



ORACLE

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

Brussels – Oct 10, 2023

The Oracle logo is displayed in a bold, orange, sans-serif font. It is positioned in the upper left quadrant of the slide, partially overlapping a decorative graphic consisting of a red and orange patterned shape and a teal shape with white wavy lines.

ORACLE

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

Utrecht – Oct 11, 2023

Safe harbor statement

The following is intended to outline our general product direction. It is intended for information purposes only, and may not be incorporated into any contract. It is not a commitment to deliver any material, code, or functionality, and should not be relied upon in making purchasing decisions. The development, release, timing, and pricing of any features or functionality described for Oracle's products may change and remains at the sole discretion of Oracle Corporation.

The materials in this presentation pertain to Oracle Health, Oracle, Oracle Cerner, and Cerner Enviza which are all wholly owned subsidiaries of Oracle Corporation. Nothing in this presentation should be taken as indicating that any decisions regarding the integration of any EMEA Cerner and/or Enviza entities have been made where an integration has not already occurred.



MIKE DIETRICH

Senior Director Product Management
Database Upgrade, Migrations & Patching



mikedietrich



@mikedietrichde



<https://mikedietrichde.com>





DANIEL OVERBY HANSEN

Senior Principal Product Manager
Database Upgrade, Migrations & Patching



dohdatabase



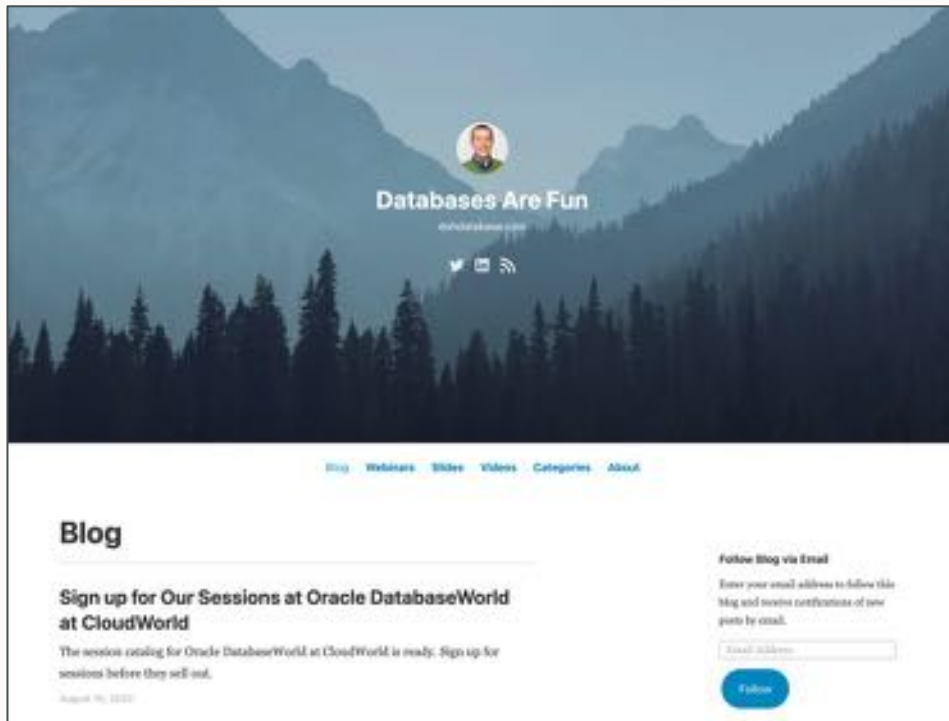
@dohdatabase



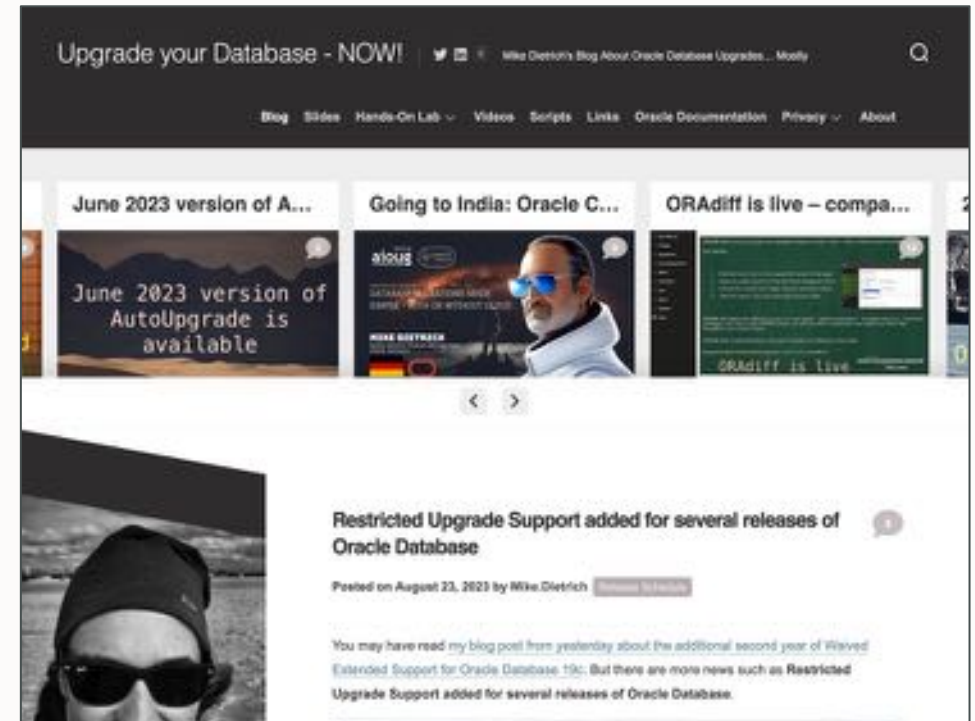
<https://dohdatabase.com>

Find slides and much more on our blogs

dohdatabase.com



MikeDietrichDE.com



Episode 1

Release and Patching Strategy

30 minutes - Apr 4, 2017



Episode 2

Auto Upgrade for Oracle Database 11g

30 minutes - Apr 10, 2017



Episode 3

Performance Statistics, Tips and Tricks and Undercoats

30 minutes - Apr 17, 2017



Episode 4

Migration to Oracle Multitenant

30 minutes - Apr 24, 2017



Episode 5

Migration Strategies - Insights, Tips and Secrets

30 minutes - Apr 24, 2017



Episode 6

Move to the Cloud - Not only for laptops

30 minutes - Apr 4, 2017



Recorded Web Seminars

<https://MikeDietrichDE.com/videos>

More than 30 hours of technical content,
on-demand, anytime, anywhere



AGENDA

09:30

Welcome
Release Strategy
Patching

11:15

Upgrade
Data Pump
Cloud Migration Advisor

13:30

Performance Stability
Insights into
development

15:15

Oracle Database 23c
What's New
What's Coming

11:00

Coffee break

12:45

Lunch

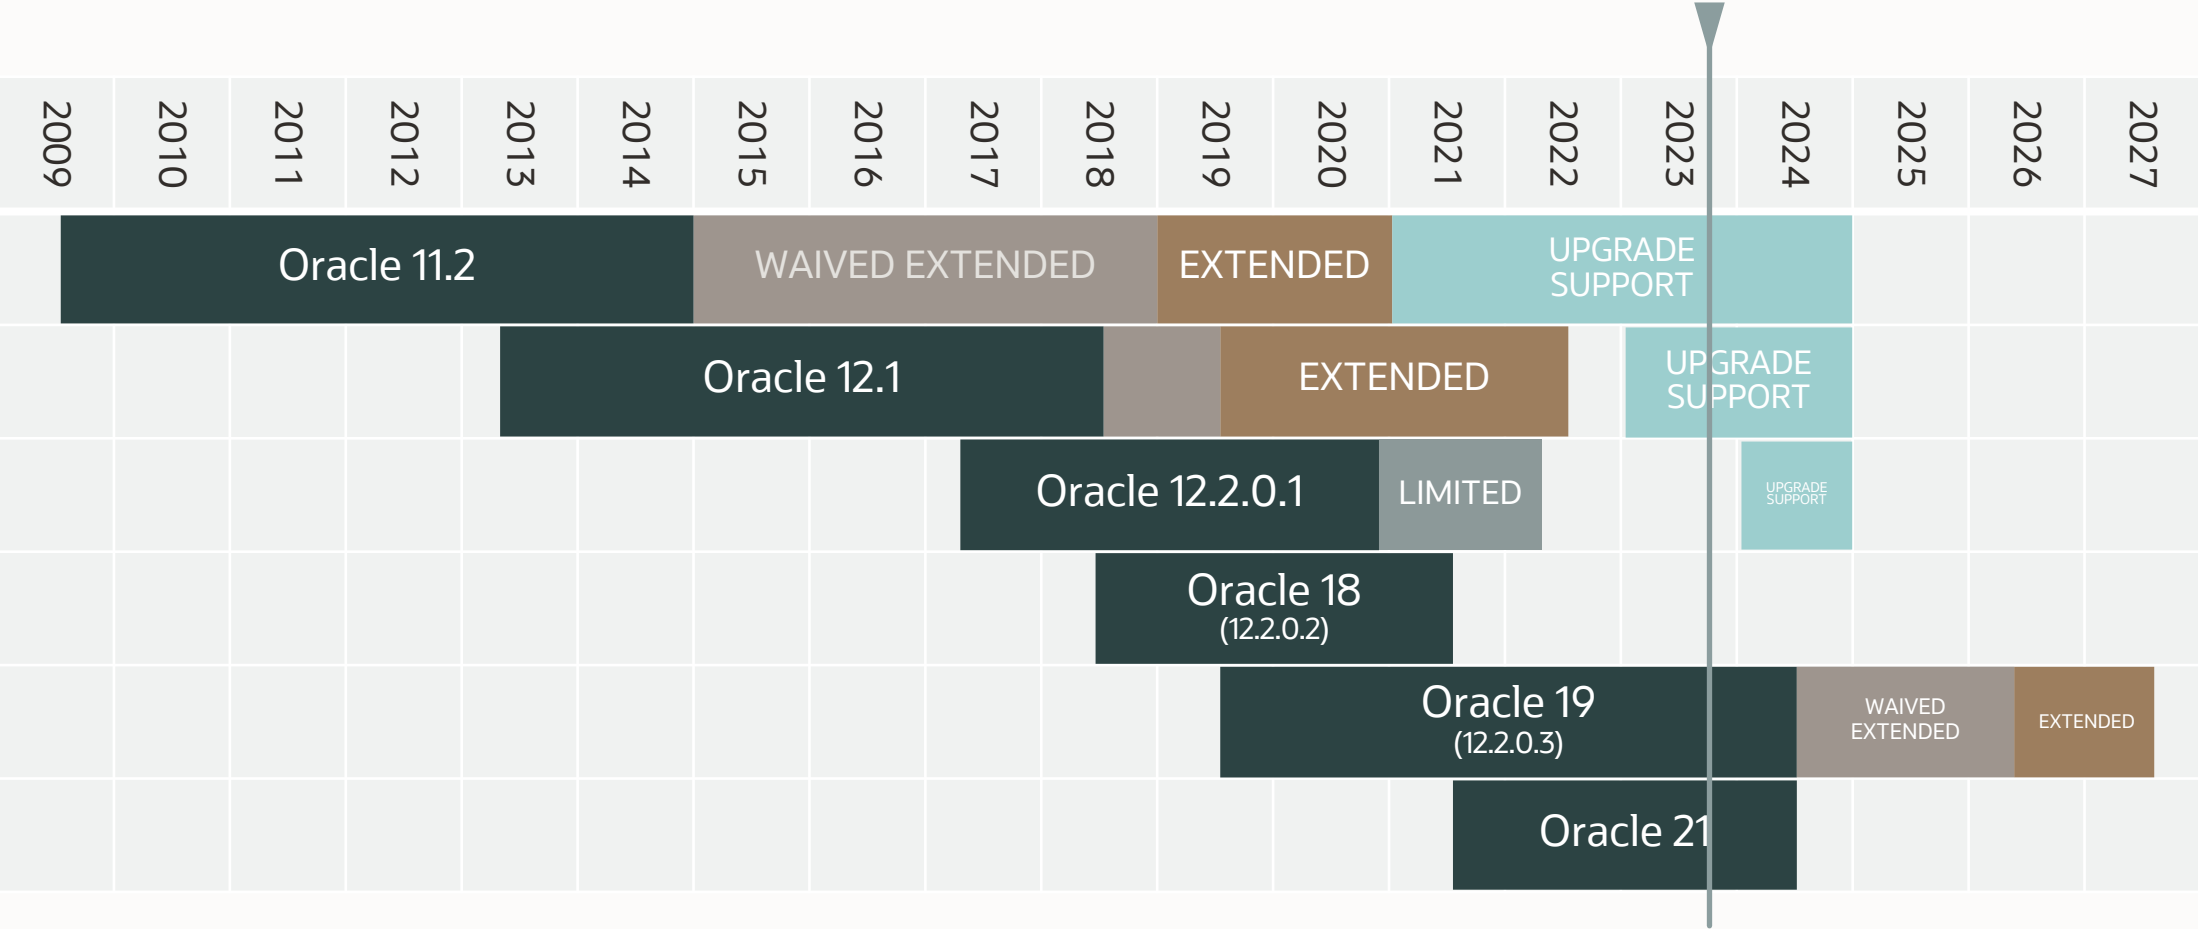
15:00

Coffee break



Release Strategy

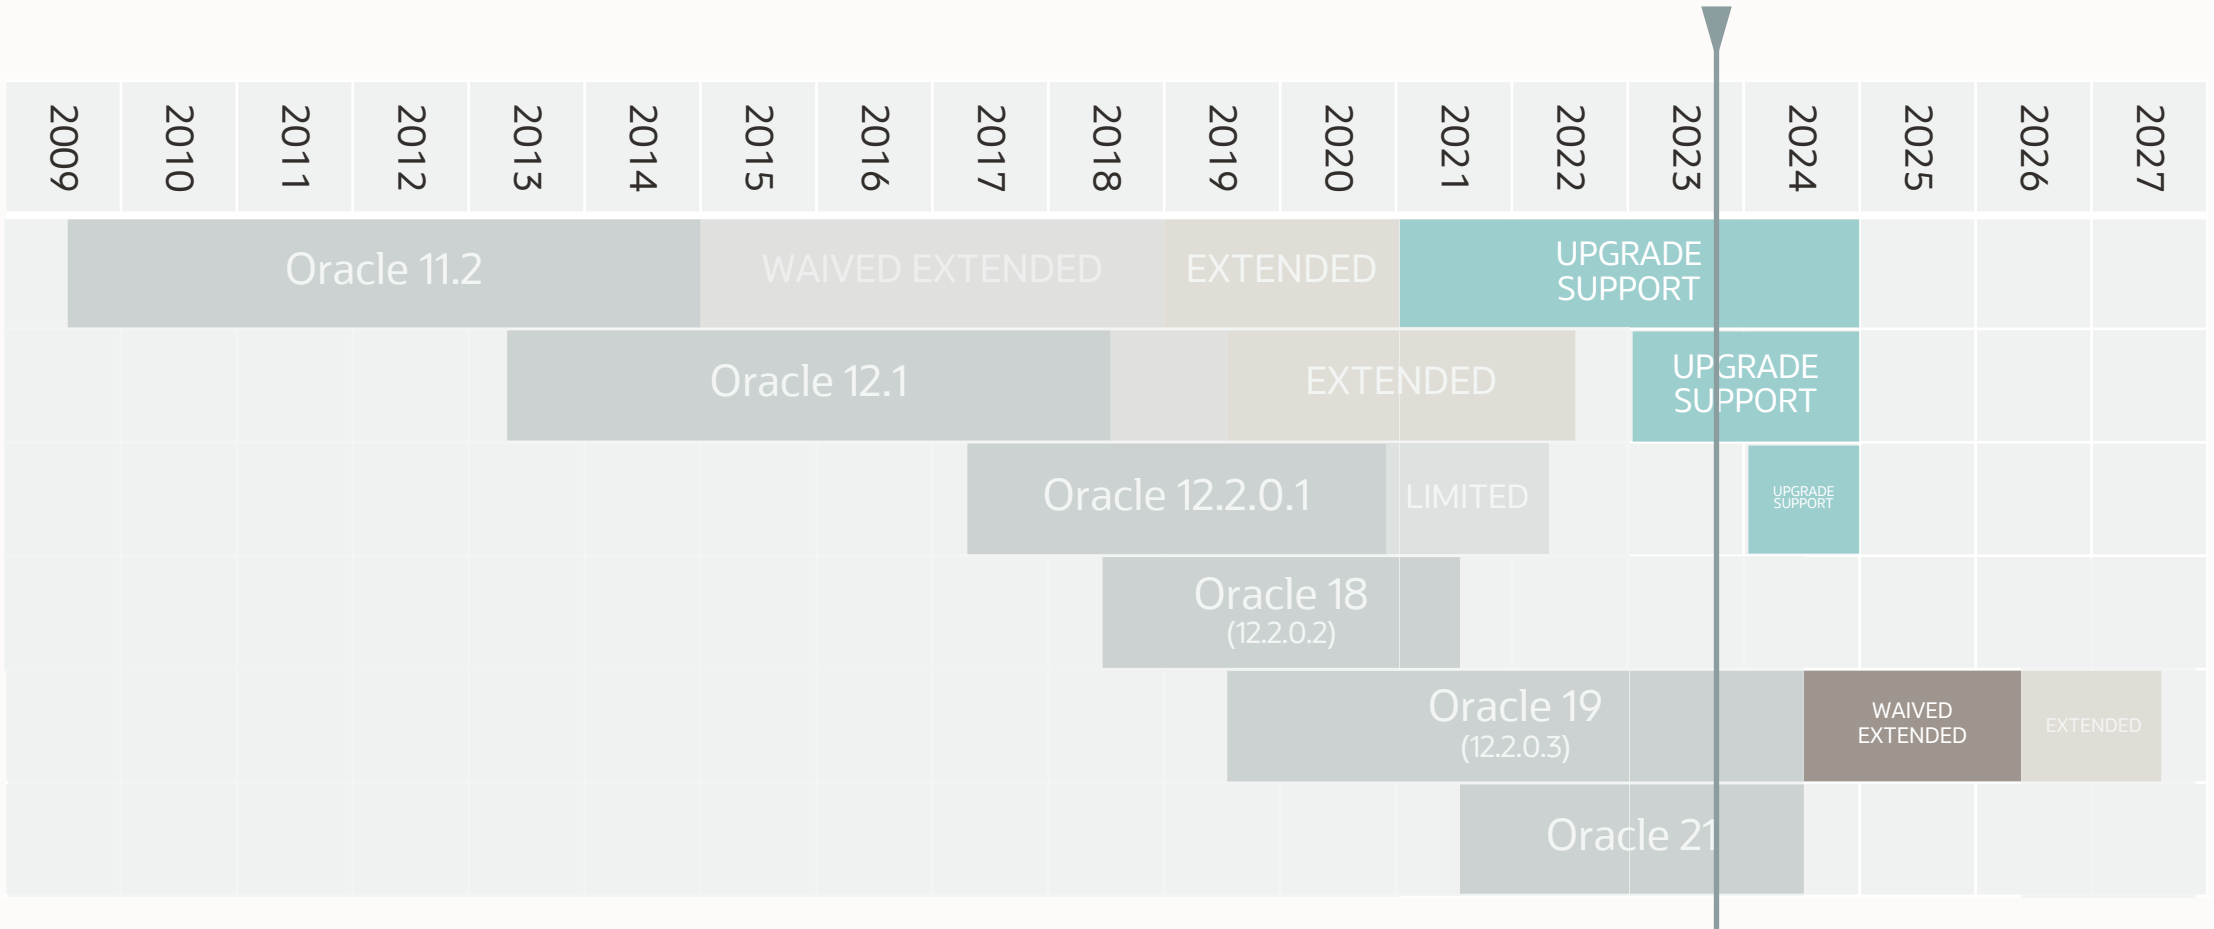
Lifetime Support Policy



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



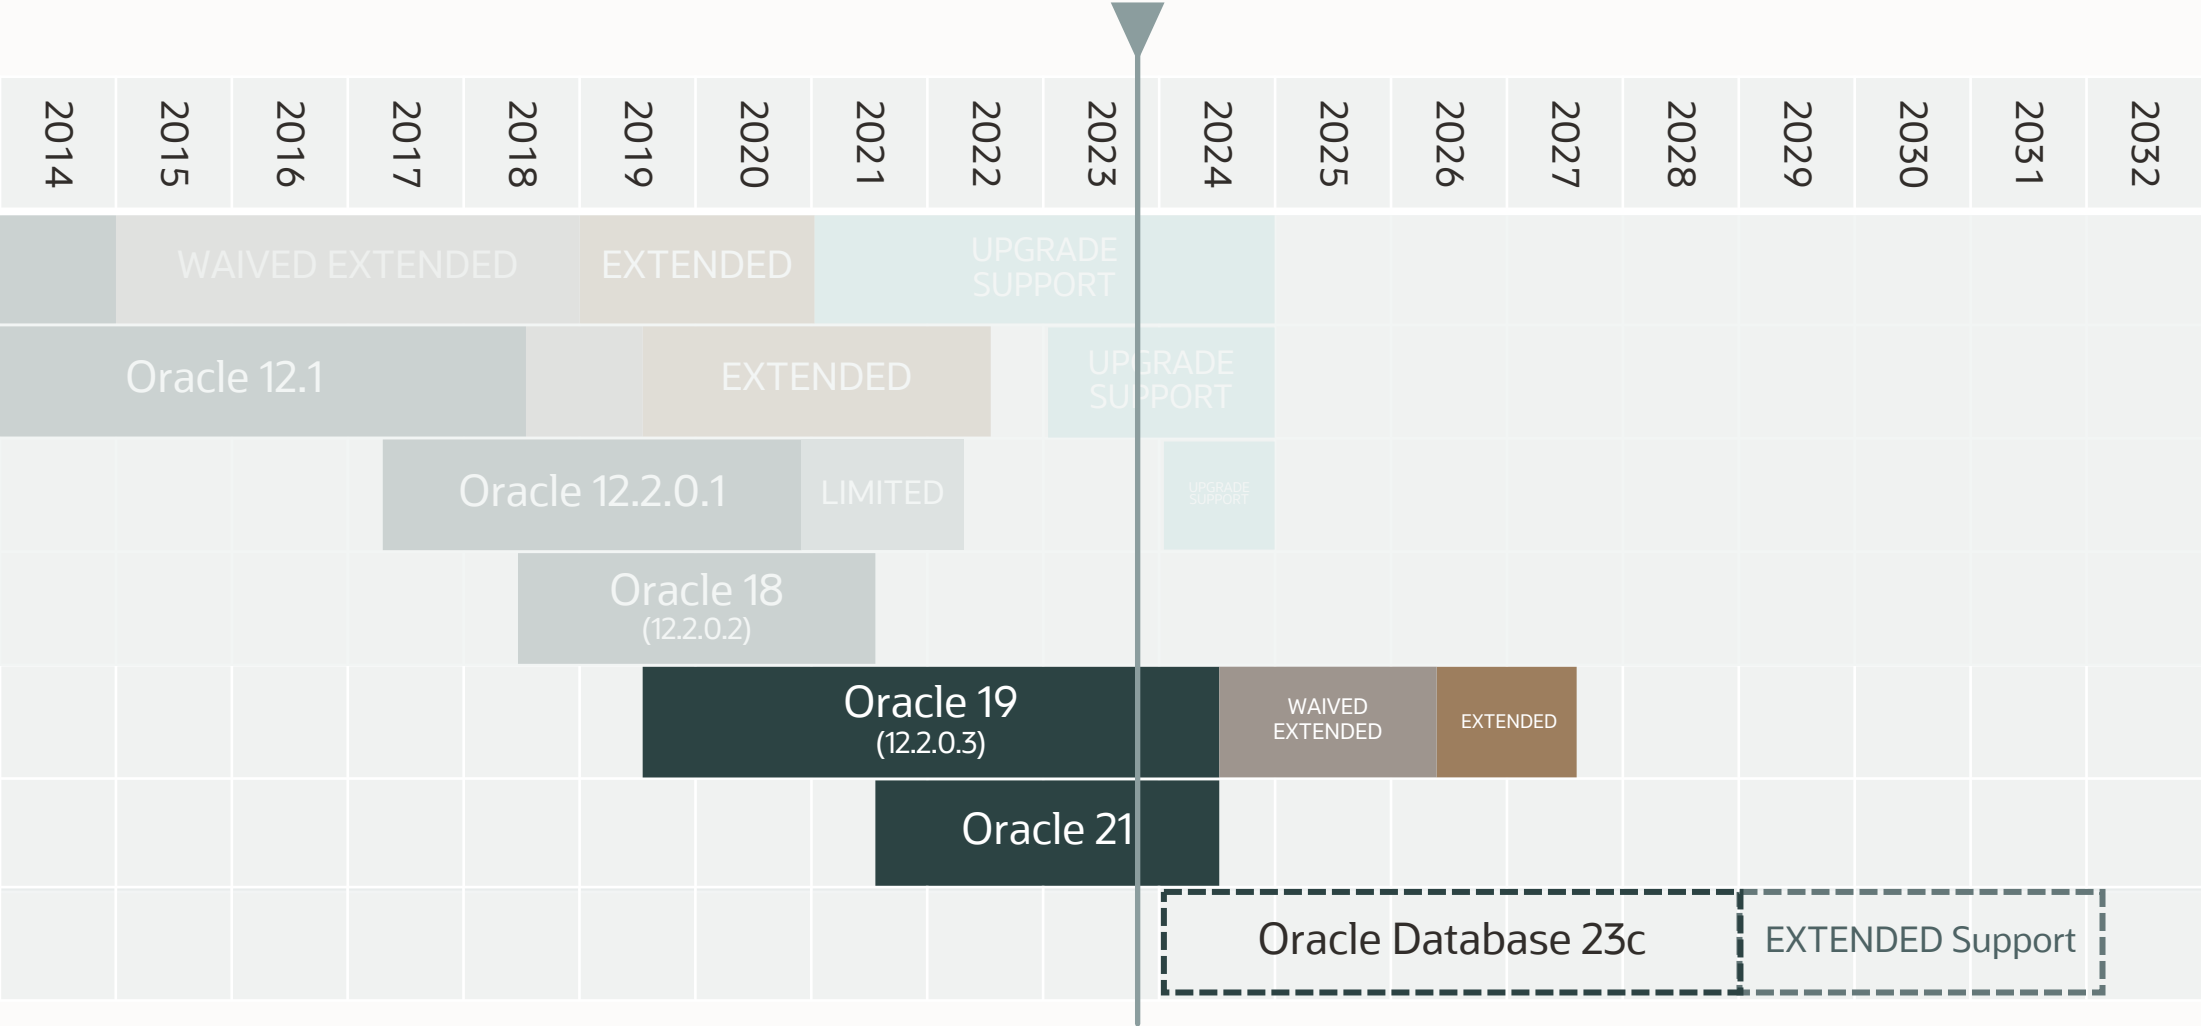
Lifetime Support Policy



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



Lifetime Support Policy





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

Next Long Term Support release

Oracle Database 23c

Upgrade possible only from:

- Oracle Database 19c
- Oracle Database 21c

Database and Grid Infrastructure Patching

Patching Best Practices

Installation

Basics

Methods

Grid Infrastructure

Datapatch

AutoUpgrade



You **always** start with
Oracle Database 19c **base release**

- Oracle Database 19.3.0

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

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

- Almost 10k fixes with 19.20.0
- Over 220 security fixes

Apply the most important patches

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

Recommended Patches for 19.20 DB Home

Below is the list of important patches to consider applying on top of 19.20. In addition to the relevant patches listed below, you should also review patches in [Database PSU/BP/Updates/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	NON ROLLING	Added
35395648		DBMRP 19.19.0.0.230815. DBMRP 19.20.0.0.230815	Domain Name Is Missing For Service Post GI Patch To 19.18	[list-patches]		23-AUG-2023
33974554			[RAM INDEX] ORA-600 [kcbp_prepare_4], Logical block corruption detected, error code 6057 against indexes	[list-patches]		19-AUG-2023
35372179		Not Applicable	[VCS] Linux: ORA-800 / Set Priority / DB Performance Merge Patch for 19.20 (Requires Root Access) - 34286265 34318125	[list-patches]		07-AUG-2023
35156218		DBMRP 19.19.0.0.230815. DBMRP 19.20.0.0.230815	[BLOCK TRACK] CTWR: Instance crashed with ORA-600 [krccso_1] when Block Change Tracking (BCT) in place	[list-patches]		13-JUL-2023
34774667			[AQ] Global-buffer-overflow in pga at kwqalockqwithinfo with ORA-00700 During PURGE_QUEUE_TABLE.	[list-patches]		11-JUL-2023
32627280		DBMRP 19.19.0.0.230815. DBMRP 19.20.0.0.230815	[DICTIONARY] Datapump Export Fails With ORA-39166 And ORA-31655 On Newly Created Tables	[list-patches]		15-JUN-2023
31061145		DBMRP 19.19.0.0.230815. DBMRP 19.20.0.0.230815	[RAC] Remove requirement to set TARGET_PDBS for RAC on multitenant deployments in 18c onwards	[list-patches]		03-JUN-2023

Monthly Recommended Patches

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

Quarterly Release Updates

	2021				2022				2023				2024		
	January	April	July	October	January	April	July	October	January	April	July	October	January	April	July
19c	19.10.0	19.11.0	19.12.0	19.13.0	19.14.0	19.15.0	19.16.0	19.17.0	19.18.0	19.19.0	19.20.0	19.21.0	19.22.0	19.23.0	19.24.0
21c		21.3.0	21.4.0	21.5.0	21.6.0	21.7.0	21.8.0	21.9.0	21.10.0	21.11.0	21.12.0	21.13.0	21.14.0	21.15.0	



Monthly Recommended Patches

	2022			2023										
	October	November	December	January	February	March	April	May	June	July	August	September	October	November
19.17.0	19.17.0	MRP1	MRP2	MRP3	MRP4	MRP5	MRP6							
19.18.0				19.18.0	MRP1	MRP2	MRP3	MRP4	MRP5	MRP6				
19.19.0							19.19.0	MRP1	MRP2	MRP3	MRP4	MRP5	MRP6	
19.20.0										19.20.0	MRP1	MRP2	MRP3	MRP4
19.21.0													19.21.0	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: 19.18.0 MRP5 contains all previous MRPs done for Oracle 19.18.0





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 \(Blog Post\)](#)

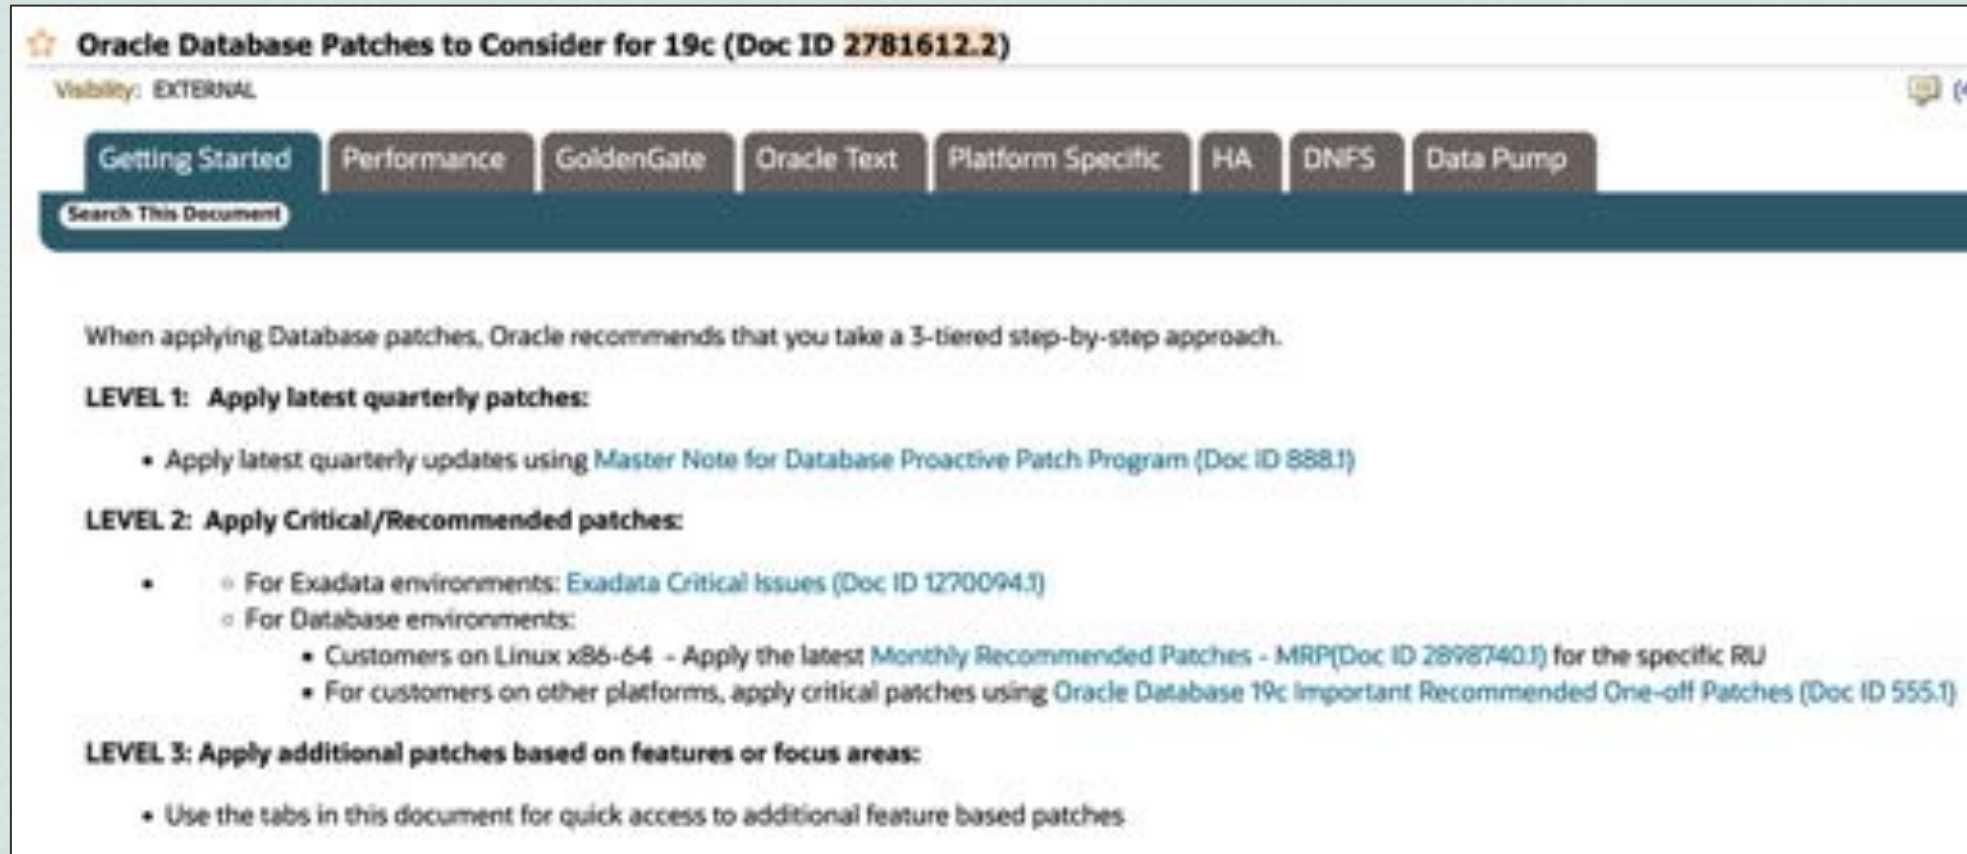
[Oracle Database 19c Important Recommended One-off Patches \(Doc ID 555.1\)](#)

[Oracle Database Patch Maintenance](#)

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

Apply additional important fixes and bundles

In addition, use Patches to Consider for 19c: [MOS Note: 2781612.2](#)



The screenshot shows the top portion of a MOS document titled "Oracle Database Patches to Consider for 19c (Doc ID 2781612.2)". The document is marked as "EXTERNAL" and has a "Visibility" of "EXTERNAL". A navigation bar contains tabs for "Getting Started", "Performance", "GoldenGate", "Oracle Text", "Platform Specific", "HA", "DNFS", and "Data Pump". Below the navigation bar is a search bar labeled "Search This Document". The main content area begins with the text: "When applying Database patches, Oracle recommends that you take a 3-tiered step-by-step approach." This is followed by three levels of patching instructions:

- 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



As of Oracle 19.18.0, **PERL patches** are included in Release Updates

- Consider patching JDK as well (Doc ID 2584628.1)



As of Oracle 19.18.0, **Time Zone patches** are included in Release Updates

- Be aware when you create a new database



Always use the latest OPatch

- [Patch 6880880](#)



Binary patching runtime gets shorter when you use **a brand new home**

- Avoid cloned Oracle Homes and In-Place Patching

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
```

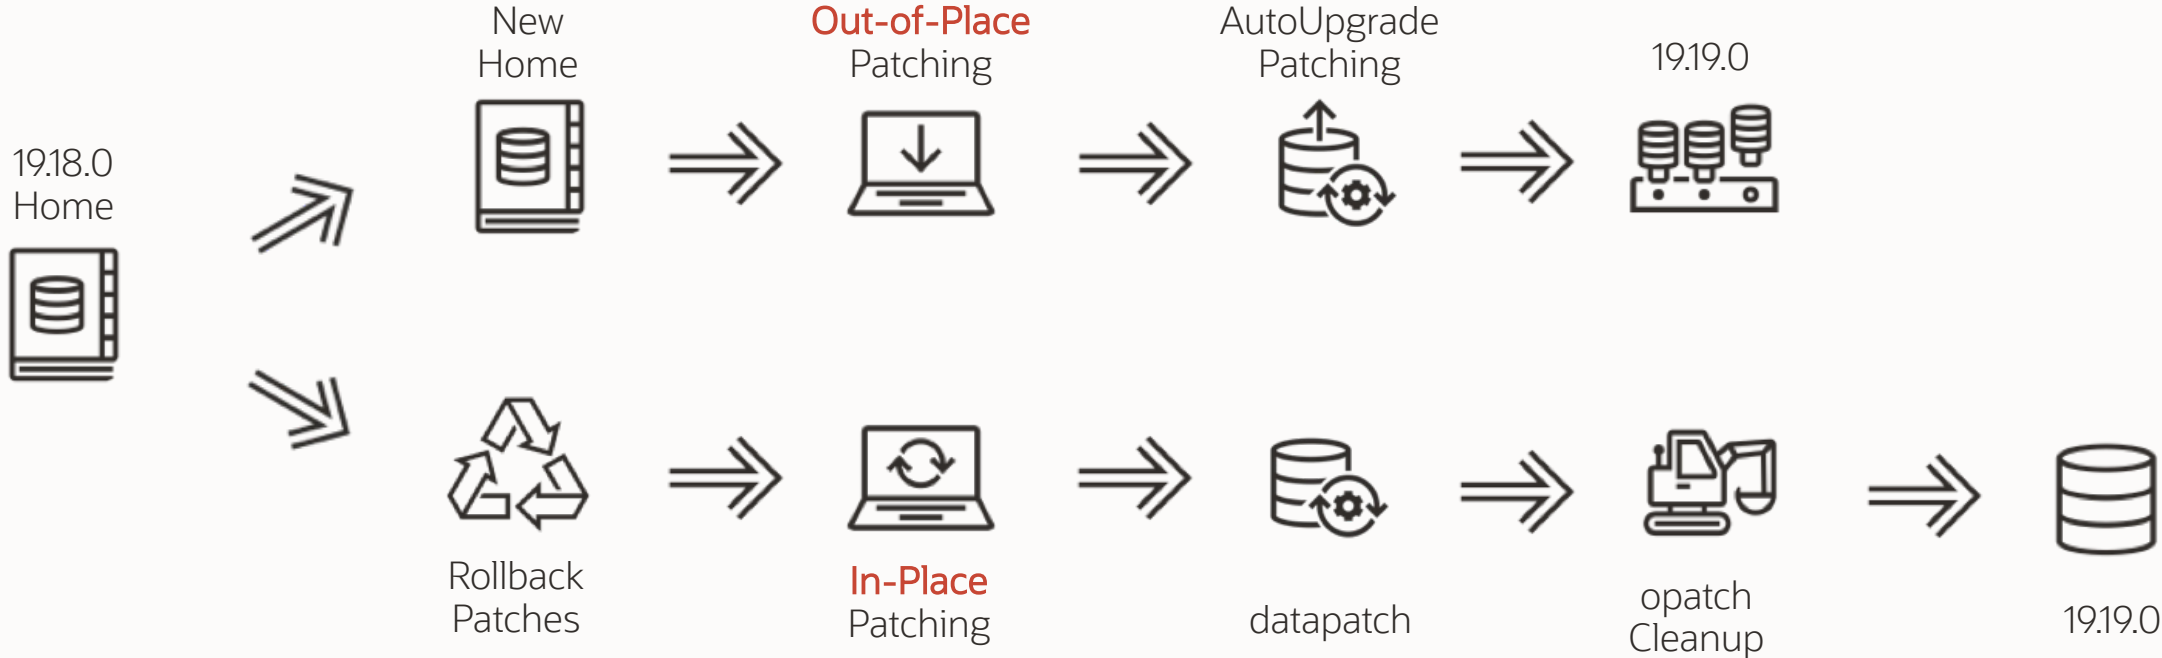
ONE SINGLE COMMAND

```
./runInstaller
  -applyRU /home/oracle/stage/RU/35042068
  -applyOneOffs ...
```

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>



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_34.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]$ $ORACLE_HOME/OPatch/datapatch -verbose
```

Out-of-Place Patching

Oracle Home, 19.19.0



```
SQL> SHUTDOWN IMMEDIATE
```



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

New or Cloned
Oracle Home, 19.20.0



Tim Hall ∞ 🧑 +∞ 🗨️
@oraclebase



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

In-Place = Current ORACLE_HOME
Out-Of-Place = New ORACLE_HOME

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





Always patch Out-of-Place


- Don't argue with us 😊



Safely test and verify patches with Standby-First Patch Apply

- [Data Guard Standby-First Patch Apply \(Doc ID 1265700.1\)](#)

Standby-First Patching



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



Patch must be labeled as Standby-First applicable

- Check the patches' README



Execute datapatch on the primary database

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




Find additional restrictions in [Data Guard Standby-First Patch Apply \(Doc ID 1265700.1\)](#)



Avoid database downtime with RAC Rolling Patch Apply

RAC Rolling Patching



```
$ $ORACLE_HOME/OPatch/opatchauto
```

- Clone 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 gridSetup



Use Out-Of-Place Patching

- Minimize downtime
- Minimize risk during outage
- Easier rollback

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

Patching GI and DB Homes together?

Option 1

TOGETHER

One maintenance window

Longer, single patching window

Several changes

Option 2

SEPARATELY

Two maintenance windows

Shorter window, but longer overall patching

One change at a time



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 19.16.0 and DB 19.20.0
- Node 1 with GI 19.16.0, node 2 with GI 19.18.0
- 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




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




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




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



DRAINING

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

Drain Timeout



Setting `drain_timeout`
very **low**?

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



Setting `drain_timeout`
very **high**?

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



Comply with Maximum Availability Architecture (MAA) principles

- [Continuous Availability - MAA Checklist for Applications for the Oracle Database](#)

top tips

PATCHING SUCCESS

Cluster Verification Utility

Patch Level

Application Continuity

OPatch

Use CVU before and after patching

Preferably through EXAchk or ORAchk

Identifies potential issues

Light-weight, non-intrusive

Always use the latest version

top tips

PATCHING SUCCESS

Cluster Verification Utility

Patch Level

Application Continuity

OPatch

Apply patches regularly

Apply recent Release Updates

Apply MRPs

Keep GI and DB patch levels in sync

top tips

PATCHING SUCCESS

Cluster Verification Utility

Patch Level

Application Continuity

OPatch

Completely hide interruptions from users

Hides planned and unplanned events

Comply with MAA guidelines

See also Transparent Application Continuity

top tips

PATCHING SUCCESS

Cluster Verification Utility

Patch Level

Application Continuity

OPatch

Always use the latest version of OPatch

Use in GI and DB homes

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
- Work in progress (fully implemented in 19.21.0)

```
$ ./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.
```

```
:
```

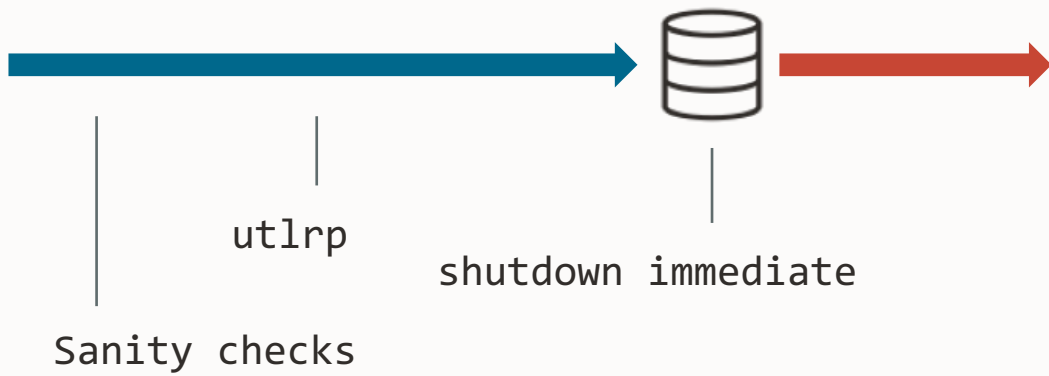
```
JOB_NAME,NEXT_RUN_DATE,SCHEMA_NAME,STATE
```

```
CLEANUP_TRANSIENT_PKG,23-MAY-23 11.08.53.000000 AM +01:00,APPUSER,SCHEDULED
```



Recompile invalid objects
before invoking `datapatch`

Patching Timeline






The database must be open
Only open PDBs are patched

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

Patching Timeline

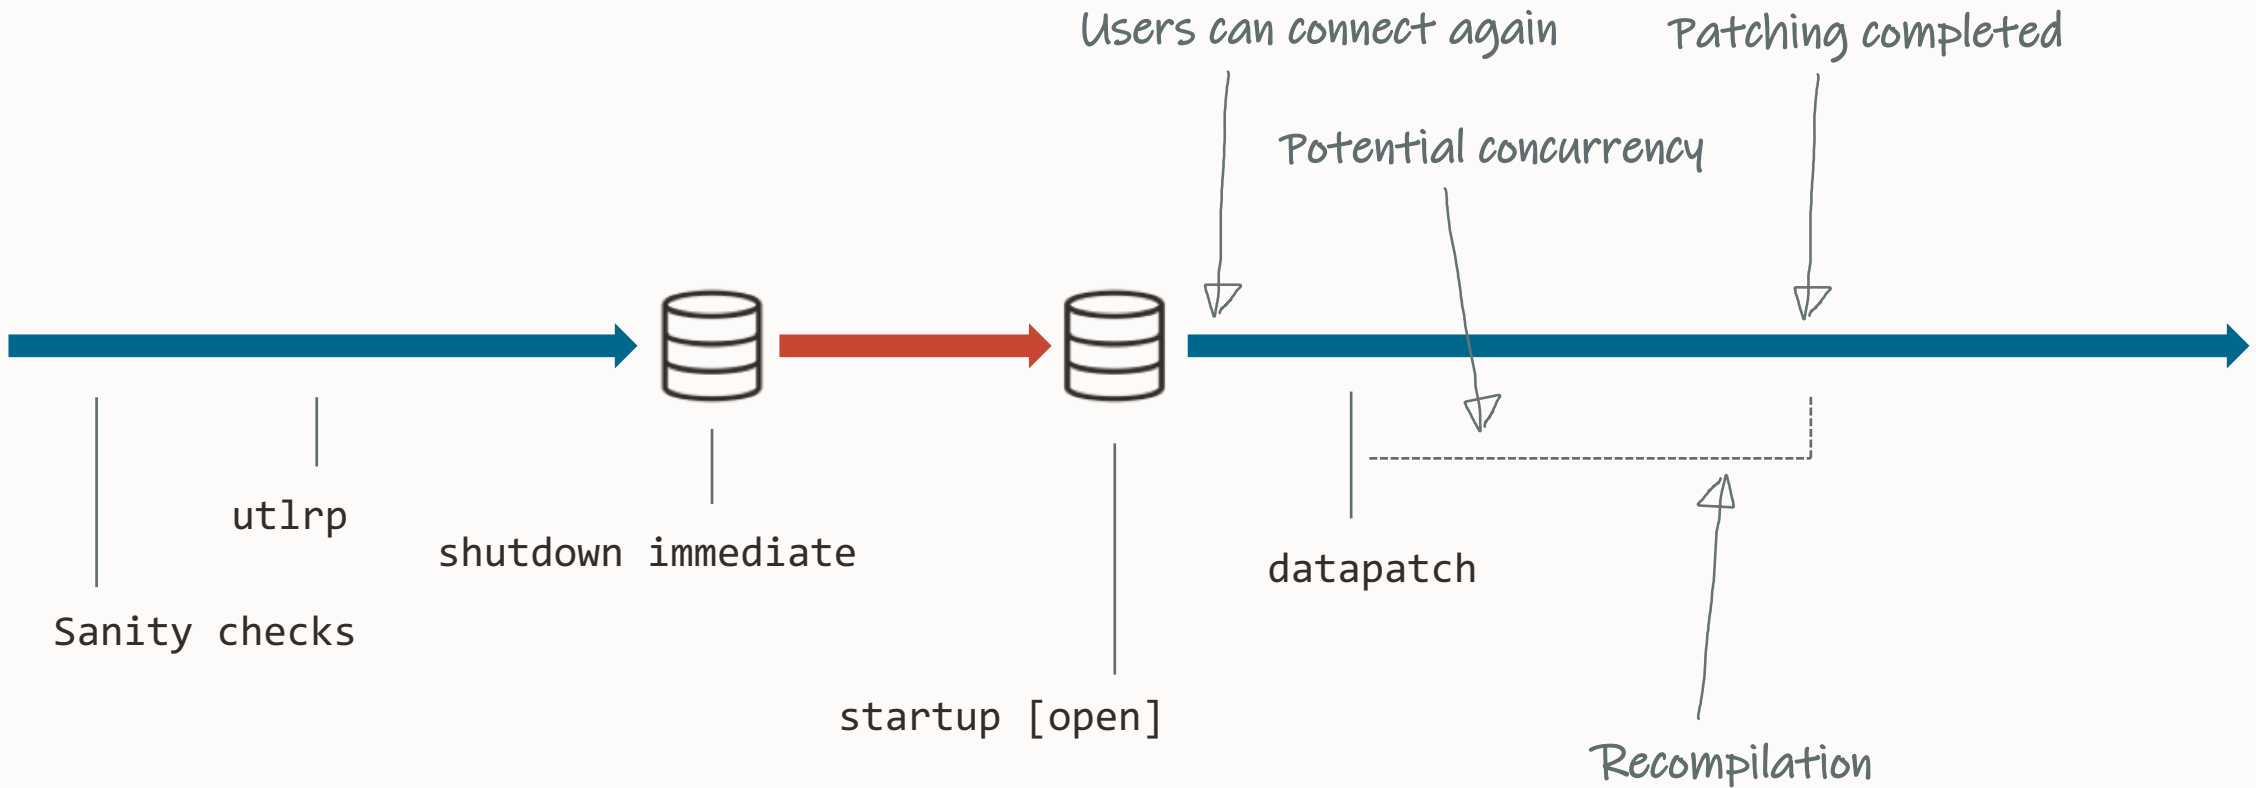





You can run **datapatch** while users are connected to the database

- Details in [blog post](#)

Patching Timeline



Patching Timeline



Concurrency

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



Disable the database scheduler (`job_queue_processes=0`)

- Also disables refresh of materialized views

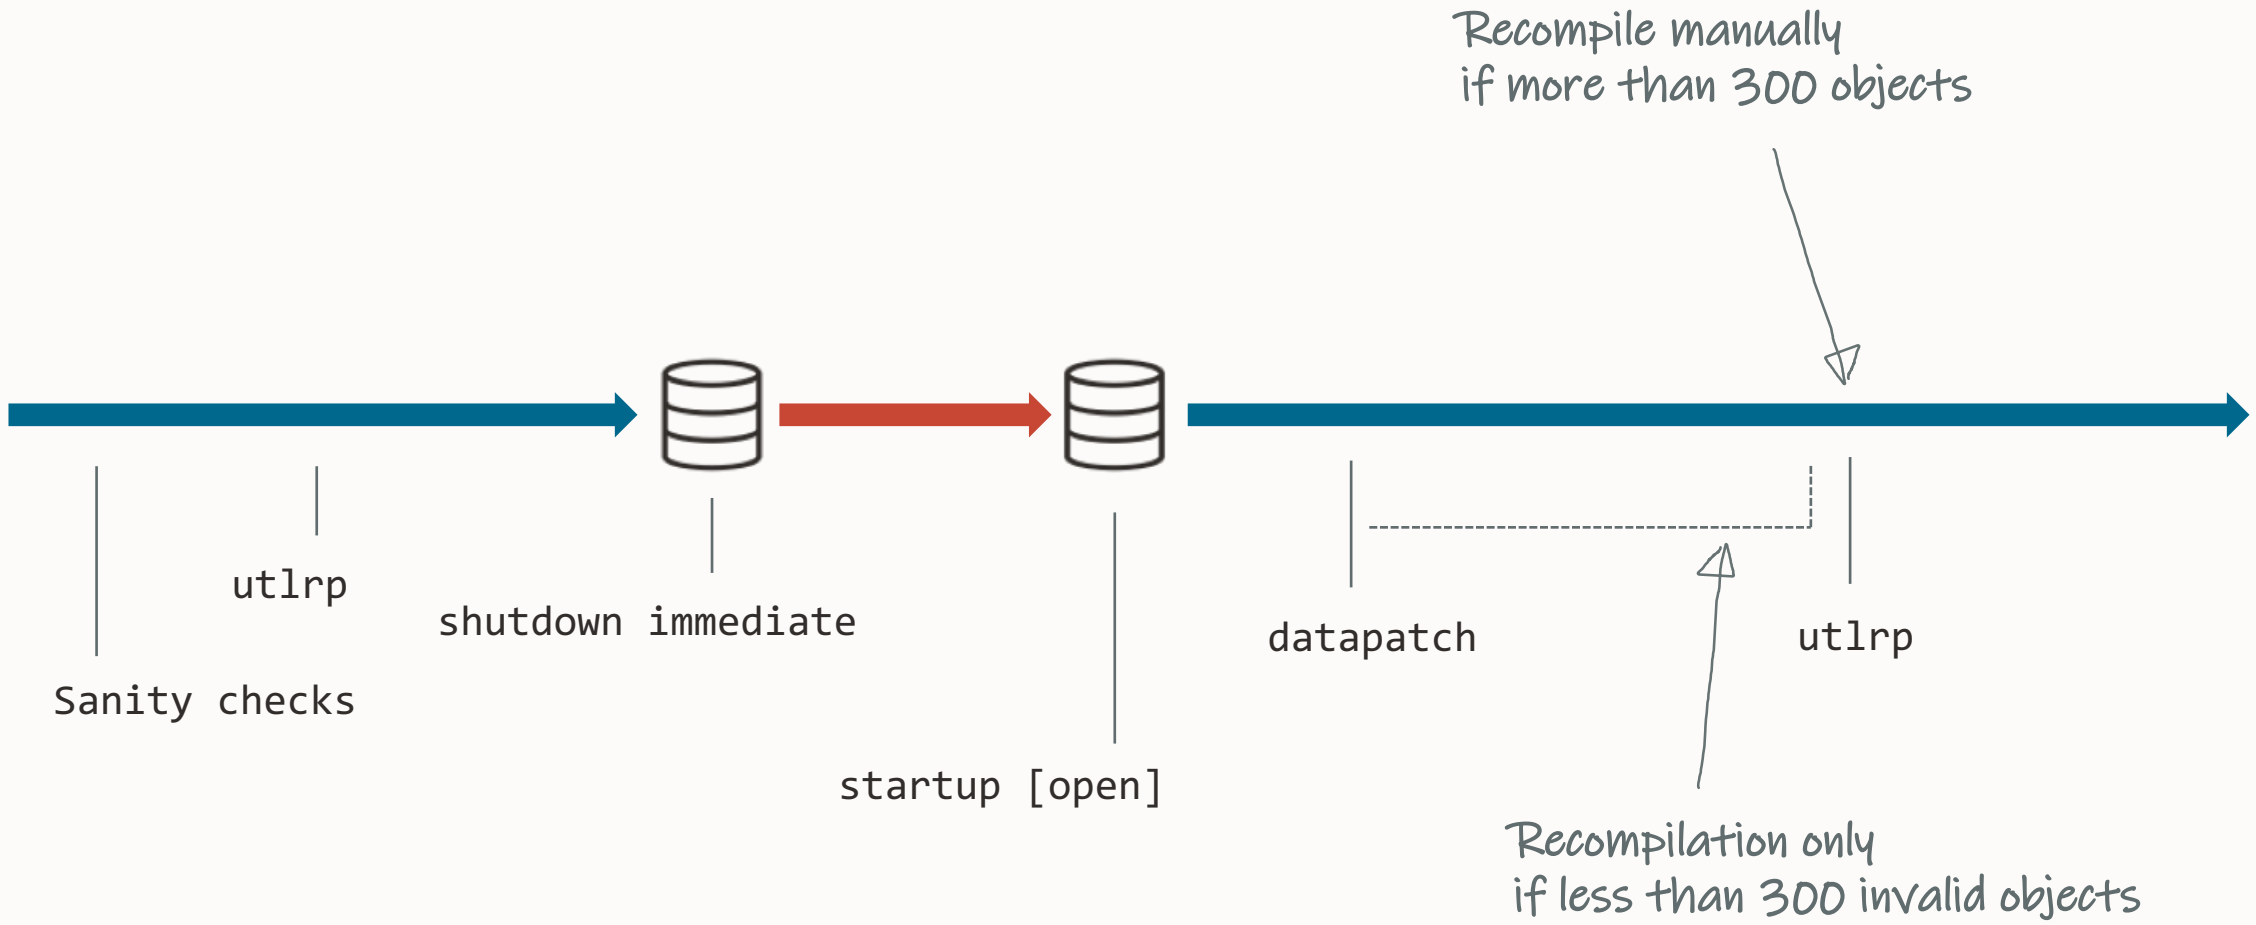


Postpone RMAN backups



Stop Oracle GoldenGate
while you are running datapatch

Patching Timeline



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** 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
```





Less installed components
lead to **faster patching**

- Typical candidates: JAVAVM, SDO

Patching Best Practices

Installation

Basics

Methods

Grid Infrastructure

Datapatch

AutoUpgrade

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

AutoUpgrade functionality extended to patching

```
$ cat DB19.cfg
```


```
patch1.source_home=/u01/app/oracle/product/19/dbhome_19_18_0
```

```
patch1.target_home=/u01/app/oracle/product/19/dbhome_19_20_0
```

```
patch1.sid=DB19
```

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

Fleet Patching



AutoUpgrade

Automate your patching process and benefit from the familiar AutoUpgrade



Fleet Patching and Provisioning

Go fleet scale with FPP and benefit from additional functionality like deployment of Oracle Home



Break

We'll start again in 15 minutes

Before upgrade

How to upgrade and convert

After upgrade



Do you want to upgrade?

Oracle Database 11.2.0.4

Oracle Database 12.1.0.2

Oracle Database 12.2.0.1

Oracle Database 18c

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

⇒⇒ Oracle Database 19c ⇒⇒ **Oracle Database 23c**



Oracle Database 23c supports the multitenant architecture only

- You must convert your database to a PDB

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

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



Ensure your clients can connect to Oracle Database 23c

- Upgrade your clients well in advance of the upgrade

Client / Server Interoperability

Client Version	Server Version						
	23c	21c	19c	18c	12.2.0	12.1.0	11.2.0
23c ^{#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
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	Was	Yes ^{#12}	Yes ^{#9}

[MOS Note: 207303.1 - Client / Server Interoperability Support Matrix](#)



JDBC Interoperability

Currently, only JDBC driver 23c can connect to Oracle Database 23c
(subject to change)

For up-to-date information:

[Starting With Oracle JDBC Drivers - Installation, Certification, and More! \(Doc ID 401934.1\)](#)



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/23c
```

```
upg1.target_cdb=CDB1
```

```
upg1.run_dictionary_health=full
```

```
#To run only the critical checks
```

```
#upg1.run_dictionary_health=critical
```

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

```
dbms_dictionary_check on 05-SEP-2023 09:41:30
-----
Catalog Version 19.0.0.0.0 (1900000000)
db_name: DB19
Is CDB?: NO
Trace File: /opt/oracle/diag/rdbms/db19/DB19/trace/DB19_ora_25104_DICTCHECK.trc

Procedure Name          Catalog      Fixed
                        Version      Vs Release   Timestamp    Result
-----
.- OIDOnObjCol          ... 1900000000 <= *All Rel* 09/05 09:41:30
PASS
.- ObjNotInObj          ... 1900000000 <= *All Rel* 09/05 09:41:30
PASS
.- SourceNotInObj      ... 1900000000 <= *All Rel* 09/05 09:41:30
PASS
.- OversizedFiles      ... 1900000000 <= *All Rel* 09/05 09:41:30
PASS
.- PoorDefaultStorage  ... 1900000000 <= *All Rel* 09/05 09:41:30
PASS
.- PoorStorage         ... 1900000000 <= *All Rel* 09/05 09:41:30
PASS

*** 2023-09-05T09:41:30.934258+00:00
PASS
.- TabPartCountMismatch ... 1900000000 <= *All Rel* 09/05 09:41:30
PASS
.- TabComPartObj       ... 1900000000 <= *All Rel* 09/05 09:41:30
PASS
.- Mview               ... 1900000000 <= *All Rel* 09/05 09:41:30
PASS
.- ValidDir            ... 1900000000 <= *All Rel* 09/05 09:41:30
PASS
.- DuplicateDataobj    ... 1900000000 <= *All Rel* 09/05 09:41:30
PASS
.- ObjSyn              ... 1900000000 <= *All Rel* 09/05 09:41:31
PASS
.- ObjSeq              ... 1900000000 <= *All Rel* 09/05 09:41:31
PASS
```



Ensure dictionary and fixed objects statistics are accurate

- Save downtime by gathering in advance

```
begin
```

```
    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 23c



Container Database



AutoUpgrade



Installation of Oracle Home is simpler

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



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 23c 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 23c



AutoUpgrade

Create Container Database



1 Character set

2 Components

3 COMPATIBLE

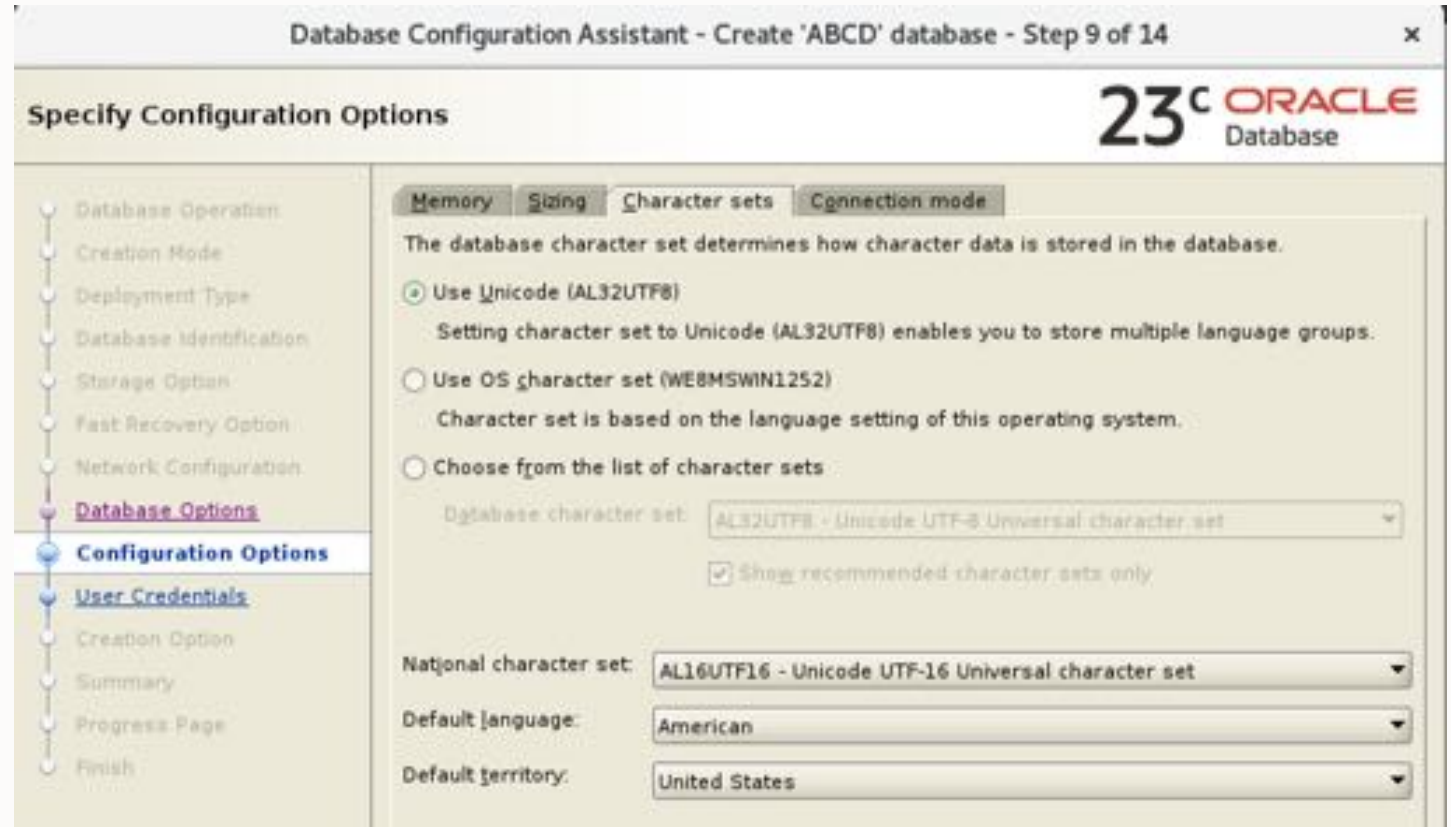
Create Container Database

1 Character set

- Always choose AL32UTF8
- Allows PDBs with any character set

2 Components

3 COMPATIBLE



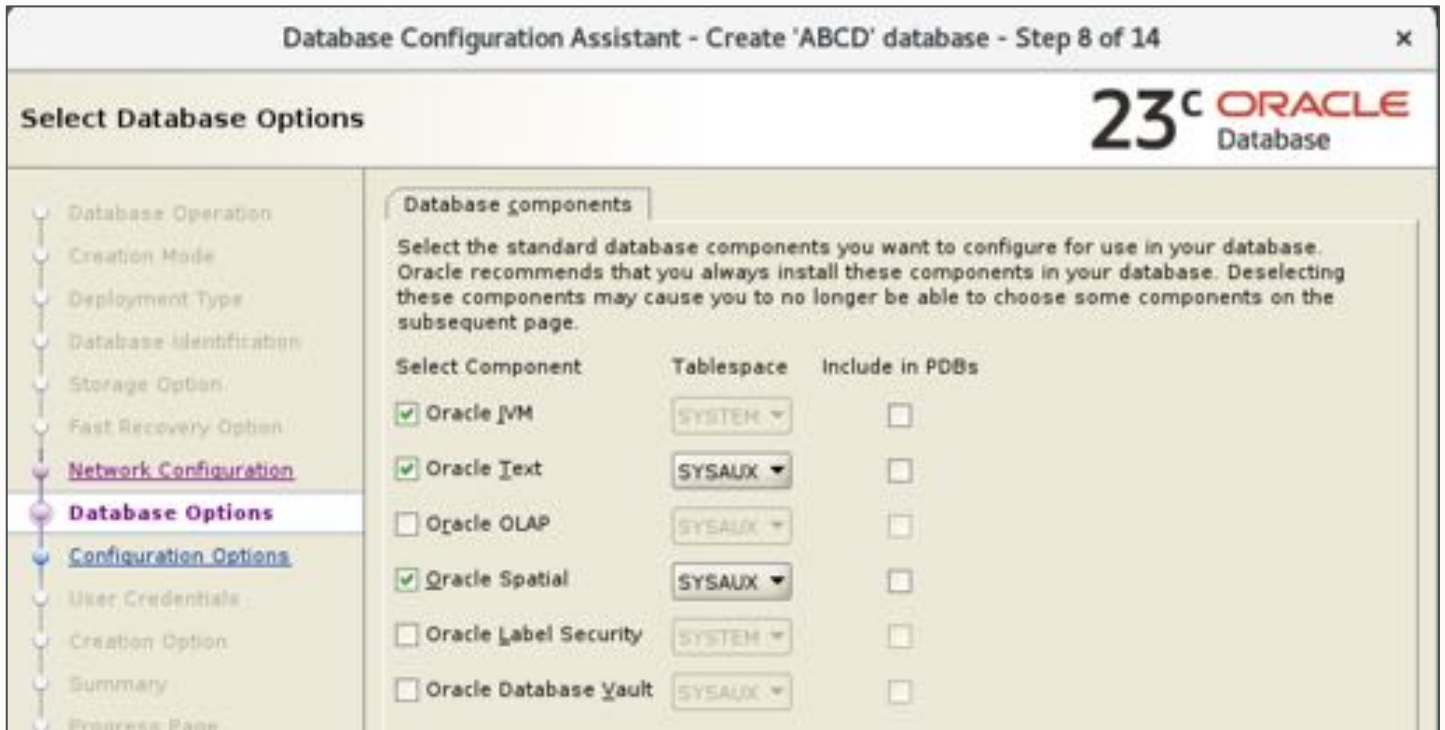
Create Container Database

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

Name ▲	Value	Include in spfile	Category
cluster_database	FALSE	<input type="checkbox"/>	Cluster Database
compatible	19.0.0	<input checked="" type="checkbox"/>	Miscellaneous
control_files	(*/u02/oradata/{DB_UNI...	<input checked="" type="checkbox"/>	File Configuration
db_block_size (bytes)	8192	<input checked="" type="checkbox"/>	Cache and I/O
db_create_file_dest	/u02/oradata/{DB UNIQUE ...	<input type="checkbox"/>	File Configuration

3 COMPATIBLE

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



--Allows CDB views to include information on PDB\$SEED objects.
--By default, such information is hidden.
--https://mikedietchde.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 24854571

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

```
build.version 23.3.230829
```

```
build.date 2023/08/29 19:47:13 -0700
```

```
build.hash 8ab1875
```

```
build.hash_date 2023/08/29 03:54:58 -0700
```

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

```
build.type test
```

```
build.label (HEAD, origin/devel)
```



Flow

1

Plug in

2

Upgrade

3

Convert



23°C

*Irreversible!
Flashback no good*

Demo

Upgrade to Oracle Database 23c

- Using AutoUpgrade
- Including PDB conversion

[Watch on YouTube](#)

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

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';
```

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
```

```
--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
```

```
--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=22/10/2023 02:00:00
--Specify relative start time
--upg1.start_time=+1h30m
```



Refreshable Clone



`autoupgrade.jar ... -mode deploy`

`upg1.start_time=22/10/2023 02:00:00`





The source non-CDB stays intact
to allow rollback



Works for unplug-plug upgrades as well



Zürcher
Kantonalbank

Customer Case | Zürcher Kantonalbank

Customer

Project

Constraints

Preparation

Migration

Success?

Remarks

A reliable partner for over 150 years

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



Customer Case | Zürcher Kantonalbank

Customer

Project

Constraints

Preparation

Migration

Success?

Remarks

Current situation

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

Planned solution: AutoUpgrade

Customer Case | Zürcher Kantonalbank

Customer

Project

Constraints

Preparation

Migration

Success?

Remarks

Test setup

- 3 non-CDB databases of different size

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

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

Customer Case | Zürcher Kantonalbank

Customer

Cloning user

Project

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

Constraints

Permissions

Preparation

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

Migration

Success?

Database link

Remarks

```
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

```
(Job#|DB_NAME| STAGE(OPERATION)| STATUS|START_TIME| UPDATED| MESSAGE|
-----|-----|-----|-----|-----|-----|-----|
| 106| | CLOSING|EXECUTING| RUNNING| 09:50:39|!!!17381s ago|Creating pluggable database|
| 107| | COMPLETED| STOPPED|FINISHED| 09:50:40| | |
| 108| | COMPLETED| STOPPED|FINISHED| 09:50:40| | |
-----|-----|-----|-----|-----|-----|
Total jobs 3
The command is running every 60 seconds. PRESS ENTER TO EXIT
-----|-----|-----|-----|-----|-----|
(Job#|DB_NAME| STAGE(OPERATION)| STATUS|START_TIME| UPDATED| MESSAGE|
-----|-----|-----|-----|-----|-----|
| 106| | CLOSING|EXECUTING| RUNNING| 09:50:39|!!!17441s ago|Creating pluggable database|
| 107| | COMPLETED| STOPPED|FINISHED| 09:50:40| | |
| 108| | COMPLETED| STOPPED|FINISHED| 09:50:40| | |
-----|-----|-----|-----|-----|-----|
Total jobs 3
The command is running every 60 seconds. PRESS ENTER TO EXIT
-----|-----|-----|-----|-----|-----|
(Job#|DB_NAME| STAGE(OPERATION)| STATUS|START_TIME| UPDATED| MESSAGE|
-----|-----|-----|-----|-----|-----|
| 106| | CLOSING|EXECUTING| RUNNING| 09:50:39|!!!17501s ago|Creating pluggable database|
| 107| | COMPLETED| STOPPED|FINISHED| 09:50:40| | |
| 108| | COMPLETED| STOPPED|FINISHED| 09:50:40| | |
-----|-----|-----|-----|-----|-----|
Total jobs 3
The command is running every 60 seconds. PRESS ENTER TO EXIT
```

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

```
The command let is running every 40 seconds, PRESS ENTER TO EXIT
-----
|Job#|DB_NAME| STAGE|OPERATION| STATUS| START_TIME|LOCATION| MESSAGE|
-----
| 106| |CLOSEPJOB| STOPPED| ERROR|Oct-03 09:50| |10PO-6016|
| 107| |COMPLETED| STOPPED|FINISHED|Oct-03 09:50| | |
| 108| |COMPLETED| STOPPED|FINISHED|Oct-03 09:50| | |
-----
Total jobs: 3

The command let is running every 40 seconds, PRESS ENTER TO EXIT
-----
|Job#|DB_NAME| STAGE|OPERATION| STATUS| START_TIME|LOCATION| MESSAGE|
-----
| 106| |CLOSEPJOB| STOPPED| ERROR|Oct-03 09:50| |10PO-6016|
| 107| |COMPLETED| STOPPED|FINISHED|Oct-03 09:50| | |
| 108| |COMPLETED| STOPPED|FINISHED|Oct-03 09:50| | |
-----
Total jobs: 3

The command let is running every 40 seconds, PRESS ENTER TO EXIT

app>
app> exit
There is 1 job in progress, if you exit it will stop.
Are you sure you wish to leave? [y|n] y
----- Final Summary -----
Number of databases: [ 3 ]
Time finished: [2]
Jobs failed: [1]
```

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



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 23c CDB



Backup your database after migration

- Level 0
- Practice restore with pre-plugin backups

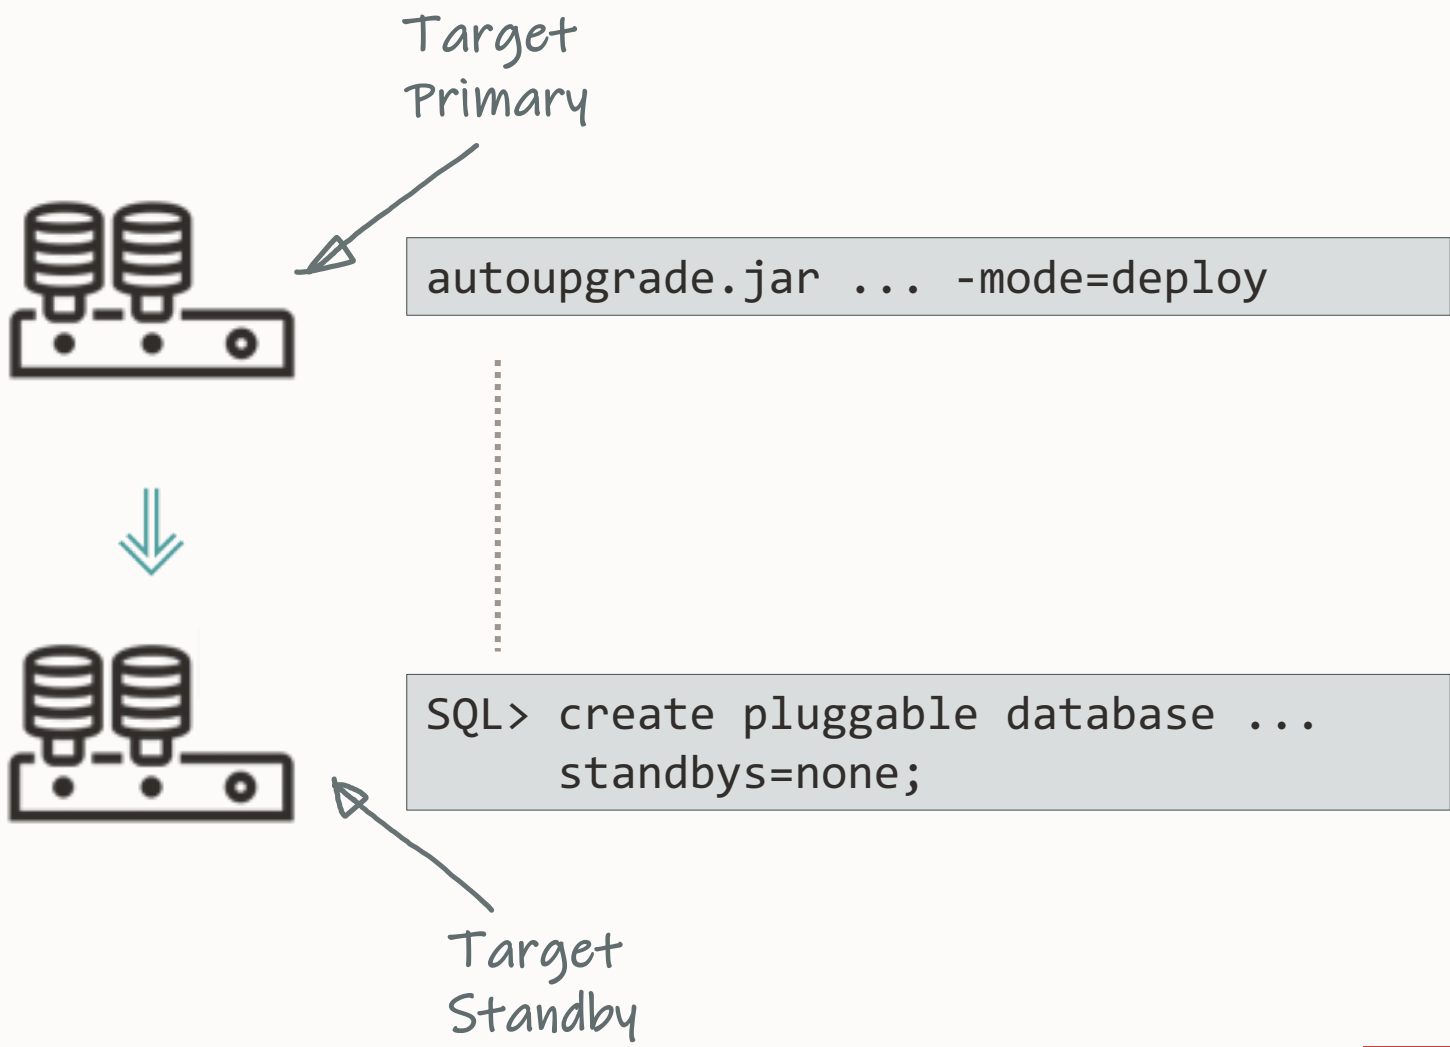




Check your standby databases

- Special attention is needed for standby databases

- Default config file setting
- PDB is not fully created on standbys
- Recovery is needed

`upg1.manage_standbys_clause=standbys=none`






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



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



*PDB created
Data files missing*





```
SQL> show pdbs
```

CON_NAME	OPEN	MODE
PDB1	READ	WRITE

```
SQL> show pdbs
```

CON_NAME	OPEN	MODE
PDB1	MOUNTED	





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

NAME	RECOVERY_STATUS
PDB1	DISABLED





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

```
SQL> alter pluggable database  
enable recovery;
```

```
SQL> alter database datafile  
... online;
```


Data Guard | Multiple Standbys

Avoid overloading the primary database when restoring data files

On standby 1

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

On standby 2

```
RMAN> restore pluggable database ...  
      from service <standby 1> ;
```

On standby 3

```
RMAN> restore pluggable database ...  
      from service <standby 1> ;
```

On standby 4

```
RMAN> restore pluggable database ...  
      from service <standby 2> ;
```

Data Guard | STANDBYS=NONE

[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.

- To change the default behavior
- Take care - don't break your standby database
- Works only for databases without ASM or OMF

`upg1.manage_standbys_clause=standbys=all`

Data Guard | Re-use Data Files

12.2.0.1
Non-CDB
Primary



12.2.0.1
Non-CDB
Standby



```
SQL> select name from v$datafile where con_id=3;
```

```
NAME
```

```
-----  
+DATA/DB_BOSTON/DD934E8207292138E053E801000A8351/DATAFILE/system.269.1103046537  
+DATA/DB_BOSTON/DD934E8207292138E053E801000A8351/DATAFILE/sysaux.270.1103046537  
+DATA/DB_BOSTON/DD934E8207292138E053E801000A8351/DATAFILE/users.273.1103046827
```

```
SQL> select name from v$datafile where con_id=3;
```

```
NAME
```

```
-----  
+DATA/DB_FRA27D/DD934E8207292138E053E801000A8351/DATAFILE/system.265.1103050007  
+DATA/DB_FRA27D/DD934E8207292138E053E801000A8351/DATAFILE/sysaux.266.1103050007  
+DATA/DB_FRA27D/DD934E8207292138E053E801000A8351/DATAFILE/users.269.1103050009
```

Data Guard | Re-use Data Files

12.2.0.1
Non-CDB
Primary



12.2.0.1
Non-CDB
Standby



The manifest file contains

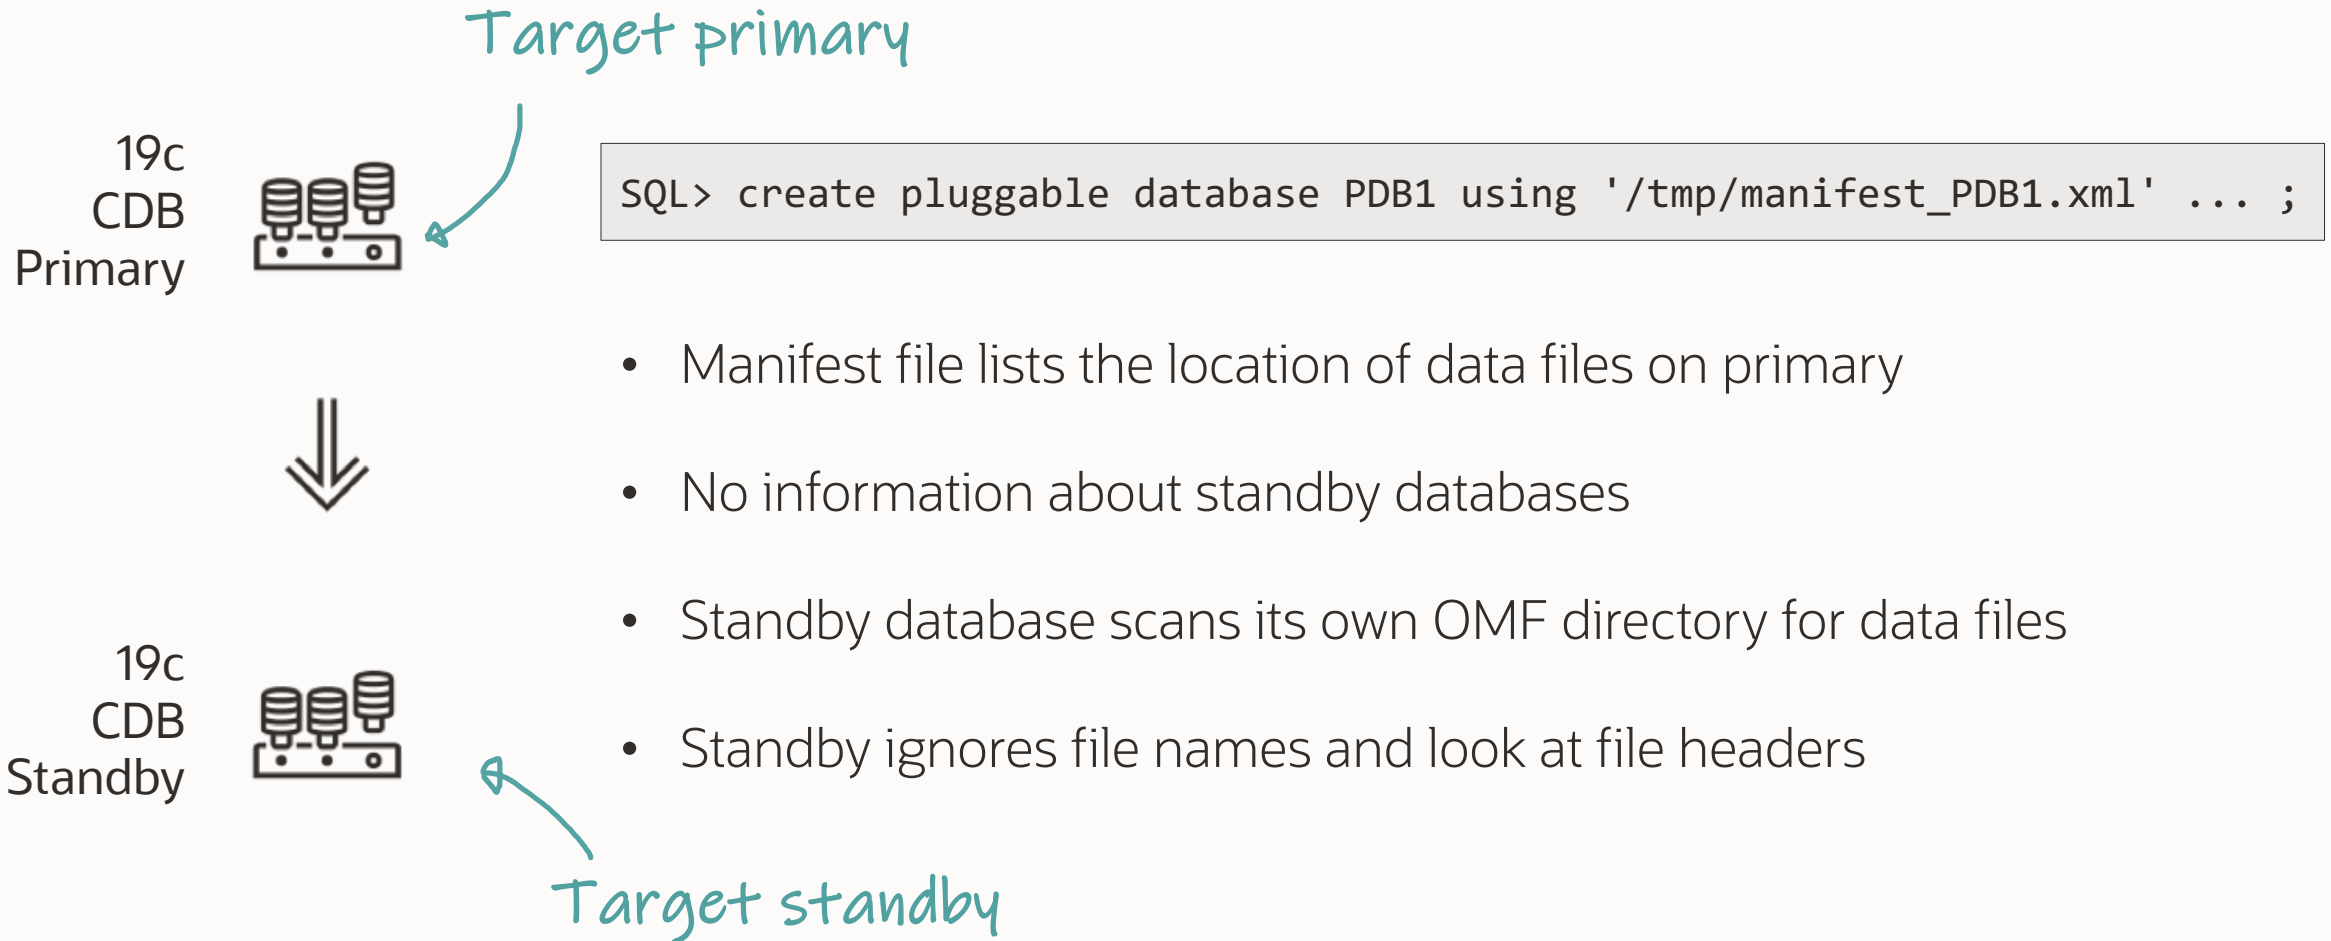
```
SQL> alter pluggable database PDB1 plug into '/tmp/manifest_PDB1.xml';
```

- Not standby database

```
<?xml version="1.0" encoding="UTF-8"?>
<PDB>
  <xmlversion>1</xmlversion>
  <pdbname>PDB1</pdbname>
  ...
  <guid>DDB49CFEFD8ED4FCE053E801000A078C</guid>
  ...
  <tablespace>
    <name>USERS</name>
    ...
    <file>
      <path>+DATA/DB_BOSTON/DD934E8207292138E053E801000A8351/DATAFILE/users.273.1103046
      827</path>
```



Data Guard | Re-use Data Files





I'll just copy the file in ASM

```
ASMCMD> cp users.269.1103050009  
+DATA/DB_FRA27D/.../users.273.1103046827
```


```
ASMCMD-8016: copy source '+DATA/DB_FRA27D/.../users.269.1103050009' and target  
'+DATA/DB_FRA27D/.../users.273.1103046827' failed
```

```
ORA-15056: additional error message
```

```
ORA-15046: ASM file name 'users.273.1103046827' is not in single-file creation form
```

```
ORA-06512: at "SYS.X$DBMS_DISKGROUP", line 617
```

```
ORA-06512: at line 3 (DBD ERROR: OCIStmtExecute)
```

Only a database can produce files
with ASM/OMF data file names

Data Guard | Re-use Data Files

19c
CDB
Primary



19c
CDB
Standby



ASM [alias](#) to the rescue!

- On standby, create aliases for the primary data files

```
ASMCMD> alter diskgroup data add alias '...' for '...' ;
```

- Plug in PDB, standby will find aliases and find the real file locations
From alert log

```
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_FRA27D/.../DATAFILE/users.269.1103050009'
```



Move data files into proper OMF location
using online data file move

- Remove aliases when no longer needed

Data Guard | STANDBYS=ALL

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

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

In this Document

- [Goal](#)
- [Solution](#)
- [Prerequisites](#)
- [Steps](#)
- [Resolving Errors](#)
- [References](#)

APPLIES TO:

Oracle Database Cloud Service - Version N/A and later
Oracle Database Exadata Express 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
Gen 1 Exadata Cloud at Customer (Oracle Exadata Database Cloud Machine) - Version N/A and later
Information in this document applies to any platform.

GOAL

To plug in an existing 12.1.0.2 or later PDB residing in a CDB as part of a Data Guard configuration into another CDB that is part of a different Data Guard configuration where the current Primary CDB and the target CDB both have standby databases and allow you to use the original Standby database's data files to update the destination CDB's Standby.

This note describes a multitenant migration option for maintaining standby databases when the source database is a PDB. If your source database is a non-CDB, please see [Document 2273204.1](#).

For Oracle RDBMS 19.15 and later, the Data Guard broker MIGRATE command has been enhanced to execute the steps contained in this document. It will manage configurations of the destination CDB containing a single physical standby database and will handle TDE enabled databases. Please see [High Availability Overview and Best Practices - PDB Switchover and Failover in a Multitenant Configuration](#) for more information on this feature.

Always test the steps in a dev/test environment prior to using in production. Since the original files are being modified directly by the plugin on the primary and by the consumption

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.



Confused? We explain all the details
in our [AutoUpgrade 2.0](#) webinar

```
--You can combine the two approaches.  
--Specify the standby databases where you have ASM aliases,  
--here std2, std4  
--The rest are treated as STANDBYS=NONE
```

```
create pluggable database ... standbys=std2,std4;
```



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;
```

```
--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,  
--Auto-Indexing and In-Memory.  
--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$  
--Value of "0" fully disables tracking.
```

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



We need real-world experience with 23c

- 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

The Data Pump LOB Mystery

—
And how to solve it

A short history of **binary data types**

v4

LONG and LONG RAW

8.0

CLOB and BLOB

11g

SecureFile LOBs

v4

LONG and LONG RAW

8.0

BasicFile LOBs

11g

SecureFile LOBs

v4

LONG and LONG RAW

- Only 1 column per table
- Max size: 2GB - 1

8.0

BasicFile LOBs

- Performance constraints
- No Parallel DML allowed
- Max size: (4GB - 1) * DB_BLOCK_SIZE

11g

SecureFile LOBs

- Improved performance
- Data Pump can use multiple workers or Parallel Query
- Deduplication, encryption and more
- Max size: same as with CLOB/BLOB



As of today, all legacy binary data types should have been migrated to **SecureFile LOBs**

```
--Always convert LOBs to SecureFile on import  
--Imports into SecureFile is always the fastest approach  
--For further info, see Doc ID 490228.1
```

```
impdp ... transform=lob_storage:securefile
```

Different LOB types

Internal LOBs stored **inside** the database

- CLOB
- NCLOB
- BLOB

External LOBs stored **outside** the database

- BFILE

Initialization Parameter

DB_SECUREFILE

- NEVER
- PERMITTED
- **PREFERRED** ← LOBs are created as SecureFile LOBs unless explicitly stated
- ALWAYS
- IGNORE


Tablespace must use Automatic Segment Space Management (ASSM)

Data Pump & LOBs

Things to know and consider



No parallelism with BasicFile LOBs

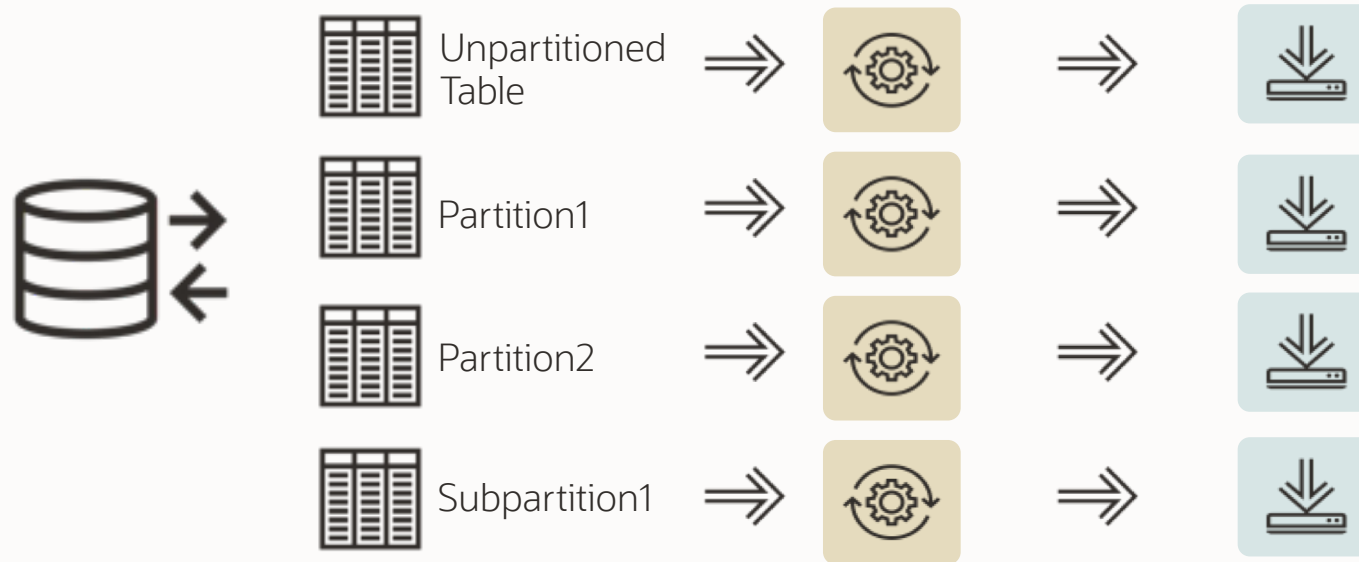


Always use **SecureFile LOBs**

*But why is there only **one worker**?*

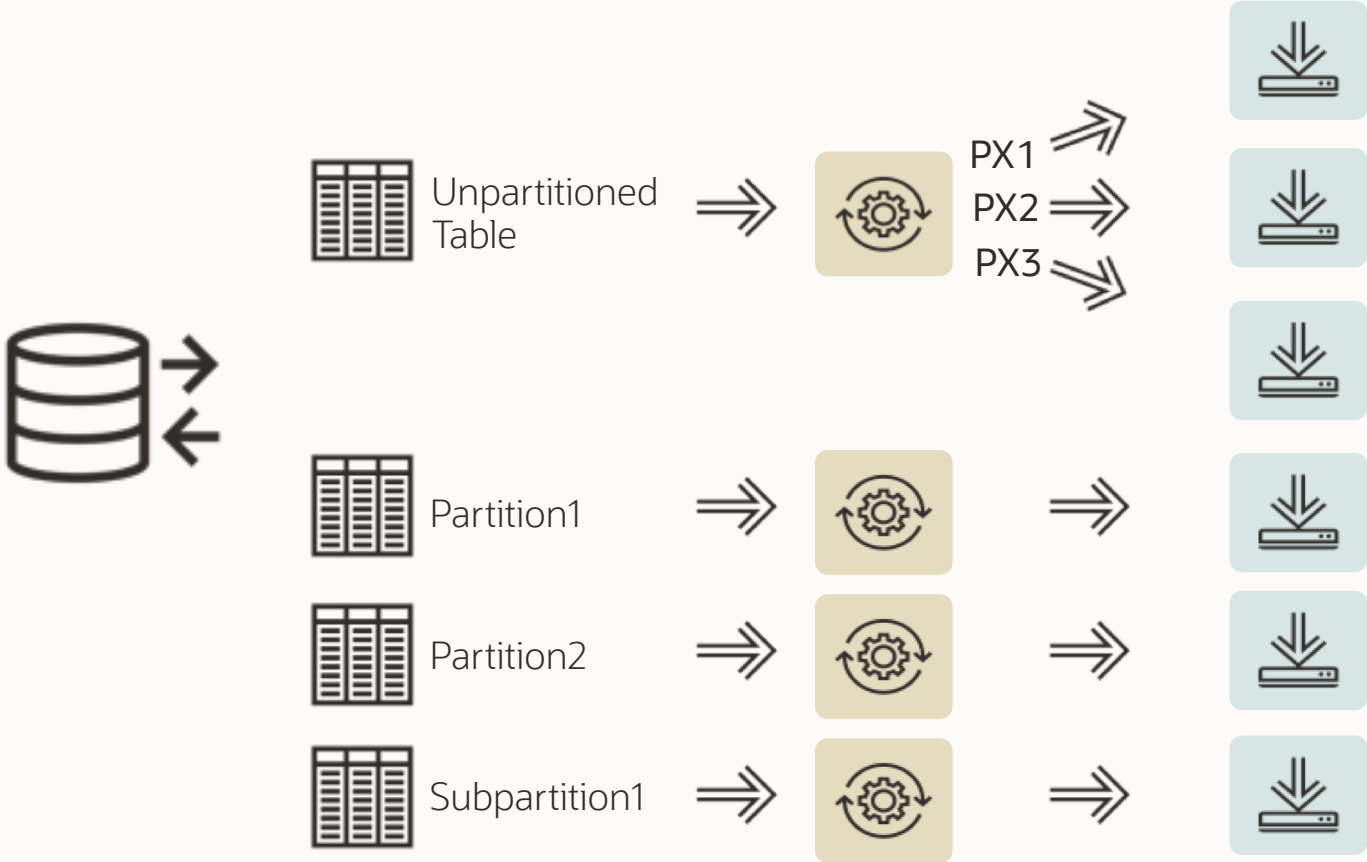
Data Pump | Parallel Worker Activity

Data Pump employs one worker per table data object



Data Pump | Worker as PQ Coordinator

If a table data object is >250MB, Data Pump can invoke parallel query



LOB Export | Example Table



```
CREATE OR REPLACE DIRECTORY BLOB_DIR AS '/tmp/mydir';
```



```
CREATE TABLE tab1 ( id NUMBER, blob_data BLOB )  
  LOB (blob_data) store as securefile;
```



```
BEGIN ... DBMS_LOB.LOADBLOBFROMFILE ...
```



```
exec DBMS_STATS.GATHER_TABLE_STATS ('HUGO', 'TAB1');
```

For a complete example, please visit oracle-base.com



LOB data is stored out-of-row
in a separate LOB segment

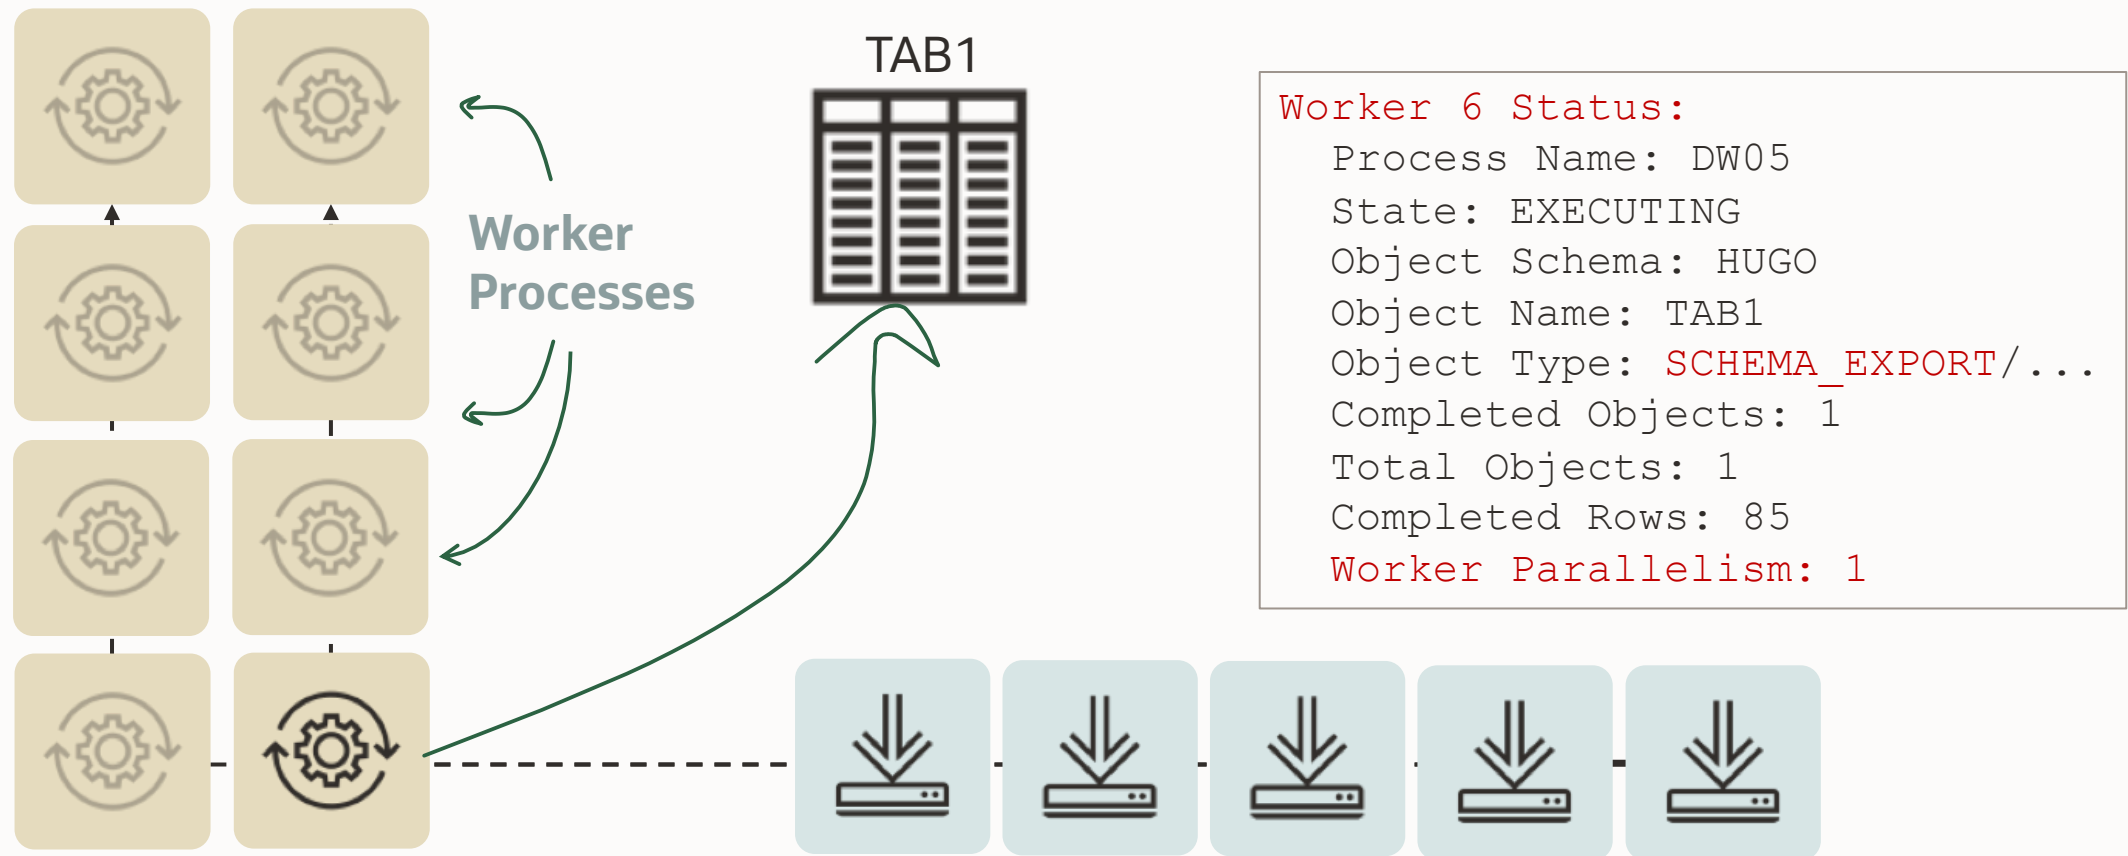
- Smaller LOBs less than 4000 byte store in-row
- Oracle Database 23c holds 8000 bytes in-row

Starting Data Pump – Test:

```
DIRECTORY=DATA_PUMP_DIR  
DUMPFILE=MYDUMP%L.DMP  
LOGFILE=MYDUMP01.LOG  
SCHEMAS=HUGO  
LOGTIME=ALL  
METRICS=YES  
PARALLEL=8
```

LOB Export | Lazy Workers?

8 workers, 5 dump files – and only 1 worker exports TAB1





Maybe the table's PARALLEL DEGREE
is too low?

LOB Export | Parallel Degree



```
select degree  
from DBA_TABLES  
where table_name='TAB1';
```

DEGREE

1

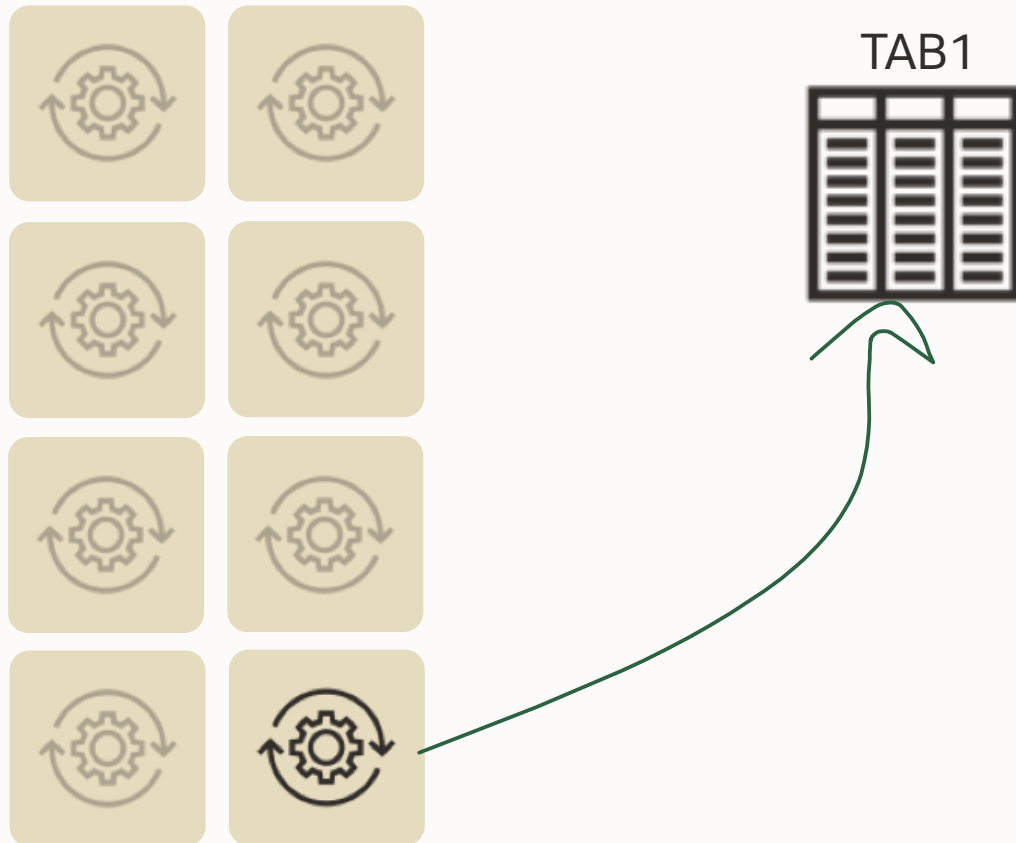
```
select degree  
from DBA_TABLES  
where table_name='TAB1';
```

DEGREE

8

LOB Export | Parallel Degree

8 workers, 5 dump files – and only 1 worker exports TAB1



Worker 1 Status:

```
Process Name: DW08
State: EXECUTING
Object Schema: HUGO
Object Name: TAB1
Object Type: SCHEMA_EXPORT/...
Completed Objects: 1
Total Objects: 1
Completed Rows: 85
Worker Parallelism: 1
```

LOB Export | Table Segments and Extents

Segments



Extents

```
select BYTES, BLOCKS, EXTENTS
from   DBA_SEGMENTS
where  SEGMENT_NAME = 'TAB1'
and    OWNER = 'HUGO';
```

BYTES	BLOCKS	EXTENTS
131072	16	2

```
select ROUND(SUM(BYTES)/1024/1024/1024,2) "GB"
from   DBA_EXTENTS
where  SEGMENT_NAME IN
      (select SEGMENT_NAME
       from   DBA_LOBS
       where  TABLE_NAME = 'TAB1'
       and    OWNER = 'HUGO');
```

10.31

LOB Export | Table Statistics

Table



Columns

```
select NUM_ROWS, BLOCKS, AVG_ROW_LEN  
from DBA_TAB_STATISTICS  
where TABLE_NAME = 'TAB1';
```

```
select COLUMN_NAME, NUM_DISTINCT,  
SAMPLE_SIZE, AVG_COL_LEN  
from DBA_TAB_COL_STATISTICS  
where TABLE_NAME='TAB1';
```

NUM_ROWS	BLOCKS	AVG_ROW_LEN
85	13	720

COLUMN_N	NUM_DIST	SAMPLE_SIZE	AVG_COL_LEN
ID	1	85	3
BLOB_DATA	0	85	717





It looks like as if Data Pump does not know anything about the dimensions of the LOB segment

LOB Export | User Objects



```
select OBJECT_NAME, OBJECT_TYPE from DBA_OBJECTS  
where OWNER = 'HUGO';
```

OBJECT_NAME	OBJECT_TYPE
TAB1	TABLE
SYS_IL0000070285C00002\$\$	INDEX
SYS_LOB0000070285C00002\$\$	LOB



Is it possible to *analyze* the LOB segment?



This is **not possible**.
So what's next?


LOB Export | Manipulating Statistics



```
begin
  DBMS_STATS.SET_TABLE_STATS (
    ownname => 'HUGO',
    tabname => 'TAB1',
    numrows => 10000000,
    numblks => 1000000);
end;
/
```

LOB Export | Parallel Degree

Relief! Workers do PQ now!



Worker 2 Status:

```
Process Name: DW01
State: EXECUTING
Object Schema: HUGO
Object Name: TAB1
Object Type: SCHEMA_EXPORT/...
Completed Objects: 1
Total Objects: 1
Completed Rows: 85
Worker Parallelism: 7
```



How about another approach ...

LOB Export | ESTIMATE=BLOCKS

```
expdp hugo/oracle ESTIMATE=BLOCKS...
```

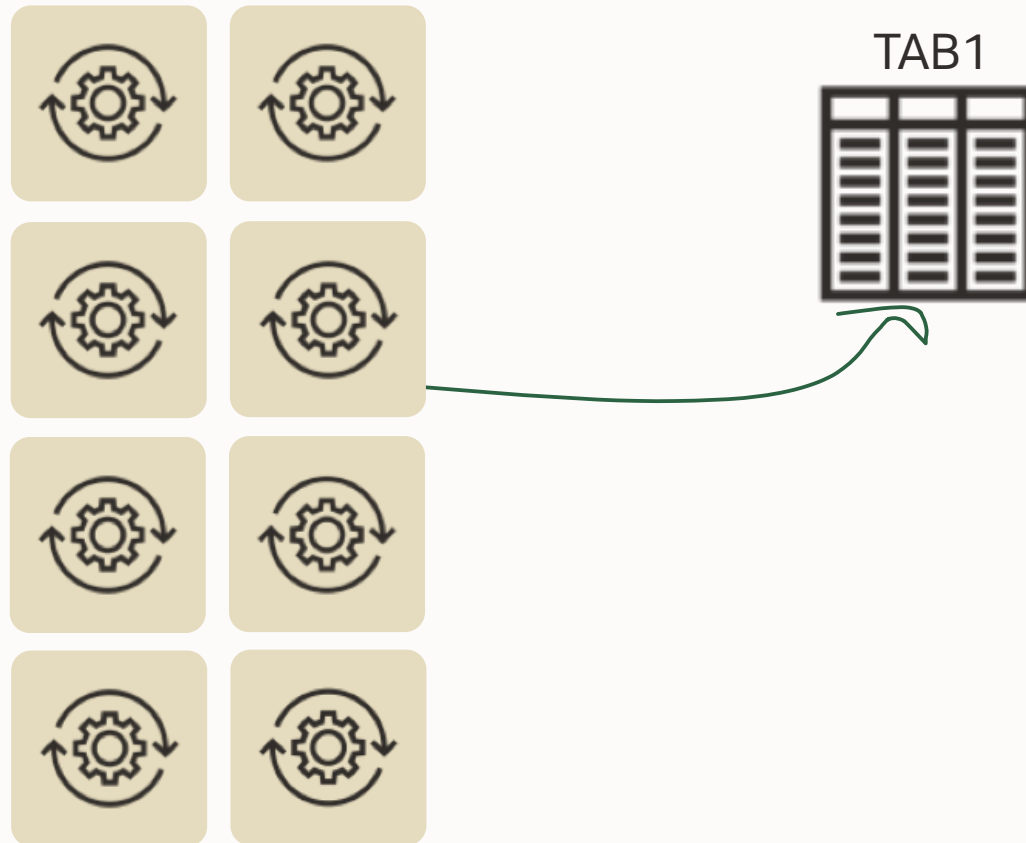


```
12-SEP-23 15:50:30.288: W-7 Startup took 0 seconds
12-SEP-23 15:50:31.409: W-1 Processing object type SCHEMA_EXPORT/TABLE/TABLE_DATA
12-SEP-23 15:50:31.735: W-1 . estimated "HUGO"."TAB1"          10.24 GB
12-SEP-23 15:50:31.735: W-1 . estimated "HUGO"."T1"           11 MB
12-SEP-23 15:50:31.735: W-1 . estimated "HUGO"."T4"           7 MB
12-SEP-23 15:50:31.735: W-1 . estimated "HUGO"."T2"           4 MB
12-SEP-23 15:50:31.735: W-1 . estimated "HUGO"."T5"          256 KB
12-SEP-23 15:50:31.735: W-1 . estimated "HUGO"."T3"           64 KB
.
.
.
```

Tech Tip: Make sure you are on 19.18 or later with the Data Pump Bundle Patch installed!

LOB Export | ESTIMATE=BLOCKS

Relief **again!** Workers do PQ now, too!



Worker 2 Status:

```
Process Name: DW01
State: EXECUTING
Object Schema: HUGO
Object Name: TAB1
Object Type: SCHEMA_EXPORT/...
Completed Objects: 1
Total Objects: 1
Completed Rows: 85
Worker Parallelism: 7
```



How do we get more workers to export data?



Boost parallelism by using partitioned tables

LOB Export | Which Approach is Better?

Faking Statistics

- Must be done for each table
- Requires testing to get best result
- Could be overwritten by stats gathering

ESTIMATE=BLOCKS

- Just one parameter for the whole export
- Estimate phase adds time to export
- Requires patch applied in 19c

We are working on an approach that combines the best of both. Stay tuned to the upgrade blog!

"And BFILE LOBs?"

BFILE LOBs

External LOBs stored **outside** the database

Full export:

- Directory definition gets exported/imported
- You must copy the files

Schema export:

- You must create the directory within the database
- You must copy the files

Table export:

- You must create the directory within the database
- You must copy the files





Save downtime by copying the external files in advance

- **BFILES** are always read-only



If the directory path changes, ensure to update the **DIRECTORY** object

LOB Export | What if you still have BasicFiles LOBs?

Option 1

- Convert to SecureFile LOBs during impdp
 - `TRANSFORM=LOB_STORAGE:SECUREFILE`
 - `DB_SECUREFILE=ALWAYS`

Option 2


- Multiple Data Pump jobs in parallel exporting subsets of rows
 - `expdp parallel=1 table=1 query="where <subset 1>"`
 - `expdp parallel=1 table=1 query="where <subset 2>"`
 - `expdp parallel=1 table=1 query="where <subset 3>"`
 - `expdp parallel=1 table=1 query="where <subset 4>"`

Pro Tip: More details
in this [blog post](#)



Apply the Data Pump Bundle Patch

- <https://support.oracle.com/epmos/faces/DocumentDisplay?id=2819284.1>



Fewer Bugs


Important patches are included.
Monitor for bugs that affects many customers.



Faster Patching

The bundle patch changes the way Data Pump is patched. Subsequent patches apply faster.





Speeds up future patching significantly

- Can be applied online
- For RAC Rolling you must ensure to have no active Data Pump jobs

Data Pump Best Practices



Real World Checkpoint



Data Pump New Features and Best Practice [LRN3201]

OCW 2023, Las Vegas
September 21th, 2023

Moritz Werning, Solution Architect Banking, Swisscom

swisscom

C1 public




Speaker



Moritz Werning

- Solution Architect Banking @ Swisscom AG
- Master's degree in computer science (University of Tübingen in Germany)
- 15 years experience with the Oracle Database, Oracle Engineered Systems and DWH





Swisscom AG



- **Swisscom**, Switzerland's leading telecom company and one of its leading IT companies
- 2022 over 19,000 employees generated sales of CHF 11.1 billion
- It is 51% confederation-owned
- One of Switzerland's most sustainable and innovative companies

Swisscom Oracle Footprint:

- 24 data centers in Switzerland
- ~20 PB data online
- Oracle Exadata
- Oracle Exadata Cloud@Customer
- Oracle Private Cloud Appliance
- Oracle ZDLRA
- Oracle Database Appliance
- Oracle Exalogic
- Oracle Exalytics
- Oracle Big Data
- Real Application Clusters
- Active Data Guard
- Advanced Security
- Advanced Compression
- Golden Gate
- Data Integrator
- APEX & ORDS





Project: Oracle DB migration Exadata to ExaC@C via Data Pump



Project: Oracle DB migration BI Solution EXA on-premise to EXA@CC

Source

RDBMS 19.16 (non-CDB)

- 15 TB data
- 13 schema
- data of 8 banks per schema (VPD)



PDWH18



PDWHZ



PMTDFB

Exadata X4-2



Firewall

Target

On-Premises

RDBMS 19.18 (Multitenant)

CDB DWH PROD (Prim)

PDB (Prim)

PDWH1BIG

PDWH2

PDWH3

PDWH4

PDWH5

PDWH6

PDWH7

PDWH8

PDWHZ

PMTDFB



ExaC@C Gen 2 X9M QR - VM Cluster 1

BI Solutions

(peripheral system to the Finnova Core Banking System)

Challenges

1. No diagnostics & no insights
2. Firewall / network bottleneck
3. ORA-39358: Export dump file version 19.7.0.0.0
not compatible with target version
4. PARALLEL – don't mess with PARALLEL
5. A lot of LOBs (some classic)
6. Operational unbundling of schema data (VPD) into PDBs




Data Pump parameter (.par) files: How we started...

expdp parfile

```
CLUSTER=N  
DIRECTORY=DUMPLOCATION_EXACC  
DUMPFILE=PDWH18_XYZ%U.dmp  
LOGFILE=PDWH18_XYZ.log  
PARALLEL=10  
SCHEMAS=XYZ
```

impdp parfile

```
CLUSTER=N  
DIRECTORY=DUMPLOCATION  
DUMPFILE=PDWH01_XYZ%U.dmp  
LOGFILE=PDWH01_XYZ_IMP.log  
PARALLEL=10  
SCHEMAS=XYZ  
CONTENT=DATA_ONLY
```




Problem 1: No diagnostics & no insights

Problem:

- No diagnostics & no insights
 - Export 1:

```
exported "DWHAC"."OUT_DWH_T":"SYS_P572604" 1.436 GB 5186115 rows
```

Solution 1: Add diagnostics & set job name

expdp parfile

```
CLUSTER=N  
DIRECTORY=DUMPLOCATION_EXACC  
DUMPFILE=PDWH18_XYZ%U.dmp  
LOGFILE=PDWH18_XYZ_EXP.log  
JOB_NAME=expjob_XYZ  
LOGTIME=ALL  
METRICS=YES
```

impdp parfile

```
CLUSTER=N  
DIRECTORY=DUMPLOCATION  
DUMPFILE=PDWH1_XYZ%U.dmp  
LOGFILE=PDWH1_XYZ_IMP.log  
JOB_NAME=impjob_PDWH1_XYZ  
LOGTIME=ALL  
METRICS=YES
```




```
Job "SYS"."SYS_EXPORT_SCHEMA_01" successfully completed at Fri Apr 14 18:50:59 2023 elapsed 0 10:48:07
```

Problem 2: Firewall / network bottleneck

Problem:

- First export was too slow
- Firewall logs showed that we hit the maximum throughput

Job "SYS"."SYS_EXPORT_SCHEMA_01" successfully completed at Fri Apr 14 18:50:59 2023 elapsed 0 10:48:07





Solution 2: Add compression

Solution:

- Use compression to speed up your export
- Export size reduced by factor 4 and time has been halved
 - Size: 1.7 TB **uncompressed** => **423 GB compressed**
 - Time: 10:48:07 => **05:25:37**

```
Job ""SYS"". ""EXPJOB_XYZ"" successfully completed at Sat May 20 12:30:54 2023 elapsed 0 05:25:37
```

Export **without compression**:

```
exported "DWHC"."OUT_DWH_T":"SYS_P572604" 1.436 GB 5186115 rows
exported "DWHA"."HLP_STGFCT":"SYS_SUBP397268" 1.670 GB 7501886 rows
```

Export **with compression**:

```
18-MAY-23 10:49:21.430: W-6 . . exported "DWHC"."OUT_DWH_T":"SYS_P572604" 157.9 MB 5186115 rows in 24 seconds using direct_path
29-APR-23 10:40:28.465: W-1 . . exported "DWHA"."HLP_STGFCT":"SYS_SUBP397268" 213.0 MB 7501886 rows in 58 seconds using direct_path
```



Solution 2: Add compression

```
SQL> show parameter COMPRESS
NAME                VALUE
-----                ---
COMPRESS             ALL
COMPRESS_FOR         ALL
COMPRESS_FOR_INDEXES OFF
```

expdp parfile

```
CLUSTER=N
DIRECTORY=DUMPLOCATION_EXACC
DUMPFILE=PDWH18_XYZ%U.dmp
LOGFILE=PDWH18_XYZ.log
JOB_NAME=expjob_XYZ
LOGTIME=ALL
METRICS=YES
COMPRESSION=ALL
COMPRESSION_ALGORITHM=MEDIUM
SCHEMAS=XYZ
```

impdp parfile

Data Pump Compression

Benchmark by Oracle - 12.2 EBS Database export

	FILE SIZE MB	RATIO	TIME
NONE	5.500	1,0	4m 54s
ALL BASIC	622	8,9	4m 58s
ALL LOW	702	7,8	5m 24s
ALL MEDIUM	567	9,7	4m 55s
ALL HIGH	417	13,2	5m 13s

	FILE SIZE MB	RATIO	TIME
NONE	5.800	1,0	2m 33s
ALL BASIC	705	8,2	3m 03s
ALL LOW	870	6,6	8m 11s
ALL MEDIUM	701	8,2	3m 01s
ALL HIGH	509	11,3	12m 16s





Problem 3: ORA-39358: Export dump file version 19.7.0.0.0 not compatible with target version

Problem:

- First import attempt ... error... **ORA-39358**

```
Master table "SYS"."SYS_IMPORT_SCHEMA_01" successfully loaded/unloaded  
ORA-39358: Export dump file version 19.7.0.0.0 not compatible with target version 19.1.0.0.0
```

- Technical depts
 - Compatible parameter source DB **19.7**



Solution 3: Use the VERSION Parameter for the export



Solution:

- Use the VERSION parameter for the export to get rid of technical depts with compatible




Solution 3: Add VERSION=19.0.0

expdp parfile

```
CLUSTER=N  
DIRECTORY=DUMPLOCATION_EXACC  
DUMPFILE=PDWH18_XYZ%U.dmp  
LOGFILE=PDWH18_XYZ.log  
JOB_NAME=expjob_XYZ  
LOGTIME=ALL  
METRICS=YES  
COMPRESSION=ALL  
COMPRESSION_ALGORITHM=MEDIUM  
VERSION=19.0.0  
SCHEMAS=XYZ
```

impdp parfile





Problem 4: PARALLEL – Don't mess with PARALLEL

Problem 4.1:

- Export and import were slow and didn't use the system resources efficiently

Problem 4.2:

- Doubled **PARALLEL** degree for import
- **TABLE_DATA** import 2h faster , **but REF_CONSTRAINT** import 15 h **slower**, had to **abort** job

```
07-MAY-23 04:23:52.071: ORA-39082:Object type REF_CONSTRAINT:"DWH"."KD_VERMO.._FK02" created with compilation warnings
07-MAY-23 09:01:03.058: Job "SYS"."IMPJOB_DWHODS" stopped due to fatal error at Sun May 7 09:01:03 2023 elapsed 022:33:27
```





Solution 4: PARALLEL – follow the best-practices

Solution 4.1 :

- Follow PARALLEL best-practices
 - On-prem (x86-64) => CPU cores x 2
 - ExaCC / OCI => Number of OCPU per VM

Solution 4.2 :

- Don't mess with PARALLEL



Solution 4: Set best-practices PARALLEL degree

expdp parfile

```
CLUSTER=N
DIRECTORY=DUMPLOCATION_EXACC
DUMPFILE=PDWH18_XYZ%U.dmp
LOGFILE=PDWH18_XYZ.log
JOB_NAME=expjob_XYZ
LOGTIME=ALL
METRICS=YES
PARALLEL=24
COMPRESSION=all
COMPRESSION_ALGORITHM=medium
VERSION=19.0.0
SCHEMAS=XYZ
```

impdp parfile

```
CLUSTER=N
DIRECTORY=DUMPLOCATION
DUMPFILE=PDWH1_XYZ%U.dmp
LOGFILE=PDWH1_XYZ_IMP.log
JOB_NAME=impjob_PDWH1_XYZ
LOGTIME=ALL
METRICS=YES
PARALLEL=20
```



Problem 5: A lot of LOBs (some classic)

Problem 5.1:

- Source DB had a lot of "old/classic" LOB columns

Problem 5.2:

- Import of tables with LOB columns was slow



Solution 5: A lot of LOBs (some classic)

Solution 5.1:

- Convert old basic LOBs to SecureFile LOBs => SecureFile transformation

Solution 5.2:

- Set statistics on tables, partitions, subpartitions containing LOB

```
begin
  DBMS_STATS.SET_TABLE_STATS (
    ownname => 'DWH',
    tabname => 'BK_BENUTZER_T',
    partname  => 'SYS_SUBP489366',
    numrows  => 10000000,
    numblks  => 1000000);
end;
/
```




Solution 5: TRANSFORM LOB_STORAGE:SECUREFILE

expdp parfile

impdp parfile

```
CLUSTER=N  
DIRECTORY=DUMPLOCATION  
DUMPFILE=PDWH1_XYZ%U.dmp  
LOGFILE=PDWH1_XYZ_IMP.log  
JOB_NAME=impjob_PDWH1_XYZ  
LOGTIME=ALL  
METRICS=YES  
PARALLEL=20  
TRANSFORM=LOB_STORAGE:SECUREFILE
```



Recap: Oracle DB migration BI solution EXA on-premise to EXA@CC


Source

RDBMS 19.16 (non-CDB)


- 15 TB data
- 13 schema
- data of 8 banks per schema (VPD)



PDWH18



PDWHZ



PMTDFB

Exadata X4-2



Firewall

Target

RDBMS 19.18 (Multitenant)

On-Premises

BI Solutions

(peripheral system to the Finnova Core Banking System)



expdp



impdp

CDB DWH PROD (Prim)

PDB (Prim)

PDWH1BIG

PDWH2

PDWH3

PDWH4

PDWH5

PDWH6

PDWH7

PDWH8

PDWHZ

PMTDFB




Exadata C@C

ExaC@C Gen 2 X9M QR - VM Cluster 1



Problem 6: Operational unbundling of schema data (VPD) into PDBs



Problem:

- Operational unbundling
- 8 banks and 13 schema
- VPD
- Non-CDB into 8 PDBs



Solution: Oracle DB Migration BI Solution EXA on-premise to EXA@CC


Source

RDBMS 19.16 (non-CDB)

- 15 TB data
- 13 schema
- data of 8 banks per schema (VPD)



PDWH18



PDWHZ



PMTDFB



RA
C

Exadata X4-2



Exadata



Firewall

Target

RDBMS 19.18 (Multitenant)

On-Premises



ACFS




*.dmps
*.logs
*.par



Exadata C@C

ExaC@C Gen 2 X9M QR - VM Cluster 1

impdp



CDB DWH PROD (Prim)

PDB (Prim)

PDWH1BIG
PDWH2
PDWH3
PDWH4
PDWH5
PDWH6
PDWH7
PDWH8
PDWHZ
PMTDFB



RA
C

BI Solutions

(peripheral system to the Finnova Core Banking System)



Solution 6: Split Schema Consolidated Data into PDBs

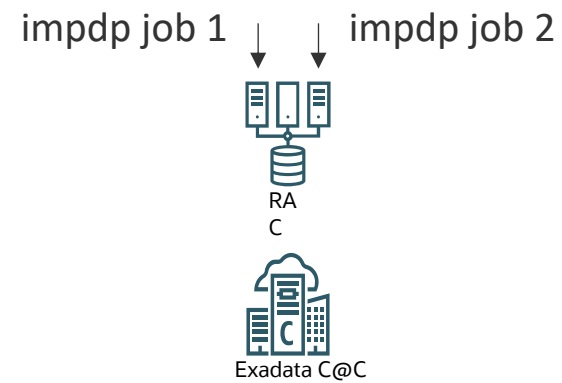
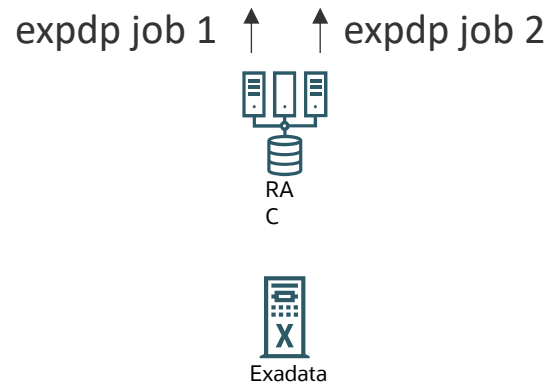
Solution 6.1:

- 1 x Datapump export per schema, without query filter to the ACFS of EXACC
- 8 x Datapump imports per schema with query filter, so one import per schema for each of the 8 PDBs

Solution 6.2:

Use the RAC nodes to distribute export & import jobs between the RAC nodes (Cluster N)

- Export some schema from node 1, some from node 2
- Import with two node in parallel one PDB of the big bank on node 1, small banks sequential on node 2





BI Solution: How we went live...

expdp parfile

```
CLUSTER=N
DIRECTORY=DUMPLOCATION_EXACC
DUMPFILE=PDWH18_XYZ%U.dmp
LOGFILE=PDWH18_XYZ_EXP.log
JOB_NAME=expjob_XYZ
LOGTIME=ALL
METRICS=YES
PARALLEL=24
COMPRESSION=all
COMPRESSION_ALGORITHM=medium
VERSION=19.0.0
SCHEMAS=XYZ
EXCLUDE=STATISTICS
EXCLUDE=RLS_POLICY
EXCLUDE=TABLE:"IN('BR...
```

impdp parfile

```
DIRECTORY=DUMPLOCATION
DUMPFILE=PDWH1_XYZ%U.dmp
LOGFILE=PDWH1_XYZ_IMP.log
JOB_NAME=impjob_PDWH1_XYZ
LOGTIME=ALL
METRICS=YES
PARALLEL=20
SCHEMAS=XYZ
REMAP_TABLESPACE=XYZ:DWH
REMAP_TABLESPACE=DWH_IDX:DWH
TRANSFORM=LOB_STORAGE:SECUREFILE
QUERY=XYZ.TAB1:"where BANKNR IN (1)"
QUERY=XYZ.TAB2:"where BANKNR IN (1)"
...
```



References

1. Data Pump Extreme – Deep Dive [Virtual Classroom Seminar 13: Data Pump Extreme - Deep Dive \(mikedietrichde.com\)](#)
2. Data Pump Best practice (<https://www.oracle.com/a/ocom/docs/oracle-data-pump-best-practices.pdf>)
3. Oracle Support Document 864582.1 (Examples Using DataPump VERSION Parameter And Its Relationship To Database COMPATIBLE Parameter) can be found at: <https://support.oracle.com/epmos/faces/DocumentDisplay?id=864582.1>
4. [Oracle Support Document 2819284.1 \(Data Pump Recommended Proactive Patches For 19.10 and Above\)](https://support.oracle.com/epmos/faces/DocumentDisplay?id=2819284.1) can be found at: <https://support.oracle.com/epmos/faces/DocumentDisplay?id=2819284.1>
5. [Oracle Support Document 555.1 \(Oracle Database 19c Important Recommended One-off Patches\)](https://support.oracle.com/epmos/faces/DocumentDisplay?id=555.1) can be found at: <https://support.oracle.com/epmos/faces/DocumentDisplay?id=555.1>



Contact information

Contact information

Swisscom AG

Data Management & Analytics

Moritz Werning

Pfingstweidstrasse 51

8005 Zürich

Phone +41-58-223 44 11

Email moritz.werning@swisscom.com

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**?

Roy Swonger



Daniel Overby Hansen



>100 years of database upgrade and migration experience

Combined in one single tool

CMA

Klaus Gronau



William Beauregard



Mike Dietrich



Cloud Migration Advisor

Rodrigo Jorge



Cloud Migration Advisor

Oracle Cloud Migration Advisor

Oracle Cloud Migration Advisor

Welcome to the Cloud Migration Advisor (CMA)

Oracle Cloud Migration Advisor brings you the expert technical knowledge of Oracle Database upgrade and migration-development teams, combined with more than a century of combined real-world experience with customer migrations, to give your customer the **best possible migration advice**.

With **Guided Mode**, CMA will quickly tell you:

- Which databases can be most easily migrated to Oracle Autonomous Database, or
- What is the **best migration method** to move chosen databases to a desired Oracle Cloud platform?

For more options, you can:

- Create and configure migration scenarios using **Create Project**, or
- Activate your existing migration projects and scenarios with **My Available Projects**.

Now it is time to start - let's move to the Oracle Cloud!

[Guided Mode](#) [My Available Projects](#) [Create New Project](#)

CMA - Oracle Cloud Migration Advisor

www.oracle.com/goto/upgrade

Watch on YouTube



Step 1

Customer collects
estate information

Cloud Migration Advisor

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



Customer
Fleet



Download
Extractor



Collect
Information



Secure Upload
to Oracle

Start here:

www.oracle.com/goto/upgrade



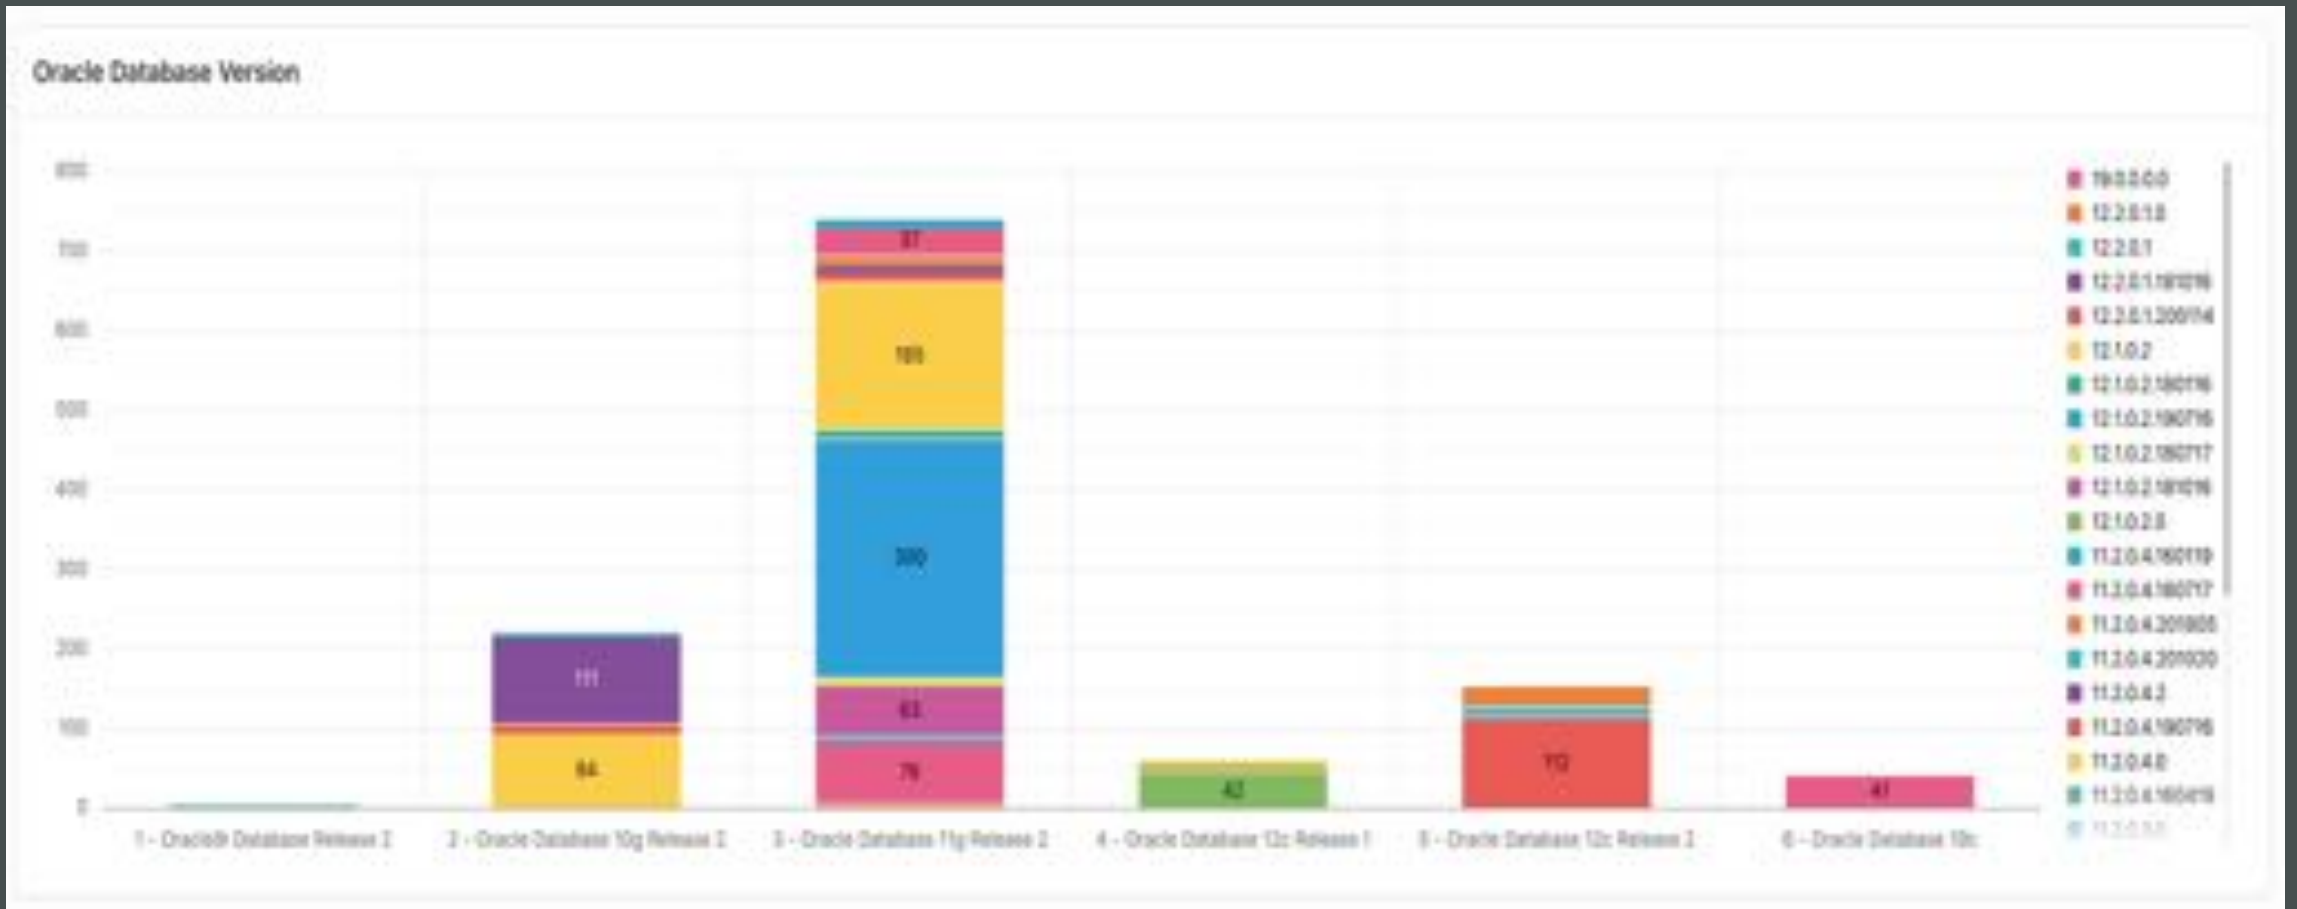
Step 2


**Load estate information
into CMA**



Cloud Migration Advisor

Landscape Overview





Step 3

**Add additional
information and
constraints**

Cloud Migration Advisor | Add New Target Server

Target Servers

Display Name
Exadata Cloud at Customer No.1

Is Source Host?

Host Name
ourexacc1

Operating System
-

Database Edition
Oracle Database Exadata Cloud Service

Cancel Create

-
- ADWD - Autonomous Data Warehouse Dedicated
- ADWS - Autonomous Data Warehouse Serverless
- ATPD - Autonomous Transaction Processing Dedicated
- ATPS - Autonomous Transaction Processing Serverless
- ✓ ExaCC - Exadata Cloud@Customer**
- ExaCS - Exadata Cloud Service
- Exadata On-premises
- Generic Server
- Oracle BaseDB System (VM, BM)



Step 4

Group the databases
into scenarios and
customize methods

Cloud Migration Advisor | Scenarios

Scenarios 1-Click Solution

To build a migration plan, the next step is to create a scenario where you will define the migration settings. Use the **"1-Click Solution"** or the **"1-Click Solution ADB"** buttons to let CMA automatically create one for you.

Search: All Text Columns Go Edit Save Reset Layout

Delete Selected

(Open)	Scenario Name	Assigned Databases
No data found		


Scenarios

To build a migration plan, the next step is to create a scenario where you will define the migration settings. Use the **"1-Click Solution"** or the **"1-Click Solution ADB"** buttons to let CMA automatically create one for you.

Search: All Text Columns Go

Delete Selected

(Open)	Scenario Name
	Can't be upgraded
	Must be upgraded



Step 5

Create solution

Cloud Migration Advisor | Solution - Methods

Migration Methods

<p>Data Pump Conventional Export/Import Complexity: 35 Downtime: Medium</p> <p>You can use this method regardless of the endian format and database character set of the source database. You can also use Data Pump to migrate data between different versions of Oracle Database. This method is simple to implement, provides the broadest cross-platform support and enables you to physically re-organize your target database.</p> <p><input type="checkbox"/> Select <input type="button" value="Read Instructions"/></p>	<p>Data Pump Import with DBLink Complexity: 35 Downtime: Medium</p> <p>The Oracle Data Pump Import command-line mode <code>NETWORK_LINK</code> parameter enables an import from a source database identified by a valid database link. The data from the source database instance is written directly back to the connected database instance.</p> <p><input type="checkbox"/> Select <input type="button" value="Read Instructions"/></p>	<p>Remote Cloning Non-CDB + Upgrade + Convert Complexity: Downtime: Medium</p> <p>Since Oracle 12.2.0.1 or higher (or with undo restrictions in 12.1.0.2 already), you have the freedom to plug a non-CDB at first via a database link, and then upgrade and adjust it.</p> <p><input checked="" type="checkbox"/> Select <input type="button" value="Read Instructions"/></p>
<p>Unplugging/Plugging Non-CDB + Upgrade + Convert Complexity: Downtime: Medium</p> <p>You can use the unplug/plug method to migrate an Oracle Database non-CDB to a PDB. This method provides a way to consolidate several non-CDB databases into a single Oracle Database multitenant database on the Database service. Finally, you run <code>dbupgrade</code> to bring it to the target version.</p> <p><input type="checkbox"/> Select <input type="button" value="Read Instructions"/></p>	<p>Data Pump Full Transportable Complexity: Downtime: Medium</p> <p>You can use the Data Pump full transportable method to copy an entire database from your on-premises host to the database on a Database service deployment.</p> <p><input checked="" type="checkbox"/> Select <input type="button" value="Read Instructions"/></p>	<p>Database Migration Workbench Physical Online Complexity: Downtime: Medium</p> <p>MWB uses RMAN Transportable Tablespace. TTS provides a Near Zero Downtime (NZDT) way to create a duplicate database from incremental backups of the source database while retaining the original target database.</p> <p><input type="checkbox"/> Select <input type="button" value="Read Instructions"/></p>



What else?

Important Features

Cloud Migration Advisor | Features


- Obfuscate hostnames and database names
- Configure individual weight and scoring
- Automatic purge
- Strict user isolation

Cloud Migration Advisor | Coming soon

Downloadable version

- Virtual Box image
- Marketplace image

CMA | Workflow



**Create
Project**

**Extract
& Load**

**Update
Data**

**Define
Scenario**

**Select
Target**

**Solution
& Success**

Lunch Break

We start again at 13:30

AGENDA

09:30

Welcome
Release Strategy
Patching

11:15

Upgrade
Data Pump
Cloud Migration Advisor

13:30

Performance Stability
Insights into
development

15:15

Oracle Database 23c
What's New
What's Coming

11:00

Coffee break

12:45

Lunch

15:00

Coffee break

Performance Stability Prescription



A person is standing in the ocean, with their right hand raised towards a cloudy sky. The water is dark and rippled. The overall mood is contemplative and dramatic.

what's your

biggest fear

when making changes?



A hammock is strung between two palm trees on a beach. The scene is captured at sunset, with a warm, golden glow in the sky and the ocean. The hammock is empty and hangs in a gentle curve. The palm trees are silhouetted against the bright sky. The overall mood is peaceful and relaxing.

Performance Stability Prescription

Performance Stability Prescription



Collect workload information

- Sample from cursor cache
- Gather from AWR

Performance Stability Prescription



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

SQL Tuning Set | Definition

SQL statement



Context



Statistics



Plans





Gather at least a full month of workload data

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

Workload Information



AWR – Automatic Workload Repository

Change the retention to a minimum of 40 days

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



Collect SQL statements and plans

Use AWR as main source

Capture from Cursor Cache for OLTP

Collect statements, plans and stats in SQL Tuning Sets

Performance Stability Prescription



Upgrade test database

Load workload data
(SQL Tuning Set)

Performance Stability Prescription



AWR Diff Report

SQL Performance Analyzer tests your workload

Report with all regressing statements

AWR | Diff Report

Use script `awrddrpt.sql`

Top Timed Events

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

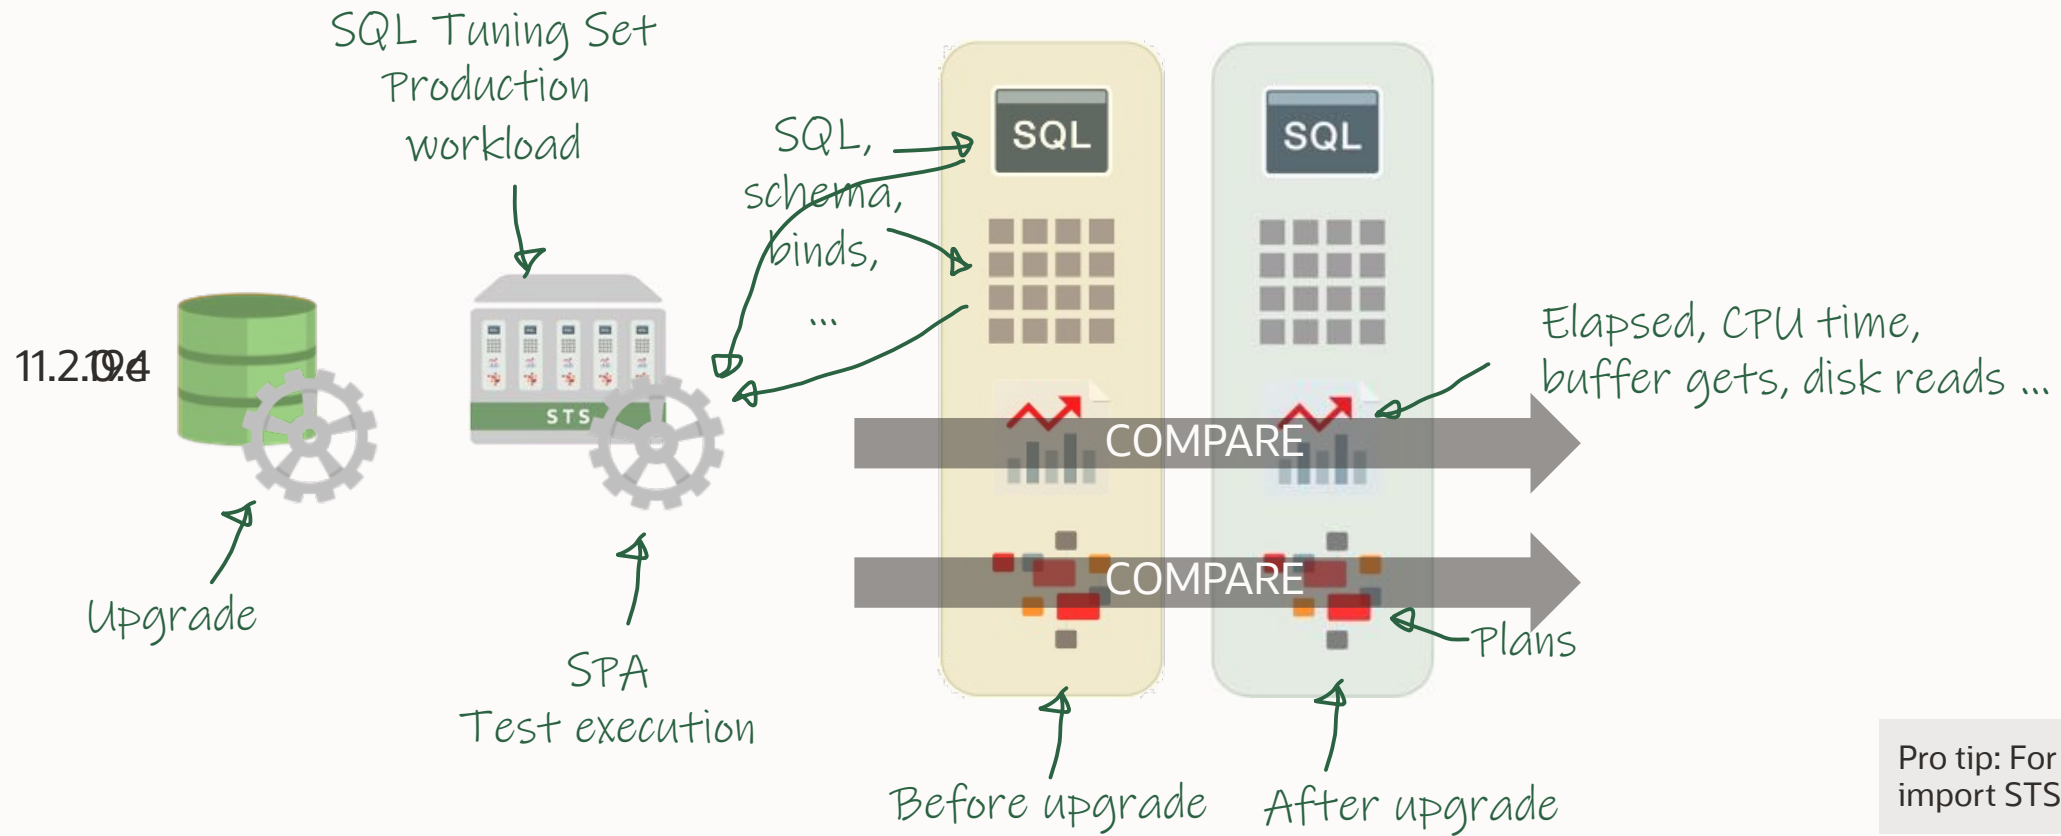
1st						2nd					
Event	Wait Class	Waits	Time(s)	Avg Time(ms)	%DB time	Event	Wait Class	Waits	Time(s)	Avg Time(ms)	%DB time
CPU time			68,289.05		43.73	db file sequential read	User I/O	22,193,998	114,919.21	5.18	23.17
db file sequential read	User I/O	6,688,953	37,737.81	5.64	24.17	enq: SS - contention	Configuration	3,913	98,997.90	25,299.74	19.96
gc buffer busy	Cluster	12,508,244	23,886.55	1.91	15.30	CPU time			73,786.55		14.88
TCP Socket (KGAS)	Network	680,629	12,514.65	18.39	8.01	row cache lock	Concurrency	73,940	48,472.30	655.56	9.77
db file scattered read	User I/O	1,572,296	4,271.68	2.72	2.74	reliable message	Other	41,148	47,600.87	1,156.82	9.60

Requires Enterprise Edition + Diagnostic pack

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



SQL Performance Analyzer | Concept



Pro tip: For migrations, import STS into target database



SQL Performance Analyzer | Report

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

object_id	sql_id	Impact on Workload	Execution Frequency	Metric Before	Metric After	Impact on SQL	Plan Change
52	csy0xdm9c394f	4.02%	3262	6149.0885959534	4208	31.57%	n
41	7m5h0wl6stg0q	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	88fggncchy6wg	-.41%	325424	6.00316202861498	8	-33.26%	n
57	g5u7xuchhfu62	.39%	32790	26.6833180847819	8	70.02%	n



SQL Performance Analyzer | Report



Execution Plan Before Change:

Plan Hash Value : 3642382161

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



Execution Plan After Change:

Plan Id : 138

Plan Hash Value : 1075826057

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



SQL Performance Analyzer | Report

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

From production workload

From test execution

SQL Performance Analyzer | Report

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



SQL Performance Analyzer | Report

Regressed SQL Statements						
	SQL ID	Net Impact on Workload (%)	Buffer Gets		Net Impact on SQL (%)	New Plan
			SQL Trial 1	SQL Trial 2		
↓	3fv28qfu9y0aq	-0.050	26,504	29,573	-11.580	Y
↓	czzzubf8fjz96	-0.030	1,410	1,981	-40.500	Y

SQL Details: czzzubf8fjz96

Parsing Schema: APPS Execution Frequency: 3

> SQL Text

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

Single Execution Statistics

	Execution Statistic Name	Net Impact on Workload (%)	Execution Statistic Collected		Net Impact on SQL (%)
			SQL Trial 1	SQL Trial 2	
↓	Elapsed Time (sec)	-0.240	0.112	0.364	-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		
↓	3fv28gfu9y0aq	-0.050	26,504	29,573	-11.580	Y
↓	czzzubf8fjz96	-0.030	1,410	1,981	-40.500	Y

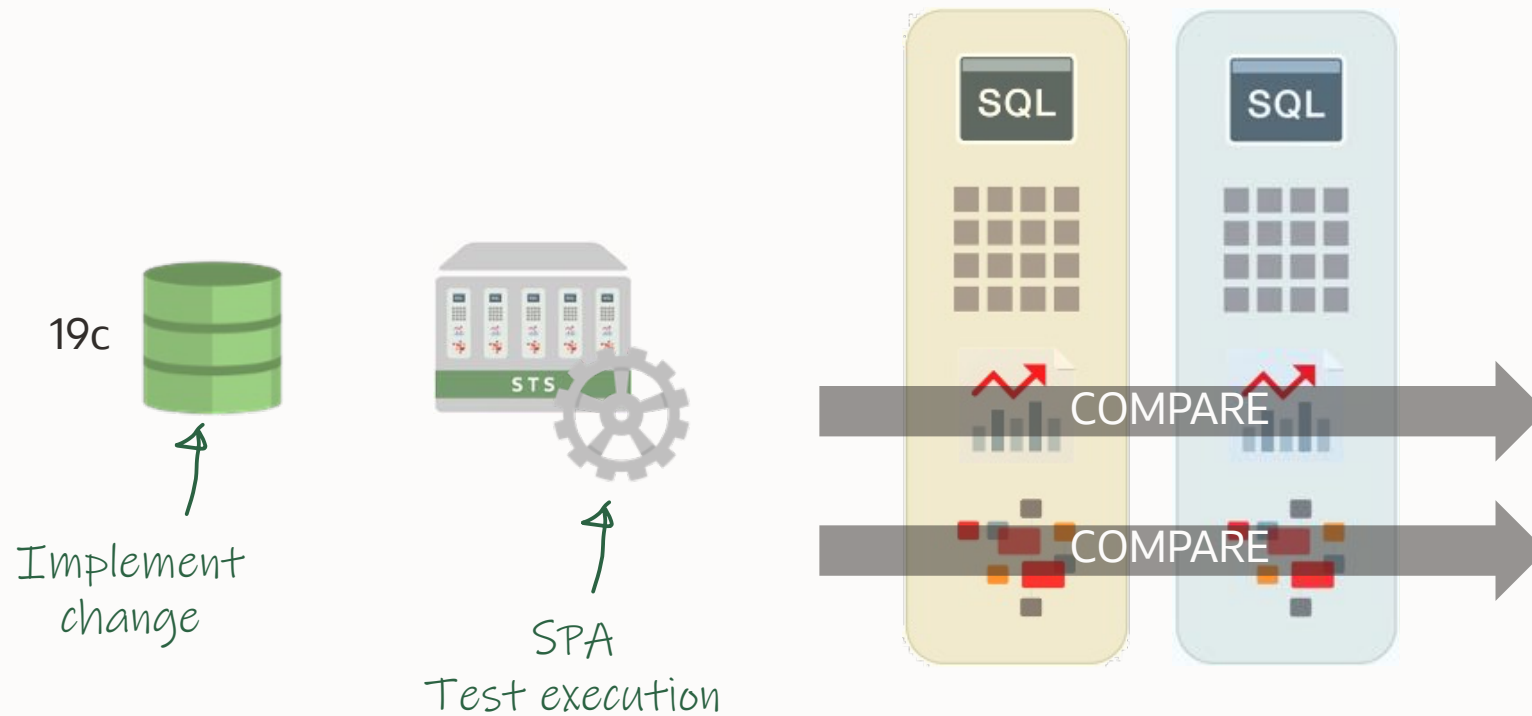
Plan Comparison
 SQL_TRIAL_1353942463446
 Plan Hash Value: 1165613724

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".PER...
INDEX RANGE SCAN	5	APPS.EXES_080F0004	1	2	T4".EXPORT_LIC_NR"=14659
HASH JOIN	6		14,210	9,792	T1".SKU_NR"=T2".SKU_NR" AN...



SPA | Continuous Improvement



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

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

Performance Stability Prescription



Tune SQLs with regressed plans

Create SQL Plan Baselines

Transport to production database

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

    dbms_sqltune.execute_tuning_task(l_task);

    l_report := dbms_sqltune.report_tuning_task(l_task);

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

SQL Tuning Advisor | Example

FINDINGS SECTION (8 findings)

1- Statistics Finding

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

Recommendation

- Consider collecting optimizer statistics for this table.

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

Rationale

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

SQL Tuning Advisor | Example

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

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

Recommendation (estimated benefit: 67.2%)
-----
- Consider accepting the recommended SQL profile.
  execute dbms_sqltune.accept_sql_profile(task_name => 'TASK_21944',
    task_owner => 'SYS', replace => TRUE);
-----
Validation results
-----
The SQL profile was tested by executing both its plan and the original plan
and measuring their respective execution statistics. A plan may have been
only partially executed if the other could be run to completion in less time.
```

SQL Tuning Advisor | Example

7- Index Finding (see explain plans section below)

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

```
DBMS_SQLTUNE.REPORT_TUNING_TASK(:STMT_TASK)
```

Recommendation (estimated benefit: 88.23%)

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

Rationale

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

6- SQL Profile Finding

A potentially better

Recommendation (est

- Consider acceptin
execute dbms_sqlt
task_owne

Validation results

The SQL profile was
and measuring their
only partially exec

FINDINGS SECTION (in paragraphs)

1- Statistics Finding

Optimizer statistics for S

Recommendation

- Consider collecting appr
execute dbms_sqlt
"STATS", user
method_opt = 'F'

Rationale

The optimizer requires a
select a good execution





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



SQL SQL

← Repeatable SQL



Filtering changed

- Slow statistics
- Slow parameters
- Upgrade
- Action

OPTIMIZER_USE_SQL_PLAN_BASELINES=TRUE

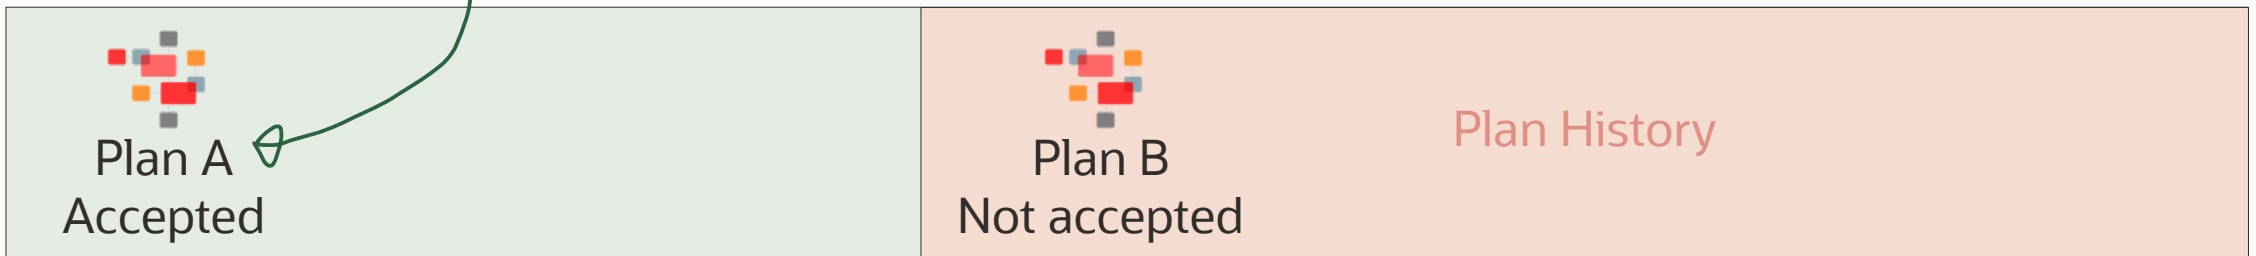
No plans in baseline

Plan B Filter

Plan A is used

Plan history

Dedicated part of SQL plan baseline





SPM | Evolve



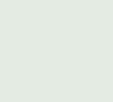
Result:
Performance **better**

Test execute Test execute
Plan stays




Plan A
Accepted


Plan C
Accepted


Plan D
Accepted


Plan B
Not accepted


Plan D
Not accepted

Plan History



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

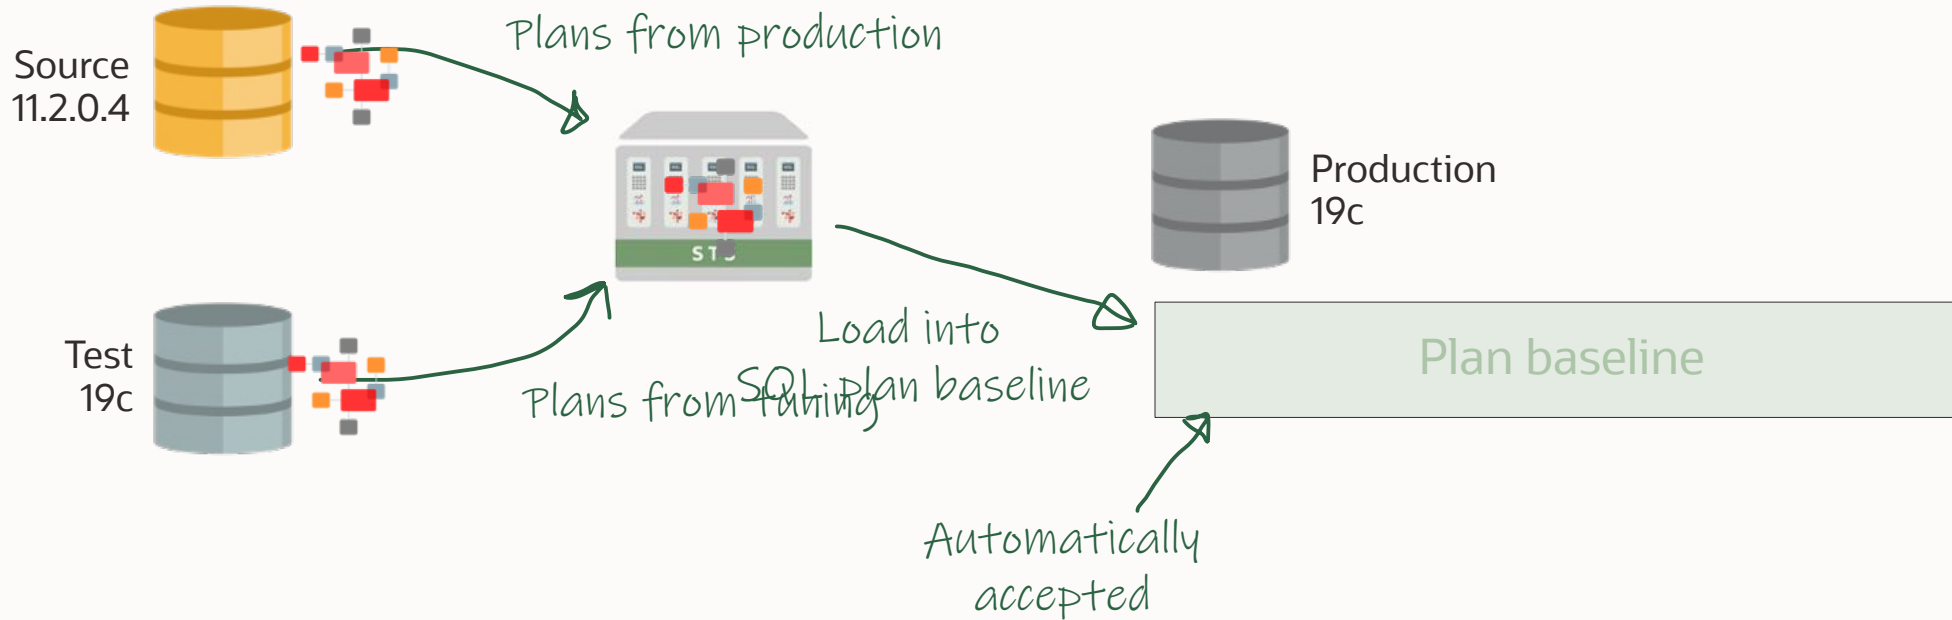


Plan B
Not accepted

Plan History

Automatically
accepted

SPM | Use Case




SPM | What if ... literals

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

- Many distinct statements
- `CURSOR_SHARING = FORCE?` No!

Optimal solution: Change your application to use bind variables



Use SQL Profiles for statements with literals

- Part of Tuning Pack

```
dbms_sqltune.accept_sql_profile(..., category=>'TEST_ENV');  
  
alter session set sqltune_category='TEST_ENV';
```



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

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

```
--After testing, fully enabled profile
```

```
dbms_sqltune.alter_sql_profile(..., attribute_name=>'CATEGORY', value=>'DEFAULT');
```

SQL Profiles | Facts

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

There is only one tool to ensure plan stability:

SQL Plan Management

Don't use

- OPTIMIZER_FEATURES_ENABLE
- COMPATIBLE



try it out for free

IT'S EASY



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

- Relevant options and packs included in most cloud offerings

Use OCI for performance testing



Performance Tips & Tricks



Use as few initialization parameters as possible

- Stick to the defaults
- Stick to vendor recommendations



Only use underscores and events to solve specific situations

- Only under guidance of Oracle Support

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



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`



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



BIG PROBLEM!!

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

... you hit a bug!!



A photograph of two men laughing joyfully while holding a large, inflatable character. The character is yellow and red, resembling a stylized figure. The background shows an outdoor setting with buildings and other people, suggesting a festival or public event.

But ... don't worry!

It is fixed! In Oracle 23c!

A person wearing a black cap with a crest, looking thoughtful against a brick wall. The person is wearing a black t-shirt and has their hand near their chin. The background is a brick wall with a red section on the left and a tan section on the right.

Ok ...??

You mean in
2023?

THREE PARTIES

- You
- Oracle Support
- Oracle Development

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 41886745 : ORA-600 [happens_on_mondays_only] before java is brewed

Bug Attributes

Type	B - Defect	Fixed in Product Version	
Severity	2 - Severe Loss of Service	Product Version	19.17.0.0.0
Status	10 - New	Platform	226 - Linux x86-64
Created	Oct 17, 2022	Platform Version	ORACLE LINUX 7
Updated	Oct 17, 2022	Base Bug	N/A
Database Version	19.17.0.0.0	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: 41886745 : 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



ERROR



FIX



BACKPORT



INCLUSION

- Developer verifies the fix
- 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



ERROR



FIX



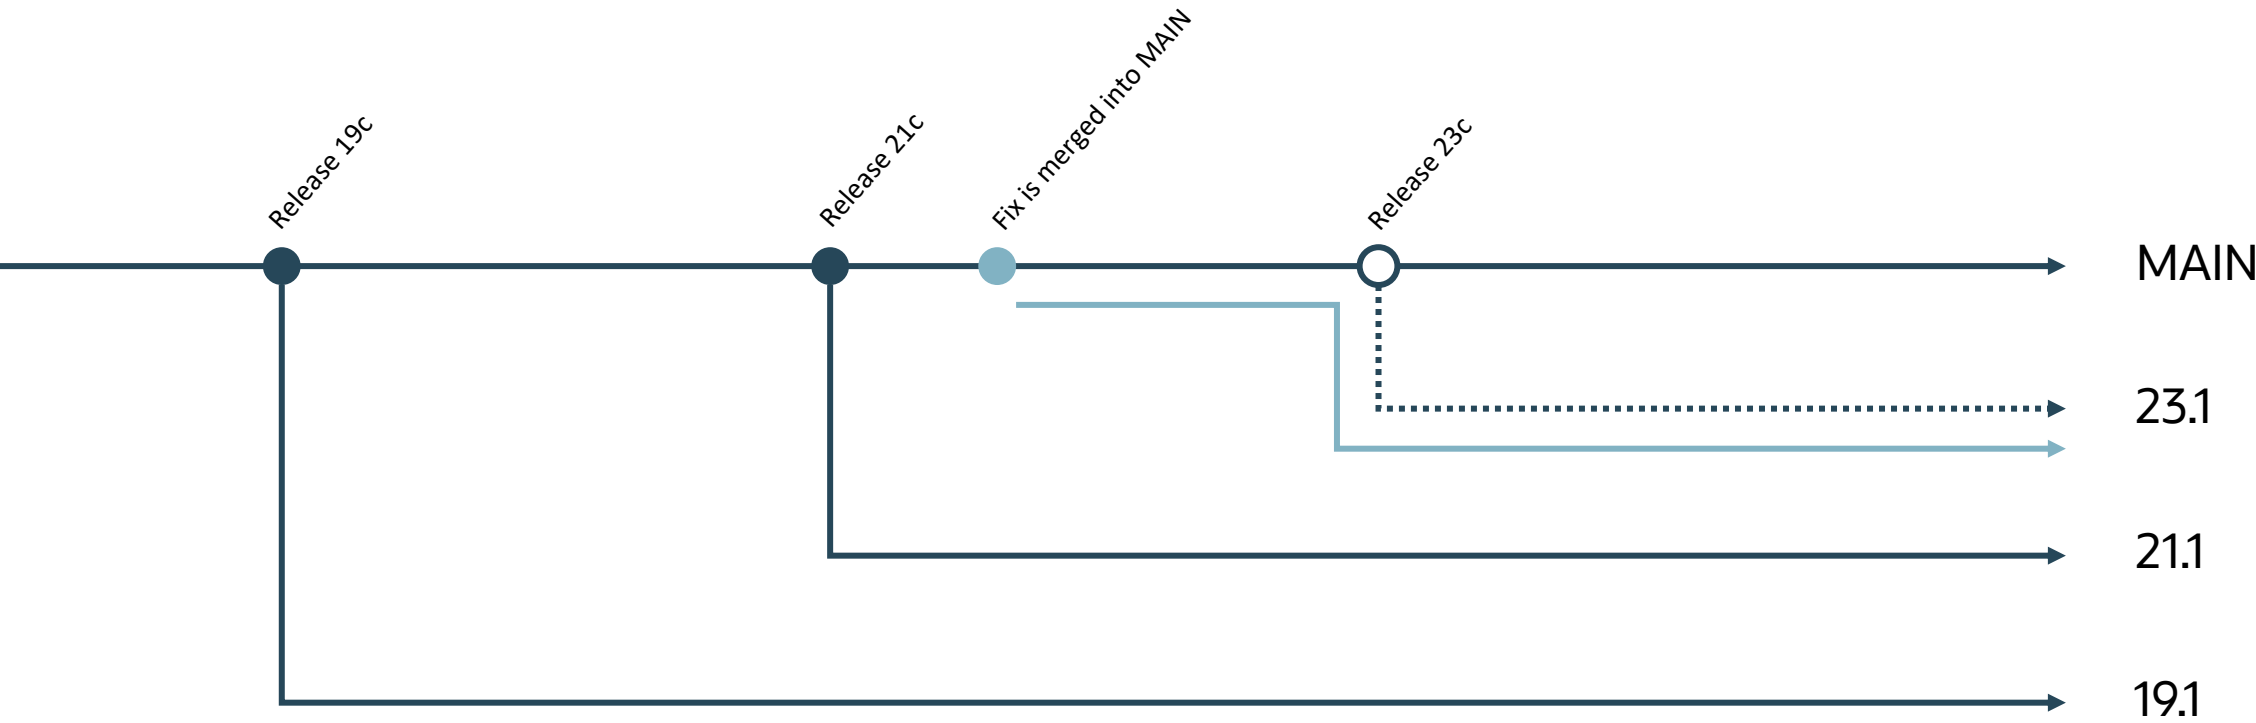
BACKPORT



INCLUSION

- 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 Attributes**

Type	B - Defect	Fixed in Product Version	23.1
Severity	2 - Severe Loss of Service	Product Version	19.17.0.0.0
Status	80 - Development to QA/Fix Delivered Internal	Platform	226 - Linux x86-64
Created	Oct 17, 2022	Platform Version	ORACLE LINUX 7
Updated	Oct 17, 2022	Base Bug	N/A
Database Version	19.17.0.0.0	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: 41886745 : 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 23c now?**

But there is only the 23c Free Edition??

A photograph of two women sitting on a yellow leather sofa in a cozy, wood-paneled room. The woman on the left, wearing a white sweater and blue jeans, is laughing heartily while holding a colorful map. The woman on the right, wearing a blue beanie, glasses, and a green sweater, is smiling and looking at a blue laptop. The scene is warm and inviting, with patterned pillows and a white door in the background.

No worry ...

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

Insights into the Patching Process



ERROR



FIX



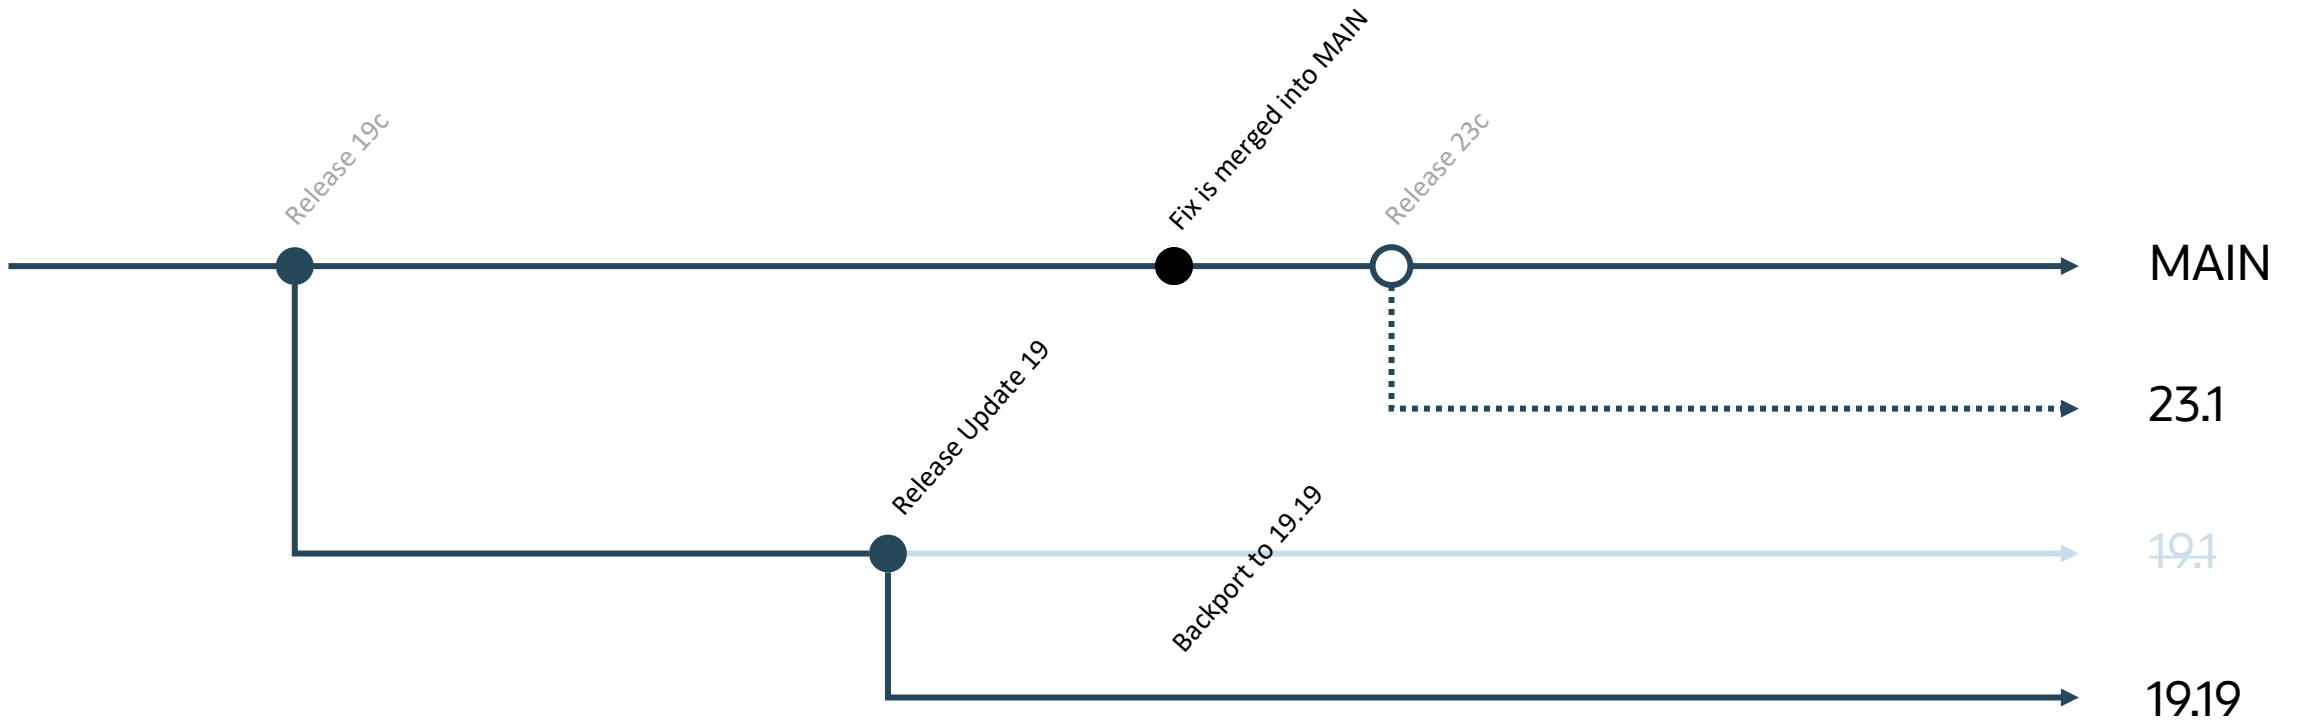
BACKPORT



INCLUSION

- Backport created on request
- Backports are specific for a release and usually also for a patch level, e.g., on top of 19.19.0
- `opatch lsinventory`

Backport



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)



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 19.18 DB Home

Below is the list of important patches to consider applying on top of 19.18. 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	NON ROLLING	Added
35037877 (replaces 20289608)			[SECURITY] EM patching may fail with ORA-4067	[list: patches]		20-APR-2023
32727143	19.19		[SQL EXECUTION] Transaction-level content isolation for transaction-duration global temporary tables	[list: patches]		20-APR-2023
34557500 (replaces 31544097)	19.19	DBMRP 19.17.0.0.230321 DBMRP 19.18.0.0.230321	[BLOCK TRACK] CTWR caused multiple instances to hung state on the RAC Standby DB	[list: patches]		20-APR-2023
34340632			[AQ] Smart Monitoring & Resiliency in AQ KGL Memory Usage To Help Message Cleanup And Prevent ORA-600 [KGL-HEAP-SIZE-EXCEEDED].	[list: patches]		14-APR-2023
35246710 (replaces 33803836)	19.19	DBMRP 19.18.0.0.230418	[BUFFER CACHE] High "Direct Path Read" Waits After 19.18 DBRU Patching	[list: patches]		08-APR-2023
34832725 (replaces			[SHRD CRSRS] ORA-4031 and / or High Shared Pool Latch Contention During Session Creation in	[list: patches]		25-APR-2023





Request for Inclusion

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

[Edit this bug](#) [Bug Tree](#) [ARU Report](#)

Bug No: 32765738 (Bug)

Filed By: [REDACTED]

Updated By: [REDACTED]

Sup Rep: [REDACTED]

Customer: [REDACTED]

Status: 80 - Development to QA/Fix Delivered Internal



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

CHG: BPS:No of customers->179->180

CHG: SR:3-3045-3045-1 (ORION) Attached



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

Close

 **Patch 30978304: ORA-20000 DURING IMPDP WITH STATS AND THE UNIQUE INDEX FOR THE PK IS NOT CREATED**

Last Updated: Jun 2, 2021 2:04 PM (2+ months ago)

Product: Oracle Database - Enterprise Edition (More...)	Size: 1.3 MB
Release: Oracle Database 19.12.0.0.0 DBRU	Download Access: Software
Platform: Generic Platform	Classification: General
	Patch Tag

Prerequisite Patches

32904851	DATABASE RELEASE UPDATE 19.12.0.0.0
----------	-------------------------------------

Bugs Resolved by This Patch

30978304	ORA-20000 DURING IMPDP WITH STATS AND THE UNIQUE INDEX FOR THE PK IS NOT CREATED
----------	--

[View Related Knowledge to this Patch](#)

Release: Oracle Database 19.12.0.0.0 DBRU

Platform: Generic Platform

Language: American English

[Read Me](#) [Download](#) [Add to Plan](#)

[Analyze with OPatch...](#)

[Discuss this patch in the community](#)



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 Updates and Revisions regression tests

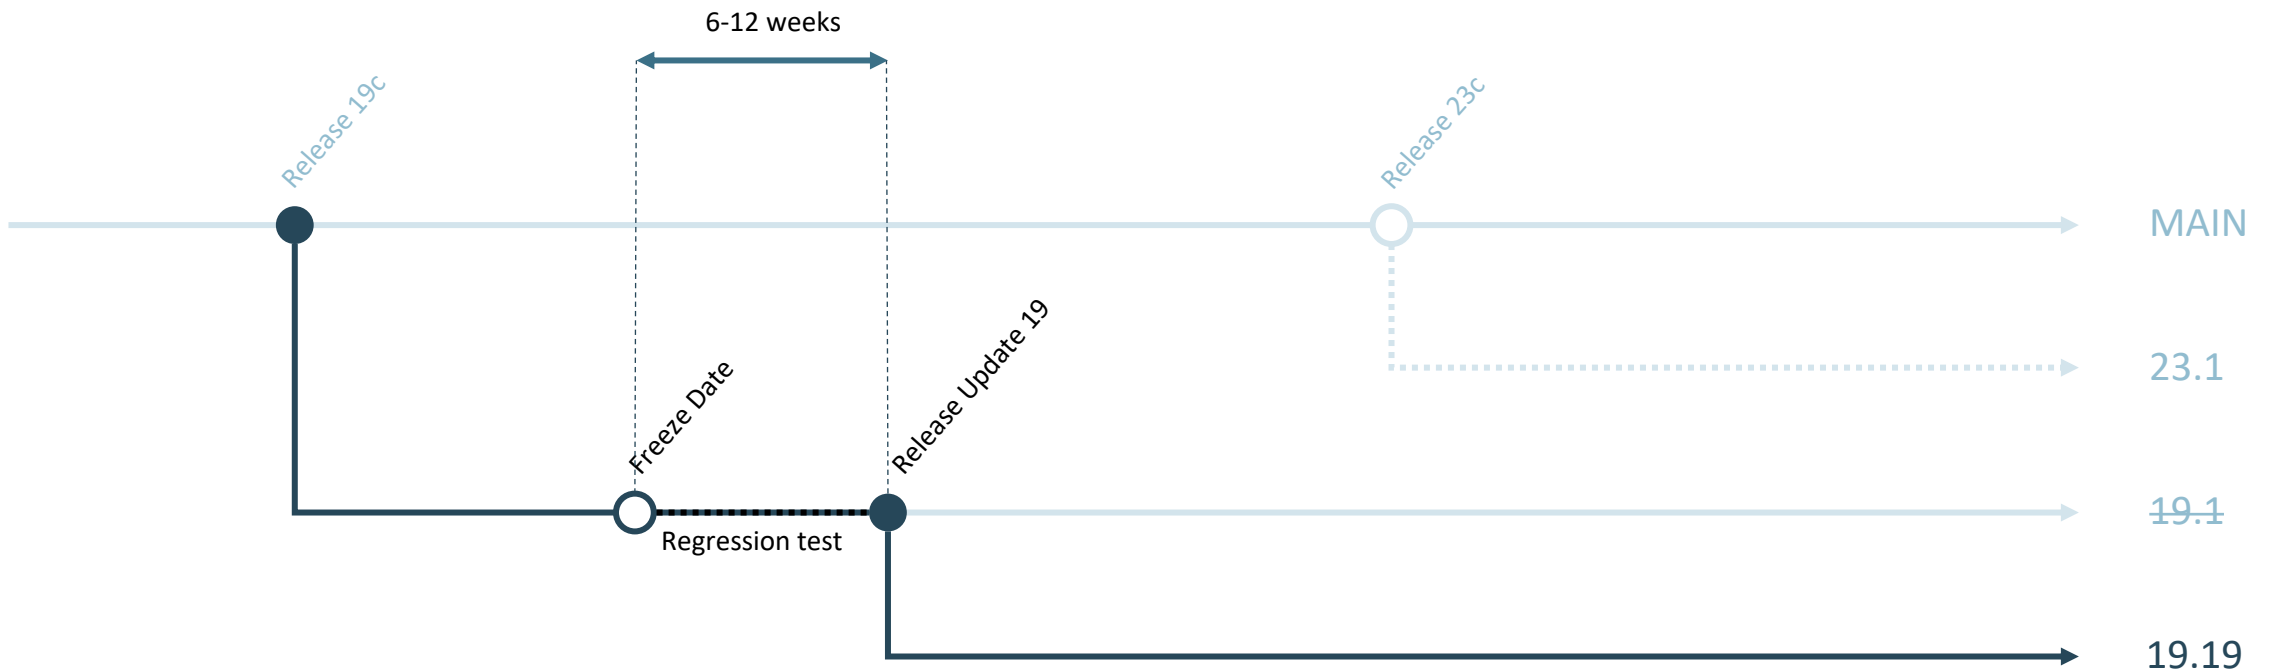
Limitation



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



Limitation



WHY DO YOU ALWAYS ASK FOR

OPATCH Isinventory



TO AVOID

CONFLICTS



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

Invoking prereq “checkconflictagainsthwithdetail”

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

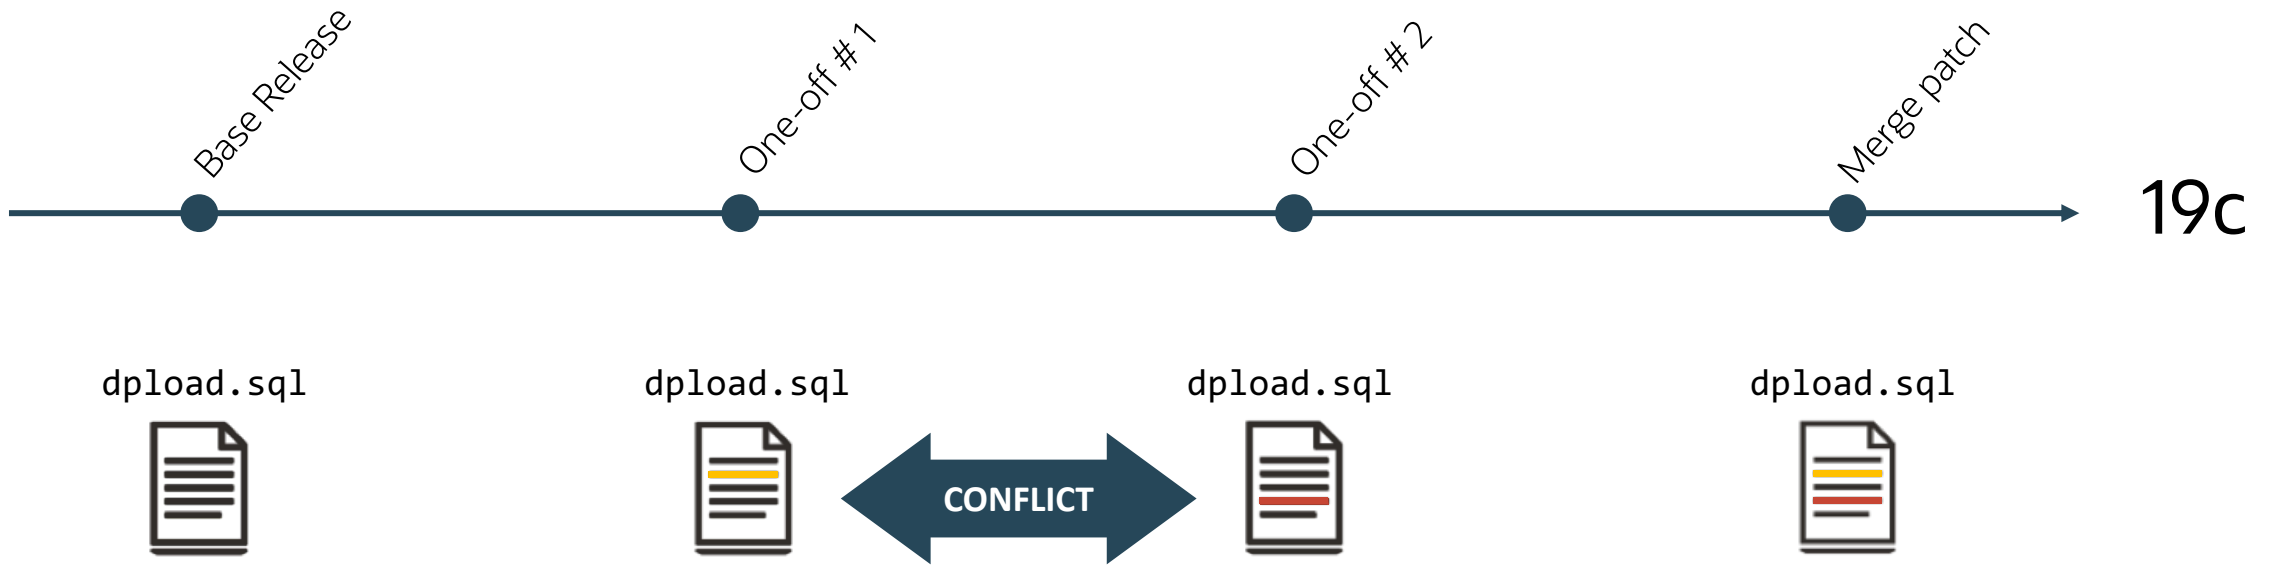
Conflict with 35012562

Conflict details:

/u01/app/oracle/product/19.19.0/db_1/lib/libserver19.a:kko.o

OPatch succeeded.

Basic Facts | Conflicts



19c





Always attach the output of
`opatch lsinventory` to your Service Request





Don't name it

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





Just name it `opatch_1sinventory.txt`

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



5031
Meth...

Re
Road

Question

Scary

+
Equation &
Algebra
Calculus
... ..

Windows is different ...



Please wait while we install a system update



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?

Knowledge

RCA

DWB

Customer Exceptions

CMS

More...



Copyright (c) 2023, Oracle. All rights reserved. Oracle Confidential.

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?

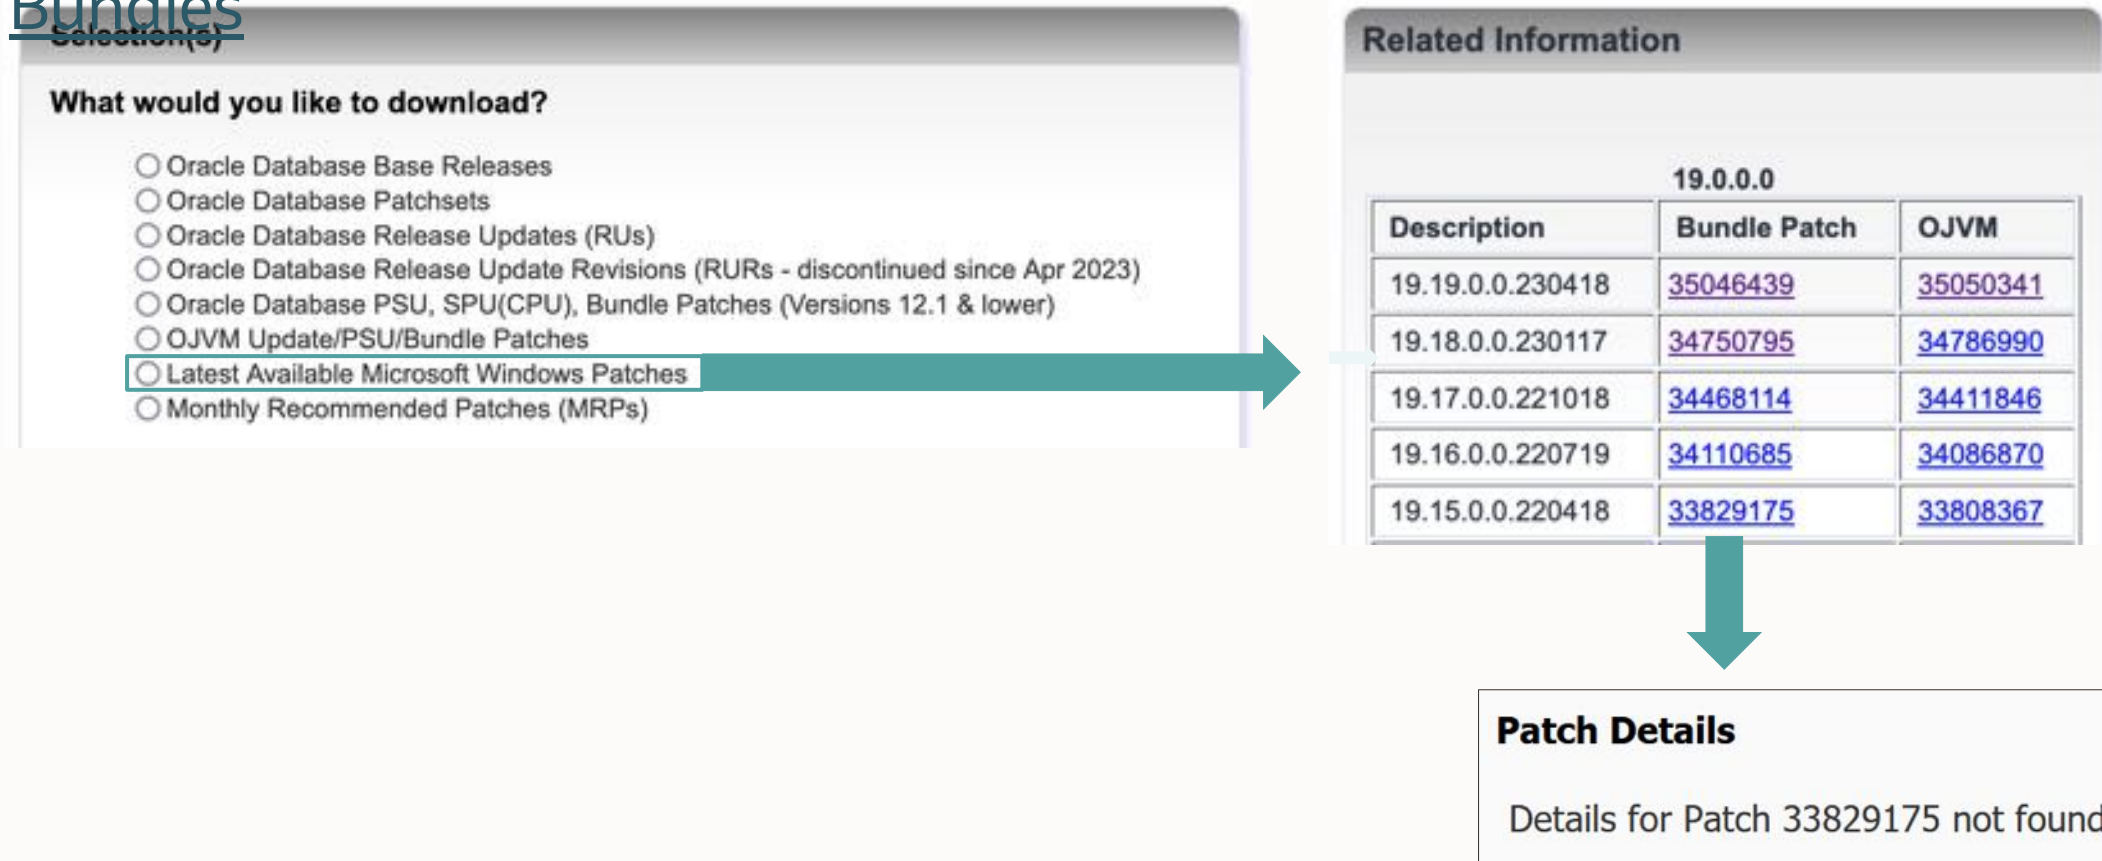
- Oracle Database Base Releases
- Oracle Database Patchsets
- Oracle Database Release Updates (RUs)
- Oracle Database Release Update Revisions (RURs - discontinued since Apr 2023)
- Oracle Database PSU, SPU(CPU), Bundle Patches (Versions 12.1 & lower)
- OJVM Update/PSU/Bundle Patches
- Latest Available 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



Basic Facts | Patch Availability

MOS Note: 2844795.1 - Apr 2022 Patch Availability Document (DB-only) (Doc ID 2844795.1)

Released April 19, 2022

This document contains the following sections:

- [Critical Patch Update April 2022 Patch Availability Document \(PAD\)](#)
 - [1 Overview](#)
 - [1.1 How To Use This Document](#)
 - [1.2 Terminology in the Tables](#)
 - [1.3 On-Request Patches](#)
 - [1.4 CPU Program and My Oracle Support Patch Recommendations](#)
 - [1.5 My Oracle Support \(MOS\) Conflict Checker Tool](#)
 - [2 What's New in April 2022](#)
 - [2.1 "Final CPU Information \(Error Correction Policies\)"](#)
 - [2.2 "Post Release Patches"](#)
 - [2.3 "Separate PADs for Separate Products"](#)

Patch Details

Details for Patch 33829175 not found.

Microsoft Windows BP 21.0.0.220419	Patch 33829143	MS-Windows	14-Jun-2022
Microsoft Windows BP 19.15.0.0.220419	Patch 33829175	MS-Windows	17-May-2022



Find the PAD by searching for
Patch Availability Document *Apr 2023*



... Oracle on Windows is different



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

A photograph of a penguin standing on a large, jagged piece of ice in a snowy, hazy landscape. The penguin is positioned in the center of the frame, facing away from the viewer. The ice is white and blue, with sharp edges and a textured surface. The background is a soft, hazy expanse of snow and sky.

Oracle Linux Development Platform

Porting



Oracle Linux x86-64

Porting

- Microsoft Windows
- 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	Tekelec	Netra X3-2 for Acme Packet	Novell NetWare
Oracle JRockit Virtual Edition x86	Qualcomm Brew MP	Oracle Solaris on SPARC (64-bit)	Linux on IBM Z
SunOS	Netra Server X5-2 for Communications	IBM S/390 Based Linux (31-bit)	Data General
Monta Vista x86	HP NonStop S-series (Guardian)	Acme Packet 3900	Pyramid
Acme Packet 1100	HP OpenVMS Itanium	SPARC	Talari
iTron	Monta Vista x86-64	Fujitsu MSP-EX	Palm Computing
Embedded Linux on cnMIPS	OpenSolaris	Trusted Solaris	HP NonStop (OSS) on x86
Embedded Linux SH4	SCO Unix	Net-Net 4250	Unisys OS 2200
HP NonStop Itanium (OSS)	Net-Net 9200	HP OpenVMS VAX	HP OpenVMS Alpha
QNX Unix	Symbian EPOC	HP-UX PA-RISC (32-bit)	Acme Packet 3820
Acme Packet 6100	Linux ARM 32-bit VFP HardFP ABI	Acme Packet 6300	FreeBSDx86
Linux MIPS 64-bit	SGI Irix	Microsoft Windows CE	Oracle Solaris Express
Fujitsu BS2000/OSD (SQ series)	ia64	IBM z/OS on System z	VxWorks
Mediatek MTZ	Linux SPARC	StorageTek Hardware	Microsoft Windows (32-bit)
HP NonStop (Guardian) on x86	Oracle Solaris on SPARC (32-bit)	Oracle Solaris on x86 (32-bit)	Sequent
Fujitsu BS2000	HP NonStop Itanium (Guardian)	Fujitsu BS2000/OSD (SX series)	Windows NT
HP Tru64 UNIX	RIM BlackBerry	Linux ARM 64-bit	nCube

Break

We start again at 15:15



Oracle Database 23c

—
What's Changing



Consult the [Upgrade Guide](#) for changes, desupports, and deprecations



Traditional Auditing is desupported in Oracle Database 23c

- Migrate your policies to Unified Auditing



TRADITIONAL AUDITING

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

UNIFIED AUDITING

- Use syntax converter script (MOS Doc ID [2909718.1](#))
- Do it before the upgrade





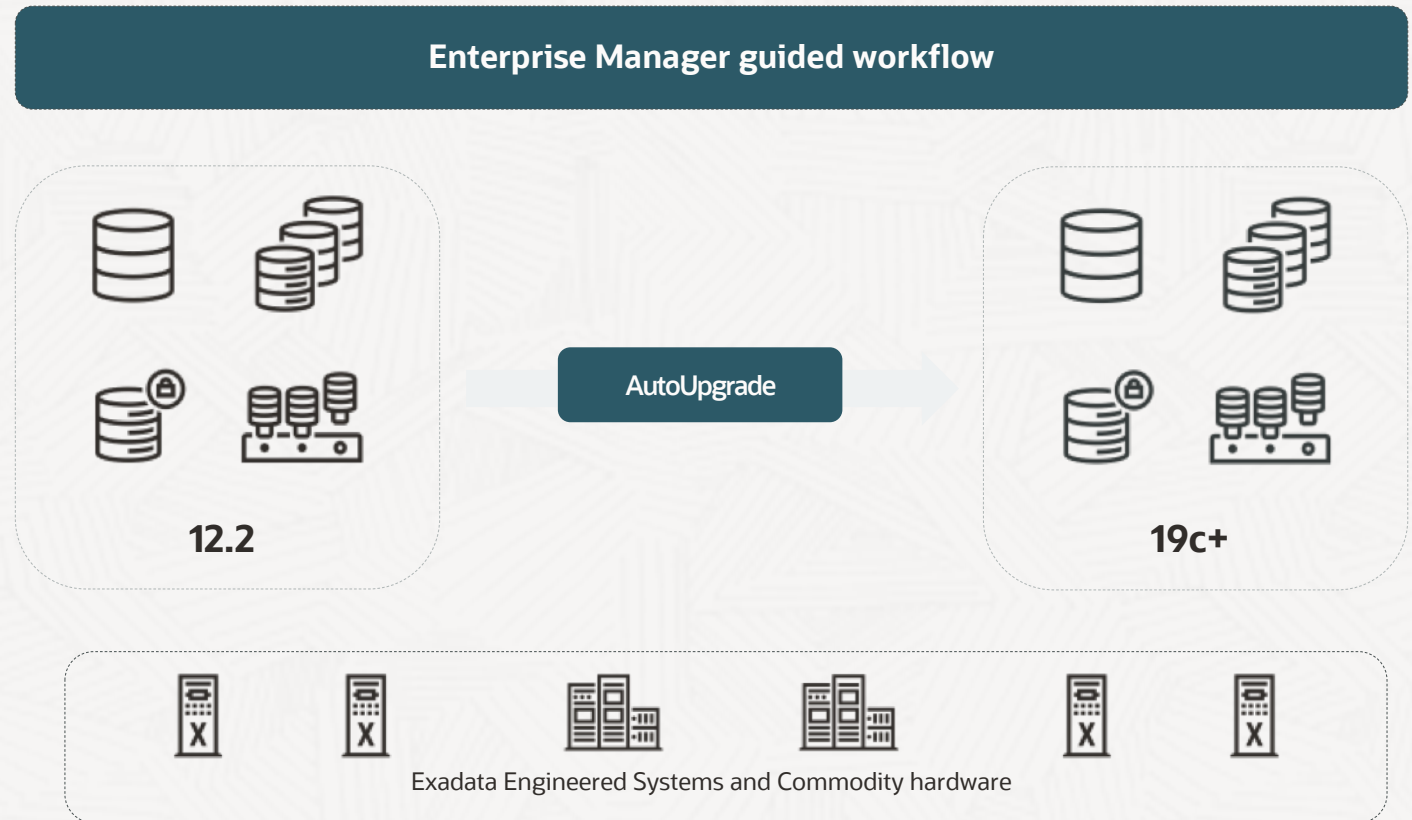
Soon, AutoUpgrade is **the only supported tool to upgrade** your Oracle Database

- DBUA and command line upgrade will get **desupported**

Enterprise Manager

AutoUpgrade integration into EM CC

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



Oracle Database 23c



What's New



Examine Oracle Database changes using ORAdiff

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



ORAdiff

ORAdiff - Find the differences between two Oracle Database releases

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

Get Started

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



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

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

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

ORAdiff

ORAdiff - Find the differences between two Oracle Database releases

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

Get Started

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

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

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

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

[Watch on YouTube](#)






Oracle Autonomous Database speaks *human*

 ORACLE AUTONOMOUS DATABASE

Autonomous Database Speaks “Human”





Get the best of both worlds
with JSON Duality Views

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

COURSE				
CID	CLASS	ROOM	TIME	TCHID
C02	MATH 201	A102	14:00	T543
C45	SCIENCE 102	B405	16:00	T789
...
...

STUDENT COURSES	
STUID	CID
S3245	C02
...	...
S3245	C45
...	...

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



```

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

```

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

Oracle Database
secures your data



Data Guard

RAC

RMAN

TDE Tablespace Encryption

Network Encryption


Auditing

SQL Firewall

ACID


No data duplication

... and so much more



Easily convert your existing application with Oracle Database API for MongoDB

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



Quickly resolve poor performance caused by change in execution plan

- Simplified use of SQL Plan Management

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

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




Allow the database to fix regressing plans automatically

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



Automatically recover from certain **ORA-00600** errors

- Automatic Error Mitigation



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

```
ERROR at line 1:
```

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

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

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

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



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.
```



Even wider tables

- Up to 4096 columns

```
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 23c

—
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 ...`



By default even **stronger** encryption algorithms are used

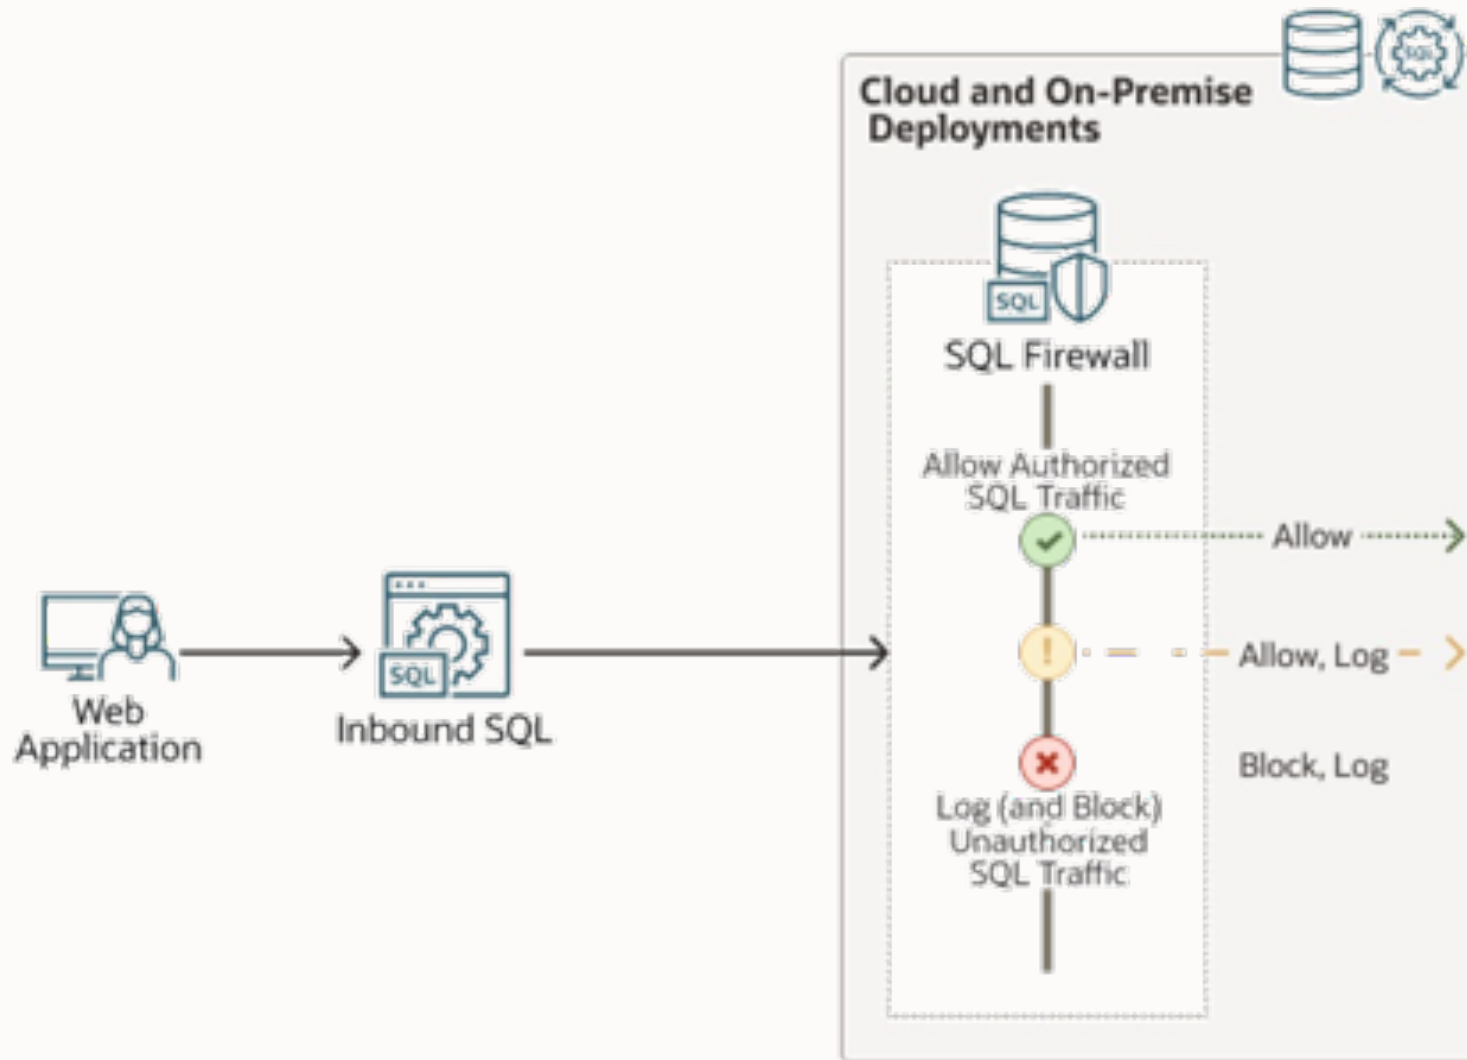
- RMAN backups
- TDE Tablespace Encryption



Tighten security with SQL Firewall

- Prevent SQL injection attacks and unauthorized queries

SQL Firewall



SQL Firewall



Oracle Database 23c



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 dropped
```

```
--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/23c/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/23c/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/23c/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.
```

```
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

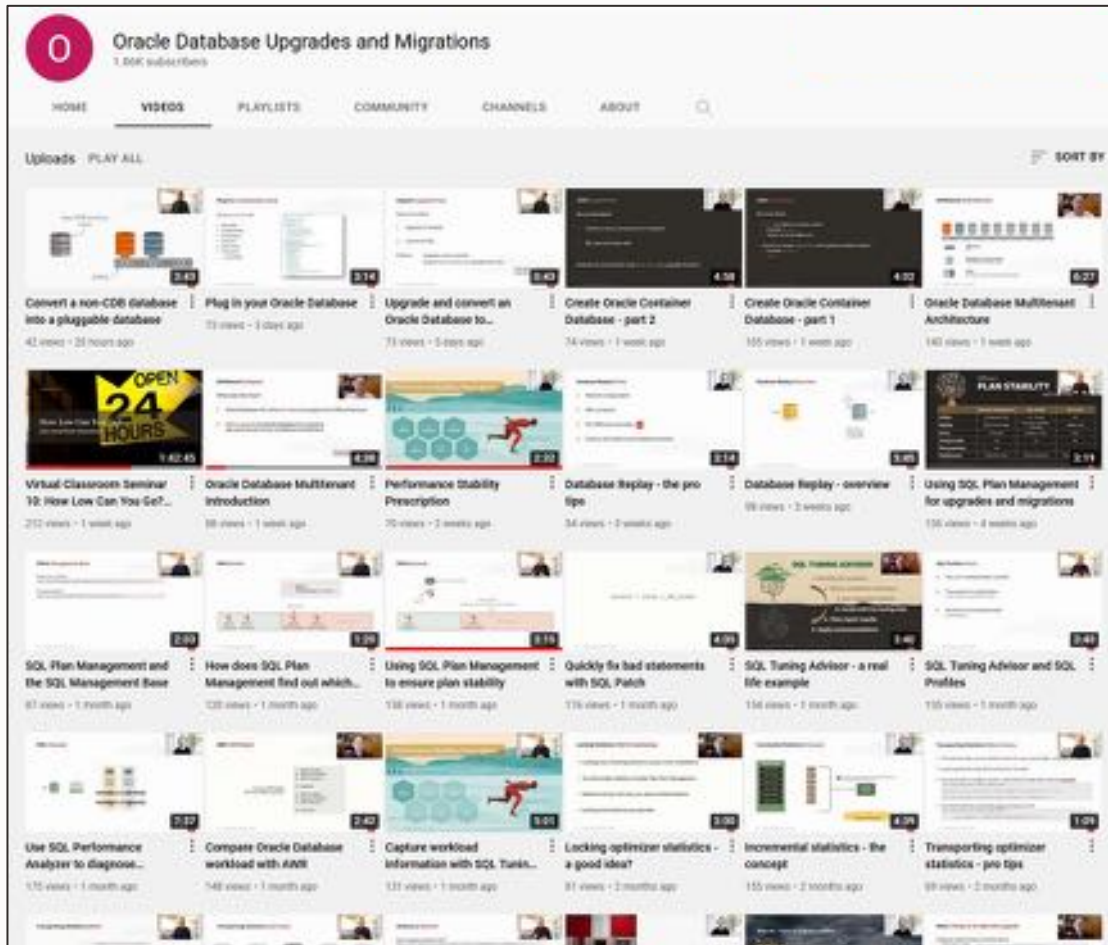
the most
important change

--No longer need to select from dual

~~select sysdate from dual;~~

select sysdate;

YouTube | Oracle Database Upgrades and Migrations

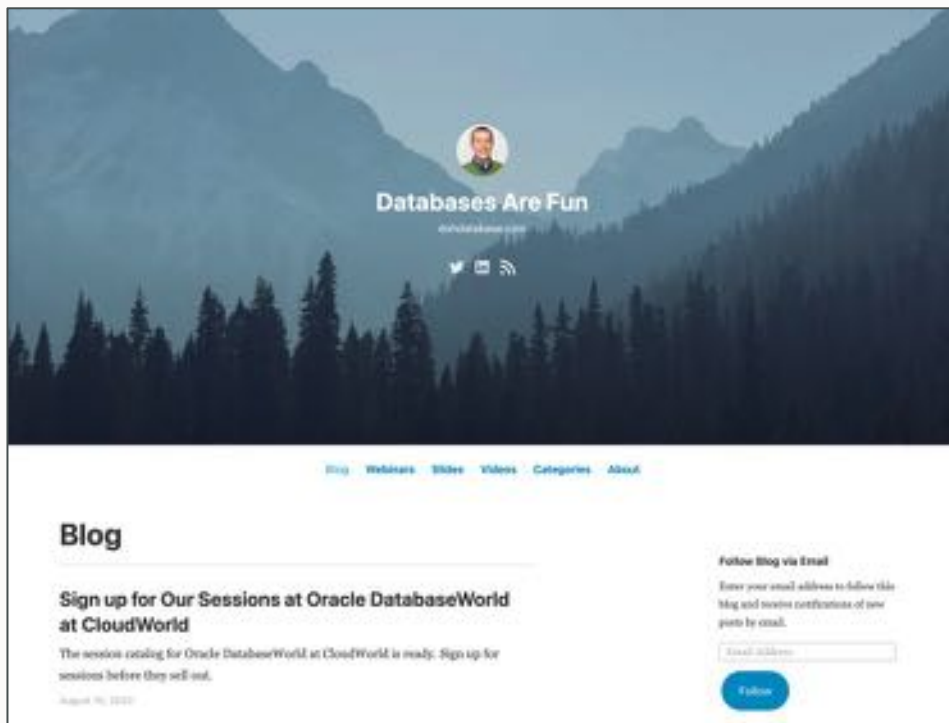


- 300+ videos
- New videos every week
- No marketing
- No buzzword
- All tech

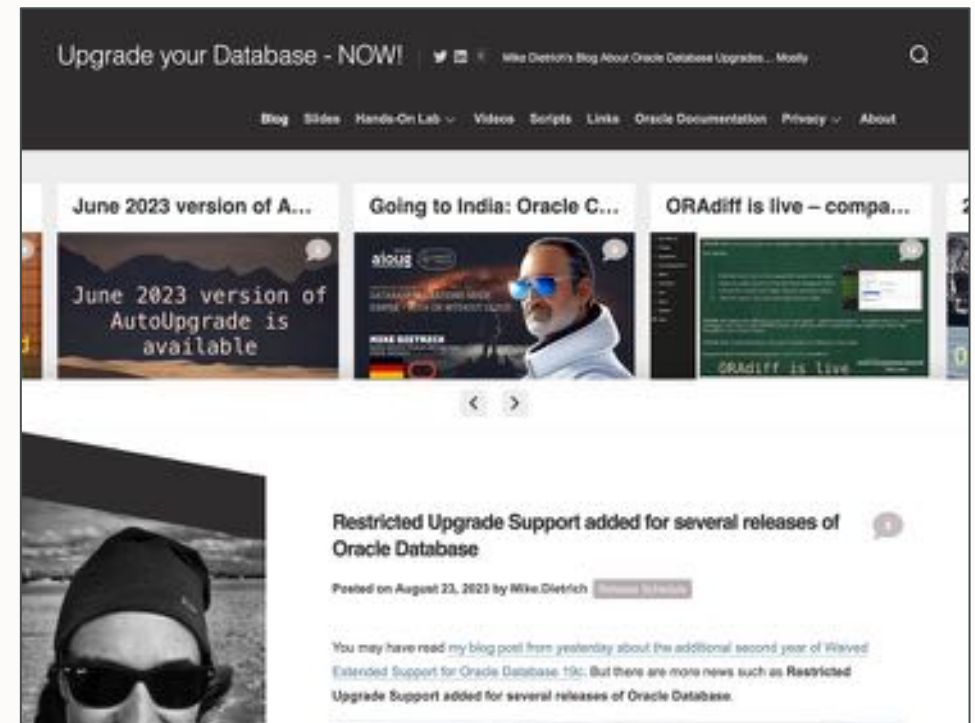


Find slides and much more on our blogs

dohdatabase.com



MikeDietrichDE.com



Thank You

