ORACLE

Real World Oracle Database Upgrade and Migration 19c & 23ai

Prague, September 2025

Photo by Denis Poltoradnev on Unsplash

Oracle

DBAs

run the world

MIKE DIETRICH
Vice President
Database Upgrade, Migrations & Patching

mikedietrich

@mikedietrichde.com

https://mikedietrichde.com

DANIEL OVERBY HANSEN
Distinguished Product Manager
Database Upgrade, Migrations & Patching

dohdatabase

@dohdatabase.com

https://dohdatabase.com

Find Slides and Much More on Our Blogs

MikeDietrichDE.com

Mike.Dietrich@oracle.com

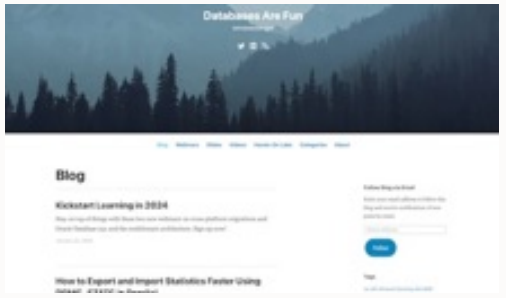

dohdatabase.com

Daniel.Overby.Hansen@oracle.com

DBArj.com.br

Rodrigo.R.Jorge@oracle.com

AlexZaballa.com

Alex.Zaballa@oracle.com

Recorded Web Seminars

https://MikeDietrichDE.com/videos

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

AGENDA

09:30
Welcome
Release Strategy
Patching

10:45
Coffee break

11:00
Upgrade
Multitenant

12:00
Lunch

13:00
Data Pump
Performance Stability
Insights into development

15:00
Coffee break

15:15
Autonomous Database
Oracle Database 23ai
What's New, What's Coming

Release Strategy

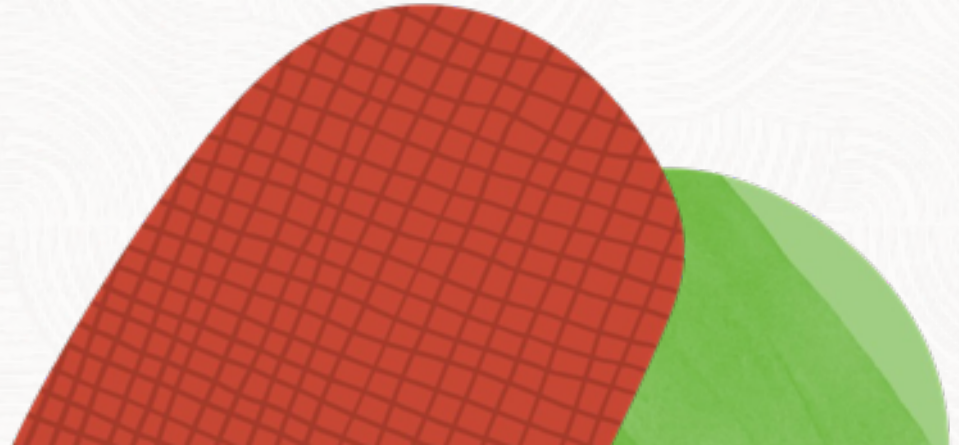

Lifetime Support Policy

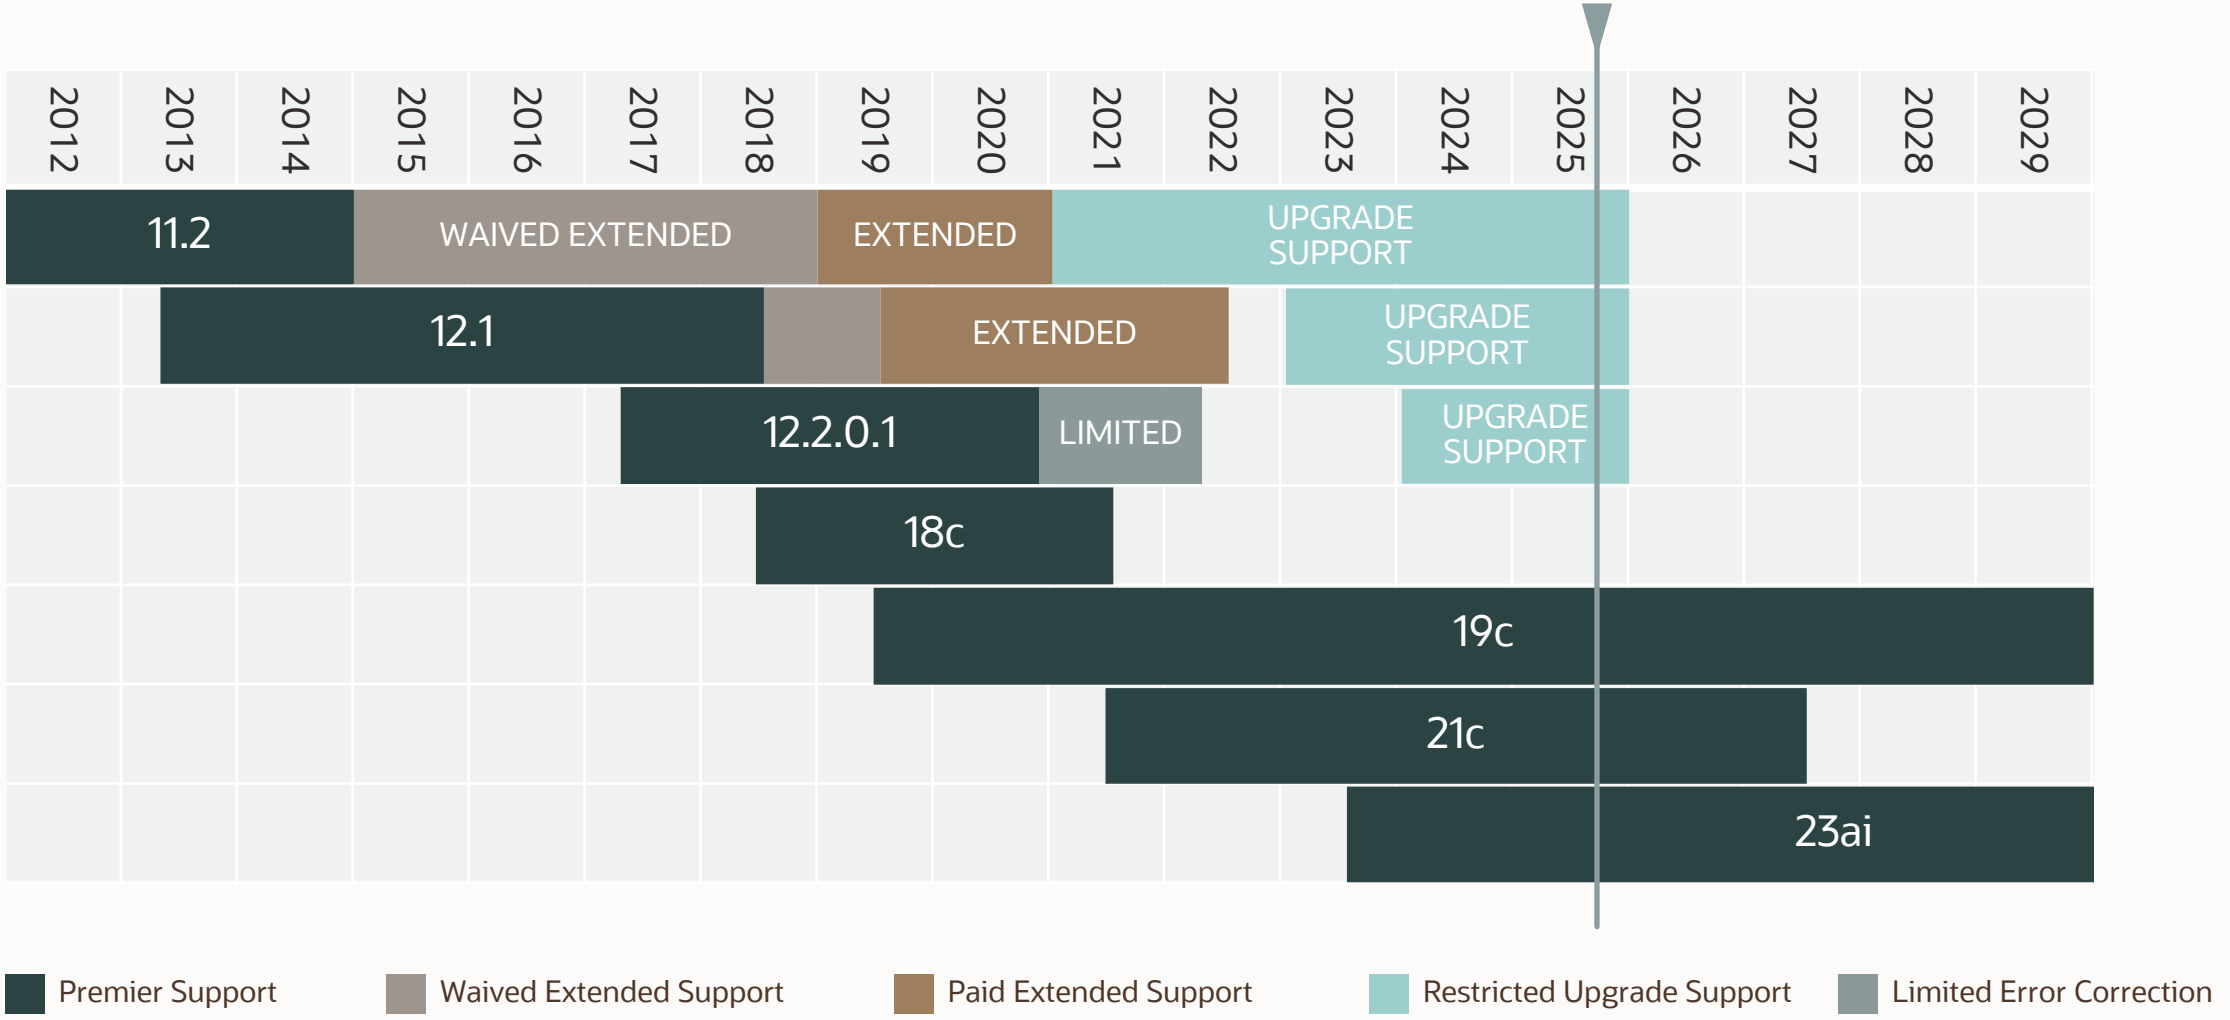

Lifetime Support Policy

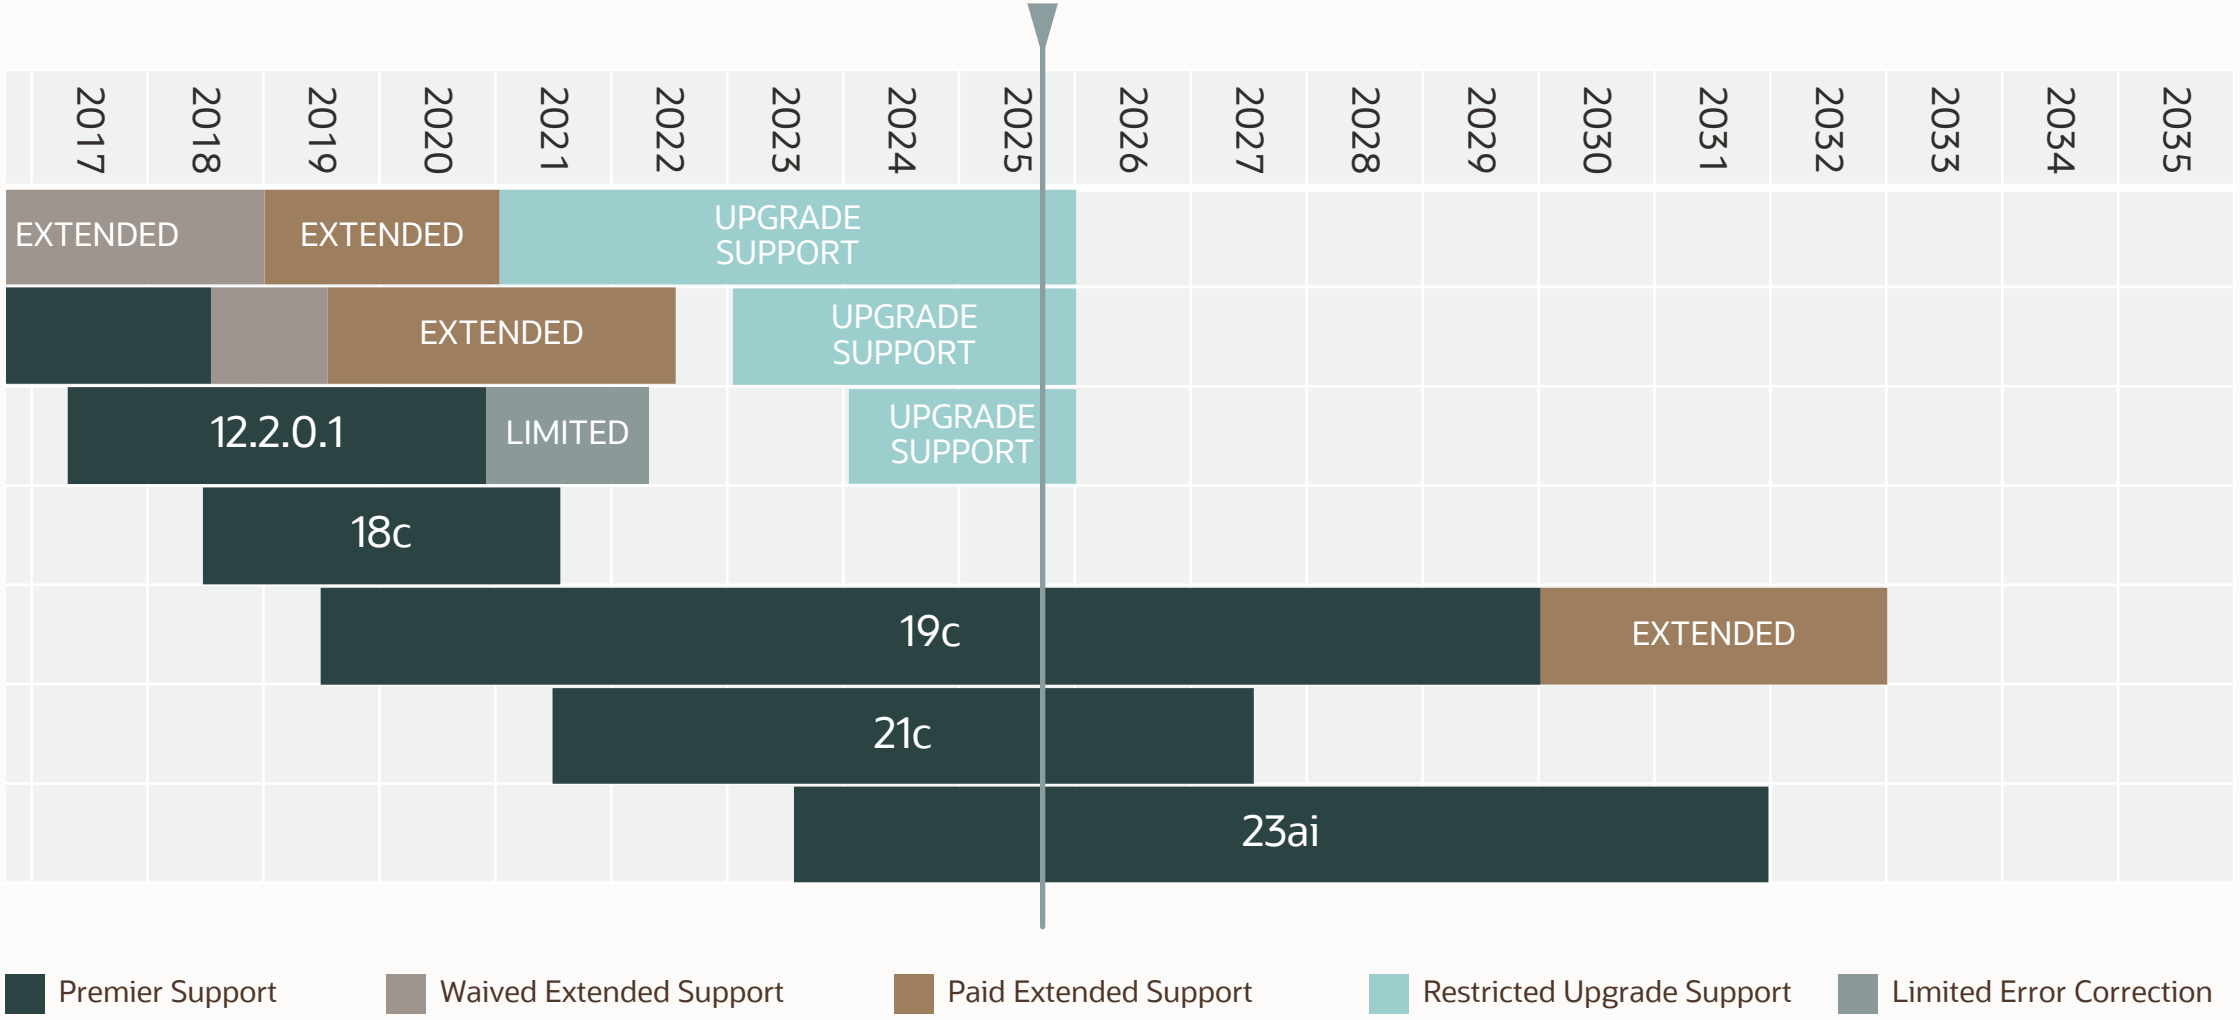

Oracle Database 23ai

Production

- Oracle Database Cloud Services
- Oracle DB@Azure
- Oracle DB@Google Cloud
- Oracle DB@AWS
- Exadata Cloud@Customer
- Compute Cloud@Customer
- Exadata Database Machine
- Database Appliance
- Private Cloud Appliance

Development / Testing

- Oracle Database Free for Linux x86
- Oracle Database Free for ARM
- Oracle Database Free for Windows
- Oracle Autonomous Database
- Oracle Exadata Exascale
- Oracle Base Database
- DIY on Oracle Cloud Infrastructure
- Oracle Database 23ai beta program

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

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

Next Long Term Support release

Oracle Database 23ai

Upgrade possible only from:

- Oracle Database 19c
- Oracle Database 21c

Database and Grid Infrastructure Patching

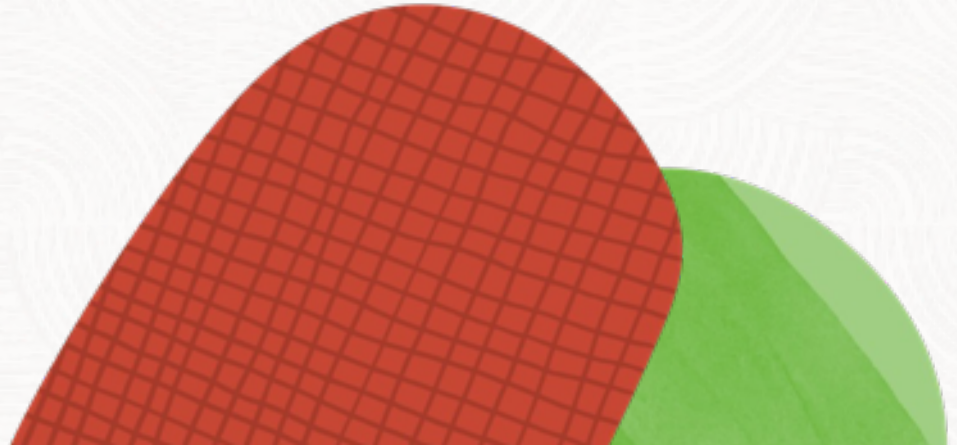

Patching Best Practices

Installation

Basics

Methods

Grid Infrastructure

Datapatch

AutoUpgrade

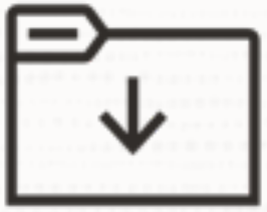

Until Oracle Database 19c, you always start with the base release

- Oracle Database 19.3.0

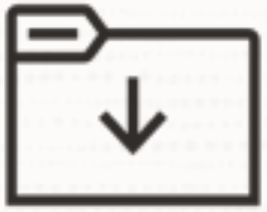

From Oracle Database 23ai onwards, you download a Gold Image

- Updated with latest Release Update

Always Apply the Most Recent RU

Use the Patch Download Assistant MOS Note: 2118136.2

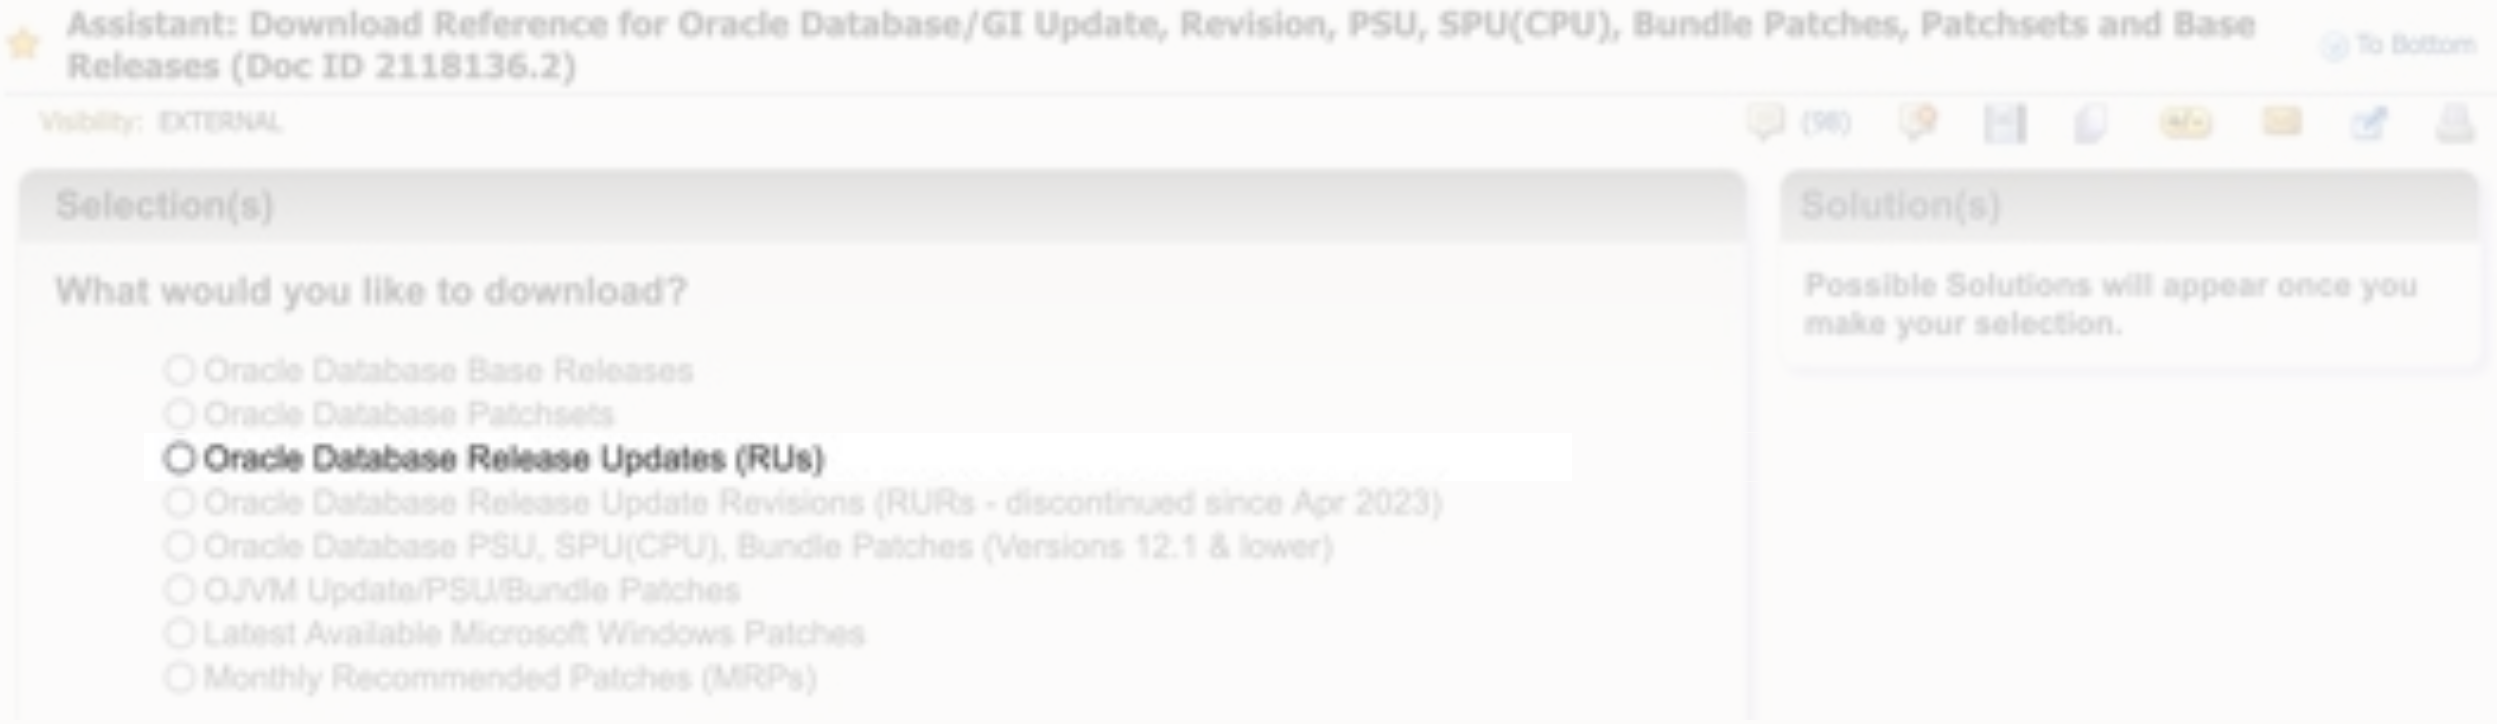

Release Update Contents

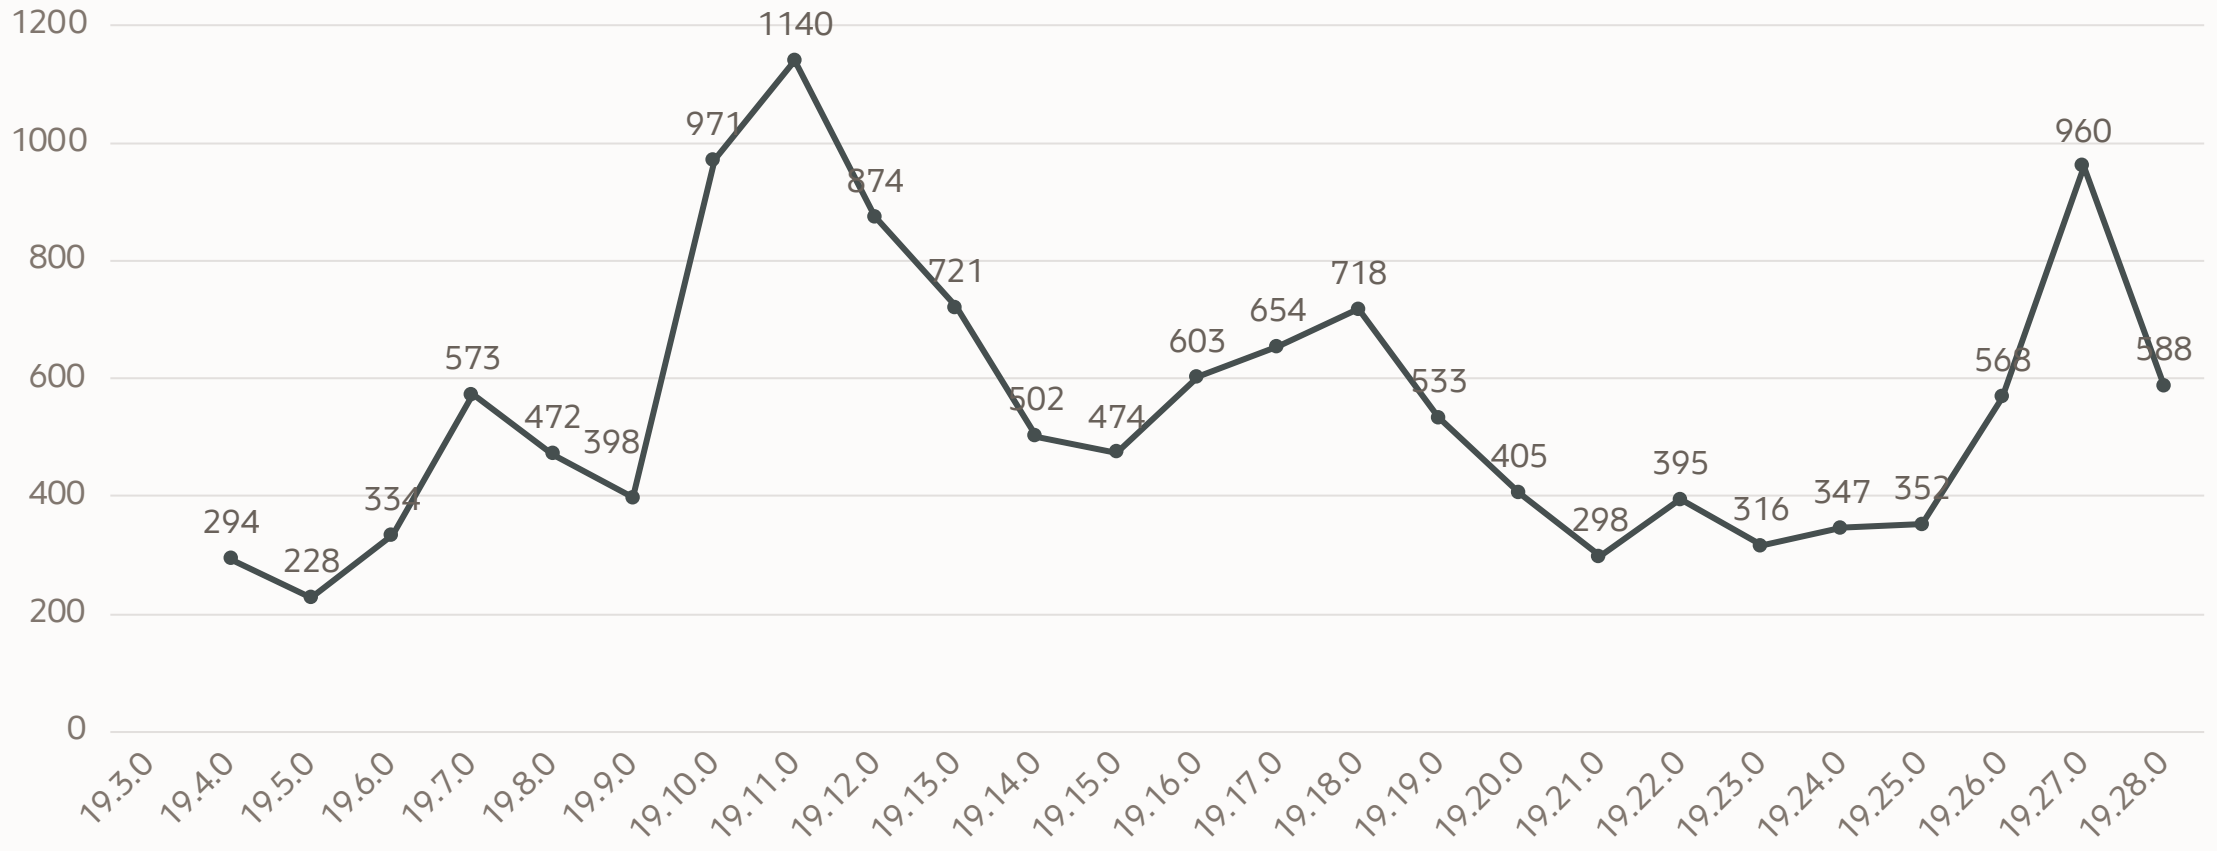

Database 19 Release Updates and Revisions Bugs Fixed Lists (Doc ID 2523220.1)

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

- Almost 14k fixes with 19.28.0
- More than 750 security fixes

Apply the Most Important Patches (555.1)

Recommended Patches for 19.27 DB Home

Below is the list of important patches to consider applying on top of 19.27.

Bug	Fixed in RU	Fixed in MRP	Description	Patches	RAC Rolling Installable	Database Online Installable	Added
34728524	19.28	DBMRP 19.26.0.0.250520, DBMRP 19.27.0.0.250617	[SQL EXECUTION] Wrong Result Or ORA-600 [evaopn2.h:kaf_qeecol] In JPPD With Function Based Indexes	[list-patches]	YES		18-JUL-2025
36501568 (replaces 31061145)	19.28	DBMRP 19.27.0.0.250819, GIMRP 19.27.0.0.250819	[RAC] REMOVE REQUIREMENT TO SET TARGET_PDBS EVEN FOR SMALL SGA SYSTEMS	[list-patches]	YES		18-JUL-2025
37450190	19.28	DBMRP 19.26.0.0.250715, DBMRP 19.27.0.0.250715	[RAM SECUREFILE] Securefile LOB Update Shows High Level of Write Complete Waits	[list-patches]	YES		03-JUL-2025
37693383 (replaces 36760879)	19.28	DBMRP 19.26.0.0.250617, DBMRP 19.27.0.0.250617	[OSS SRV LAYER] ORA-600 [kcfis_get_disk_for_translation:volume lookup] - kcfis_get_disk_for_translation	[list-patches]	YES		03-JUL-2025
37104857	19.28	DBMRP 19.26.0.0.250715, DBMRP 19.27.0.0.250715	[SQL EXECUTION] Hitting ORA-600 [kxhrpack:cvl2] for UPDATE statement	[list-patches]	YES	YES	03-JUL-2025

Monthly Recommended Patches

A collection of recommended one-off fixes
provided at monthly intervals
via a 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					
23ai							23.5.0	23.6.0	23.7.0	23.8.0	23.9.0	23.10.0	23.11.0	23.12.0	23.13.0

Monthly Recommended Patches

	2024			2025									
	October	November	December	January	February	March	April	May	June	July	August	September	October	November
19.25.0	19.25.0	MRP1	MRP2	MRP3	MRP4	MRP5	MRP6							
19.26.0				19.26.0	MRP1	MRP2	MRP3	MRP4	MRP5	MRP6				
19.27.0							19.27.0	MRP1	MRP2	MRP3	MRP4	MRP5	MRP6	
19.28.0										19.28.0	MRP1	MRP2	MRP3	MRP4
19.29.0													19.29.0	MRP1

Sorry, but there is more to talk about …

Apply Additional Important Fixes and Bundles

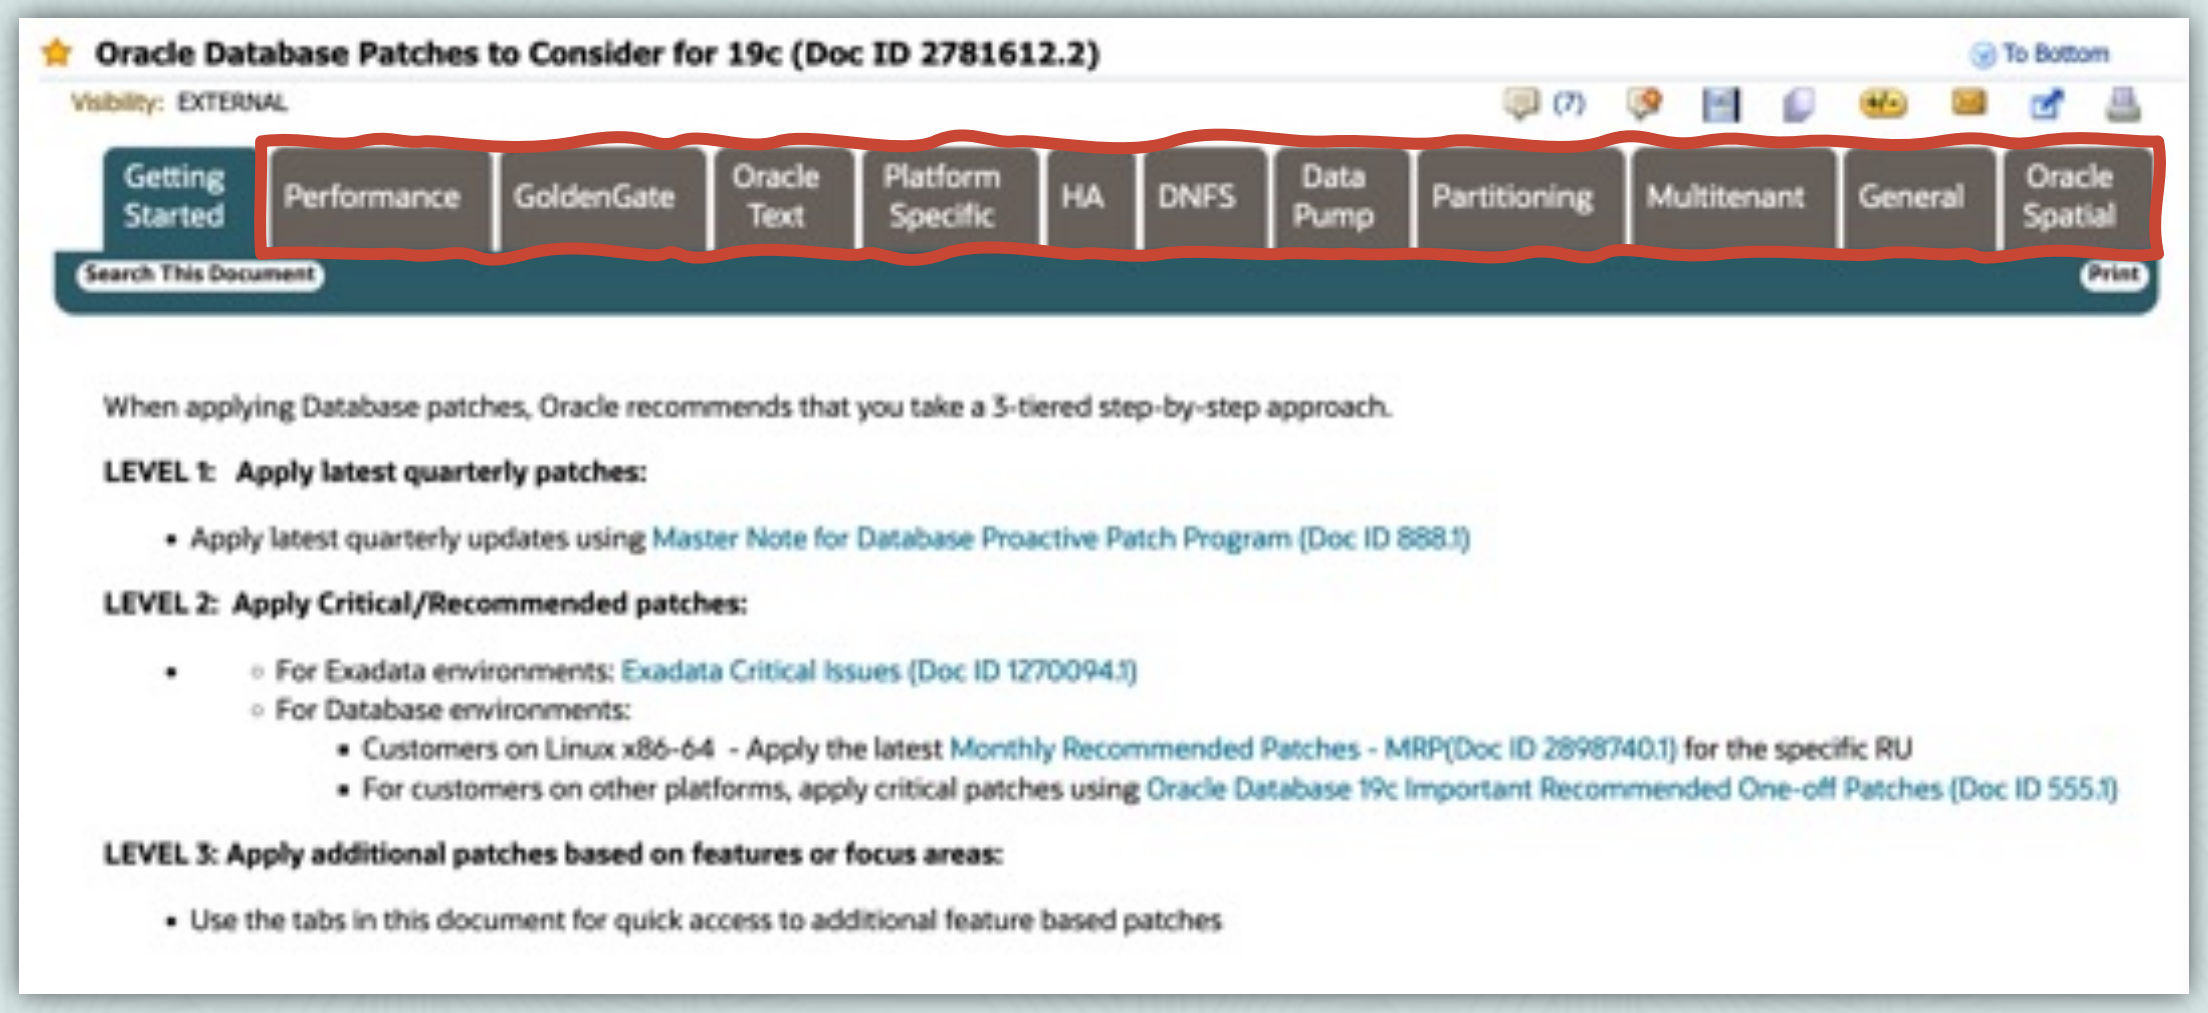

Oracle Database Patches to Consider for 19c (Doc ID 2781612.2)

To Bottom

Visibility: EXTERNAL

(7)

Getting Started | Performance | GoldenGate | Oracle Text | Platform Specific | HA | DNFS | Data Pump | Partitioning | Multitenant | General | Oracle Spatial

Search This Document | Print

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 888.1)

LEVEL 2: Apply Critical/Recommended patches:

- For Exadata environments: Exadata Critical Issues (Doc ID 1270094.1)
- For Database environments:
  - Customers on Linux x86-64 - Apply the latest Monthly Recommended Patches - MRP(Doc ID 2898740.1) for the specific RU
  - For customers on other platforms, apply critical patches using Oracle Database 19c Important Recommended One-off Patches (Doc ID 555.1)

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 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 (Doc ID 1530108.1)

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, 19.28.0

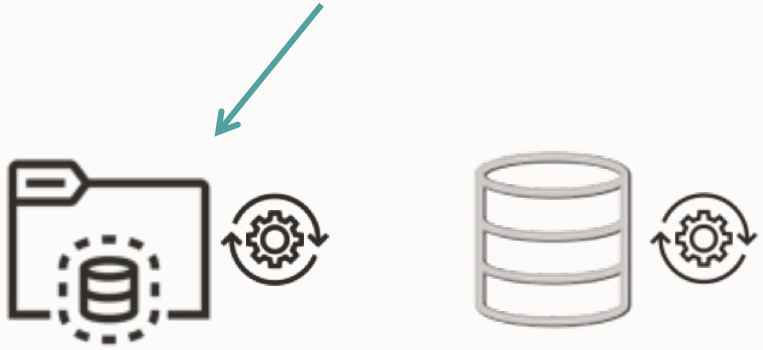

```
$ORACLE_HOME/OPatch/opatch rollback -id ...
```

Out-of-Place Patching

Oracle Home, 19.27.0

```
SQL> SHUTDOWN IMMEDIATE
```

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

New Oracle Home
Oracle Home, 19.28.0

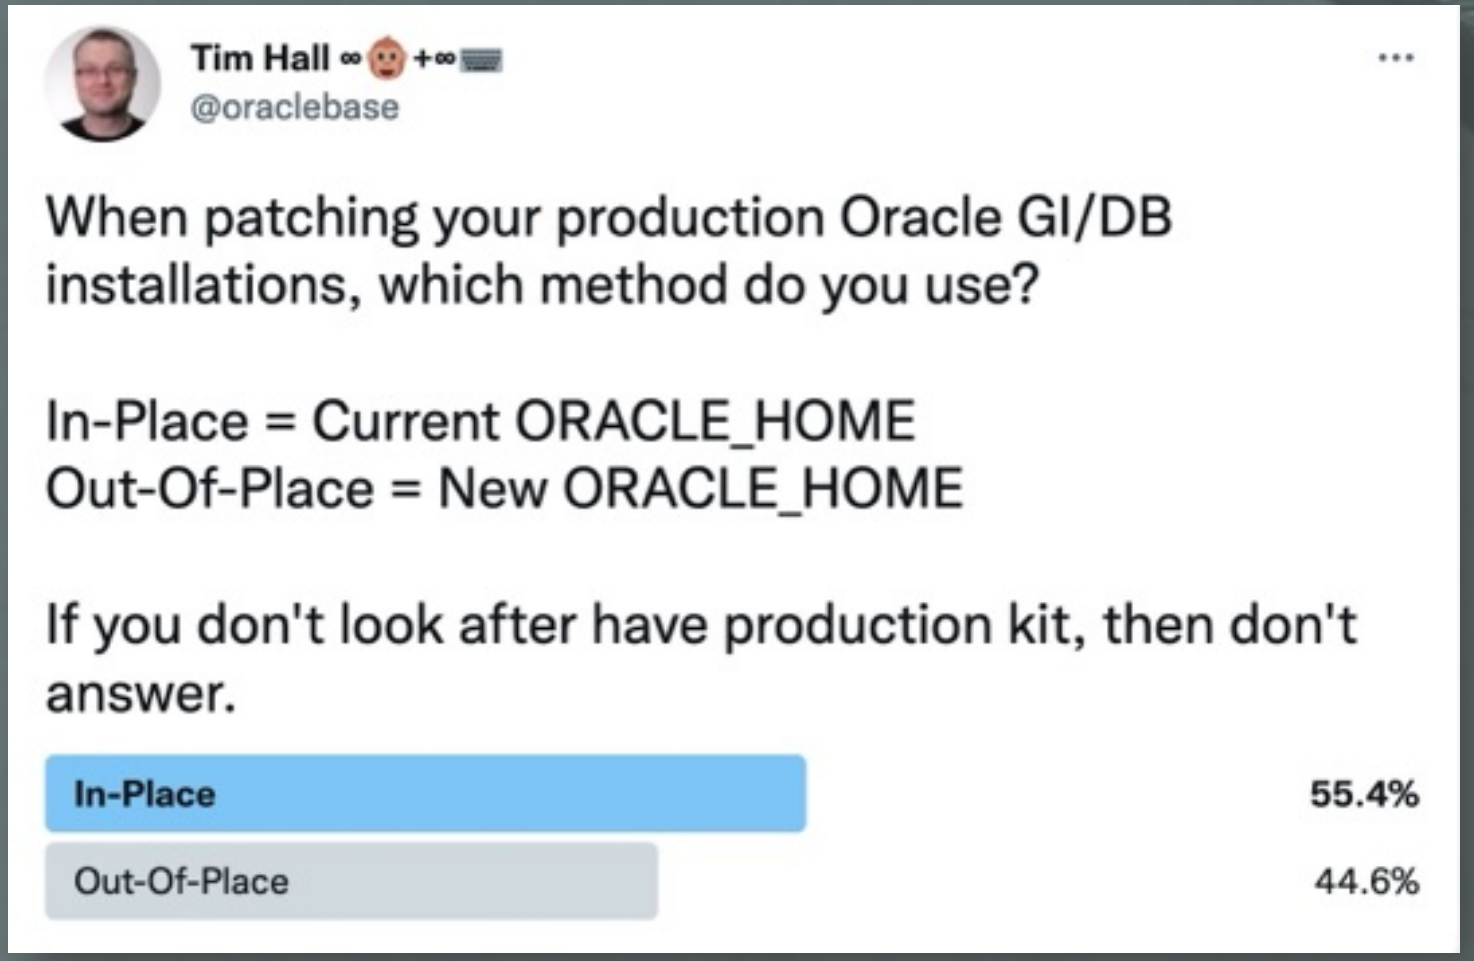

Always patch Out-of-Place

- Don't argue with us ☺

Reduce downtime to the time it takes to perform a switchover

- Data Guard Standby-First Patch Apply (Doc ID 1265700.1)

Safely test and verify patches with Standby-First Patch Apply

- Data Guard Standby-First Patch Apply (Doc ID 1265700.1)

Standby-First Patching

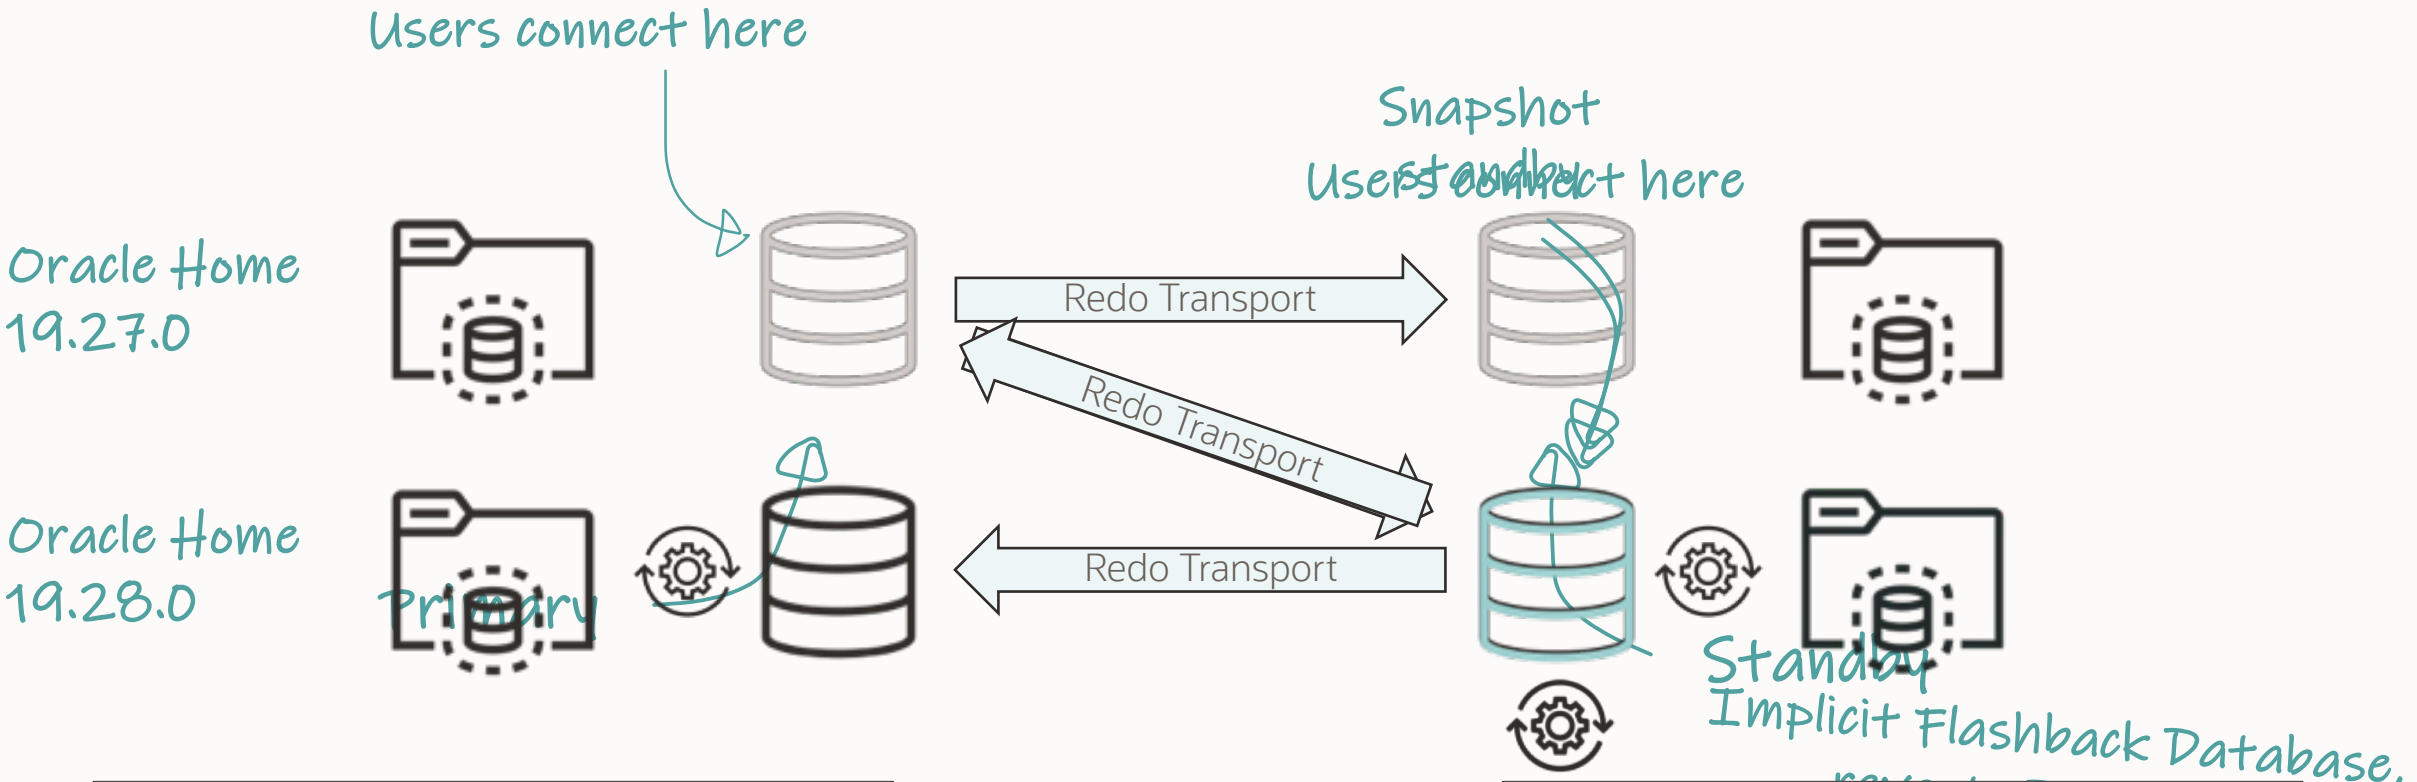


```
java -jar autoupgrade.jar \
   -config ... \
   -mode deploy
```

Patch/datapatch

```
java -jar autoupgrade.jar \
   -config ... \
   -mode deploy
```

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

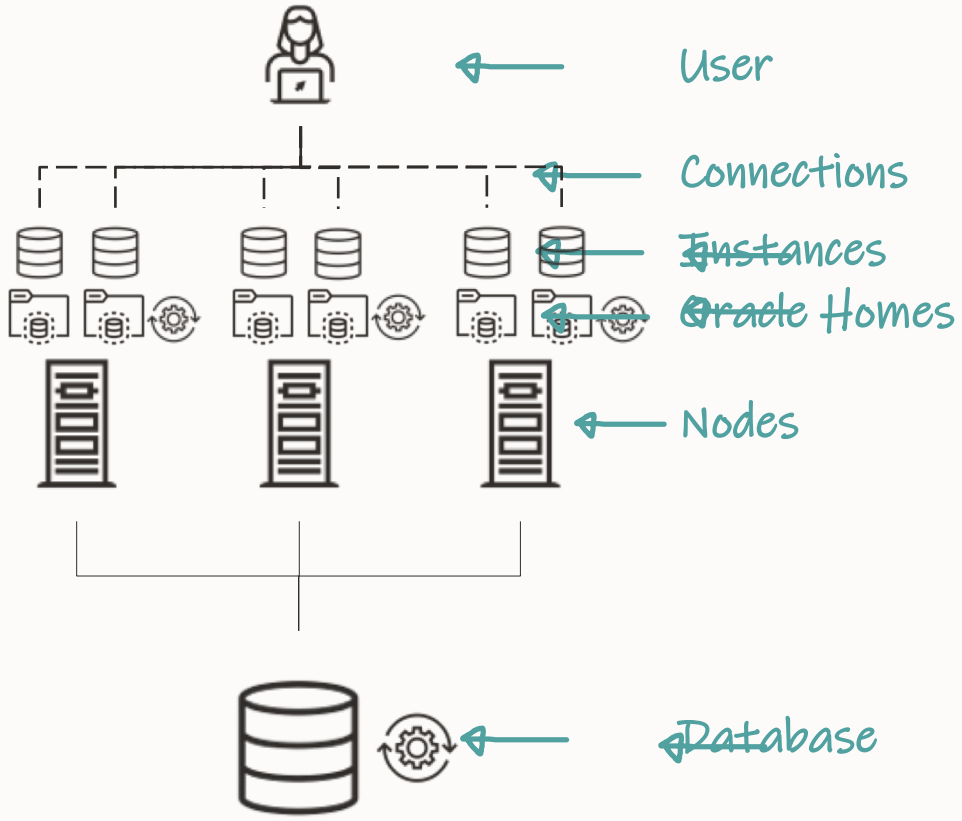

- New Oracle Home
- Patch Oracle Home
- Move to new Oracle Home
- Execute Datapatch

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

1 In-place

Replaces existing Oracle Home

Uses `opatchauto`

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 ... \
               ...
```

```
--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

23ai 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

- RAC: Frequently Asked Questions (Doc ID 220970.1)

Unusual combinations are supported, but we strongly advice 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

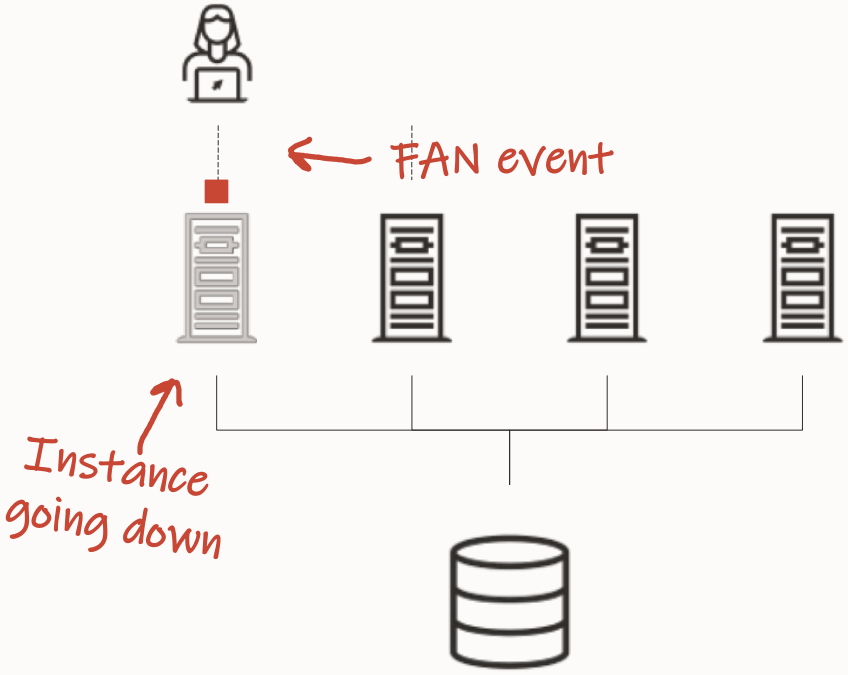

DRAINING

- 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 cautions 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 (Doc ID 2680521.1)

Analyze the database for patching readiness using Datapatch Sanity Checks

- Datapatch User Guide (Doc ID 2680521.1)
- Executed by AutoUpgrade in analyze mode

Recompile invalid objects
before invoking `datapatch`

Patching Timeline

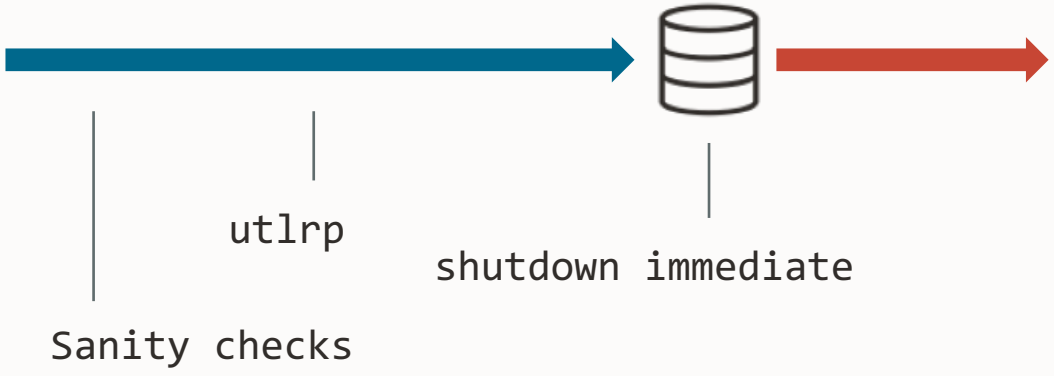

The database must be open
Only open PDBs are patched

- Upgrade mode or restricted session is not needed

Patching Timeline

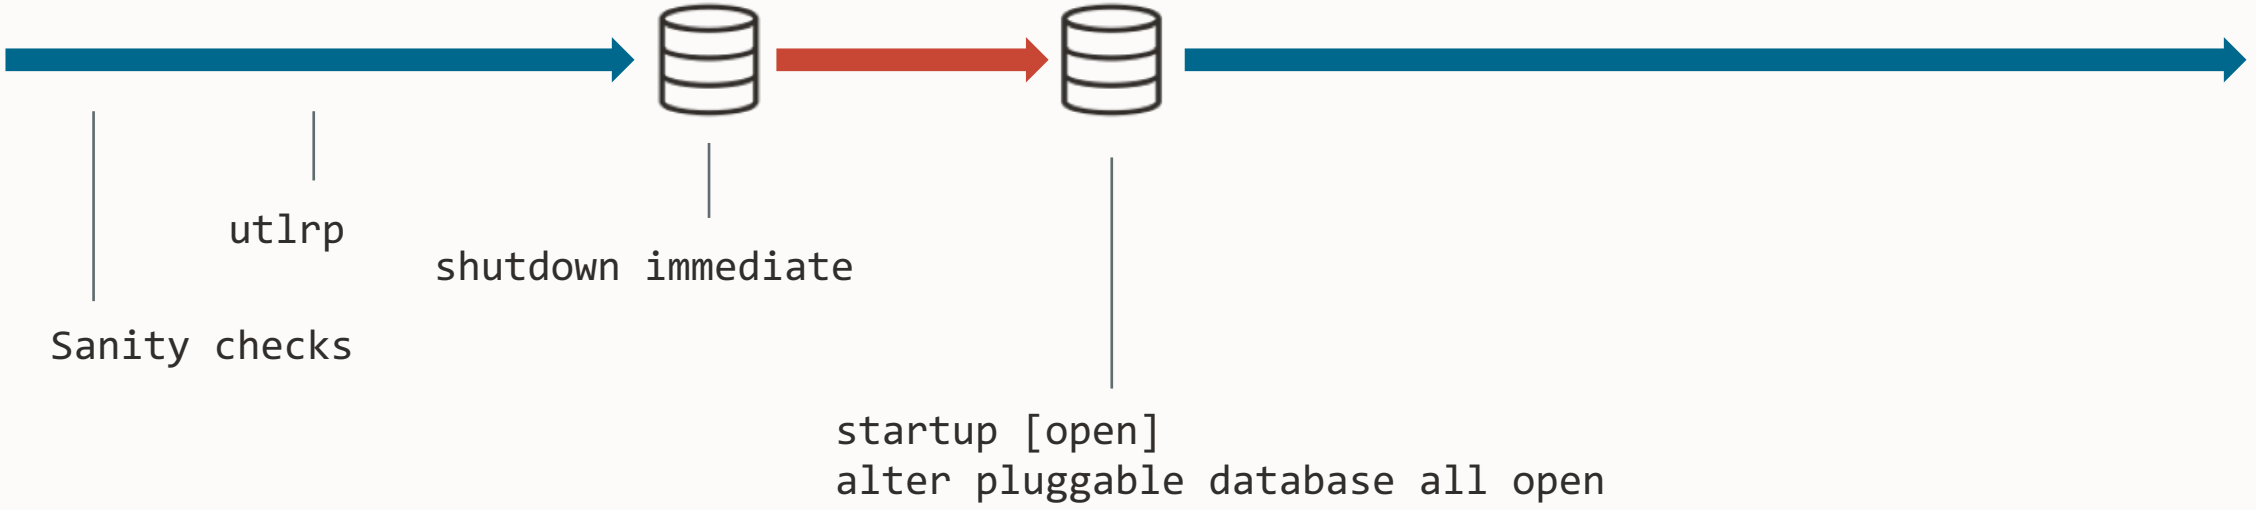

You can run `datapatch` while users are connected to the database

- Details in blog post

Patching Timeline

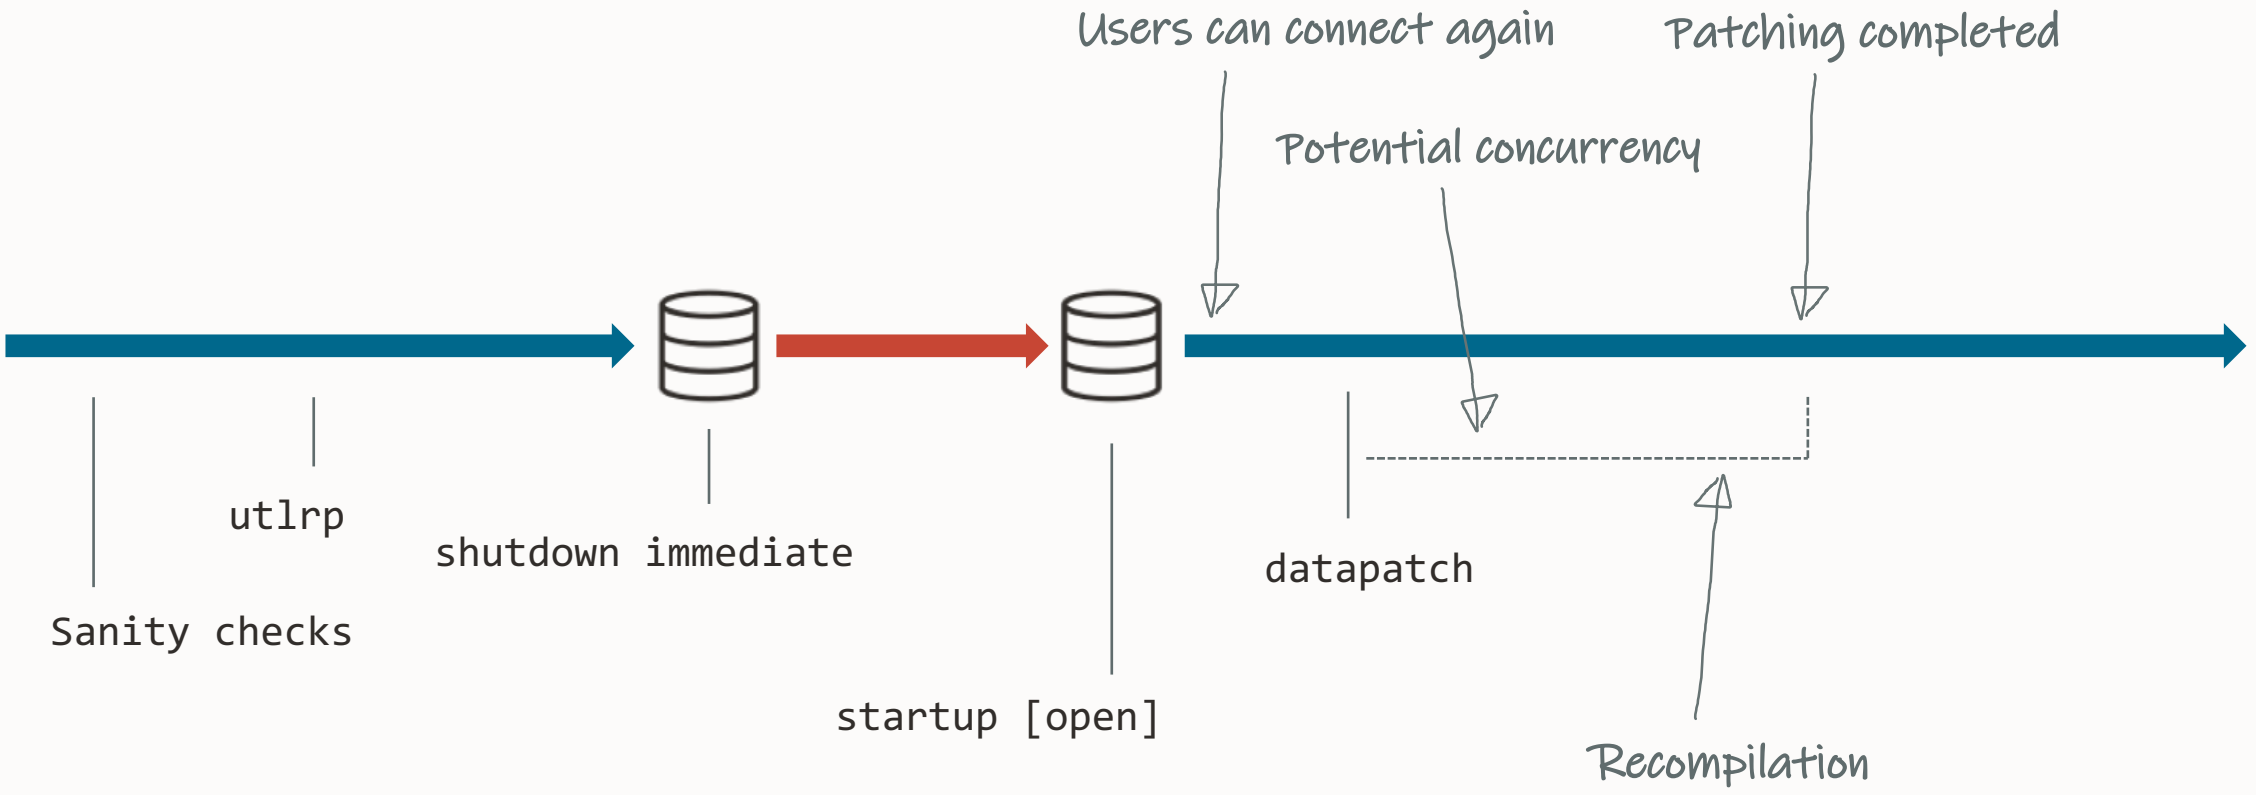

Patching Timeline

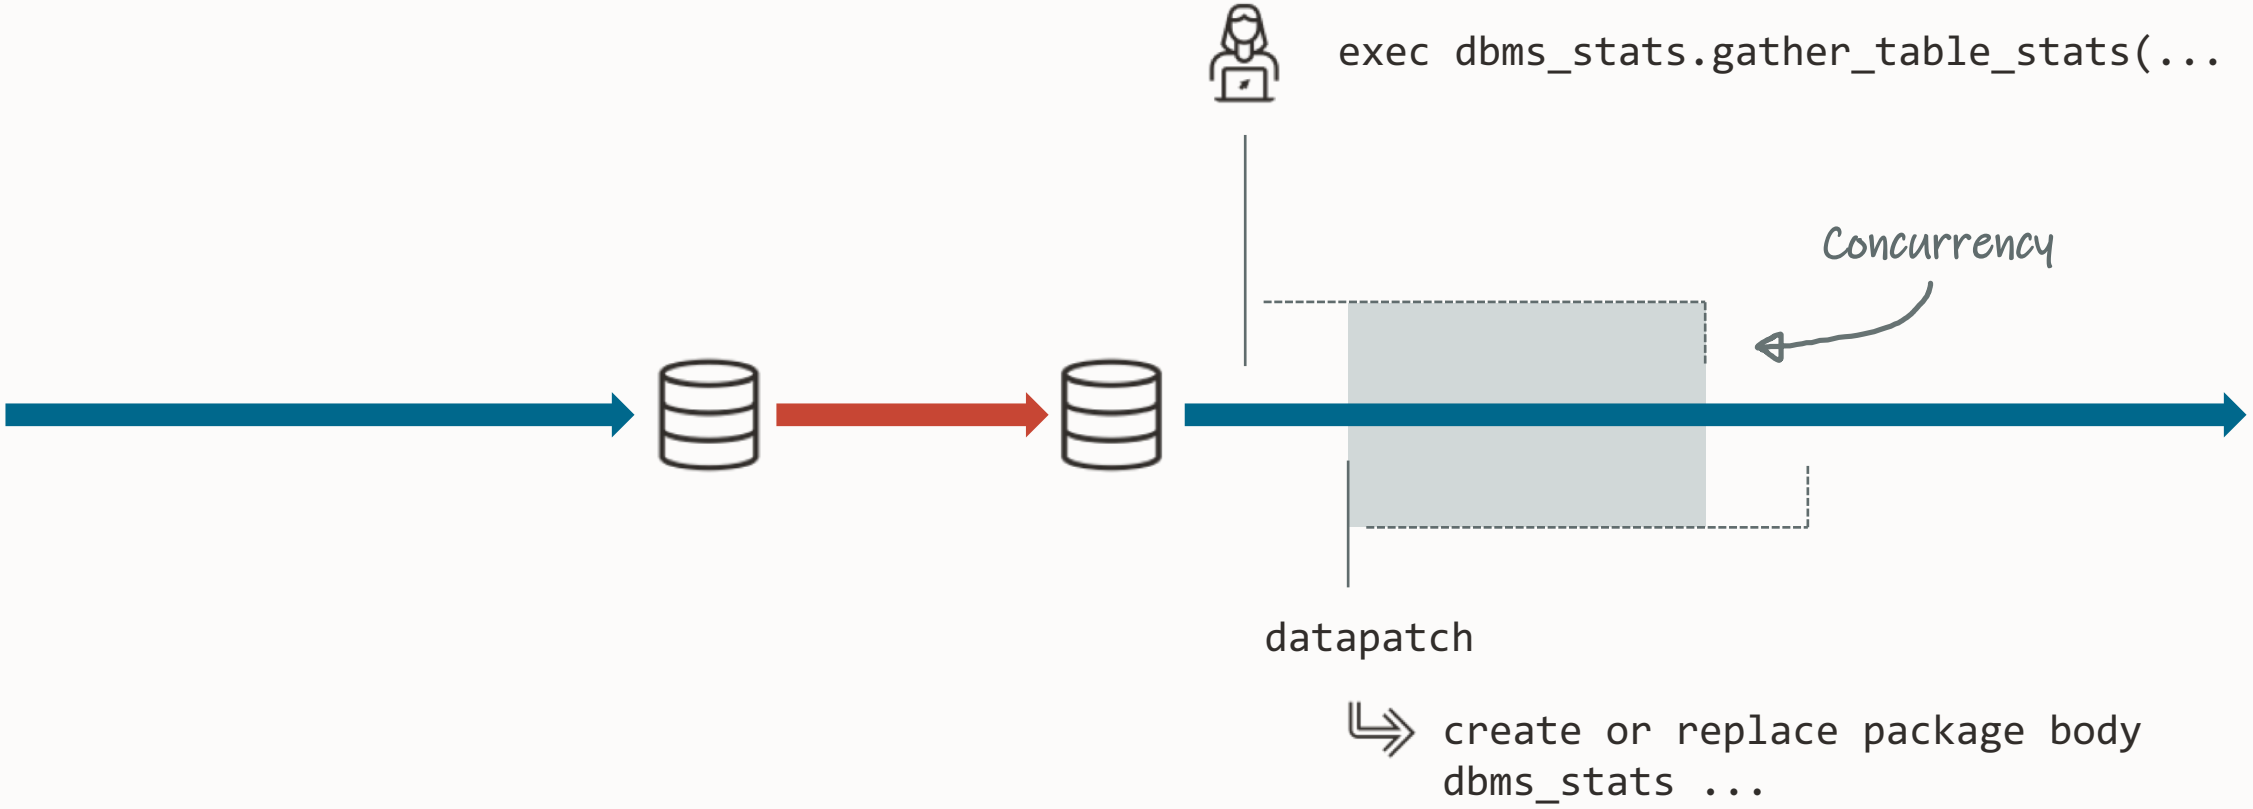

Concurrency

- 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

- Increase timeout using `-ddl_lock_timeout` *`<time-in-seconds>`*

Patching Timeline

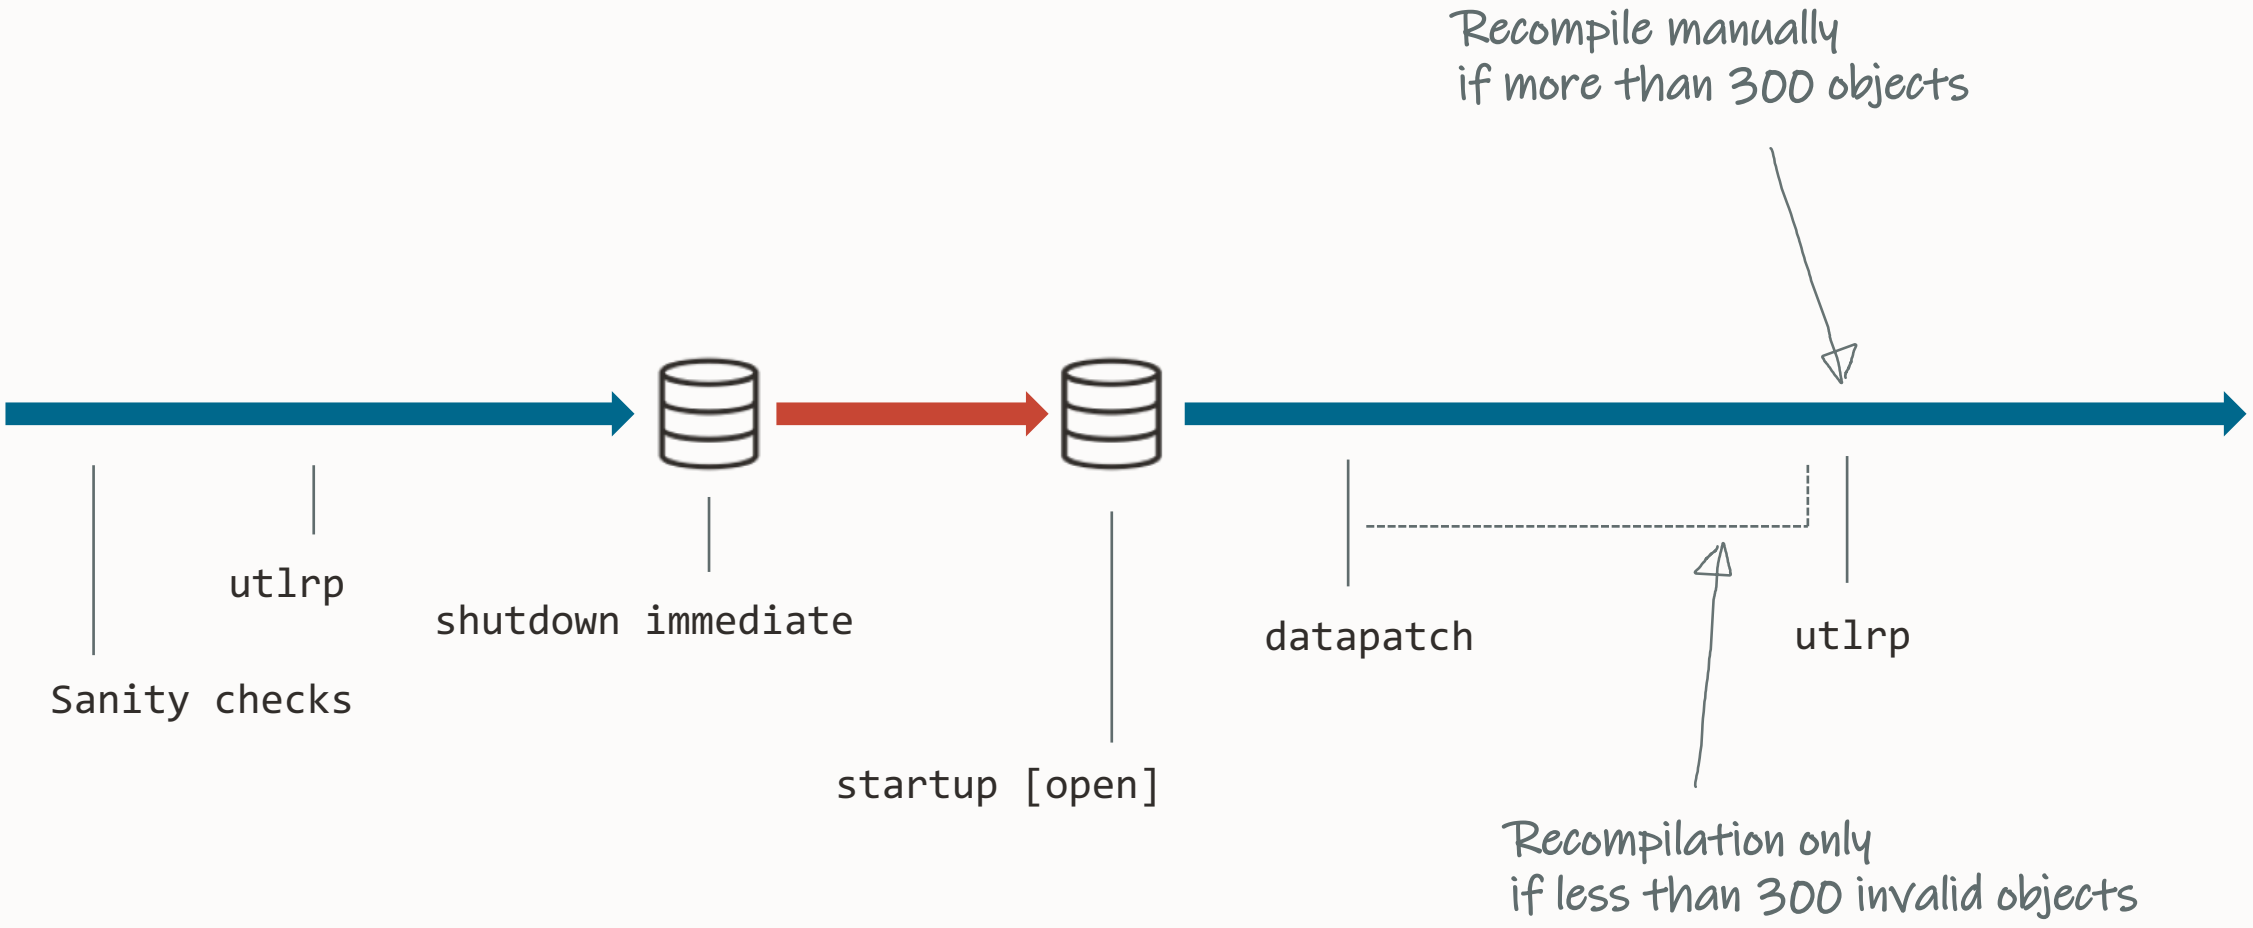

```
$ ./datapatch
SQL Patching tool version 19.19.0.0.0 Production on Sun Jun 25 07:12:19 2023

.
.
.
.

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 and Patching

1 

Download

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

2 

Install

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

3 

Patch

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

AutoUpgrade and Patching

1 

Download

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

2 

Install

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

3 

Patch

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

```
$ cat 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=RU,OPATCH,OJVM,DPBP
```

```
$ 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=RU,OPATCH,OJVM,DPBP

$ java -jar autoupgrade.jar -config DB19.cfg –patch -mode deploy
```


Demo

One-Button Patching

- Download patches
- Install Oracle home
- Patch database

Watch on YouTube

AutoUpgrade Patching

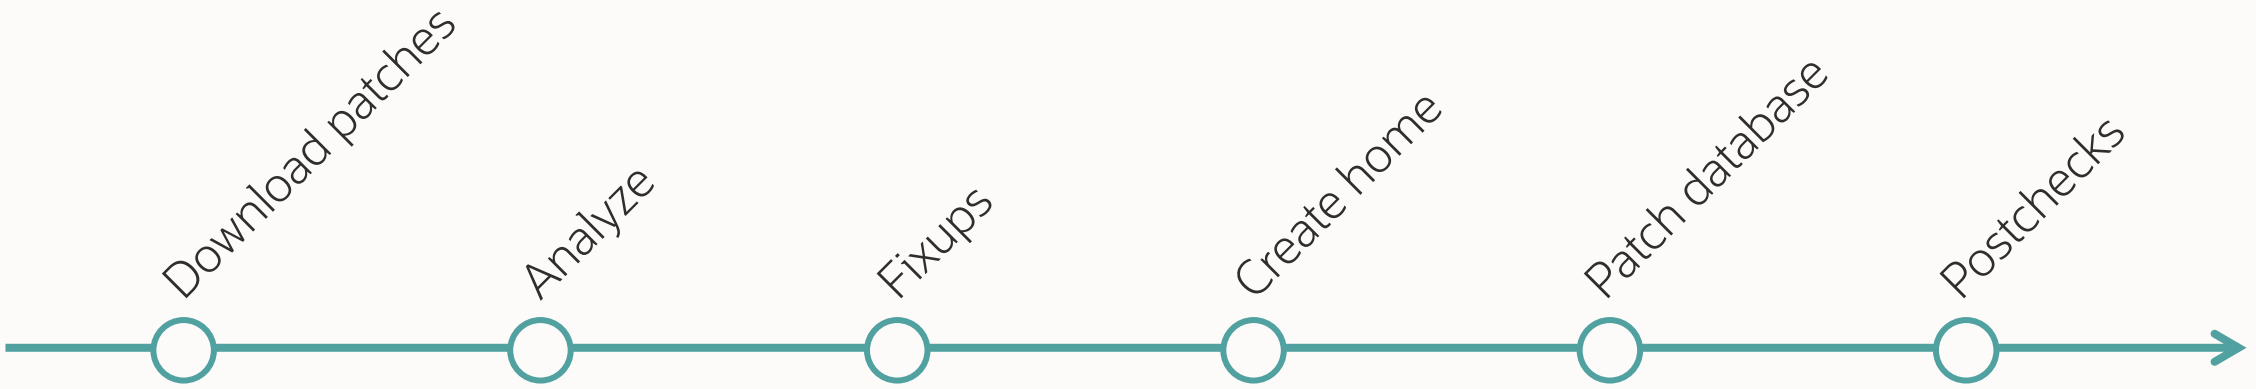


```
-mode deploy
```

AutoUpgrade Patching

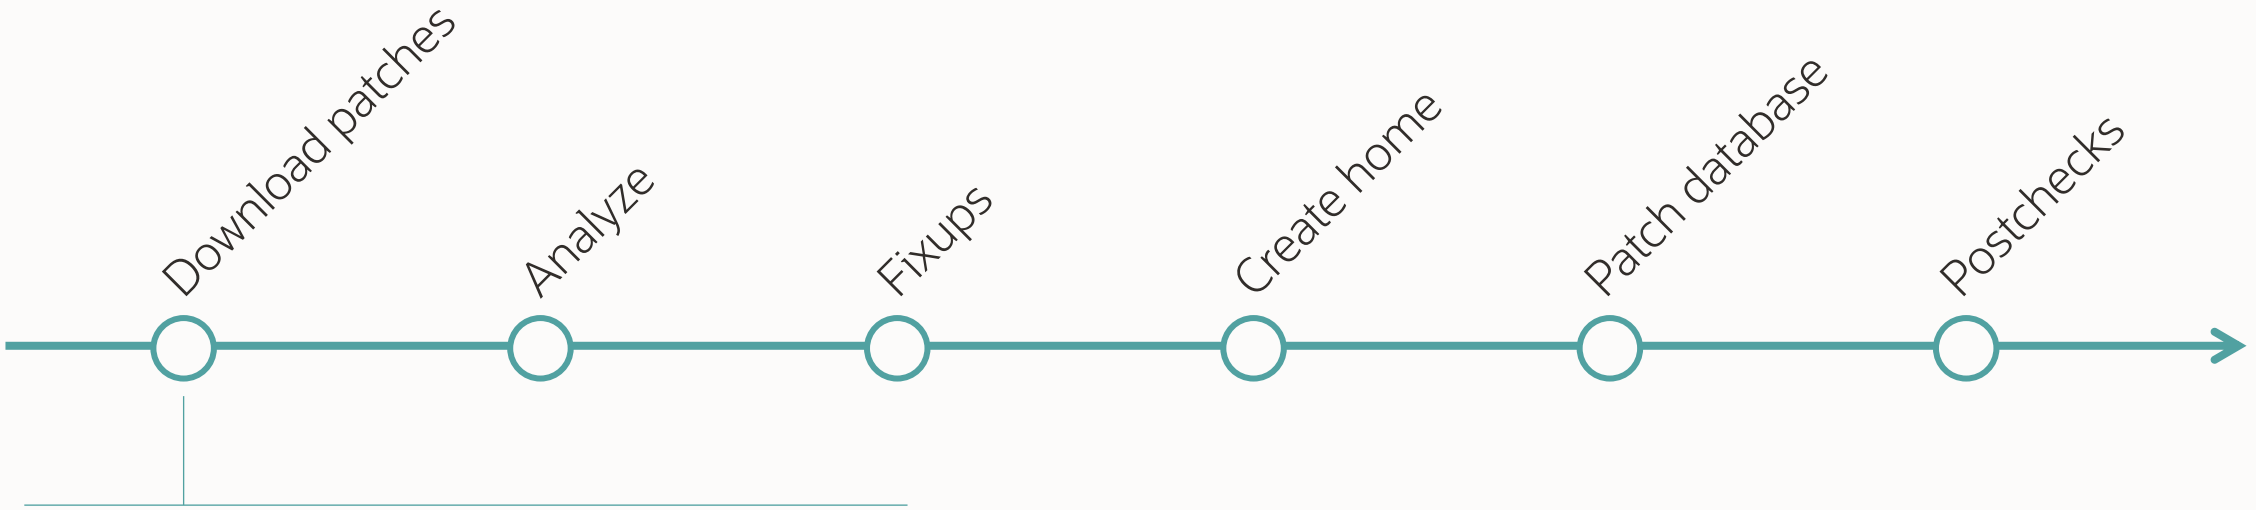

- Finds and downloads patches
- Uses your MOS credentials
- Chooses the right platform
- Stores patches in local repository

AutoUpgrade Patching

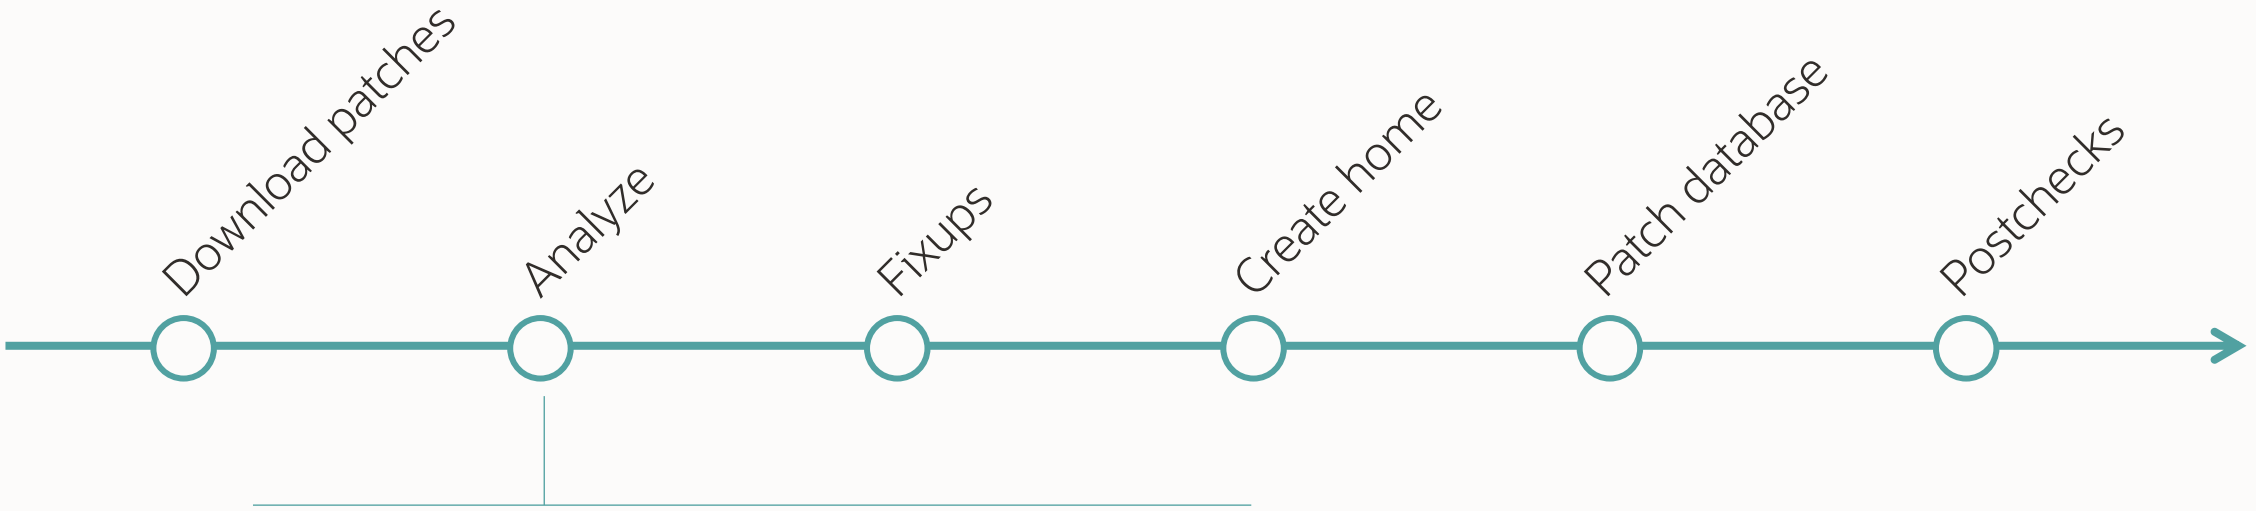

- Analyzes database for patch readiness
- Datapatch Sanity Checks
- Lightweight
- Non-intrusive
- Recommended, not required

AutoUpgrade Patching

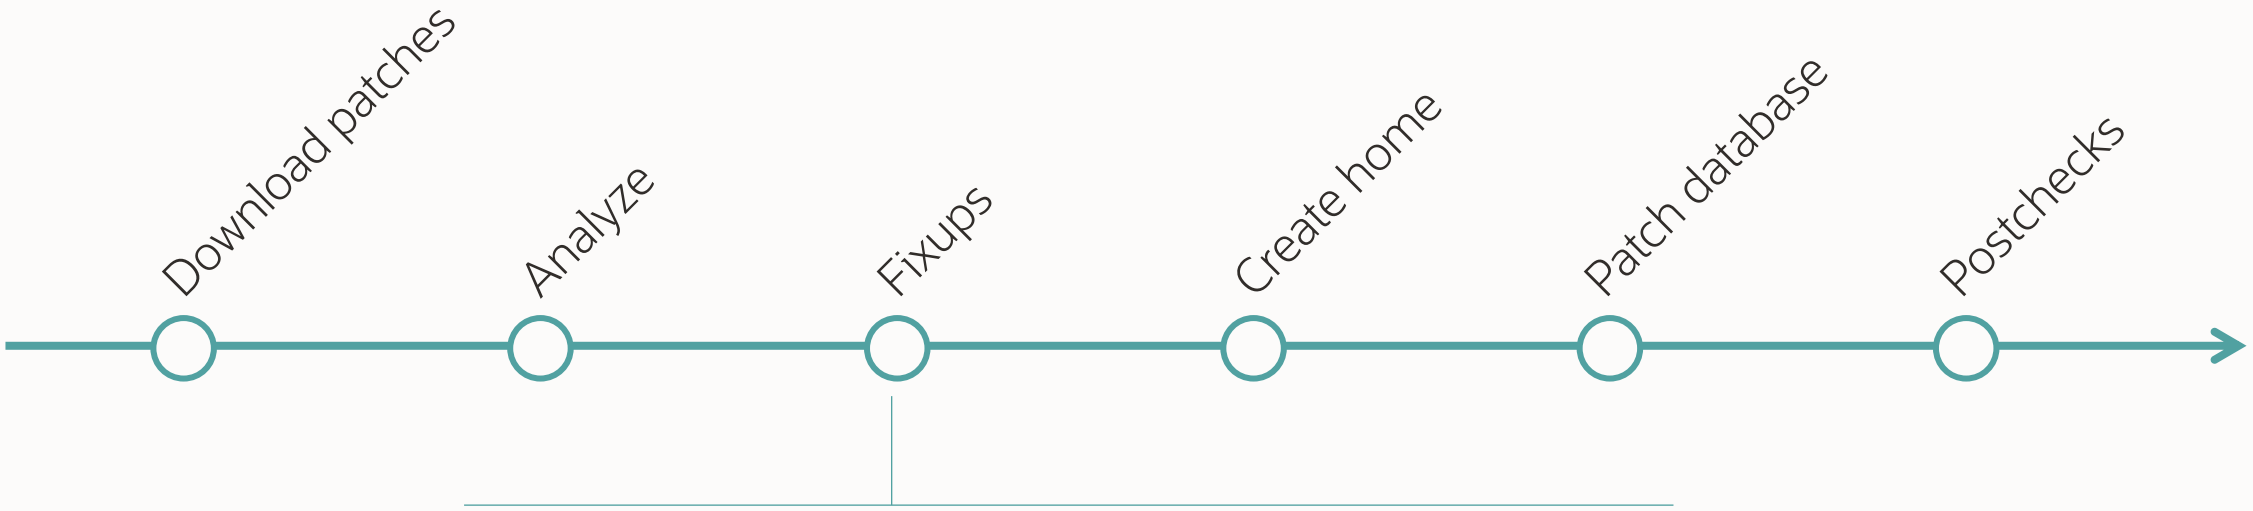

- Gathers dictionary statistics if needed
- Recompiles Oracle-maintained objects if needed
- Executes checks - see Doc ID 2380601.1

AutoUpgrade Patching

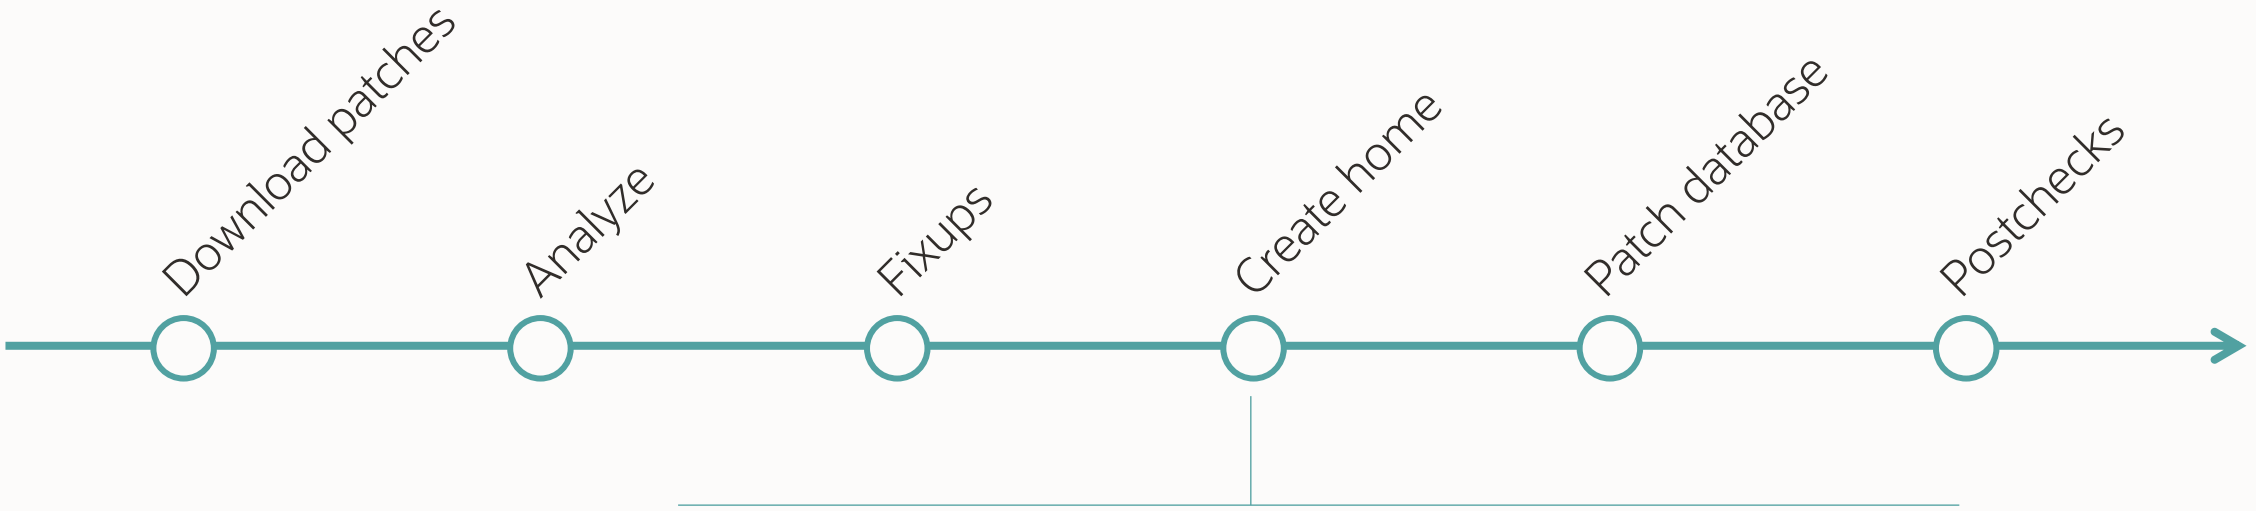

- Out-of-place patching
- Creates a brand-new Oracle home
- Uses *runInstaller settings* and binary options from source Oracle home
- Execute `root.sh` via sudo

AutoUpgrade Patching

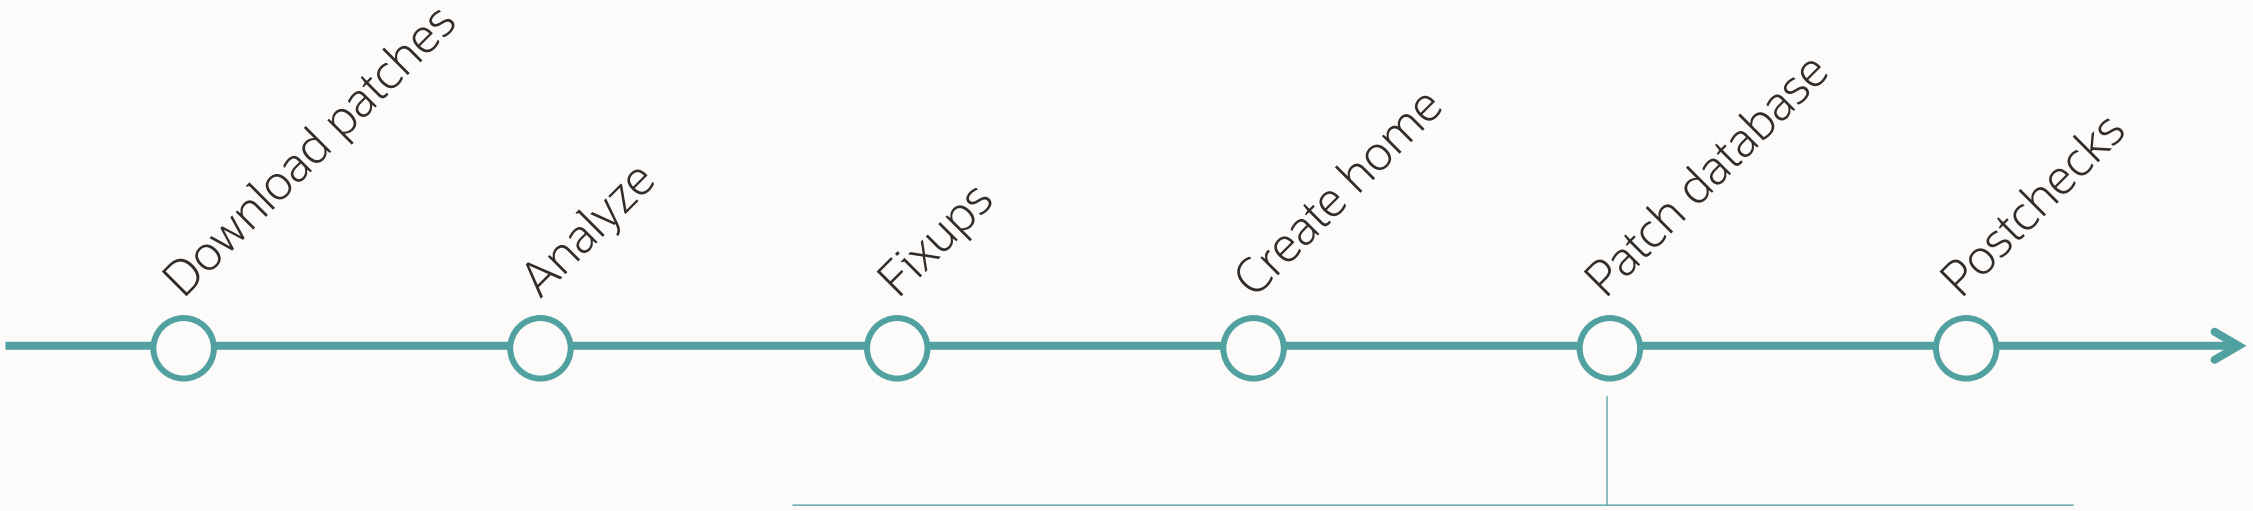

- Moves database instance to new Oracle home
- Moves configuration files
- Executes Datapatch
- Updates system directories (`utlfixdirs.sql`)
- Updates `/etc/oratab`
- Supports read-only Oracle home

AutoUpgrade Patching

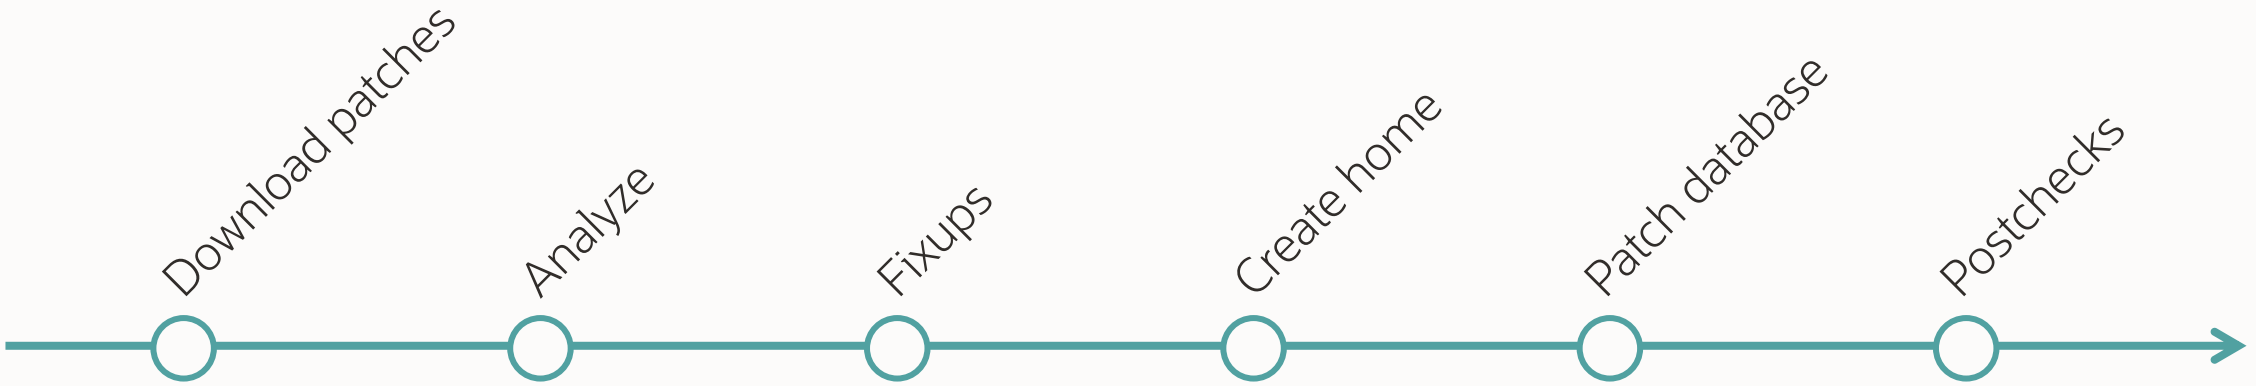

- Post-patching fixups
- Recompiles if needed
- Cleans up

Patching Modes

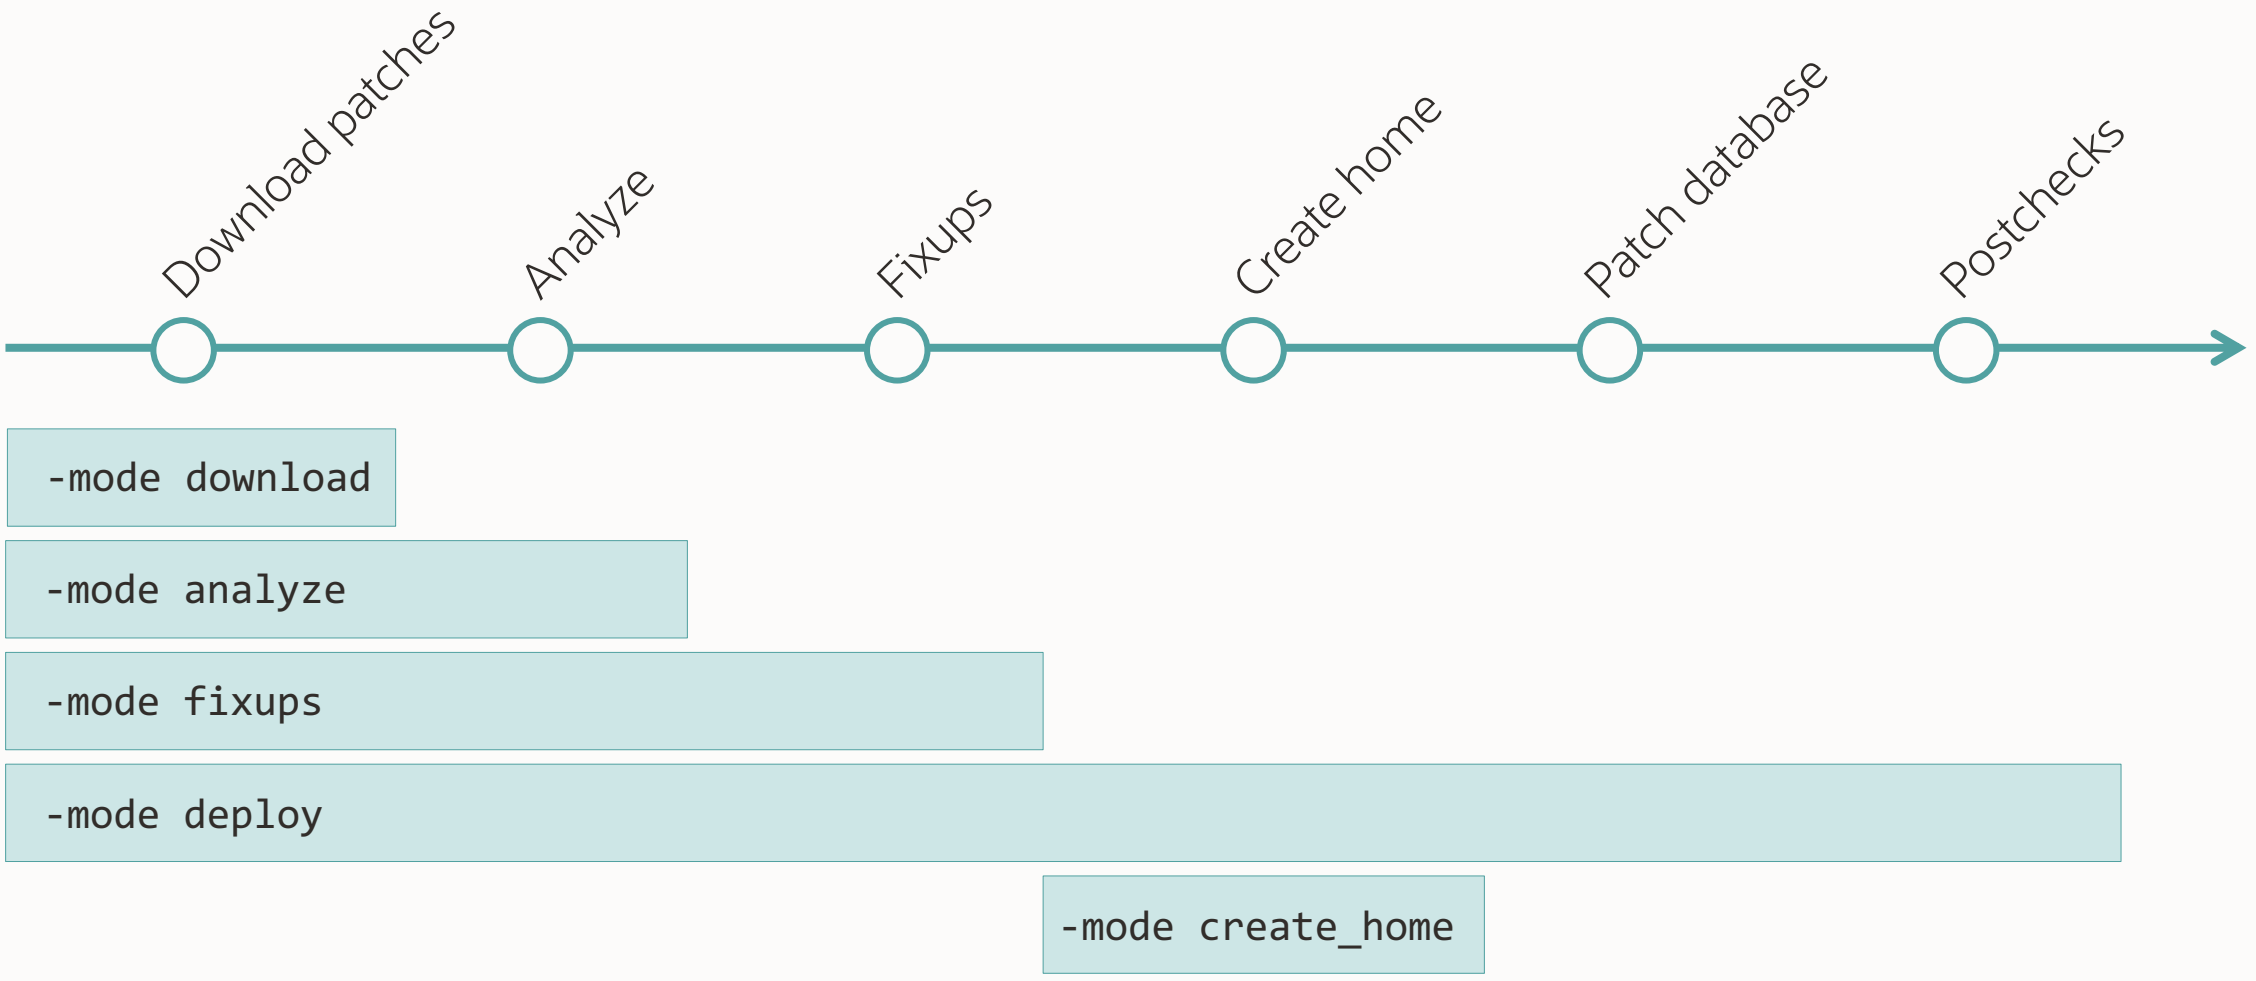

Which patches should you install?

```
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

Important MOS Notes

Download Assistant

MOS Note: 2118136.2

Download Release Update, OJVM, Bundles

Important One-Offs

MOS Note: 555.1

Add critical fixes on top for known issues

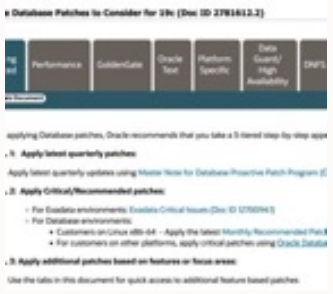

Additional Fixes

MOS Note: 2781612.2

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

SQL Performance

MOS Note: 2773715.1

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

```
$ 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
```

MOS Note: 555.1

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
```

Demo

Applying recommended patches

- Specifying patches
- Patch database

Watch on YouTube

Your database host doesn't have internet access?

Using Download Mode

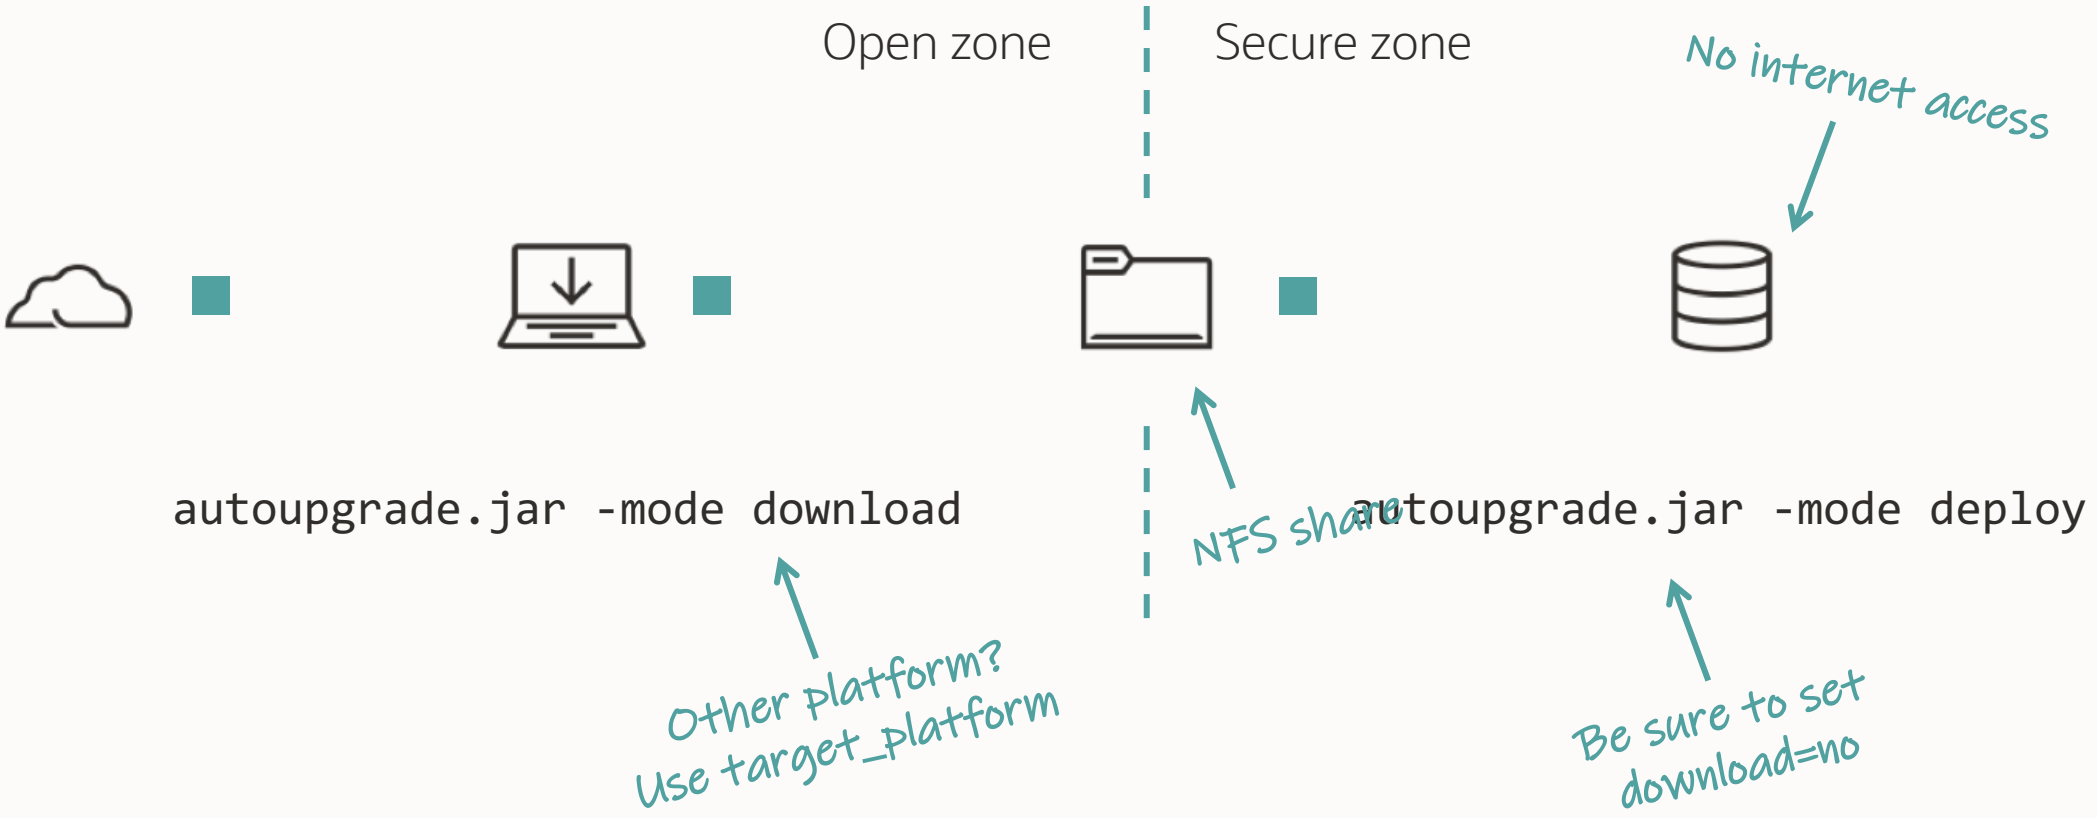

Without internet, AutoUpgrade applies the recent-most RU found in *patch* folder

- 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

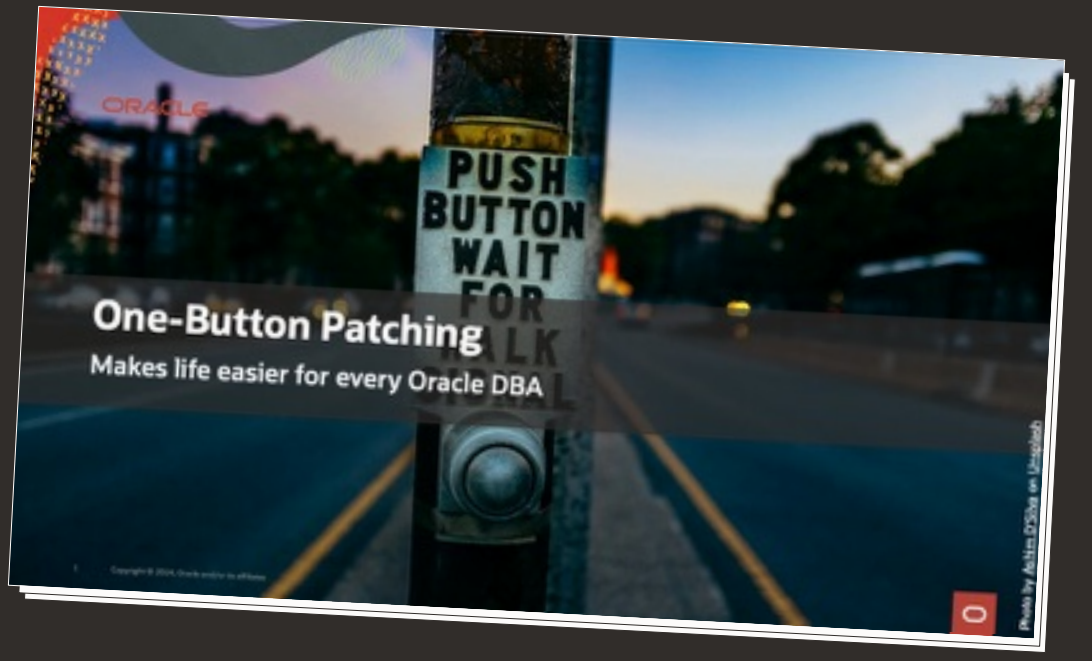

Recording on YouTube
Get the slides

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

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 23ai supports the multitenant architecture only

- You must convert your database to a PDB

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


```
--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_pdbs=3;
```


Update Exadata System Software to the latest version, minimum 24.1.6

- Exadata Database Machine and Exadata Storage Server Supported Versions (Doc ID 888828.1)

Upgrade to Oracle Grid Infrastructure 23ai

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

Raise ASM disk group compatibility

- Set `compatible.rdbms` to `19.0.0.0.0`

Ensure your clients can connect to Oracle Database 23ai

- Upgrade your clients well in advance of the upgrade

Client / Server Interoperability

Client Version	Server Version						
	23ai	21c	19c	18c	12.2.0	12.1.0	11.2.0
23ai[#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]

MOS Note: 2073031 - Client / Server Interoperability Support Matrix

Client / Server Interoperability

Client Version	Server Version						
	23ai	21c	19c	18c	12.2.0	12.1.0	11.2.0
23ai[#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]

MOS Note: 2073031 - Client / Server Interoperability Support Matrix

```
--List current connections and their driver details
--Join to gv$session for more details.
--https://dohdatabase.com/2024/03/19/are-your-oracle-database-clients-ready-for-the-next-database-upgrade/

select * from gv$session_connect_info;
```


On important databases, execute a dictionary check before upgrade

- Formerly known as *Health Check*
- MOS Doc ID 1366971

```
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 Database 23ai

Container Database

AutoUpgrade

Installation of Oracle Home is simpler

- Gold images with recent Release Update
- Available for Oracle Database 23ai

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 Database 23ai

AutoUpgrade

Create Container Database

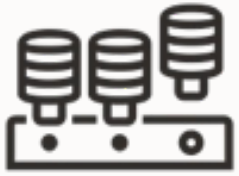

1 Character set

2 Components

3 `COMPATIBLE`

Create Container Database

1 Character set

- Always choose AL32UTF8
- Allows PDBs with any character set

2 Components

3 COMPATIBLE

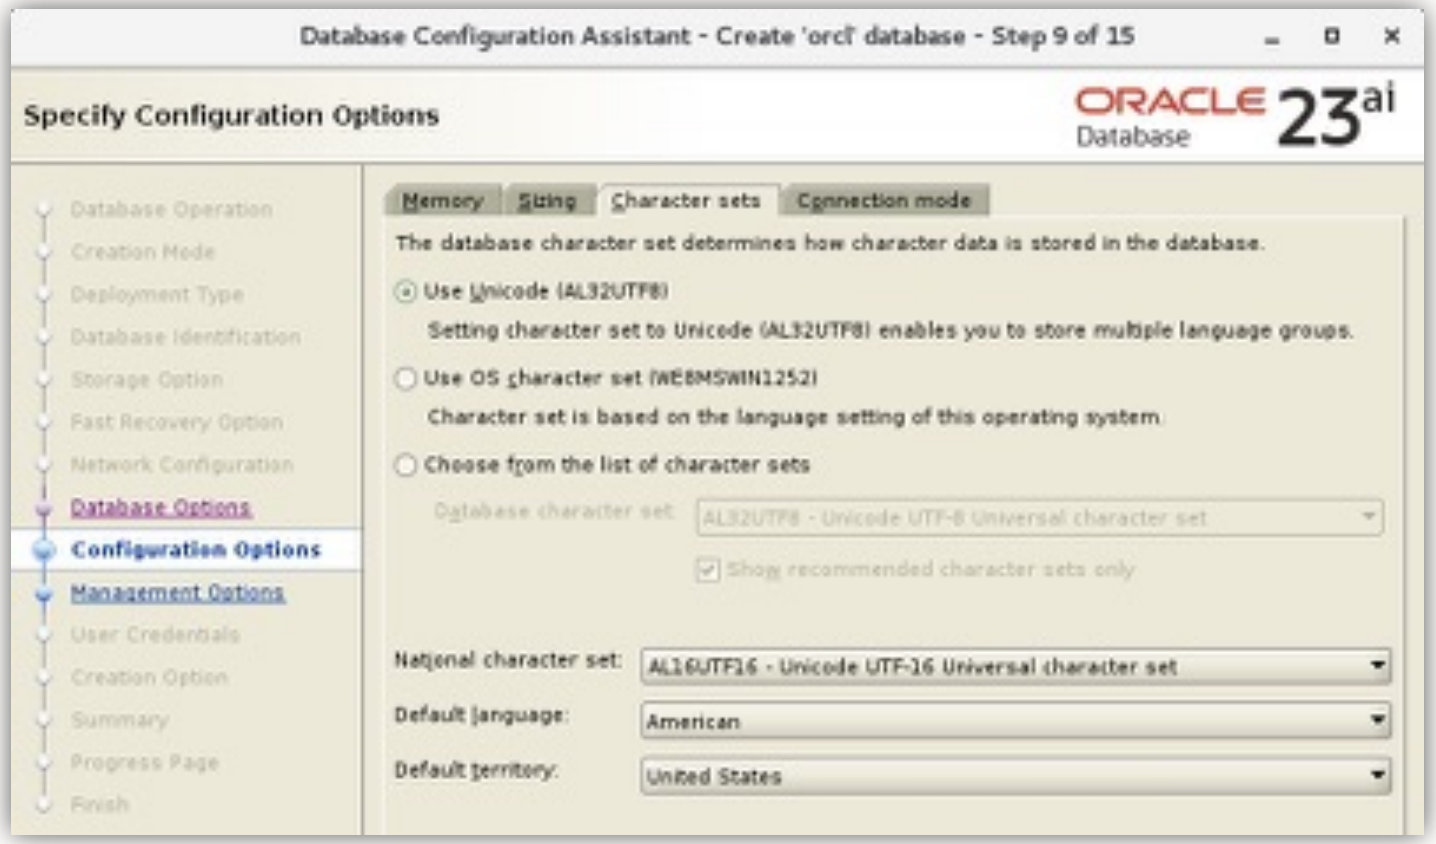

Create Container Database

1 Character set

2 Components

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

3 `COMPATIBLE`

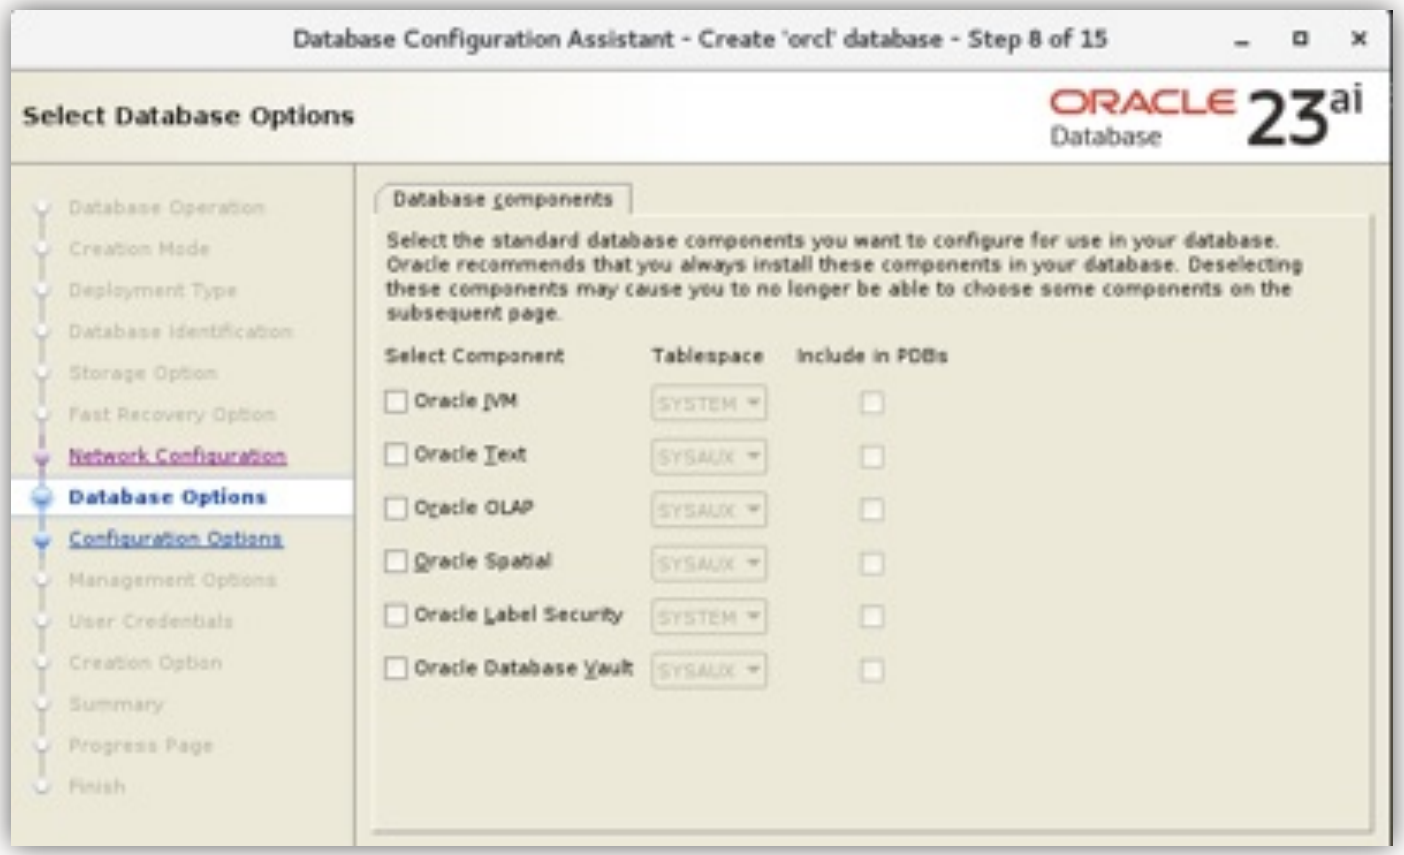

Create Container Database

1 Character set

2 Components

3 COMPATIBLE

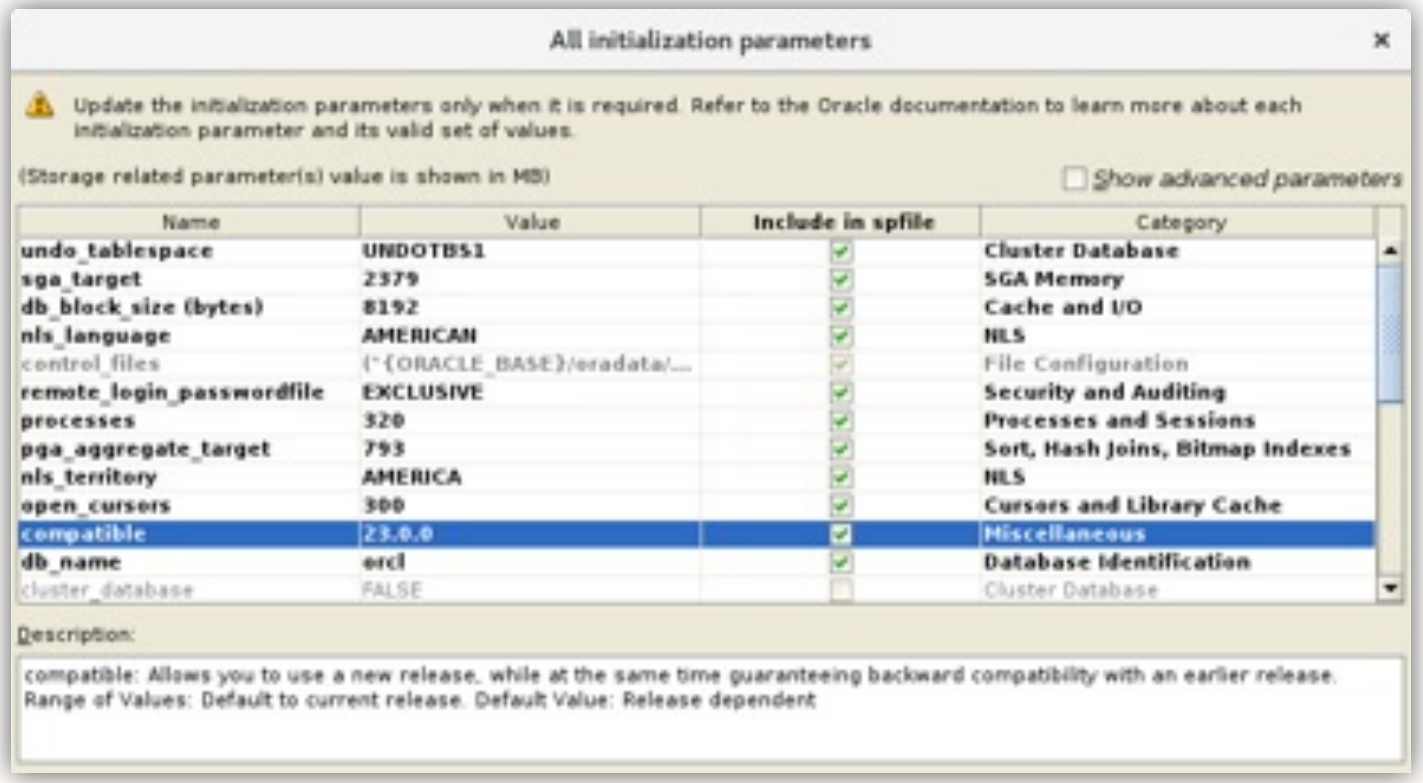

- Keep at the default setting, *23.0.0*
- Unless you want the option of downgrade

Compatible

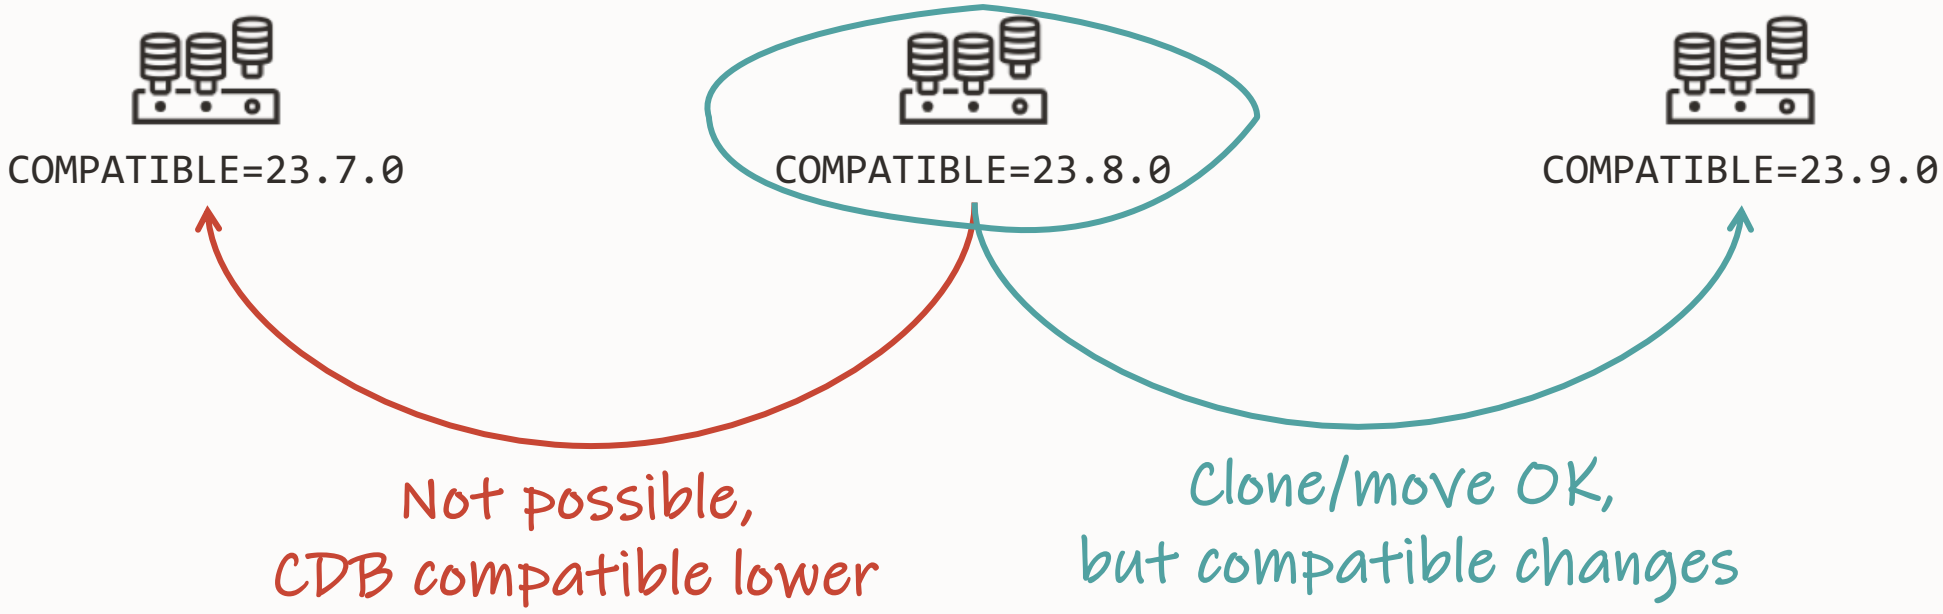


```
--Allows CDB views to include information on PDB$SEED objects.
--By default, such information is hidden.
--https://mikedietrichde.com/2017/07/21/why-exclude_seed_cdb_view-is-now-an-underscore-in-oracle-12-2/

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


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 25.4.250724
build.date 2025/07/24 18:03:59 +0000
build.hash 46b63c27c
build.hash_date 2025/07/22 19:14:35 +0000
build.supported_target_versions 12.2,18,19,21,23
build.type production
build.label (HEAD -> jira-aupg-4916, origin/jira-aupg-4916)
build.MOS_NOTE 2485457.1
build.MOS_LINK https://support.oracle.com/.../?id=2485457.1
```


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

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


Flow

1 Plug in

2 Upgrade

3 Convert

⇒ 23ai

Demo

Upgrade to Oracle Database 23ai

- Using AutoUpgrade
- Including PDB conversion

Watch on YouTube

```
SQL> select version_full from v$instance;

VERSION_FULL
------------
23.5.0.24.07
```


Non-CDB to PDB conversion is irreversible

What are your rollback options?

ROLLBACK

Backup / restore

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

Copy data files

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

ROLLBACK

Backup / restore

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

Copy data files

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

Refreshable clone PDB

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

Requires Oracle Database 12.2 or newer

Refreshable Clone PDB

CREATE

Create PDB from non-CDB over a database link

REFRESH

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

OUTAGE

Disconnect users and refresh PDB for the last time

CONVERT

To become a proper PDB, it must be converted

Refreshable Clone PDB

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

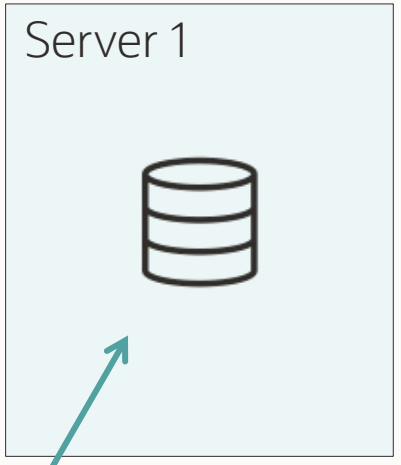

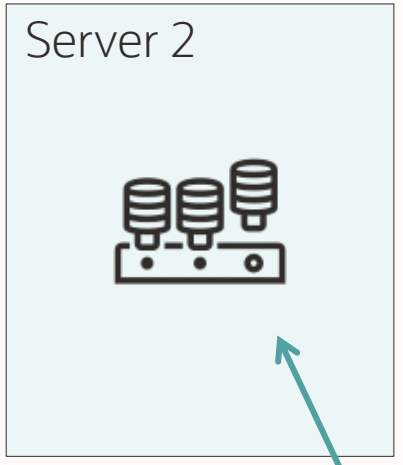

Source non-CDB

Target CDB

Refreshable Clone PDB

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

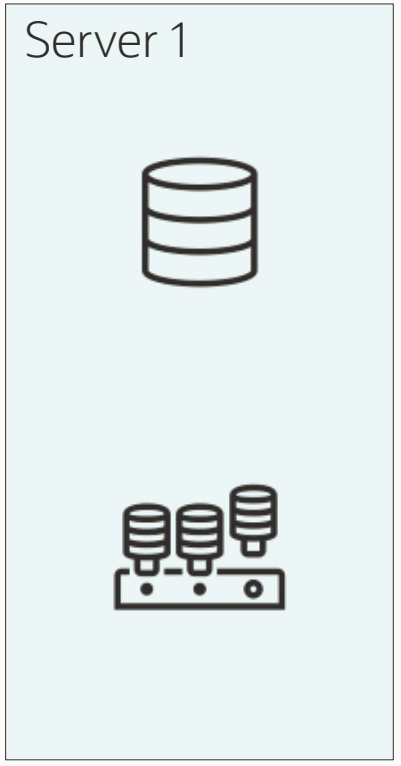

Could be same server as well

Refreshable Clone

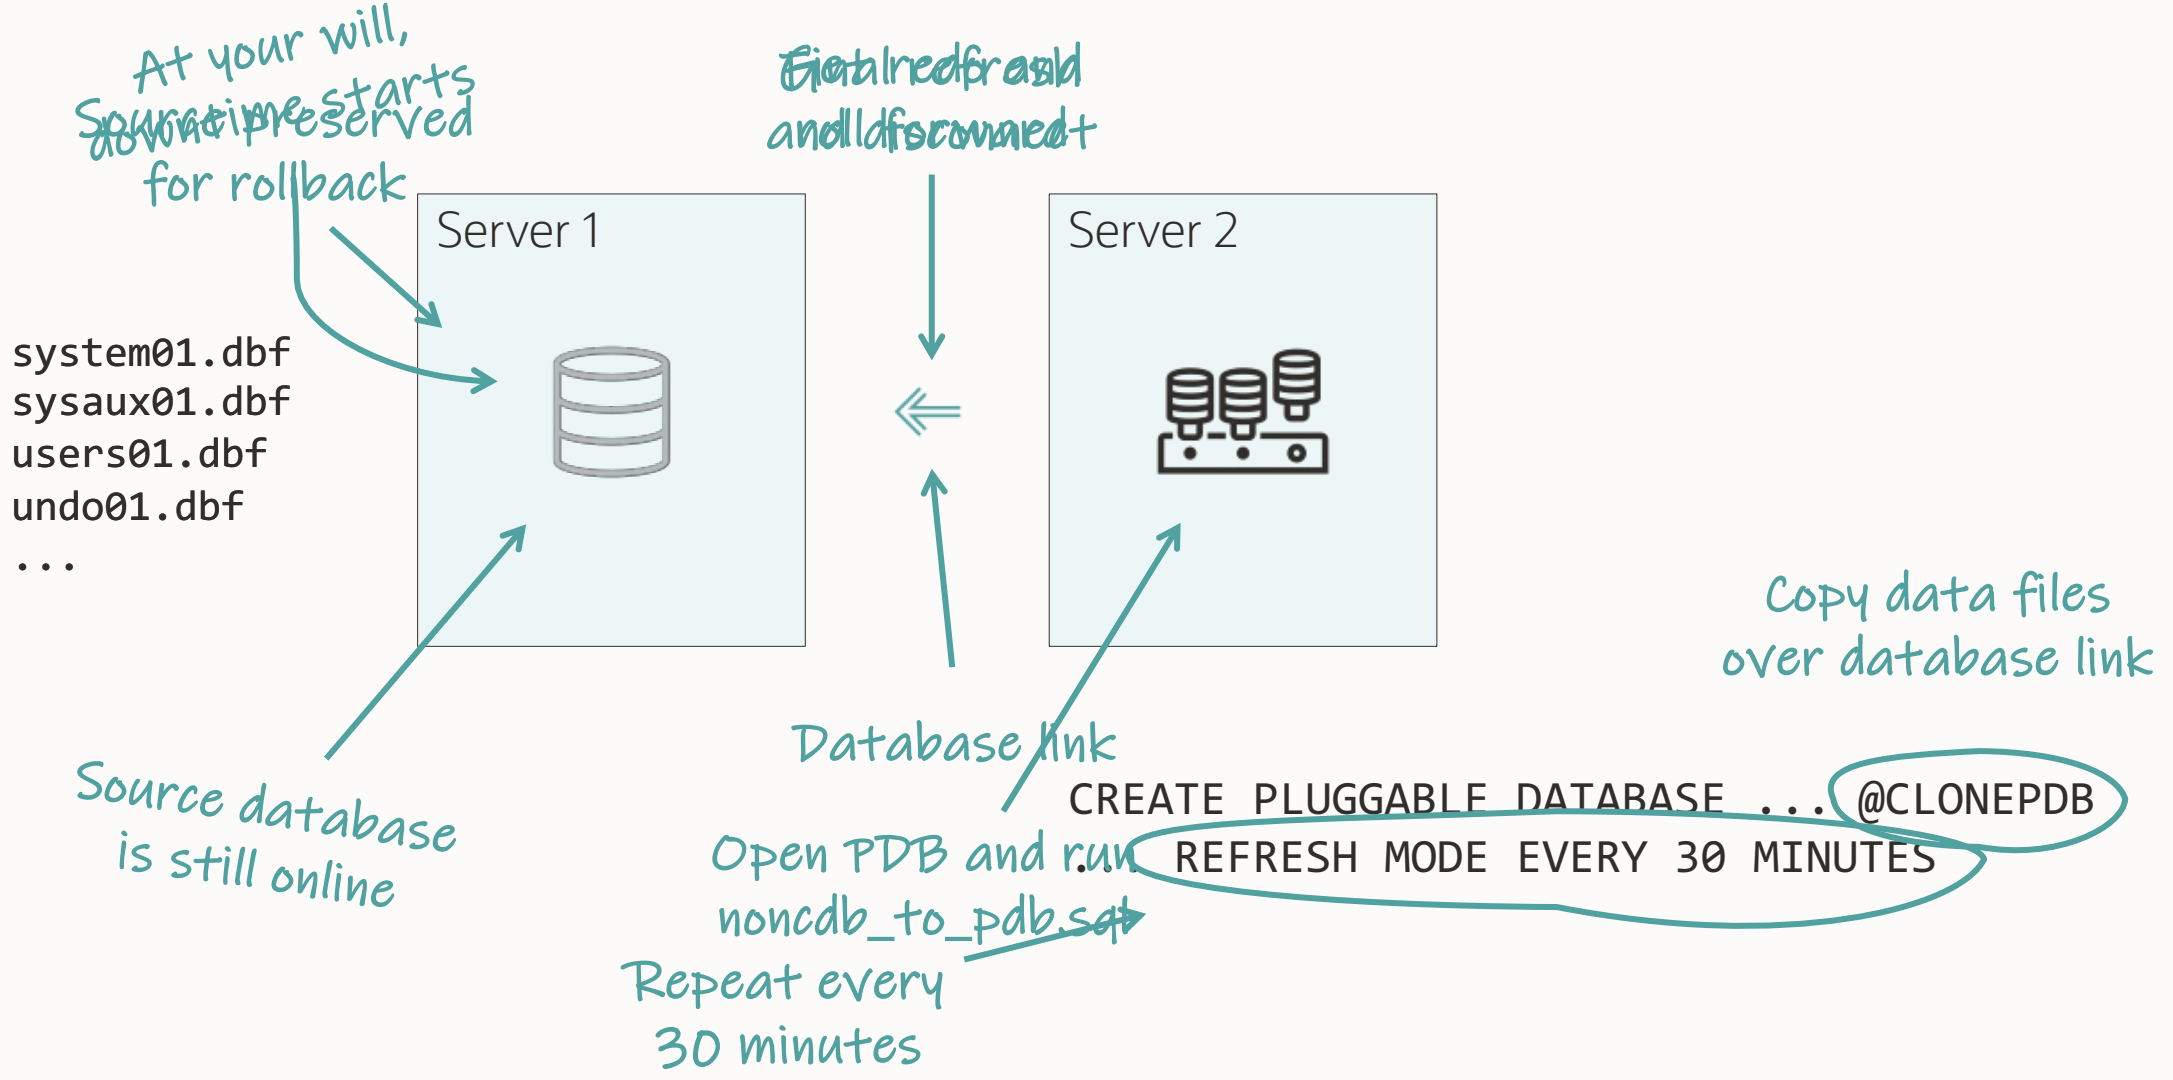

Refreshable Clone PDB

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


```
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

Refreshable clone works only with deferred recovery on standby database

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

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


Zürcher
Kantonalbank

Customer Case | Zürcher Kantonalbank

Customer

Project

Constraints

Preparation

Migration

Success?

Remarks

A reliable partner for over 150 years

- The bank for the people of Zurich since 1870
- With over 5’100 employees one of the largest employers in the canton of Zurich
- Globally networked full-service bank with strong regional and local roots

Customer Case | Zürcher Kantonalbank

Customer

Project

Constraints

Preparation

Migration

Success?

Remarks

Current situation

- Oracle databases on old OS and on Oracle Exadata
- 2023:
  - Migrate everything to Exadata until end of 2023
  - Consolidation to Multitenant and to the next long-term support release

Planned solution: AutoUpgrade

Customer Case | Zürcher Kantonalbank

Customer

Project

Constraints

Preparation

Migration

Success?

Remarks

Test setup

- 3 non-CDB databases of different size

Source	Size / GB
TEST40 (108)	165
TEST42 (107)	555
TEST41 (106)	18’496

- Exadata X6-2 compute node
- 7 storage cells (2x X6-2L / 3x X7-2L / 2x X8-2L)
- Oracle Database 19.15.0
- No additional options

Customer Case | Zürcher Kantonalbank

Customer

Project

Constraints

Preparation

Migration

Success?

Remarks

Cloning user

```
create user dblinkuser identified by Oracle_4UOracle_4U;
```

Permissions

```
grant CONNECT, RESOURCE, CREATE PLUGGABLE DATABASE,
      SELECT_CATALOG_ROLE to dblinkuser;
grant ALL ON SYS.ENC$ to dblinkuser;
```

Database link

```
create database link TEST42.DOMAIN connect to dblinkuser
identified by oracle_4uoracle_4u using ‘test42.domain’;
```

Customer Case | Zürcher Kantonalbank

- Customer
- Project
- Constraints
- Preparation
- **Migration**
- Success?
- Remarks

Migration in progress

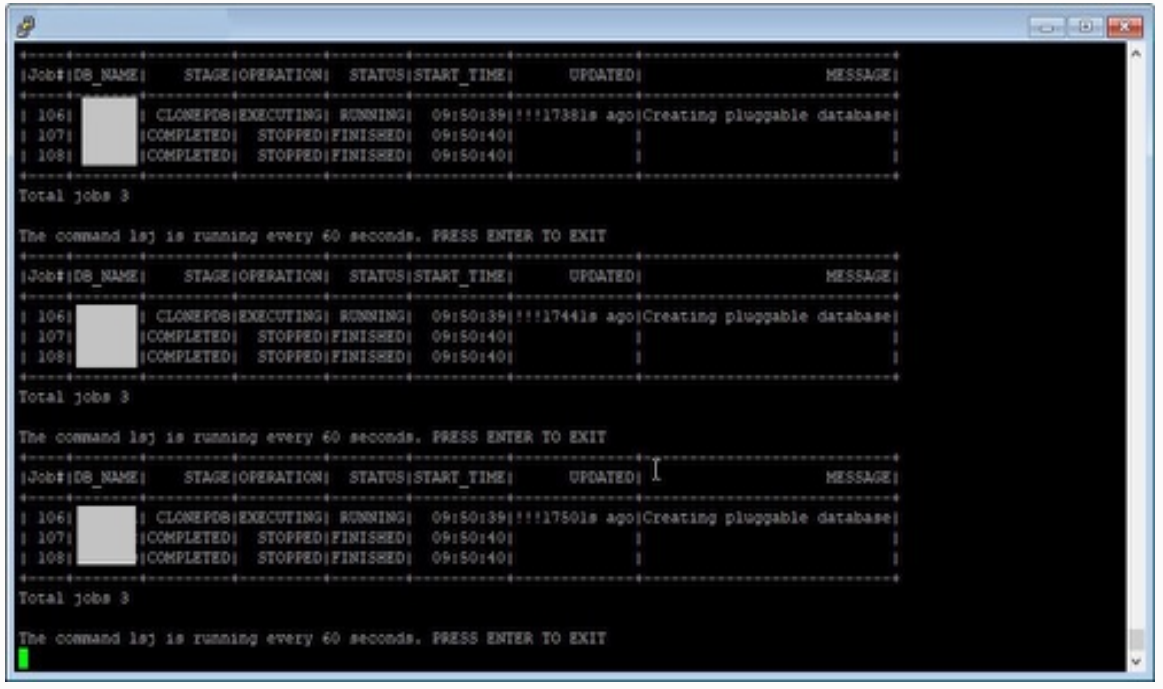

Source	Runtime/Min
TEST40 (108)	26
TEST42 (107)	ongoing
TEST41 (106)	ongoing

Customer Case | Zürcher Kantonalbank

Customer

Project

Constraints

Preparation

Migration

Success?

Remarks

Migration completed

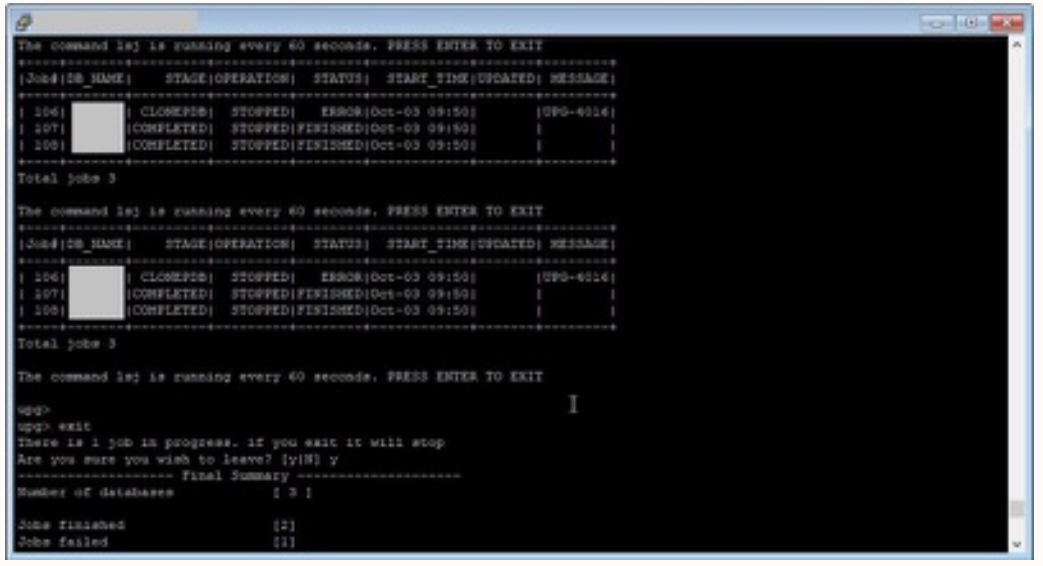

Source	Runtime/Min
TEST40 (108)	26
TEST42 (107)	226 (~3.5h)
TEST41 (106)	1770 (29h)

Customer Case | Zürcher Kantonalbank

Customer

Project

Constraints

Preparation

Migration

Success?

Remarks

First non-CDBs migrated successfully

- Project is ongoing

Customer Case | Zürcher Kantonalbank

Customer

Project

Constraints

Preparation

Migration

Success?

Remarks

For large databases, make sure archives aren't cleaned up

- Solution: restore archivelogs from backup

User profile with IDLE_TIME lead to kill of the session

- Solution: assign a different profile to the clone user

Summary

- Very comfortable to use
  - Everything happens automatically
  - Does not require user interaction
- Simple syntax
- No license costs associated
- Perfect for pre-migration test

- Very Stable

Christian Marquardt
Database Specialist Expert & Solution Architect

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 23ai 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

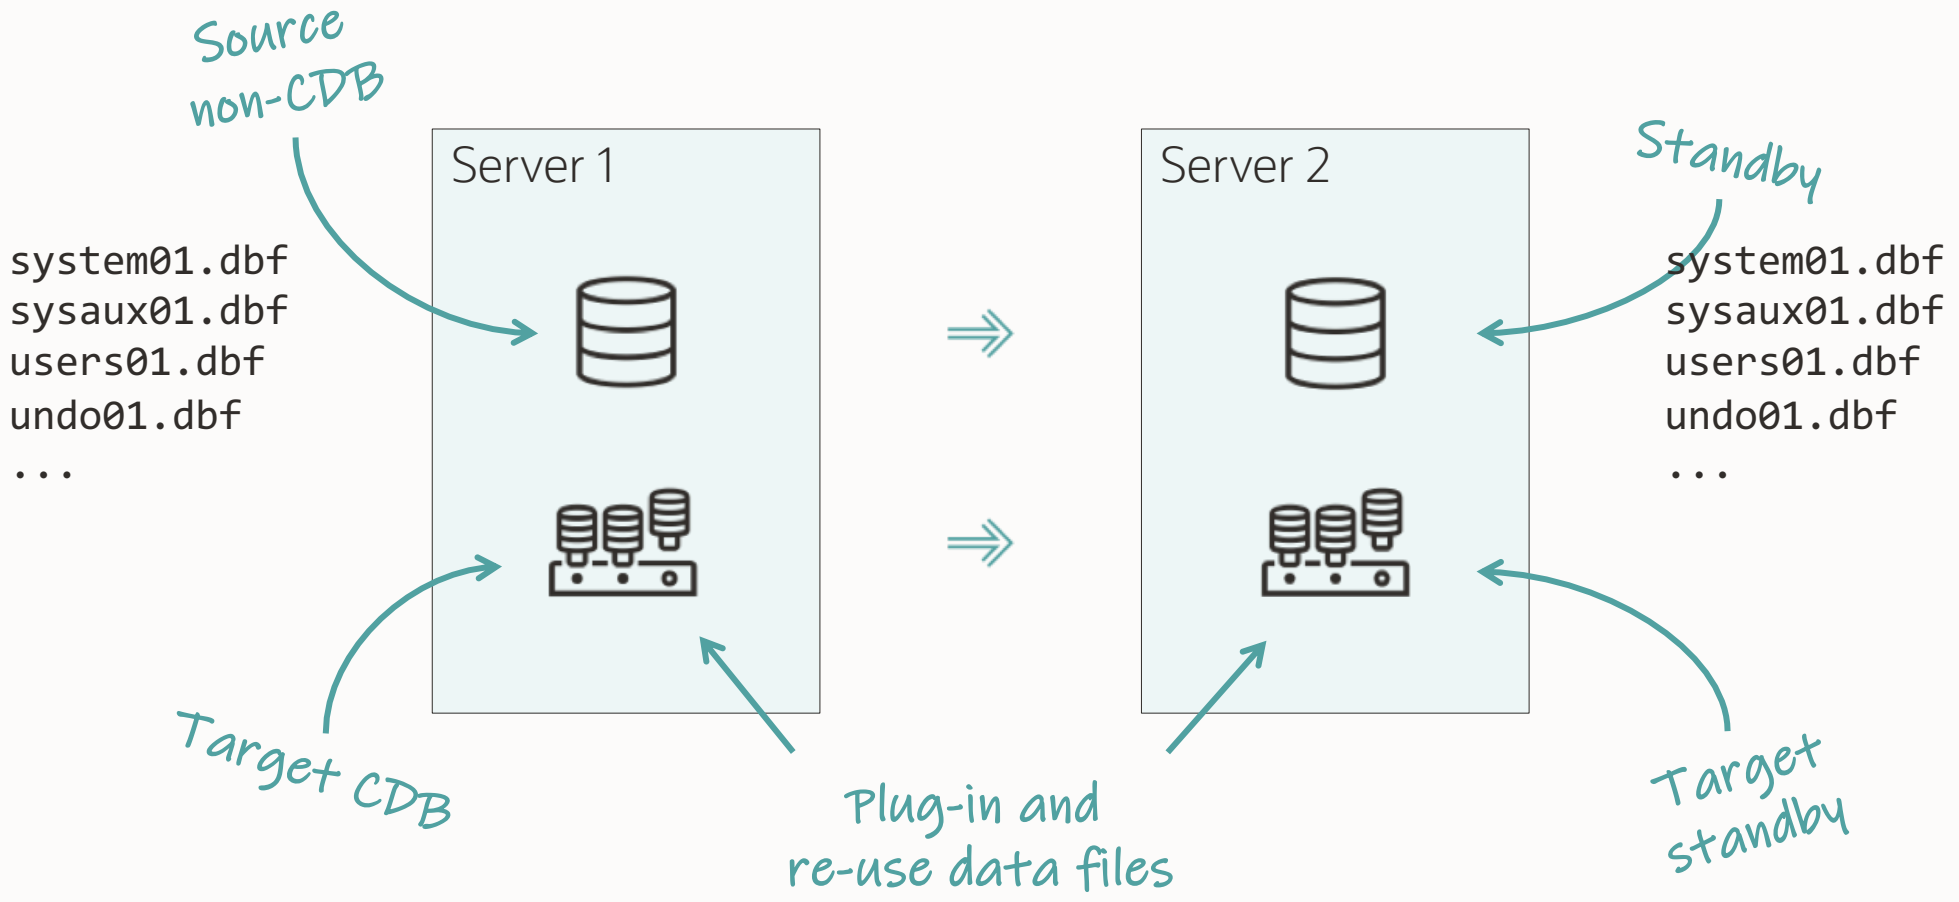

All data files on primary and standby must be at the same SCN

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

Enabled Recovery

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

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

Enabled Recovery

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

1 Regular files

2 OMF in regular file system

3 ASM

Enabled Recovery

1 Regular files

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

Enabled Recovery

2 OMF in regular file system

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

Enabled Recovery

3 ASM

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

Enabled Recovery | ASM

Primary

Standby

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

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

```
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

- File path on primary database only
- Not standby database

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

```
<?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

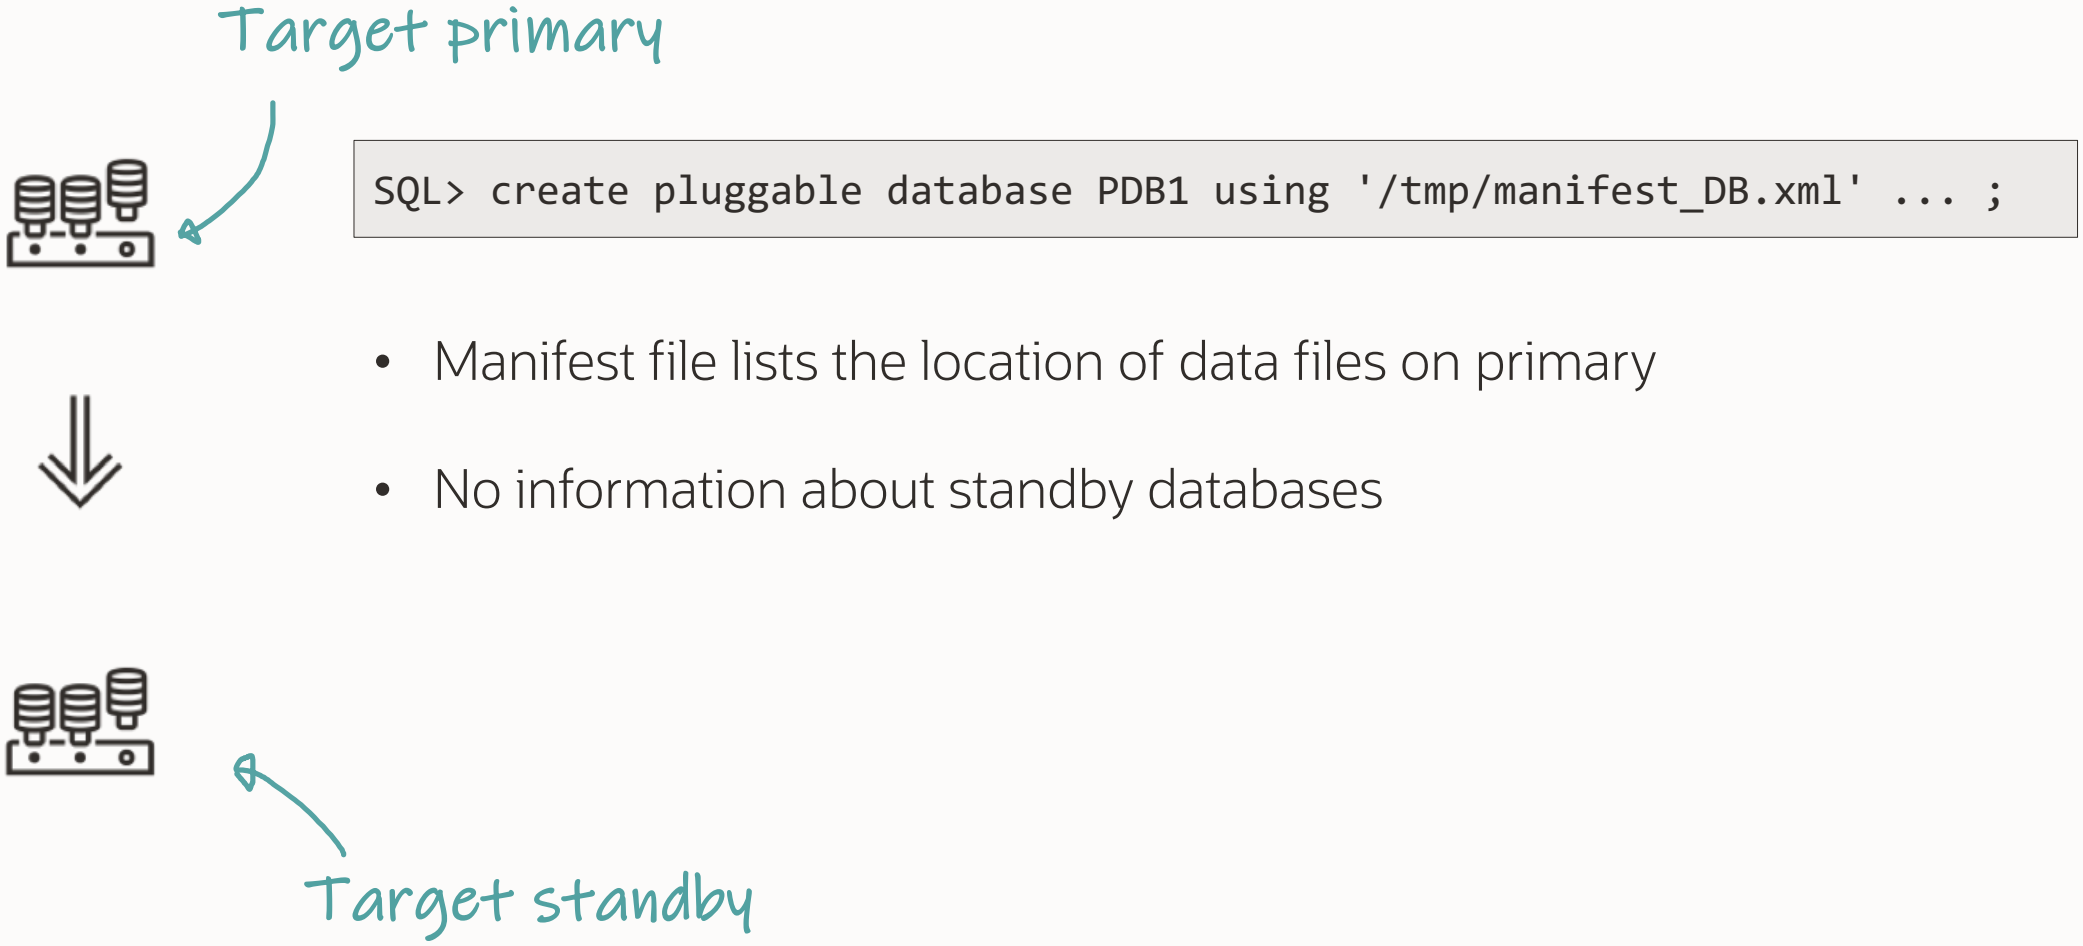

- Manifest file lists the location of data files on primary
- No information about standby databases

Enabled Recovery | ASM

23ai
CDB
Primary

+DATA/DB_BOSTON/DATAFILE/users.273.1103046827

Redo record says:
Plug in this data file

No good, data file
has a different name

23ai
CDB
Standby

+DATA/DB_CHICAGO/DATAFILE/users.269.1103050009

Enabled Recovery | ASM

23ai
CDB
Primary

+DATA/DB_BOSTON/DATAFILE/users.273.1103046827

23ai
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

23ai
CDB
Primary

+DATA/DB_BOSTON/DATAFILE/users.273.1103046827

OK, let's check the OMF directory

23ai
CDB
Standby

+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 (24.1) does not support plugging in with enabled recovery

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

Deferred Recovery

1

Enabled recovery

`create pluggable database ... standbys=all`

Standby records PDB creation

Standby locates data files

MRP applies redo to PDB

PDB is immediately protected

2

Deferred recovery

`create pluggable database ... standbys=none`

Standby records PDB creation

Standby ignores data files

MRP skips redo

PDB protected after restore

Deferred Recovery

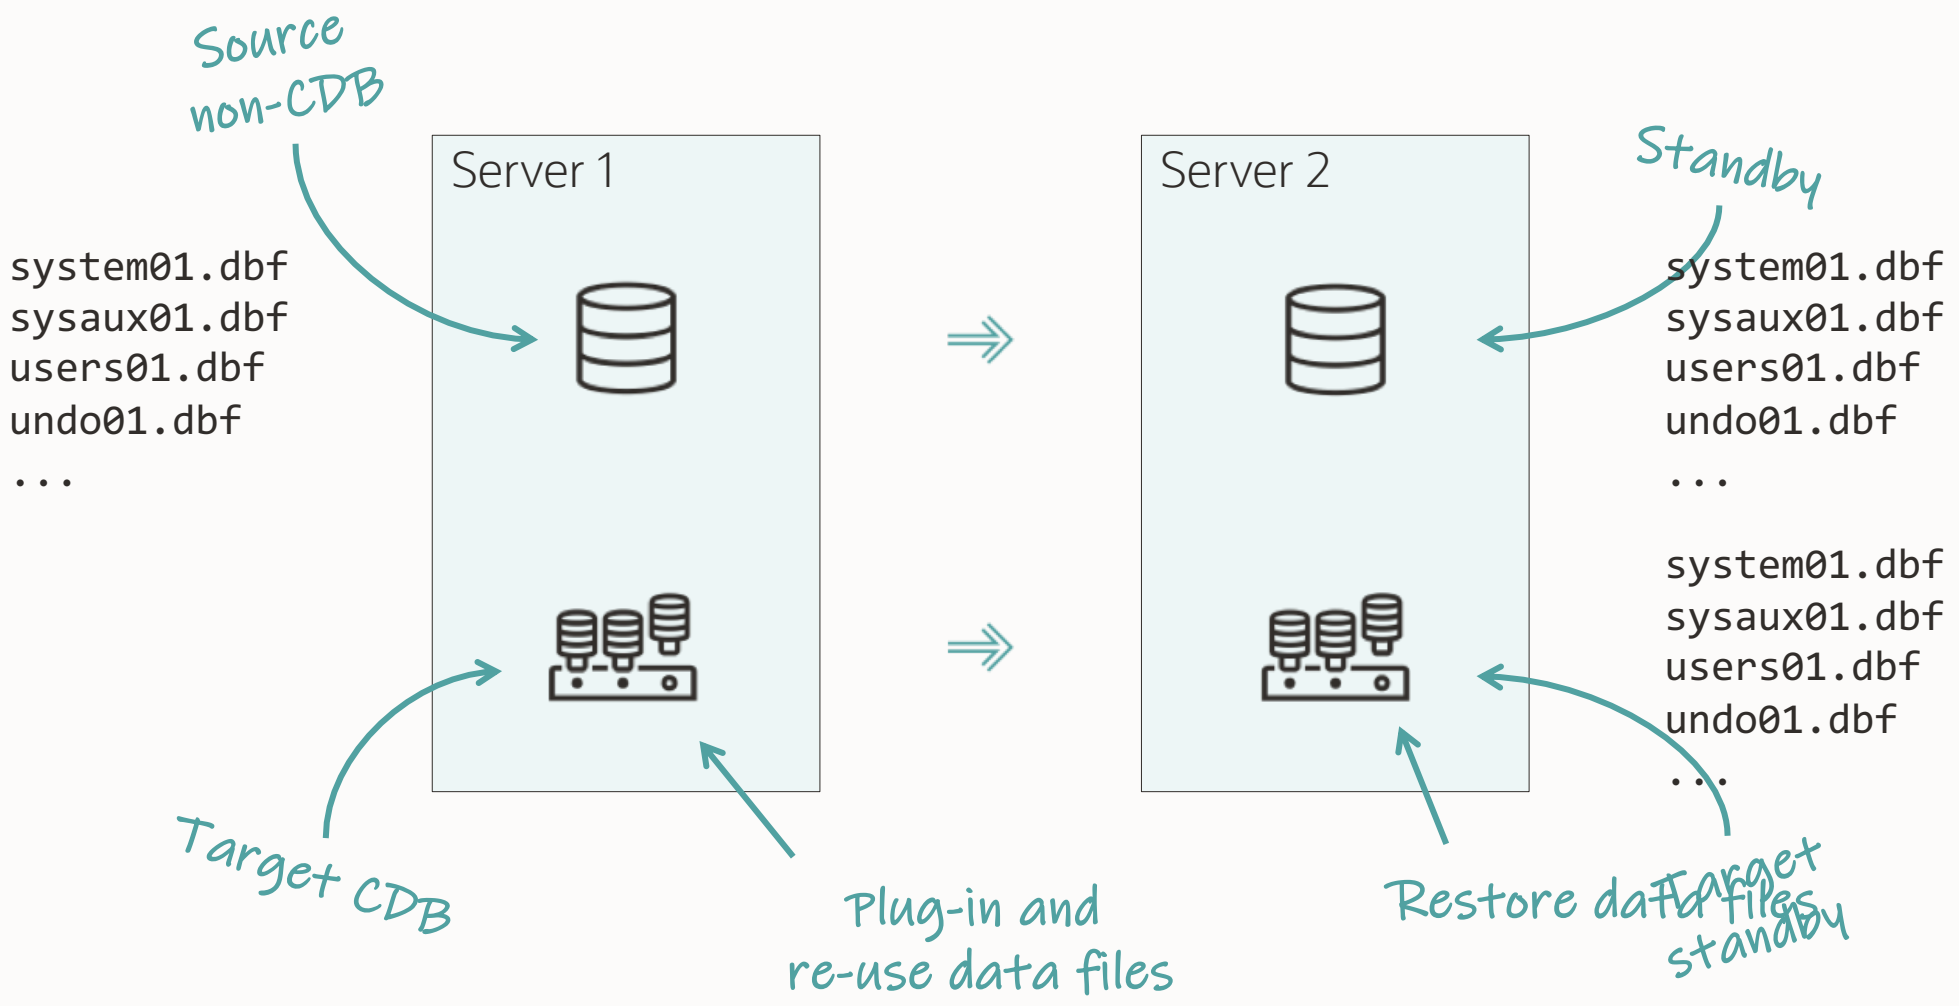

Deferred Recovery

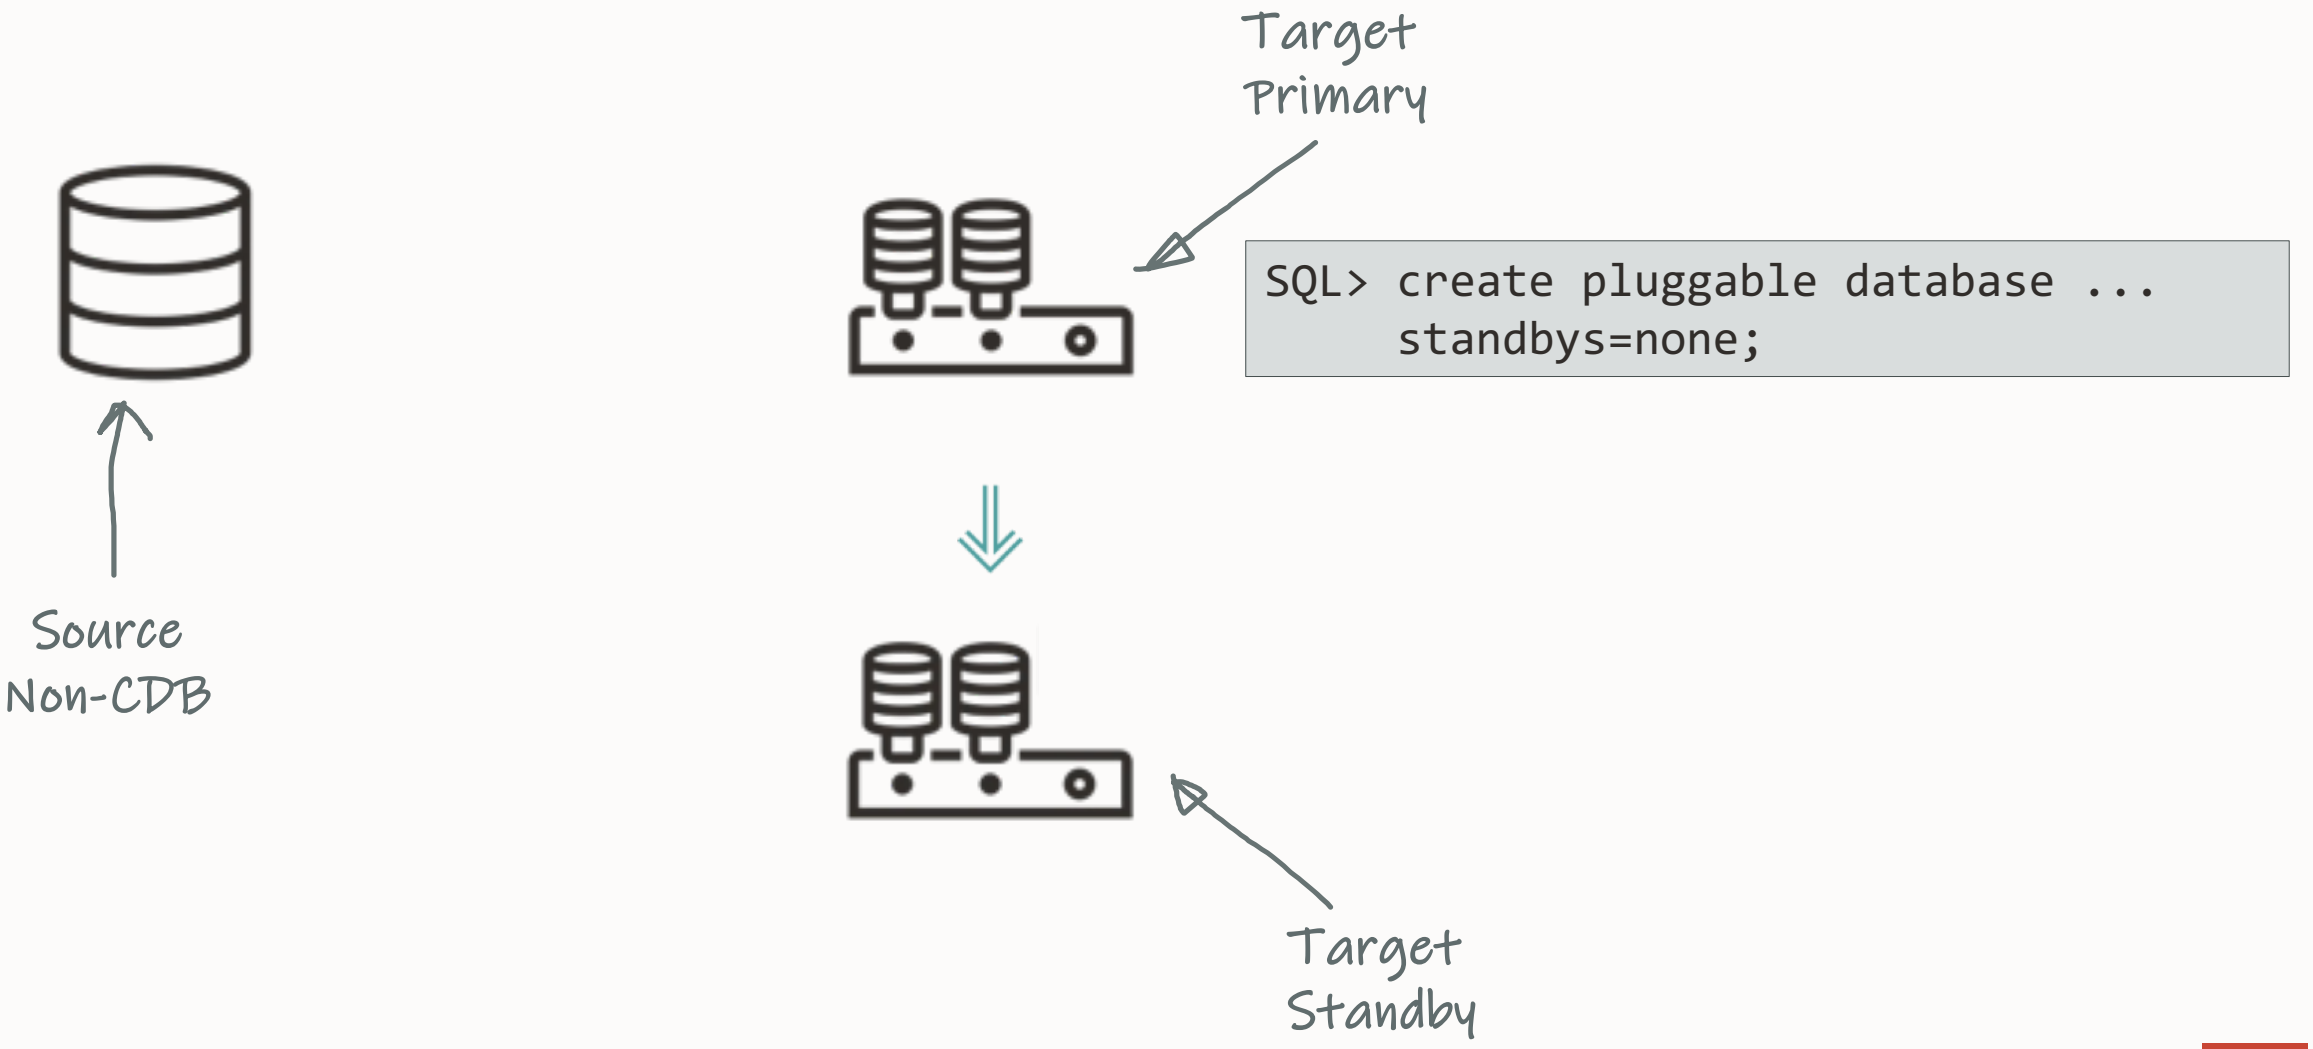

Deferred Recovery

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

PDB created
Data files missing

Deferred Recovery

```
SQL> show pdbs

CON_NAME     OPEN MODE
PDB1         READ WRITE
```

```
SQL> show pdbs

CON_NAME     OPEN MODE
PDB1         MOUNTED
```

Deferred Recovery

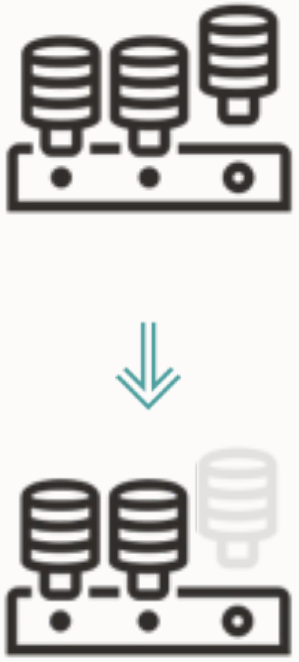

```
SQL> select name, recovery_status
     from v$pdbs;

NAME     RECOVERY_STATUS
PDB1     DISABLED
```

Deferred Recovery

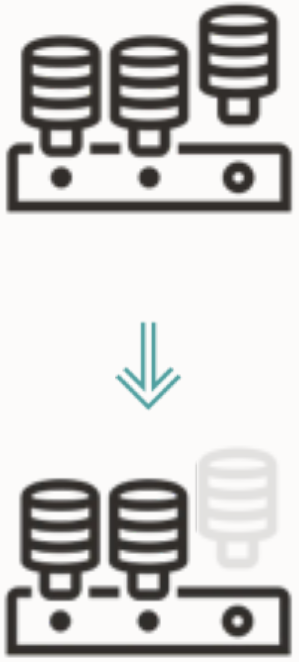

```
RMAN> restore pluggable database
      ... from service ... ;

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

Deferred Recovery

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

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

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

Further Details

Webinar:

Move to Oracle Database 23ai –
Everything you need to know about Oracle Multitenant – Part 1

Don't jeopardize your Data Guard

- Test the procedure and verify your environment

```
--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 SQL Plan Directives even when adaptive
--statistics are off.
--If you do not use Adaptive Statistics (optimizer_adaptive_statistics)
--then turn it completely off as described in MOS 2209560.1

alter system set "_sql_plan_directive_mgmt_control"=0;
```


```
--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

Lunch Break

We start again at 13:30

Data Pump Top Tips

Supercharge data loading/unloading

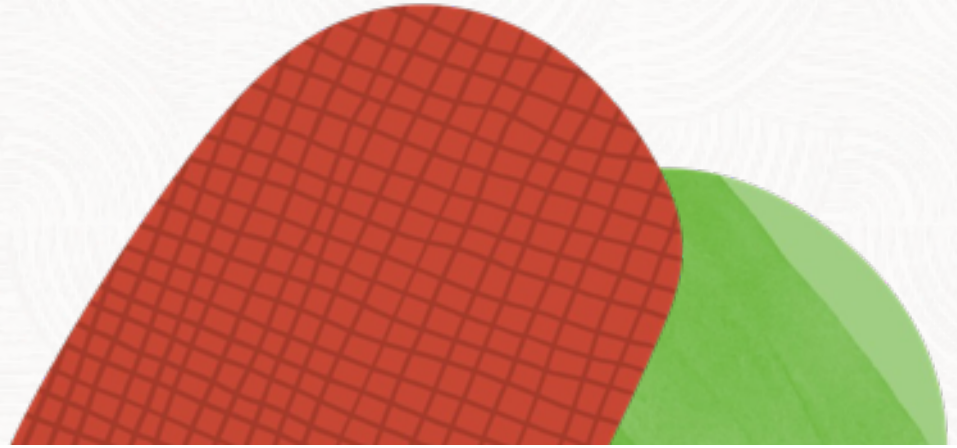

Always use
Data Pump Bundle Patch

More than 230 functional and performance fixes

- Data Pump Recommended Proactive Patches For 19.10 and Above (Doc ID 2819284.1)

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

Always 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=10G
```


```
-- 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 in Oracle Database 21c

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

expdp ... full=y transportable=always parallel=16
expdp ... tablespace=<list> parallel=16

impdp ... parallel=16
```


Parallel Transportable | Benchmark

Oracle E-Business Suite database

600.000+ objects

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

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

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_tl" 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_tl" 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

Also available in Oracle Database 19c via 19.23.0 Data Pump Bundle Patch

- Plus patch 37280692 - or be on 19.27

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`

Index Creation

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

From 12.1

Really good for many small indexes

Worker 1	`CREATE INDEX .... PARALLEL 1`
Worker 2	`CREATE INDEX .... PARALLEL 1`
...	`CREATE INDEX .... PARALLEL 1`
Worker 16	`CREATE INDEX .... PARALLEL 1`

Index Creation

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

From 23

The best of both worlds

Worker 1	`CREATE INDEX .... PARALLEL 1`
Worker 2	`CREATE INDEX .... PARALLEL 8`
Worker 3	`CREATE INDEX .... PARALLEL 4`
Worker 4	`CREATE INDEX .... PARALLEL 3`

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
```

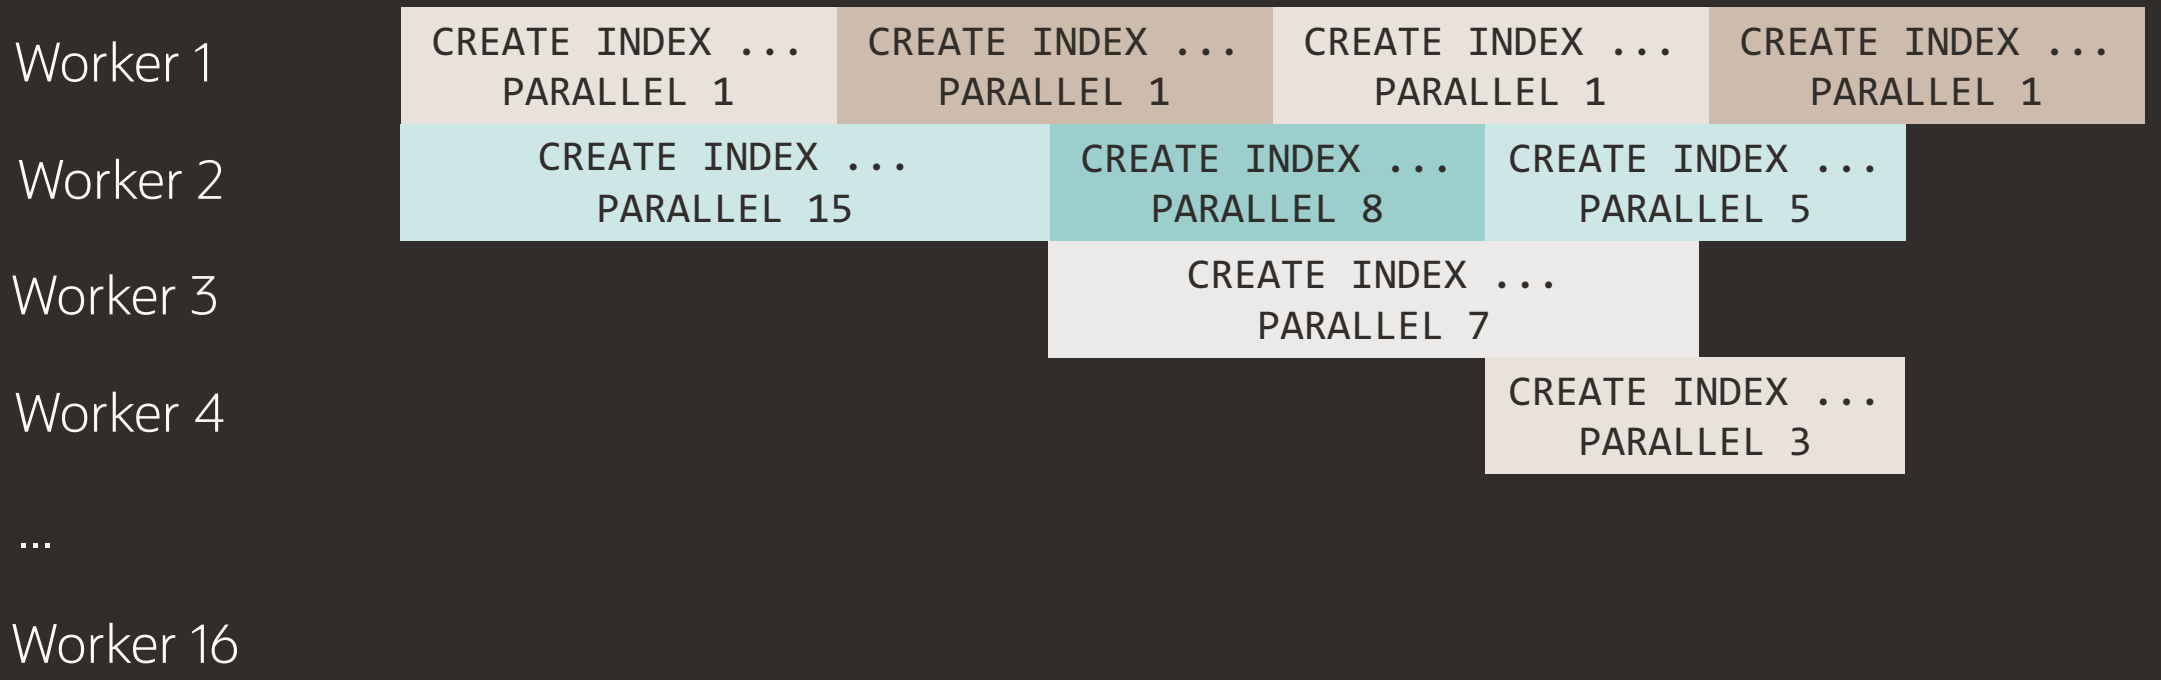

Benchmark

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

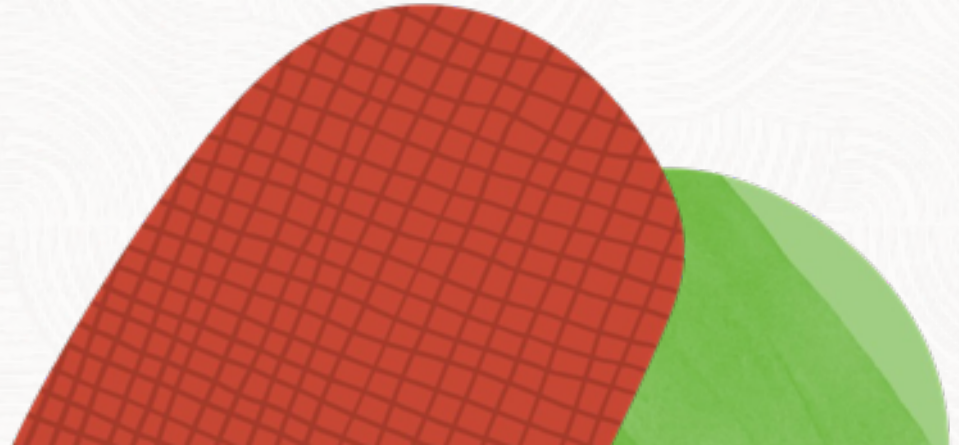

Photo by nikko macaspac on Unsplash

Photo by Dave on Unsplash

Performance Stability Prescription

Photo by Urip Dunker on Unsplash

Performance Stability Prescription

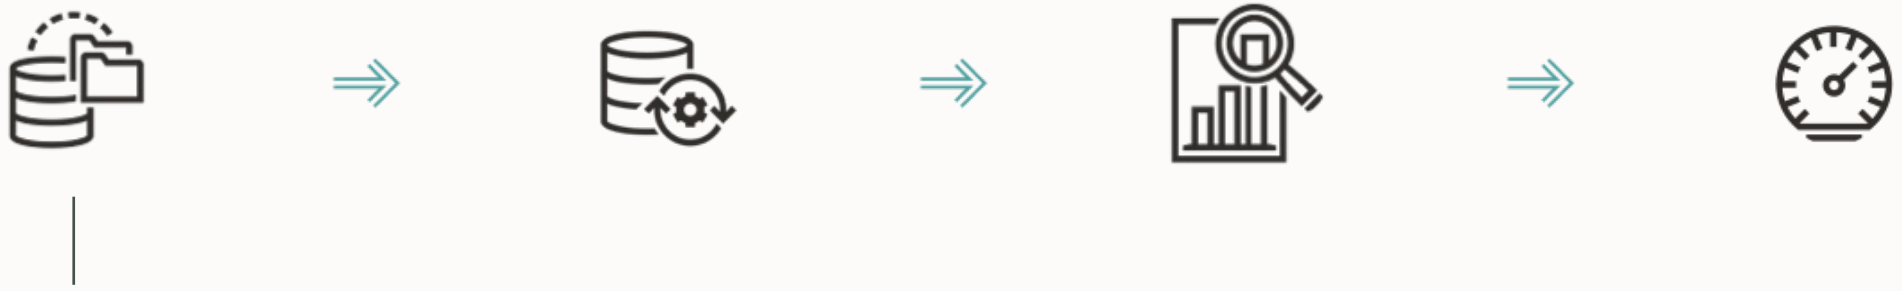

Collect workload information

- Sample from cursor cache
- Gather from AWR

Performance Stability Prescription

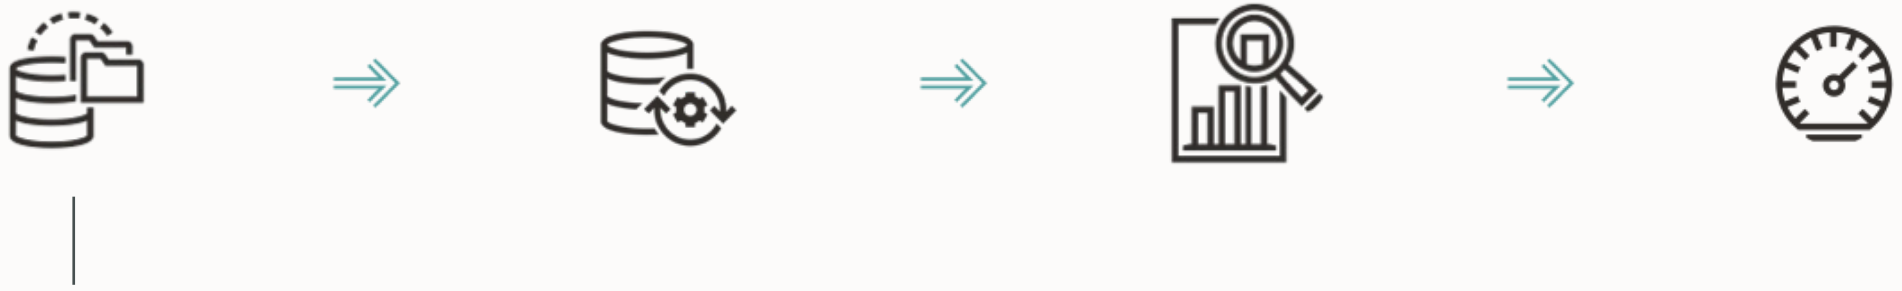

```
SQL> --Load statements from cursor cache
SQL> exec dbms_sqlset.capture_cursor_cache_sqlset( ...
SQL>
SQL> --Load statements from AWR
SQL> open ... table(dbms_sqltune.select_workload_repository) ..
SQL> dbms_sqlset.load_sqlset(...', cur);
```

SQL Tuning Set | Definition

SQL statement

Context

Statistics

Plans

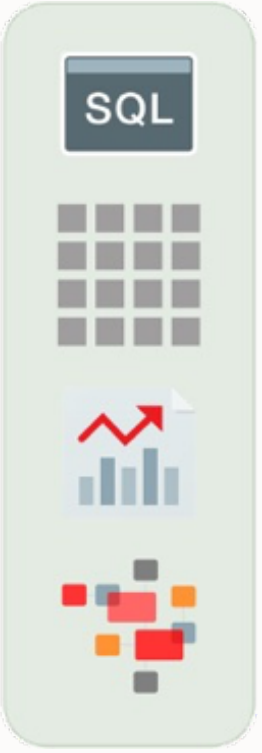

Gather at least a full month of workload data

- 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

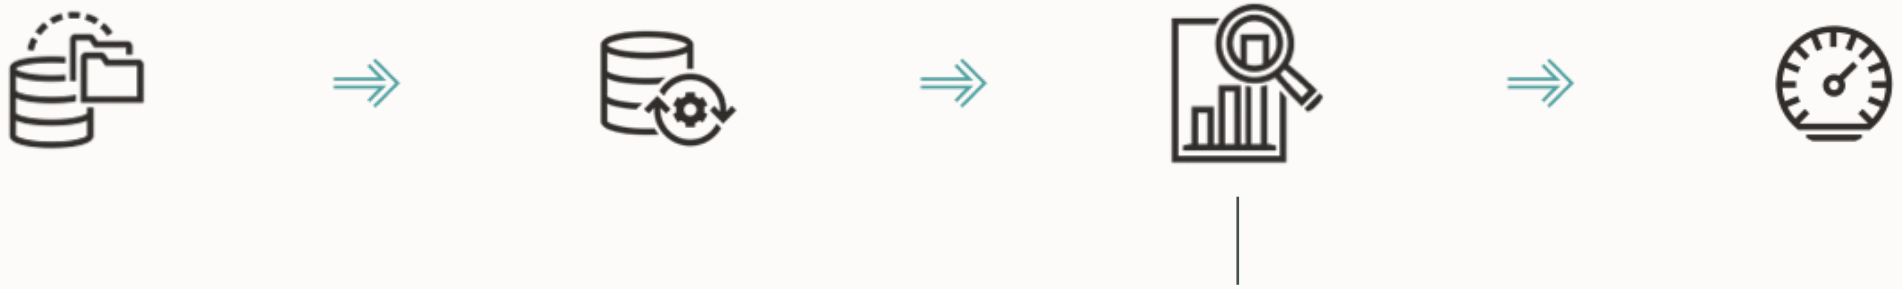

AWR Diff Report

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

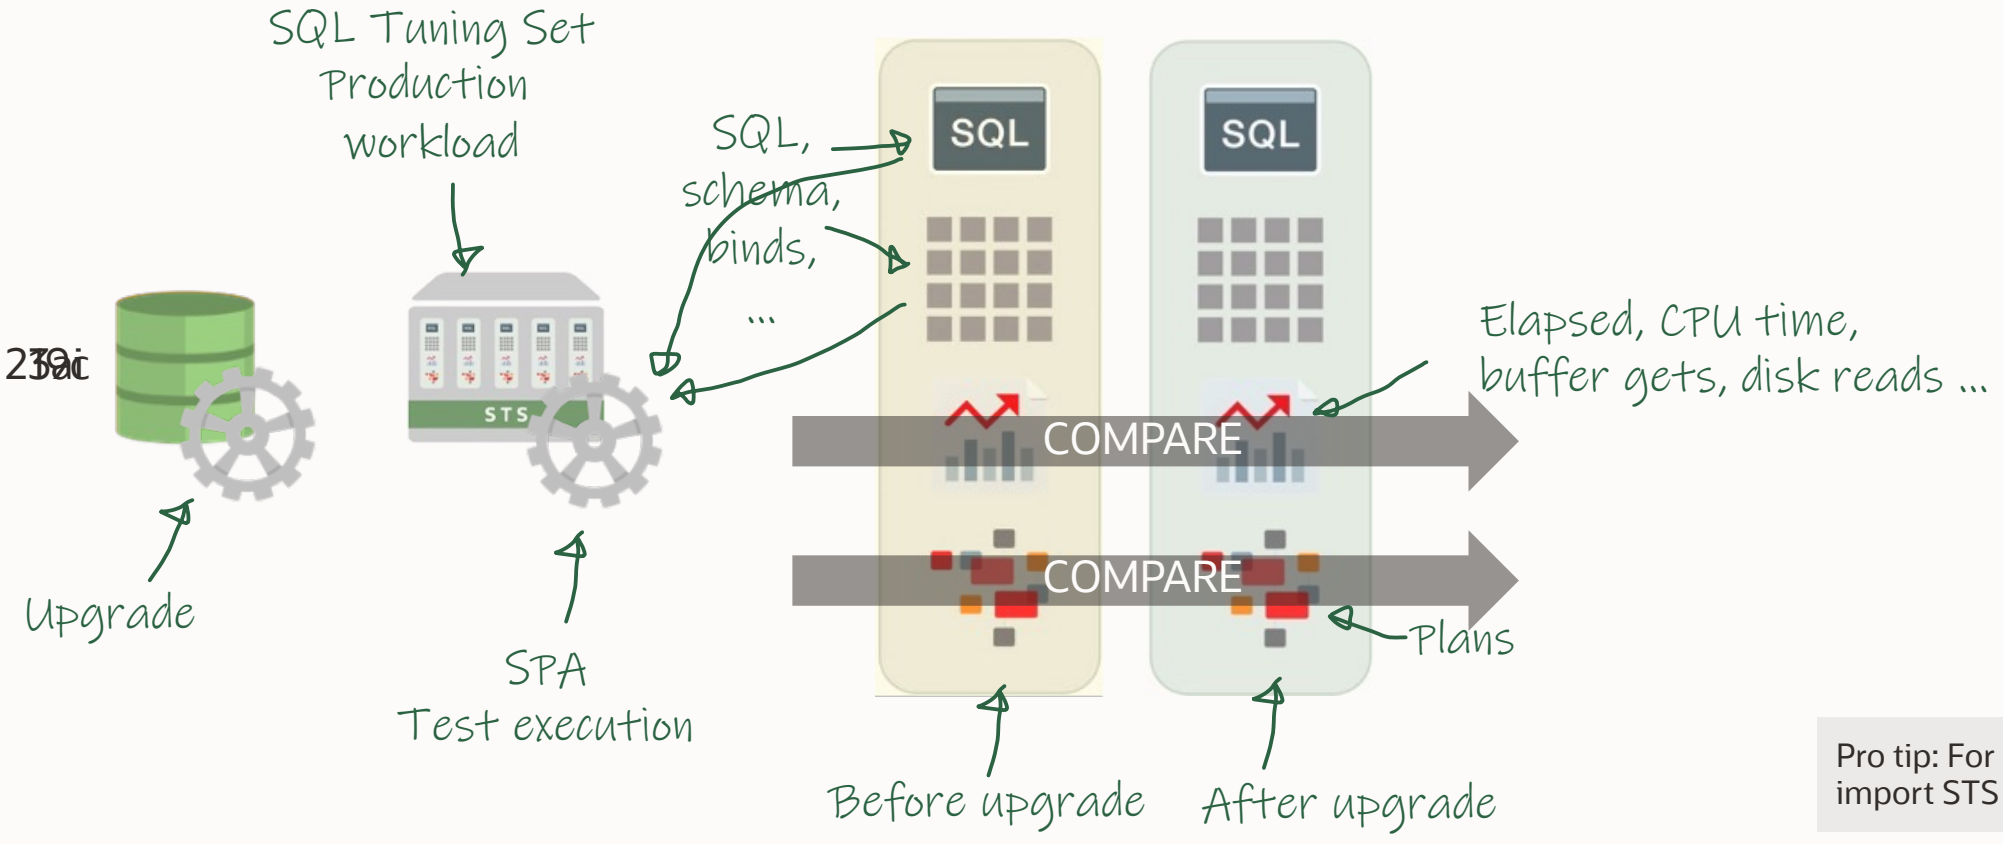

Pro tip: For migrations,
import STS into target database

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	7m5h0wf6stq0g	2.79%	21694	692.311883470084	490	29.22%	y
34	4wg725nwpxb1z	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		
⇩	3fv28gfu9y0ag	-0.050	26,504	29,573	-11.580	Y
⇩	czzzubf8fjz96	-0.030	1,410	1,981	-40.500	Y

SQL Details: czzzubf8fjz96

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		
⇩	3fv28gfu9y0ag	-0.050	26,504	29,573	-11.580	Y
⇩	czzzubf8fjz96	-0.030	1,410	1,981	-40.500	Y

Plan Comparison

SQL_TRIAL_13539424634446

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

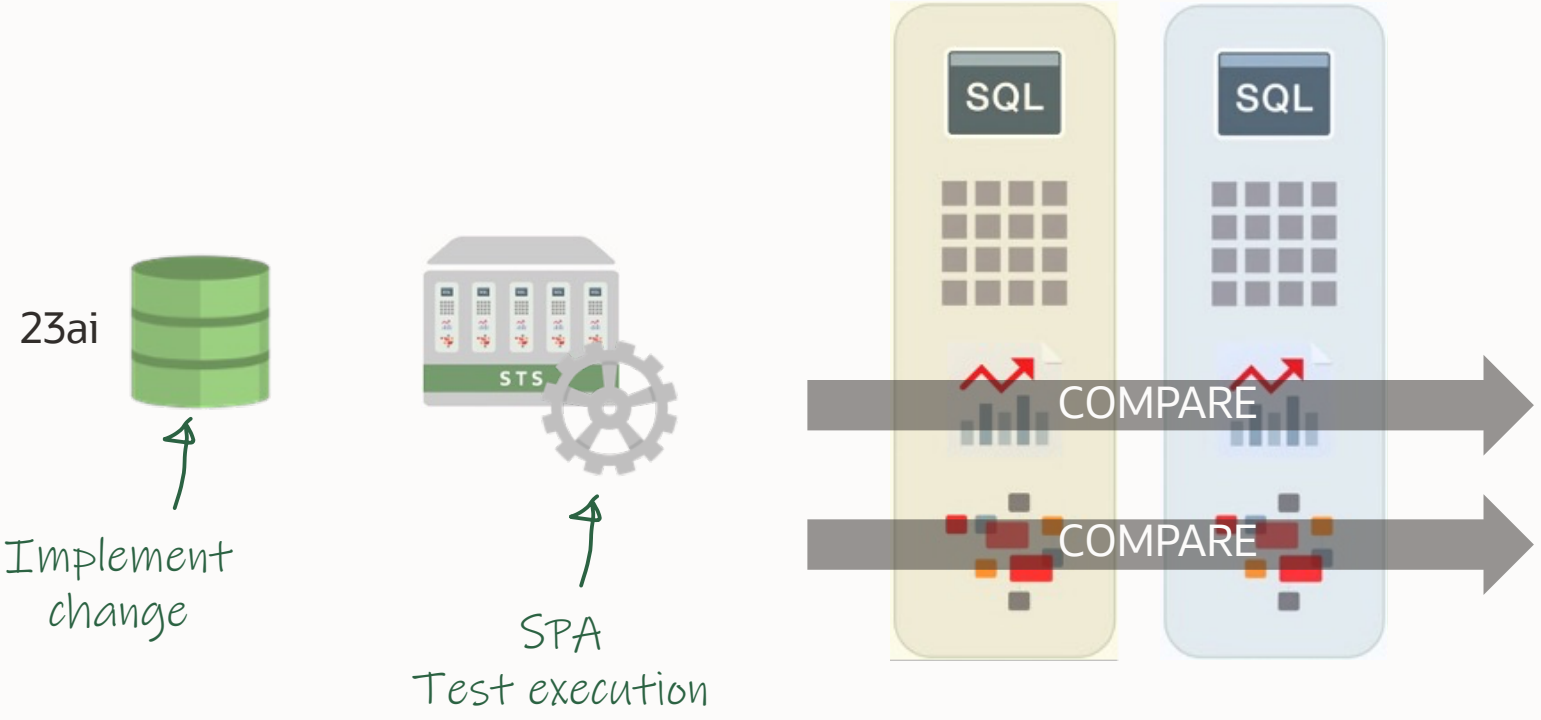

```
--If your queries have a lot of binds, you may tweak the below
--underscore. It defines how much bind data will be kept in memory
--for each query. Default at 400 is often too low.
--Find procedures and more on:
--https://blogs.oracle.com/coretec/post/spa-in-autonomous-database

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_
            'TABPART$', 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

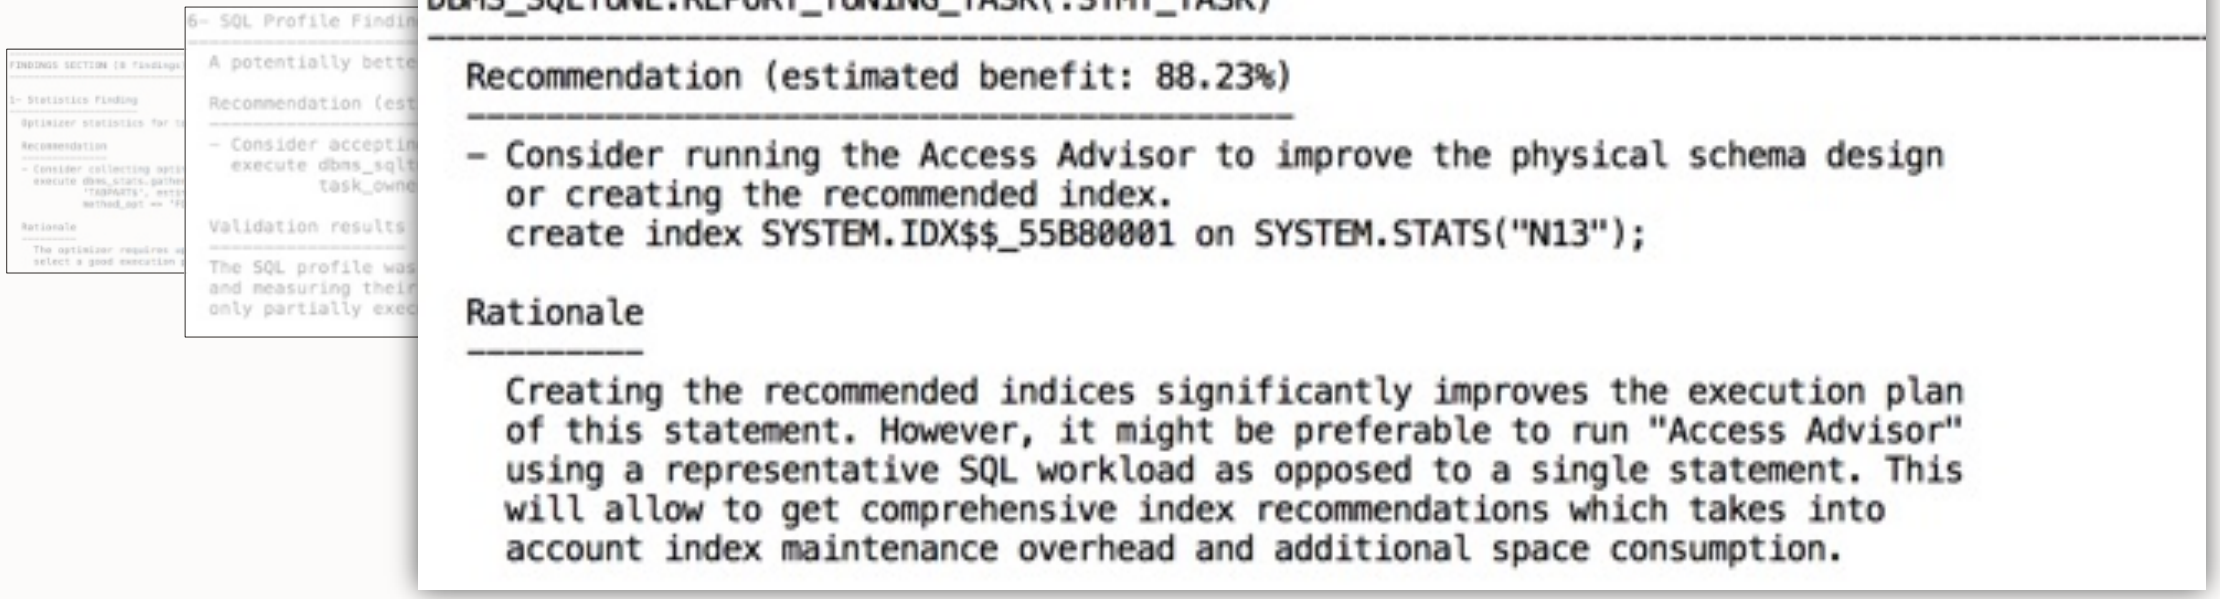

```
7- Index Finding (see explain plans section below)
--------------------------------------------------
  The execution plan of this statement can be improved by creating one or more
  indices.


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.
```

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

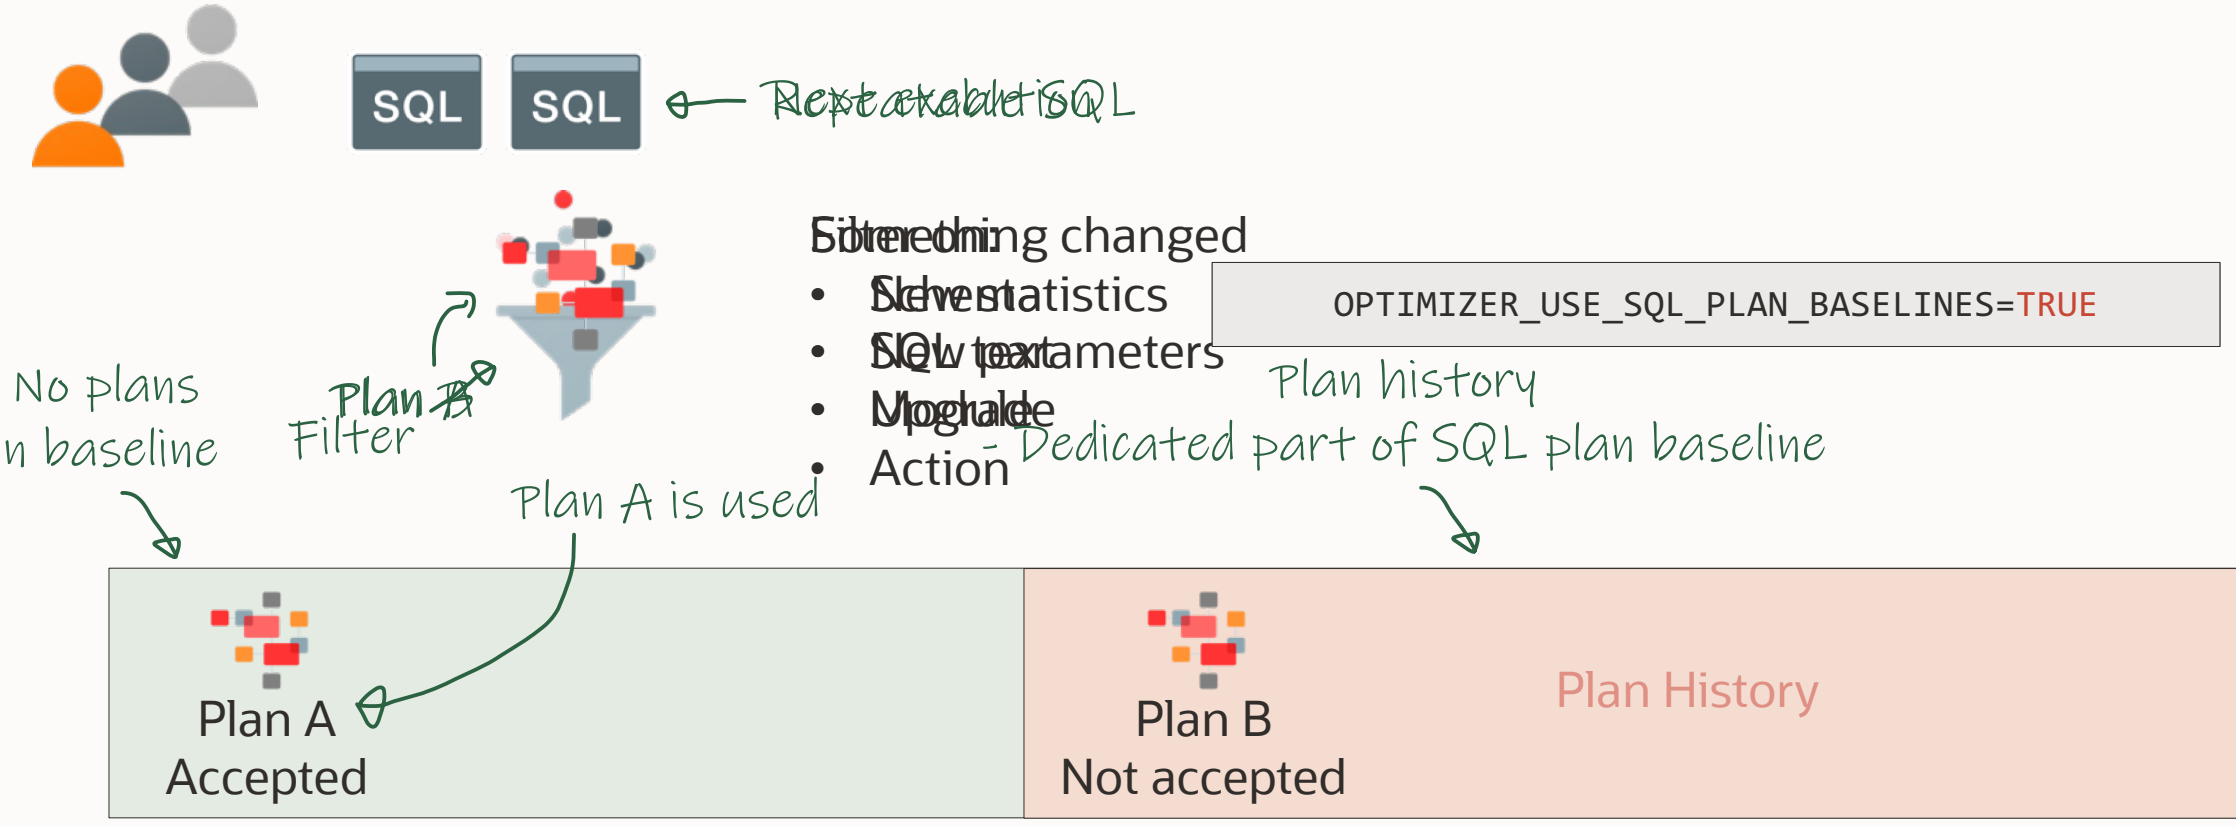

SPM | Evolve

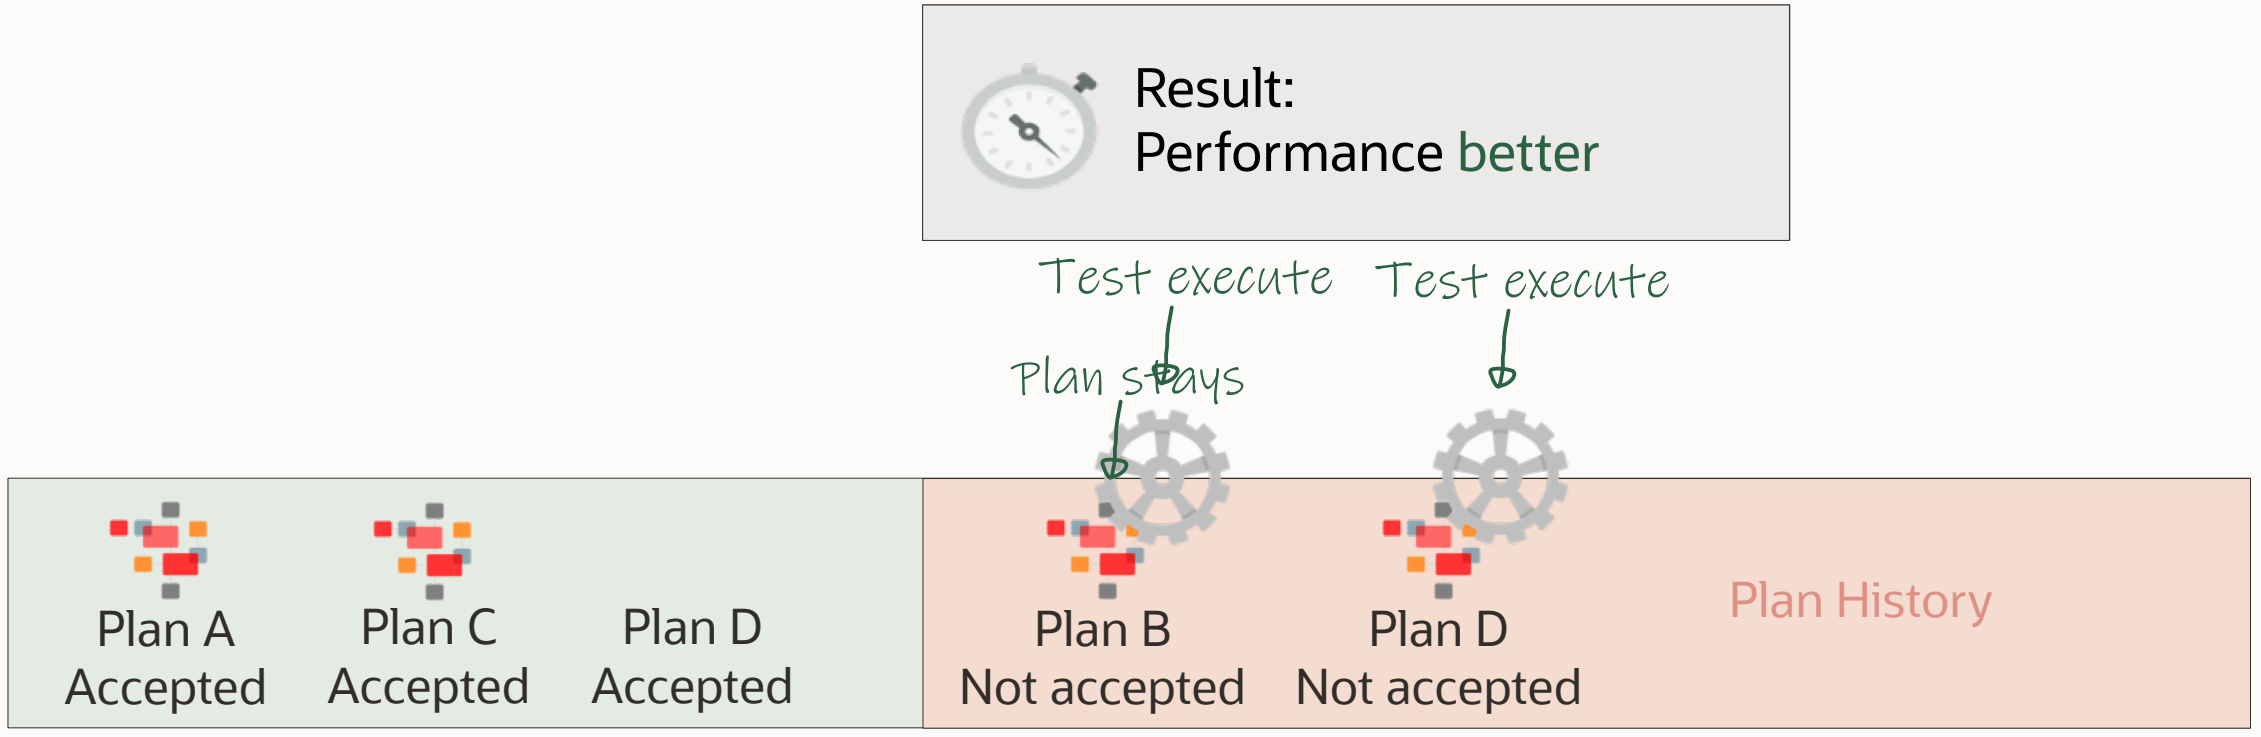

SPM | Load from STS


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

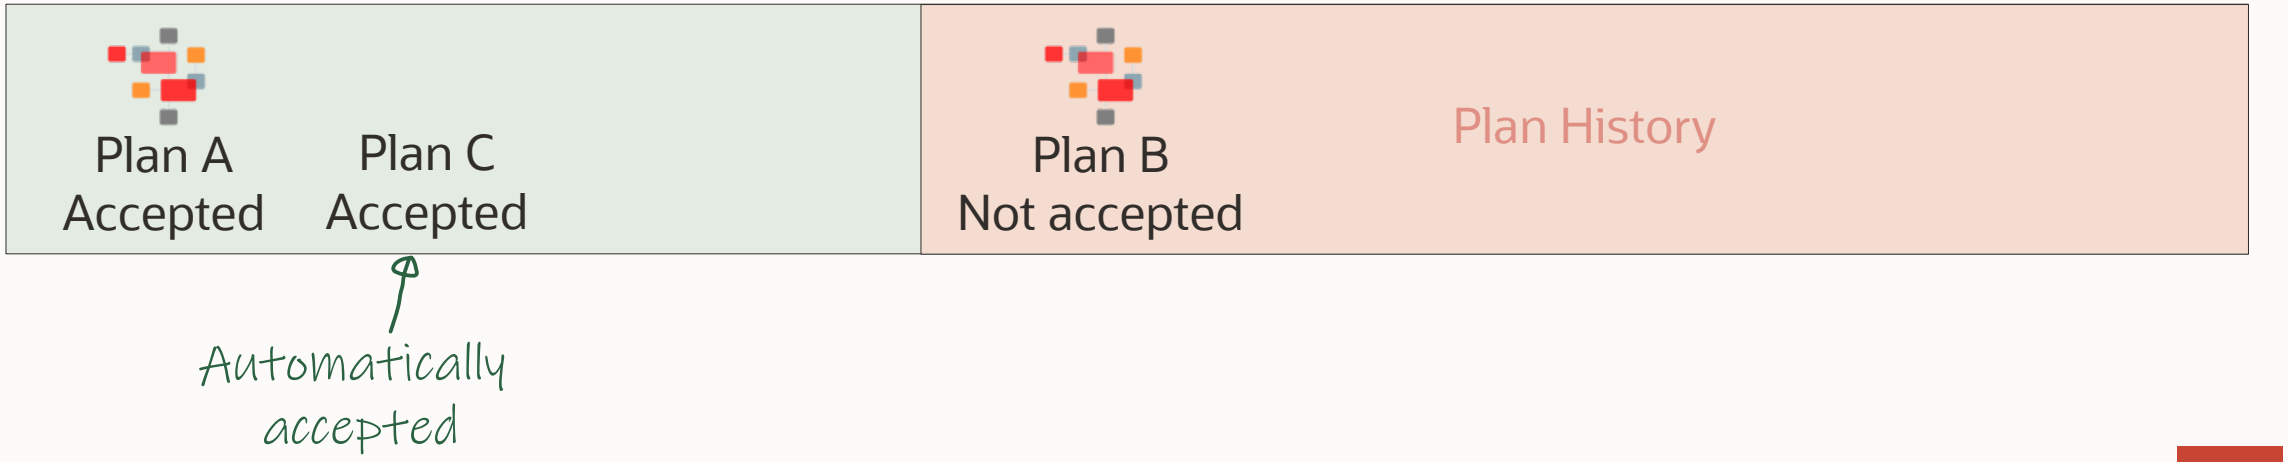

SPM | Use Case

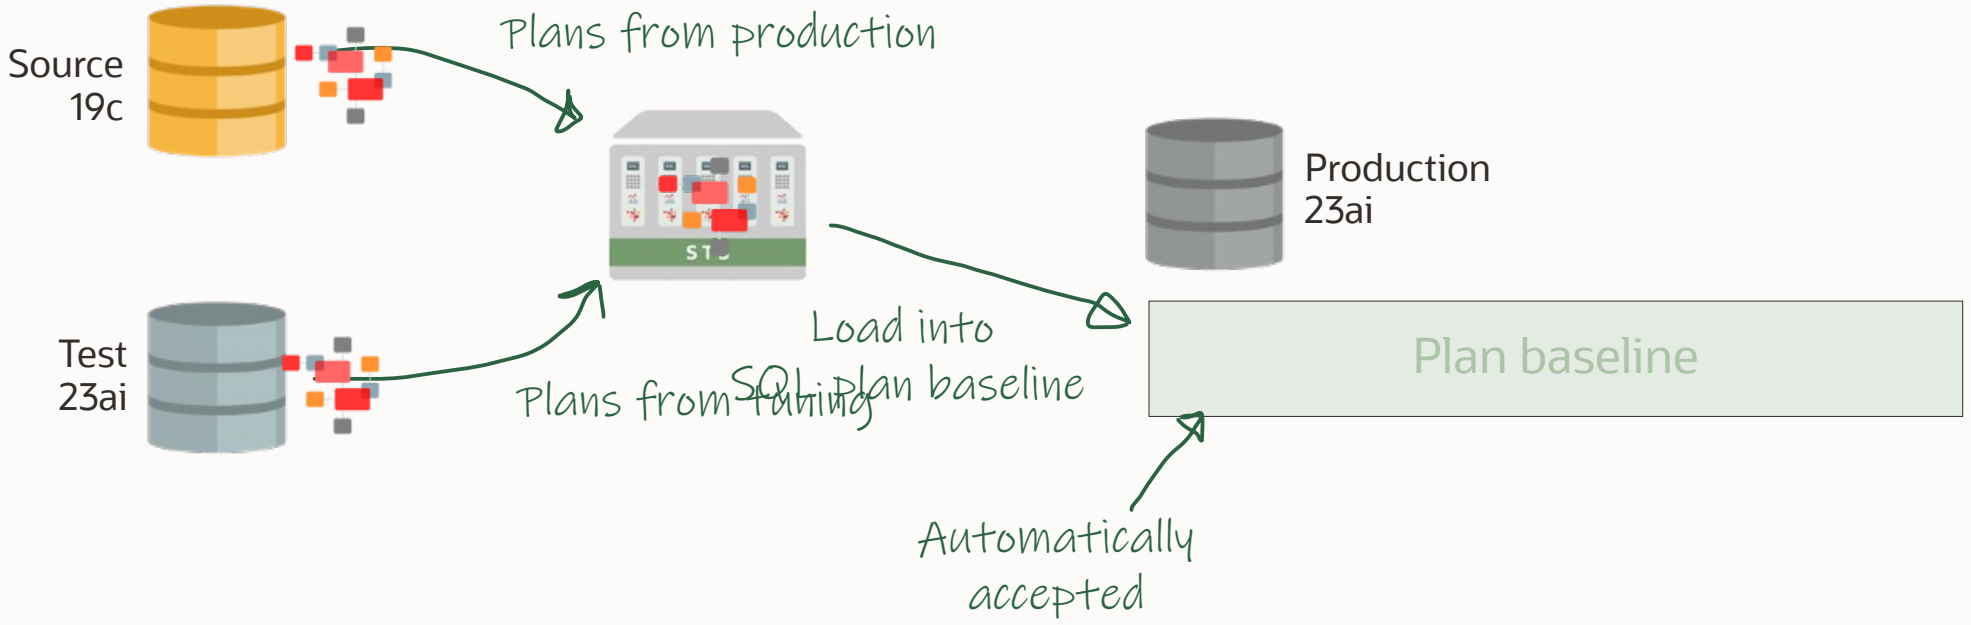

SPM | What If … Literals

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_sqltune.accept_sql_profile(..., category=>'TEST_ENV');

alter session set sqltune_category='TEST_ENV';

--After testing, fully enabled profile
dbms_sqltune.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 Database 23ai

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

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`

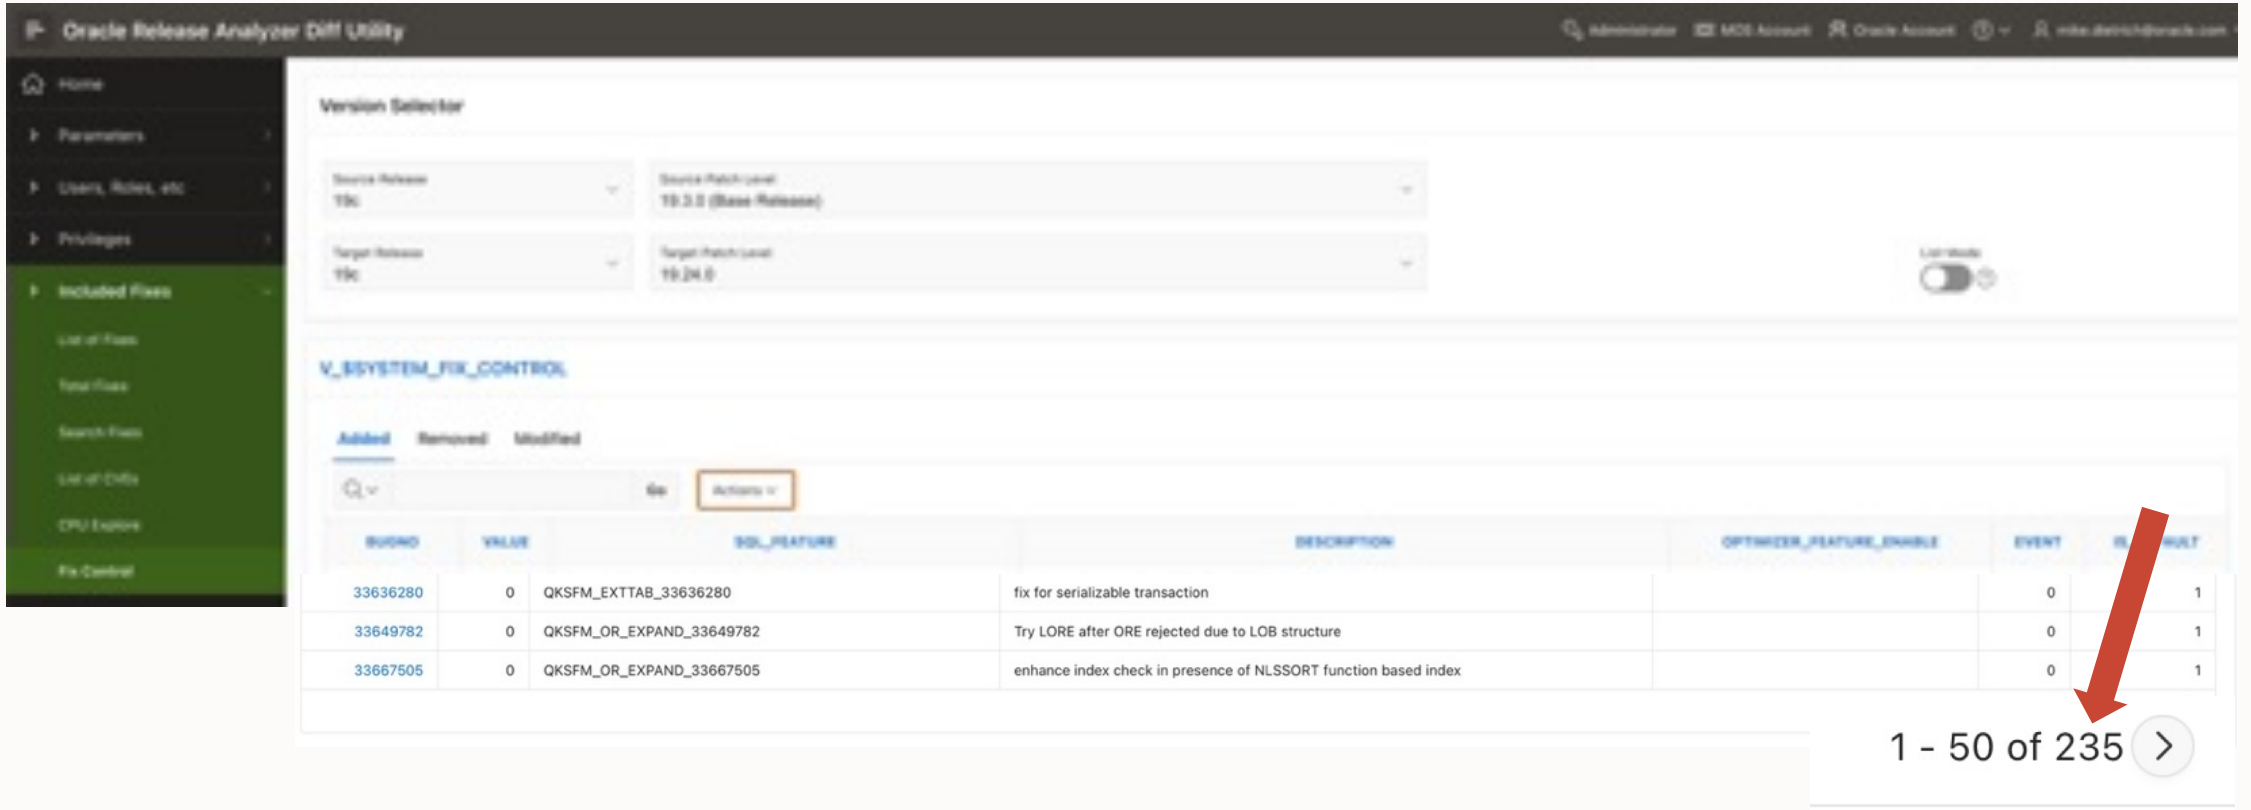

Should You Enable Optimizer Fixes?

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

Insights into the Oracle Database Development Process

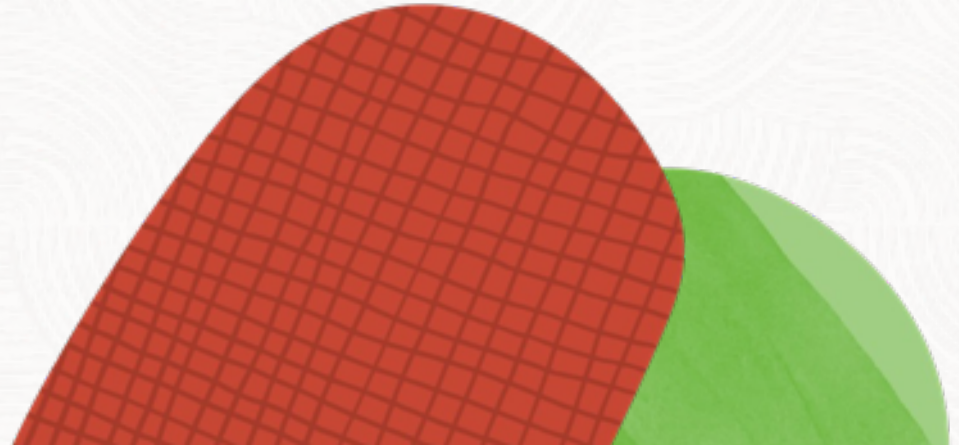

REALITY CHECK

A true customer story

Photo by Tobias Bjerknes on Unsplash

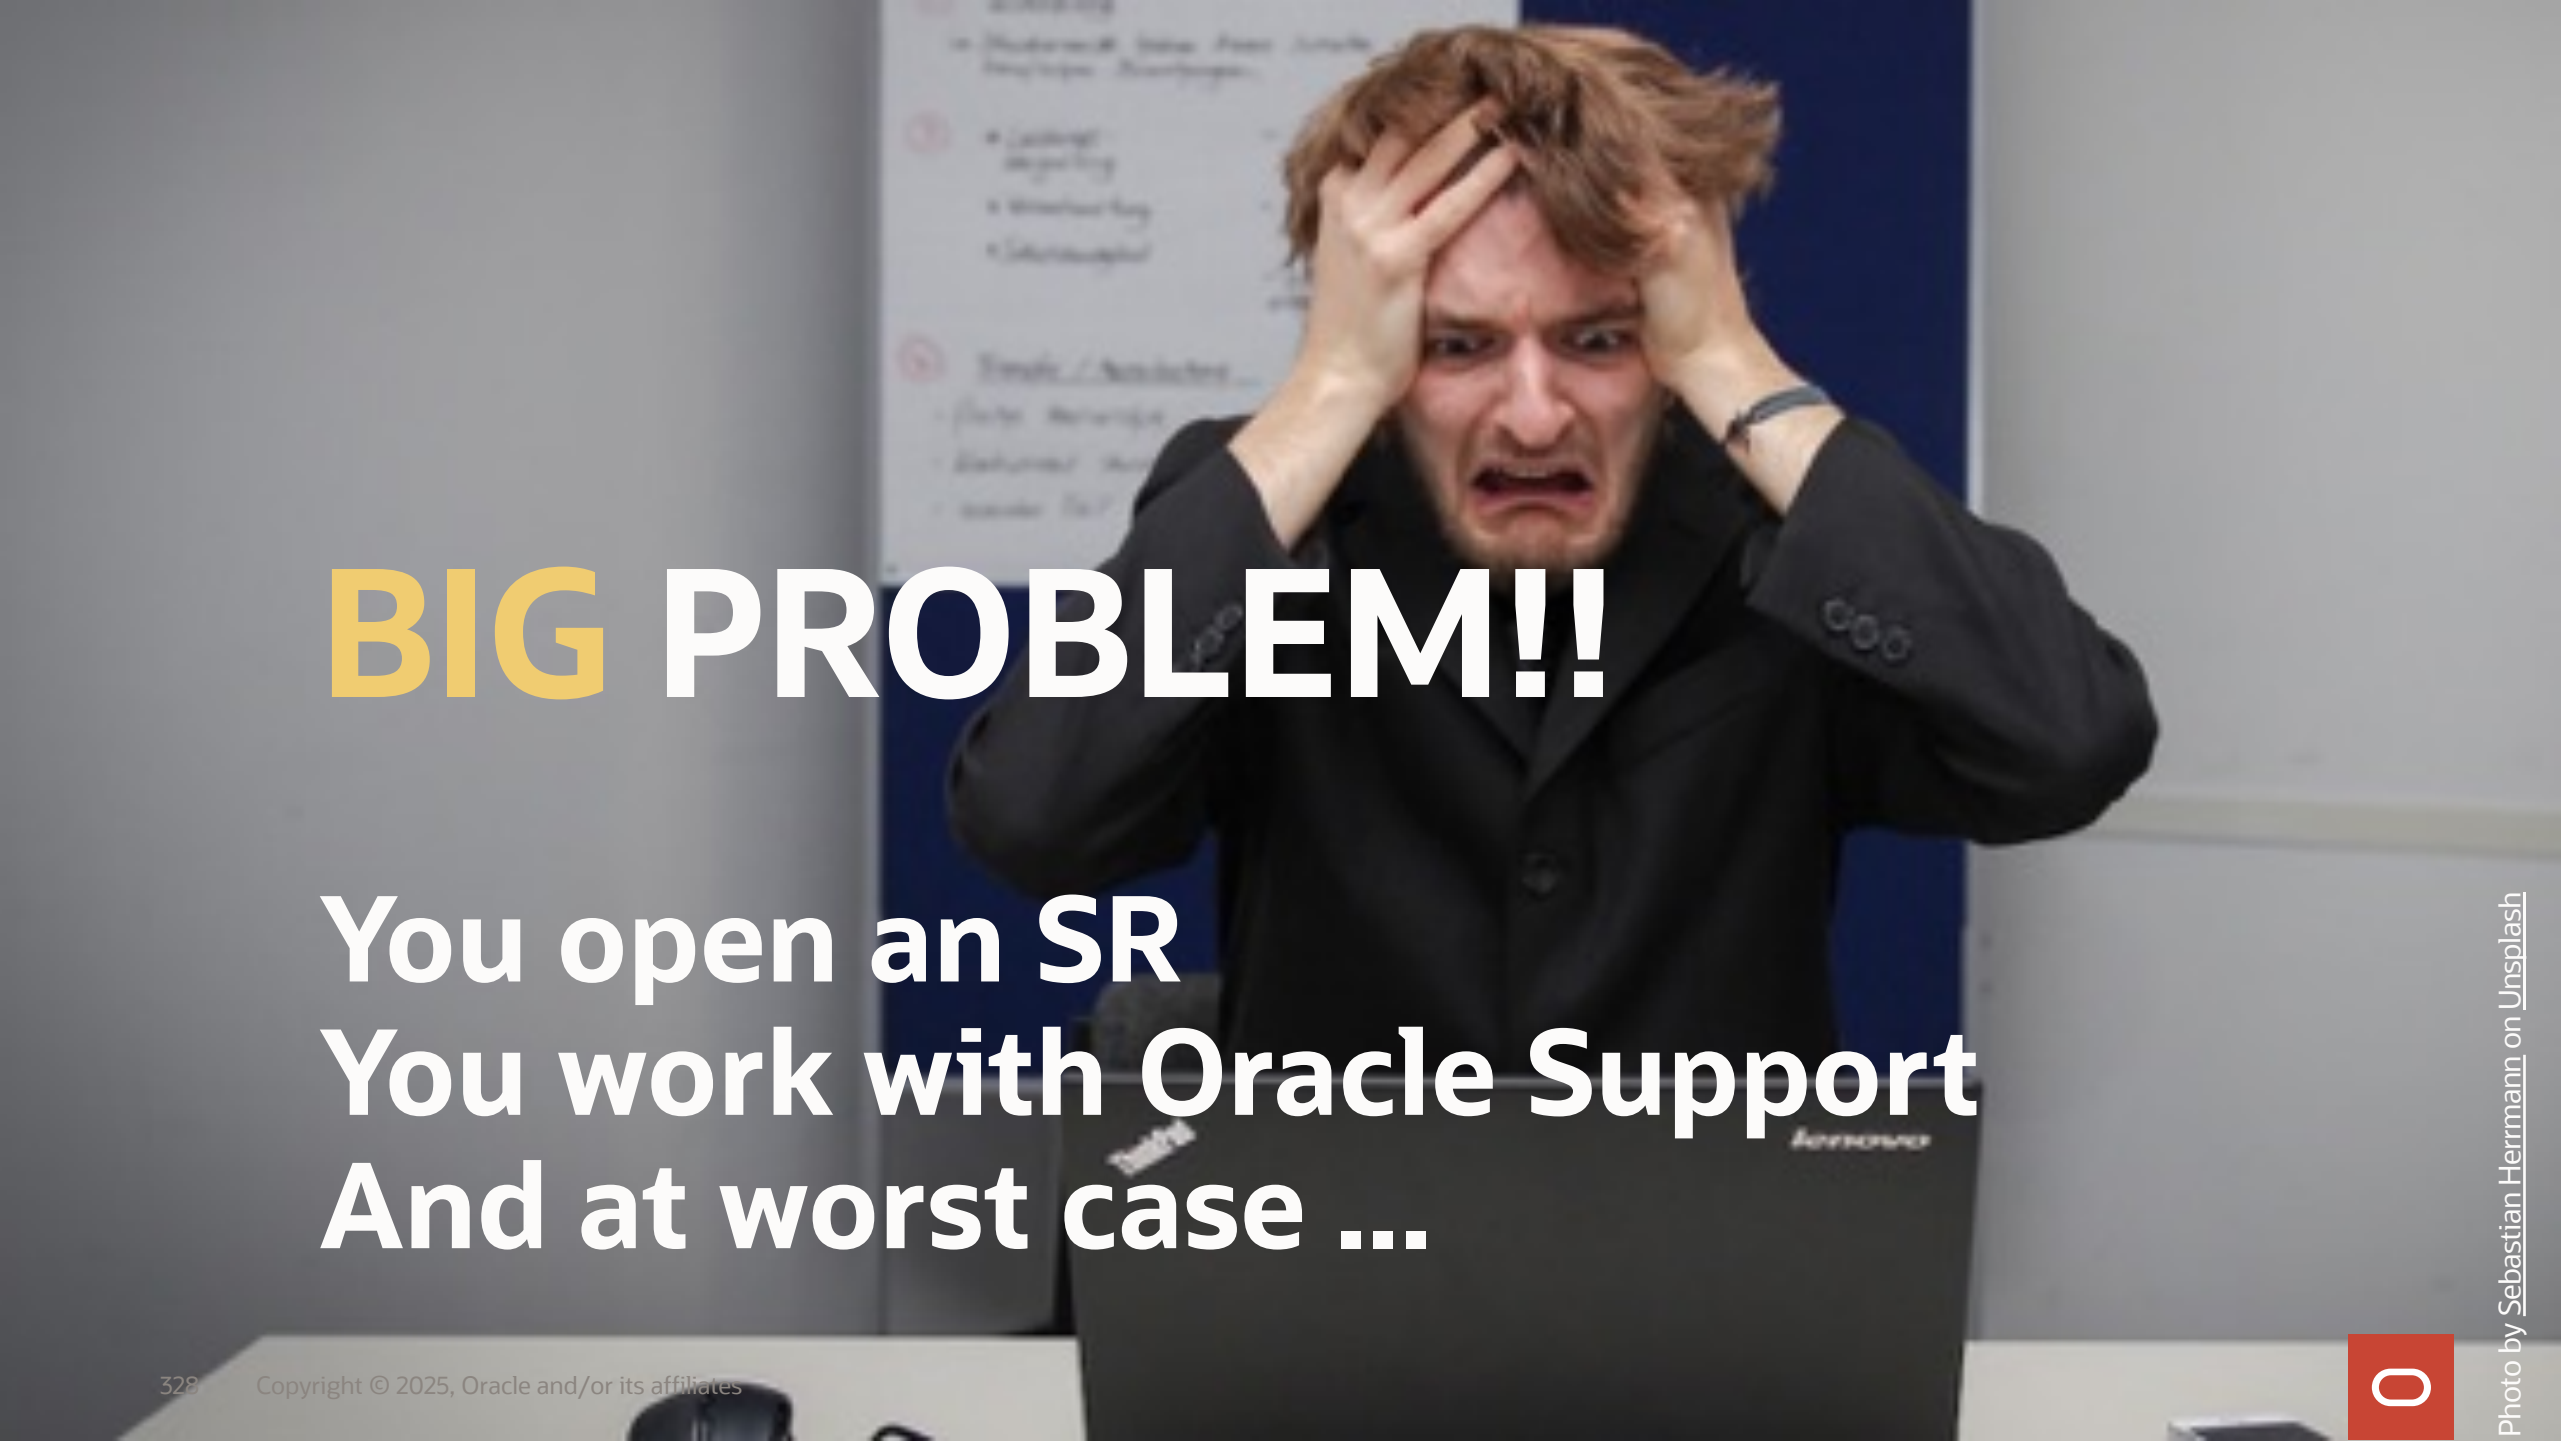

BIG PROBLEM!!

You open an SR
You work with Oracle Support
And at worst case ...

Photo by Sebastian Herrmann on Unsplash

… you hit a bug!!

Photo by Tobias Bjerknes on Unsplash

Photo by Viktor Forgacs on Unsplash

Photo by Andre Hunter on Unsplash

THREE PARTIES

- You
- Oracle Support
- Oracle Development

Photo by Joris Berthelot on Unsplash

Insights into the Patching Process

- Report via My Oracle Support
- Service Request inspected by automation

How many Service Requests are solved by automation within one day?

60 % of Service Requests get solved by automation within one single day

Insights into the Patching Process

- Support Engineer creates a bug

- Help us help you:
  - Clear description
  - Reproducible test case
  - Logs, traces, dumps
  - Various release tests

Bug 34774667 : ORA-600 [happens_on_mondays_only] before java is brewed

Bug Attributes

Type	8 - Defect	**Fixed in Product Version**	25.1
Severity	2 - Severe Loss of Service	**Product Version**	23.1
Status	80 - Development to QA/ Fix Delivered Internal	**Platform**	226 - Linux x86-64
Created	Nov 7, 2022	**Platform Version**	ORACLE LINUX GENERIC
Updated	Apr 30, 2024	**Base Bug**	33786319
Database Version	23.1	**Affects Platforms**	Generic
Product Source	Oracle	Knowledge, Patches, Service Requests and Bugs related to this bug	

Related Products

Line	Oracle Database Products	**Family**	Oracle Database Suite
Area	Oracle Database	**Product**	5 - Oracle Database - Enterprise Edition

```
Hdr: 34774667 : ORA-600 [happens_on_mondays_only] before java is brewed
Abstract: Happens on most important sales table. Sales are yelling ... loudly
```

When **your case** needs further attention

1

Update the SR
and raise severity

2

Call Oracle Support

Blog Post: Request Management Attention for your SR

Insights into the Patching Process

- Developer creates a fix
- Always in MAIN branch first

Insights into the Patching Process

- Developer creates a fix
- Always in MAIN branch first
- Regression tests on Farm in OCI
  - Subset of tests before merge
  - All tests run daily, ~1 million functional tests (runtime ~25.000 hours)

Insights into the Patching Process

- In case, fix isn't working as expected?
  - Fix the fix
  - Regression tests again
- Fix is ok?
  - Merge into MAIN
  - Merge can only happen with no failing tests

Branches

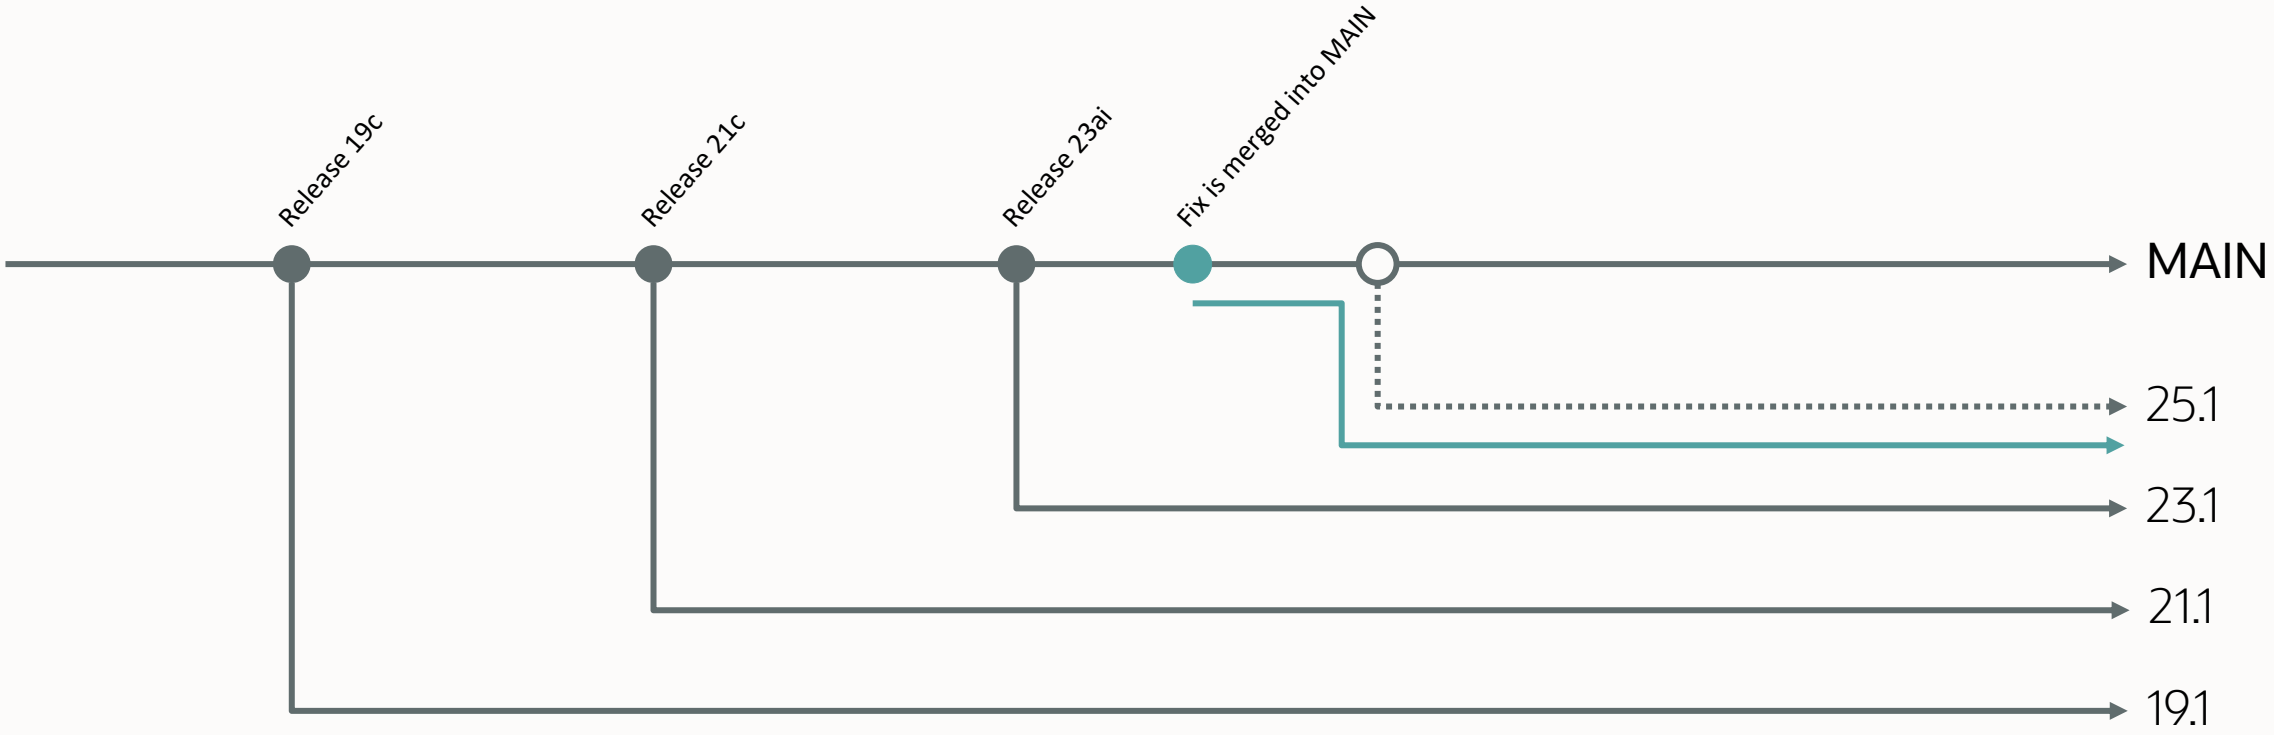
Release 19c
Release 21c
Release 23ai
Fix is merged into MAIN
MAIN
25.1
23.1
21.1
19.1

Bug 34774667 : ORA-600 [happens_on_mondays_only] before java is brewed

Bug Attributes

Type	8 - Defect	**Fixed in Product Version**	**25.1**
Severity	2 - Severe Loss of Service	Product Version	23.1
Status	80 - Development to QA/ Fix Delivered Internal	Platform	226 - Linux x86-64
Created	Nov 7, 2022	Platform Version	ORACLE LINUX GENERIC
Updated	Apr 30, 2024	Base Bug	33786319
Database Version	23.1	Affects Platforms	Generic
Product Source	Oracle	Knowledge, Patches, Service Requests and Bugs related to this bug	

Related Products

Line	Oracle Database Products	Family	Oracle Database Suite
Area	Oracle Database	Product	5 - Oracle Database - Enterprise Edition

```
Hdr: 34774667 : ORA-600 [happens_on_mondays_only] before java is brewed
Abstract: Happens on most important sales table. Sales are yelling ... loudly
```

Photo by Andre Hunter on Unsplash

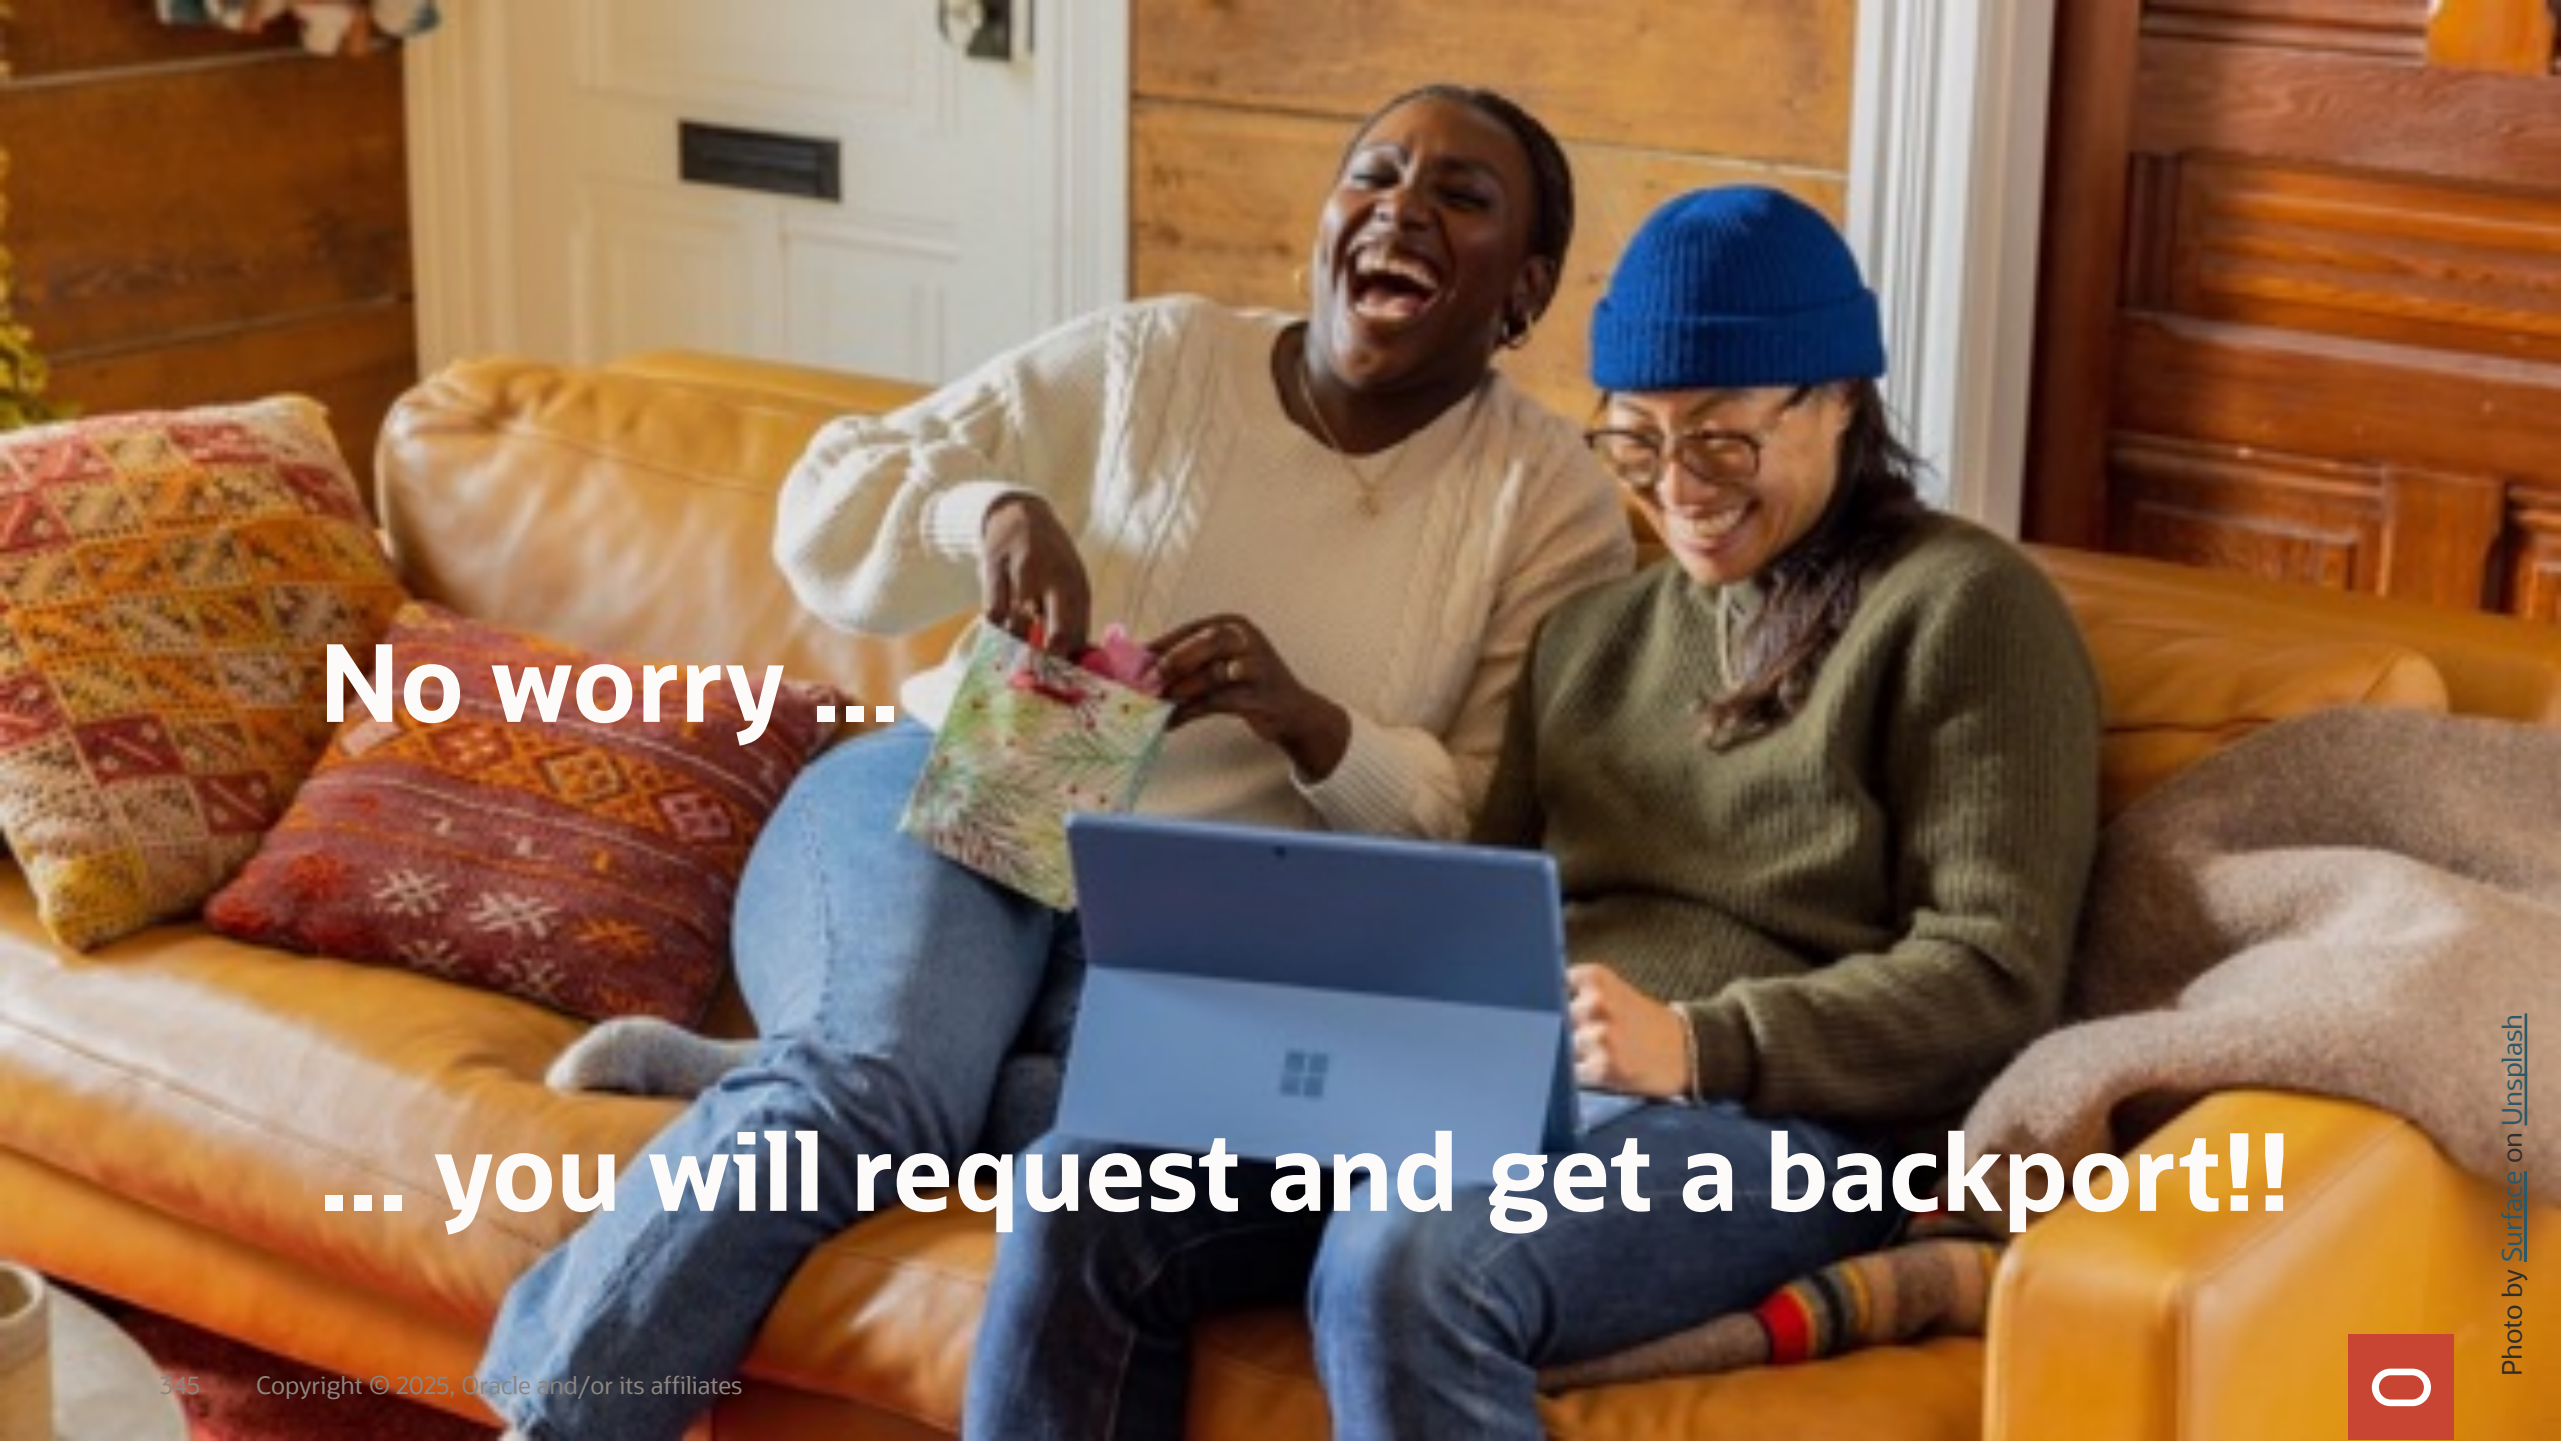

Photo by Surface on Unsplash

Insights into the Patching Process

- Backport created mostly on request
- Around 2.500 backports created proactively
- Backports are specific for a release
  and usually also for a patch level, e.g., on top of 19.28.0
- `opatch lsinventory`

Branches

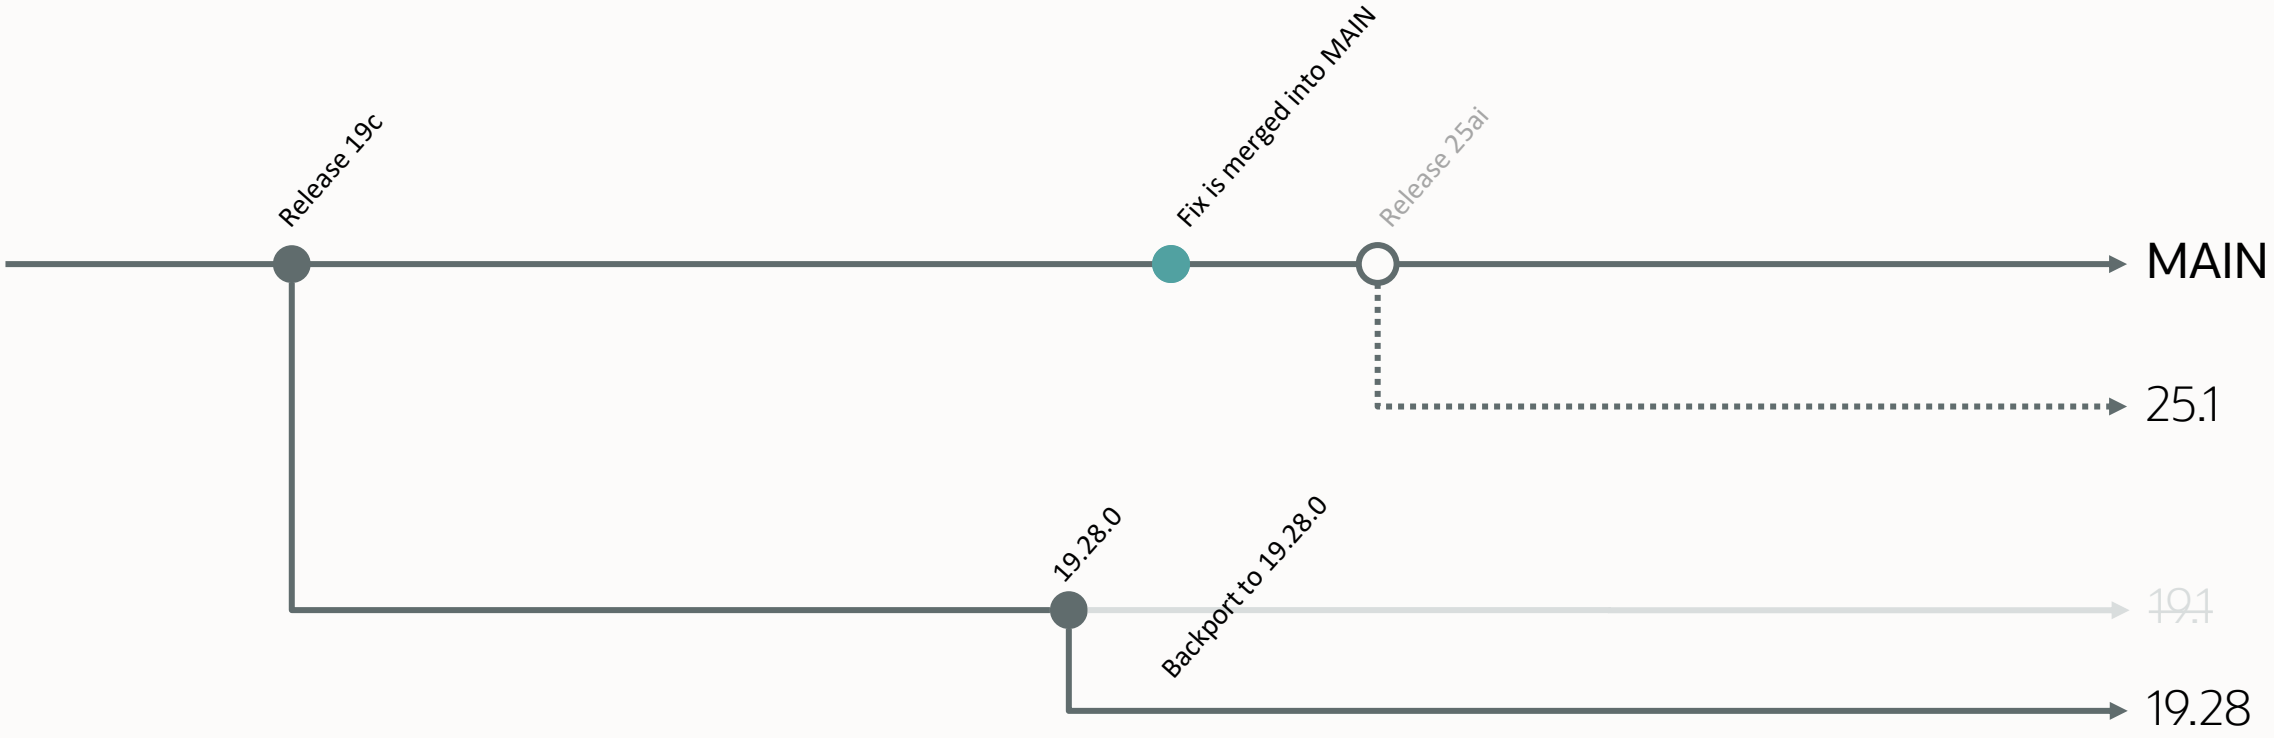
Release 19c
Fix is merged into MAIN
Release 25ai
MAIN
25.1
19.28.0
Backport to 19.28.0
19.1
19.28

Insights into the Patching Process

- Automated flow, manual handling might be required
  - Some conflicts
  - Dependencies
- Manual handling may postpone patch delivery

Patch Simple Search Results

Filters: Patch Name or Number is 31517417; Platform is Linux x86-64; Edit Search

Table ▾ View ▾ Detach Share Link

Patch Name	Description	Release	Platform (Language)
31517417	ORA-22308: OPERATION NOT ALLOWED ON EVOLVED TYP (Patch)	19.16.0.0.0DBRU	Generic Platform (American English)
31517417	ORA-22308: OPERATION NOT ALLOWED ON EVOLVED TYP (Patch)	19.15.0.0.0DBRU	Generic Platform (American English)
31517417	ORA-22308: OPERATION NOT ALLOWED ON EVOLVED TYP (Patch)	19.14.0.0.0DBRU	Generic Platform (American English)
31517417	ORA-22308: OPERATION NOT ALLOWED ON EVOLVED TYP (Patch)	19.13.0.0.0DBRU	Generic Platform (American English)
31517417	ORA-22308: OPERATION NOT ALLOWED ON EVOLVED TYP (Patch)	19.12.0.0.0DBRU	Generic Platform (American English)
31517417	ORA-22308: OPERATION NOT ALLOWED ON EVOLVED TYP (Patch)	19.11.0.0.0DBRU	Generic Platform (American English)
31517417	ORA-22308: OPERATION NOT ALLOWED ON EVOLVED TYP (Patch)	19.10.0.0.0DBRU	Generic Platform (American English)
31517417	ORA-22308: OPERATION NOT ALLOWED ON EVOLVED TYP (Patch)	19.9.0.0.0DBRU	Generic Platform (American English)
31517417	ORA-22308: OPERATION NOT ALLOWED ON EVOLVED TYP (Patch)	19.8.0.0.0DBRU	Generic Platform (American English)
31517417	ORA-22308: OPERATION NOT ALLOWED ON EVOLVED TYP (Patch)	19.6.0.0.0DBRU	Generic Platform (American English)
31517417	ORA-22308: OPERATION NOT ALLOWED ON EVOLVED TYP (Patch)	19.4.0.0.0DBRU	Generic Platform (American English)

19.7.0 missing

Can I Request a Backport?

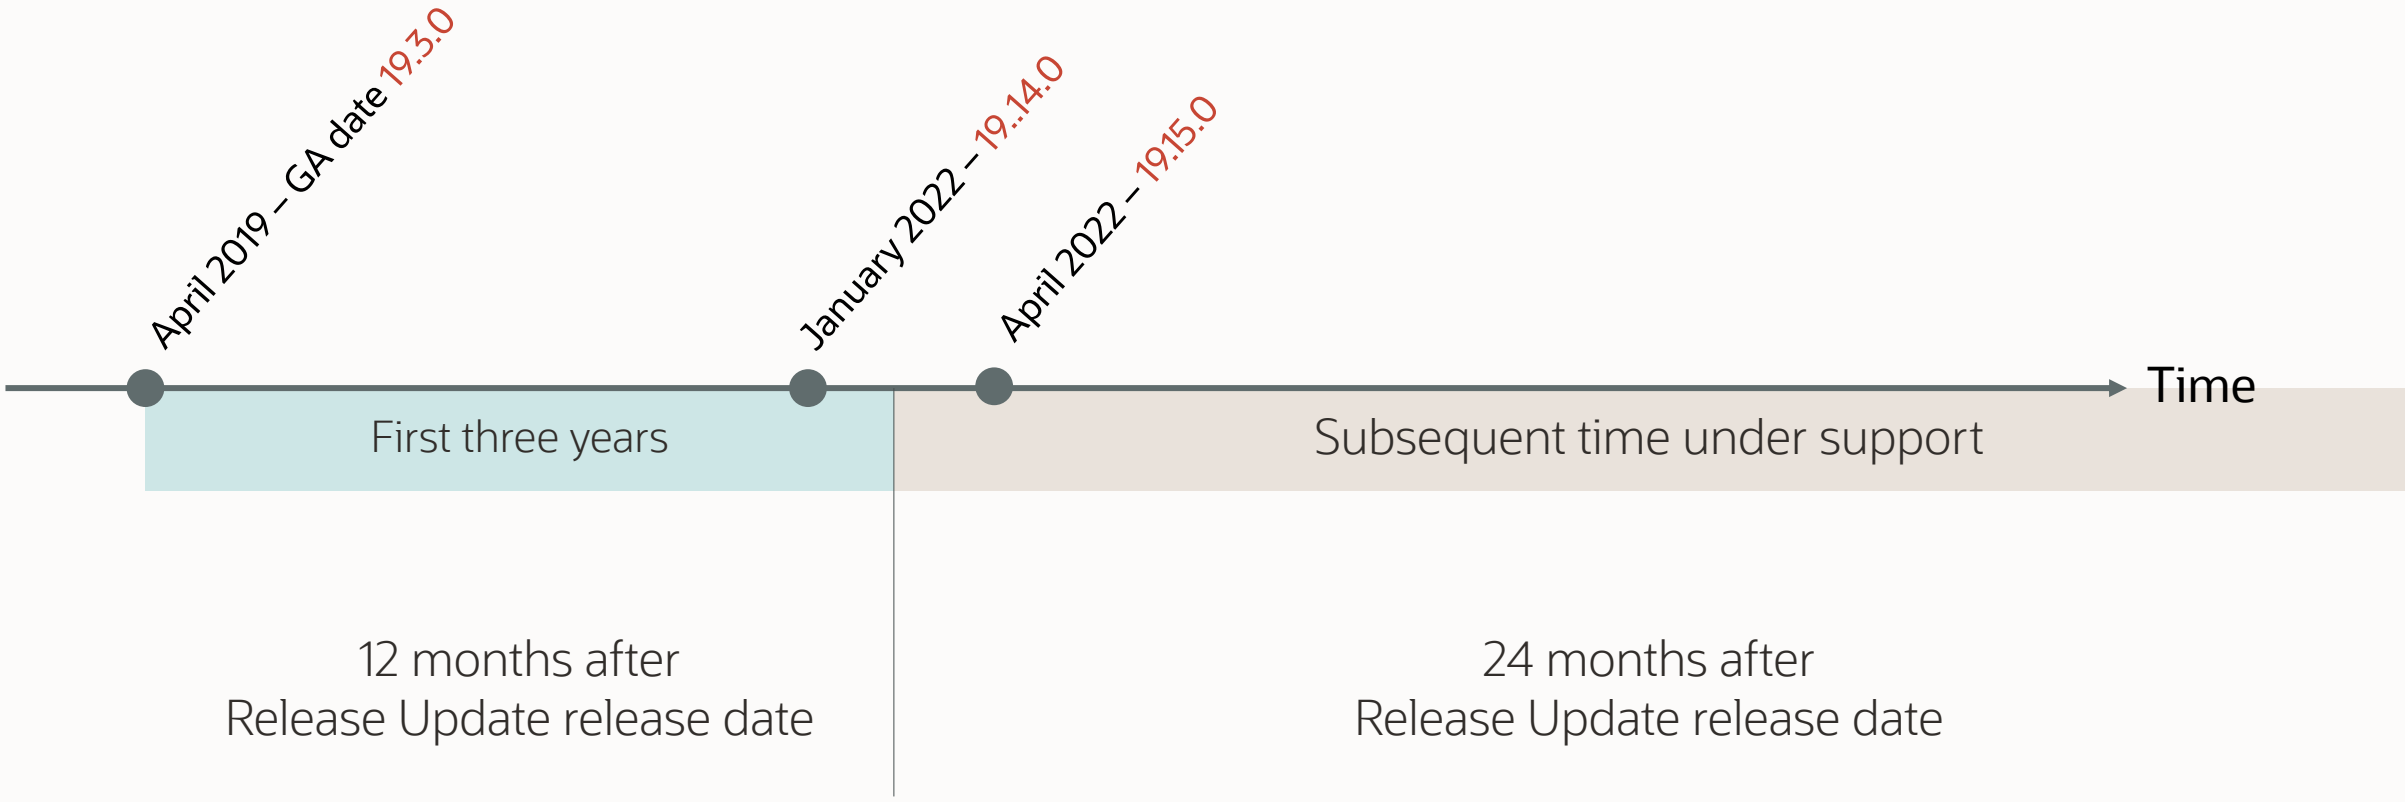

Can I Request a Backport? Examples

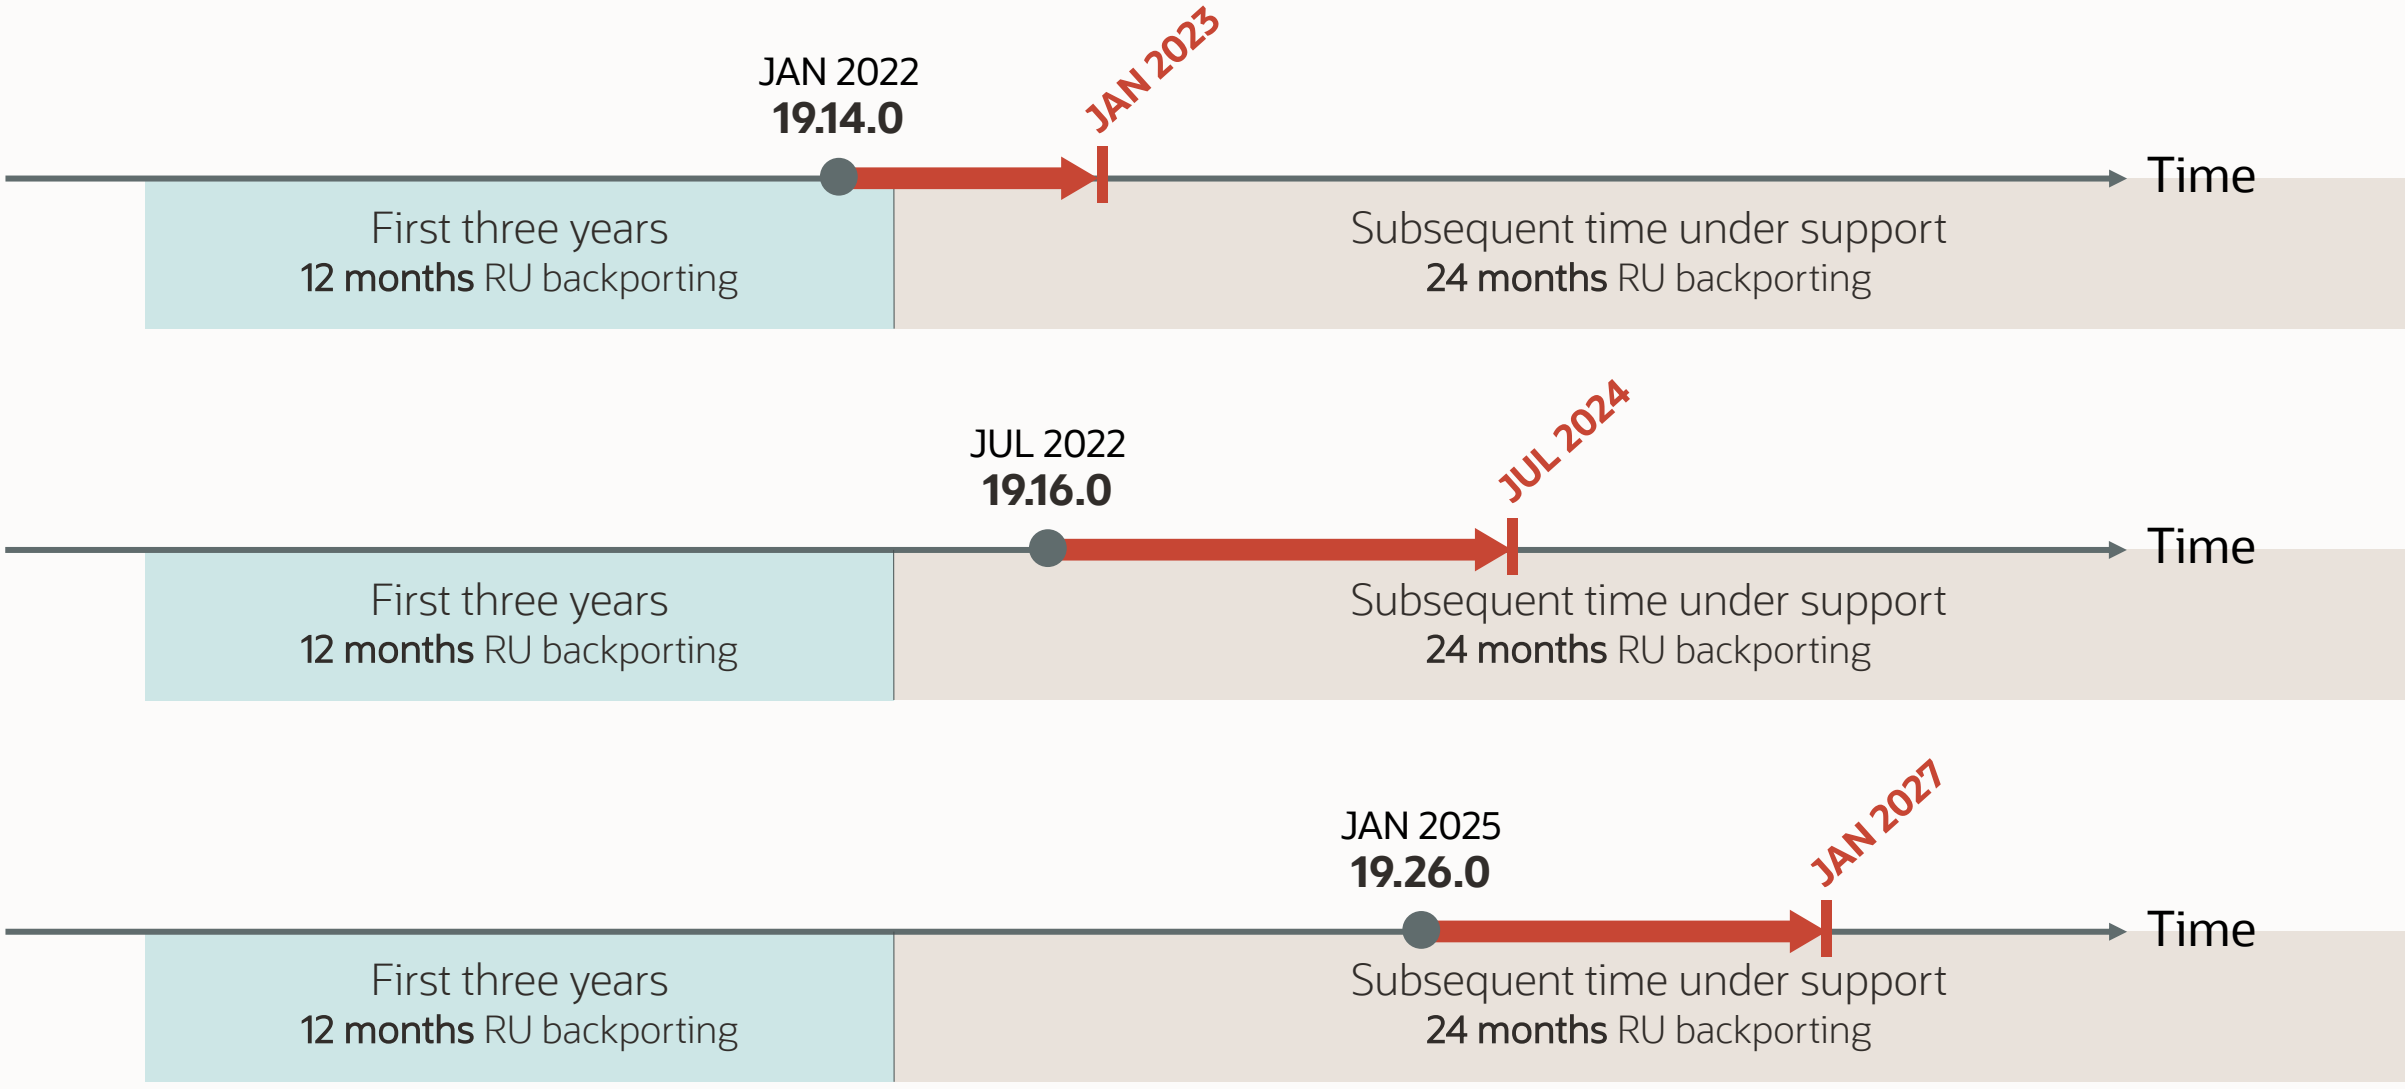

Why is the patch not in the next Release Update?

Insights into the Patching Process

- What goes into a Release Update?
- Screening and monitoring for candidates
- Support can request inclusion via base bug

Request for Inclusion

Recommended Patches for 19.22 DB Home

Below is the list of important patches to consider applying on top of 19.22. 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
36273767 (replaces 35733946)			ORA-1578: oracle data block corrupted on tempfile even after 35904282, 35733946	[list-patches]	YES	YES	30-APR-2024
35286895	19.23		[KPDB] Switchover/Failover Failing for Backup-Based Cadg : ORA-1113: File 3013 Needs Media Recovery	[list-patches]	YES	YES	29-APR-2024
36480774			[RECOVERY] Slow Opening of database in RAC database for other instance	[list-patches]	YES	YES	27-APR-2024
36366069	19.23	DBMRP 19.21.0.0.240319, DBMRP 19.22.0.0.240319	CPU spinning on CTWR and reports ORA-32701 / instance crash post 19.21 DBRU on standby	[list-patches]	YES	YES	28-MAR-2024
35998116 (replaces 35037877)	19.23	DBMRP 19.21.0.0.240319, DBMRP 19.22.0.0.240319	[DBSEC_PRIVS] PLS-00801: internal error [pgm.c:pgmrcm 4] from internal trigger compilation	[list-patches]	YES	YES	27-MAR-2024

Request for Inclusion

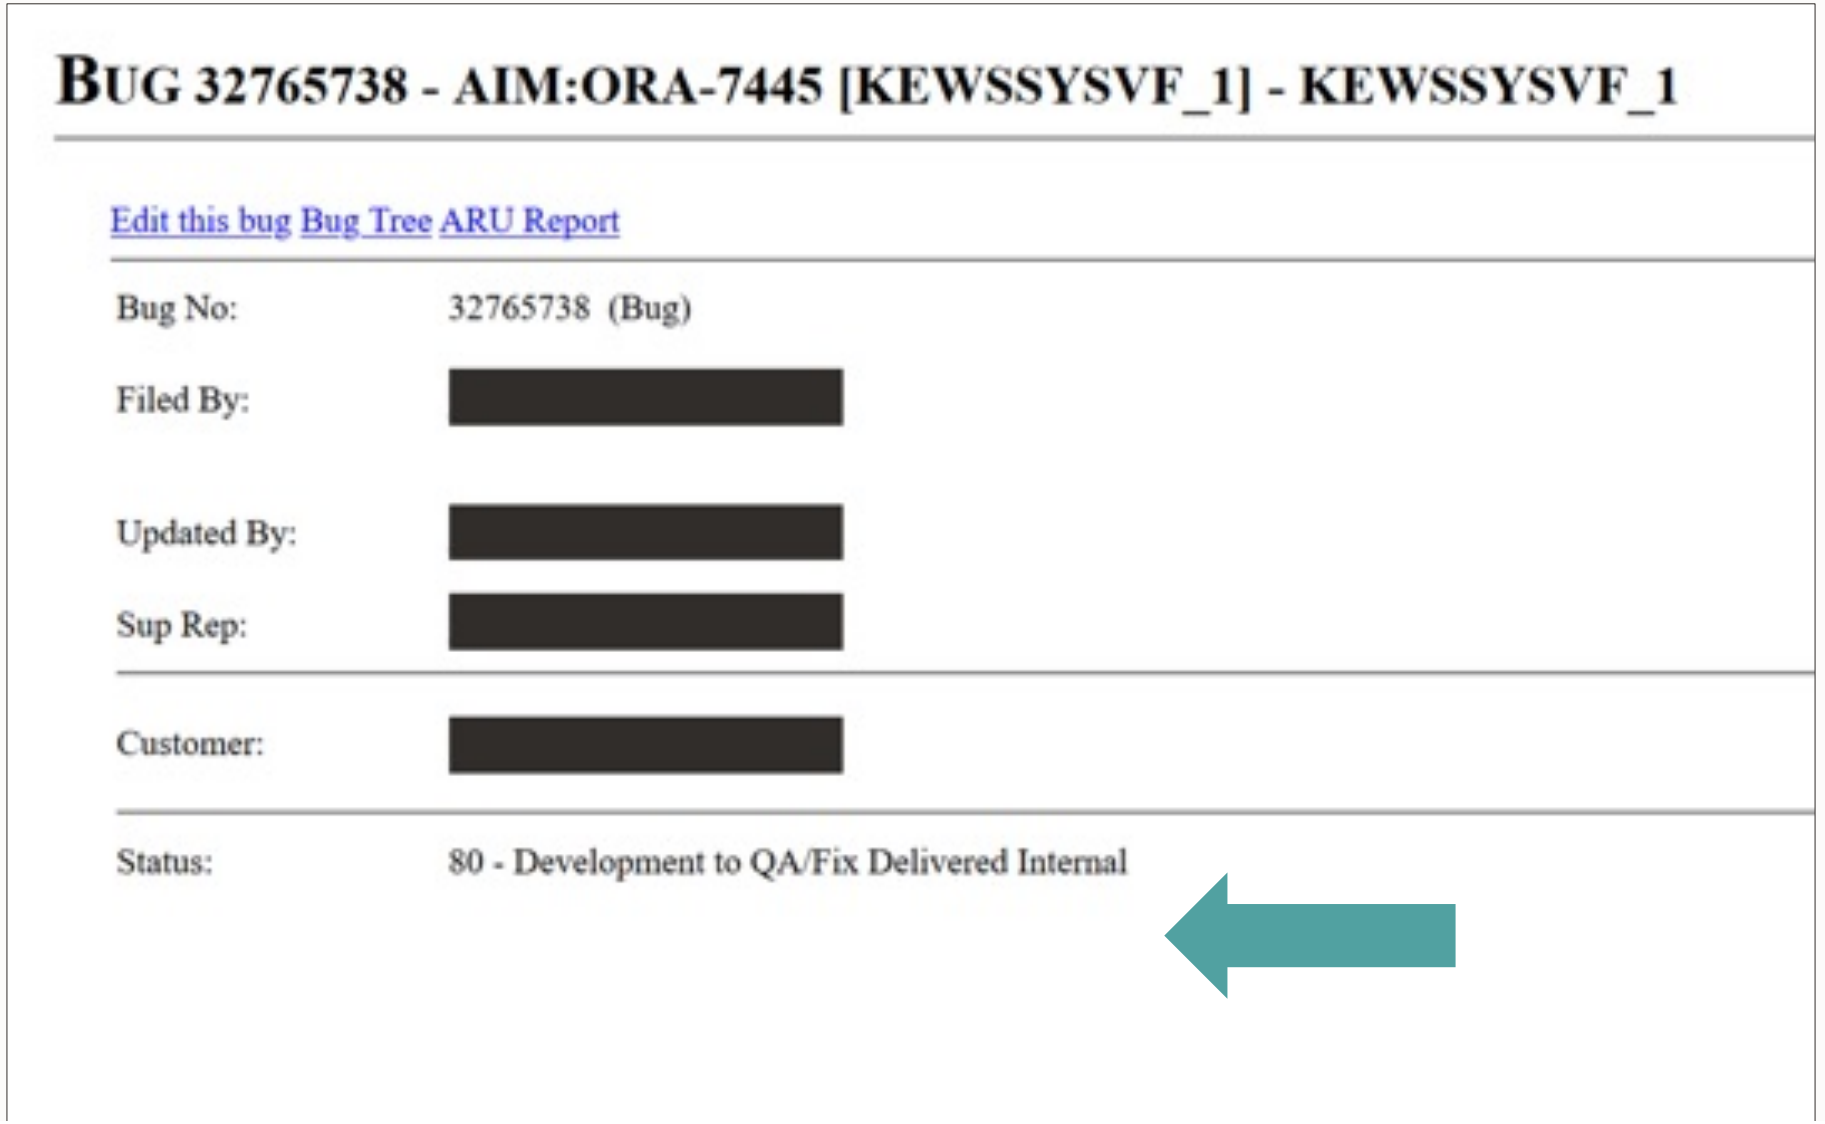

BUG 32765738 - AIM:ORA-7445 [KEWSSYSVF_1] - KEWSSYSVF_1

Edit this bug Bug Tree ARU Report

Bug No:	32765738 (Bug)
Filed By:	[redacted]
Updated By:	[redacted]
Sup Rep:	[redacted]
Customer:	[redacted]
Status:	80 - Development to QA/Fix Delivered Internal

BUG 32765738 - AIM:ORA-7445 [KEWSSYSVF_1] - KEWSSYSVF_1

Edit this bug Bug Tree ARU R

Bug No: 32765738

Filed By:

Updated By:

Sup Rep:

Customer:

Status: 80 - Development to QA/Fix Delivered Internal

SRs in Bug Hierarchy: Open: 12 Others: 274 (Show Details)

Automatic Incident Management

Automatic Incident Management

- Scans for incidents
- Monitors also shared autonomous databases
- Automatically creates bugs
- Report and fix bugs before customers

Ensure your Service Request is associated with the corresponding bug

Insights into the Patching Process

- Evaluated for Release Update inclusion
- Strict requirements apply
  - RAC Rolling
  - Standby-First

Inclusion Evaluation and Criteria

Inclusion Evaluation and Criteria

```
Oracle Database 19 Release 19.12.0.0.210720DBRU

ORACLE DATABASE Patch for Bug# 30978304 for Generic Platforms

This patch is non-RAC Rolling Installable.

This patch is non-Data Guard Standby-First Installable - Please read My Oracle Support Note
1265700.1 https://support.us.oracle.com/oip/faces/secure/km/DocumentDisplay.jspx?id=1265700.1
Oracle Patch Assurance - Data Guard Standby-First Patch Apply for details on how to remove risk
and reduce downtime when applying this patch.

(1) Prerequisites
--------------------
Before you install or deinstall the patch, ensure that you meet the following requirements:
```

NEVER INCLUDED

Insights into the Patching Process

- Limited number of fixes per Release Update
- Release Update regression tests

Release Update Contents

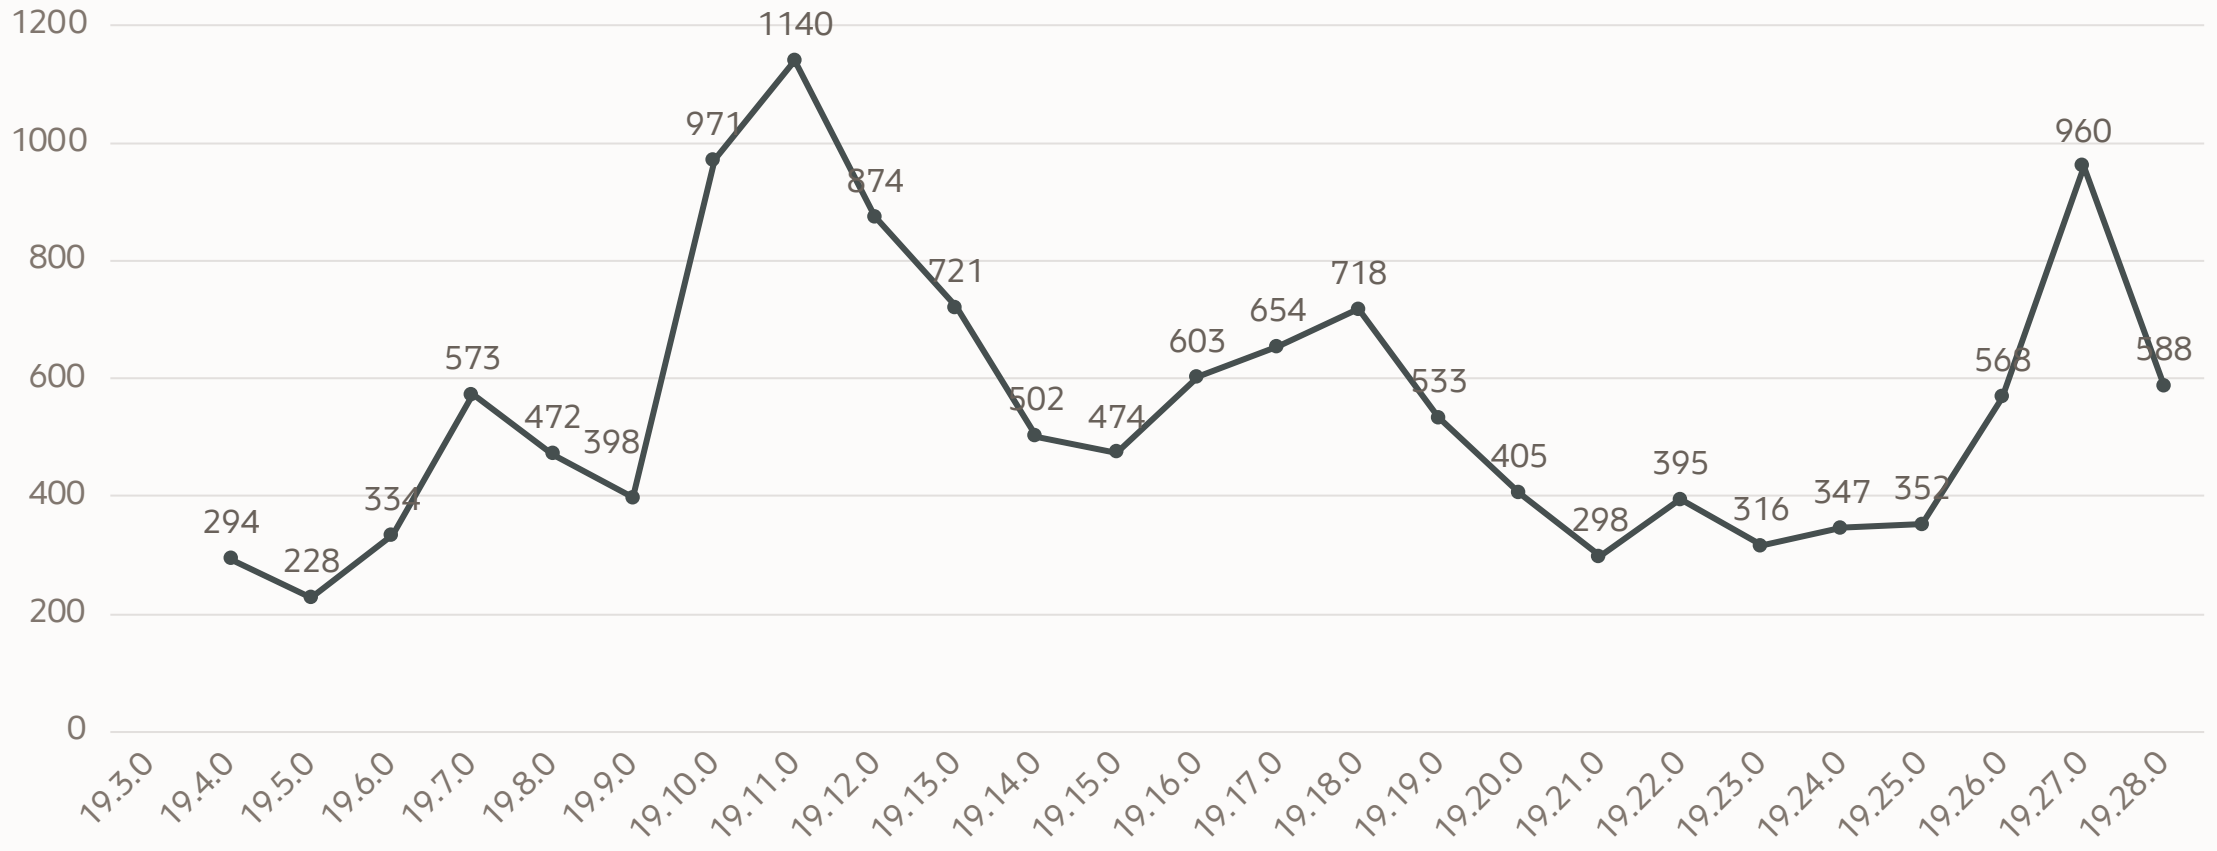

Database 19 Release Updates and Revisions Bugs Fixed Lists (Doc ID 2523220.1)

Limitation

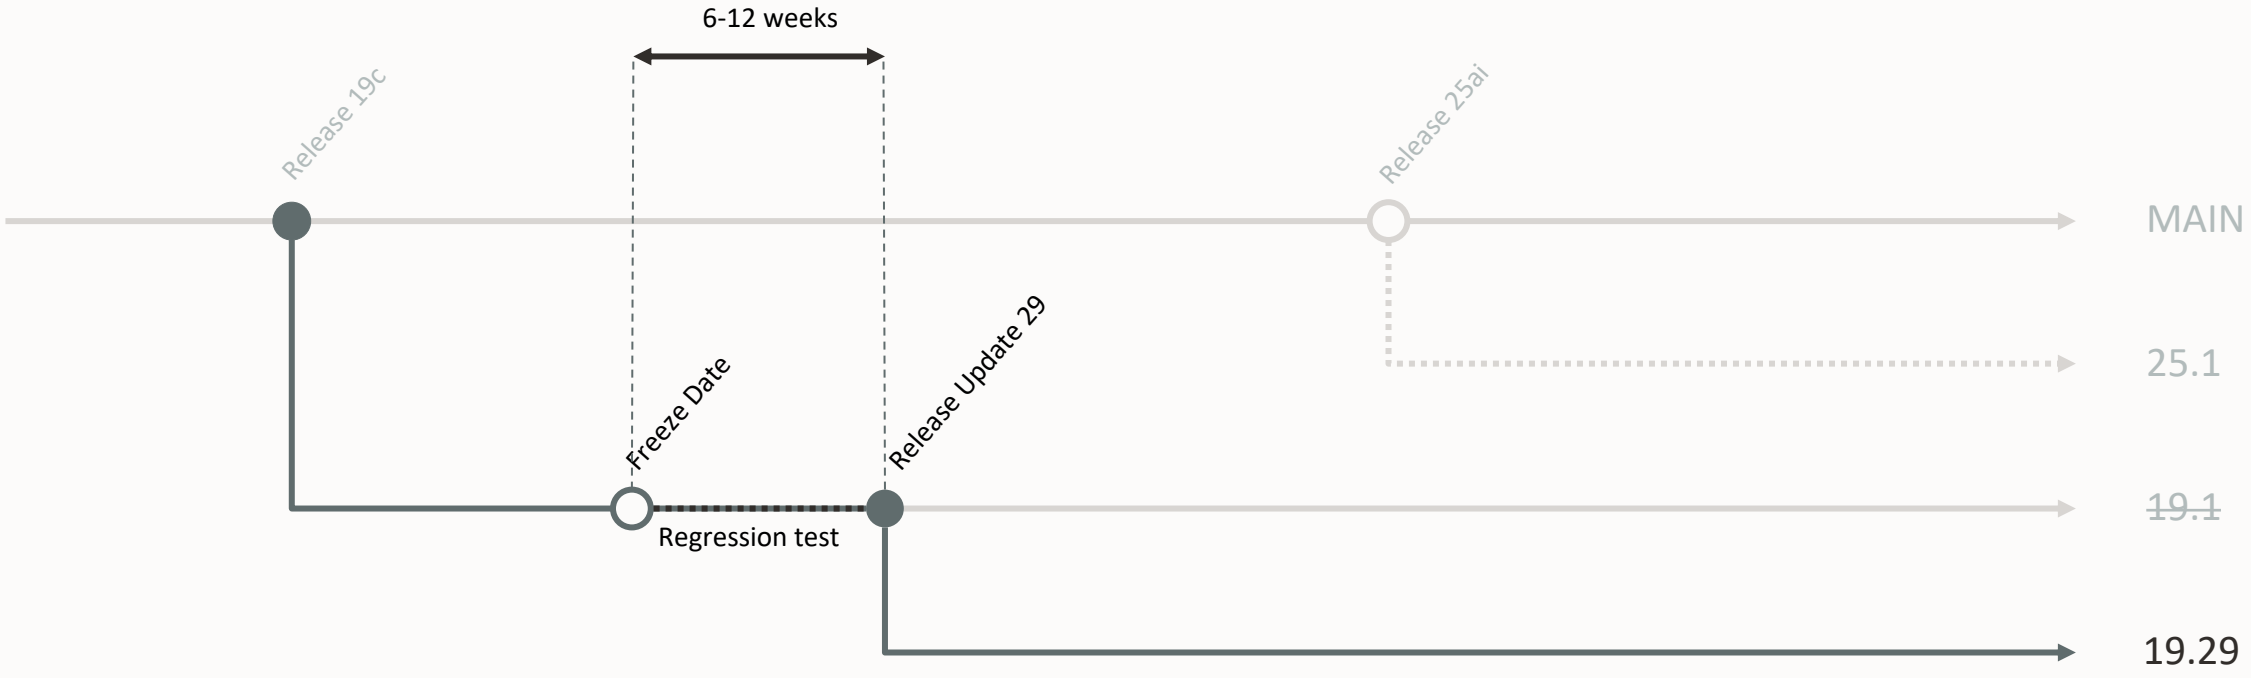
6-12 weeks
Release 19c
Release 25ai
MAIN
25.1
Freeze Date
Release Update 29
19.1
Regression test
19.29

WHY DO YOU ALWAYS ASK FOR

OPATCH lsinventory

Photo by Joel J. Martínez on Unsplash

TO AVOID CONFLICTS

Photo by Glomad Marketing on Unsplash

```
$ORACLE_HOME/OPatch/opatch prereq CheckConflictAgainstOHWithDetail -ph ./

Invoking prereq "checkconflictagainstohwithdetail"
ZOP-40: The patch(es) has conflicts with other patches installed in the
Oracle Home (or) among themselves.
Prereq "checkConflictAgainstOHWithDetail" failed.
Summary of Conflict Analysis:
There are no patches that can be applied now.
Following patches have conflicts. Please contact Oracle Support and get
the merged patch of the patches :
35012562, 35095748
Conflicts/Supersets for each patch are:
Patch : 35095748
Conflict with 35012562
Conflict details:
/u01/app/oracle/product/19.19.0/db_1/lib/libserver19.a:kko.o

OPatch succeeded.
```

Basic Facts | Conflicts

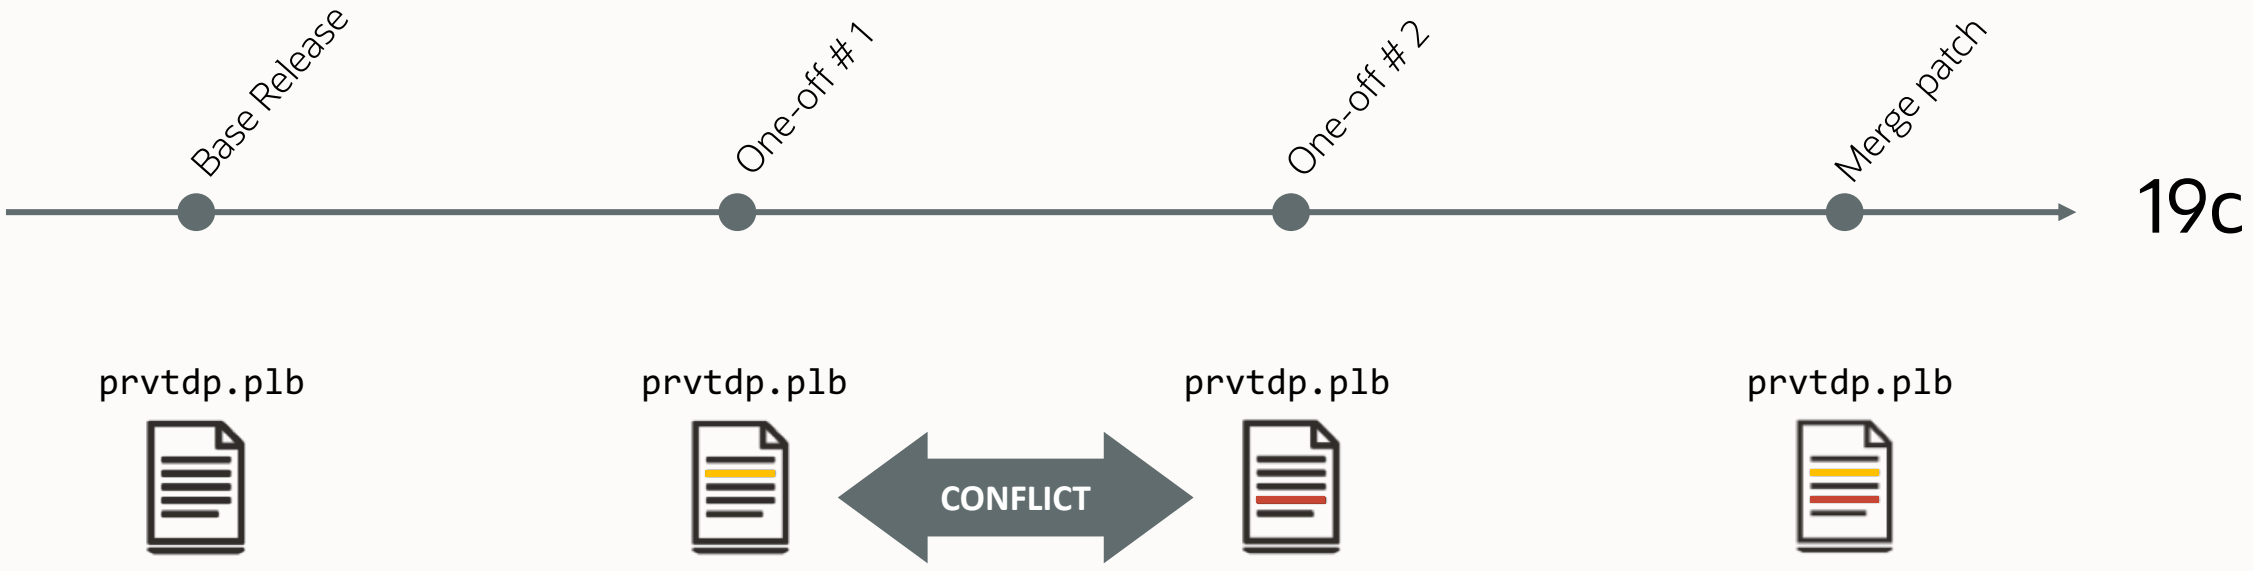

On your Service Request, always attach the output of `opatch lsinventory`

Don't name it

- `optch_lsinvt.txt`
- `alle_mine_rettelser.txt`
- `alle_meine_patches.txt`
- `všechny_mé_záplaty.txt`

Just name it `opatch_lsinventory.txt`

But … I am on MS Windows …

Photo by Sebastian Herrmann on Unsplash

Windows is different …

Photo by Clint Patterson on Unsplash

Windows Bundle Patches is the term used for Release Updates on Windows

A Windows Bundle Patch and a Release Update from the same quarter, e.g., April 2023, do not contain the same patches

Comparing Release Updates and Bundle Patches

19.19.0 Linux Release Update

- 533 bug fixes
- 5 were not in the Windows bundle patch

19.19.0 Windows Bundle Patch

- 153 additional bug fixes
- 528 bug fixes in common

Where do I find the Windows Bundle Patches?

Assistant: Download Reference for Oracle Database/GI Update, Revision, PSU, SPU(CPU), Bundle Patches, Patchsets and Base Releases (Doc ID 2118136.2)

Visibility: EXTERNAL (94)

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 - discontinued since Apr 2023)
- ○ Oracle Database PSU, SPU(CPU), Bundle Patches (Versions 12.1 & lower)
- ○ OJVM Update/PSU/Bundle Patches
- ◯ Latest Available Microsoft Windows Patches
- ○ Monthly Recommended Patches (MRPs)

Solution(s)

Possible Solutions w
you make your selec

Basic Facts | Patch Availability

MOS Note: 2118136.2 - Assistant: Download Reference for DB and GI Patch Bundles

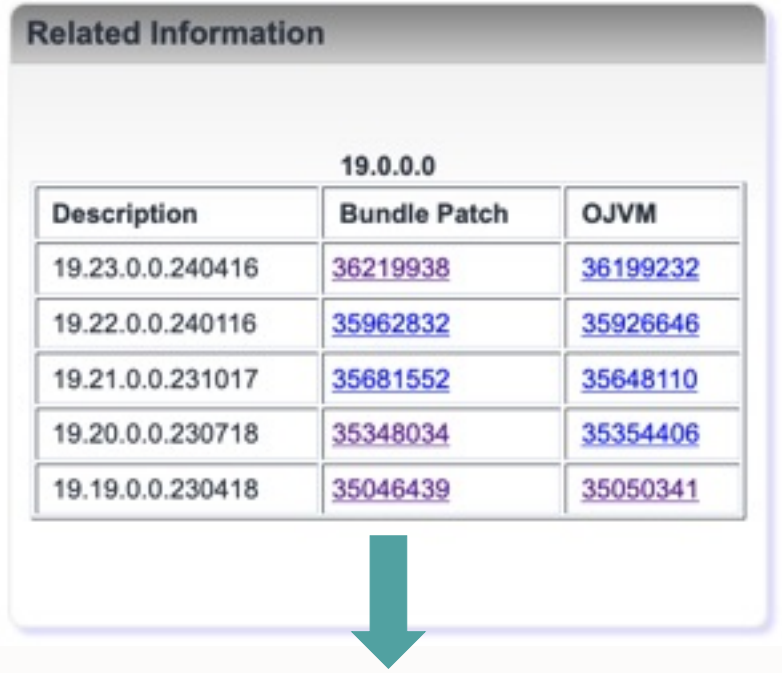

Patch Details

Details for Patch 36219938 not found.

Basic Facts | Patch Availability

Critical Patch Update Program Apr 2024 Patch Availability Document (DB-only)

My Oracle Support Note 3000005.1

Released April 16, 2024

This document contains the following sections:

- Critical Patch Update April 2024 Patch Availability Document (PAD)
  - 1 Overview
    - 1.1 How To Use This Document
    - 1.2 Terminology in the Tables
    - 1.3 On-Request Patches
    - 1.4 CPU Program and My Oracle Support Patch Recommendations
    - 1.5 My Oracle Support (MOS) Conflict Checker Tool
  - 2 What's New in April 2024
    - 2.1 "Final CPU Information (Error Correction Policies)"
    - 2.2 "Post Release Patches"
    - [illegible] Products"

Basic Facts | Patch Availability

2.2 Post Release Patches

Oracle strives to complete preparations and testing of each Quarterly Security Patch for each platform by the quarterly release date. Occasionally, circumstances beyond Oracle's control can require that a particular patch must be released a few days after the quarterly release date. The following table lists any current patch delays and the estimated date of availability.

Patch	Patch Number	Platform	Availability
21.14.0.0.240416 DB RU	Patch 36352352	Linux x86-64	Available
		HP-UX Itanium	07-May
21.14.0.0.240416 GI RU	Patch 36352207	Linux x86-64	Available
		HP-UX Itanium	07-May
21.14.0.0.240416 WIN BP	Patch 36219877	WINDOWS x64, NT	Available
19.23.0.0.240416 DB RU (& assoc. COMBO)	Patch 36233263 (& Patch 36209492)	Linux x86-64, Solaris.x64, Solaris SPARC 64-Bit, and AIX.PPC64	Available
		HP-UX Itanium, LINUX.ZSERIES64, and LINUX.ARM64	07-May
19.23.0.0.240416 GI RU (& assoc. COMBO)	Patch 36233126 (& Patch 36209493)	Linux x86-64, Solaris.x64, Solaris SPARC 64-Bit, and AIX.PPC64	Available
		HP-UX Itanium, LINUX.ZSERIES64, and LINUX.ARM64	07-May
19.23.0.0.240416 WIN BP	Patch 36219938	WINDOWS.X64	23-Apr
		WINDOWS NT	07-May
19.23.0.0.240416 OJVM	Patch 36199232	All except Linux x86-64	07-May

… Oracle on Windows is different

Photo by Clint Patterson on Unsplash

Oracle Linux Development Platform

Photo by Danielle Barnes on Unsplash

Porting

Oracle Linux x86-64

Porting

- Microsoft Windows
- Linux ARM (aarch64)
- SPARC Solaris
- Intel Solaris x86-64
- IBM AIX
- IBM zLinux
- HP UX Itanium
- BS2000

Some of the platforms from our bug tracking system

Blast from the past

Blast from the Past | Platforms

Microsoft Windows Phone
Oracle JRockit Virtual Edition x86
SunOS
Monta Vista x86
Acme Packet 1100
iTron
Embedded Linux on cnMIPS
Embedded Linux SH4
HP NonStop Itanium (OSS)
QNX Unix
Acme Packet 6100
Linux MIPS 64-bit
Fujitsu BS2000/OSD (SQ series)
Mediatek MTZ
HP NonStop (Guardian) on x86
Fujitsu BS2000
HP Tru64 UNIX

Tekelec
Qualcomm Brew MP
Netra Server X5-2 for Communications
HP NonStop S-series (Guardian)
HP OpenVMS Itanium
Monta Vista x86-64
OpenSolaris
SCO Unix
Net-Net 9200
Symbian EPOC
Linux ARM 32-bit VFP HardFP ABI
SGI Irix
ia64
Linux SPARC
Oracle Solaris on SPARC (32-bit)
HP NonStop Itanium (Guardian)
RIM BlackBerry

Netra X3-2 for Acme Packet
Oracle Solaris on SPARC (64-bit)
IBM S/390 Based Linux (31-bit)
Acme Packet 3900
SPARC
Fujitsu MSP-EX
Trusted Solaris
Net-Net 4250
HP OpenVMS VAX
HP-UX PA-RISC (32-bit)
Acme Packet 6300
Microsoft Windows CE
IBM z/OS on System z
StorageTek Hardware
Oracle Solaris on x86 (32-bit)
Fujitsu BS2000/OSD (SX series)
Linux ARM 64-bit

Novell NetWare
Linux on IBM Z
Data General
Pyramid
Talari
Palm Computing
HP NonStop (OSS) on x86
Unisys OS 2200
HP OpenVMS Alpha
Acme Packet 3820
FreeBSDx86
Oracle Solaris Express
VxWorks
Microsoft Windows (32-bit)
Sequent
Windows NT
nCube

Key Learnings

1. Use Data Pump Bundle Patch
2. Save workload into SQL Tuning Sets
3. Use SQL Plan Management

Break

We start again at 15:15

Autonomous Database

A migration approach

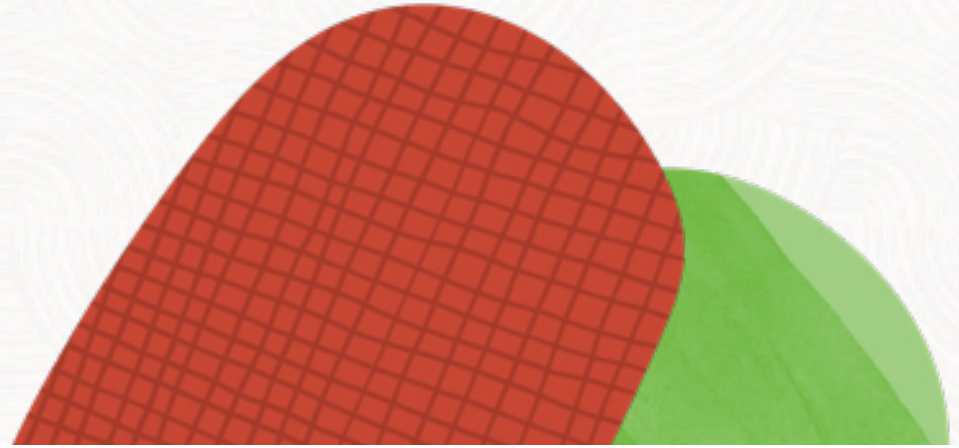

What's the story?

Flashback to October 2017

The idea?

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

KEYNOTE PRESENTATION

Oracle OpenWorld San Francisco 2017


OOW 2017
ORACLE

Will the DBA's be fired?

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

Photo by LARAM on Unsplash

Autonomous Database – Where?

Public cloud

Autonomous Database

Cloud@Customer

Autonomous Database in a VM environment

One Autonomous Database – Two Deployment Choices

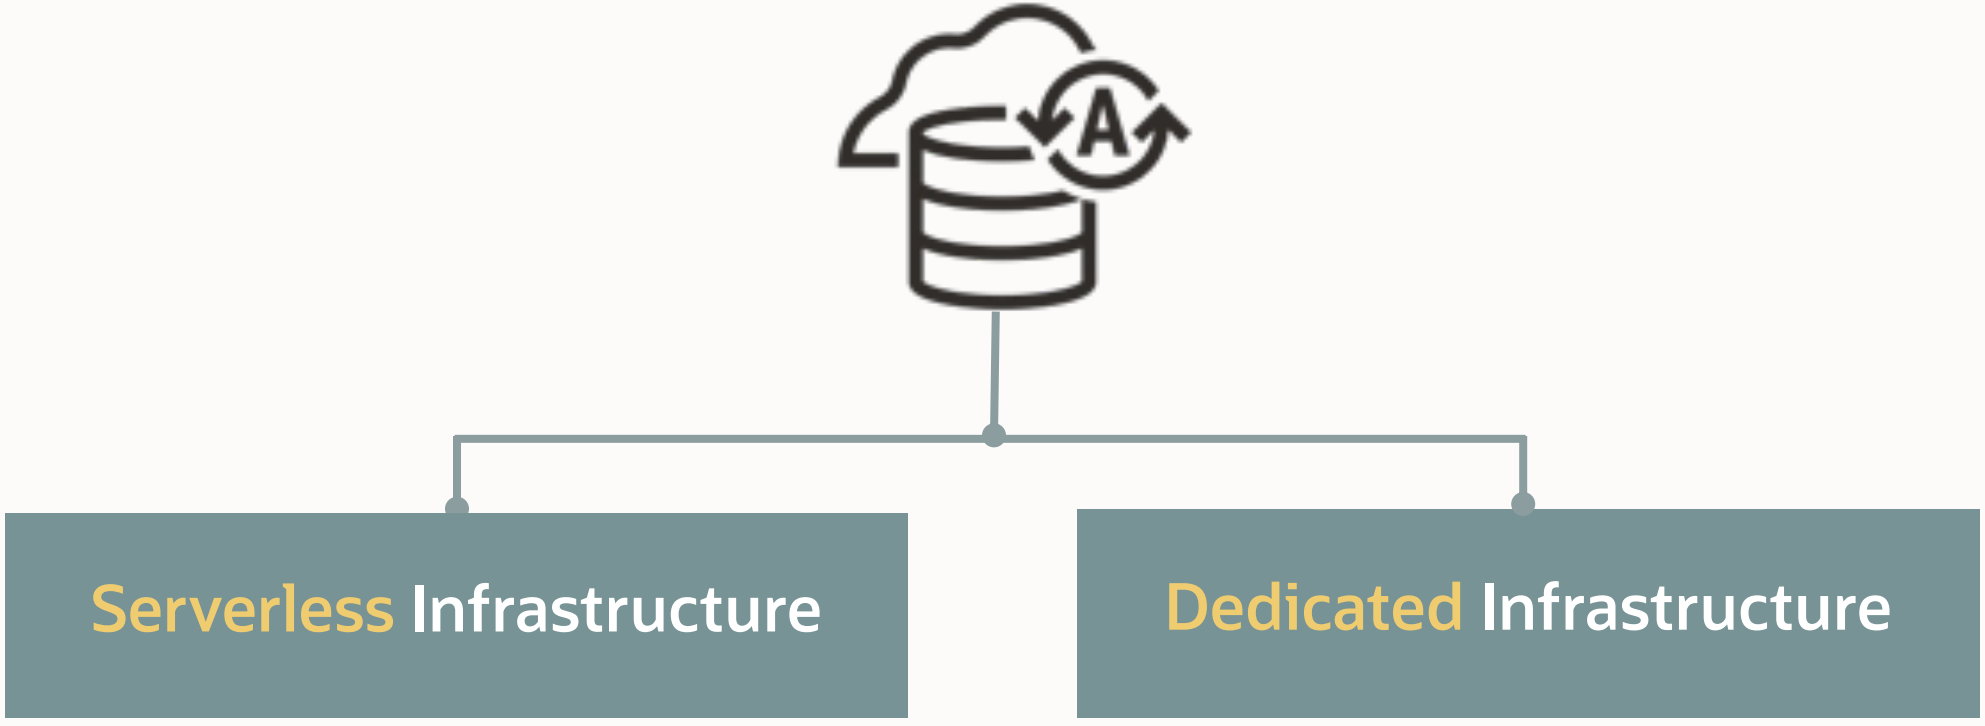

One Autonomous Database – Workload Choices

Autonomous Data Warehouse (ADW)	Autonomous Transaction Processing (ATP)
• Data Warehouse, Data Mart • Data Lake, Machine Learning	• Transactions, Batch, Reporting, IoT • Application Dev, Machine Learning

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 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

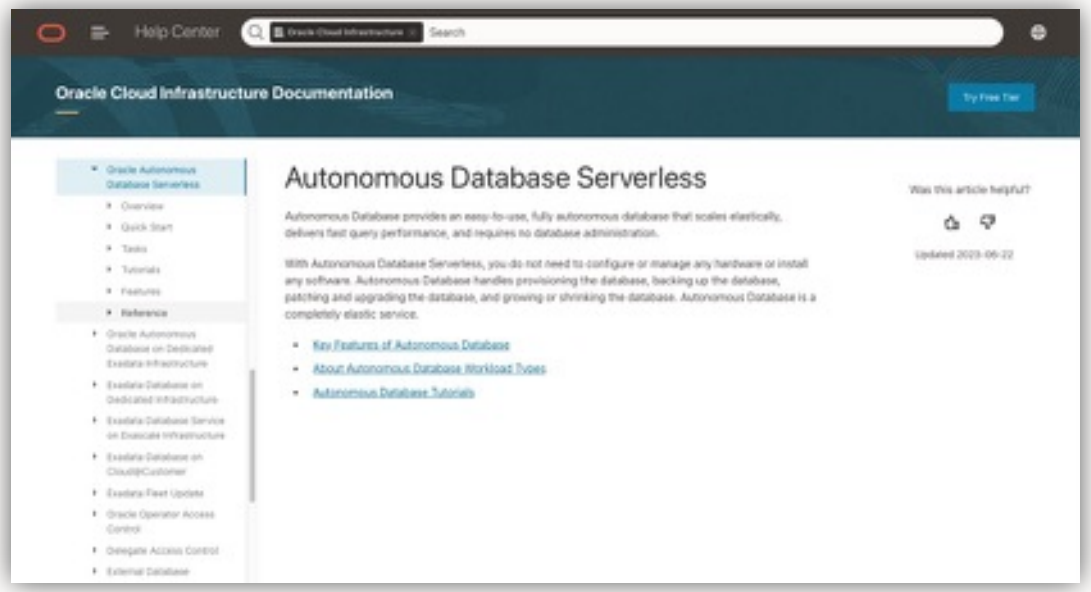

Outside the database,
check OCI Documentation

Example: Deploy, start, stop, scale

Essentials

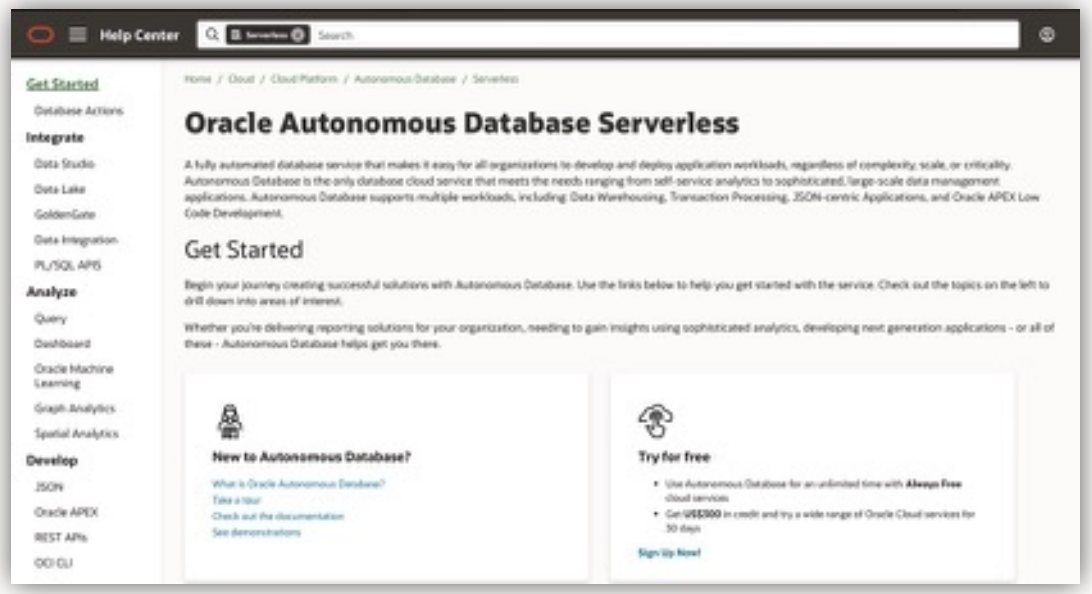

Inside the database,
check Database Documentation

Example: Schema, capabilities, connecting

Planning

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

- And which ones are good candidates?

Getting an Overview

1 Estate Explorer

2 Cloud Premigration Advisor Tool

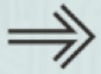

3 Cloud Migration Advisor

Getting an Overview

1 Estate Explorer

2 Cloud Premigration Advisor Tool

3 Cloud Migration Advisor

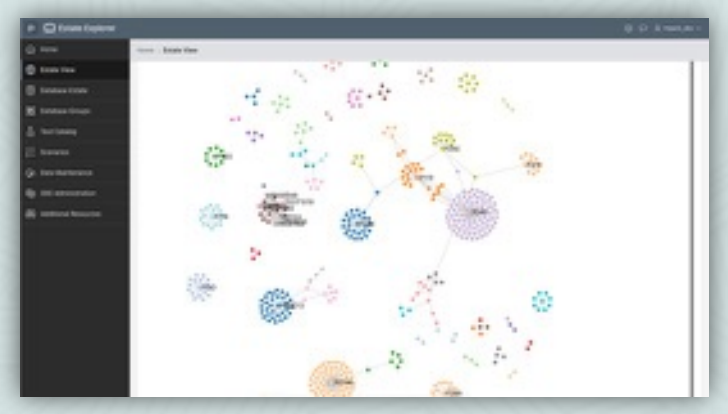

Getting an Overview

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 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

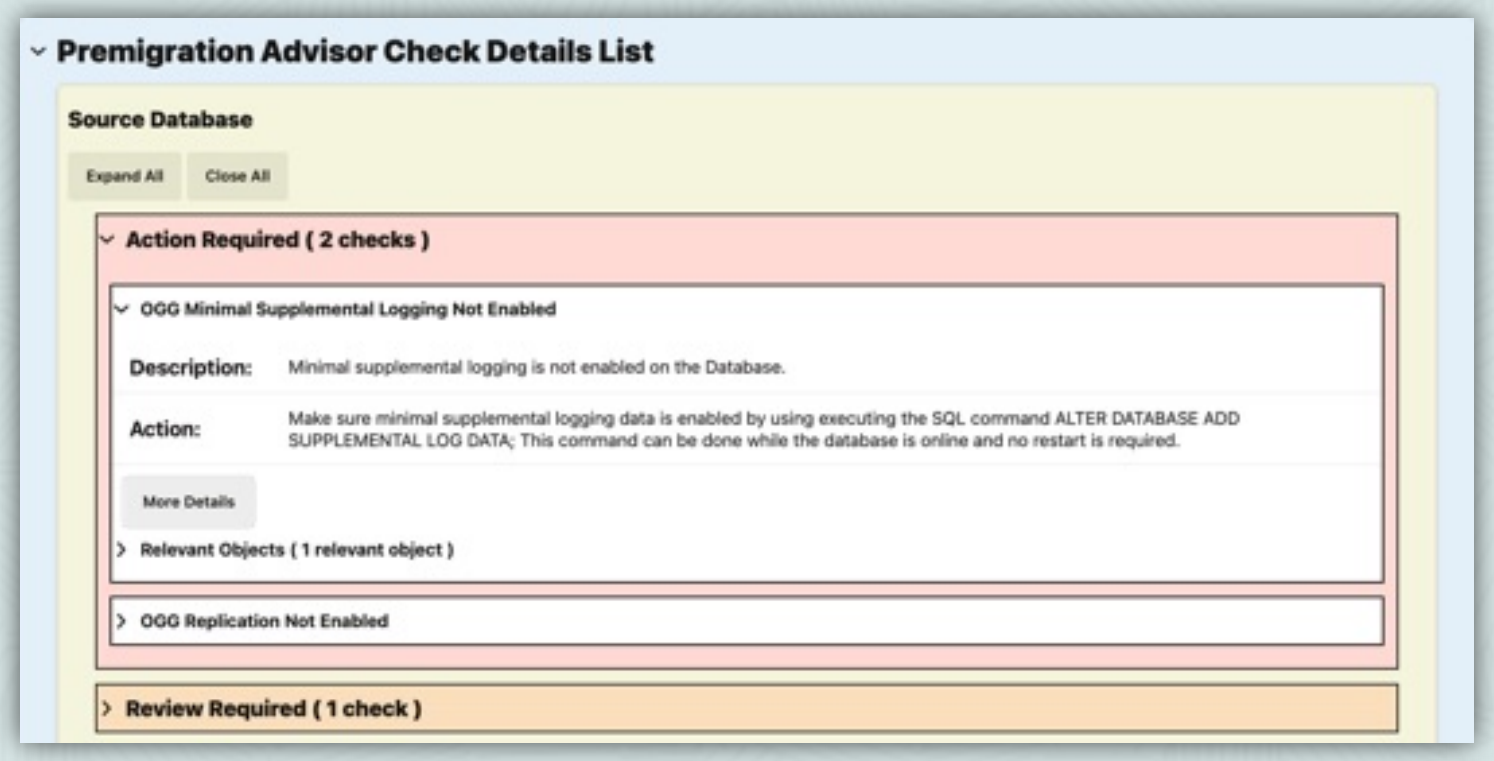

Getting an Overview

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

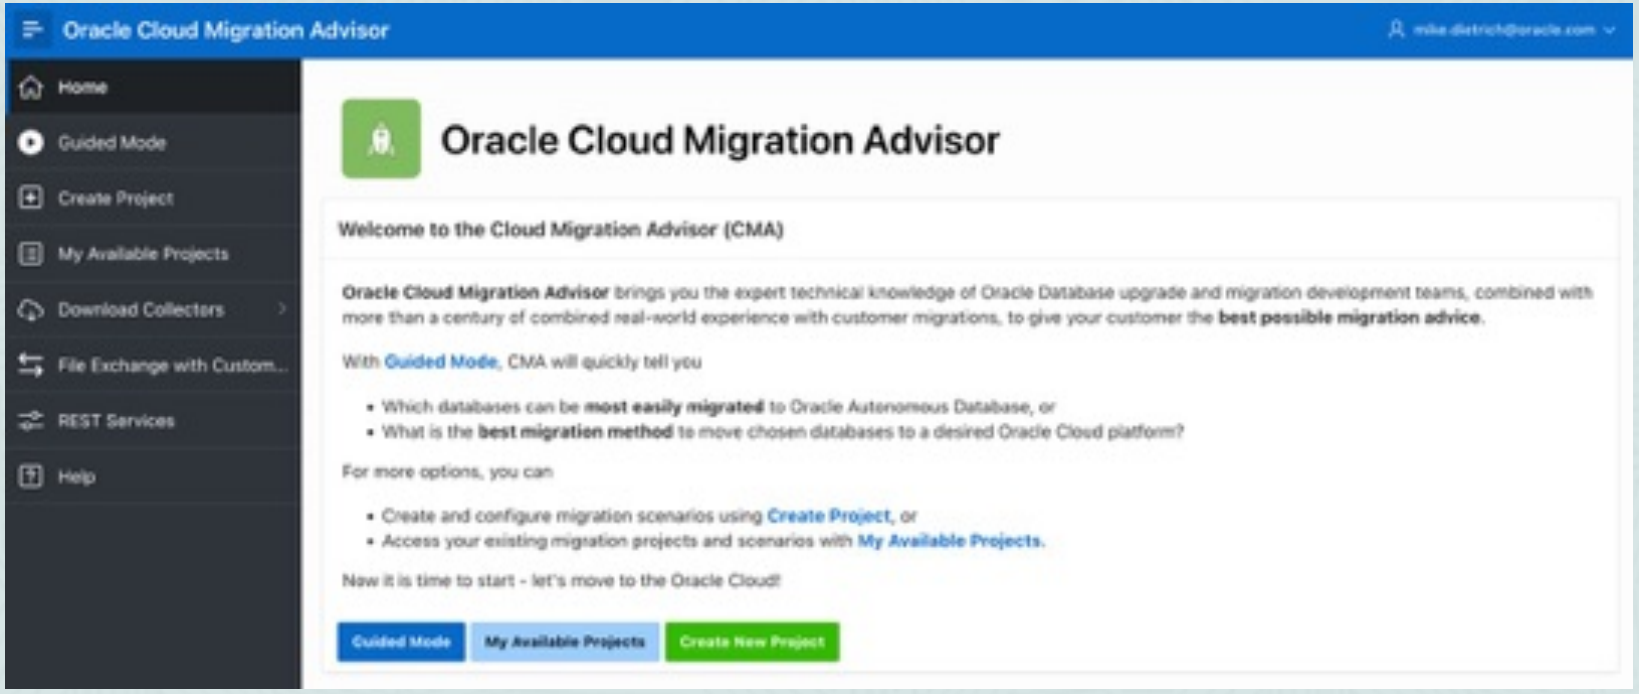

Getting an Overview

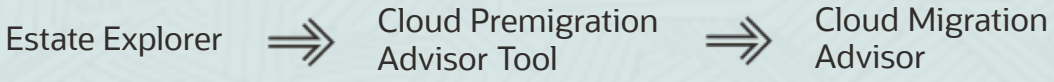

- OCI Database Migration Service
- Zero Downtime Migration
- Autonomous Migration Automation
- Data Pump
- GoldenGate
- O2O / OOO

Cloud Premigration Advisor Tool

Evaluate an Oracle Database for compatibility with Autonomous Database

- Use Cloud Premigration Advisor Tool (CPAT)

Overview

Connects → Checks → Reports → Fixes (optional)

Download CPAT from MOS Note: 2758371.1

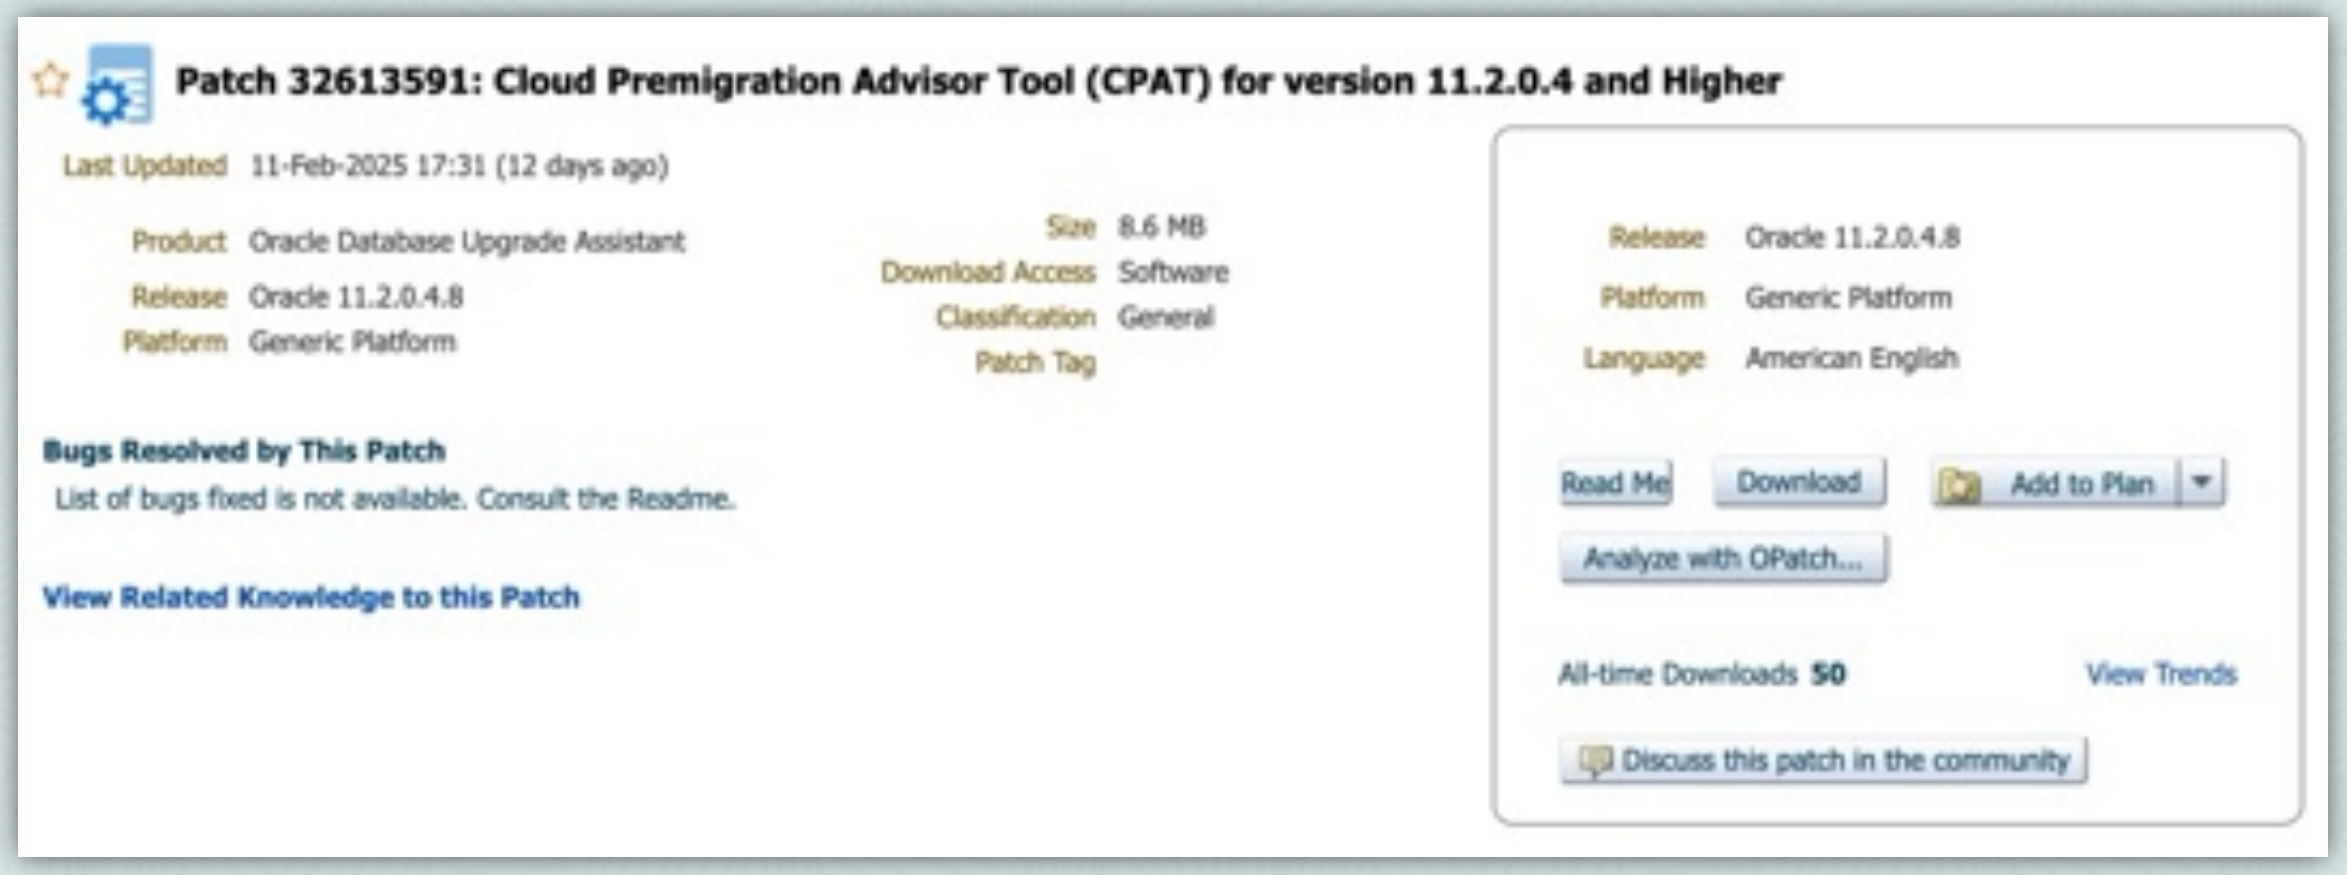

```
# One or more report formats separated by spaces
# json html text

./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

macsdata articles tools tags about

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
```

ORACLE

[illegible]

Cloud Premigration Advisor Tool (CPAT) Report

CPAT Version: [illegible]
Version Date: [illegible]
Days Since Last CPAT Update: [illegible]

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	[illegible]
[illegible]	[illegible]

- Premigration Advisor Report Summary
- Premigration Advisor Check Details List
- Report Legend

Expand All | Close All

Premigration Advisor Report Summary

Report Result	Action Required
Number of schemas analyzed	1
List of schemas analyzed	[redacted]

Report Results Overview

Source Database		Target Database		Migration Method		Additional Tasks	
Action Required	1	Action Required	10	Action Required	11	Action Required	1
Review Required	1	Review Required	11	Review Required	12	Review Required	1
Review Suggested	1	Review Suggested	4	Review Suggested	11	Review Suggested	4
Passed	10	Passed	17	Passed	8	Passed	10

Return to Table of Contents

Report Details

Return to Table of Contents

Report Details

[illegible]	[illegible]
[illegible]	[illegible]
[illegible]	[illegible]
[illegible]	[illegible]
[illegible]	[illegible]
[illegible]	[illegible]

[illegible]

Return to Table of Contents

Report Analysis Notes

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

Preupgrade Advisor Check Details List

Source Database

Source Database Details

Source Cloud Vendor:	Oracle Cloud Infrastructure (Database)
Source Database Host Name:	dbsystemaz
Source Oracle SID:	ORCL
Source Database Created Date:	Fri Jan 24 22:23:51 UTC 2025
Source Database DBID:	1719058167
Source Database Unique Name:	ORCL_Str_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 Database 23ai

```
--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 *(June 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 Database Migration Automation (AMA)

- Simple migration solution for ADB Serverless
- Script based
- Single configuration file
- Migrates in phases
- Can act fully automated

- Not a new product, just a solution to ease migrations

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

You won't start with filming right away

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

Now you can start filming your scenes

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

AMA Workflow

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

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

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

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

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


Parameter File

Adjust;

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

Documentation

AMA Documentation is available at request

- Documents the entire flow and all options and parameters

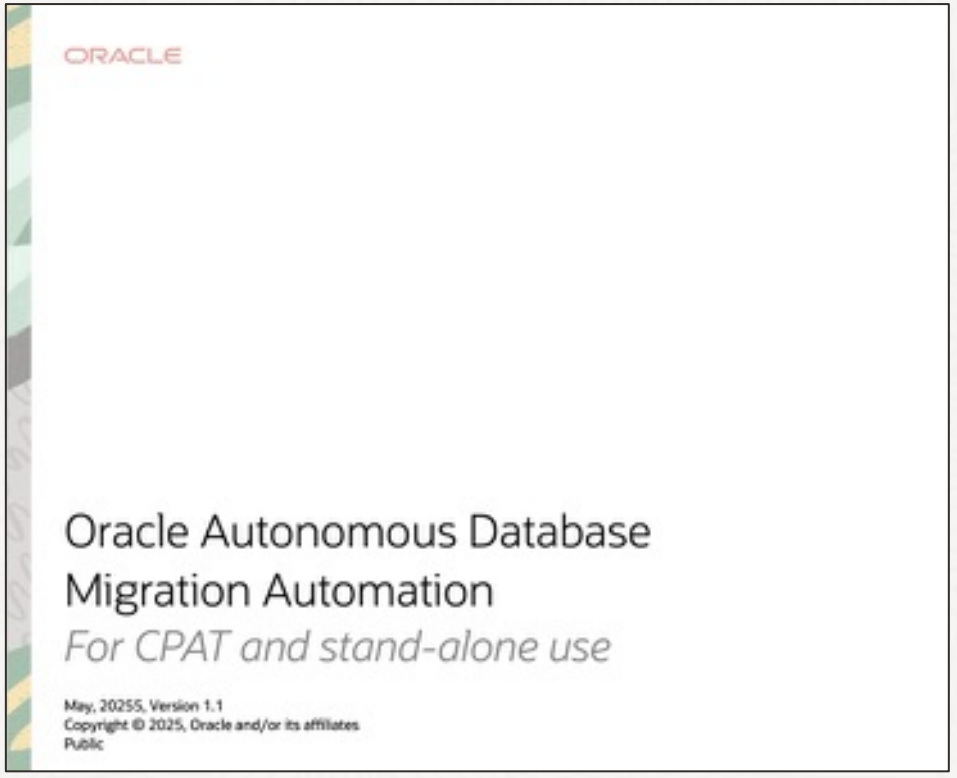

AMA Workflow

- Java and OLAP (ADB)
- Migration user
- Statistics scripts
- Quiesce scripts

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.ctl

---
---
--- PLAN PHASE ---
--- All steps in this phase affect the source database ---
--- * You can collect the statistics or create the migration user in advance  ---
--- * Get familiar with the restricted session privilege and how to prepare it  ---
--- * shortly before the migration starts make sure no unwanted user is connected  ---
---   to the source database, turn on restricted session and disable the scheduler ---
---
---

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

PLAN TARGET 00001 01 ..................................................................... 00001_US3BLDW_SQL_ENABLE_OLAP_JAVA.sh
PLAN SOURCE 00002 01 00002_US3BLDW_SQL_OPTIONAL_SOURCE_STATS.sh
PLAN SOURCE 00003 01 00003_US3BLDW_SQL_CREATE_MIG_USER_SRC.sh
PLAN SOURCE 00004 01 00004_US3BLDW_SQL_SET_JOB_QUEUE_PROCESSES.sh
PLAN SOURCE 00005 01 00005_US3BLDW_SQL_ENABLE_RESTRICTED_SESSION.sh
---
```

AMA Workflow

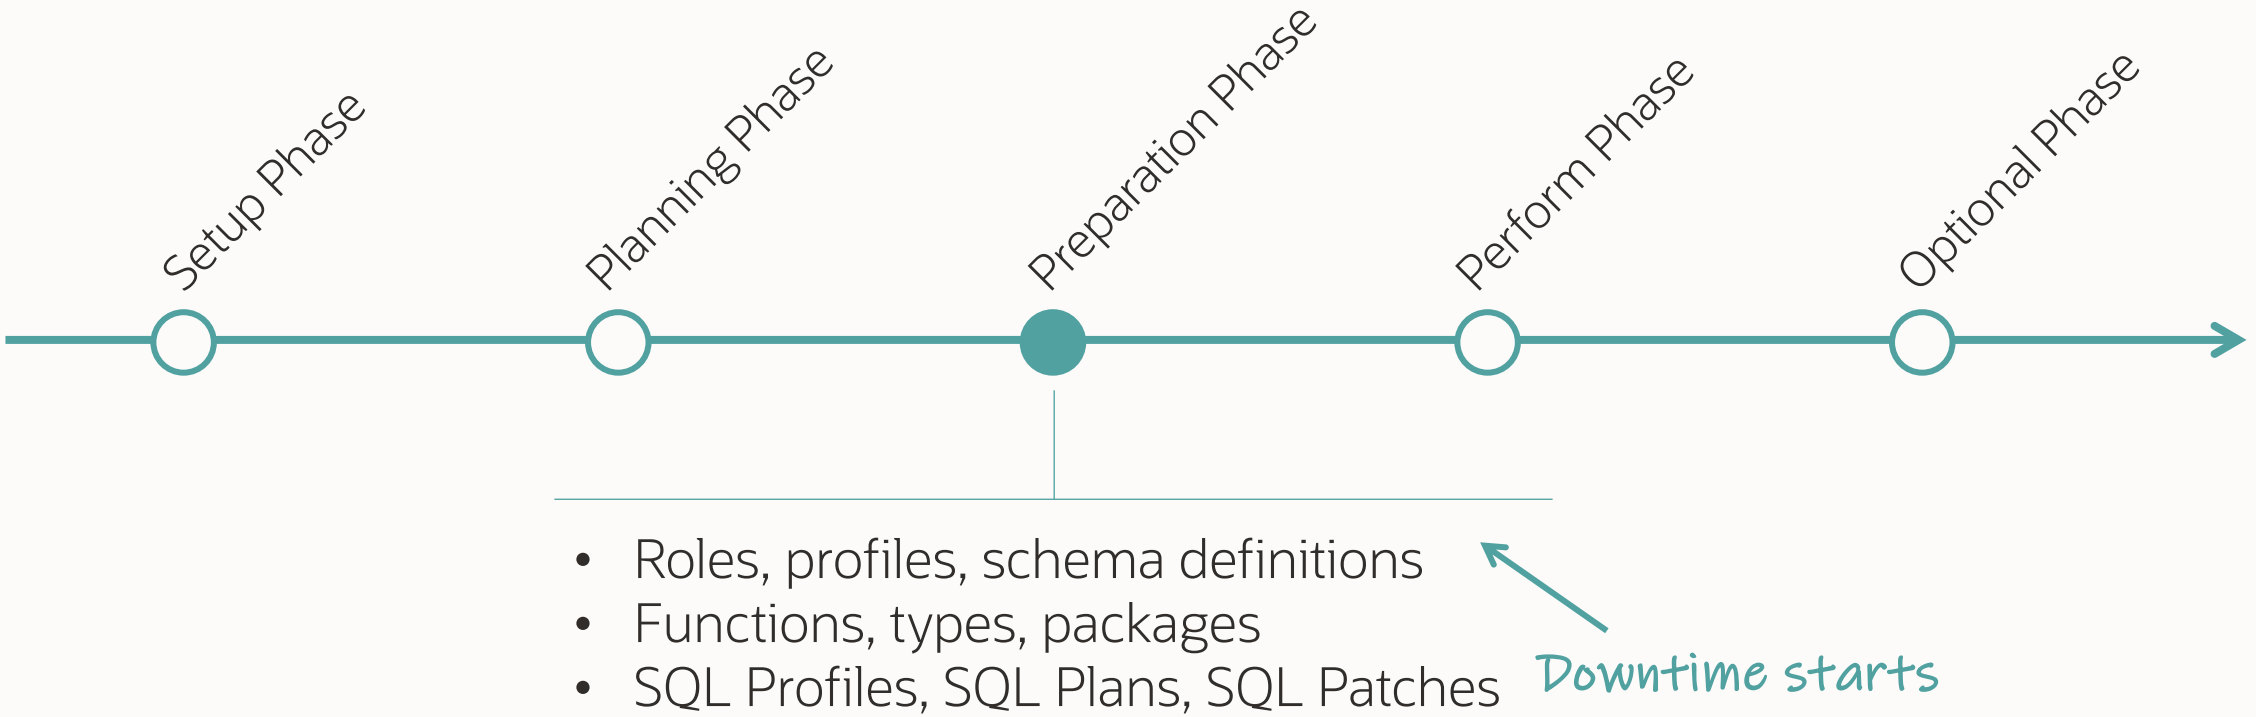

AMA | Preparation Phase

On-Prem - Source	ADB-S - Target
Collect allowed ROLES	
	Create ROLES
Collect PROFILES	
	Create PROFILES
	Create storage credential (NFS, Object Store)
Export schema definition	
	Import schema definition
Export FUNCTIONS, TYPES, PACKAGES	
	Import FUNCTIONS, TYPES, PACKAGES
	Granting migration privileges
	Alter user profiles
Collect SQL Profiles, SQL Plans, SQL Patches	
	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 ---
---
---
```

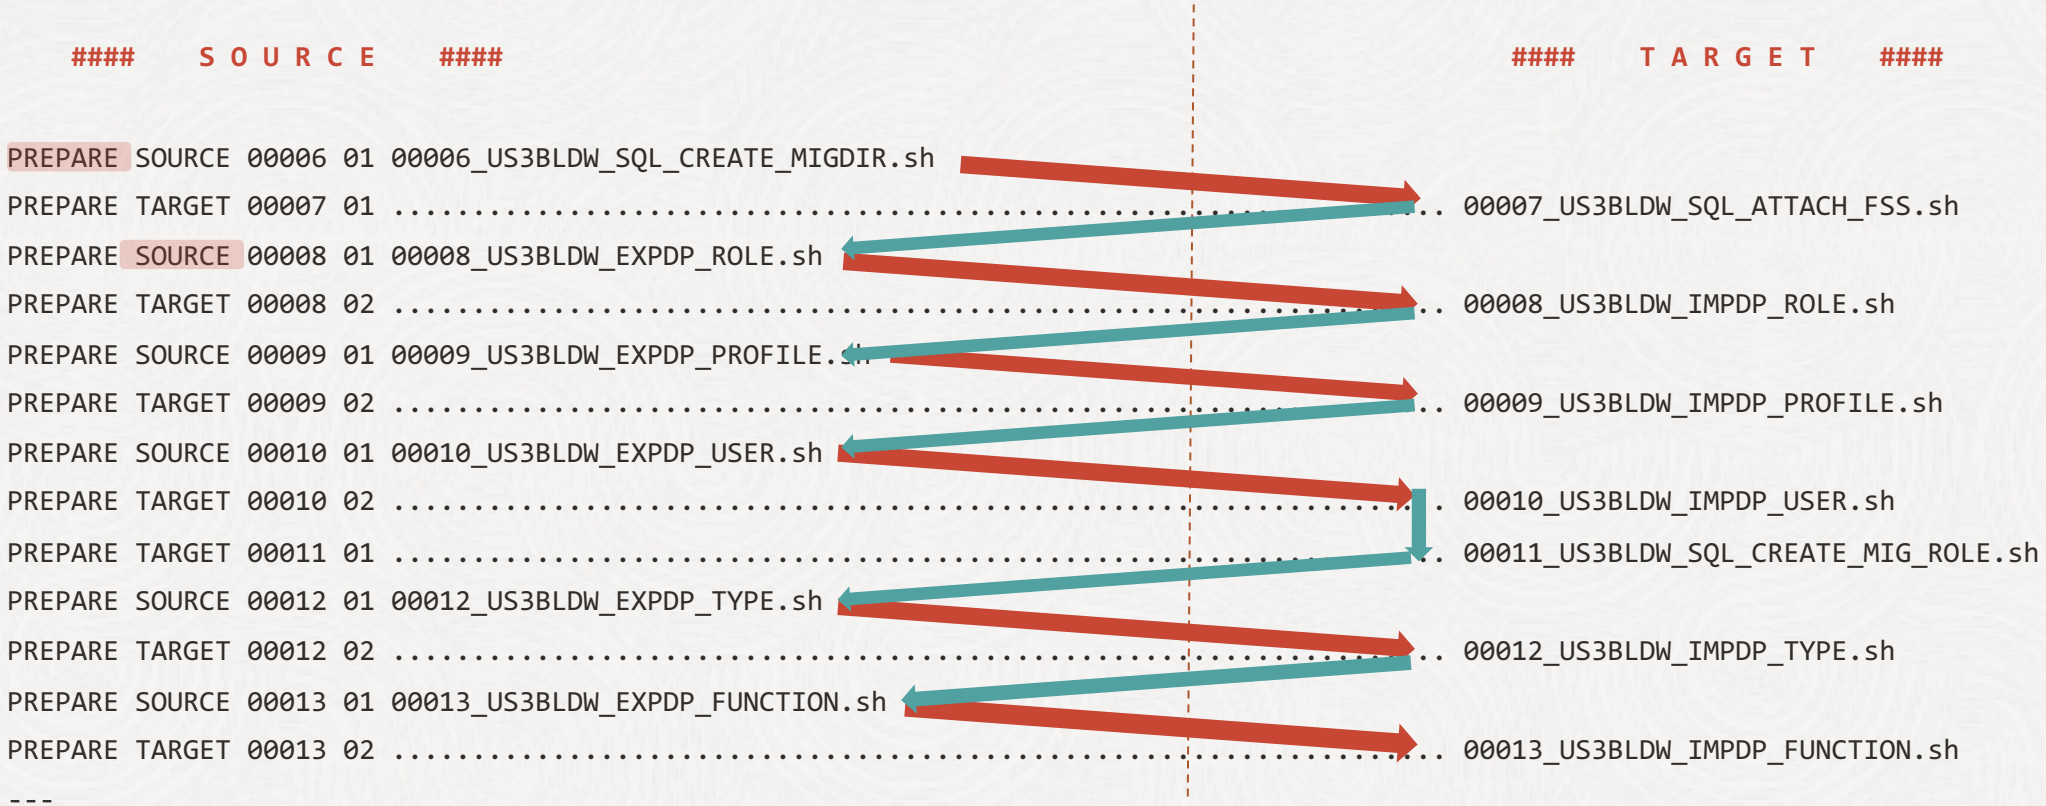


```
---
```

AMA Workflow

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

AMA | Perform Phase

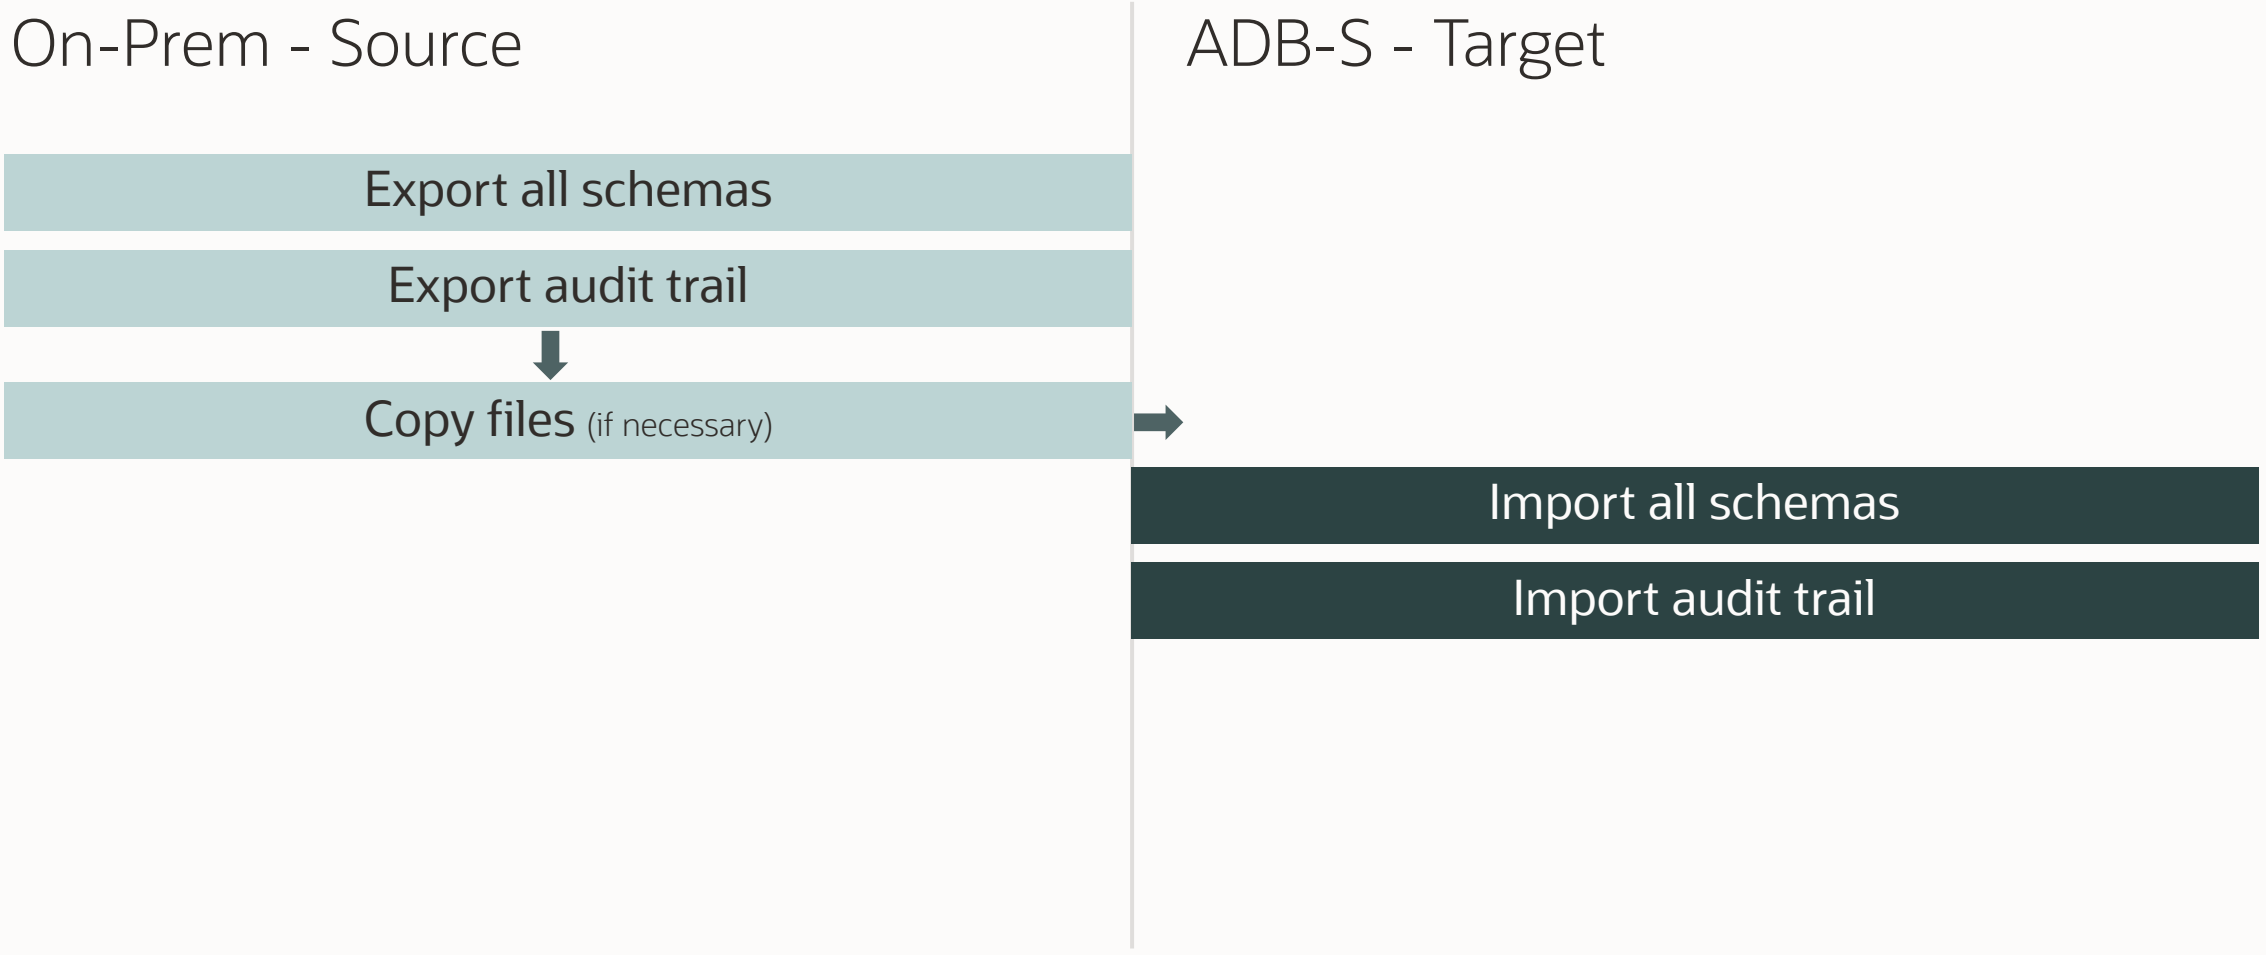

AMA Workflow

- 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

Export network ACLs

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

Import network ACLS

Enable Advanced Queues

Recompilation

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

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

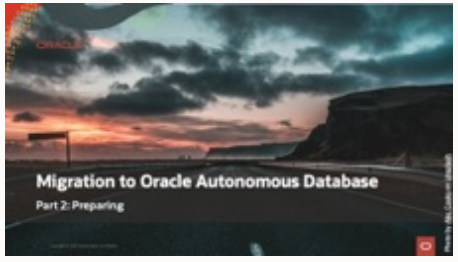

2 PREPARING

Watch recording
Get slides

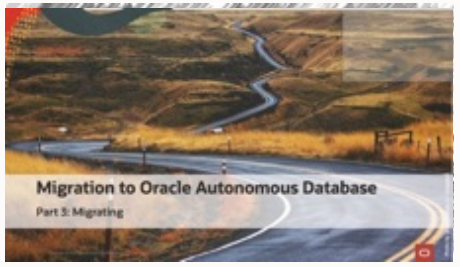

3 MIGRATING

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

Oracle Database 23ai

What's Changing

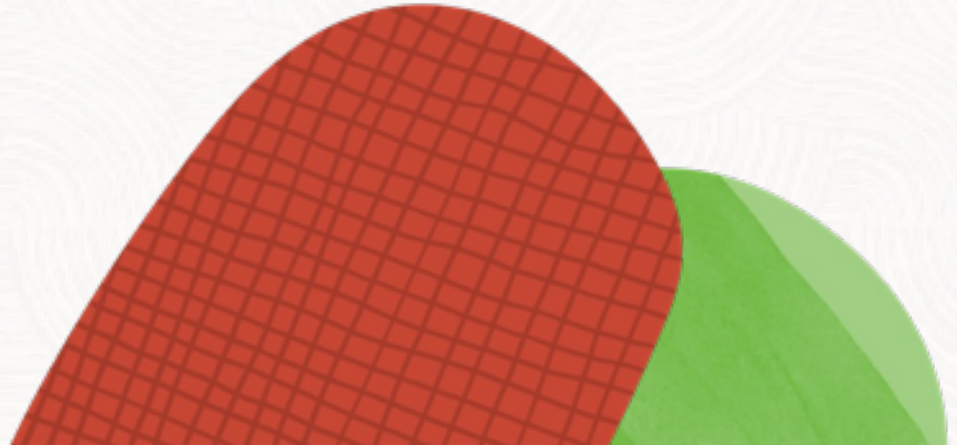

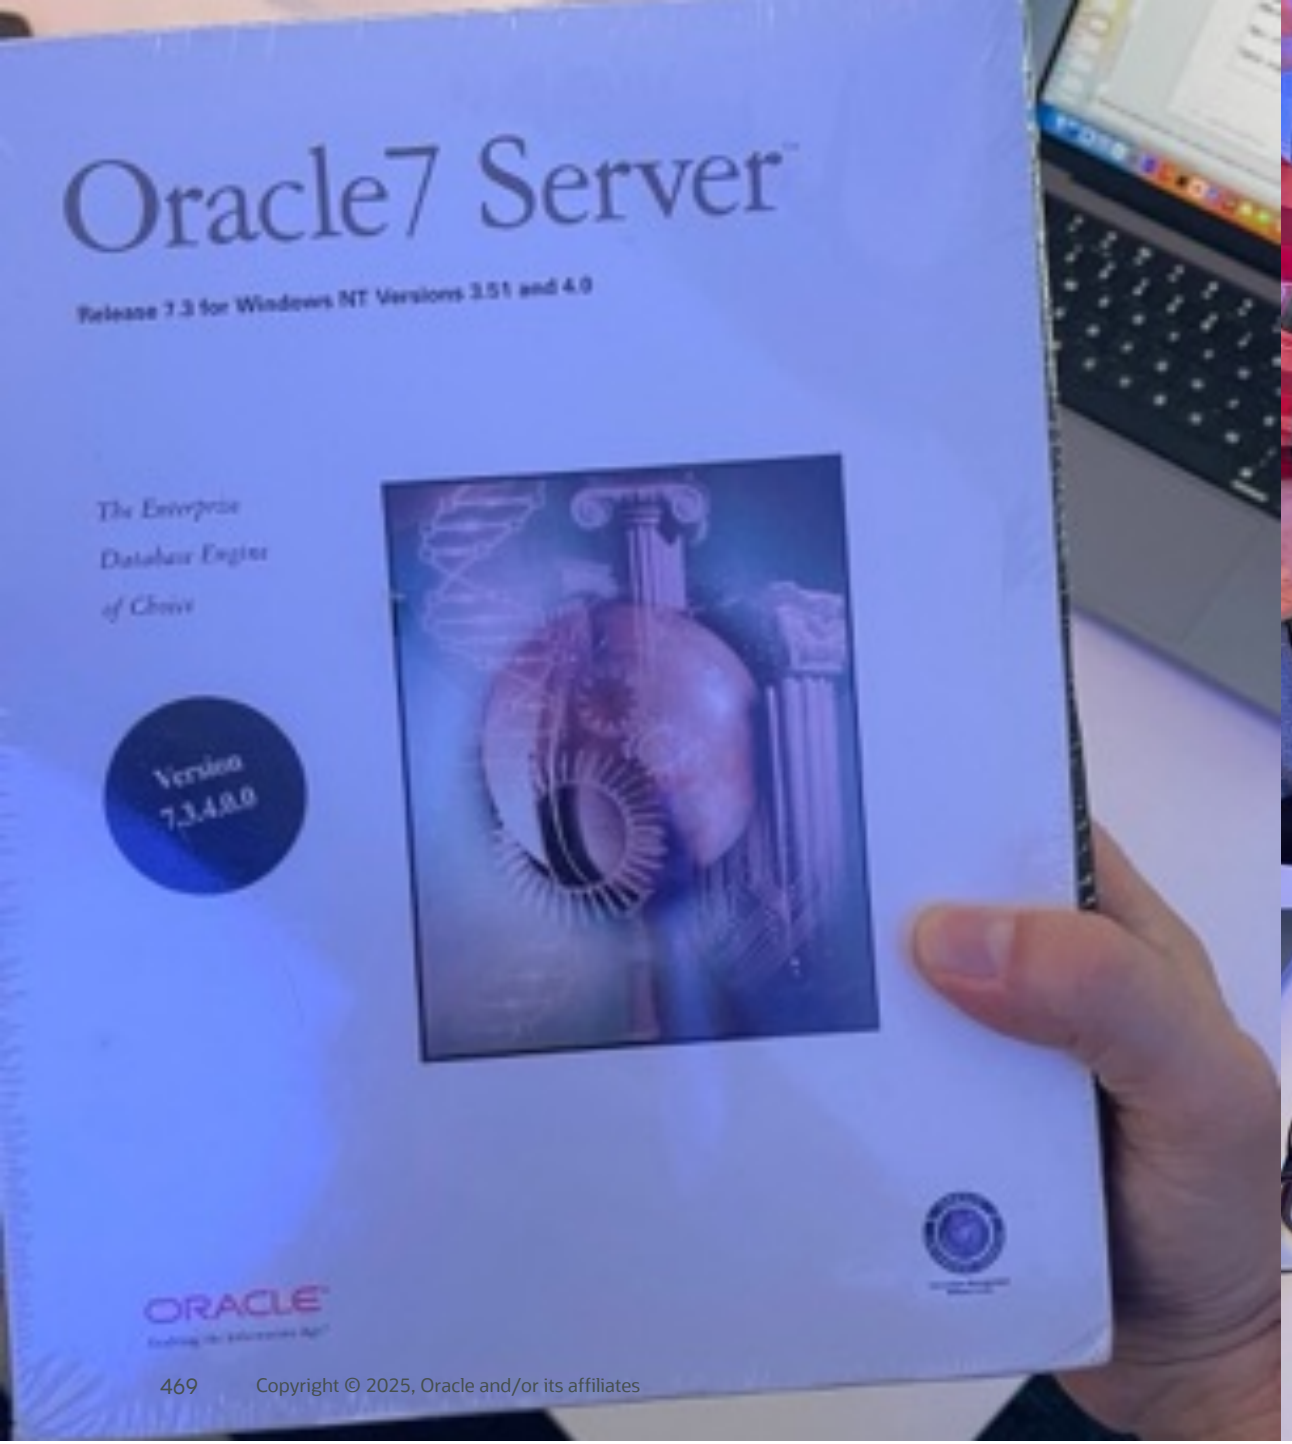

Consult the Upgrade Guide for changes, desupports, and deprecations

Traditional Auditing is desupported in Oracle Database 23ai

- 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 Doc ID 2909718.1)
- Do it before the upgrade

AutoUpgrade is the only supported tool to upgrade your Oracle Database

- DBUA is **desupported**

Enterprise Manager

AutoUpgrade integration into EM CC

- Fleet scale
- Non-CDB to PDB
- Non-rolling logical standby
- TDE
- REST and `emcli` automation

- MOS Note: 2978593.1 EM 13c: Is Agent 13.5 Certified on RHEL 9/OL9?)

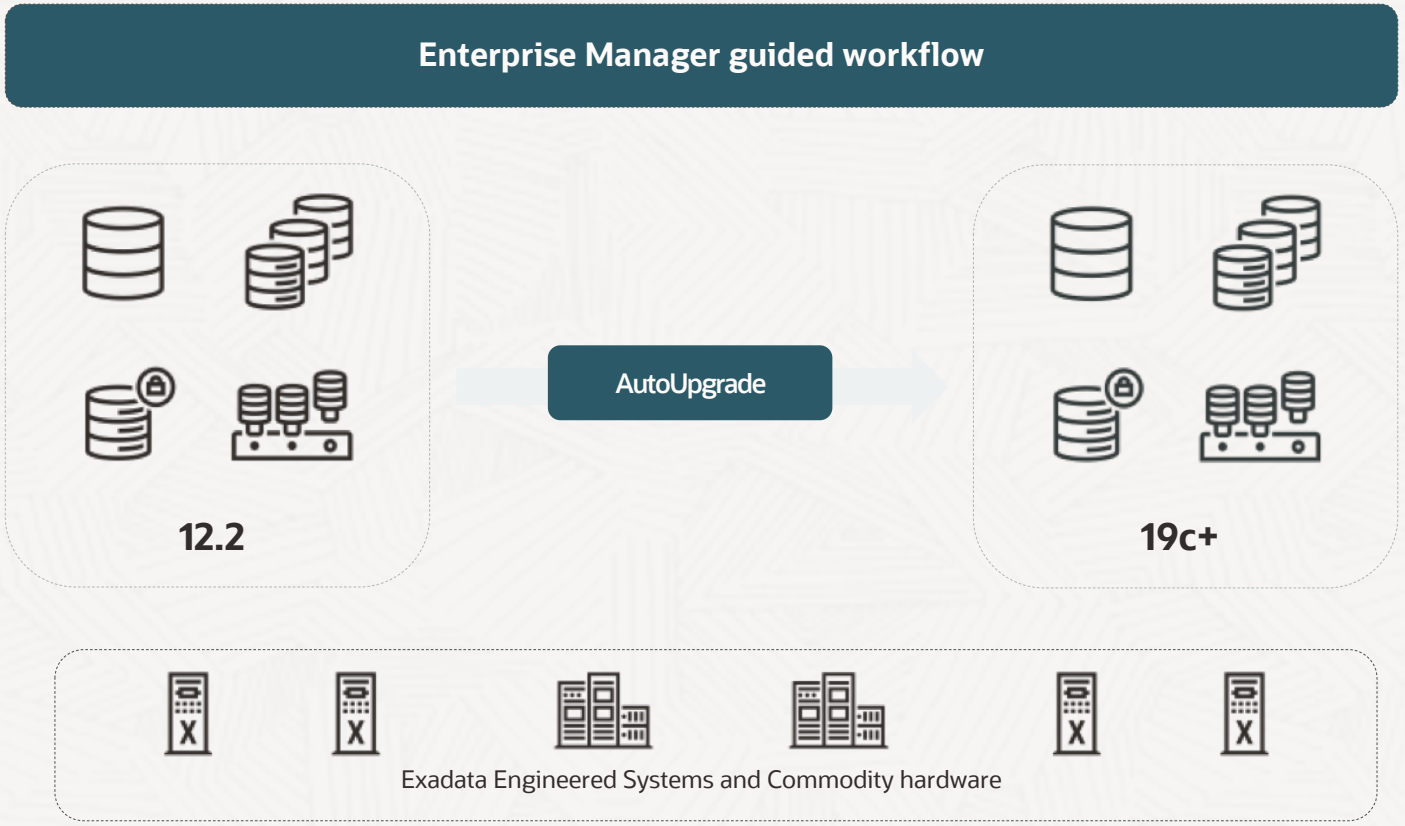

Oracle Database 23ai

What's New

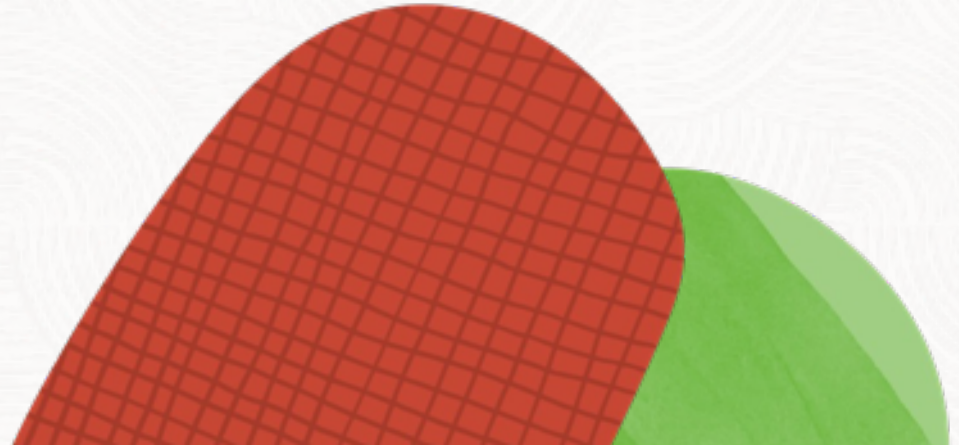

Examine Oracle Database changes using ORAdiff

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

ORAdiff | MOS Account | Oracle Account | daniel.overby.hansen@oracle.com

- Home
- Parameters
- Users, Roles, etc
- Privileges
- Included Fixes
- Oracle Database Home
- Objects
- Fixed Objects
- Audit
- System
- Scheduler
- Report

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

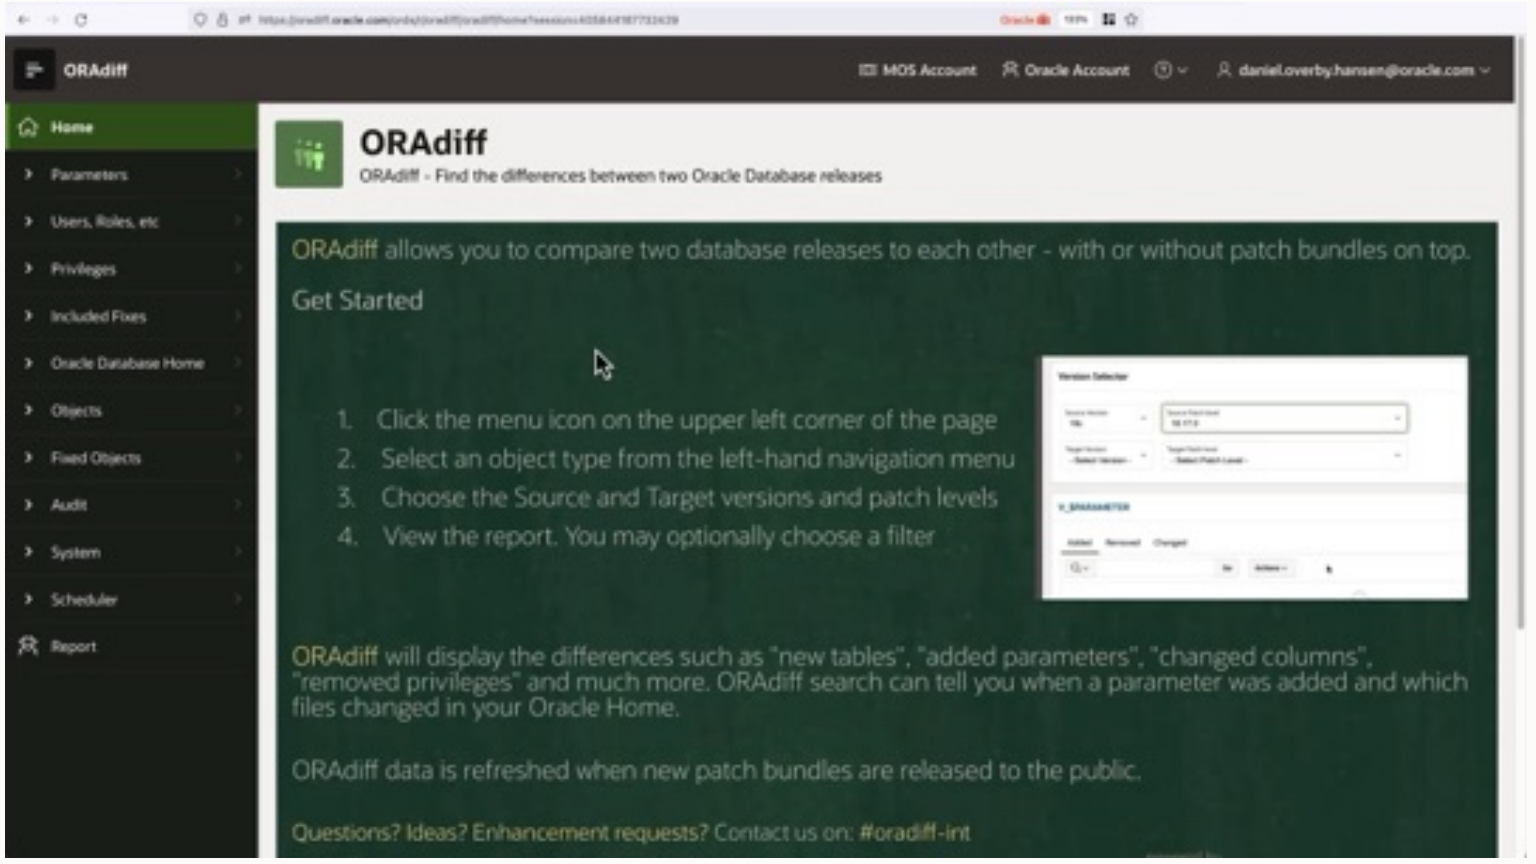

[Watch on YouTube]

Accelerate your applications 10x with True Cache

True Cache

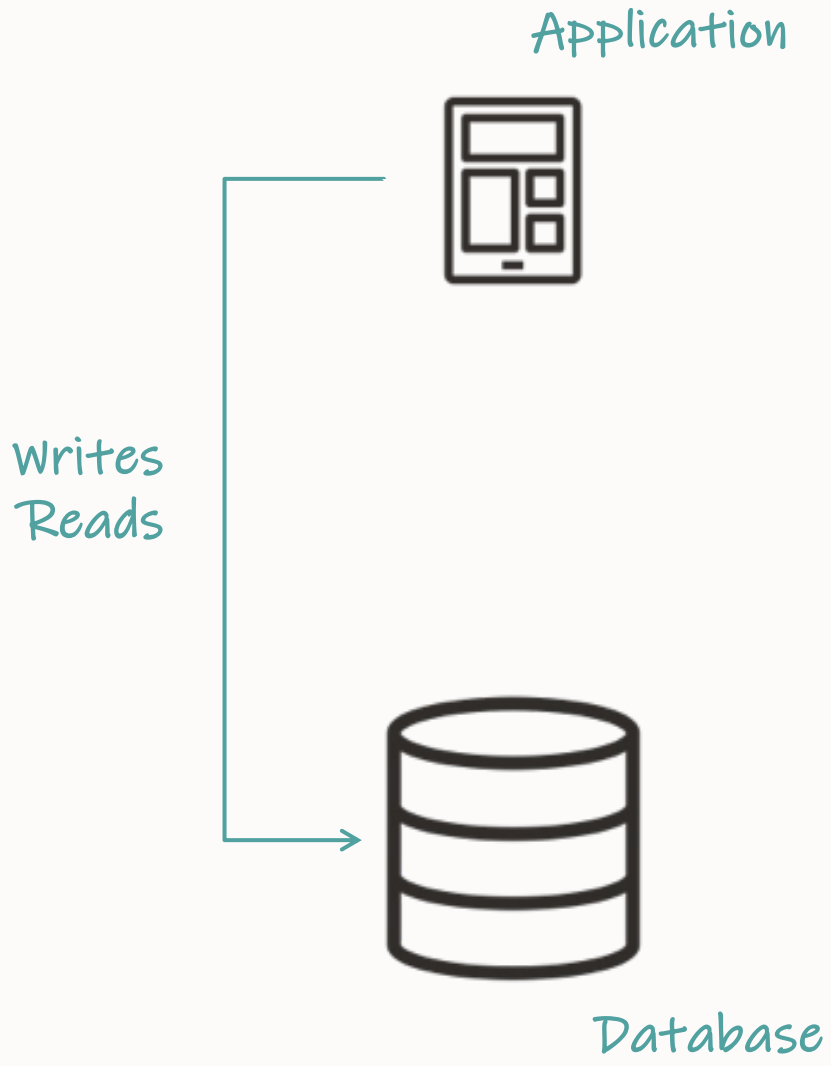

We have a performance problem;
we need a cache!

True Cache

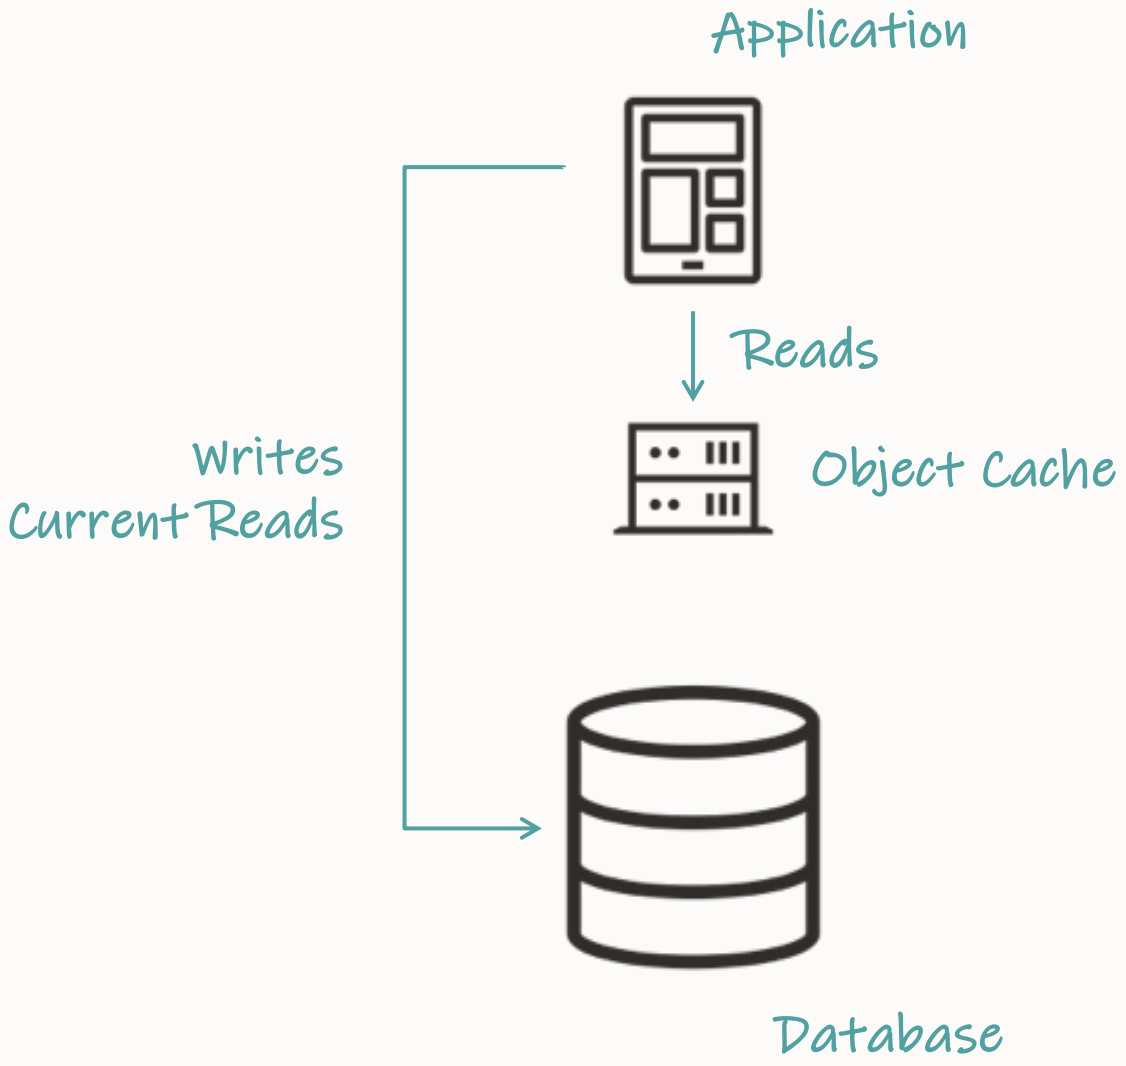

True Cache

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

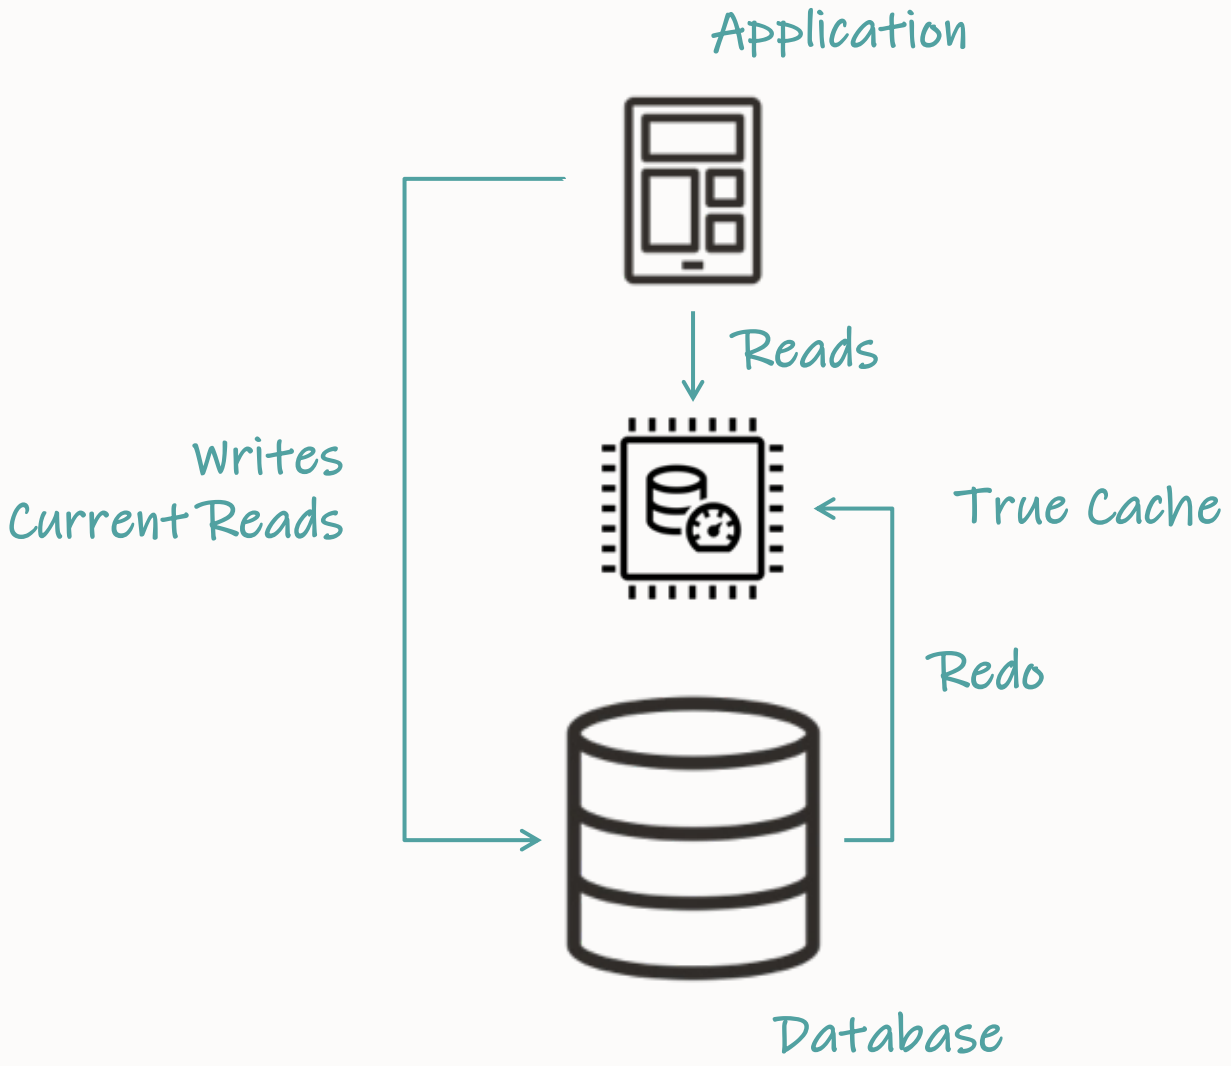

Oracle True Cache is part of Enterprise Edition

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>');
```

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
```


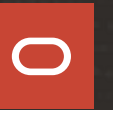

Even wider tables

- 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

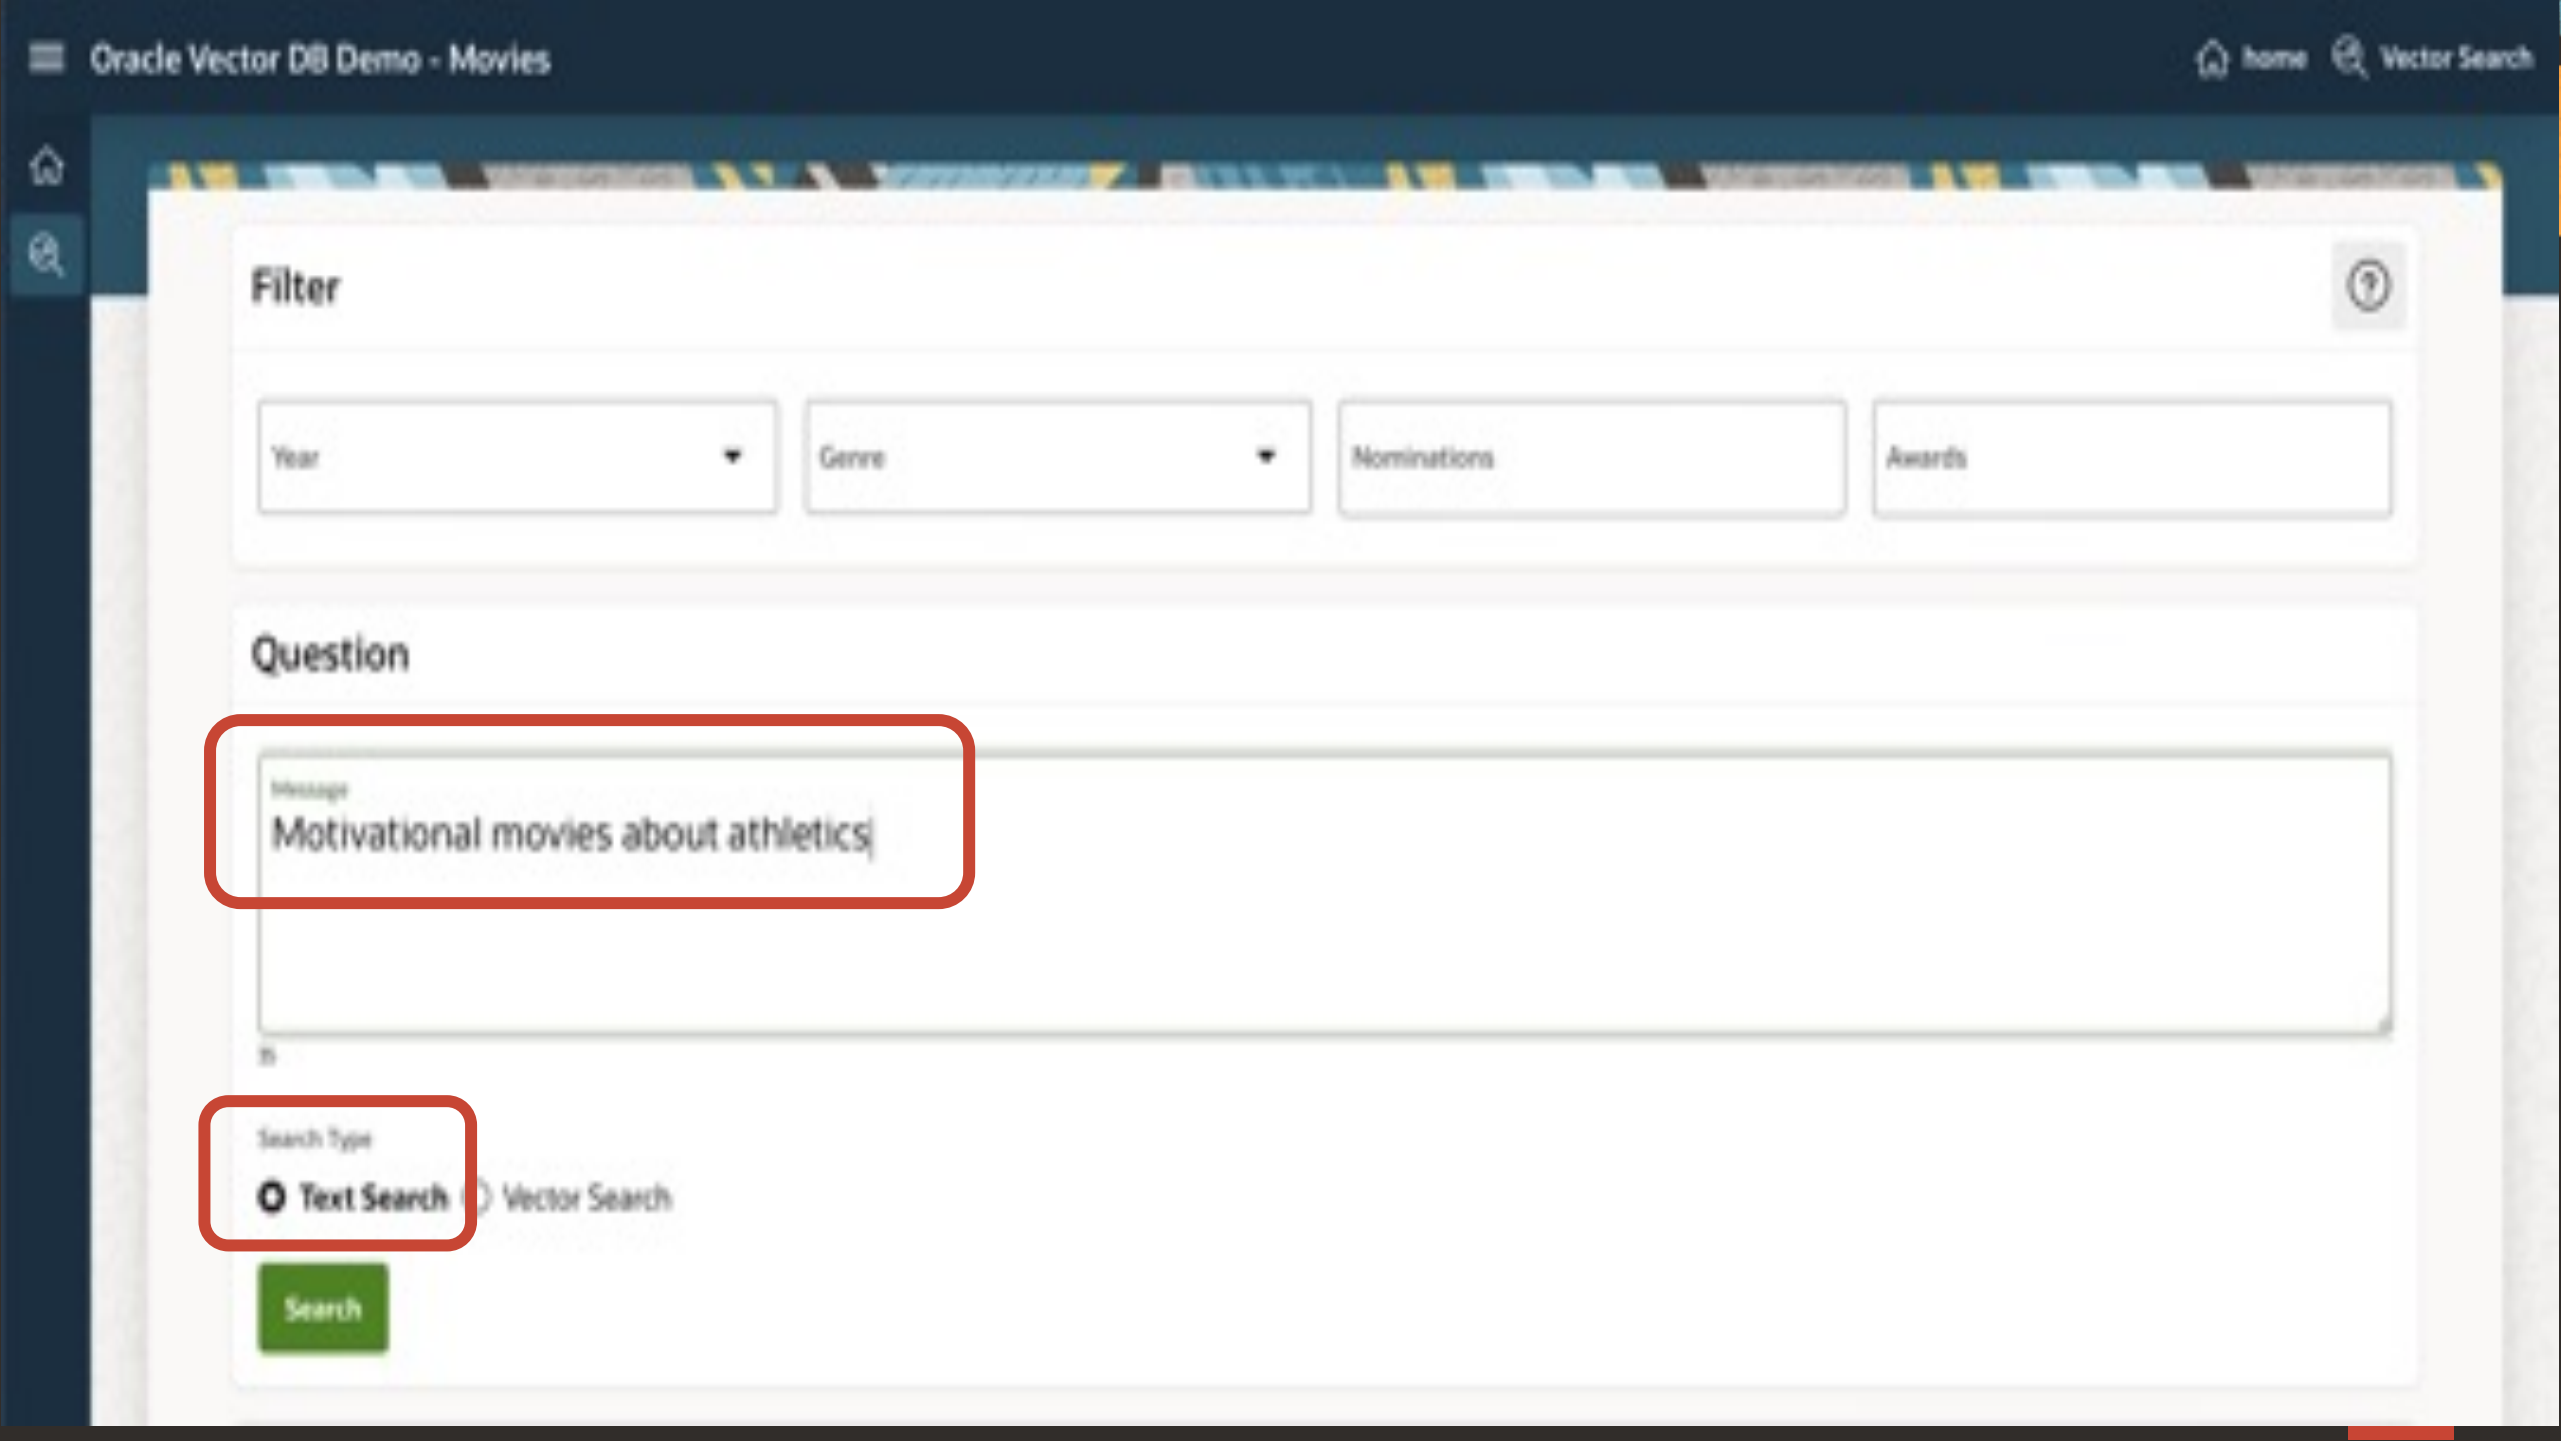
Oracle Vector DB Demo - Movies
home
Vector Search
Filter
Year
Genre
Nominations
Awards
Question
Motivational movies about athletics
Search Type
Text Search
Vector Search
Search

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).

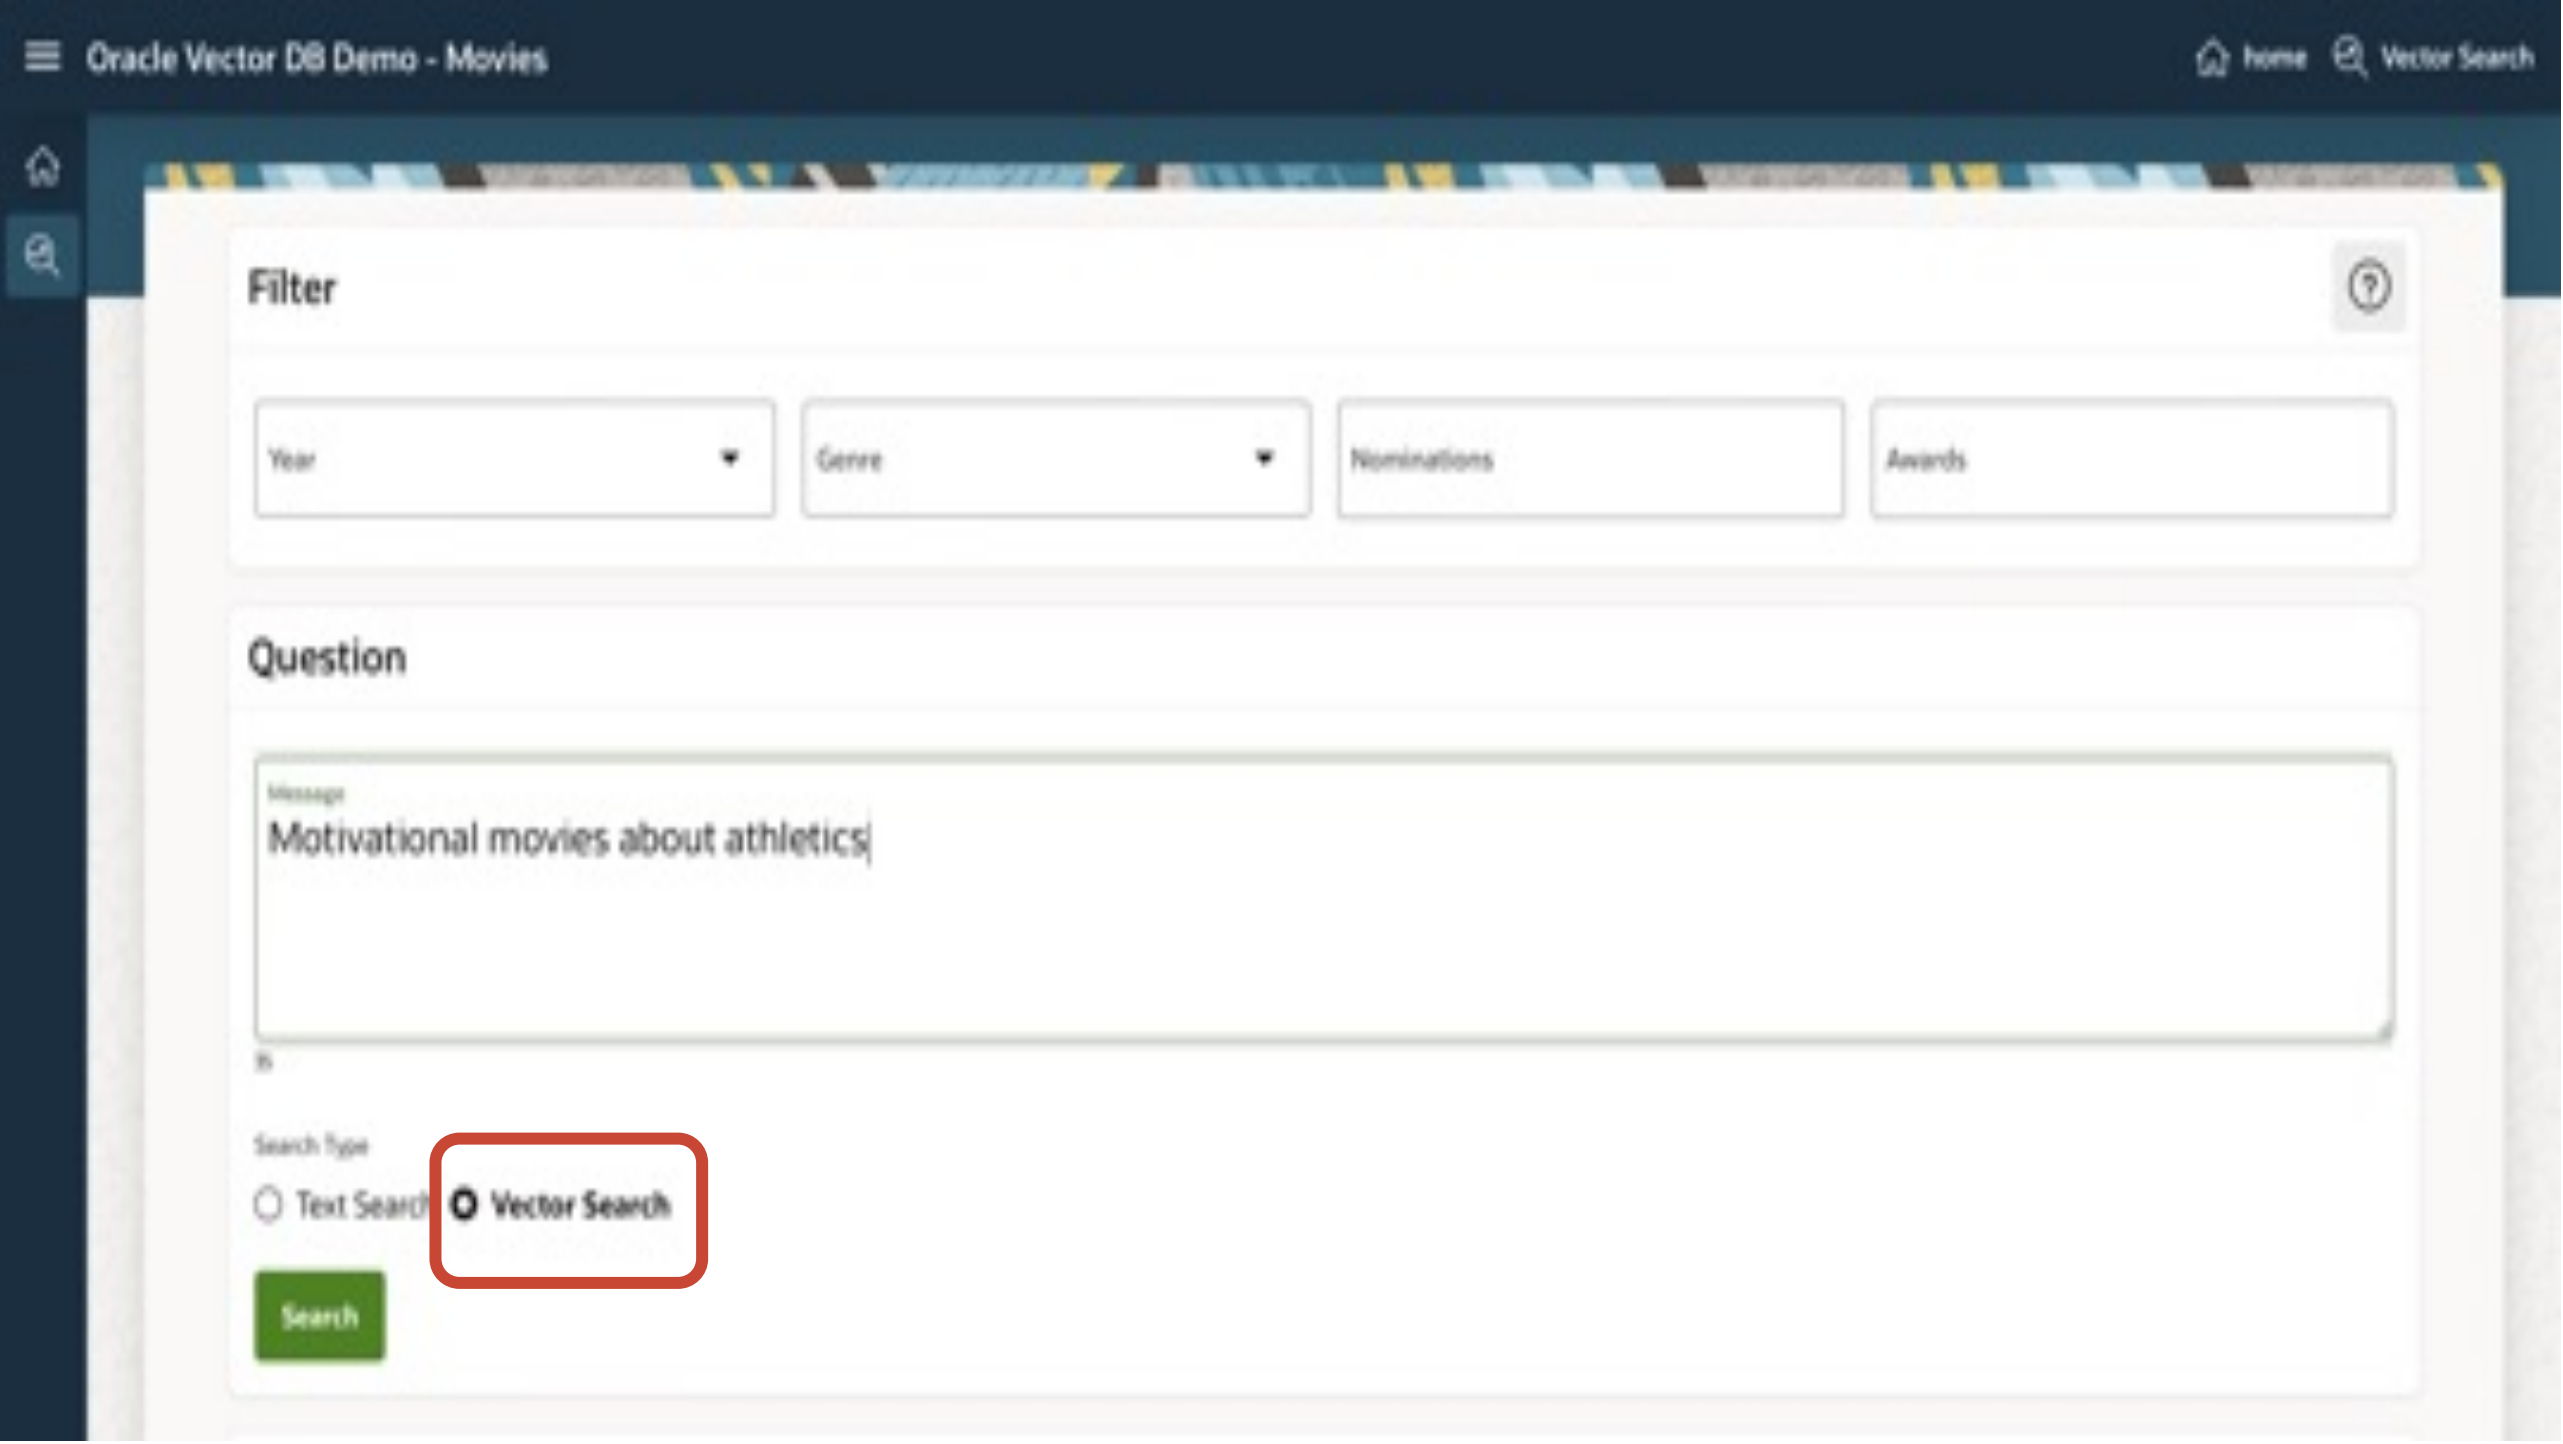
Oracle Vector DB Demo - Movies
home
Vector Search
Filter
Year
Genre
Nominations
Awards
Question
Message
Motivational movies about athletics
Search Type
Text Search
Vector Search
Search

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

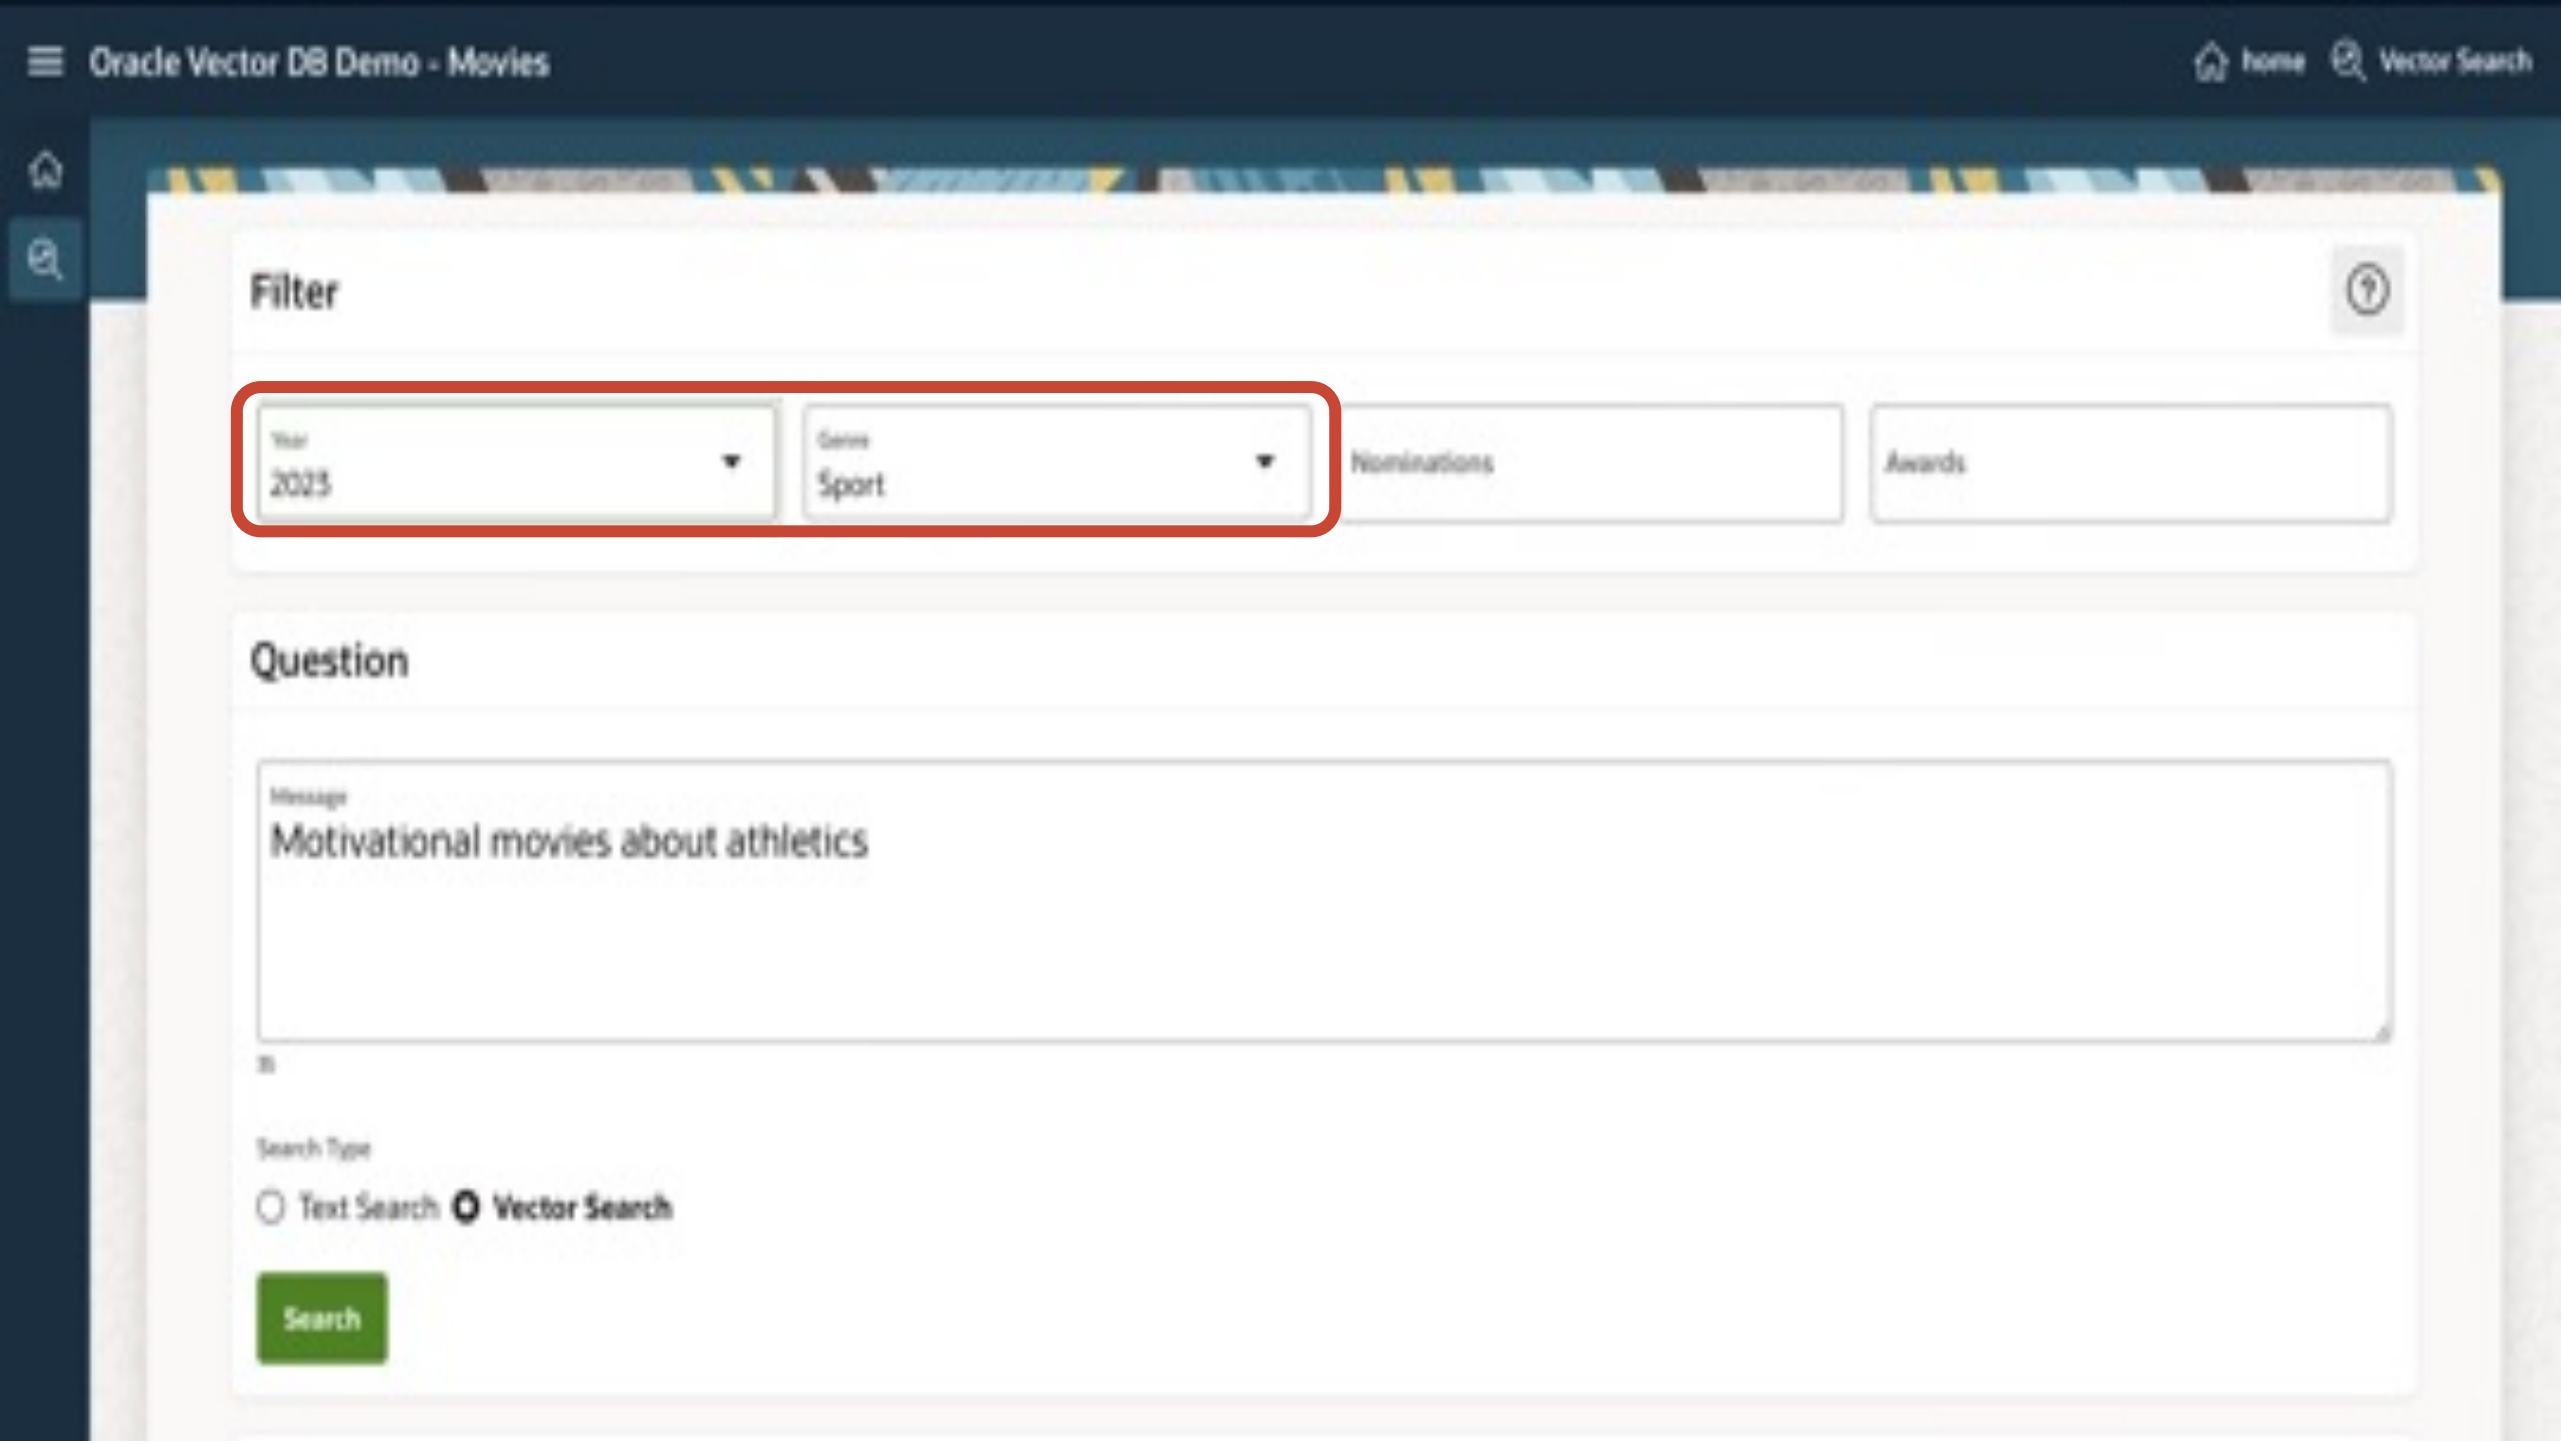
Oracle Vector DB Demo - Movies
home
Vector Search
Filter
Year
2023
Genre
Sport
Nominations
Awards
Question
Message
Motivational movies about athletics
Search Type
Text Search
Vector Search
Search

Results

1 - 3 of 3

	Title	Awards	Year	Nominations	Genre
	Champions	NONE	2025	NONE	Comedy,Sport
	Sweetwater	NONE	2025	NONE	Biography,Sport
	80 for Brady	NONE	2025	NONE	Sport,Comedy

1 - 3 of 3

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?

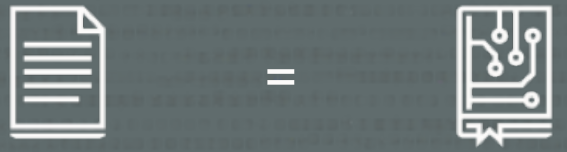

How to search for search complex data types?

Vector Database

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

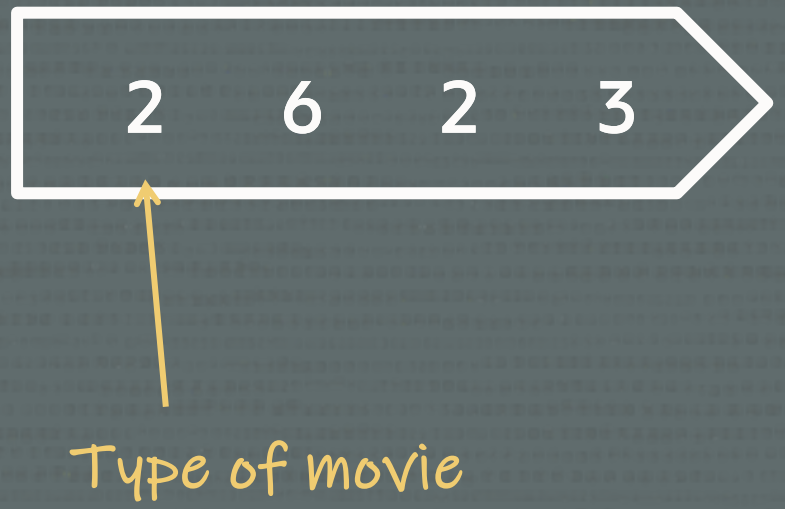

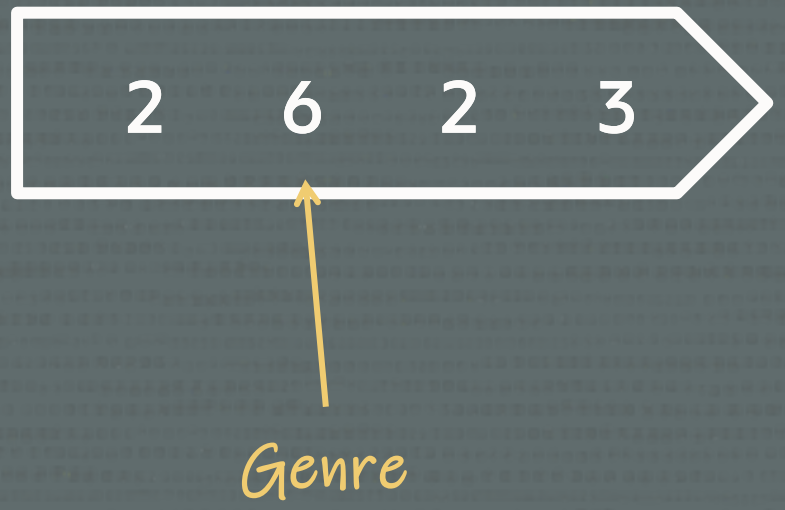

Distance (Euclidean Squared)
$= (3-2)^2+(1-6)^2+(2-2)^2+(8-3)^2$

Vector Database


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

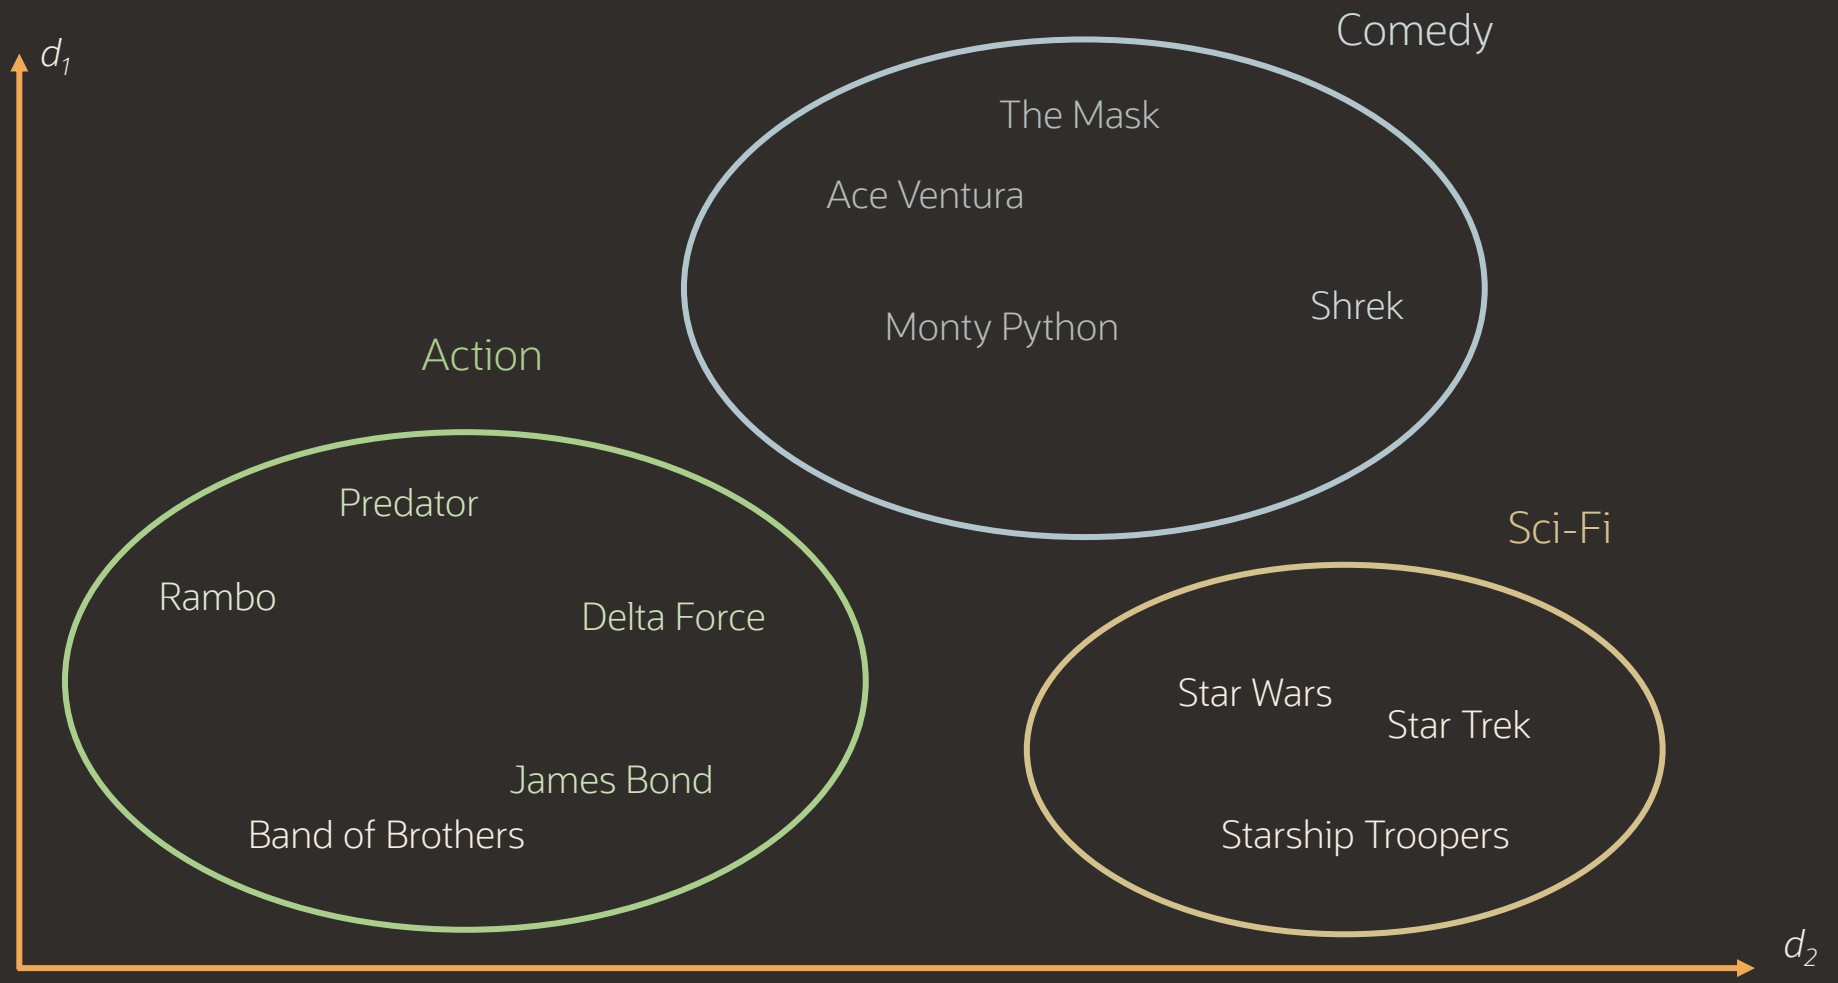

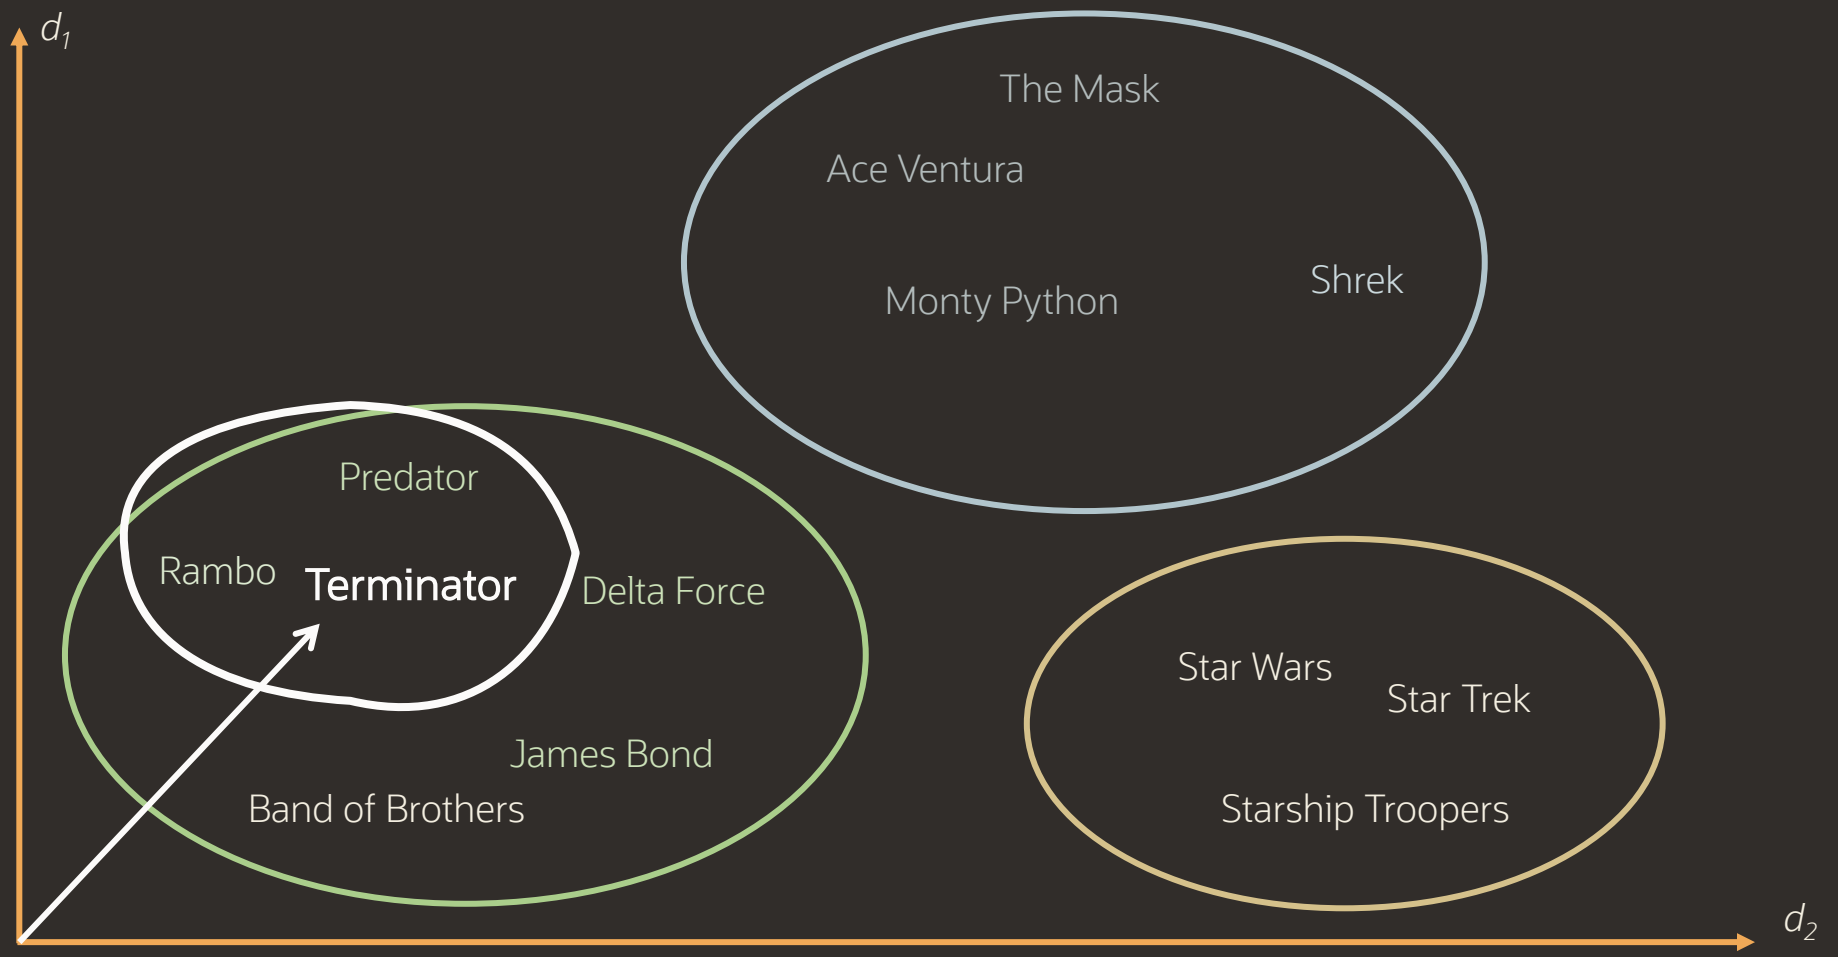

How do you convert your data into vectors?

Vector Database

- Embedding models transform your data into a vector

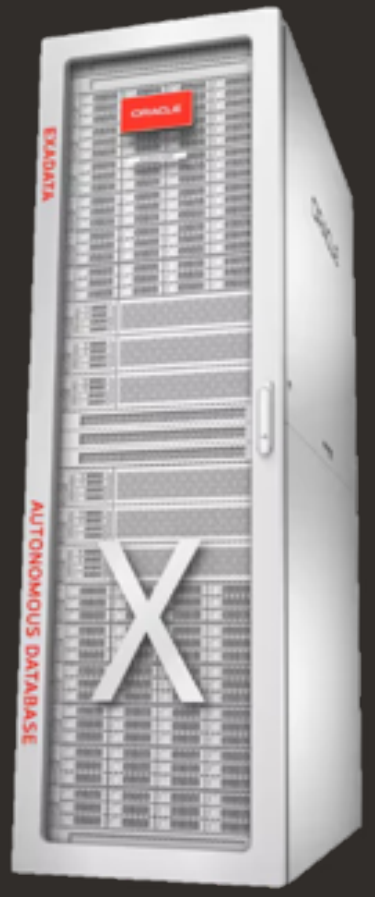

10,000 GPU ?

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


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

Vector Database

Vector Database

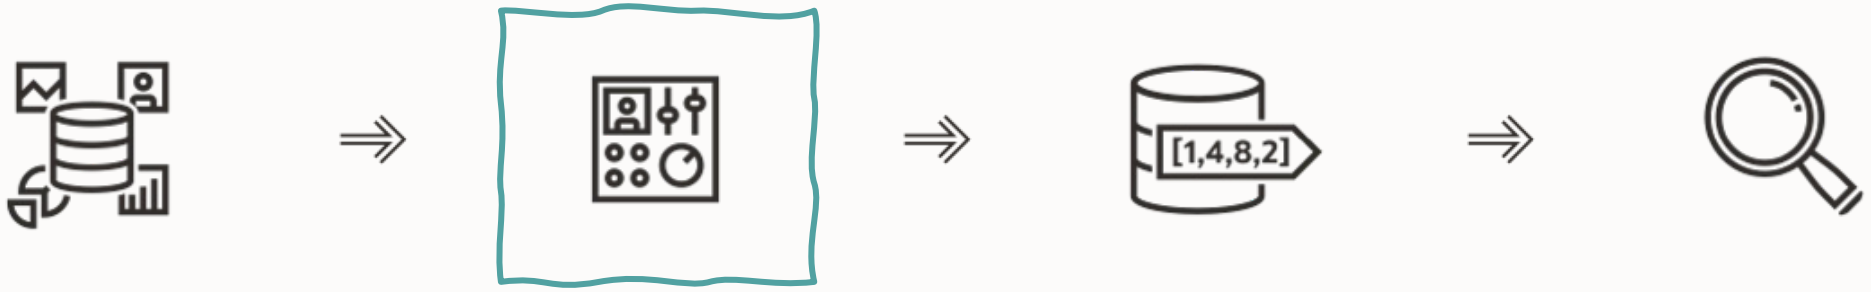

Vector Database

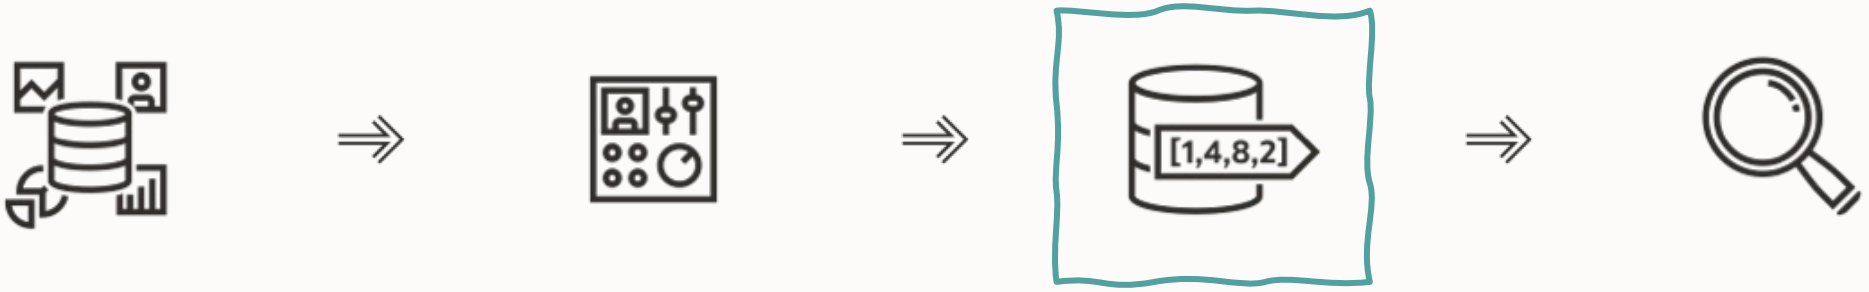

Vector Database

Included in any edition of Oracle Database 23ai

- No extra license required

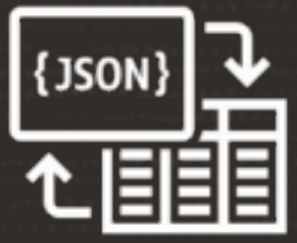

Get the best of both worlds with JSON Duality Views

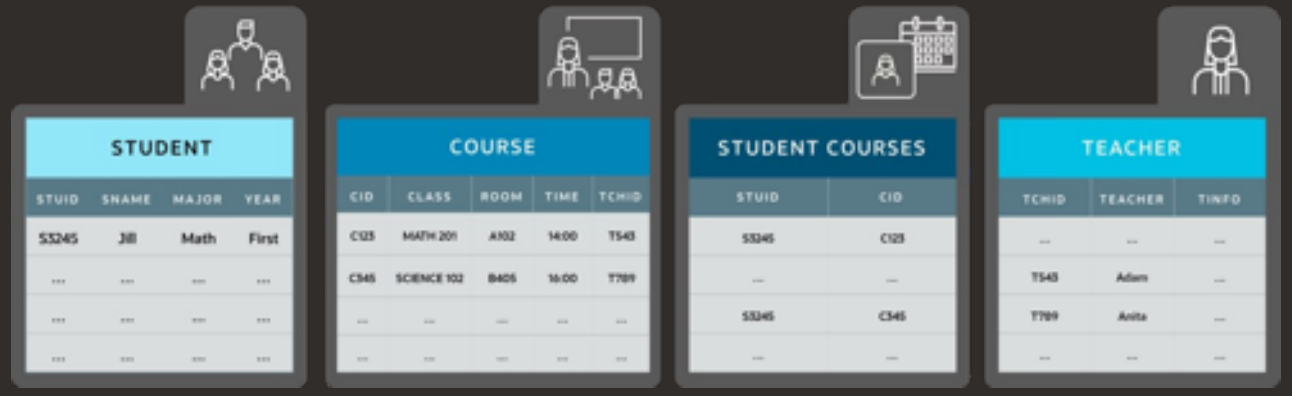

SCHEDULE FOR: JILL

```
"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

Oracle Database
secures your data

Data Guard

RAC

RMAN

TDE Tablespace Encryption

Network Encryption

Auditing

SQL Firewall

ACID

No data duplication

... and so much more

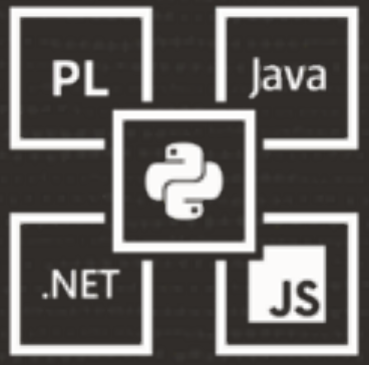

Easily convert your existing application with Oracle Database API for MongoDB

- Using Oracle Database API for MongoDB

Oracle Database 23ai

Even More Secure

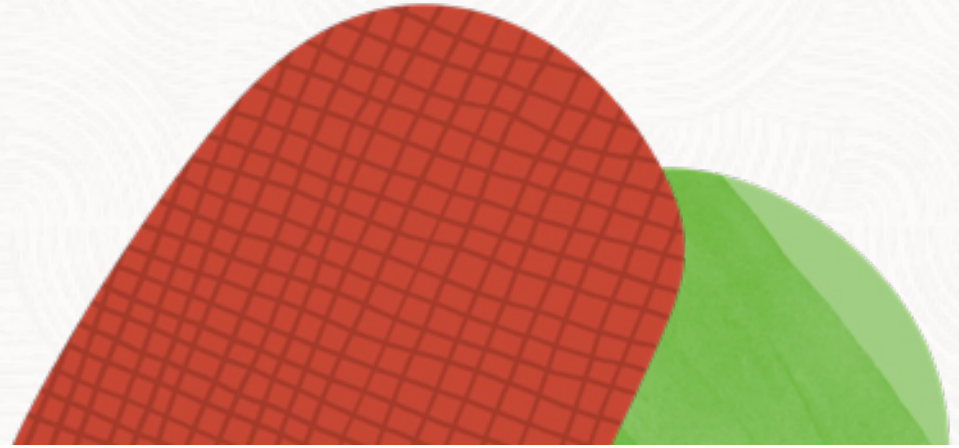

No more insecure case insensitive password

- 10G password verifies are no longer accepted

Grant privileges to an entire schema in one command

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

Oracle Database 23ai

Small, but useful

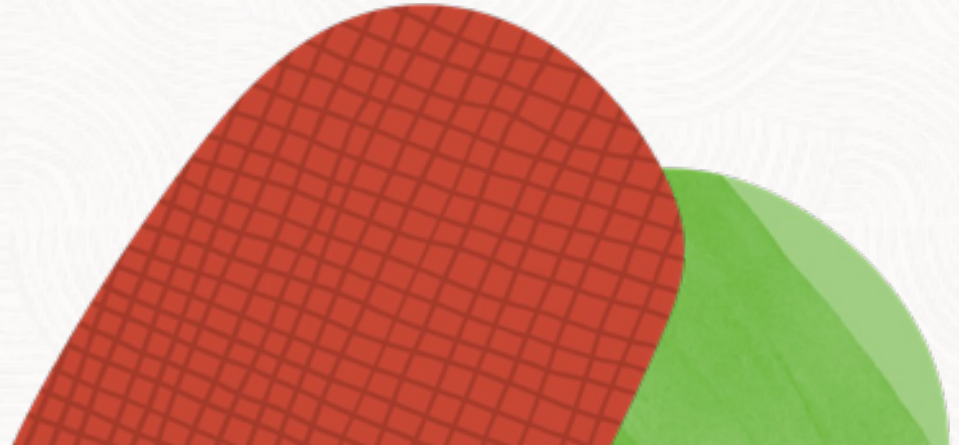

```
--Regardless of whether the object exists or not,
--the DROP command don't produce an error


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

```
--No longer need to select from dual

select sysdate from dual;

select sysdate;
```


Key Learnings

1. Patch regularly and out-of-place
2. Use AutoUpgrade
3. Use SQL Plan Management

Oracle

DBAs

run the world

YouTube | Oracle Database Upgrades and Migrations

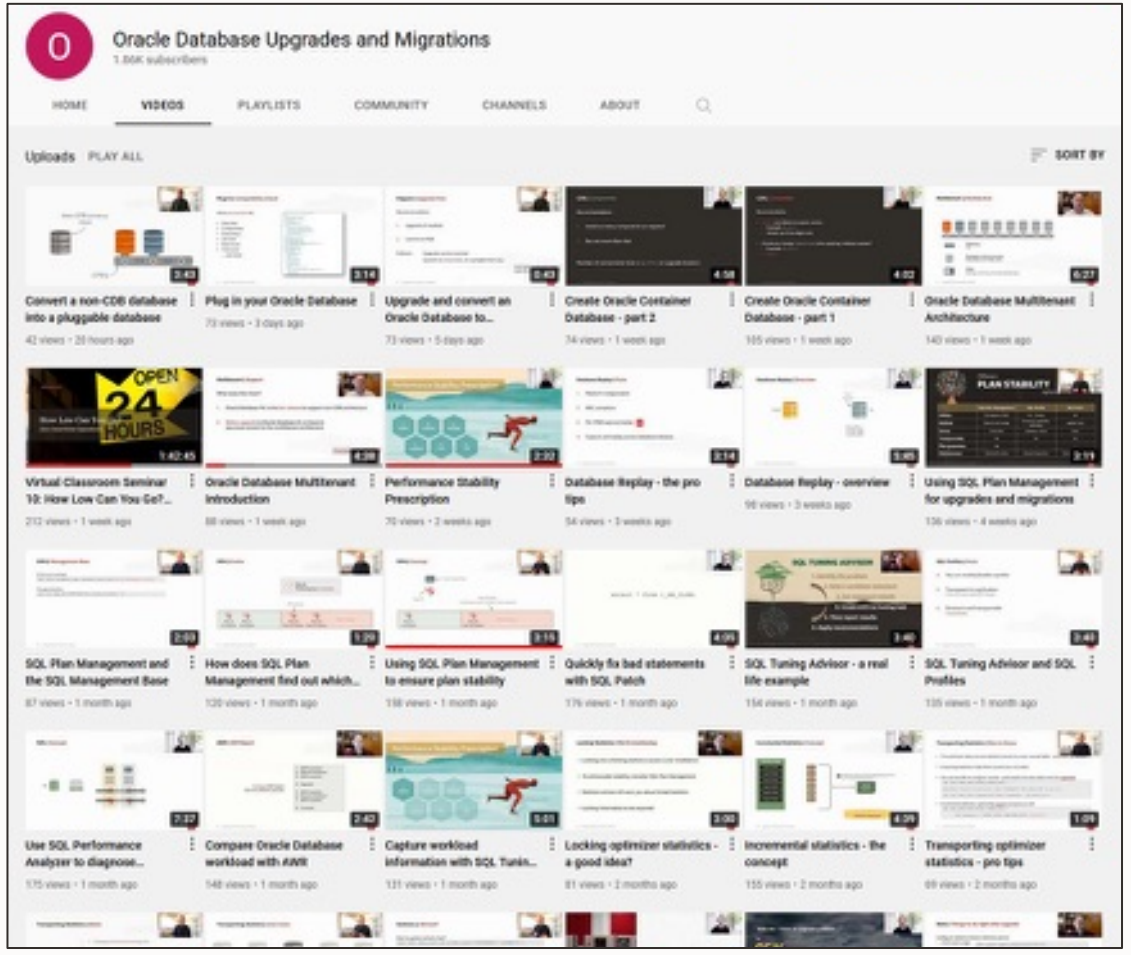

https://www.youtube.com/@upgradenow

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

SCAN ME

Thank You