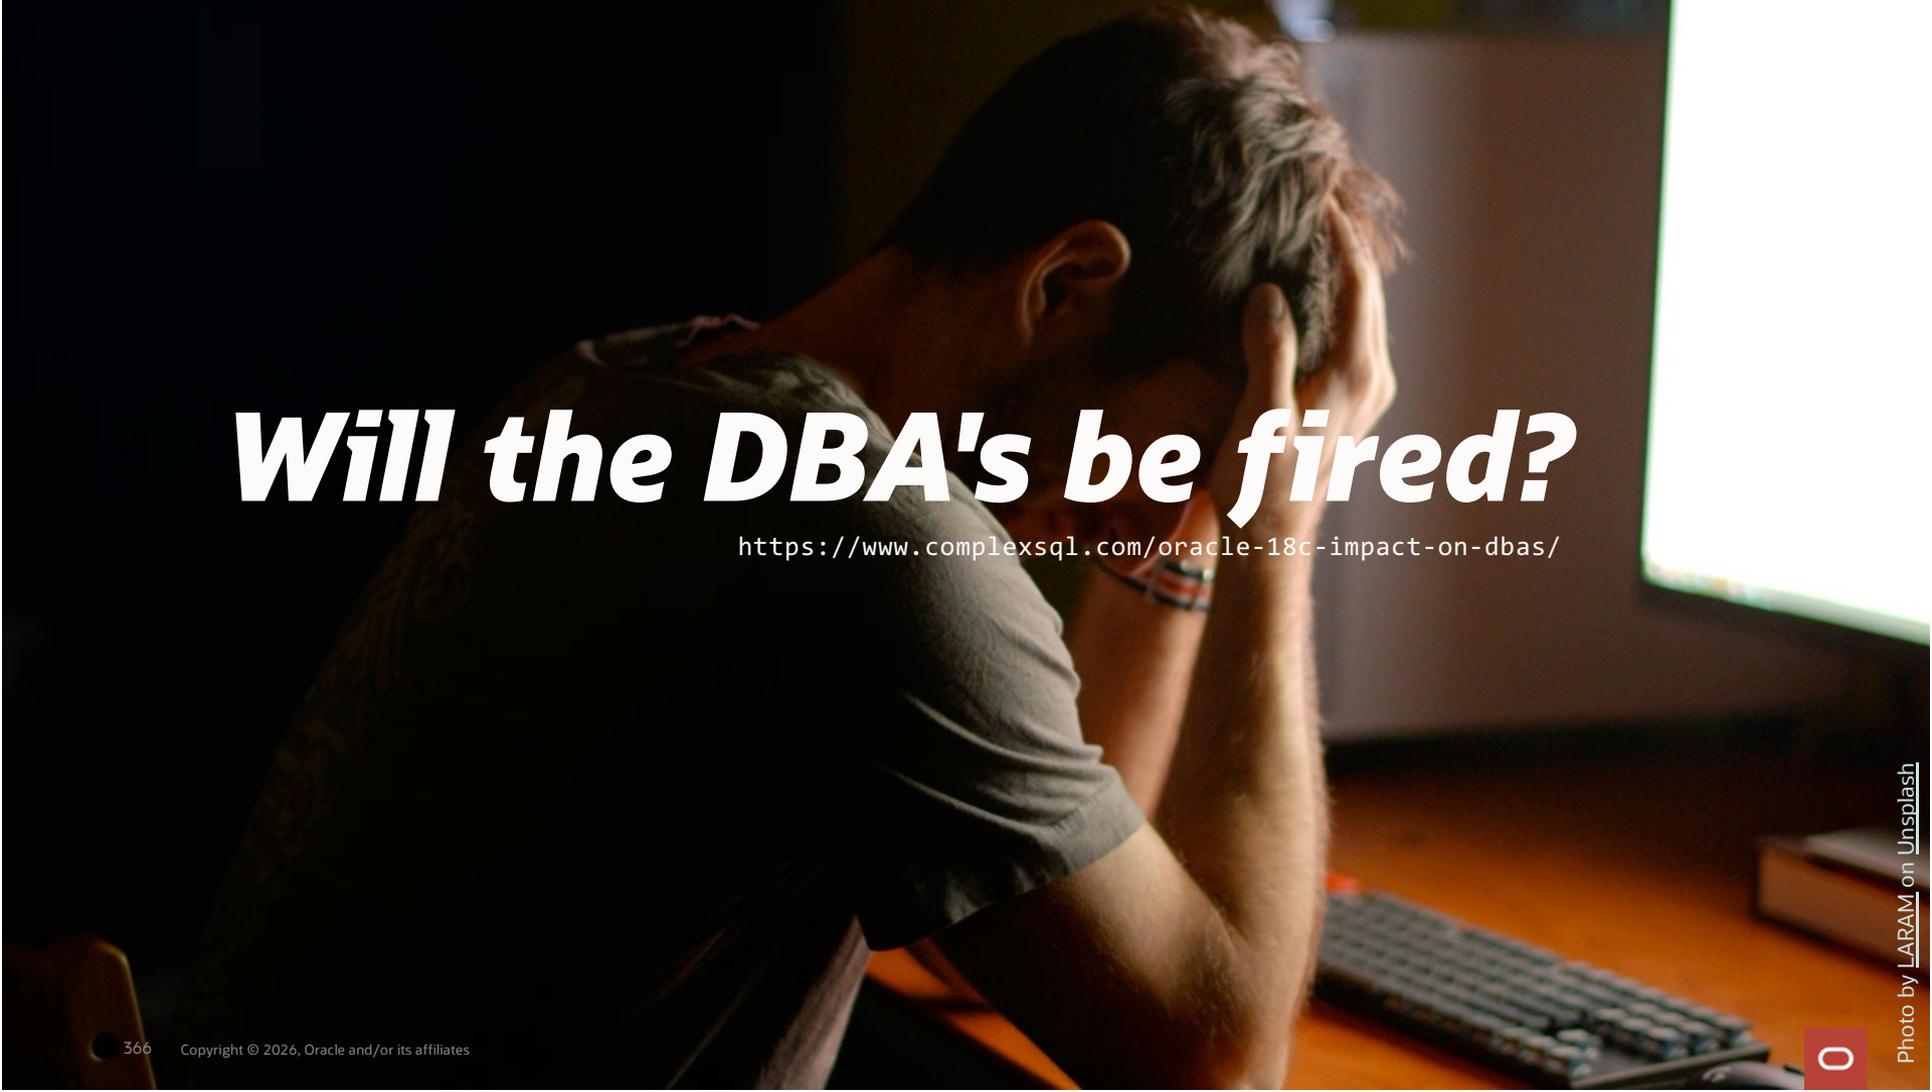
Oracle OpenWorld San Francisco 2017



by **LESDN**

OOOW 2017

ORACLE®



Will the DBA's be fired?

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



"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

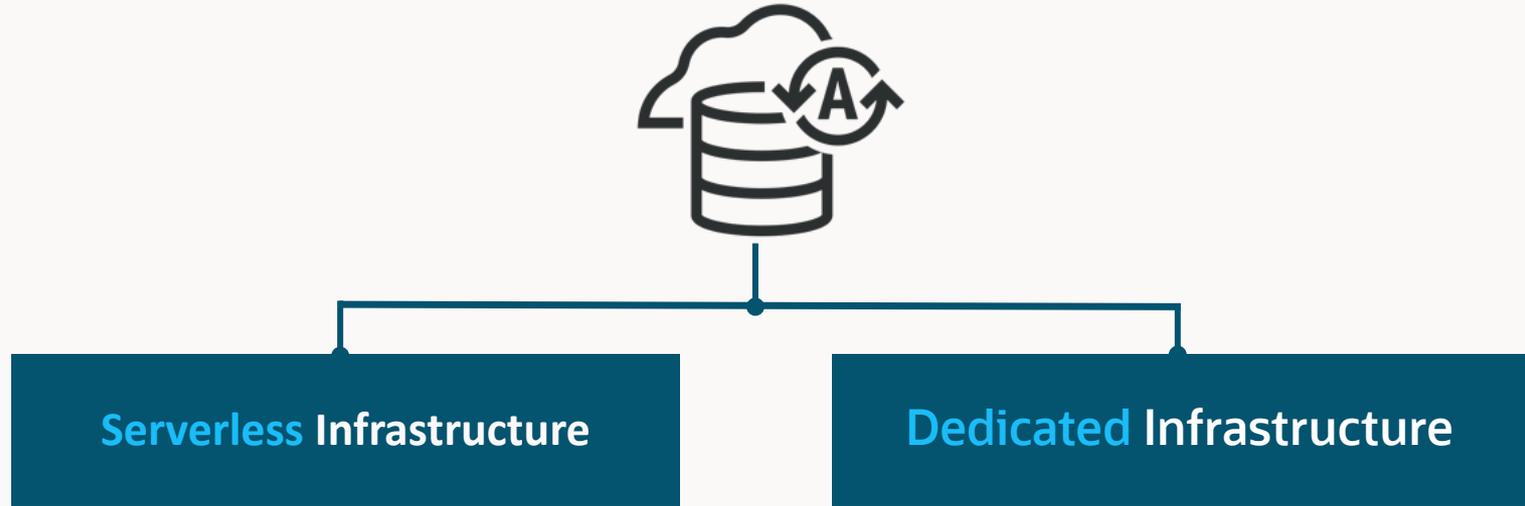


Cloud@Customer

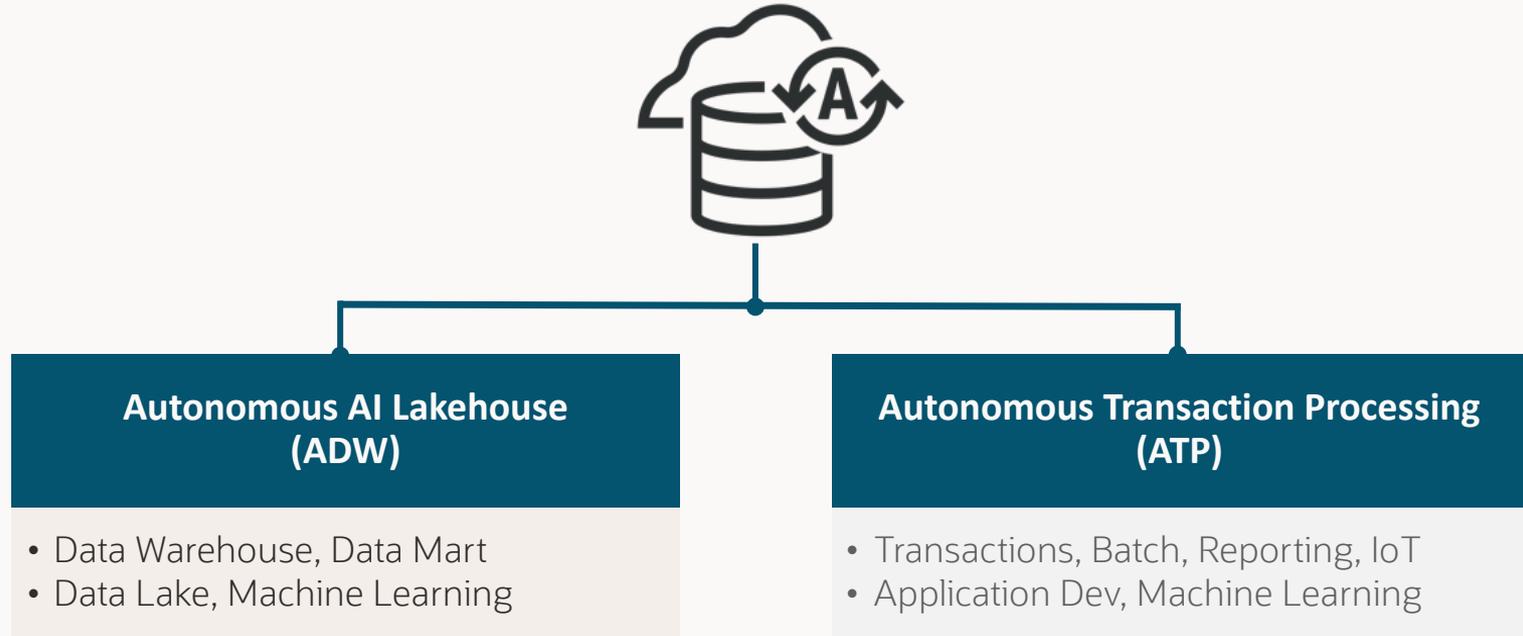
Autonomous AI Database in a VM environment



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

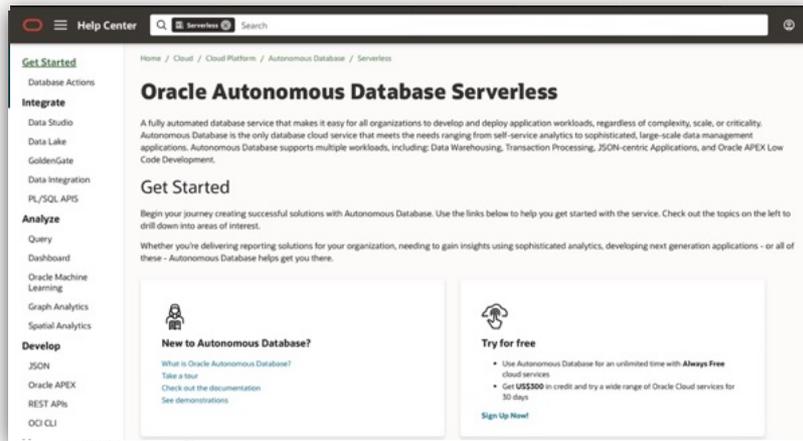


APEX

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



Essentials

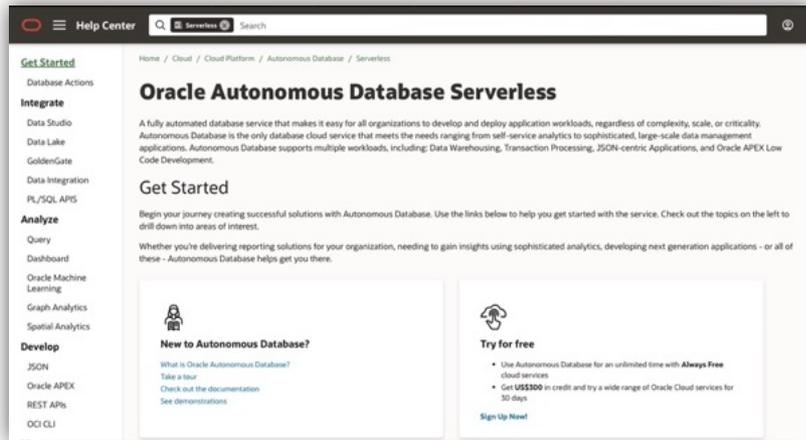


Outside the database,
check [OCI Documentation](#)

Example: Deploy, start, stop, scale



Essentials



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



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

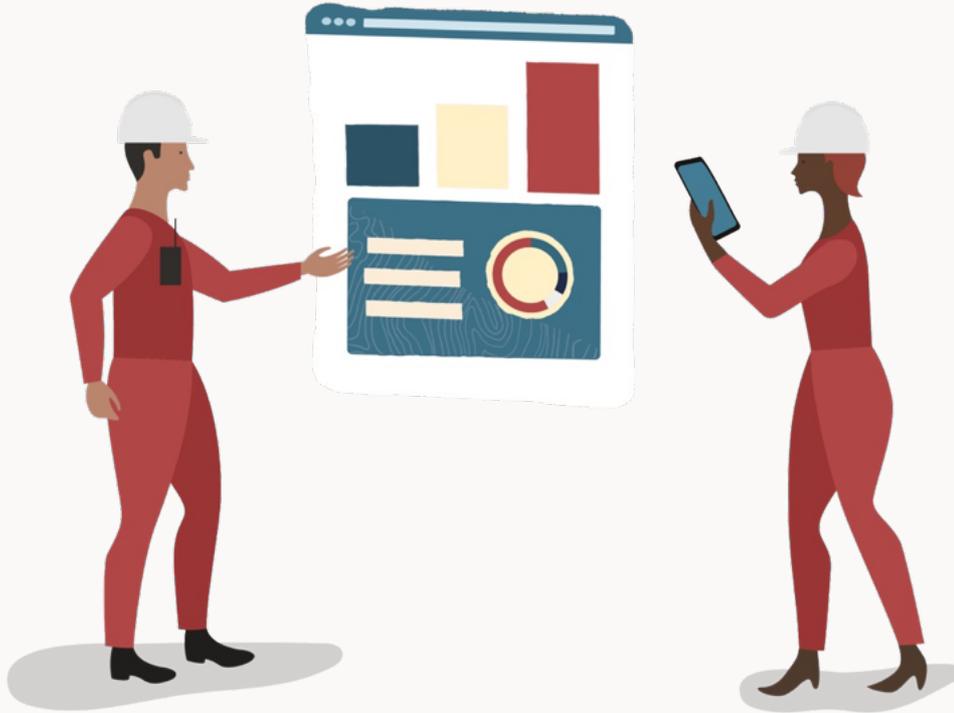
Cloud Migration
Advisor



Getting an Overview



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



Cloud Premigration Advisor Tool





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

- Use Cloud Premigration Advisor Tool (CPAT)



Overview



Connects



Checks



Reports



Fixes
(optional)



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

Release Oracle 11.2.0.4.8

Platform Generic Platform

Size 8.6 MB

Download Access Software

Classification General

Patch Tag

Release Oracle 11.2.0.4.8

Platform Generic Platform

Language American English

Bugs Resolved by This Patch

List of bugs fixed is not available. Consult the Readme.

View Related Knowledge to this Patch

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

Reset

Search

About

Valid

Input Form

★ Recommended Default

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

Output



★ Set Recommended

Use Placeholder

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

List of schemas analyzed:

(DBUSER_00_UR_MYAPP)

- [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

[Return to Table of Contents](#)

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:@dbssystemaz: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](#)

[Return to Table of Contents](#)

Report Analysis Notes

[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:	dbssystemaz
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

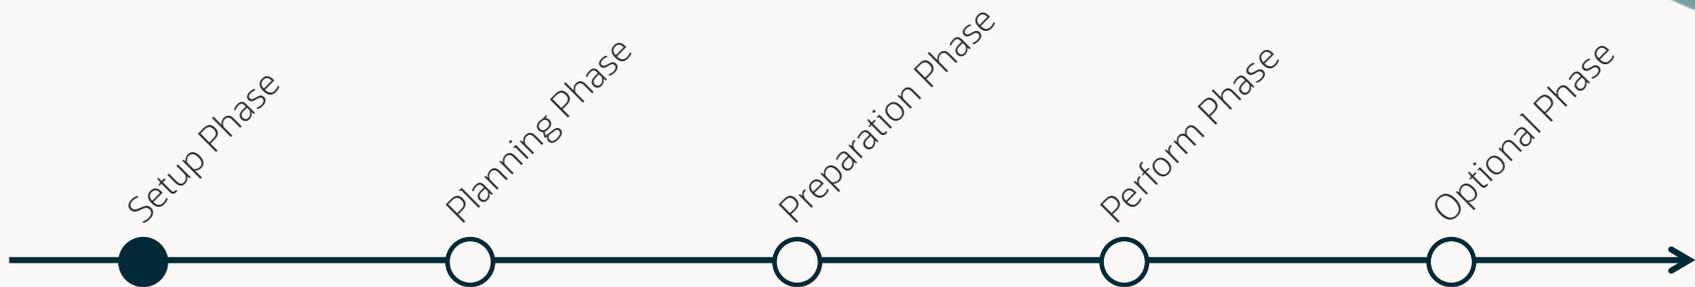
*You need a script book
You need to cast actors
You need a film set
You need ...*

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:

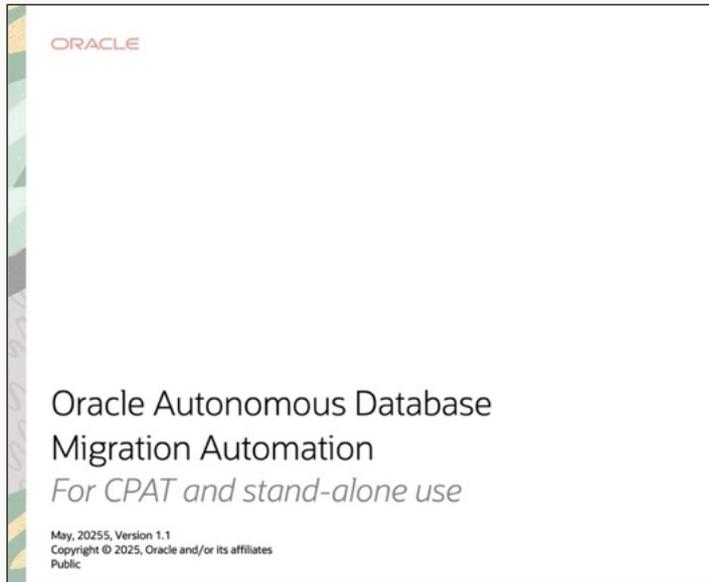
- Connect strings source and target
- Data Pump encryption
- Storage (FSS or Object Store)
- Format: TAB or SCRIPT



Documentation

AMA Documentation is available at request

- Documents the entire flow and all options and parameters

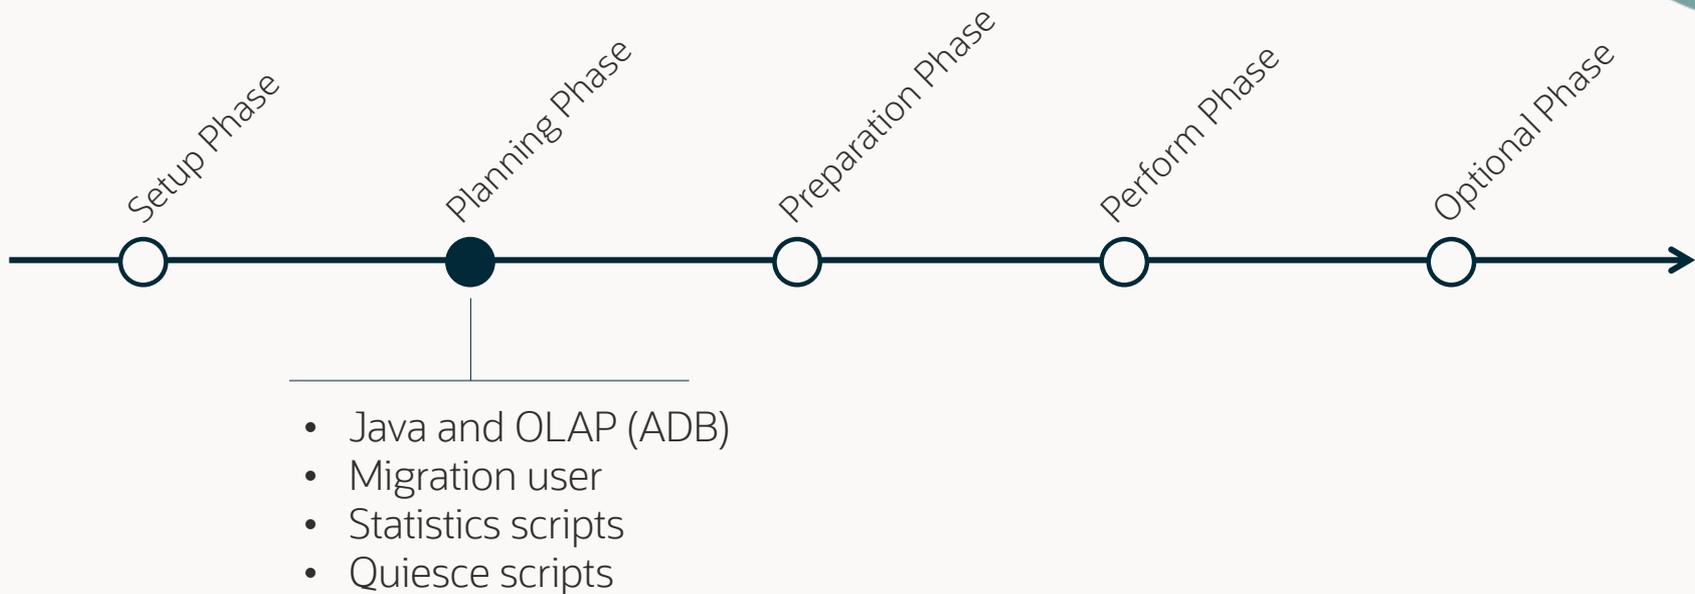


The image shows the table of contents for the Oracle Autonomous Database Migration Automation documentation. The Oracle logo is at the top left. The title "Table of contents" is prominently displayed in the center. The table of contents lists the following sections and their corresponding page numbers:

Introduction	3
Current Restrictions	3
Workflow	4
Setup	4
Planning Phase	4
Prepare Phase	4
Perform Phase	5
Post Phase	5
Optional Phase	6
How to use AMA	7
Setup and execute AMA	7
How to setup the shared storage	10
Setting up an NFS Share for the migration	10
OCI Console	10
Associate Mount Target	13
Linux	13
Windows	13
Setting up an Object Storage Bucket for the migration	13
Pre-Authenticated URL	14
APPENDIX A - AMA Migration Parameters	16
CONNECT_SRC	16
CONNECT_TGT	16
UID	16
PWD	16
EXCLUDE_USER	16
DP_ENCRYPTION_PWD	16
USE_FSS_CURL	16
DUMP_OUTPUT_PATH	17
ADB Dump File Storage Related Parameters	17
CONTROL_FILE_FORMAT	17
USE_DP_APR TRUE	17
CPAT_OUTPUT_DIR	18
APPENDIX B - The AMA Configuration File	19
APPENDIX C - AMA Walkthrough including Output (Linux)	21



AMA Workflow



AMA | Planning Phase

On-Prem - Source

Gather stats for SYS / SYSTEM

Create Migration user

Enable restricted session

Set JOB_QUEUE_PROCESSES=0

ADB-S - Target

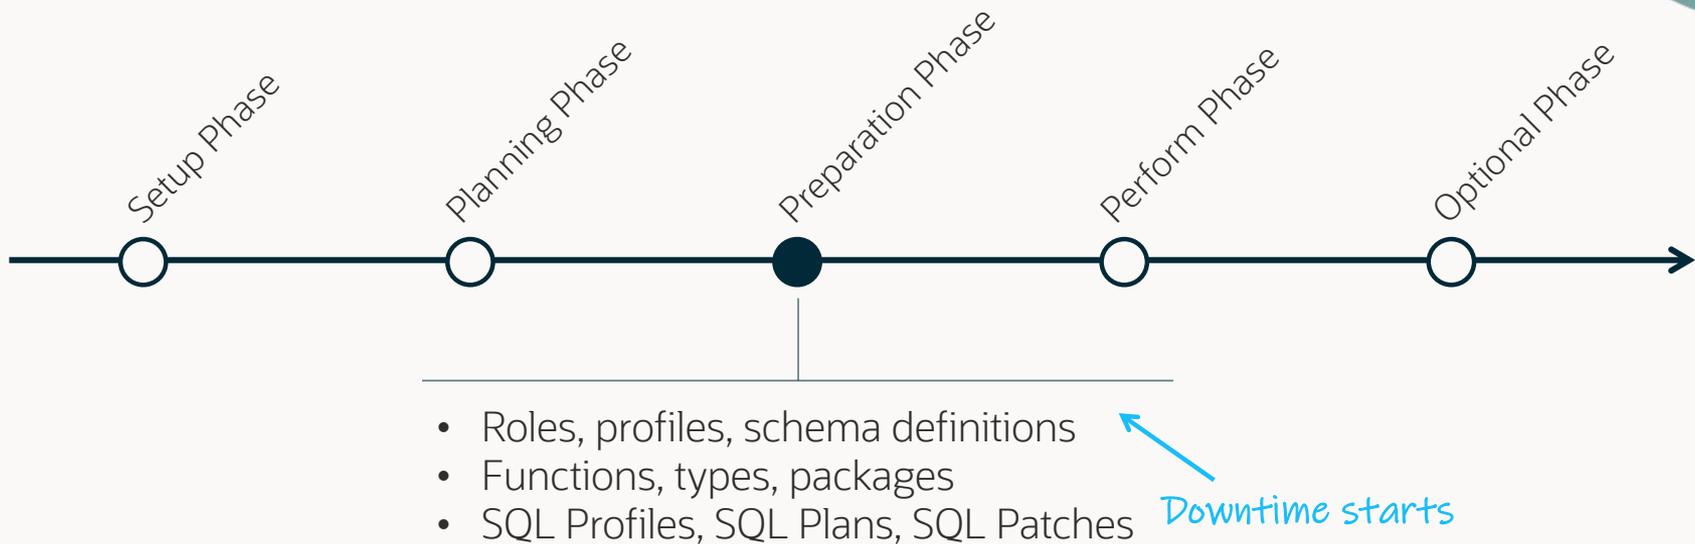
Enable OLAP / JAVA in ADB-S

```
[oracle@ephx31vm1-jlosd1 OUTPUT]$ cat __US3BLDW_MIGRATION_CONTROL_FILE.ct1
```

```
---  
---  
--- 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

Collect allowed ROLES

Collect PROFILES

Export schema definition

Export FUNCTIONS, TYPES, PACKAGES

Collect SQL Profiles, SQL Plans, SQL Patches

ADB-S - Target

Create ROLES

Create PROFILES

Create storage credential (NFS, Object Store)

Import schema definition

Import FUNCTIONS, TYPES, PACKAGES

Granting migration privileges

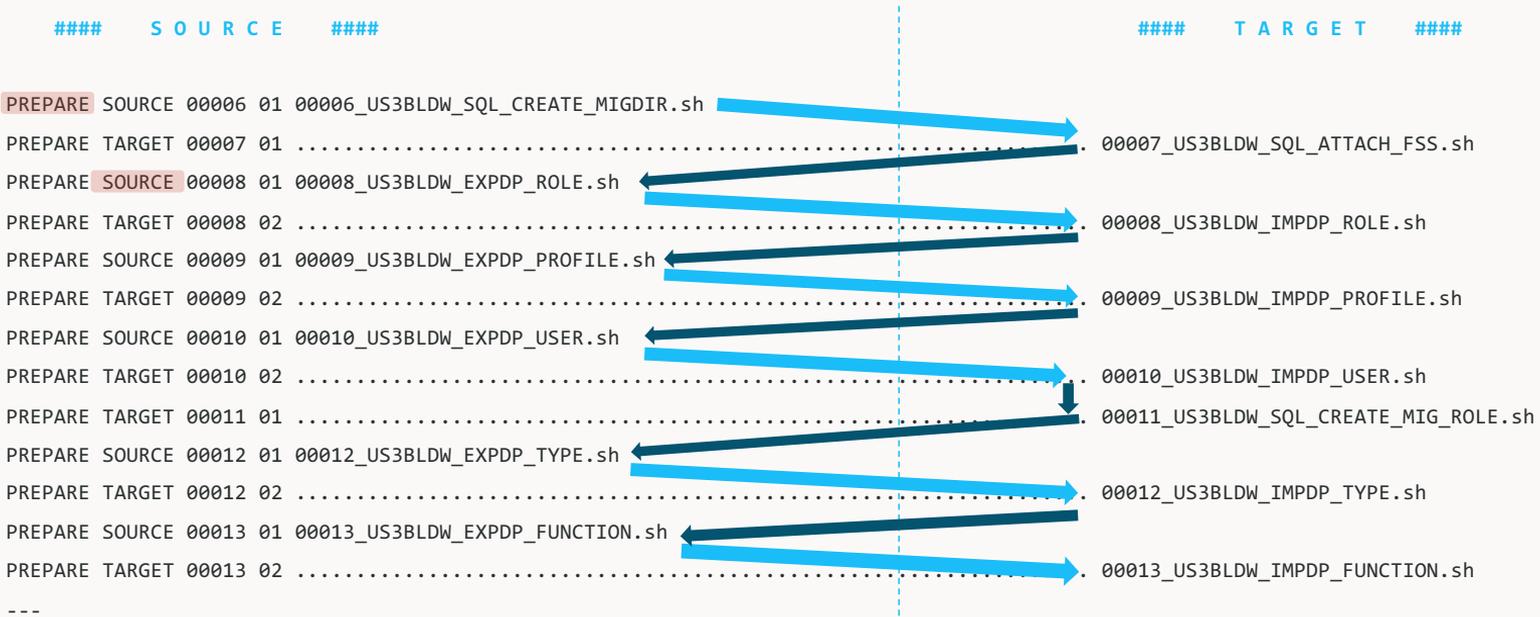
Alter user profiles

Create SQL Profiles, SQL Plans, SQL Patches

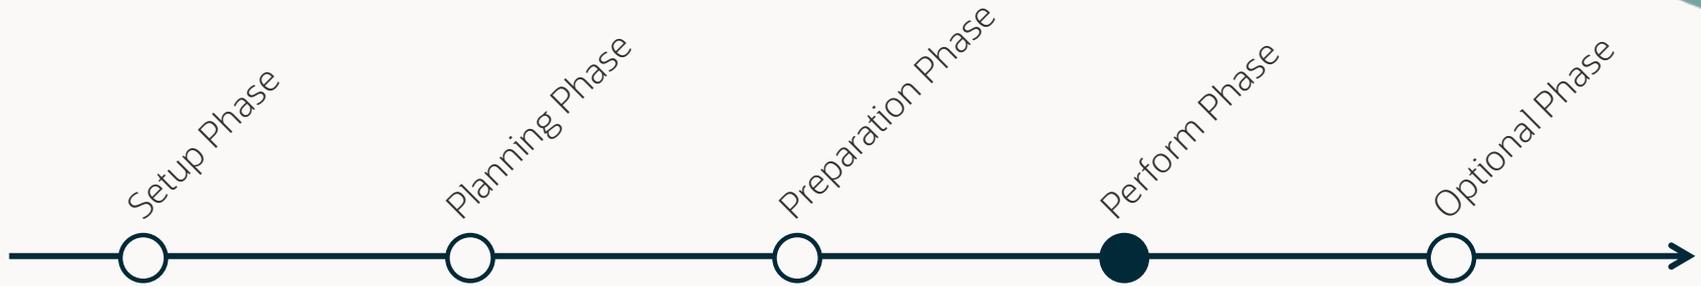
--- PREPARATION PHASE ---

--- 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 ---



AMA Workflow



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



AMA | Perform Phase

On-Prem - Source

Export all schemas

Export audit trail



Copy files (if necessary)

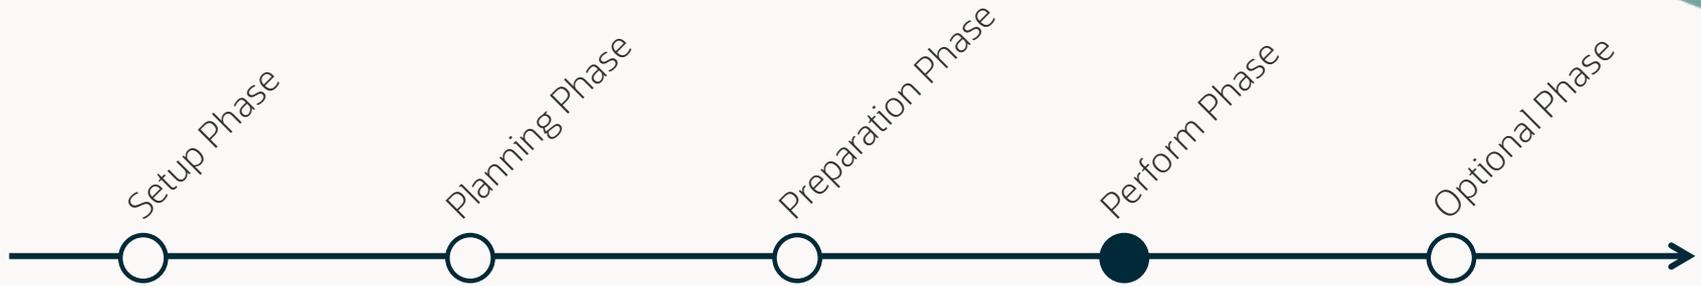


ADB-S - Target

Import all schemas

Import audit trail

AMA Workflow



- Cross-schema objects
- Privileges
- Profile adjustments
- Advanced queues
- Recompilation



--- 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 ---

####	S O U R C E	####	####	T A R G E T	####
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

ADB-S - Target

FOREIGN KEYS cross-schemas

INDEXES cross-schemas

FUNCTIONAL INDEXES enableing

REVOKE transition privileges

GRANT privs SYS, SYSTEM, CTXSYS, objects

Restore final profiles

Set tablespace quotas

Export network ACLs

Import network ACLS

Enable Advanced Queues

Recompilation



--- POST PHASE ---

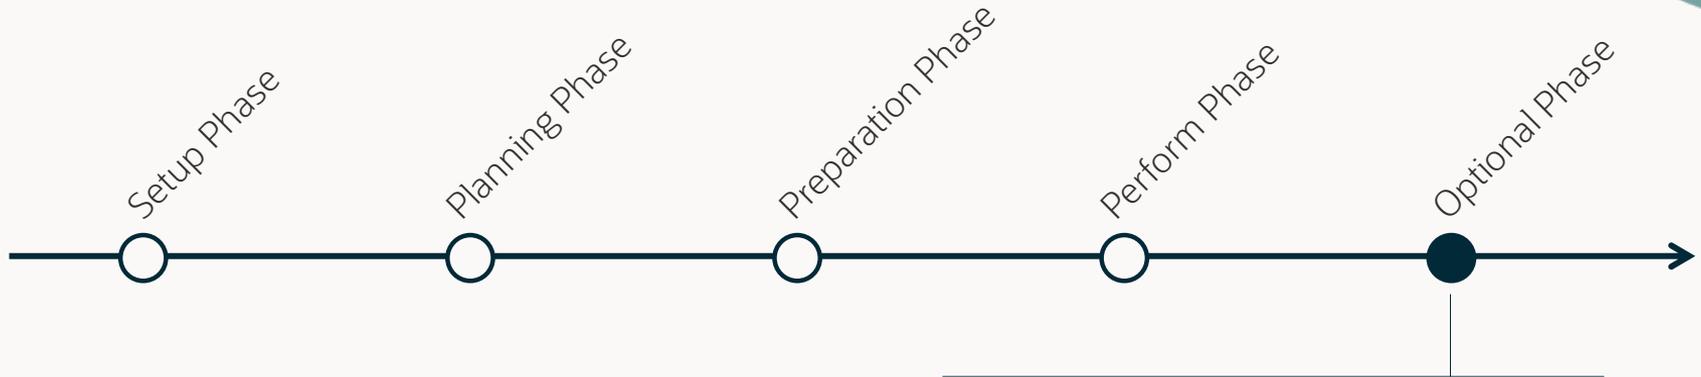
--- Here execute again all scripts one after the other as they might have dependencies again ---

####	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

Wach [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



Oracle AI Database 26ai

What's Changing

Oracle7 Server™

Release 7.3 for Windows NT Versions 3.51 and 4.0

The Enterprise
Database Engine
of Choice

Version
7.3.4.0.0



ORACLE®
Enabling the Information Age™





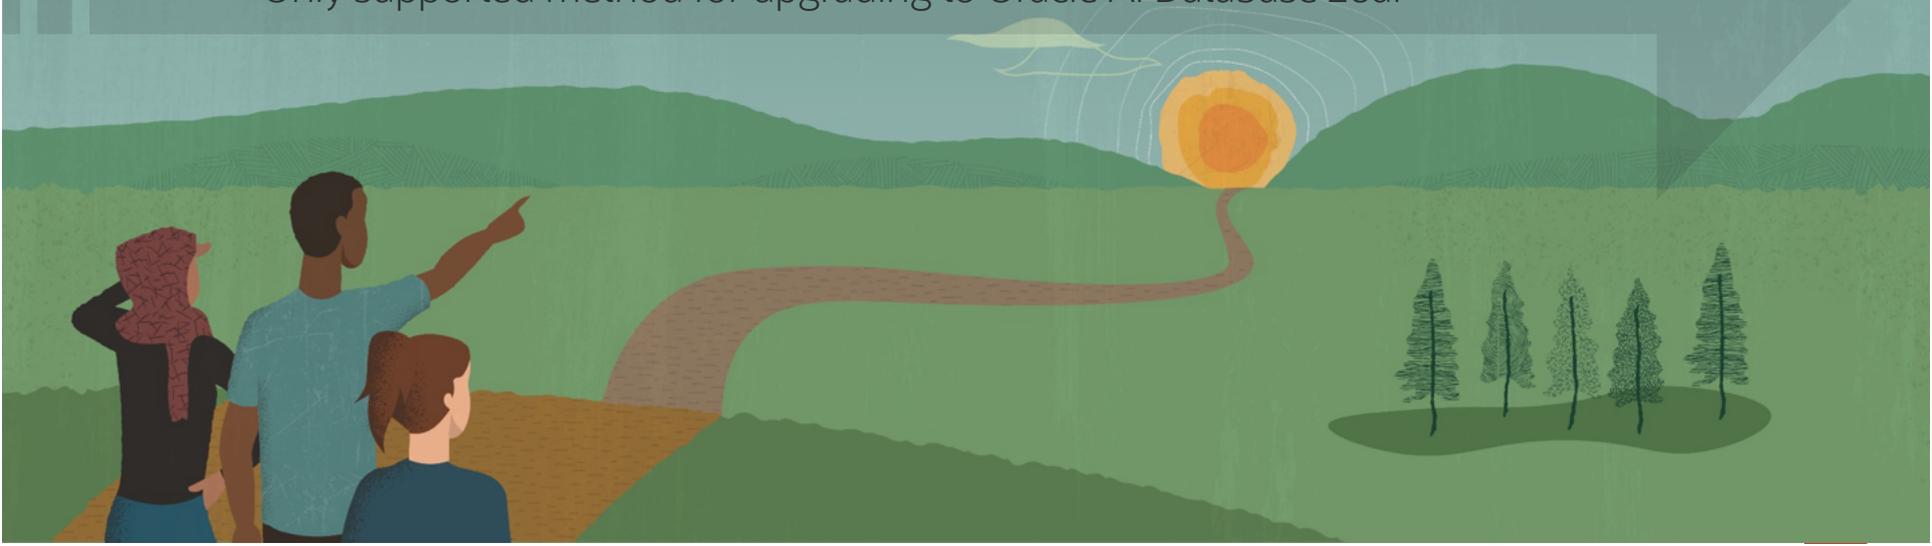
Familiarize yourself with the changes in Oracle AI Database

- Check the [Upgrade Guide](#)

your key to
Successful Database Upgrades

AutoUpgrade

Only supported method for upgrading to Oracle AI Database 26ai





Oracle AI Database 26ai

What's New



STANDARD TELEVISION

OMEGA MART

WELCOME
To Fabulous
LAS VEGAS
NEVADA

Relive the highlights of Oracle AI World

Join us next year: October 26-29, 2026 | Las Vegas

Oracle AI World is officially a wrap! Thanks to everyone who joined us in person or On Air. Missed something or want to revisit key moments? Keynotes and select sessions are available on Oracle AI World On Air until Nov 17.

Subscribe for updates and get \$100 off Oracle AI World 2026—keep learning and stay connected.

[Get your free pass](#)

[Subscribe for Oracle AI World 2026](#)

World

Thank you
for attending

Join us next year!
October 26 – 29, 2026

Deloitte

www.oracle.com/ai-world/

Keynotes on demand

Missed the keynotes live? Catch up anytime with Oracle AI World On Air.

Watch on demand



Oracle AI: Powering Your Business

Hear Oracle Chief Executive Officer Mike Sicilia discuss how Oracle embeds AI into every layer of our technology stack. Oracle's vision is AI built end-to-end, from the data platform to cloud infrastructure to the applications that run your business.



Oracle Vision and Strategy

Larry Ellison, chairman of the board and CTO of Oracle, is here live to talk about Oracle's latest AI innovations across databases, infrastructure, and applications. Hear what Oracle's doing today and what it plans for the future of AI to help move your business forward.



The "AI for Data" Revolution is here— How to Survive and Thrive

Juan Loaiza, EVP of Oracle Database Technologies, and T.K. Anand, EVP of Oracle Healthcare and Analytics, discuss how AI is completely reshaping data and apps—right now. Get a glimpse into what the AI-driven future looks like, and learn how you can thrive in the AI for data era.



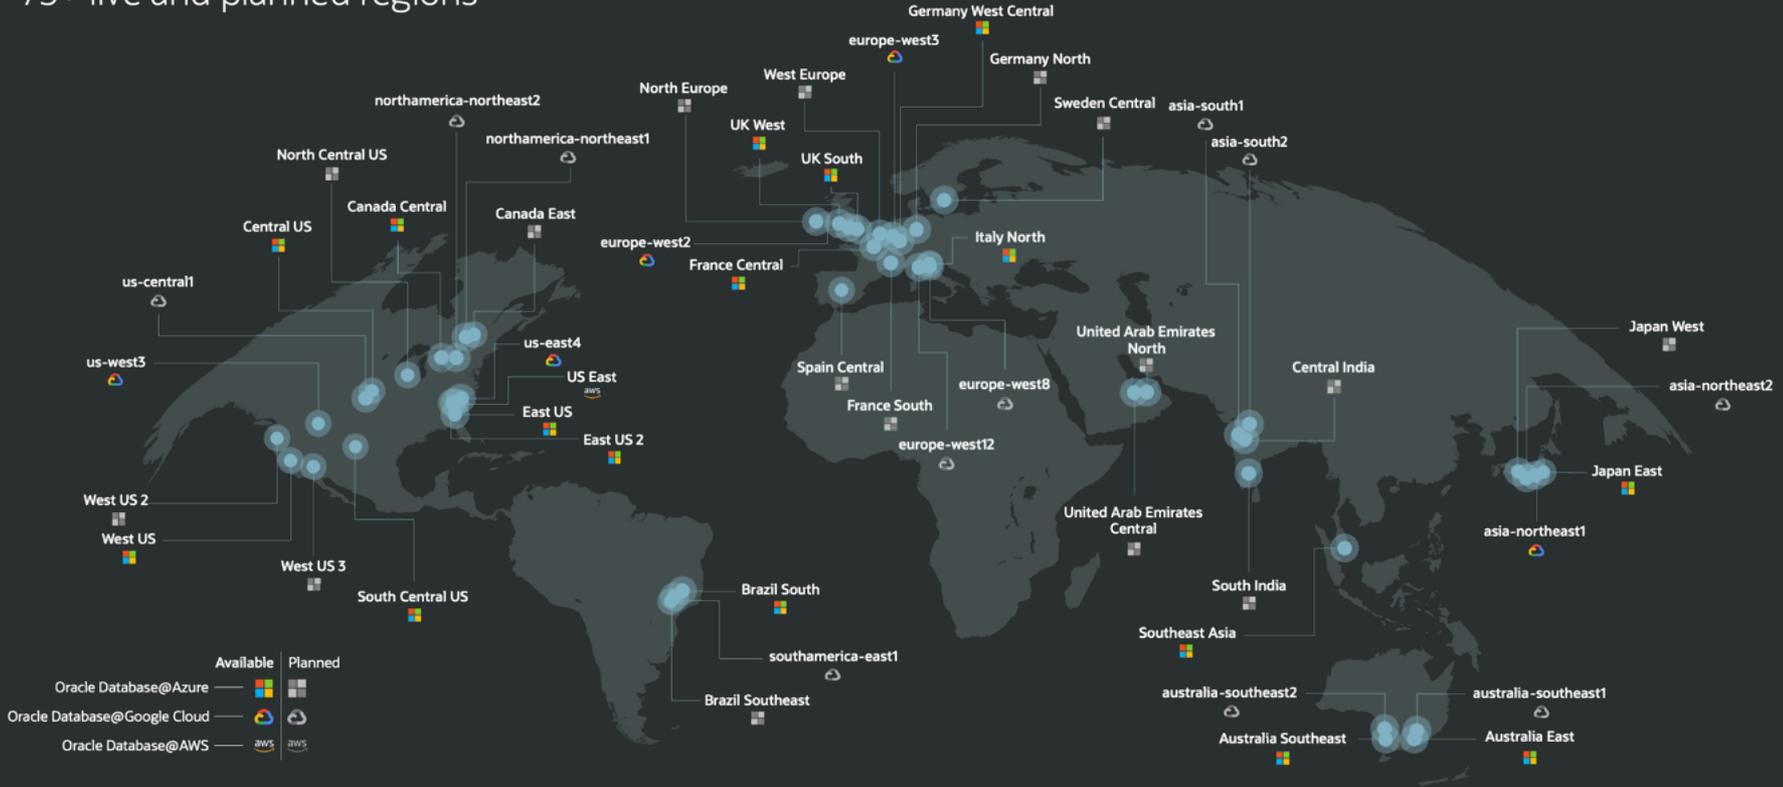
Building the Cloud for You

Whether you need big AI solutions, the smallest sovereign footprint, security, or the ability to get your data in any cloud with Oracle-level performance, come hear Oracle Chief Executive Officer, Clay Magouyrk, talk about the enterprises who are winning and innovating with Oracle.



Oracle Cloud Infrastructure multicloud regions

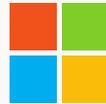
75+ live and planned regions¹



OCI Multicloud

Choice of
Deploying Oracle Database in any Cloud

ORACLE
CLOUD
Infrastructure



aws





Examine Oracle Database changes using ORAdiff

- Oracle Release Analyzer Diff Utility
- <https://oradiff.oracle.com>



- Home
- Parameters
- Users, Roles, etc
- Privileges
- Included Fixes
- Oracle Database Home
- Objects
- Fixed Objects
- Audit
- System
- Scheduler
- Release & Updates Timeline
- Report
- Custom Patch Sets



Oracle Release Analyzer Diff Utility

Find the differences between two Oracle Database releases

Oracle Release Analyzer Diff Utility (ORAdiff) allows you to compare two database releases to each other - with or without patch bundles on top.

Get Started

1. Click the menu icon on the upper left corner of the page
2. Select an object type from the left-hand navigation menu
3. Choose the Source and Target versions and patch levels
4. View the report. You may optionally choose a filter



The tool will display the differences such as "new tables", "added parameters", "changed columns", "removed privileges" and much more. It search can tell you when a parameter was added and which files changed in your Oracle Home.

ORAdiff

ORAdiff - Find the differences between two Oracle Database releases

ORAdiff allows you to compare two database releases to each other - with or without patch bundles on top.

Get Started

1. Click the menu icon on the upper left corner of the page
2. Select an object type from the left-hand navigation menu
3. Choose the Source and Target versions and patch levels
4. View the report. You may optionally choose a filter

ORAdiff will display the differences such as "new tables", "added parameters", "changed columns", "removed privileges" and much more. ORAdiff search can tell you when a parameter was added and which files changed in your Oracle Home.

ORAdiff data is refreshed when new patch bundles are released to the public.

Questions? Ideas? Enhancement requests? Contact us on: #oradiff-int

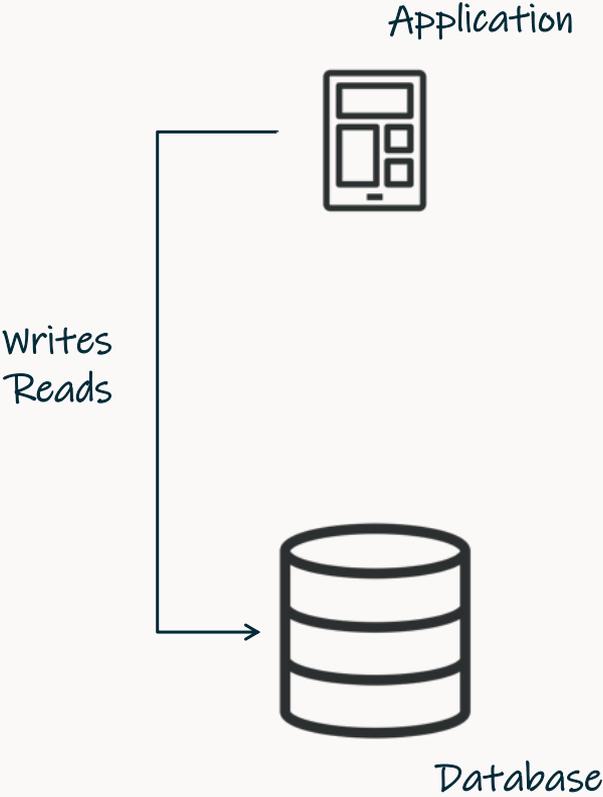
[Watch on YouTube](#)





Accelerate your applications 10x
with True Cache

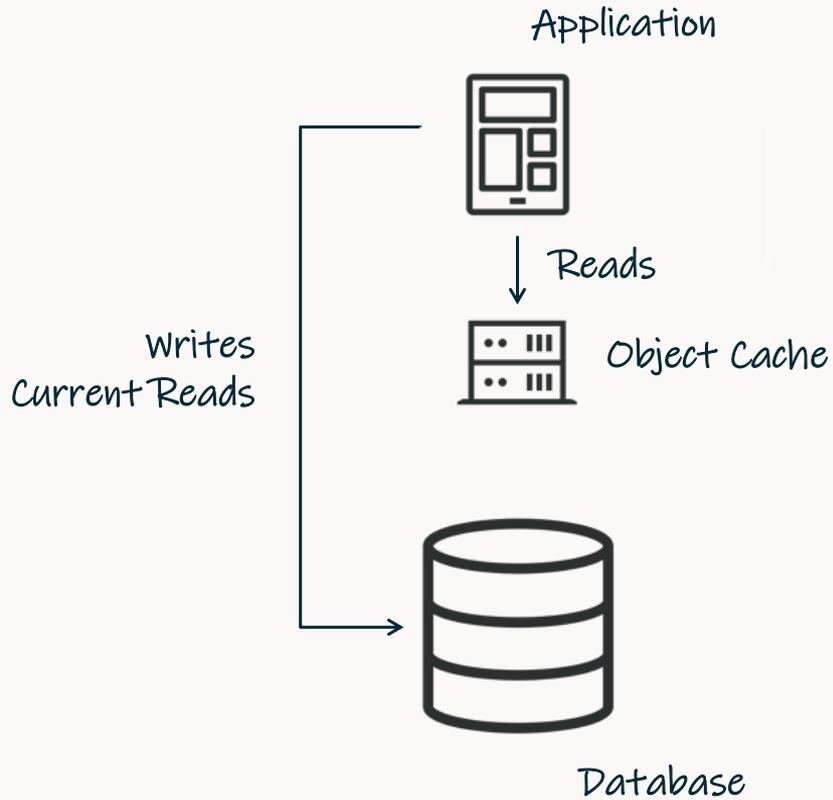
True Cache





We have a performance problem;
we need a cache!

True Cache



True Cache

Who takes care of the cache:

- Populating?
- Consistency?
- Monitoring?
- High availability?
- Security?
- Auditing?
- ...



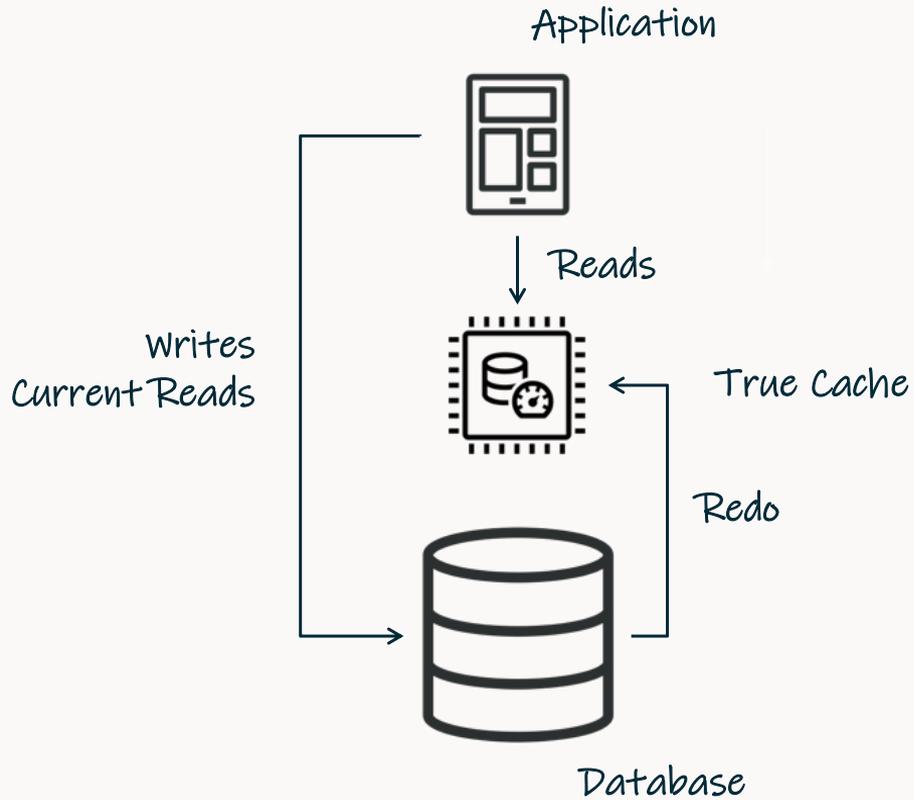


... now we're having a cache problem!



Oracle True Cache - conceptually a diskless Active Data Guard

True Cache





Oracle True Cache is part of Enterprise Edition



Quickly resolve poor performance caused by change in execution plan

- Simplified use of SQL Plan Management



```
--Loads all known plans from cursor cache, AWR,  
--and automatic SQL tuning sets into a SQL plan baseline.  
--Use Evolve Advisor to find the best plan and mark that as accepted.
```

```
select dbms_spm.add_verified_sql_plan_baseline('<sql_id>');
```

Backported to 19.22.0





Allow the database to fix regressing plans automatically

- Automatic SQL Plan Management
- Real-time SQL Plan Management



Automatically recover from certain ORA-00600 errors

- [Automatic Error Mitigation](#)

```
SQL> select ... from ...
```

```
ERROR at line 1:
```

```
ORA-00600: internal error code, arguments: [...]
```

```
SQL> select ... from ...
```

```
ERROR at line 1:
```

```
ORA-00600: internal error code, arguments: [...]
```

```
SQL> alter session set sql_error_mitigation = 'on';
```

```
SQL> select ... from ...
```

```
n rows returned
```



Even wider tables

- Up to 4096 columns
- Be aware of row chaining

```
SQL> alter system set max_columns=extended scope=spfile;
```



Ensure your database clients are updated

- Older clients **do not** support more than 1000 columns

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

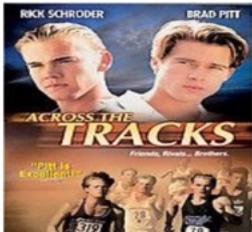


Vector Search

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

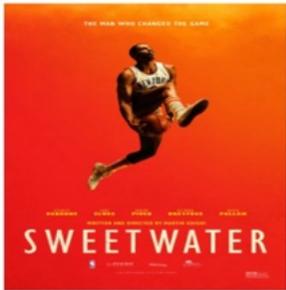
Search Type

Text Search Vector Search

Search

Results

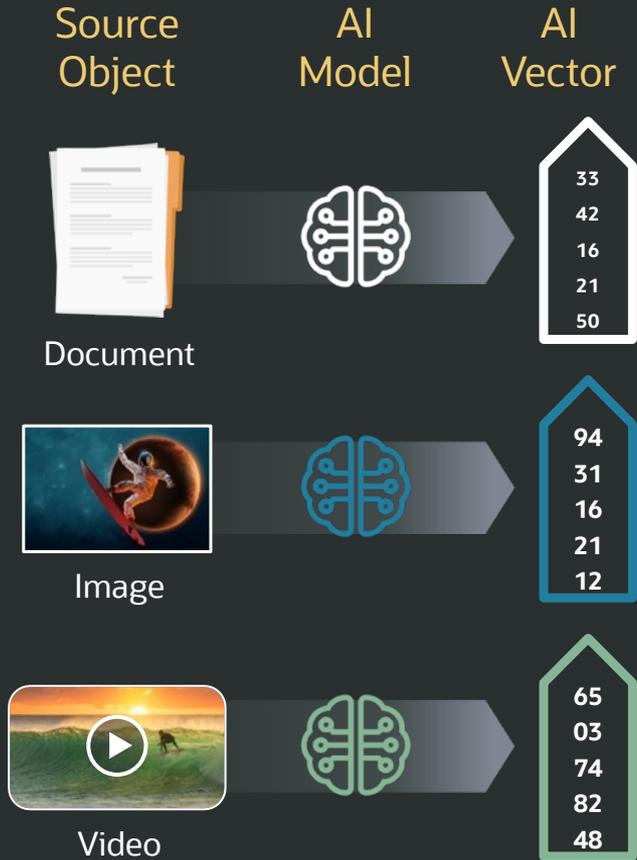
1 - 3 of 3

	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

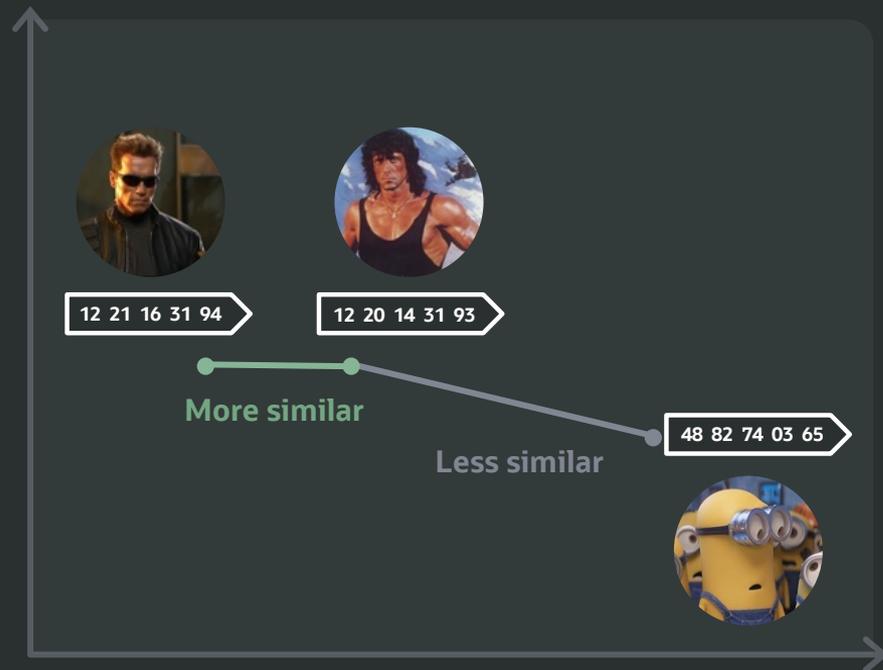
1 - 3 of 3

Key new AI data type: AI Vector

Stores semantic content
of a complex object



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**

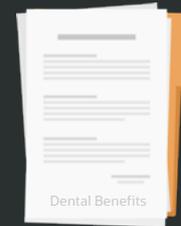


AI Database **unifies** search across traditional data and AI vectors



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



Vector



Spatial



Graph



Relational



Text

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



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

Retrieval Augmented Generation

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 ...



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 = true`

How to compare complex data types?



How to search for search complex data types?



Vector Database

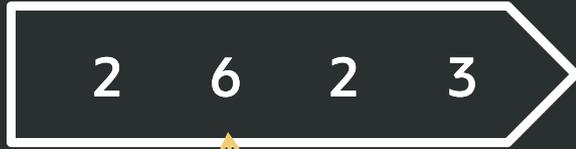


- Search vector data using vector indexes
- Calculate the vector distance



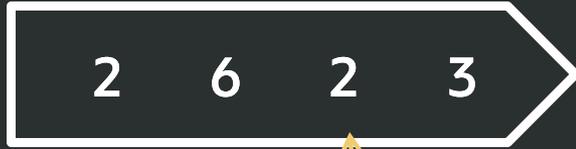


Type of movie



Genre





Mood



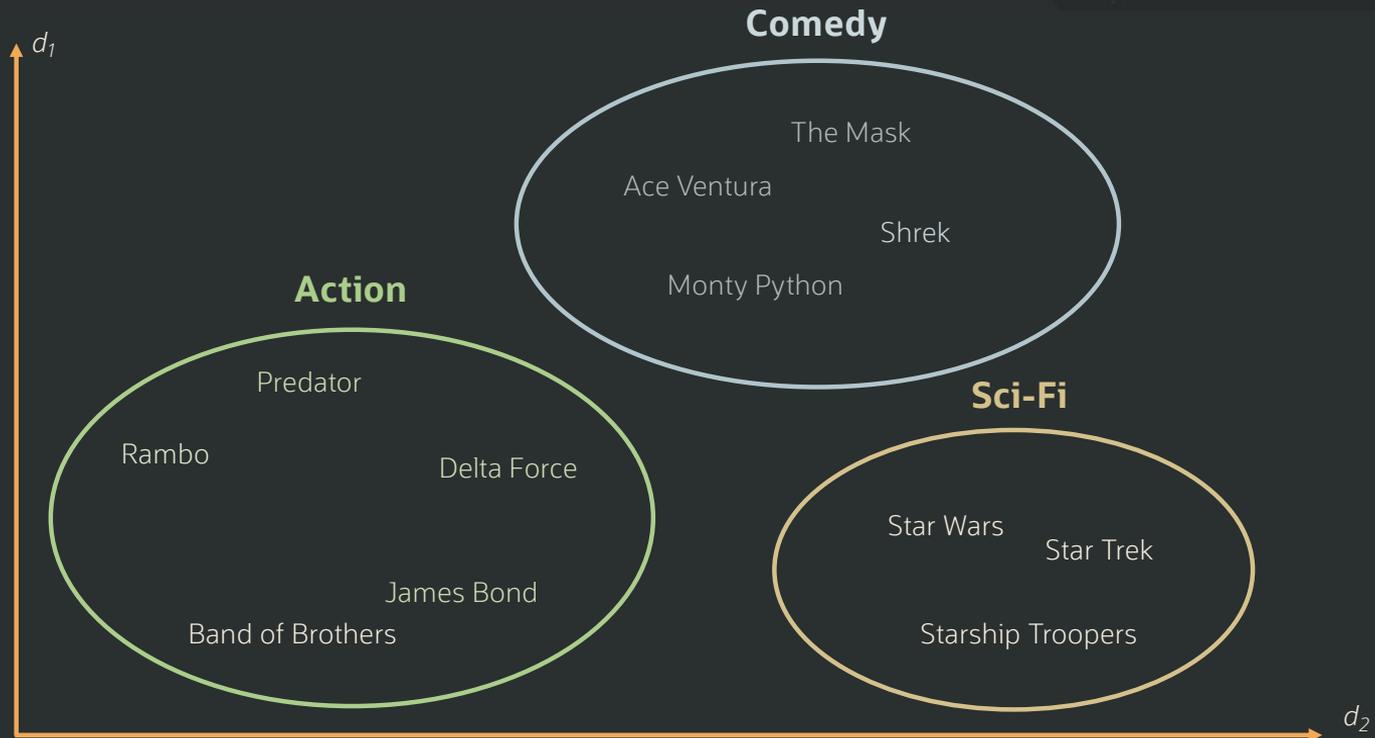
$$\begin{aligned} &\text{Distance (Euclidean Squared)} \\ &= (3-2)^2 + (1-6)^2 + (2-2)^2 + (8-3)^2 \end{aligned}$$

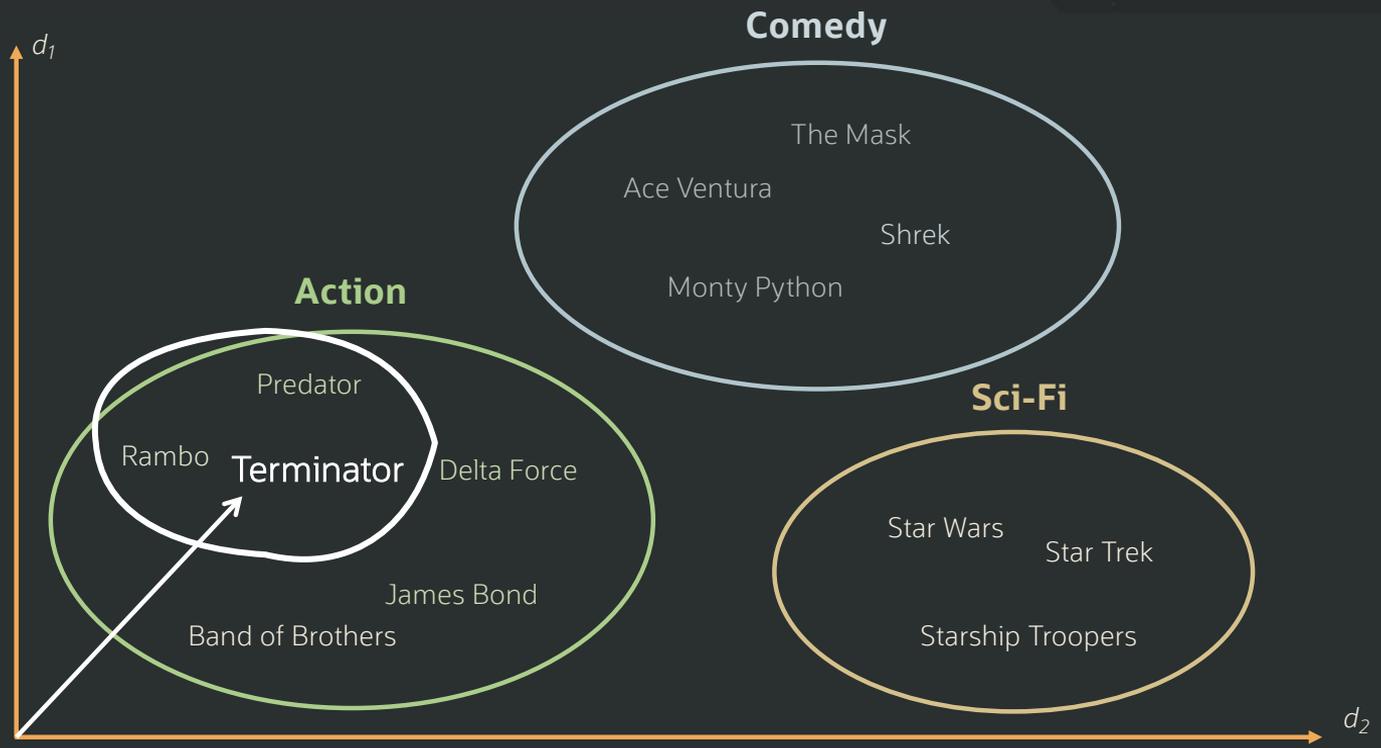
Vector Database



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

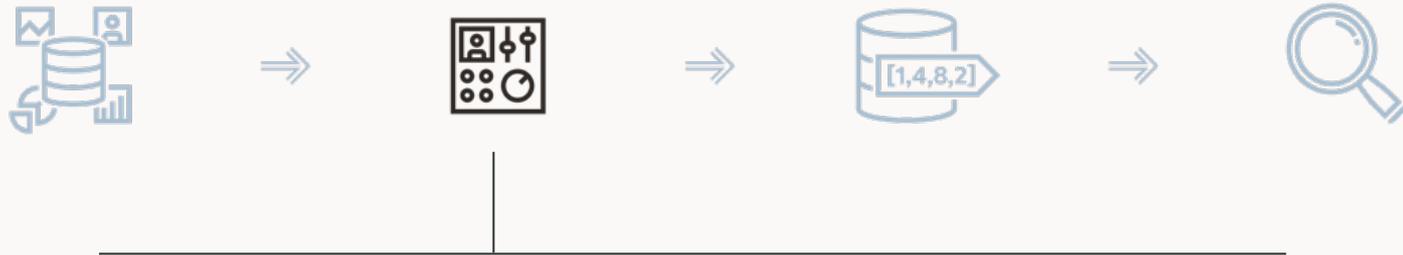






How do you convert your data
into vectors?

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



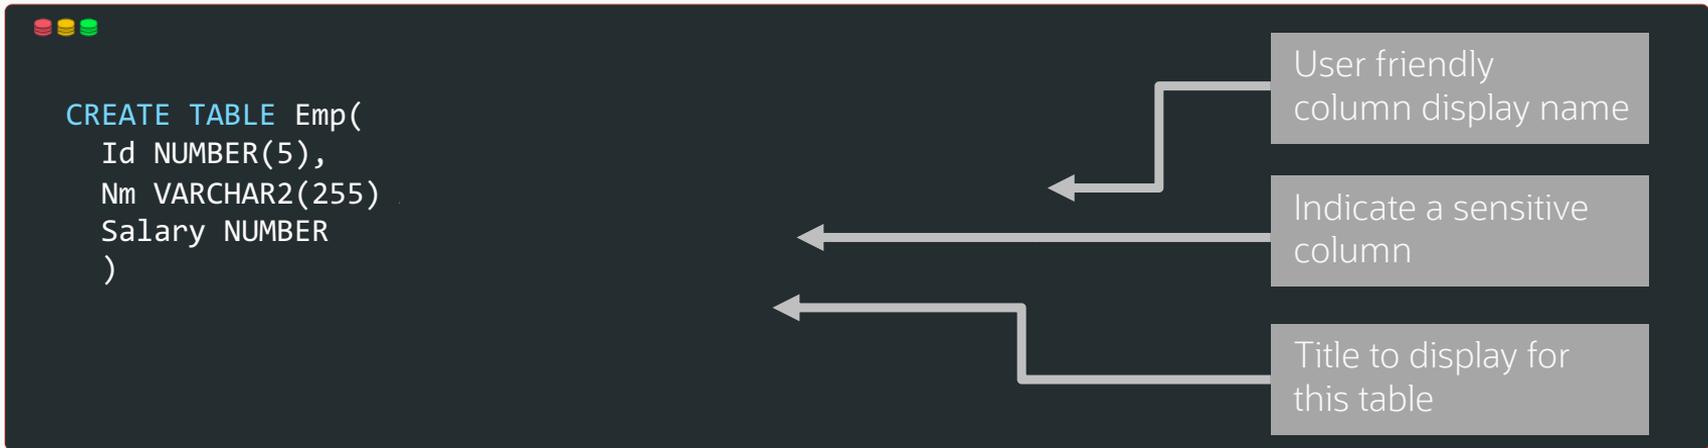


Included in **any** edition
of Oracle Database 26ai

- No extra license required

Oracle Database 26ai Annotations

Annotations add **intended usage** to schema elements such as tables, views, columns
User-friendly display names, identify sensitive data, etc.



The diagram illustrates the use of annotations in Oracle SQL. It shows a code block with the following SQL statement:

```
CREATE TABLE Emp(  
  Id NUMBER(5),  
  Nm VARCHAR2(255)  
  Salary NUMBER  
)
```

Three annotations are shown on the right, with arrows pointing to the corresponding parts of the code:

- User friendly column display name**: Points to the `Nm` column definition.
- Indicate a sensitive column**: Points to the `Salary` column definition.
- Title to display for this table**: Points to the `Emp` table name in the `CREATE TABLE` statement.

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





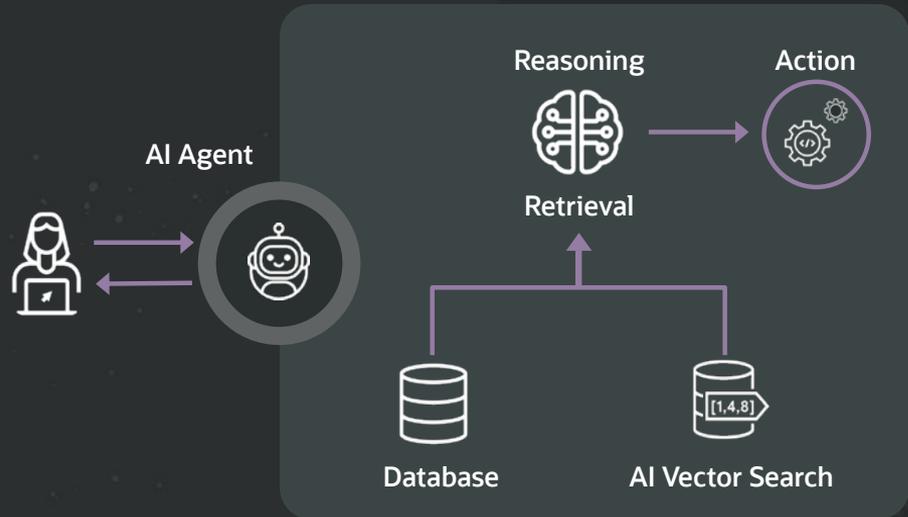
AI Database



Using AI

Architecting Agentic AI into the database

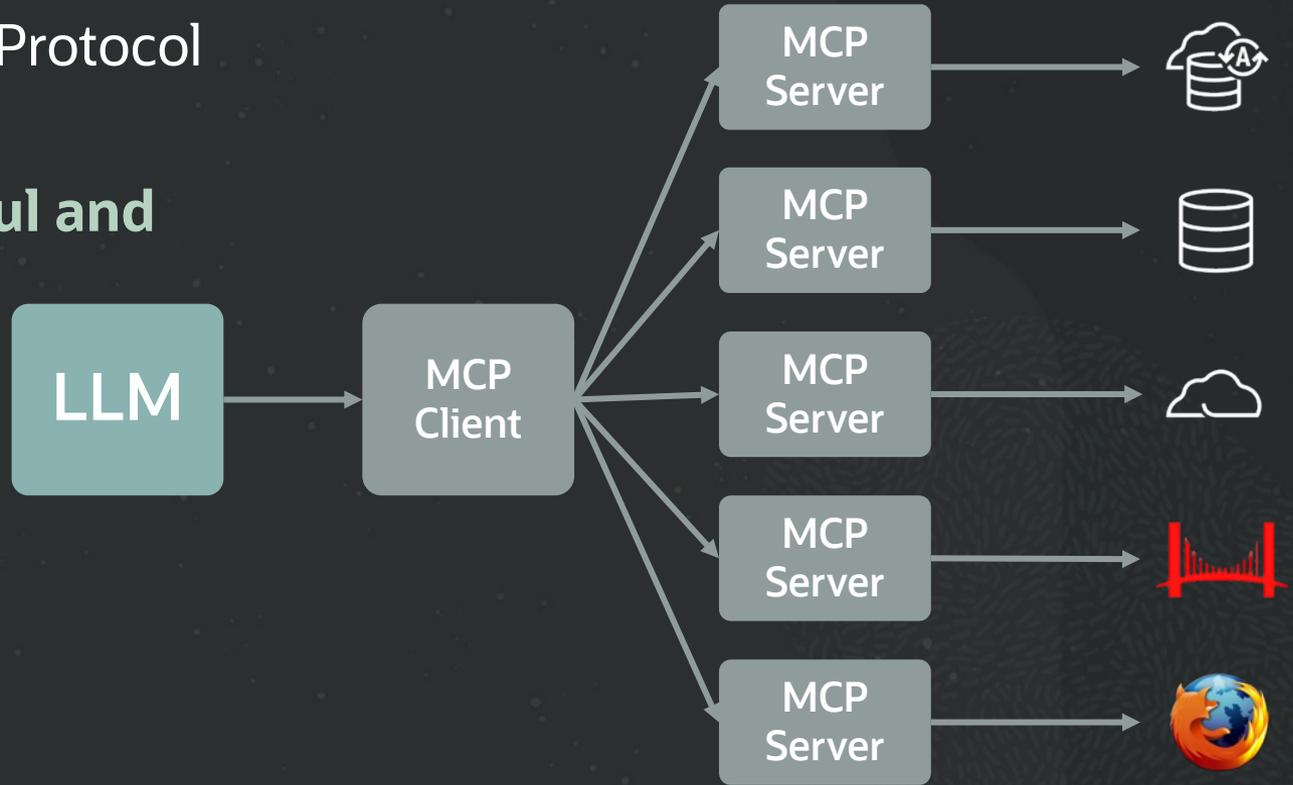
- Select AI Agent
- Private Agent Factory
- SQLcl MCP Server

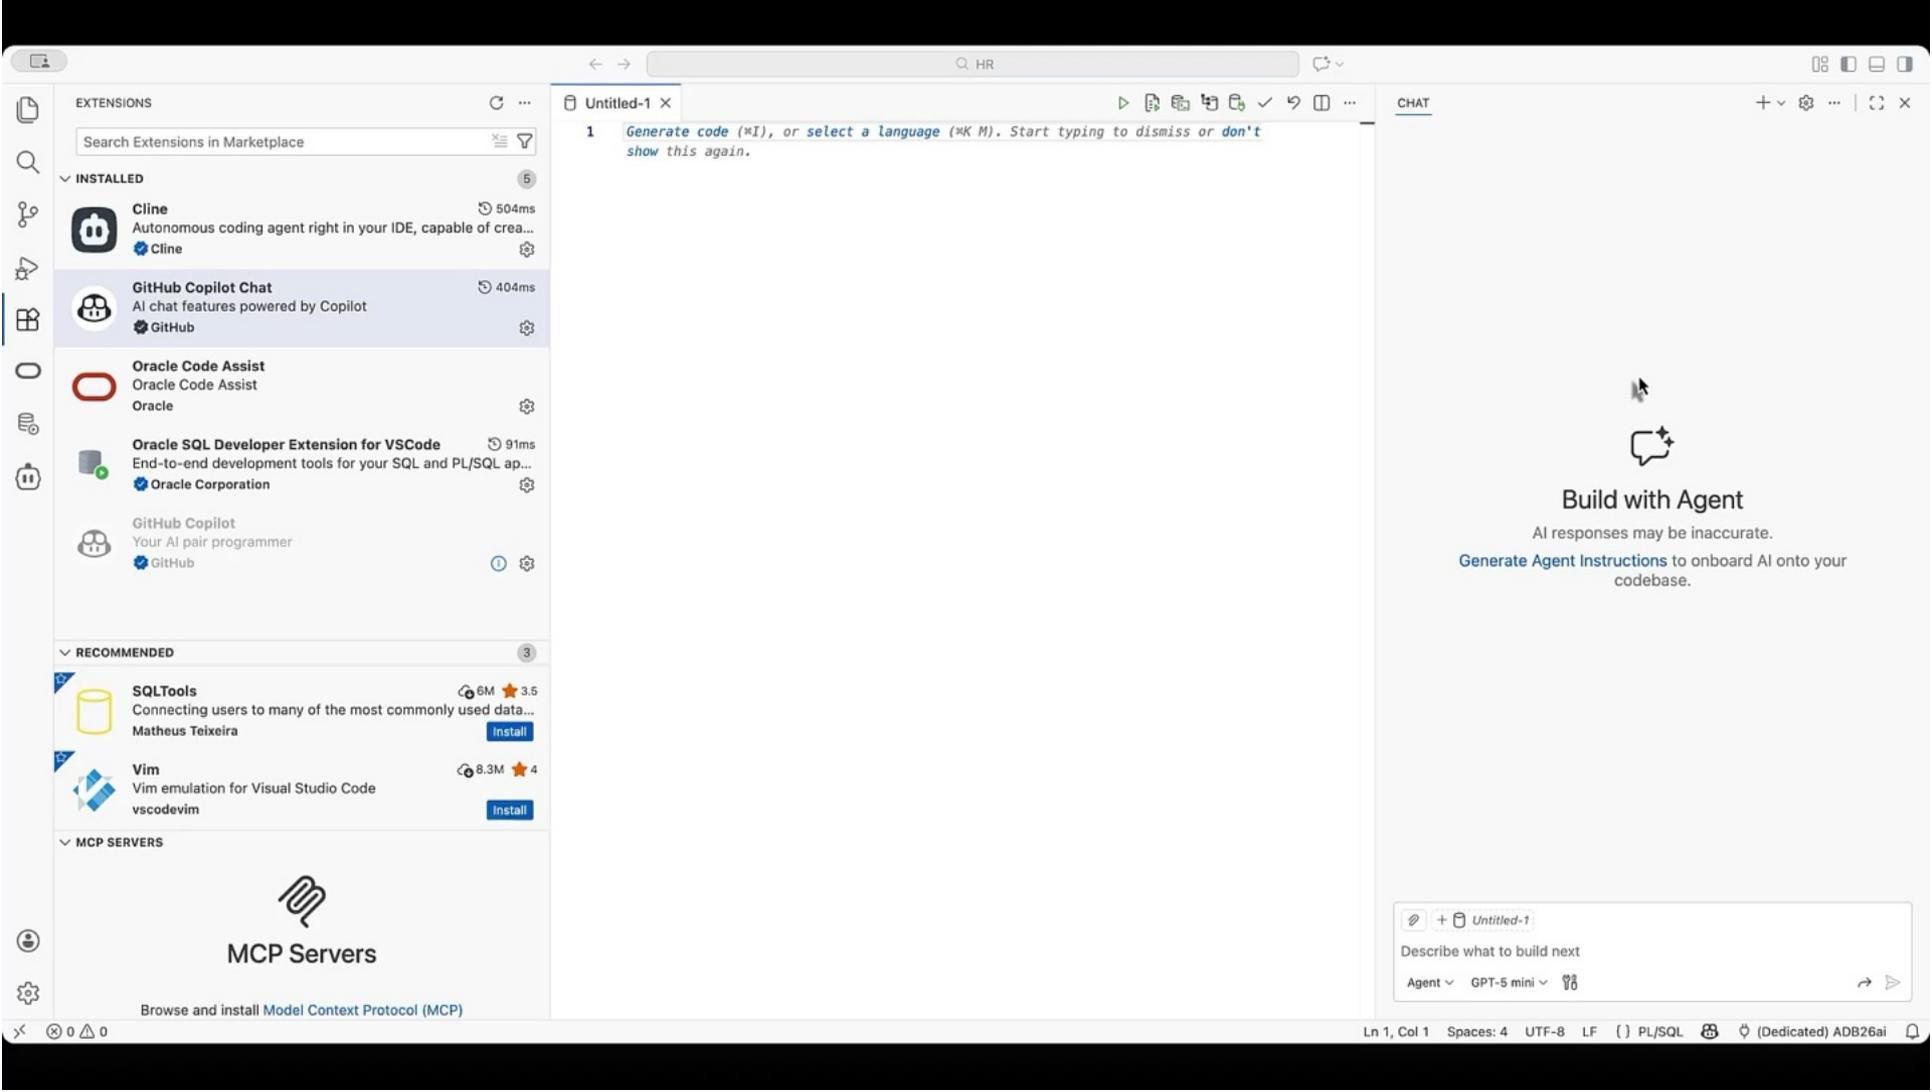


MCP Server

Model Context Protocol

Enable powerful and
simplistic AI





EXTENSIONS

Search Extensions in Marketplace

INSTALLED 5

- Cline**
Autonomous coding agent right in your IDE, capable of crea...
Cline 504ms
- GitHub Copilot Chat**
AI chat features powered by Copilot
GitHub 404ms
- Oracle Code Assist**
Oracle Code Assist
Oracle
- Oracle SQL Developer Extension for VSCode**
End-to-end development tools for your SQL and PL/SQL ap...
Oracle Corporation 91ms
- GitHub Copilot**
Your AI pair programmer
GitHub 1

RECOMMENDED 3

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

MCP SERVERS



Browse and install Model Context Protocol (MCP)

Untitled-1 X

```
1 Generate code (M), or select a language (MK M). Start typing to dismiss or don't show this again.
```

CHAT



Build with Agent

AI responses may be inaccurate.

[Generate Agent Instructions](#) to onboard AI onto your codebase.

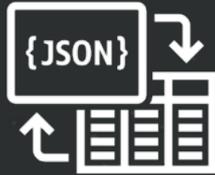
+ Untitled-1

Describe what to build next

Agent GPT-5 mini

What else?

More very important features



Get the best of both worlds
with JSON Duality Views

JSON Duality

The diagram illustrates the duality between relational and document data. On the left, four relational tables are shown: STUDENT, COURSE, STUDENT COURSES, and TEACHER. A large blue plus sign is positioned between these tables and a JSON document on the right.

STUDENT			
STUID	SNAME	MAJOR	YEAR
S3245	Jill	Math	First
...
...
...

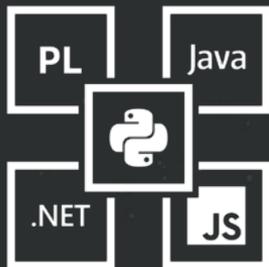
COURSE				
CID	CLASS	ROOM	TIME	TCHID
C123	MATH 201	A102	14:00	T543
C345	SCIENCE 102	B405	16:00	T789
...
...

STUDENT COURSES	
STUID	CID
S3245	C123
...	...
S3245	C345
...	...

TEACHER		
TCHID	TEACHER	TINFO
...
T543	Adam	...
T789	Anita	...
...

```
SCHEDULE FOR: JILL
{
  "student": "S3245",
  "name": "Jill",
  "major": "Math",
  "schedule": [
    {
      "time": "14:00",
      "course": "Math 201",
      "room": "A102",
      "teacher": "Adam"
    },
    {
      "time": "16:00",
      "course": "Science 102",
      "room": "B405",
      "teacher": "Anita"
    }
  ]
}
```

A single database can now support both relational and document data providing the **benefits of both**



Easily convert your existing application with Oracle Database API for MongoDB

- Using [Oracle Database API for MongoDB](#)



No more insecure case insensitive password

- 10G password verifies are no longer accepted





Grant privileges to an entire schema in one command

- `grant ... on schema ... to ...`



```
--Regardless of whether the object exists or not,  
--the DROP command don't produce an error
```

```
SQL> drop table t1;
```

```
ERROR at line 1:
```

```
ORA-00942: table or view does not exist
```

```
SQL> drop table if exists t1 ... ;
```

```
Table dropped
```



--Regardless of whether the object exists or not,
--the CREATE command don't produce an error

```
SQL> create table t1 ( ... );
```

ERROR at line 1:

```
ORA-00955: name is already used by an existing object
```

```
SQL> create table if not exists t1 (c1 number);
```

Table created



```
--Group by expression must be written in full  
--Works for HAVING clause as well
```

```
select  owner, to_char(created, 'YYYYMM'), count(*)  
from    dba_objects  
group by owner, to_char(created, 'YYYYMM');
```



```
--Group by expression must be written in full  
--Works for HAVING clause as well
```

```
alter session set group_by_position_enabled=true;
```

```
select  owner, to_char(created, 'YYYYMM'), count(*)  
from    dba_objects  
group by 1, 2;
```



```
sqlplus appuser@alias_does_not_exist
```

ERROR:

```
ORA-12154: Cannot connect to database. Could not find alias  
alias_does_not_exist in
```

```
/opt/oracle/product/23ai/dbhome_1/network/admin/tnsnames.ora.
```

```
Help: https://docs.oracle.com/error-help/db/ora-12154/
```



```
SQL> ping salesgold
```

```
Network service name mapping file:
```

```
/opt/oracle/product/23ai/dbhome_1/network/admin/tnsnames.ora
```

```
Attempting to contact: (DESCRIPTION = (CONNECT_TIMEOUT=5) (RETRY_COUNT=2) (RETRY_DELAY=3)  
(TRANSPORT_CONNECT_TIMEOUT=3) (ADDRESS_LIST = (LOAD_BALANCE=on) (ADDRESS = (PROTOCOL =  
TCP)(HOST=localhost)(PORT=1521))) (CONNECT_DATA= (SERVICE_NAME = pdb1)))
```

```
Ok (1.177 msec)
```



```
--Examines a bigfile tablespace to find objects that can be moved  
--to the start of the data files. In the end, shrink the data file  
--to release the space.  
--Tablespace MUST be ASSM managed. Does not work in SYSTEM and SYSAUX.
```

```
exec dbms_space.tablespace_shrink('USERS');
```



Tablespace Shrink

- Dependent objects and cursors might be invalidated
- Objects are moved in an online manner
- Purges recycle bin

the most
important change

```
select sysdate from dual;  
select sysdate;
```



Key Learnings



- 1 Move to Oracle AI Database 26ai
- 2 Use AutoUpgrade
- 3 Adopt AI features



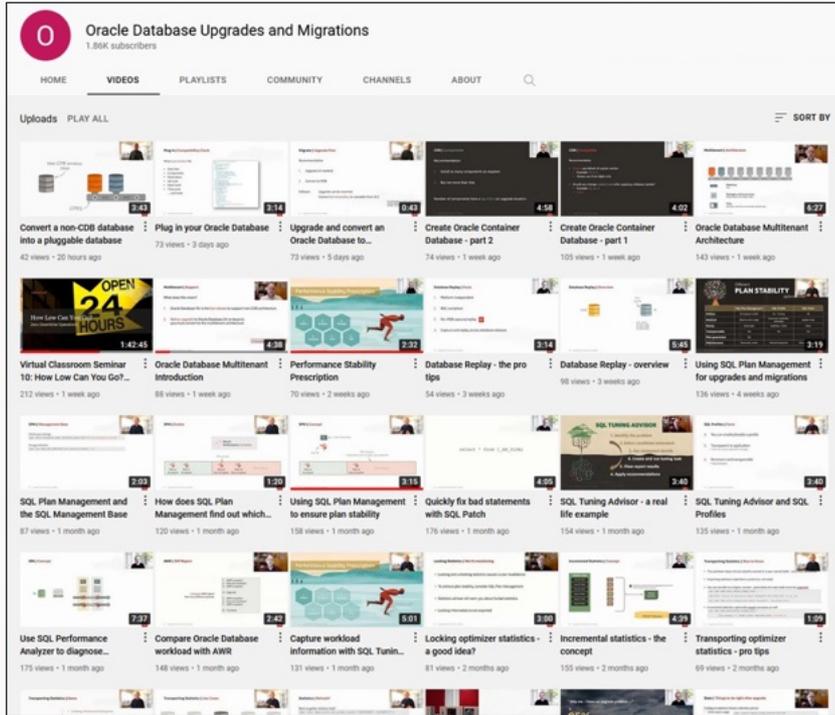
Oracle

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
