ORACLE

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

Oslo, October 2024

MIKE DIETRICH
Vice President
Database Upgrade, Migrations & Patching

mikedietrich

@mikedietrichde

https://mikedietrichde.com

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

dohdatabase

@dohdatabase

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

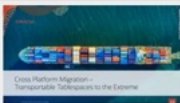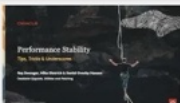| Web Seminar	Slides
Episode 16 (replaces Episode 1 from Feb 2021) Oracle Database Release and Patching Strategy for 19c and 23c *115 minutes – May 10, 2023*	
Episode 17 From SR to Patch – Insights into the Oracle Database Development process *55 minutes – June 22, 2023*	
NEW Episode 18 Cross Platform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

Recorded Web Seminars

https://MikeDietrichDE.com/videos

More than 35 hours of technical content,
on-demand, anytime, anywhere

AGENDA

09:30
Welcome
Release Strategy
Patching

11:00
Upgrade
Multitenant

13:00
Data Pump + CMA
Performance Stability
Insights into development

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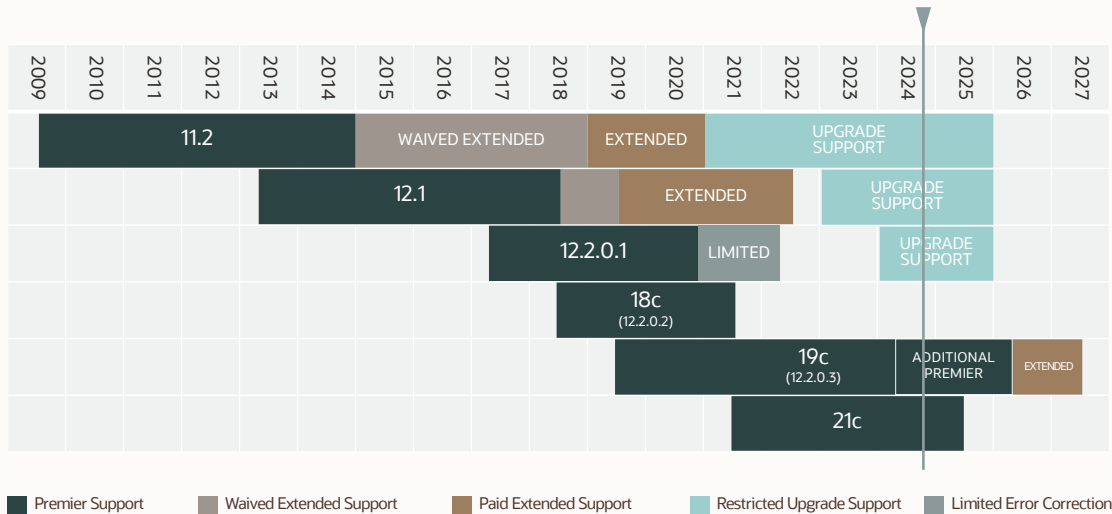What's Coming

10:45
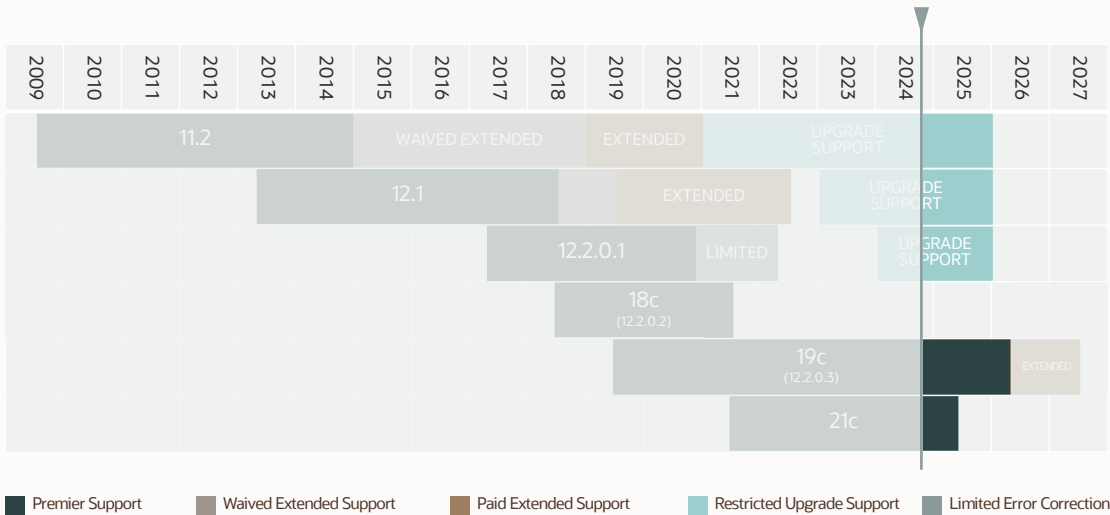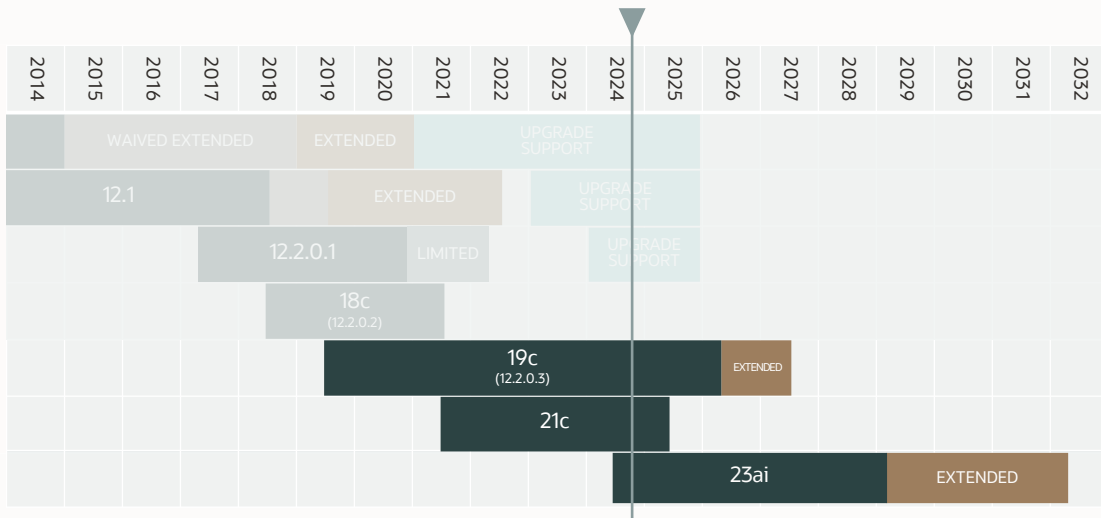Coffee break

12:00
Lunch

15:00
Coffee break

Release Strategy

Lifetime Support Policy

	2009	2010	2011	2012	2013	2014	2015	2016	2017	2018	2019	2020	2021	2022	2023	2024	2025	2026	2027

11.2 — Premier Support (2009–2014), WAIVED EXTENDED (2015–2017), EXTENDED (2018–2019), UPGRADE SUPPORT (2020–2024)

12.1 — Premier Support (2012–2017), Waived Extended (2018), EXTENDED (2018–2021), UPGRADE SUPPORT (2022–2024)

12.2.0.1 — Premier Support (2016–2020), LIMITED (2021), UPGRADE SUPPORT (2023–2024)

18c (12.2.0.2) — Premier Support (2018–2021)

19c (12.2.0.3) — Premier Support (2019–2023), ADDITIONAL PREMIER (2024–2025), EXTENDED (2025–2026)

21c — Premier Support (2021–2024)

Legend:
- ■ Premier Support
- ■ Waived Extended Support
- ■ Paid Extended Support
- ■ Restricted Upgrade Support
- ■ Limited Error Correction

Lifetime Support Policy

	2009	2010	2011	2012	2013	2014	2015	2016	2017	2018	2019	2020	2021	2022	2023	2024	2025	2026	2027
11.2							WAIVED EXTENDED				EXTENDED		UPGRADE SUPPORT						
12.1									EXTENDED					UPGRADE SUPPORT					
12.2.0.1									LIMITED				UPGRADE SUPPORT						
18c (12.2.0.2)																			
19c (12.2.0.3)																EXTENDED			
21c																			

Legend: Premier Support · Waived Extended Support · Paid Extended Support · Restricted Upgrade Support · Limited Error Correction

Lifetime Support Policy

	2014	2015	2016	2017	2018	2019	2020	2021	2022	2023	2024	2025	2026	2027	2028	2029	2030	2031	2032
		WAIVED EXTENDED			EXTENDED		UPGRADE SUPPORT												
12.1					EXTENDED			UPGRADE SUPPORT											
12.2.0.1							LIMITED		UPGRADE SUPPORT										
18c (12.2.0.2)																			
19c (12.2.0.3)										EXTENDED									
21c																			
23ai											EXTENDED								

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 <u>only</u> 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

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

To Bottom

Visibility: EXTERNAL

(98)

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 will appear once you make your selection.

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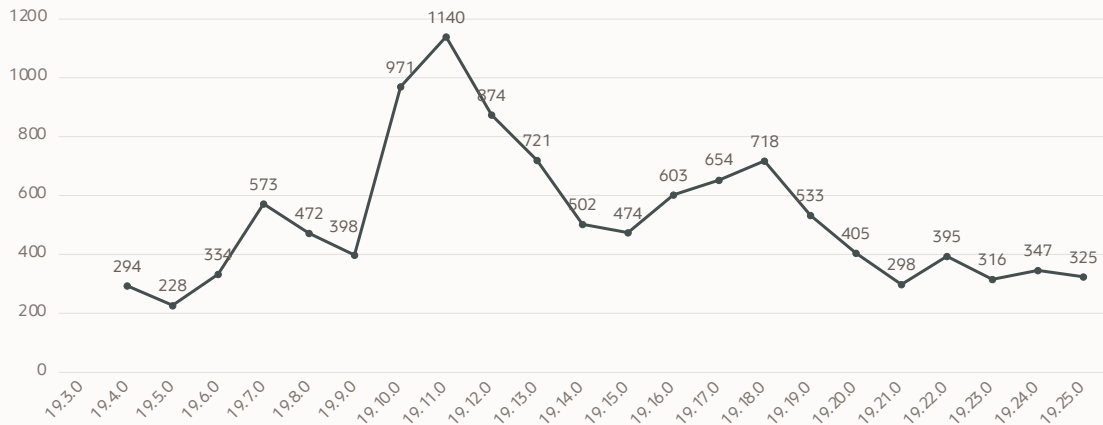
Database 19 Release Updates and Revisions Bugs Fixed Lists (Doc ID 2523220.1)

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

- More than 11k fixes with 19.24.0
- More than 300 security fixes

Apply the Most Important Patches

Always use Important Recommended One-Off Patches: [MOS Note: 555.1](#)

Recommended Patches for 19.23 DB Home

Below is the list of important patches to consider applying on top of 19.23. In addition to the relevant patches listed below, you should also review patches in the [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
[36587533](#)	19.24	[DBMRP 19.22.0.0.240618](#), [DBMRP 19.23.0.0.240618](#)	[SQL EXECUTION] Result Cache: Global Flush Should Always Clear Bypass Flag even if the Result Cache is Uninitialized	[[list-patches]](#)	YES	YES	27-SEP-2024
[32781163](#)			[QRY OPTIMIZER] Deadlock on library cache lock on MV refresh and dbms_stats gathering the same object	[[list-patches]](#)	YES		27-SEP-2024
[36908826](#) (replaces [33539779](#))			[RMAN] ORA-600[krbr1sf_no_dc_record] during RMAN Restore	[[list-patches]](#)	YES		24-SEP-2024
[35398148](#)			[VOS HEAP MGMT] SHARED POOL not shrinking during shrink of SGA	[[list-patches]](#)	YES	YES	19-SEP-2024

Monthly Recommended Patches

A collection of recommended one-off fixes
provided at monthly intervals
via a single downloadable patch

Quarterly Release Updates

	January	April	July	October	January	April	July	October	January	April	July	October	January	April	July
	2023				**2024**				**2025**				**2026**		
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	21.10.0	21.11.0	21.12.0	21.13.0	21.14.0	21.15.0	21.16.0	21.17.0	21.18.0	21.19.0					
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

	2023			2024										
	October	November	December	January	February	March	April	May	June	July	August	September	October	November
19.21.0	19.21.0	MRP1	MRP2	MRP3	MRP4	MRP5	MRP6							
19.22.0				19.22.0	MRP1	MRP2	MRP3	MRP4	MRP5	MRP6				
19.23.0							19.23.0	MRP1	MRP2	MRP3	MRP4	MRP5	MRP6	
19.24.0										19.24.0	MRP1	MRP2	MRP3	MRP4
19.25.0													19.25.0	MRP1

An MRP is an optional collection of several important one-off patches

- Delivered as a merge patch

An MRP does not change the release number

- Like `v$instance.version_full`

MRPs are cumulative
but only within one MRP line

- Example: 19.21.0 MRP6 contains all previous MRPs done for Oracle 19.21.0

MRPs are Linux only

Monthly Recommended Patches

Introducing Monthly Recommended Patches (MRPs) and FAQ (Doc ID 2898740.1)

Patching News: RURs are gone – long live MRPs (Blog Post)

Oracle Database 19c Important Recommended One-off Patches (Doc ID 555.1)

Oracle Database Patch Maintenance

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)

Visibility: EXTERNAL

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

Search This Document

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`

Installation Tip

```
/home/oracle/stage
├── DPBP
│   ├── 35261302
│   └── PatchSearch.xml
├── MRP
│   ├── 35333937
│   │   ├── 34340632
│   │   ├── 35012562
│   │   ├── 35037877
│   │   ├── 35116995
│   │   └── 35225526
│   └── PatchSearch.xml
├── OJVM
│   ├── 35050341
│   └── PatchSearch.xml
└── RU
    ├── 35042068
    └── PatchSearch.xml
```

ONE SINGLE COMMAND

```
./runInstaller \
  -applyRU /home/oracle/stage/RU/35042068 \
  -applyOneOffs /home/oracle/stage/DPBP/35261302,
                /home/oracle/stage/OJVM/35050341,
                /home/oracle/stage/MRP/34340632,
                ...
                /home/oracle/stage/MRP/35225526
```

Exercise Patching?

Use our brand new Patch Me If You Can LiveLabs
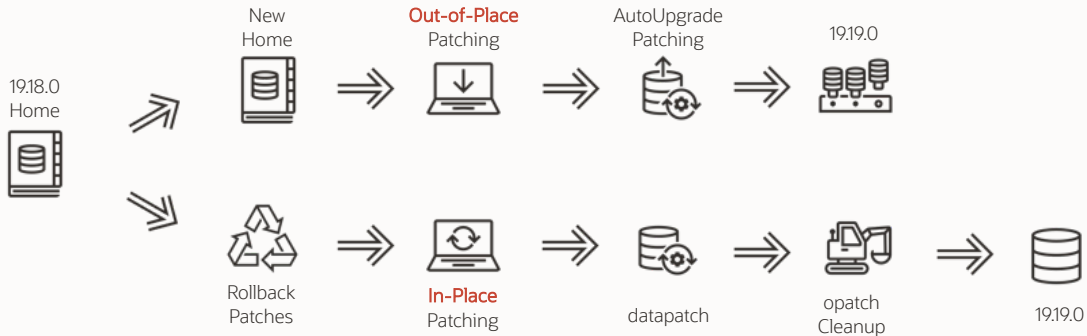
- https://apexapps.oracle.com/pls/apex/dbpm/r/livelabs/view-workshop?wid=3740

19.18.0
Home

New
Home

Out-of-Place
Patching

AutoUpgrade
Patching

19.19.0

Rollback
Patches

In-Place
Patching

datapatch

opatch
Cleanup

19.19.0

Patching Best Practices

Installation

Basics

Methods

Grid Infrastructure

Datapatch

AutoUpgrade
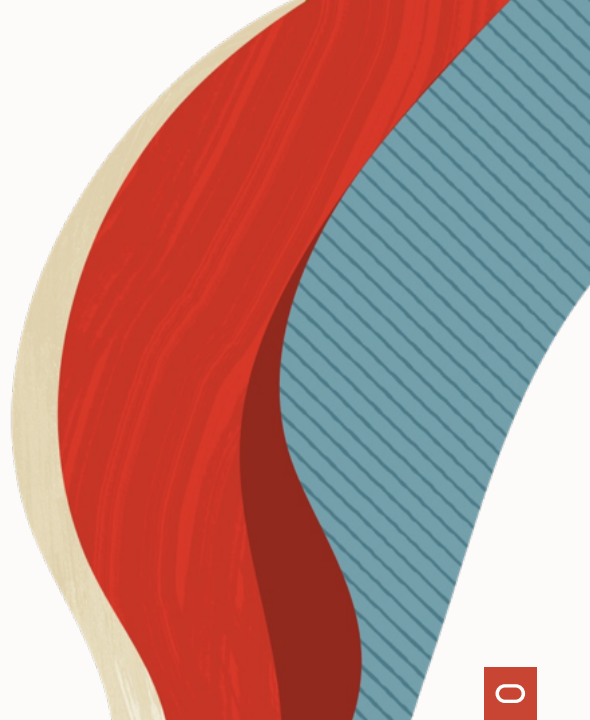
What Can Be in a Patch?

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, 19.2~~5~~.1~~0~~

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

Out-of-Place Patching

Oracle Home, 19.24.0

```
SQL> SHUTDOWN IMMEDIATE
```

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

New Oracle Home
Oracle Home, 19.25.0

👍 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

Users connect here

Snapshot Standby
Users connect here

Oracle Home 19.24.0

Redo Transport

Oracle Home 19.25.0
Oracle Home 19.24.0

Primary

Redo Transport

Redo Transport
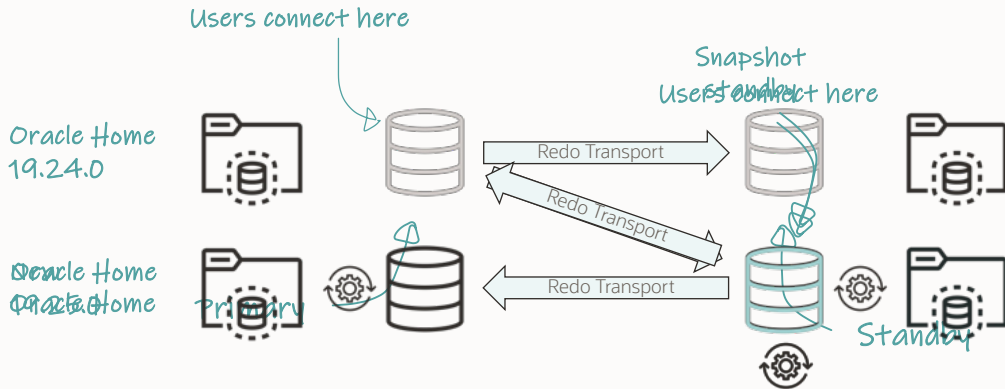
Standby

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

Patch must Standby-First installable

- Check the patch readme

Execute datapatch on the primary database

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

Find additional restrictions in <u>Data Guard Standby-First Patch Apply (Doc ID 1265700.1)</u>

Avoid database downtime with
RAC Rolling Patch Apply

RAC Rolling Patching

User

Connections

Instances

Oracle Homes

Nodes

Database
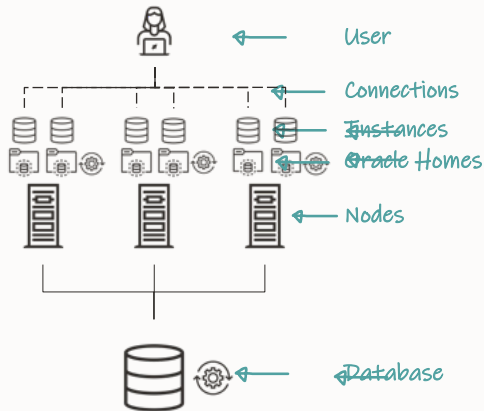
- New Oracle Home
- Patch Oracle Home
- Move to new Oracle Home
- Execute datapatch

Release updates are **always**:

Standby-First installable

RAC Rolling installable

Patching Best Practices

Installation

Basics

Methods

Grid Infrastructure

Datapatch

AutoUpgrade
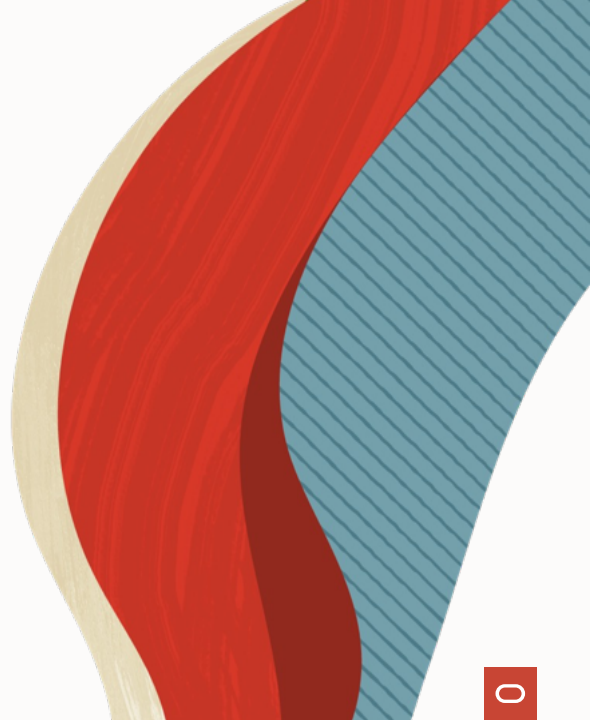
Grid Infrastructure Patching Methods

1 **In-place**
Replaces existing Oracle Home
Uses `opatchauto`

2 **Out-of-place**
Creates a new Oracle Home
Uses `opatchauto` or `gridSetup`

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`

Use Out-Of-Place Patching

- Minimize downtime
- Minimize risk during outage
- Easier rollback

23ai GI home disk space greatly reduced to 3 GB

- 12 GB in 19c

Use golden images

- [Blog post](#)

Golden Images

Unzip
base
release

Replace
OPatch
with newer
version

Deploy
all necessary
patches

Install
with
patches

Create
Gold
Image

Deploy
Gold
Image
to targets
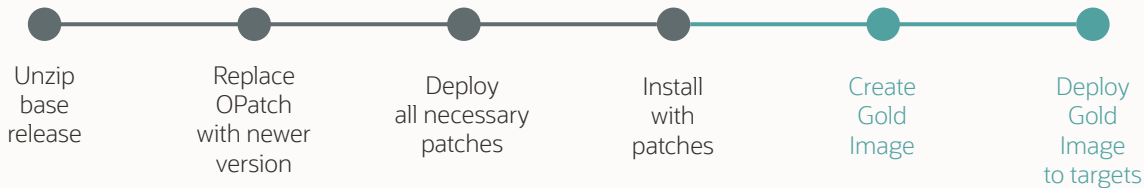
```
--Unzip base release and update OPatch

unzip -oq LINUX.X64_193000_grid_home.zip

mv OPatch OPatch_old
unzip p6880880...zip
```

```
--Install the Oracle Home and apply Release Update and one-offs
--Specify multiple one-offs using comma-separated list

./gridSetup.sh -silent -applyRU <patch_dir> \
                       -applyOneOffs <patch_dir> \
                       ...
```

```
--Always create your golden image from a "fresh" home
--Never use a production home

./gridSetup.sh -createGoldImage \
               -destinationLocation $GOLDIMAGEDIR \
               -silent
```

```
--Deploy golden image throughout your environment
--Just unzip and attach on node 1, installer copies to other nodes

unzip -oq my_golden_image.zip
./gridSetup.sh -silent \
                oracle.install.db.CLUSTER_NODES=node1,node2 \
                ...
```

Demo

Install GI home
Apply Release Update
Create golden image

Watch on YouTube

Works for database homes as well

- Use `runInstaller` instead

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 synch

- This is what we test and run in our Cloud

Unusual combinations are supported, but we strongly advice against it

- GI 19.20.0 and DB 19.23.0
- Node 1 with GI 19.23.0, node 2 with GI 19.22.0
- Patching node 1 on Monday, node 2 on Tuesday …
- Mix of GI and DB versions on various nodes

Complete a rolling patching operation always as quickly as possible

- RAC: Frequently Asked Questions (Doc ID 220970.1)

The following patching concepts
apply to Oracle Database patching as well

Grid Infrastructure Patching Concepts
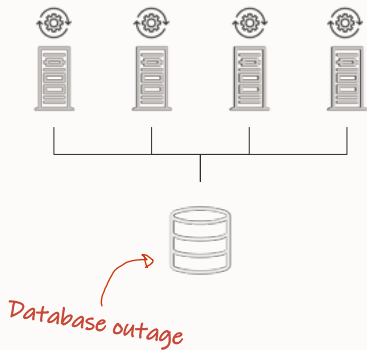
Database outage

ALL NODE

- All nodes patched at one time
- One long database outage
- Works for all patches, including non-rolling
- Cluster at full capacity except for outage

Rolling Patch - OPatch Support for RAC (Doc ID 244241.1)

Grid Infrastructure Patching Concepts
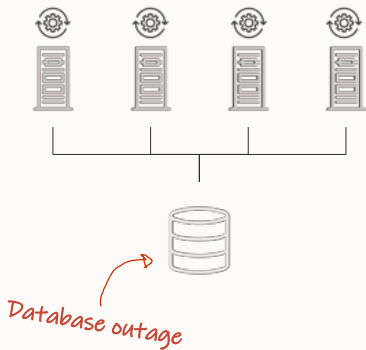
Database outage

MINIMUM DOWNTIME

- Nodes patched in two batches
- One short database outage
- Works for all patches, including non-rolling
- Other nodes must handle workload while another batch is patched

Rolling Patch - OPatch Support for RAC (Doc ID 244241.1)

Grid Infrastructure Patching Concepts
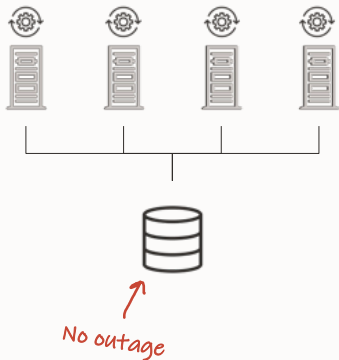
No outage

ROLLING

- Each node patched separately
- No database outage
- Patch must be RAC rolling installable
- Other nodes must handle workload while one node is patched

Rolling Patch - OPatch Support for RAC (Doc ID 244241.1)

Grid Infrastructure Patching Concepts
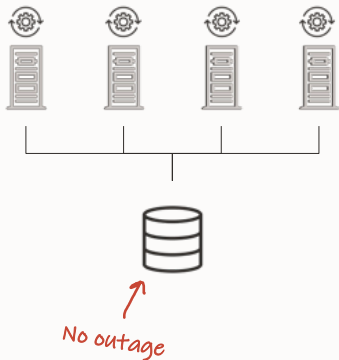
ROLLING IN GROUPS

- Patch a subset together
- Useful when draining is a problem
- No database outage
- Patch must be RAC rolling installable
- Other nodes must handle workload while one node is patched

Rolling Patch - OPatch Support for RAC (Doc ID 244241.1)

Rolling patching requires efficient draining

- Optionally, consider a *batched* approach

Draining Connections

FAN event

Instance going down
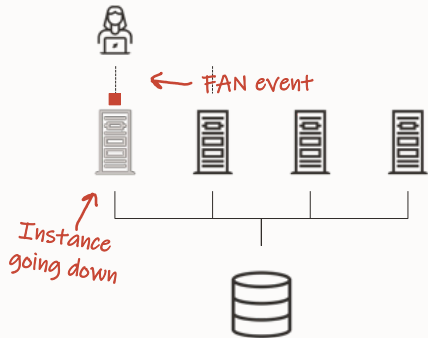
DRAINING

- 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

Comply with Maximum Availability Architecture (MAA) principles

- Continuous Availability - MAA Checklist for Applications for the Oracle Database

Patching Best Practices

Installation

Basics

Methods

Grid Infrastructure

Datapatch

AutoUpgrade
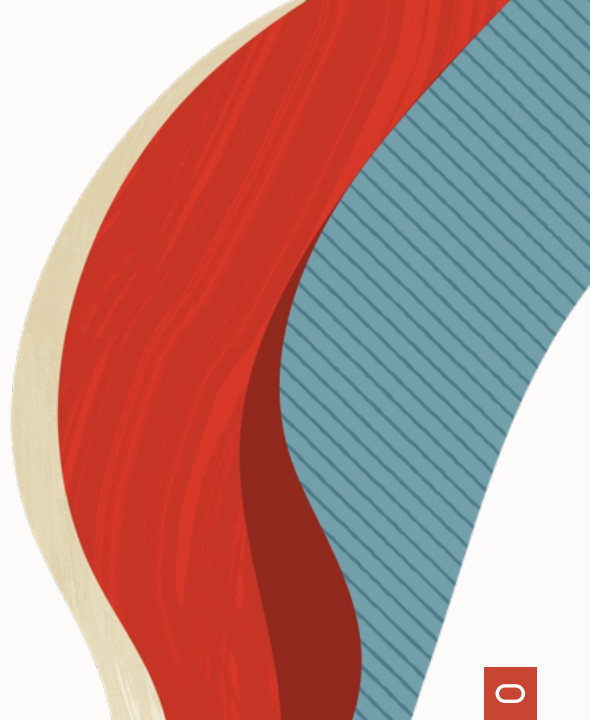
Patching a Database

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)

Patch multiple databases simultaneously
by starting multiple instances of Datapatch

- Each Datapatch works on one database
- Be careful about resource consumption
- AutoUpgrade handles it for you

Analyze the database for patching readiness using Datapatch Sanity Checks

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

```
$ ./datapatch -sanity_checks

...

Check: DB Components status - OK

Check: PDB Violations - OK
Check: System invalid objects - OK
Check: Tablespace Status - OK
Check: Backup jobs - OK
Check: Temp Datafile exists - OK
Check: Datapump running - OK
Check: Container status - OK

Check: Encryption wallet - OK
Check: Dictionary statistics gathering - OK
Check: Scheduled Jobs - NOT OK (WARNING)
  Message: There are current running or scheduled jobs set to run on the next hour.
  Scheduled jobs may have an impact when run during patching.
  :
    JOB_NAME,NEXT_RUN_DATE,SCHEMA_NAME,STATE
    CLEANUP_TRANSIENT_PKG,23-MAY-23 11.08.53.000000 AM +01:00,APPUSER,SCHEDULED
```

Recompile invalid objects
before invoking `datapatch`

Patching Timeline

utlrp

shutdown immediate

Sanity checks
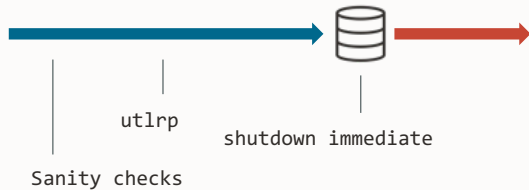
The database must be open
Only open PDBs are patched

- Upgrade mode or restricted session is not needed

Patching Timeline

utlrp

shutdown immediate

Sanity checks

startup [open]
alter pluggable database all open
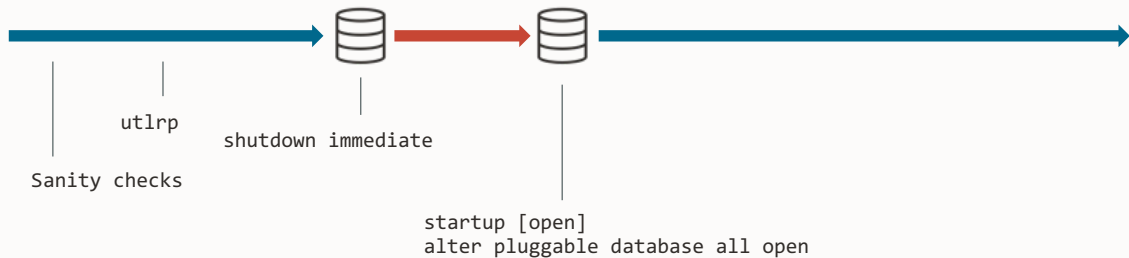
You can run `datapatch` while users are connected to the database

- Details in [blog post](#)

Patching Timeline

Users can connect again

Patching completed

Potential concurrency

shutdown immediate

utlrp

Sanity checks

startup [open]

datapatch

Recompilation

Patching Timeline

exec dbms_stats.gather_table_stats(...

Concurrency

datapatch

create or replace package body
dbms_stats ...
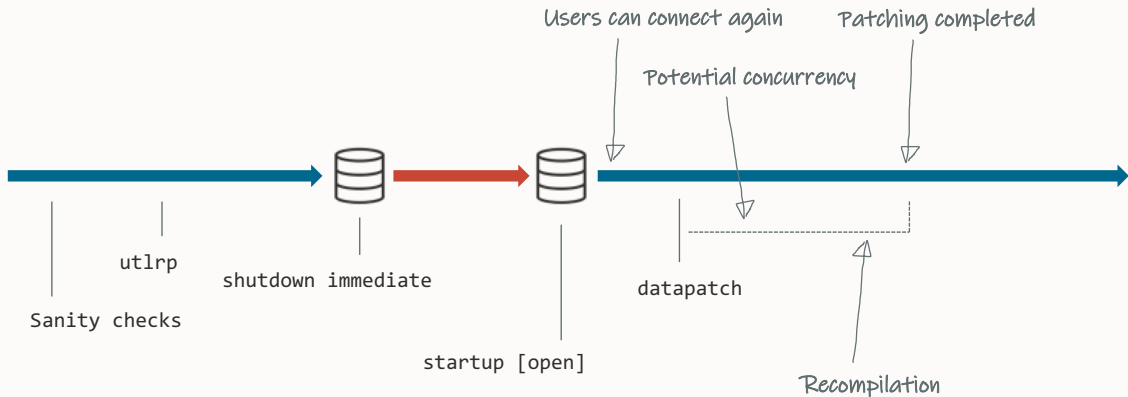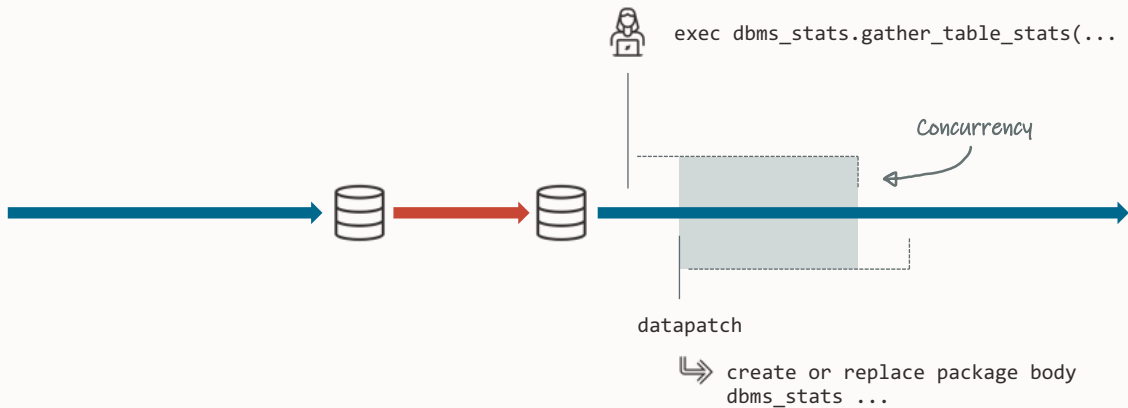
Concurrency

- 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

Recompile manually
if more than 300 objects

Sanity checks

utlrp

shutdown immediate

startup [open]

datapatch

utlrp

Recompilation only
if less than 300 invalid objects
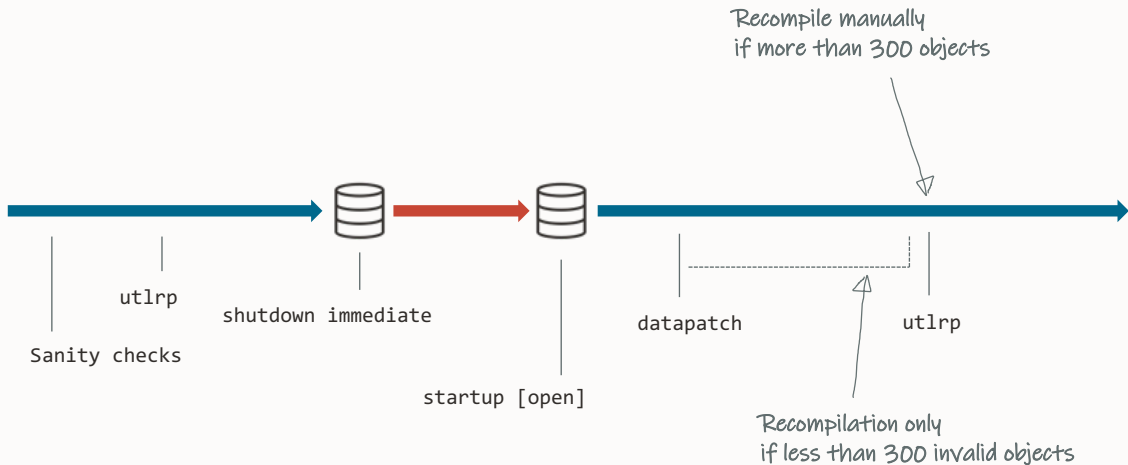
Recompilation

Datapatch recompiles objects invalidated during patching

If more than 300 objects are invalidated no recompilation takes places
- Recompile manually
- Or, objects will be recompiled on usage

Adjust the threshold
```
datapatch ... -recomp_threshold 300
```

Consider recompiling invalid objects after patching

```
$ ./datapatch -verbose
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
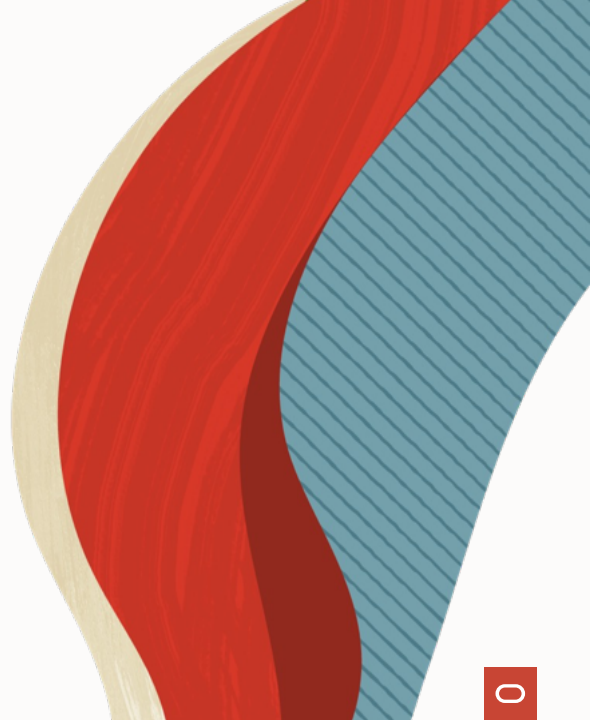
Patch

- 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_24_0
patch1.target_home=/u01/app/oracle/product/19/dbhome_19_25_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_24_0
patch1.target_home=/u01/app/oracle/product/19/dbhome_19_25_0
patch1.sid=DB19
patch1.folder=/home/oracle/autoupgrade-patching/patch
patch1.patch=RU,OPATCH,OJVM,DPBP

$ java -jar autoupgrade.jar -config DB23.cfg –patch -mode deploy
```

Demo

One-Button Patching

- Download patches
- Install Oracle home
- Patch database

Watch on [YouTube](#)

One-Button Patching

Makes life easier for every Oracle DBA

Thursday, 24 October, 14:00

Break
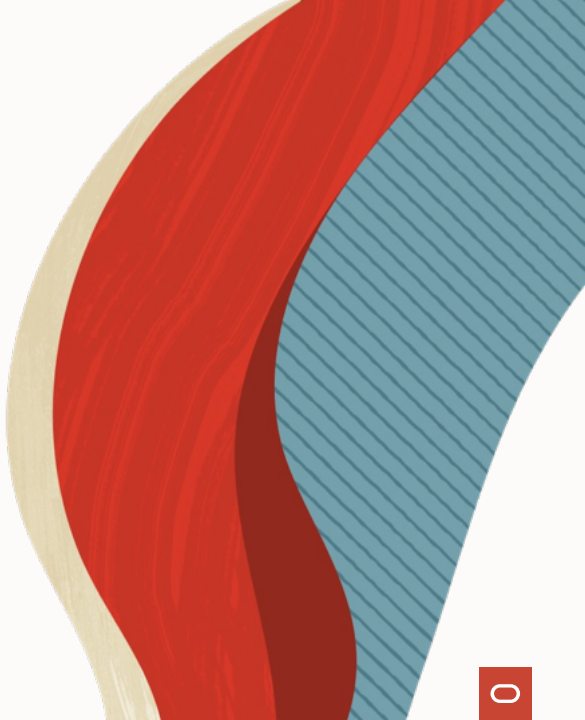
We start again at 11:00

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

$\Longrightarrow$ Oracle Database 19c $\Longrightarrow$ **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;
```

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: 207303.1 - 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: 207303.1 - 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 [136697.1](#)

```
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;
/
```

```
begin
    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"

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

In Oracle Database 23ai an Oracle Home is read-write by default

- Reverting behavior change from Oracle Database 21c
- Read-only Oracle Home is now an optional configuration

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
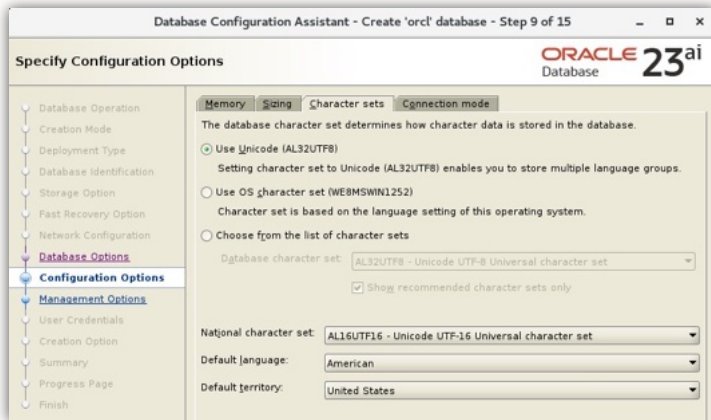
3 COMPATIBLE

Create Container Database

1 Character set

2 Components

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

3 COMPATIBLE
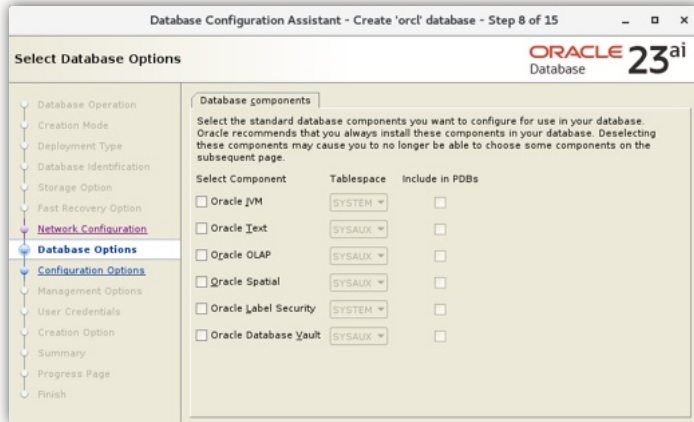
Create Container Database

1 Character set

2 Components

3 COMPATIBLE

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

Compatible

COMPATIBLE=23.4.0   COMPATIBLE=23.5.0   COMPATIBLE=23.6.0

Not possible,
CDB compatible lower

Clone/move OK,
but compatible changes

```
--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

Always download
the latest version of AutoUpgrade

- My Oracle Support Doc ID 2485457.1

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

build.version 24.6.240905
build.date 2024/09/05 11:30:40 -0400
build.hash 0ca273885
build.hash_date 2024/09/05 11:23:06 -0400
build.supported_target_versions 12.2,18,19,21,23
build.type production
build.label (HEAD, origin/devel)
build.MOS_NOTE 2485457.1
build.MOS_LINK https://support.oracle.com/.../?id=2485457.1
```

Flow

1

Plug in

2

Upgrade

3

Convert

$\Longrightarrow$

23^{ai}

Irreversible!
Flashback no good

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

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

Requires Oracle Database 12.2 or newer

Refreshable Clone

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

```
system01.dbf
sysaux01.dbf
users01.dbf
undo01.dbf
...
```
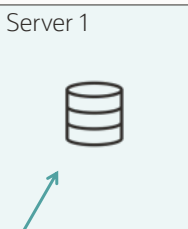
Server 1

Server 2
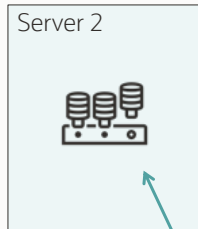
Source
non-CDB

Target
CDB

Refreshable Clone

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

Server 1

Could be same
server as well

Refreshable Clone

Server 1

Server 2

At your will,
~~Source time starts~~ Source preserved
~~down, but preserved~~
for rollback

Clone refreshed ~~Globaled name~~
and ~~nalld~~ formatted and informed

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

Copy data files
over database link

Database link

Source database
is still online

CREATE PLUGGABLE DATABASE ... @CLONEPDB
REFRESH MODE EVERY 30 MINUTES

Open PDB and run
noncdb_to_pdb.sql

Repeat every
30 minutes

Refreshable Clone

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.source_dblink.NONCDB1=CLONEPDB
upg1.target_pdb_name.NONCDB1=PDB1
```

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.source_dblink.NONCDB1=CLONEPDB 300
upg1.target_pdb_name.NONCDB1=PDB1
```

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.source_dblink.NONCDB1=CLONEPDB 300
upg1.target_pdb_name.NONCDB1=PDB1
upg1.start_time=18/05/2024 02:00:00
--Specify relative start time
--upg1.start_time=+1h30m
```

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.source_dblink.NONCDB1=CLONEPDB 300
upg1.target_pdb_name.NONCDB1=PDB1
upg1.start_time=18/05/2024 02:00:00
upg1.parallel_pdb_creation_clause=4
```

Copyright © 2024, Oracle and/or its affiliates

Refreshable Clone

1

Run on source

```
autoupgrade.jar ... –mode analyze
autoupgrade.jar ... –mode fixups
```

2

Run on target

```
autoupgrade.jar ... –mode deploy
```

Refreshable Clone

1.	2.	3.	4.	5.
PDB is created	Data files are copied	Redo is applied	Final refresh	Disconnect and convert

`autoupgrade.jar ... -mode deploy`

`upg1.start_time=18/05/2024 02:00:00`

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	A reliable partner for over 150 years
Project	
Constraints	• The bank for the people of Zurich since 1870
Preparation	
Migration	• With over 5'100 employees one of the largest employers in the canton of Zurich
Success?	
Remarks	• 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

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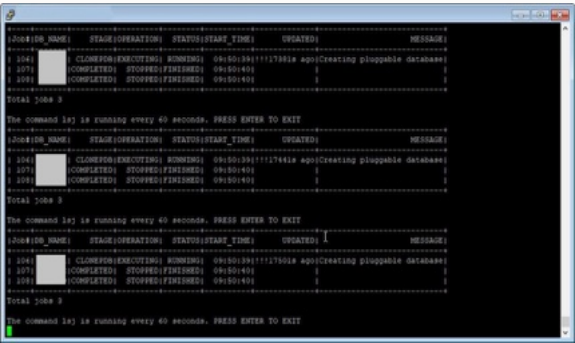
Source	Runtime/Min
TEST40 (108)	26
TEST42 (107)	ongoing
TEST41 (106)	ongoing

Customer Case | Zürcher Kantonalbank

Customer

Project

Constraints

Preparation

Migration

Success?

Remarks

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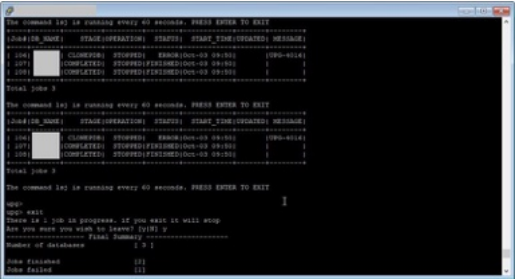
Source	Runtime/Min
TEST40 (108)	26
TEST42 (107)	226 (~3.5h)
TEST41 (106)	1770 (29h)

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

Christain Marquardt
Database Specialist Expert & Solucation Architect

Before upgrade

How to upgrade and convert

After upgrade
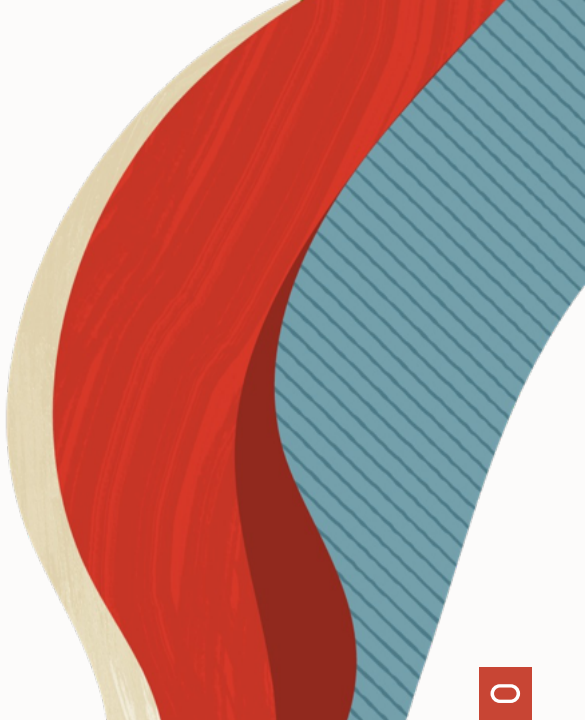
Fallback Options | After Go-Live

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

Source non-CDB

Server 1

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

Standby

Server 2

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

Target CDB

Plug-in and re-use data files

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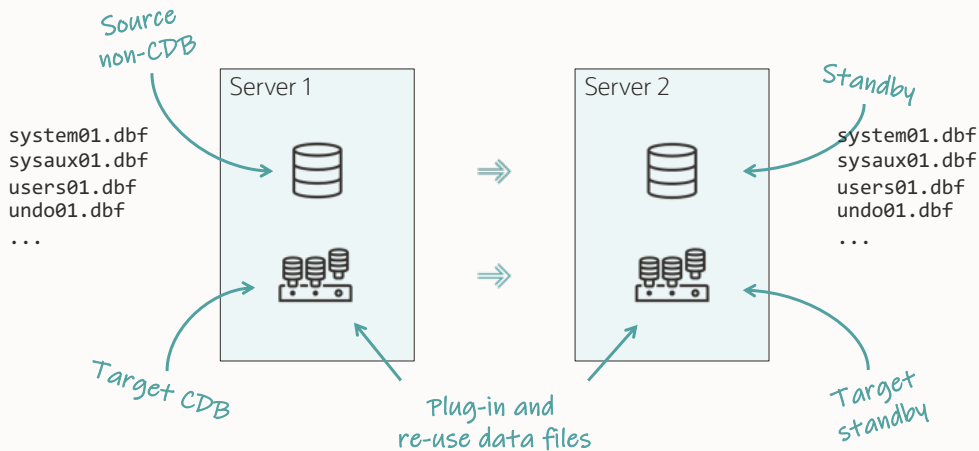
All data files on primary and standby must be at the same SCN

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

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

Enabled Recovery

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

1   Regular files

2   OMF in regular file system

3   ASM

Enabled Recovery

1   Regular files

- Standby search for data files at the same location as the primary

- Override with `DB_FILE_NAME_CONVERT`

- Or, override with `STANDBY_PDB_SOURCE_FILE_DIRECTORY`

Enabled Recovery

2   OMF in regular file system

- Standby search for data files at the OMF location (`DB_CREATE_FILE_DEST`)

- Move data files from non-CDB location into OMF location

- Or, create soft links in OMF location pointing to data file location

Enabled Recovery

3   ASM

- Standby search for data files at the OMF location (`DB_CREATE_FILE_DEST`)

- Use ASM aliases to find data files at non-CDB OMF location

Enabled Recovery | ASM

Primary

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

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

⇓

Standby

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

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

*Same file,
but different name*

```
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

Data Guard | Enabled Recovery

[Reusing the Source Standby Database Files When Plugging a PDB into the Primary Database of a Data Guard Configuration (Doc ID 2273829.1)](#)

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

Source non-CDB

Standby

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

Server 1

Server 2

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

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

Target CDB

Plug-in and re-use data files

Restore data files

Target standby
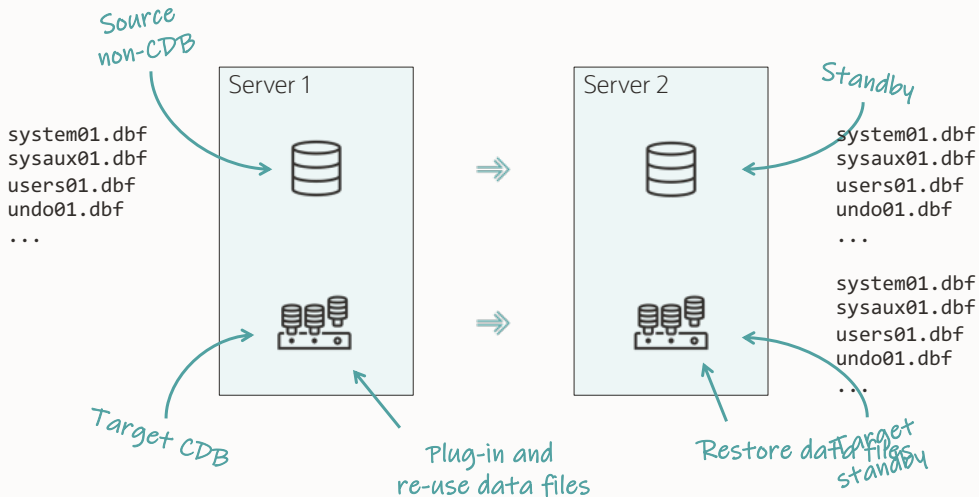
Deferred Recovery

Target
Primary

Source
Non-CDB

Target
Standby

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

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

Data Guard | Deferred Recovery

Making Use Deferred PDB Recovery and the STANDBYS=NONE Feature with Oracle Multitenant (Doc ID 1916648.1)

⭐ **Making Use Deferred PDB Recovery and the STANDBYS=NONE Feature with Oracle Multitenant (Doc ID 1916648.1)**

In this Document

Goal
Solution
  Creating a PDB with the STANDBYS=NONE clause in a Data Guard configuration with 1 physical standby
  Showing how the cloned PDB will appear in certain tables and views on the physical standby
  Performing a Data Guard Role Transition with a PDB in DISABLED RECOVERY
  The zero downtime instantiation process using RMAN for copying the files from the primary to standby
  Steps required for enabling recovery on the PDB after the files have been copied
  Steps to DISABLE RECOVERY of a Pluggable Database
  Conclusion
References

APPLIES TO:

Oracle Cloud Infrastructure - Database Service - Version N/A and later
Oracle Database Cloud Service - Version N/A and later
Oracle Database - Enterprise Edition - Version 12.1.0.2 and later
Oracle Database Cloud Schema Service - Version N/A and later
Oracle Database Exadata Express Cloud Service - Version N/A and later
Information in this document applies to any platform.

Data Guard | Additional Information

Data Guard Impact on Oracle Multitenant Environments (Doc ID 2049127.1)

The physical standby database and redo apply will normally expect a new PDB's datafiles to have been pre-copied to the standby site and be in such a state that redo received from the primary database can be immediately applied. The standby database ignores any file name conversion specification on the CREATE PLUGGABLE DATABASE statement and relies solely on the standby database's initialization parameter settings for DB_CREATE_FILE_DEST and DB_FILE_NAME_CONVERT for locations and file naming.

For these cases, Oracle recommends deferring recovery of the PDB using the STANDBYS=NONE clause on the CREATE PLUGGABLE DATABASE statement. Recovery of the PDB can be enabled at some point in the future once the PDB's data files have been copied from the primary database to the standby database in a manner similar to that documented in Document 1916648.1.

Don't jeopardize your Data Guard

- Test the procedure and verify your environment

```
--Default value is for CDBs with many PDBs
--Other places, 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 many PDBs
--Other places, 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;
```

We need real-world experience with 23ai

- We are looking for reference customers

You can also migrate with
Data Pump or Transportable Tablespaces

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

Lunch Break

We start again at 13:00

Data Pump Top Tips

Supercharge data loading/unloading

Always use
Data Pump Bundle Patch

More than 200 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
```

```
-- 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
```

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		Export parallel 16	1h 8m
Import parallel 1	6h 44m		Import parallel 16	1h 23m
Total	**8h 46m**		**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 <u>DIY parallelism</u> 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=constr_novalidate
```

```
--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

- Doc ID 2819284.1

Even faster index imports

Copyright © 2024, Oracle

New In 23ai

Use index size to determine parallel degree on index creation

- Coming in future 23ai Data Pump Bundle Patch

Index Creation

Before 12.1

Really good for few big indexes

Worker 1　　　　`CREATE INDEX .... 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

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 Create Indexes

2 Create indexes

- One worker creates small indexes (parallel 1) in large batches

- The next worker starts with the biggest index (measured by optimal parallel degree)

How Data Pump Create Indexes

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

Worker 1

| CREATE INDEX ... PARALLEL 1 | CREATE INDEX ... PARALLEL 1 | CREATE INDEX ... PARALLEL 1 | CREATE INDEX ... PARALLEL 1 |

Worker 2

| CREATE INDEX ... PARALLEL 15 | CREATE INDEX ... PARALLEL 8 | CREATE INDEX ... PARALLEL 5 |

Worker 3

CREATE INDEX ... PARALLEL 7

Worker 4

CREATE INDEX ... PARALLEL 3

...

Worker 16
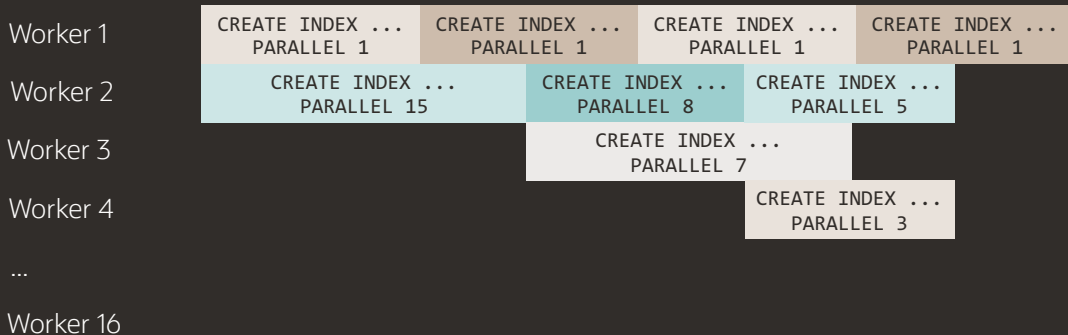
Benchmark, 1 billion rows

Importing with 19c settings constraints

```
10-AUG-24 00:55:35.830: Job "SYSTEM"."SYS_IMPORT_TABLE_01" successfully completed at Sat Aug 10 00:55:35 2024 elapsed 0 00:23:09
```

Importing NOVALIDATE constraints + new index method

```
10-AUG-24 01:48:38.844: Job "SYSTEM"."SYS_IMPORT_TABLE_01" successfully completed at Sat Aug 10 01:48:38 2024 elapsed 0 00:10:40
```
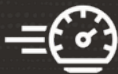
We expect much better result with more complex schemas

We'd love to see this feature in Oracle Database 19c

- Planned for future Data Pump Bundle Patch

Cloud Migration Advisor

Your ultimate migration guidance tool, not only for cloud migrations

Migration Challenges

How do you easily migrate hundreds and
thousands of databases, to the cloud, to ADB or on-prem?

36 migration methods.

Which one is the **best**?

Cloud Migration Advisor

Cloud Migration Advisor

https://www.oracle.com/goto/upgrade

Customer
Fleet

⟹

Download
Extractor
- **CPAT**
- SQL Extractor
- **OEM Extractor**
- Excel sheet

⟹

Collect
Information

⟹

Option 3:
Install CMA in
your VBox
Vagrant Build

Cloud Migration Advisor
OEM Collector

Cloud Migration Advisor
CPAT – Cloud Premigration Advisor Tool

MOS Note: 2758371.1

```
premigration.sh
premigration.cmd
bin
README.txt
misc
lib
p32613591_112048_Generic.zip
```

Shell Script

JSON

JSON

JSON

Discovers information from non-CDB
or PDB individually on a per-DB basis

Start here:

www.oracle.com/goto/upgrade

Cloud Migration Advisor

AutoUpgrade tool for
Oracle Database

**Cloud Migration Advisor
(CMA)**

Database migration
resources

Load data into Oracle
Database

Cloud Migration Advisor

The Oracle Cloud Migration Advisor (CMA) is the best tool for advising you about your migration to Oracle Autonomous Database, Oracle Exadata Cloud at Customer, Oracle Exadata Cloud Service, and other Oracle Cloud Infrastructure (OCI) Database services—and more.

CMA collects database metadata details about your current Oracle Database environment, including patch levels. (No business data or customer information is collected.) CMA uses this database metadata to provide detailed technical advice about possible migration targets and methods.

Choose and download the most appropriate metadata collection method for your situation:

- Cloud Pre-Migration Advisor Tool (CPAT) (Use Oracle SQLcl MIGRATEADVISOR Command, see Jeff Smith's blog post — or if you can't use SQLcl, Download CPAT from My Oracle Support Note 2758371.1 — Documentation (PDF))

- SQL CSV script (Download — Documentation (PDF))

- Oracle Enterprise Manager script (Download — Documentation (PDF))

- Spreadsheet (Download — Documentation (PDF))

The documentation for each method describes how to collect the metadata and share it with your Oracle advisor.
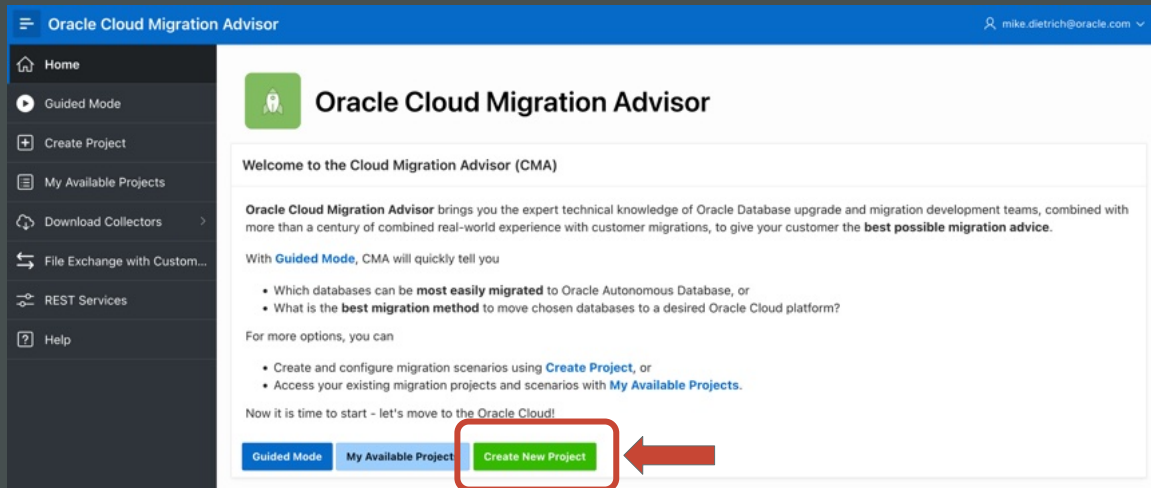
Step 2

Load estate information into CMA

Cloud Migration Advisor
Create New Project

Copyright © 2024, Oracle and/or its affiliates

Cloud Migration Advisor
Example

Databases

Here you have the source databases for this project. Click on "Manage Databases" button to add, modify or edit databases in this list.
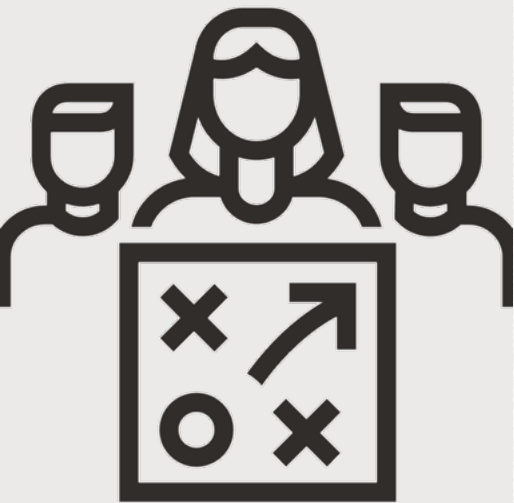
The *Upgrade* and *Accepted Downtime* columns can be used to add information about business requirements for each database.

	Database Display I	Database Version	Block Size	Platform Name	First Hostname	Instances	Upgrade	Accepted Downtime
☐	DB12	12.2.0.1.0		Linux x86 64-bit	hol.localdomain	1		
☐	FTEX	11.2.0.4.0		Linux x86 64-bit	hol.localdomain	1		
☐	PDB1 (CDB2)	19.18.0.0.0		Linux x86 64-bit	hol.localdomain	1		
☐	PDB2 (CDB3)	21.10.0.0.0		Linux x86 64-bit	hol.localdomain	1		
☐	PDB3 (CDB1)	12.2.0.1.0		Linux x86 64-bit	hol.localdomain	1		
☐	UP19	19.18.0.0.0		Linux x86 64-bit	hol.localdomain	1		
☐	UPGR	11.2.0.4.0		Linux x86 64-bit	hol.localdomain	1		
☐	V121	12.1.0.2.0		Linux x86 64-bit	hol.localdomain	1		

|< < 1 > >| 1 - 8 of 8

Cloud Migration Advisor
Example: Estate Overview

Oracle Database Version

Legend:
- 19.0.0.0.0
- 12.2.0.1.0
- 12.2.0.1
- 12.2.0.1.181016
- 12.2.0.1.200114
- 12.1.0.2
- 12.1.0.2.180116
- 12.1.0.2.190716
- 12.1.0.2.180717
- 12.1.0.2.181016
- 12.1.0.2.0
- 11.2.0.4.160119
- 11.2.0.4.180717
- 11.2.0.4.201805
- 11.2.0.4.201020
- 11.2.0.4.2
- 11.2.0.4.190716
- 11.2.0.4.0
- 11.2.0.4.160419
- 11.2.0.3.0

Bar chart values:
- 2 - Oracle Database 10g Release 2: 94, 111
- 3 - Oracle Database 11g Release 2: 78, 63, 300, 185, 37
- 4 - Oracle Database 12c Release 1: 42
- 5 - Oracle Database 12c Release 2: 112
- 6 - Oracle Database 19c: 41

X-axis categories:
1 - Oracle9i Database Release 2, 2 - Oracle Database 10g Release 2, 3 - Oracle Database 11g Release 2, 4 - Oracle Database 12c Release 1, 5 - Oracle Database 12c Release 2, 6 - Oracle Database 19c

Step 3

Add additional information and constraints

Cloud Migration Advisor | Add New Target Server

Copyright © 2024, Oracle and/or its affiliates

Step 4

Group the databases into scenarios and customize methods

Cloud Migration Advisor | Scenarios

Step 5

Create solution

Cloud Migration Advisor | Detailed Solution Report
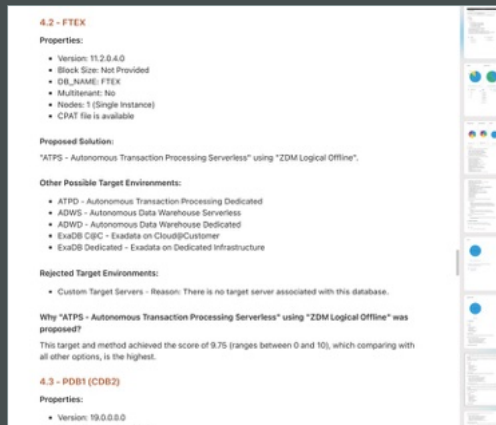
Anything else?

Additional Features

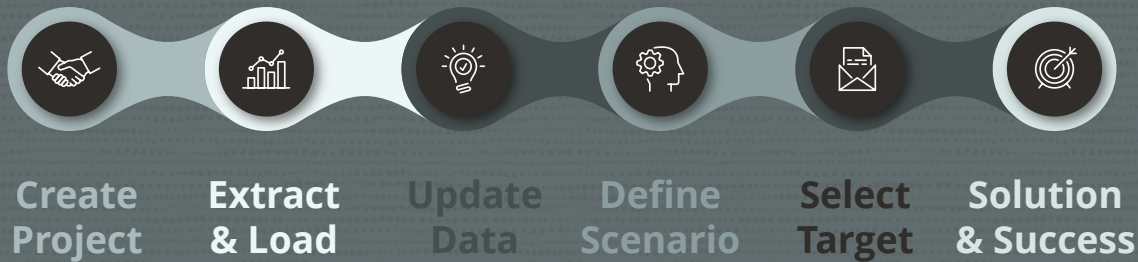# Cloud Migration Advisor | Guided Mode

Cloud Migration Advisor | How to get it?

Downloadable version

- Vagrant Build for a Virtual Box image

- ZIP File for simple deployment in AlwaysFree Autonomous DB

CMA | Workflow

Create Project · **Extract & Load** · **Update Data** · **Define Scenario** · **Select Target** · **Solution & Success**

Performance Stability Prescription

what's your **biggest** fear when making changes?

Photo by nikko macaspac 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);
```
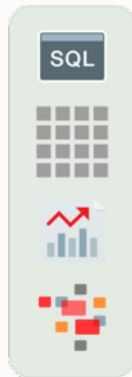
SQL Tuning Set | Definition

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

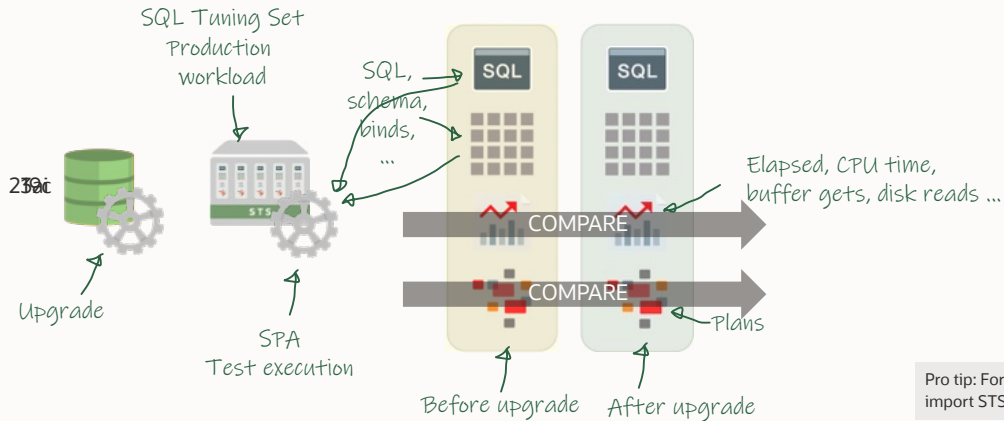Use script `awrddrpt.sql`

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

SQL Tuning Set
Production
workload

SQL, schema, binds, ...

Elapsed, CPU time, buffer gets, disk reads ...

COMPARE

COMPARE

Plans

19c
23ai

Upgrade

SPA
Test execution

Before upgrade

After upgrade

Pro tip: For migrations,
import STS into target database

SQL Performance Analyzer | Report

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

object_id	sql_id	Impact on Workload	Execution Frequency	Metric Before	Metric After	Impact on SQL	Plan Change
52	csv0xdm9c394t	4.02%	3262	6149.0885959534	4208	31.57%	n
41	7m5h0wf6stq0q	2.79%	21694	692.311883470084	490	29.22%	y
34	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					
			Buffer Gets		
SQL ID	Net Impact on Workload (%)	SQL Trial 1	SQL Trial 2	Net Impact on SQL (%)	New Plan
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

			Buffer Gets			
	SQL ID	Net Impact on Workload (%)	SQL Trial 1	SQL Trial 2	Net Impact on SQL (%)	New Plan
⬇	3fv28qfu9v0aq	-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 Collected		
	Execution Statistic Name	Net Impact on Workload (%)	SQL Trial 1	SQL Trial 2	Net Impact on SQL (%)
⬇	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		
3fy28gfu9v0ag	-0.050	26,504	29,573	-11.580	Y
czzzubf8fjz96	-0.030	1,410	1,981	-40.500	Y

Plan Comparison

SQL_TRIAL_1353942463446

Plan Hash Value 1165613724

Expand All | Collapse All

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

SPA | Continuous Improvement

23ai

Implement change

SPA
Test execution

COMPARE

COMPARE

```
--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

```
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

Replace old SQL
Reparse the SQL

Something changed
- Schema
- New statistics
- SQL text
- New parameters
- Upgrade
- Action

OPTIMIZER_USE_SQL_PLAN_BASELINES=TRUE

Plan history
Dedicated part of SQL plan baseline

No plans
in baseline

Plan B
Filter

Plan A is used

Plan A Accepted	Plan B Not accepted	Plan History

SPM | Evolve

Result:
Performance better

Test execute   Test execute

Plan stays

Plan A	Plan C	Plan D	Plan B	Plan D	Plan History
Accepted	Accepted	Accepted	Not accepted	Not accepted	

SPM | Load from STS

SQL Tuning set

Plan C

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

Plan A Accepted Plan C Accepted	Plan B Not accepted Plan History

Automatically accepted

SPM | Use Case

Source
19c

Plans from production

Production
23ai

Test
23ai

Plans from tuning

Load into
SQL plan baseline

Plan baseline

Automatically
accepted
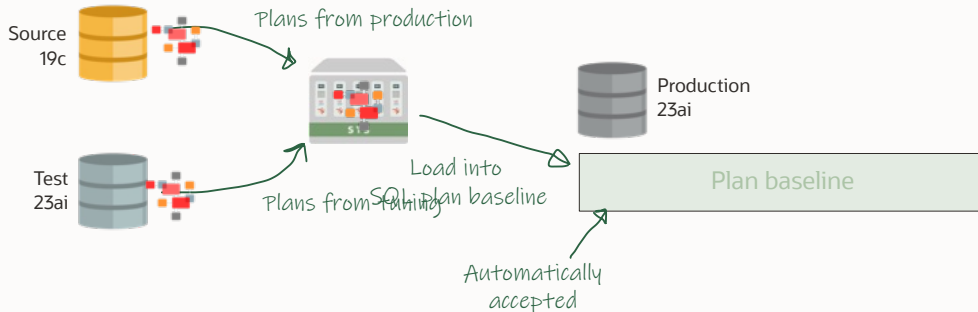
SPM | What If … Literals

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

- Many distinct statements

- `CURSOR_SHARING = FORCE?` <u>No</u>!

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

try it out for free

IT'S EASY

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

- Relevant options and packs included in most cloud offerings

Use OCI for Performance Testing

Export SQL Tuning Set

⇒

⇐

Export SQL Plan Baselines

Generate SPA report
Create SQL Plan Baselines
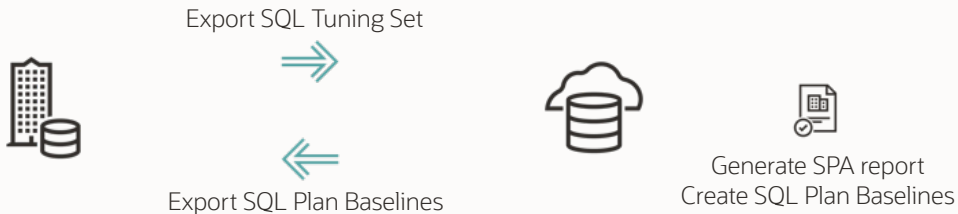
Performance Tips & Tricks

Use as few initialization parameters as possible

- Stick to the defaults
- Stick to vendor recommendations

Only use underscores and events to solve specific situations

- Only under guidance of Oracle Support

Patches For Optimal Performance

1   Install the latest Release Update

2   Install the latest Monthly Recommended Patches

3   Check for important recommended one-off patches (Doc ID 555.1)

4   Check for other SQL performance bug fixes (Doc ID 2773715.1)

Enable Optimizer Fixes

5   Selectively enable optimizer fixes using `DBMS_OPTIM_BUNDLE`

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

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

Should You Enable Optimizer Fixes?

`oradiff.oracle.com`

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

But ... don't worry!

It is fixed! In Oracle 25ai!

Photo by Viktor Forgacs on Unsplash

Hmm??

You mean in
20**25**?

Photo by Andre Hunter on Unsplash

THREE PARTIES

- **You**
- **Oracle Support**
- **Oracle Development**

Photo by Joris Berthelot on Unsplash

Insights into the Patching Process

ERROR → FIX → BACKPORT → INCLUSION

- 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

⚠ ERROR	💡 FIX	🖼 BACKPORT	🎁 INCLUSION

- 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	B - 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

ERROR	FIX	BACKPORT	INCLUSION

- 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

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

Bug Attributes

Type	B - Defect
Severity	2 - Severe Loss of Service
Status	80 - Development to QA/ Fix Delivered Internal
Created	Nov 7, 2022
Updated	Apr 30, 2024
Database Version	23.1
Product Source	Oracle

Fixed in Product Version	**25.1**
Product Version	23.1
Platform	226 - Linux x86-64
Platform Version	ORACLE LINUX GENERIC
Base Bug	33786319
Affects Platforms	Generic

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

Should I
upgrade to
Oracle Database 25ai now?

But you just released 23ai???

Photo by Andre Hunter on Unsplash

No worry ...

... you will request and get a backport!!

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 19.25.0

- `opatch lsinventory`

Branches

Release 19c

Fix is merged into MAIN

Release 25ai

MAIN

25.1

19.25.0

Backport to 19.25.0

19.1

19.25

Insights into the Patching Process

⚠️ ERROR	💡 FIX	🔄 BACKPORT	🎁 INCLUSION
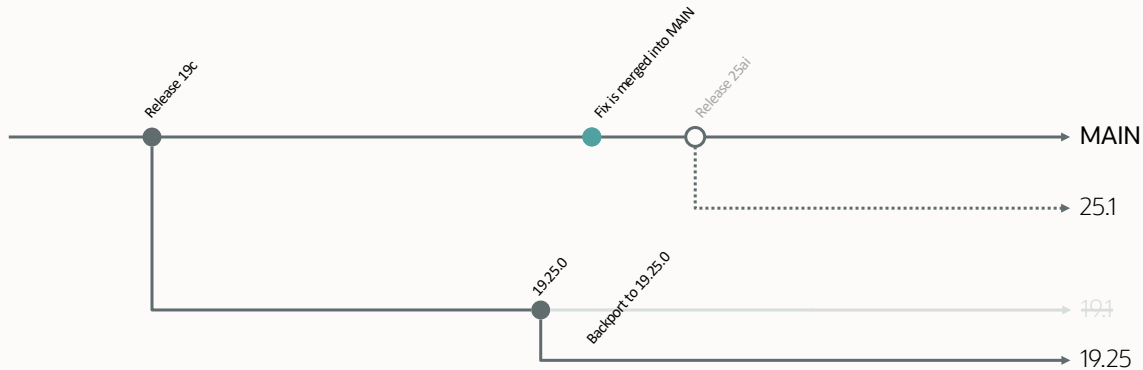
- Automated flow, manual handling might be required
  - Some conflicts
  - Dependencies

- 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?

April 2019 – GA date 19.3.0

January 2022 – 19.14.0

April 2022 – 19.15.0

Time

First three years

Subsequent time under support

12 months after
Release Update release date

24 months after
Release Update release date
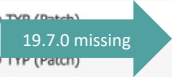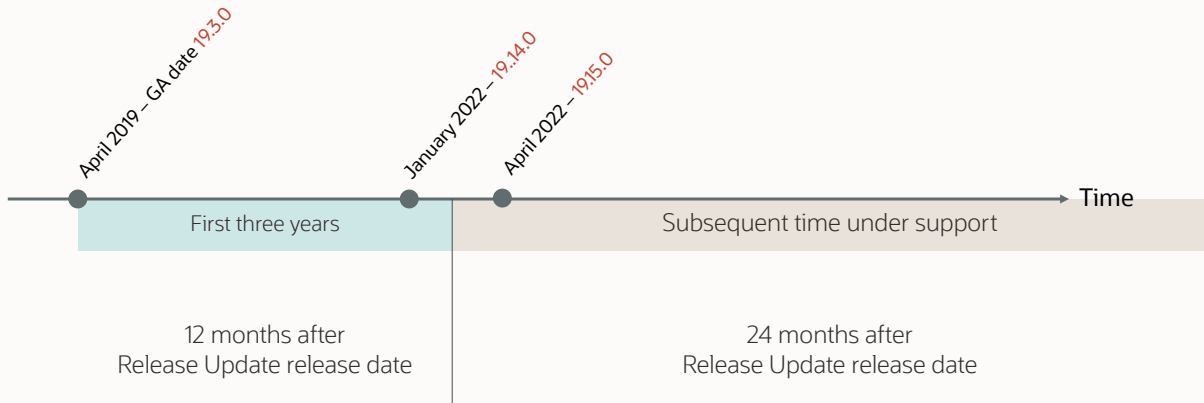
Can I Request a Backport? Examples

JAN 2022
19.14.0

JAN 2023

Time

First three years
12 months RU backporting

Subsequent time under support
24 months RU backporting

JUL 2022
19.16.0

JUL 2024

Time

First three years
12 months RU backporting

Subsequent time under support
24 months RU backporting

JAN 2024
19.22.0

JAN 2026

Time

First three years
12 months RU backporting

Subsequent time under support
24 months RU backporting

Can I Request a Backport?

Based on the Database Release's GA date:
- First three years, eligible for new interim fixes for 12 months from that Release Update's release date
- After that, eligible for new interim fixes for 24 months from that Release Update's release date

Database, FMW, Enterprise Manager, TimesTen In-Memory Database, and OCS Software Error Correction Support Policy (Doc ID 209768.1)

Why is the patch not in the next Release Update?

Insights into the Patching Process

⚠ ERROR > 💡 FIX > 🗔 BACKPORT > 🎁 INCLUSION

- 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 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

⚠ ERROR	💡 FIX	🗔 BACKPORT	🎁 INCLUSION

- 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.

NEVER INCLUDED

(1) Prerequisites

Before you install or deinstall the patch, ensure that you meet the following requirements:

Insights into the Patching Process

| ⚠ ERROR | 💡 FIX | 🗔 BACKPORT | 🎁 INCLUSION |

- Limited number of fixes per Release Update

- Release Update regression tests

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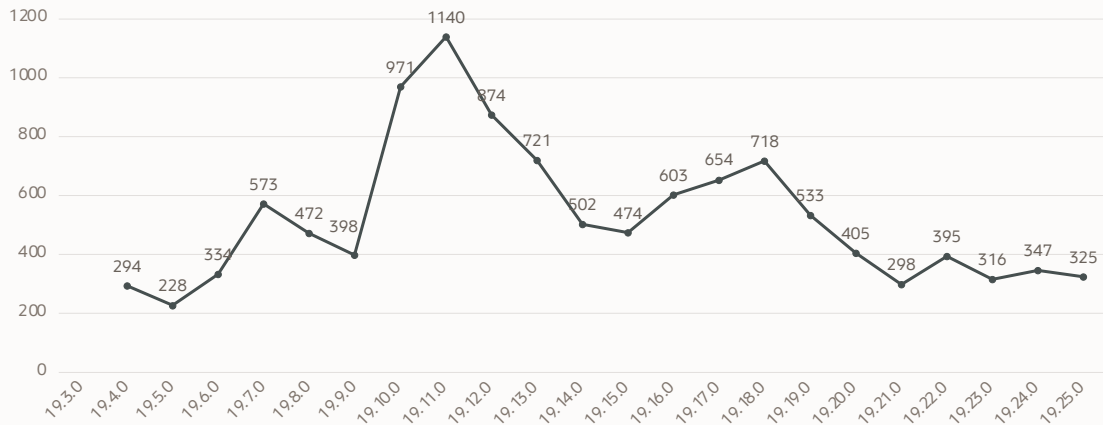
Database 19 Release Updates and Revisions Bugs Fixed Lists (Doc ID 2523220.1)

Limitation
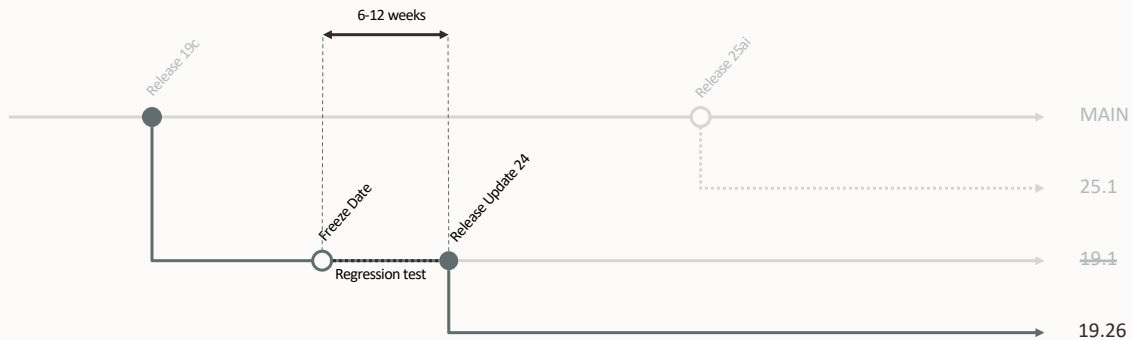
OPATCH

lsinventory

TO AVOID
CONFLICTS

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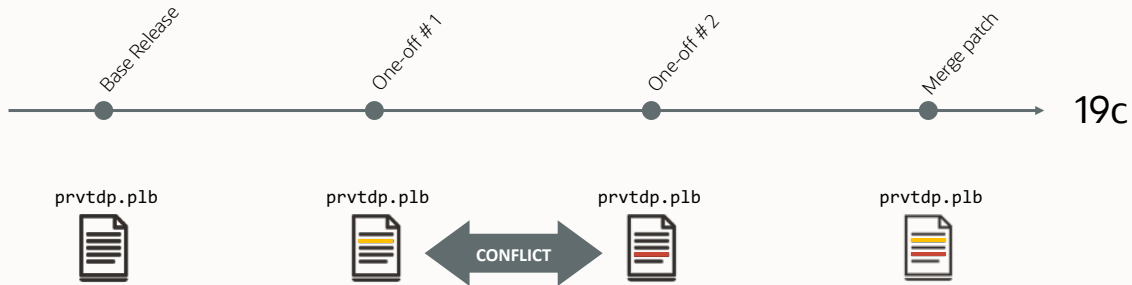
On your Service Request, always attach the output of `opatch lsinventory`

Don't name it

- optch_lsinvt.txt
- tekst_fra_opatch.txt
- alle_meine_patches.txt
- textodeopatch.txt

Just name it `opatch_lsinventory.txt`

But ... I am on MS Windows ...

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

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)

Related Information

19.0.0.0

Description	Bundle Patch	OJVM
19.23.0.0.240416	36219938	36199232
19.22.0.0.240116	35962832	35926646
19.21.0.0.231017	35681552	35648110
19.20.0.0.230718	35348034	35354406
19.19.0.0.230418	35046439	35050341

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:

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

Find the PAD by searching for *Patch Availability Document DB-only* *October 2024*

... Oracle on Windows is different

Photo by Clint Patterson on Unsplash

But this MOS note is missing ...

MOS Note is not accessible

MOS Notes into `UNDER REVIEW` state without telling you more details
- See blog post "[This MOS note is not available anymore?](#)"

Document cannot be displayed. Possible reasons are:

- The document id was entered incorrectly. Please check and try again.
- The document id does not exist (was referenced incorrectly).
- The document is not classified as publicly accessible ("non-public").
- The content is being updated and it is temporarily unavailable but will be made available again soon.

Just try it again a few days

- If urgent, open an SR and check with Oracle Support

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

Break

We start again at 15:15

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?)](#)

Enterprise Manager guided workflow

12.2 → AutoUpgrade → 19c+

Exadata Engineered Systems and Commodity hardware
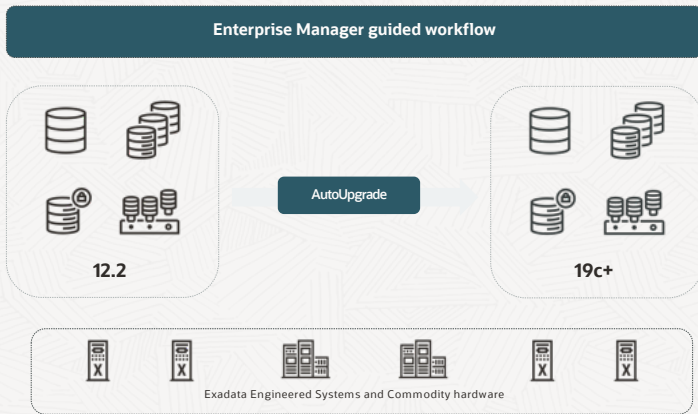
Oracle Database 23ai

What's New

Examine Oracle Database changes using ORAdiff

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

☰ ORAdiff

🖃 MOS Account   👥 Oracle Account   ⑦ ⌄   👤 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

Version Selector

Source Version   Source Patch level
19c   18.17.0

- Select Version -   Target Patch level
                      - Select Patch Level -

V_SPARAMETER

Added   Removed   Changed

🔍 ⌄   Go   Actions ⌄

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](https://youtube.com)

Filter   ?

| Year ▼ | Genre ▼ | Nominations | Awards |

Question

Message

Motivational movies about athletics

35

Search Type

● **Text Search** ○ Vector Search

Search

Results

	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*?

NJQ15253, Batman Returns, 1992, https://upload.wikimedia.org/wikipedia/en/8/83/Batman_returns_poster2.jpg, 'Film-Noir, Action, Family, Fantasy', 'Academy Award for Best Makeup and Hairstyling, Academy Award for Best Visual Effects, MTV Movie Award for Best Villain', , Batman Returns is a 1992 American superhero film directed by Tim Burton and produced by Denise Di Novi and Burton, based on the DC Comics character Batman. The sequel to the 1989 film Batman , it is the second installment of Warner Bros. initial Batman film series, and stars Michael Keaton as Bruce Wayne / Batman, alongside Danny DeVito, Michelle Pfeiffer, Christopher Walken, Michael Gough, Pat Hingle and Michael Murphy. In Batman Returns , Batman faces the Penguin, who plots to kill all of Gotham City's firstborn sons, while dealing with Catwoman, who seeks vengeance against Max Shreck, a corrupt tycoon who allies with the Penguin to bring Gotham City under his control. Burton originally did not want to direct another Batman film. Warner Bros. developed a script with Sam Hamm which had the Penguin and Catwoman going after hidden treasure. Burton agreed to return after they granted him more creative control and replaced Hamm with Daniel Waters. Wesley Strick was later chosen to do an uncredited rewrite shortly before filming. This included normalizing dialogue, fleshing out the Penguin's ==motivations== and master plan, and removing scenes due to budget concerns. Strick continued working as the on-set writer through filming. Annette Bening was originally cast as Catwoman, but became pregnant and was replaced with Pfeiffer. Batman Returns was released on June 19, 1992. It grossed $266.8 million worldwide on a total budget of $80 million and received positive reviews. Critics praised its action sequences, performances, Danny Elfman's score, effects and villains, although its dark tone and high level of violence for a PG-13 film, was criticized. The film was nominated for two Academy Awards: Best Visual Effects and Best Makeup, as well as two BAFTA awards. A stand- alone sequel, Batman Forever , was released in 1995, with Val Kilmer replacing Keaton as Batman. An alternate comic book continuation, which ignores the events of the subsequent films, will be published by DC Comics starting in July 2021. Keaton is also set to reprise the role of Batman in the DC Extended Universe beginning with The Flash (2022).

Filter   ⓘ

Year ▾	Genre ▾	Nominations	Awards

Question

Message
Motivational movies about athletics

35

Search Type

○ Text Search   ● **Vector Search**

Search

1 - 10 of 10

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

Filter

⑦

| Year 2023 ▼ | Genre Sport ▼ | Nominations | Awards |

Question

Message
Motivational movies about athletics

35

Search Type
○ Text Search ◉ Vector Search

Search

Results

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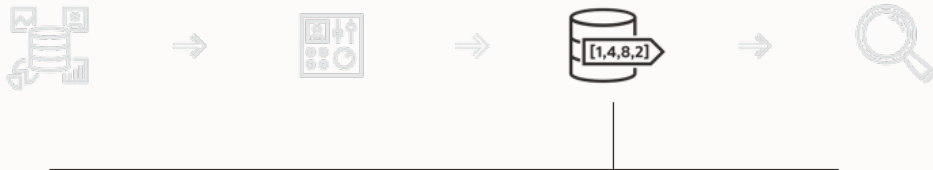
Vector Database

- New data type: `VECTOR`

- A vector is
  a numerical representation of data
  that captures key features and relationships of the data

Vector Database
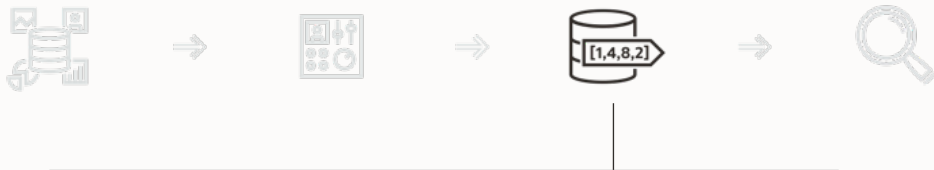

```
CREATE TABLE movies (
     id           NUMBER,
     description CLOB,
     photo        BLOB,
     my_vector    VECTOR(768, FLOAT32));
```

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 $\qquad$ `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

2 6 2 3

Type of movie

2 6 2 3

Genre

2   6   2   3

Mood

| 2 | 6 | 2 | 3 | ⟷ | 3 | 1 | 2 | 8 |

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);
```

d_1

Comedy

The Mask

Ace Ventura

Monty Python

Shrek

Action

Predator

Rambo

Delta Force

James Bond

Band of Brothers

Sci-Fi

Star Wars

Star Trek

Starship Troopers

d_2

d_1

The Mask
Ace Ventura
Monty Python
Shrek

Predator
Rambo
Terminator
Delta Force
James Bond
Band of Brothers

Star Wars
Star Trek
Starship Troopers

d_2
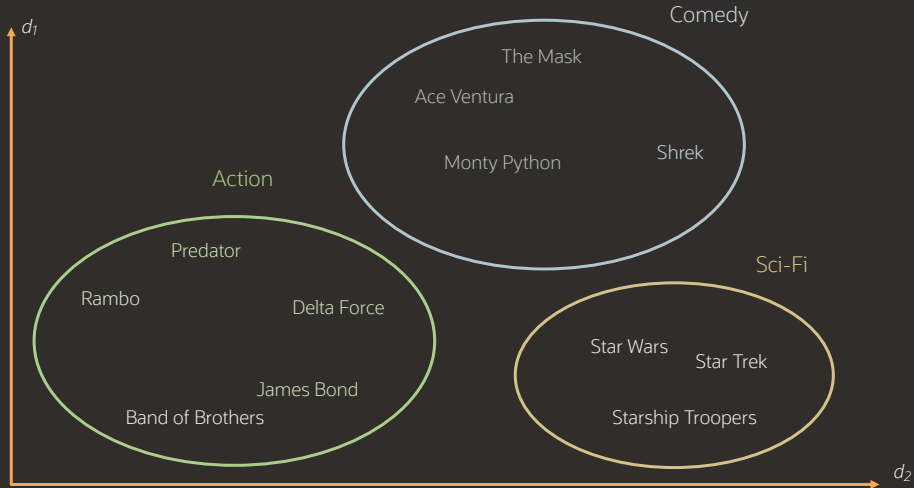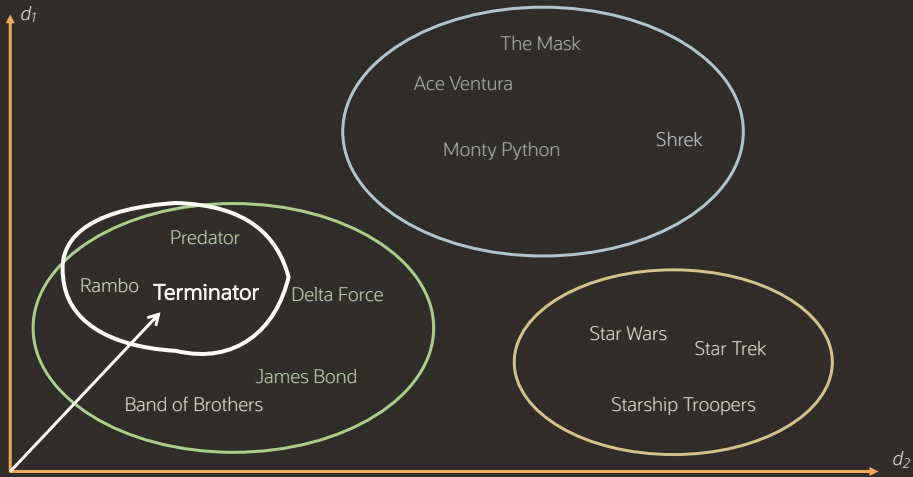
How do you convert your data into vectors?

Vector Database
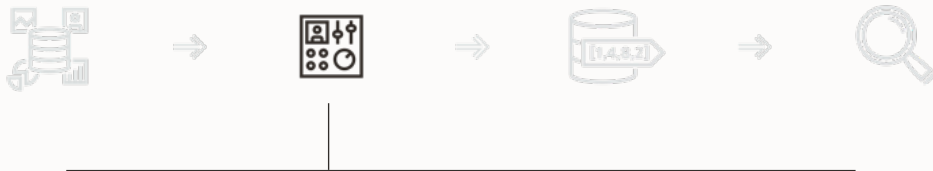
- Embedding models transform your data into a vector
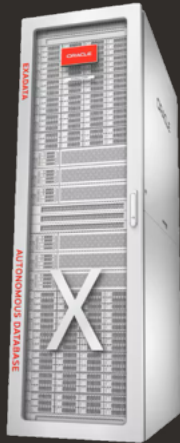
10,000 GPU ?

```
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

Copyright © 2024, Oracle and/or its affiliates
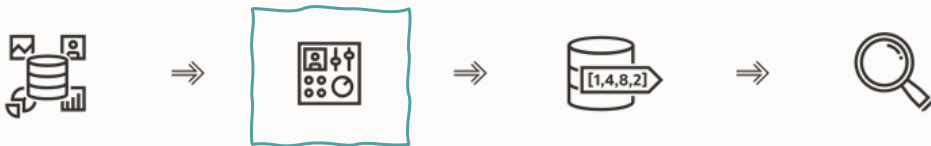
Vector Database

Included in any edition
of Oracle Database 23ai

- No extra license required

Get the best of both worlds
with JSON Duality Views

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

Oracle Database
secures your data

$\Longrightarrow$

Data Guard

RAC

RMAN

TDE Tablespace Encryption

Network Encryption

Auditing

SQL Firewall

ACID

No data duplication

... and so much more
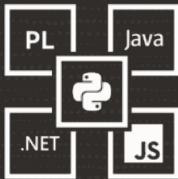
Easily convert your existing application with Oracle Database API for MongoDB

- Using Oracle Database API for MongoDB

Accelerate your applications 10x
with True Cache

True Cache

Application

Writes
Reads

Database
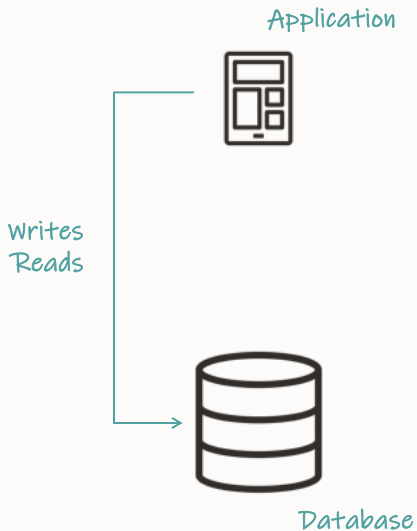
We have a performance problem;
we need a cache!

True Cache

Application

Reads

Object Cache

Writes
Current Reads

Database
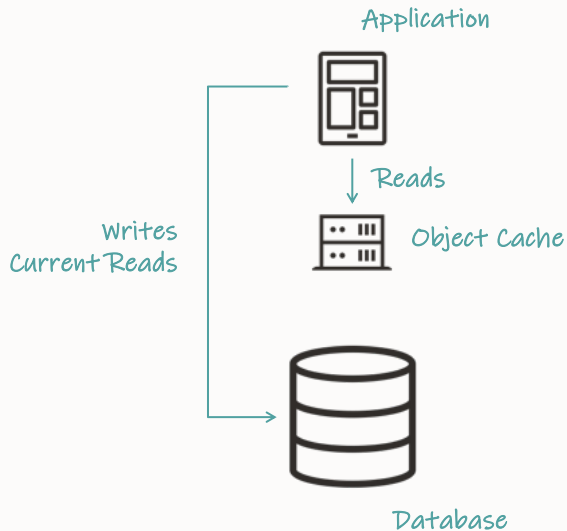
True Cache

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

Application

Reads

Writes
Current Reads

True Cache

Redo

Database

Oracle True Cache is part of Enterprise Edition

True Cache

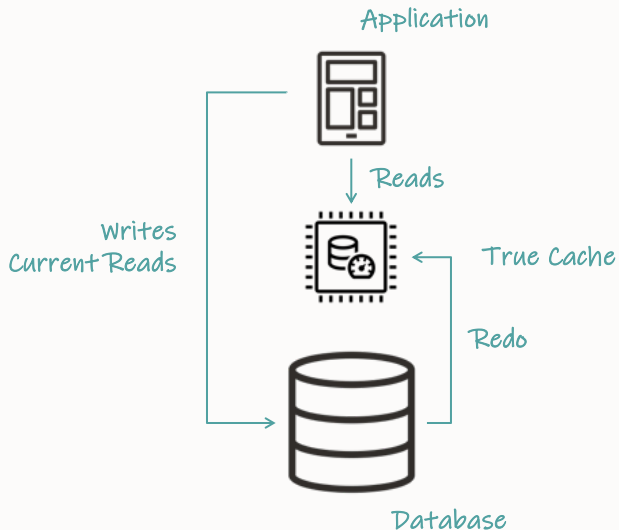| Operation	Conventional Caches	True Cache
Loading the cache	Developer responsibility	**Automated**
Cache consistency with DB	Developer responsibility	**Automated**
Cache consistency with objects in the same cache	Developer responsibility	**Automated**
Cache consistency with other caches	Developer responsibility	**Automated**
Complex data type support	Developer responsibility	**Automated**
Full JSON support	Developer responsibility	**Automated**
Comprehensive security	Developer responsibility	**In-Built**
Parallel processing	Developer responsibility	**In-Built**
High Availability	Developer responsibility	**In-Built**

True Cache

Mid-Tier Cache

Edge Cache

Cross-Region Cache

Cross-Cloud Cache

Quickly resolve poor performance caused by change in execution plan

- Simplified use of SQL Plan Management

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

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

Backported to 19.22.0

Allow the database to fix regressing plans automatically

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

Automatically recover from certain `ORA-00600` errors

- [Automatic Error Mitigation](#)

```
SQL> select ... from ...
ERROR at line 1:
ORA-00600: internal error code, arguments: [...]
```

```
SQL> select ... from ...
ERROR at line 1:
ORA-00600: internal error code, arguments: [...]

SQL> alter session set sql_error_mitigation = 'on';

SQL> select ... from ...
n rows returned
```

Automatically rollback sessions blocking high-priority sessions

- Automatic Transaction Rollback

Session 1

```
alter session set txn_priority=low;
```

Session 2

```
alter session set txn_priority=high;
```

Session 1

```
alter session set txn_priority=low;


--Updates row and goes to lunch
update t1 set c1 = 1000 where id = 1;
```

Session 2

```
alter session set txn_priority=high;
```

Session 1

```
alter session set txn_priority=low;


--Updates row and goes to lunch
update t1 set c1 = 1000 where id = 1;
```

Session 2

```
alter session set txn_priority=high;



--Session waits for row lock
update t1 set c1 = 2000 where id = 1;
```

Session 1

```
alter session set txn_priority=low;
```

```
--Updates row and goes to lunch
update t1 set c1 = 1000 where id = 1;
```

```
--After wait time session is killed
ERROR at line 1:
ORA-03113: end-of-file on communication
channel
```

Session 2

```
alter session set txn_priority=high;
```

```
--Session waits for row lock
update t1 set c1 = 2000 where id = 1;
```

Session 1

```
alter session set txn_priority=low;


--Updates row and goes to lunch
update t1 set c1 = 1000 where id = 1;




--After wait time session is killed
ERROR at line 1:
ORA-03113: end-of-file on communication
channel
```

Session 2

```
alter session set txn_priority=high;




--Session waits for row lock
update t1 set c1 = 2000 where id = 1;




--Row lock acquired
--Update completes
1 row updated.
```

```
--You can enable a tracking mode, which just records the blocking sessions
--Customize the time it takes for a session to be considered blocking

alter system set priority_txns_mode=track;
alter system set priority_txns_high_wait_target=300;
```

```
--You can enable a tracking mode, which just records the blocking sessions
--Customize the time it takes for a session to be considered blocking

alter system set priority_txns_mode=track;
alter system set priority_txns_high_wait_target=300;
```

Automatic Transaction Rollback

1. You can enable a tracking mode instead of killing sessions

2. Monitor blocking events using v$transaction

3. Information gets recorded in the alert log as well

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 Database 23ai

Even More Secure

No more insecure case insensitive password

- 10G password verifies are no longer accepted

Even stronger passwords

- Up to 1024 bytes

Get started quickly and securely using new developer role

- `DB_DEVELOPER_ROLE`

Grant privileges to an entire schema in one command

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

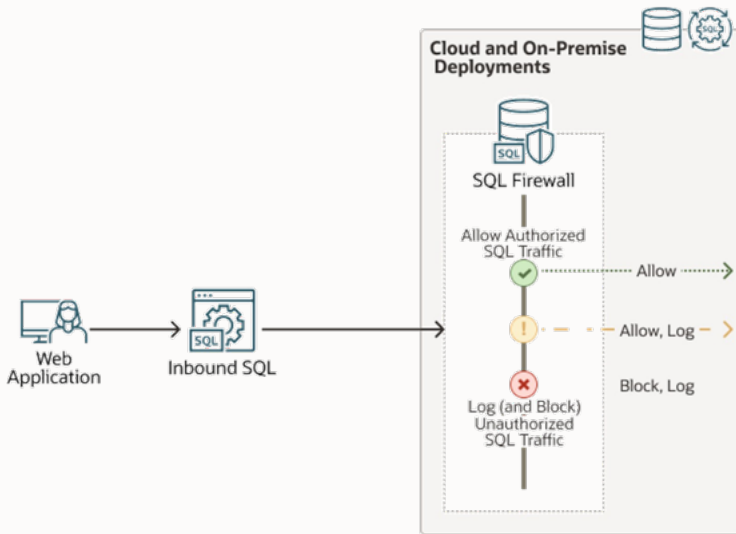By default even stronger encryption algorithms are used

- RMAN backups
- TDE Tablespace Encryption

Tighten security with SQL Firewall

- Prevent SQL injection attacks and unauthorized queries

SQL Firewall

SQL Firewall

For Every Application Account

Capture SQL Activities → Review → Generate Allow-List → Enable Allow-List → Monitor

Learning Stage
- Capture SQL Activities
- Review Captured SQL
- Generate an Allow-List Rule

Protecting Stage
- Enable the Allow-List
- Monitor Possible Violations
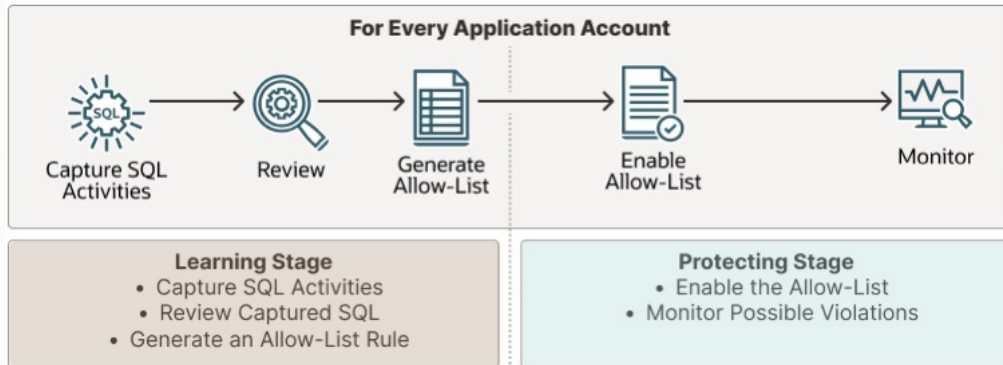
Oracle Database 23ai

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/
```

```
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> oerr ORA-12154
```

Message: "Cannot connect to database. Could not find alias %s in %s."
Help: https://docs.oracle.com/error-help/db/ora-12154/
Cause: A connection to a database or other service was requested using a
         connection alias but the alias specified could not be resolved into
         a connect descriptor using one of the configured naming methods.

Action: Do the following:
        - Check for mistakes in the connection string that you used.
        - If you have a sqlnet.ora file containing a NAMES.DIRECTORY_PATH
          parameter, then ensure that the parameter contains valid values.
        - If you are using an alias from a tnsnames.ora file:
          * Verify that the tnsnames.ora file exists, is in the proper
            directory, and is accessible.
          * Ensure that the alias exists in one of the tnsnames.ora files.
          * Ensure that there are no syntax errors anywhere in the
            tnsnames.ora file or files. Look for unmatched parentheses or
            stray characters. Ensure that magic quotes are not used.

            ...

```
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;
```
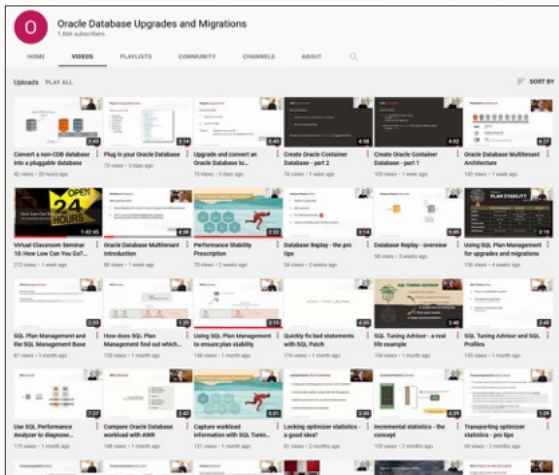
Oracle
DBAs
run the world

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