Mastering Oracle AI Database Upgrades:

Best Practices, Migration Strategies and Minimum Downtime Techniques

DOAG 2025 Konferenz + Ausstellung, November 2025

Oracle

DBAs

run the world

Mike Dietrich

Vice President

mikedietrich

@mikedietrichde.com

https://mikedietrichde.com

Daniel Overby Hansen

Distinguished Product Manager

dohdatabase

@dohdatabase.com

https://dohdatabase.com

Find Slides and Much More on Our Blogs

MikeDietrichDE.com

Mike.Dietrich@oracle.com

dohdatabase.com

Daniel.Overby.Hansen@oracle.com

DBArj.com.br

Rodrigo.R.Jorge@oracle.com

AlexZaballa.com

Alex.Zaballa@oracle.com

Virtual Classroom Seminars	Slides
Episode 16 (replaces Episode 1 from Feb 2021) Oracle Database Release and Patching Strategy for 19c and 23c *115 minutes – May 10, 2023*	
Episode 17 From SR to Patch – Insights into the Oracle Database Development process *55 minutes – June 22, 2023*	
Episode 18 Cross Platform Migration – Transportable Tablespaces to the Extreme *145 min – February 22, 2024*	
Episode 19 Move to Oracle Database 23ai – Everything you need to know about Multitenant PART 1 *145 min – May 16, 2024*	
Episode 20 Move to Oracle Database 23ai – Everything you need to know about Multitenant PART 2 *100 min – June 28, 2024*	

Recorded Web Seminars

https://MikeDietrichDE.com/videos

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

PART 1
Release strategy
Upgrade
Performance

PART 2
Multitenant
Migration
Data Guard

BREAK

PART 3
Minimal downtime
Data Pump

LUNCH

PART 4
Autonomous AI Database
Migrations
Life of a DBA

BREAK

Release Strategy

Lifetime Support Policy

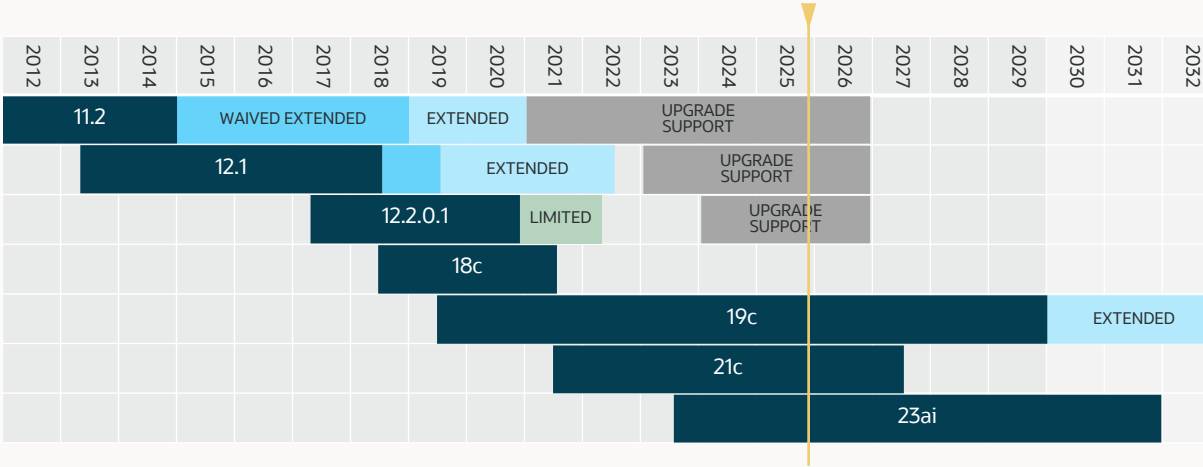

Lifetime Support Policy

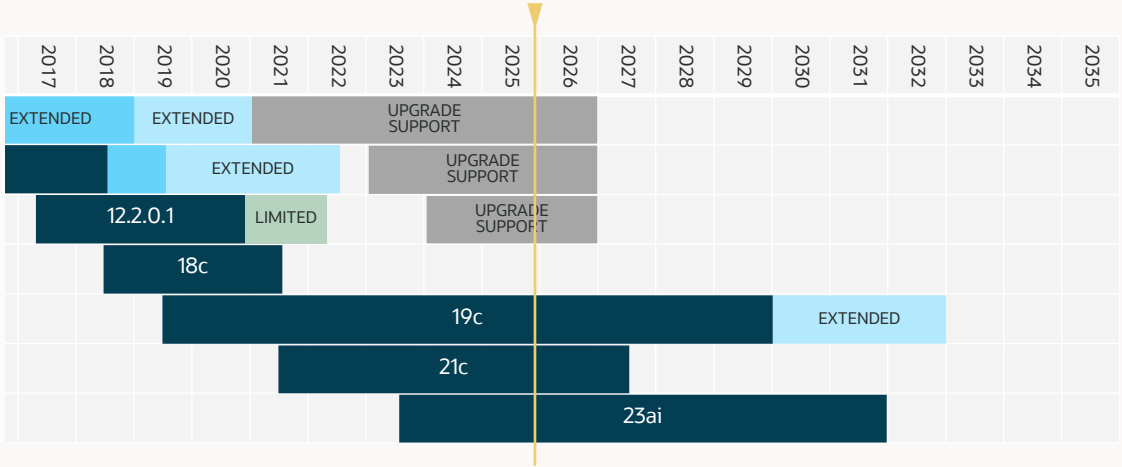

Lifetime Support Policy

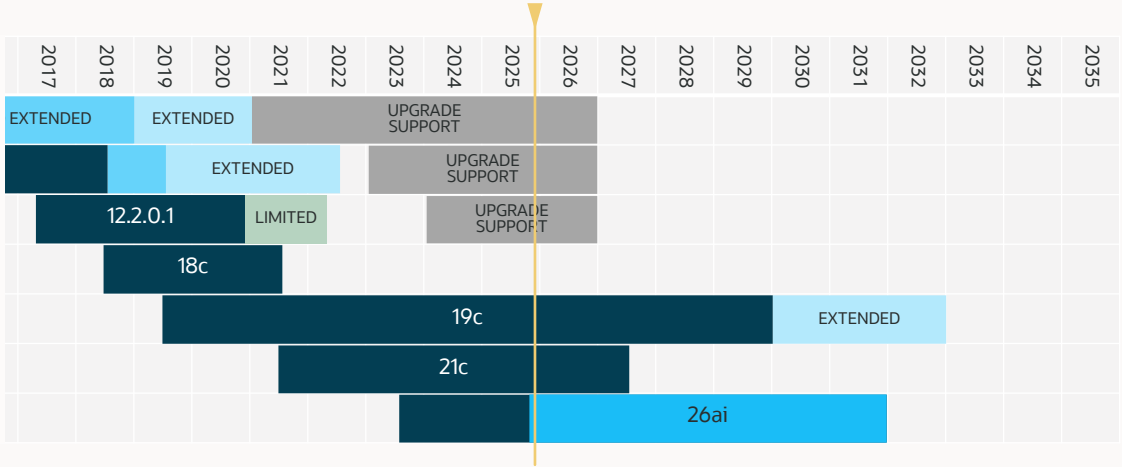

Do you want to upgrade?

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

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

Oracle Database 19c

Oracle Database 23ai

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

Oracle Database 19c

Oracle Database 23ai

Oracle AI Database 26ai

26ai

When is a database upgrade required?

Oracle Database 19c ⇒ Oracle Database 23ai ⇒ Oracle AI Database 26ai

UPGRADE

Oracle Database 19c

Oracle AI Database 26ai

UPGRADE

Oracle Database 19c ⇛ Oracle Database 23ai ⇛ Oracle AI Database 26ai

UPDATE

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

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

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

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

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

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

```
SQL> select version, version_full from v$instance

VERSION           VERSION_FULL
----------------- -----------------
23.0.0.0.0        23.26.0.0.0
```

```
SQL> select version, version_full from v$instance

VERSION           VERSION_FULL
----------------- -----------------
23.0.0.0.0        23.26.0.0.0
```

Oracle AI Database 26ai supports the multitenant architecture only

- You must convert your database to a PDB

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

Single vs. Multitenant

Single Tenant

One PDB
No extra license

Multitenant

Multiple PDBs
Extra license if more than 3 PDBs

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

alter system set max_pdbs=3;
```

Database Upgrade

Replay Upgrade

1 Plug in

2 Upgrade

3 Convert

Irreversible!
Flashback no good

```
--The database automatically starts an upgrade
--when you plug in a lower-release PDB

SQL> alter pluggable database pdb1 open;

Pluggable database altered.

Elapsed: 00:06:01.95
```

```
SQL> select property_name, property_value
     from   database_properties
     where  property_name like '%OPEN%';

PROPERTY_NAME          PROPERTY_VALUE
---------------------- --------------------
CONVERT_NONCDB_ON_OPEN true
UPGRADE_PDB_ON_OPEN    true
```

```
SQL> select property_name, property_value
     from   database_properties
     where  property_name like '%OPEN%';

PROPERTY_NAME           PROPERTY_VALUE
----------------------- --------------------
CONVERT_NONCDB_ON_OPEN  true
UPGRADE_PDB_ON_OPEN     true
```

Classic Upgrade

Phase 1
Phase 2
Phase 3
Phase 4
Phase 5
Phase 6
Phase 7
Phase 8
...
Phase *nnn*

Classic Upgrade

Phase 1
Phase 2
Phase 3
Phase 4
Phase 5

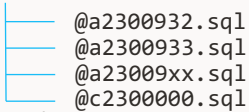

Phase 6
Phase 7
Phase 8
...
Phase *nnn*

Classic Upgrade

@a2300932.sql

```
VARIABLE initfile VARCHAR2(32)
COLUMN :initfile NEW_VALUE init_file NOPRINT;

Rem ================================================================
Rem SQLJTYPE
Rem ================================================================

BEGIN
  IF sys.dbms_registry.is_loaded('JAVAVM',sys.dbms_registry.release_version) = 1 THEN
     :initfile := 'initsjty.sql';
  ELSE
     :initfile := 'nothing.sql';
  END IF;
END;
/
SELECT :initfile FROM DUAL;
@@&init_file
```

Classic Upgrade

Comparison

Classic

Phase 1
Phase 2
Phase 3
Phase 4
Phase 5
Phase 6
Phase 7
Phase 8
...
Phase *nnn*

Replay

```
DROP INDEX SYSTEM.IDX$FLOW ...
CREATE OR REPLACE ...
ALTER TYPE ...
CREATE FUNCTION ...
CREATE TABLE SYS.T1 ...
CREATE INDEX SYS.T1I1 ...
DROP INDEX MDSYS.IDX$IK ...
DROP TABLE MDSYS.TBL$TT ...
CREATE OR REPLACE ...
ALTER TYPE ...
GRANT SELECT ON ...
CREATE VIEW ...
```

```
select sqlstmt from pdb_sync$;
```

```
ALTER SESSION SET "_oracle_script_counter"=7
alter pluggable database application app$cdb$pdbonly$ncdbtopdb begin install '1.0.upgmode'
alter session set "_enable_view_pdb"=false
alter session set NLS_LENGTH_SEMANTICS=BYTE
INSERT INTO sys.utl_recomp_skip_list select obj# from obj$ where BITAND(flags, 4194304)=0 ...
create or replace view sys.cdb$common_root_objects sharing=object as
select u.name owner, o.name object_name, o.type# object_type, o.namespace nsp,
       o.subname object_subname, o.signature object_sig,
       decode(bitand(o.flags, (65536+131072+4294967296)),
       4294967296+65536, 'EDL', 131072, 'DL', 'MDL') sharing
  from sys.obj$ o, sys.user$ u
 where o.owner#=u.user# and bitand(o.flags, (65536+131072+4294967296)) <> 0
   and bitand(o.flags,0)=0
```

(output truncated)

Classic

- Triggered by AutoUpgrade
- Runs `catalog.sql` / `catproc.sql`
- Many `CREATE OR REPLACE` statements for objects that didn't change
- Customizable
- Used by AutoUpgrade

Replay

- Triggered by `OPEN` command
- Runs the captured statements
- Only statements that actually do some change
- Automated

Comparison

Classic

```
Stages
   SETUP              <1 min
   PREUPGRADE         <1 min
   PRECHECKS          <1 min
   PREFIXUPS          <1 min
   DRAIN              <1 min
   DBUPGRADE          19 min
   DISPATCH           <1 min
   UNPLUGWORK         <1 min
   POSTCHECKS         <1 min
   POSTFIXUPS         10 min
   POSTUPGRADE        <1 min
   SYSUPDATES         <1 min
```

Replay

```
Stages
   SETUP              <1 min
   PREUPGRADE         <1 min
   PRECHECKS          <1 min
   PREFIXUPS          <1 min
   DRAIN              <1 min
   DBUPGRADE          17 min
   DISPATCH           <1 min
   UNPLUGWORK         <1 min
   POSTCHECKS         <1 min
   POSTFIXUPS         10 min
   POSTUPGRADE        <1 min
   SYSUPDATES         <1 min
```

Replay upgrade doesn't handle pre- and post-upgrade tasks

- You must run these manually
- Or use Replay Upgrade through AutoUpgrade

```
SQL> alter pluggable database pdb1 open;
alter pluggable database pdb1 open
*
ERROR at line 1:
ORA-60510: encountered an error during Replay Upgrade
```

If Replay Upgrade fails

- Check for errors:
  - `SELECT * FROM dba_replay_upgrade_errors`
  - `SELECT * FROM dba_app_errors`
  - `SELECT * FROM dba_applications WHERE app_name='APP$CDB$CATALOG';`
  - Check alert log
  - Trace files

- Revert to classic upgrade
  - Use AutoUpgrade (`upg1.replay=no`)

```
--To disable replay upgrade
ALTER DATABASE UPGRADE SYNC OFF;

--Or
ALTER DATABASE PROPERTY SET UPGRADE_PDB_ON_OPEN='false';

--To disable convert on open
ALTER DATABASE PROPERTY SET CONVERT_NONCDB_ON_OPEN='false';
```

Performance Stability Prescription


what's your
biggest fear
when making changes?
Photo by nikko macaspac on Unsplash

Photo by Dave on Unsplash

Performance Stability Prescription

Photo by Urip Dunker on Unsplash

Performance Stability Prescription

Collect workload information

- Sample from cursor cache
- Gather from AWR

SQL Tuning Set | Definition

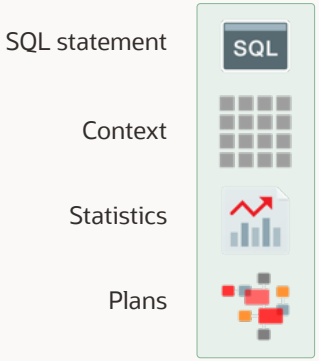

Gather at least a full month of workload data

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

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

SQL Performance Analyzer | Concept

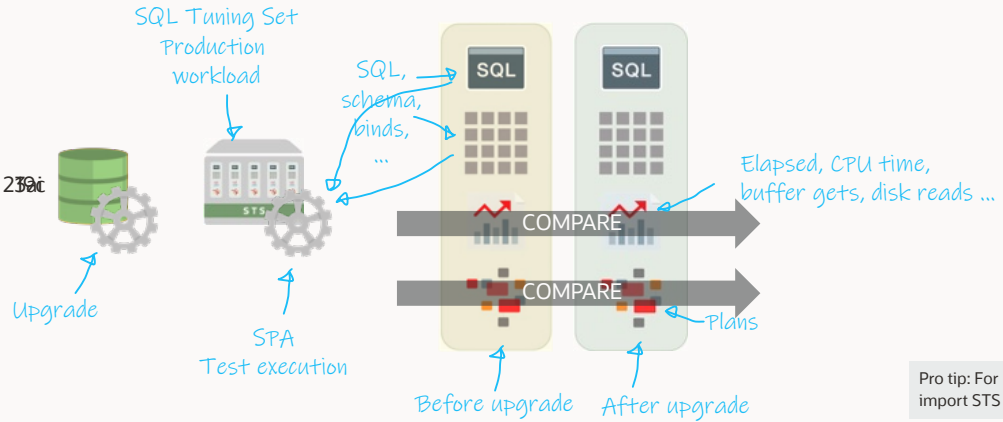

Pro tip: For migrations, import STS into target database

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		
⇩	3fv28nfu9v0aq	-0.050	26,504	29,573	-11.580	Y
⇩	czzzubf8fjz96	-0.030	1,410	1,981	-40.500	Y

SQL Details: czzzubf8fjz96

Parsing Schema APPS Execution Frequency 3

▷ SQL Text

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

Single Execution Statistics

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

SQL Performance Analyzer | Report

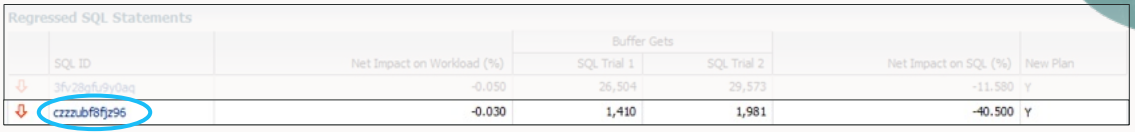

Regressed SQL Statements

	SQL ID	Net Impact on Workload (%)	Buffer Gets		Net Impact on SQL (%)	New Plan
			SQL Trial 1	SQL Trial 2		
⇩	3fv28pfu9v0aq	-0.050	26,504	29,573	-11.580	Y
⇩	czzzubf8fjz96	-0.030	1,410	1,981	-40.500	Y

Plan Comparison

SQL_TRIAL_1353942463446

Plan Hash Value 1165613724

Expand All | Collapse All

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

SPA | Continuous Improvement

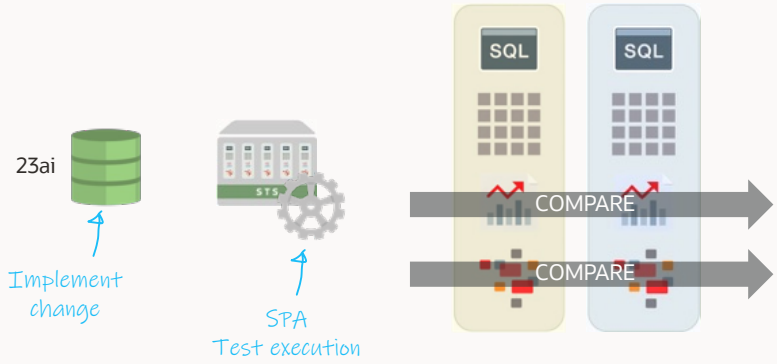

Performance Stability Prescription

Tune SQLs with regressed plans

Create SQL Plan Baselines

Transport to production database

SPM | Use Case

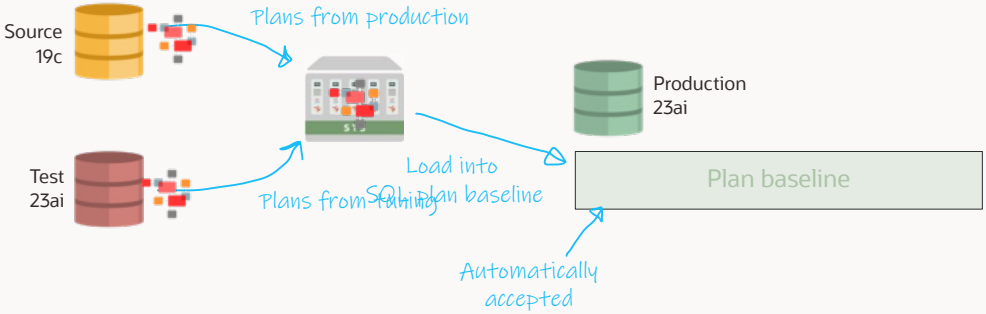

There is only one tool to ensure plan stability:

SQL Plan Management

Don't use

- `OPTIMIZER_FEATURES_ENABLE`
- `COMPATIBLE`

COMPATIBLE vs. OPTIMIZER_FEATURES_ENABLE

COMPATIBLE

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

OPTIMIZER_FEATURES_ENABLE

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

```
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
```


Speaking of changes…
Do you know ORADiff?

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

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

ORAdiff

ORAdiff - Find the differences between two Oracle Database releases

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

Get Started

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

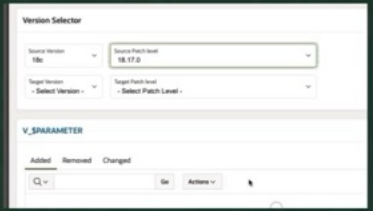

ORAdiff will display the differences such as "new tables", "added parameters", "changed columns", "removed privileges" and much more. ORAdiff search can tell you when a parameter was added and which files changed in your Oracle Home.

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

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

Performance Tips & Tricks

Use as few initialization parameters as possible

- Stick to the defaults
- Stick to vendor recommendations

Only use underscores and events to solve specific situations

- Only under guidance of Oracle Support

Patches For Optimal Performance

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

Patches For Optimal Performance

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`

Home
Parameters
Users, Roles, etc
Privileges
Included Fixes
List of Fixes
Total Fixes
Search Fixes
List of CVEs
CPU Explore
Fix Control
Oracle Database Ho...
Objects
Fixed Objects
Audit
System

Version Selector

Source Release: 19c
Source Patch Level: 19.3.0 (Base Release)
Target Release: 19c
Target Patch Level: 19.28.0
List Mode

V_$SYSTEM_FIX_CONTROL

Added | Removed | Modified

Go | Actions

BUGNO	VALUE	SQL_FEATURE	DESCRIPTION	EVENT	IS_DEFAULT
9876287	0	QKSFM_PQ_9876287	Fix perf slowdown due to wait event PXNSQ:PQLOADINFOQUERY	0	1
10123661	0	QKSFM_CURSOR_SHARING_10123661	Enable cursor sharing for AS OF queries	0	1
17295505	0	QKSFM_OBJ_EXT_17295505	allow stateful-access to image operator for ODCI table function	0	1
18101156	0	QKSFM_PARTITION_18101156	no RowCR for global index access on partitioned table	0	1
19138896	0	QKSFM_DML_19138896	allow update of join view without key preserved property	0	1
20922160	0	QKSFM_CBO_20922160	correct null selectivity adjustment in presence of NVL	0	1
22387320	0	QKSFM_EXECUTION_22387320	use HASH UNIQUE with UNION operator when applicable	0	1

BUGNO	VALUE	SQL_FEATURE	DESCRIPTION	EVENT	IS_DEFAULT
29331066	0	QKSFM_VECTOR_AGG_29331066	allow VT approximate count distinct detail support	0	1
29385774	0	QKSFM_OR_EXPAND_29385774	Allow LORE when cbqt ore rejected due to sub qb is parametrised	0	1
29435966	0	QKSFM_ALL_29435966	Allow 32K Varchar to be returned from a pipelined table function	0	1
29450812	0	QKSFM_OR_EXPAND_29450812	Allow legacy ORE for exotic query constructs	0	1
29463553	0	QKSFM_SVM_29463553	allow SPJ with outer join, CONNECT BY and one remote table	0	1
29487407	0	QKSFM_SQL_CODE_GENERATOR_29487407	Streamline traversal of CAST_PSR operator	0	1
29499077	0	QKSFM_CBO_29499077	correct sel for LIKE preds with char binds holding numeric data	0	1
29590666	0	QKSFM_ALL_29590666	Do not perform xmlexists(fn:not) rewrite	0	1
29651517	0	QKSFM_COMPILATION_29651517	do not mark load into CDT as multi-start with result cache	0	1
29653132	0	QKSFM_ALL_29653132	Force the usage of BLOB in XMLAGG and XMLELEMENT	0	1
29657973	0	QKSFM_ACCESS_PATH_29657973	for update: force VC base column into row vector for top qbc onl	0	1
29687220	0	QKSFM_CBO_29687220	improve costing for indexes with empty statistics.	0	1
29696242	0	QKSFM_ACCESS_PATH_29696242	use only the best auto index for access path analysis	0	1
29712727	0	QKSFM_DYNAMIC_SAMPLING_29712727	analysis of skipped non-empty table partitions	0	1

1 - 50 of 288

Patches For Optimal Performance

Upgrade New database	Enable optimizer fixes using `DBMS_OPTIM_BUNDLE`
Patching	Do proper testing before enabling optimizer fixes using `DBMS_OPTIM_BUNDLE`

The Lab Environment

Requirements

Laptop

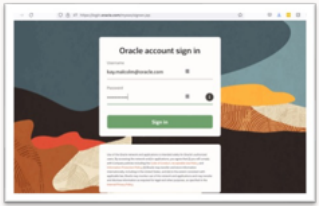

Oracle Account

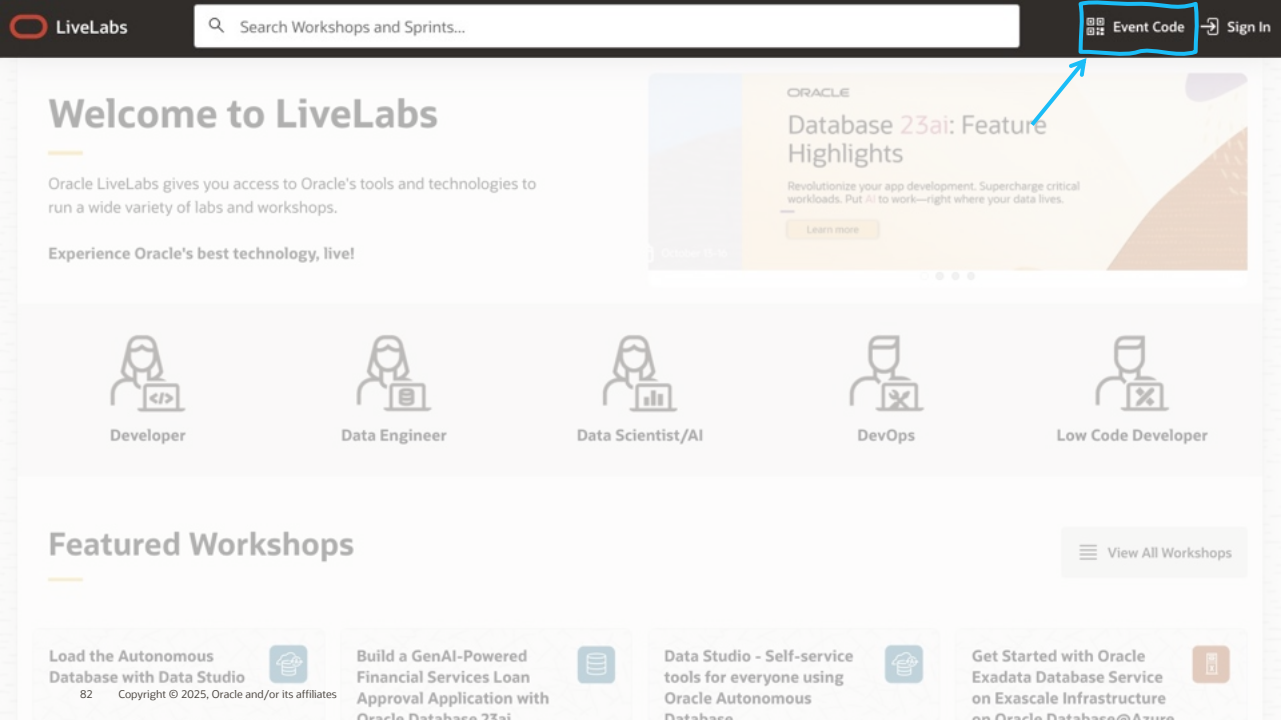

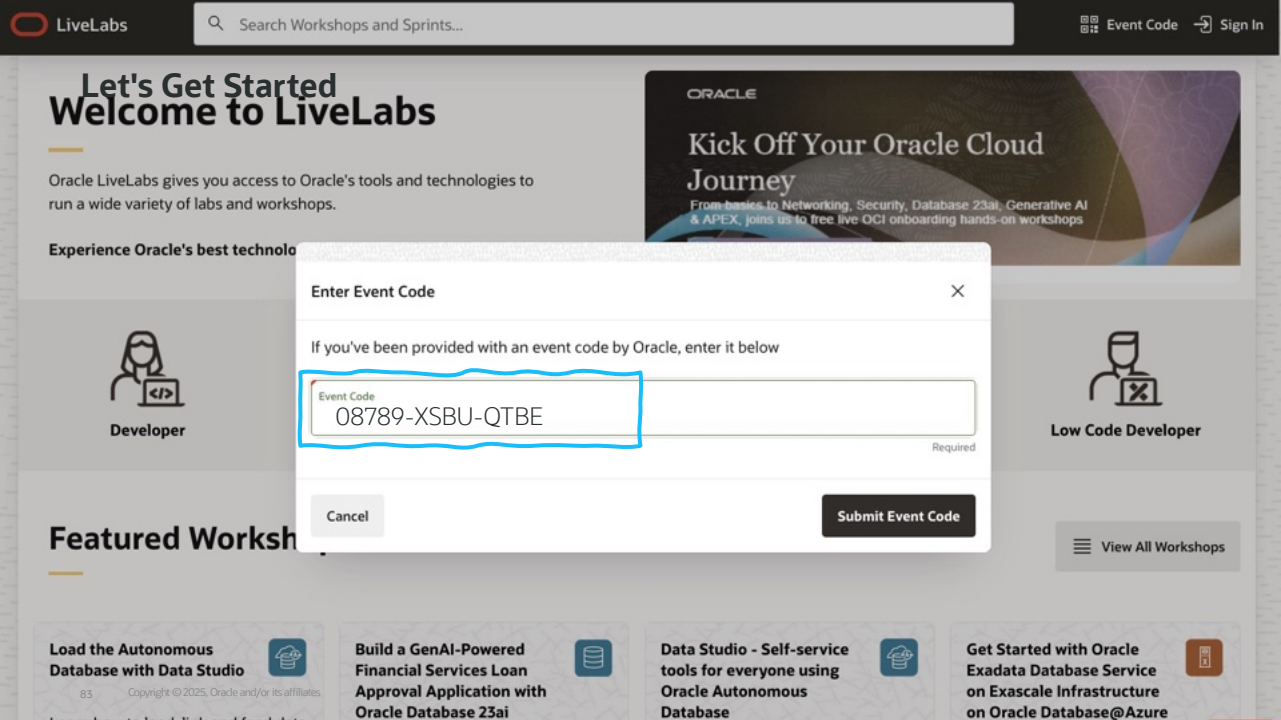
LiveLabs
Search Workshops and Sprints...
Event Code
Sign In
Let's Get Started
Welcome to LiveLabs
Oracle LiveLabs gives you access to Oracle's tools and technologies to run a wide variety of labs and workshops.
Experience Oracle's best technolo
ORACLE
Kick Off Your Oracle Cloud Journey
From basics to Networking, Security, Database 23ai, Generative AI & APEX, joins us to free live OCI onboarding hands-on workshops
Enter Event Code
If you've been provided with an event code by Oracle, enter it below
Event Code
08789-XSBU-QTBE
Required
Cancel
Submit Event Code
Developer
Low Code Developer
Featured Worksh
View All Workshops
Load the Autonomous Database with Data Studio
Build a GenAI-Powered Financial Services Loan Approval Application with Oracle Database 23ai
Data Studio - Self-service tools for everyone using Oracle Autonomous Database
Get Started with Oracle Exadata Database Service on Exascale Infrastructure on Oracle Database@Azure
83
Copyright © 2025, Oracle and/or its affiliates

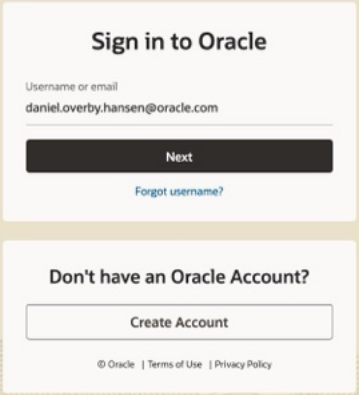
Sign in to Oracle
Username or email
daniel.overby.hansen@oracle.com
Next
Forgot username?
Don't have an Oracle Account?
Create Account
© Oracle | Terms of Use | Privacy Policy

Use your Oracle Account to log in

- This is not your Oracle Cloud Tenancy account
- It is your Oracle SSO account

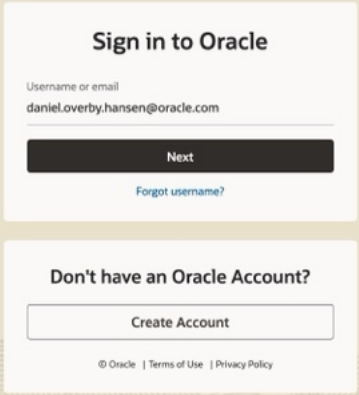

Oracle AI World 2025 - Fast Track: Upgrade to Oracle Database 23ai [PI1147]

Share
Start

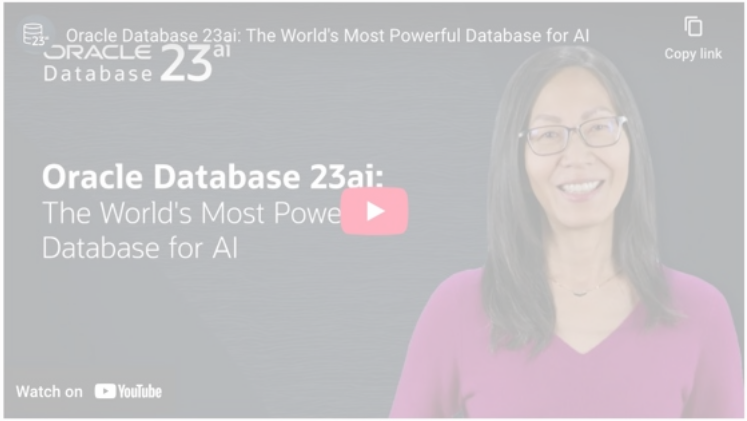

2 hours

Organizer:
Oracle

Event Date:
Oct 13 - 16

Outline

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

Hitchhiker's Guide for Upgrading to Oracle Database 23ai

Oracle AI World 2025 - Fast Track: Upgrade to Oracle Database 23ai [PI1147]

Share
Start

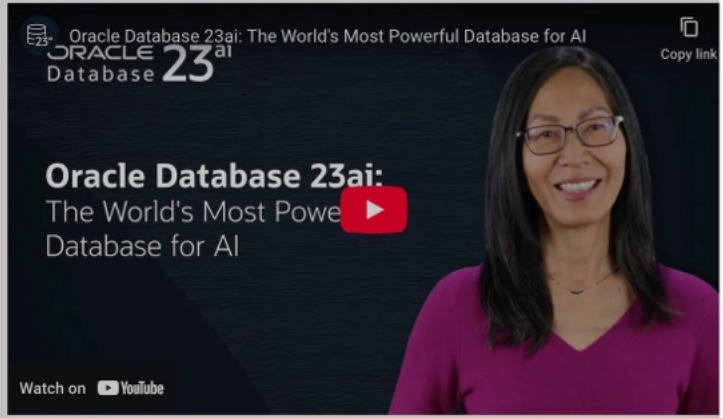

The **Run on Your Environment** button provides step-by-step instructions so you can run this workshop using your own resources!

Run on Your Environment

The **Run on LiveLabs** button will dynamically create resources in an Oracle-owned tenancy for you to use for free!
Oracle account help | ***Oracle account signup***

Run on LiveLabs Sandbox

The **Preview Sandbox Instructions** button will open a link with the Sandbox instructions for you to preview before creating a reservation.

Preview Sandbox Instructions

Hitchhiker's Guide for Upgrading to Oracle Database 23ai

- Restore failed upgrade

LiveLabs

Search for workshops and sprints...

Oracle AI World 2025 - Fast Track: Upgrad
[PI1147]

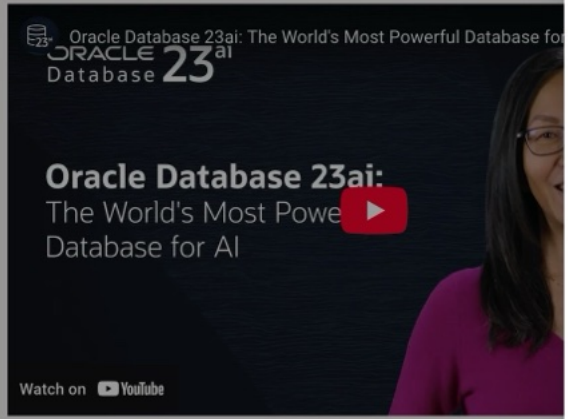

Hitchhiker's Guide for Upgrading to Oracle Data

Reserve Workshop

Attendee Email Address
daniel.overby.hansen@oracle.com

Attendee Timezone
PST (-07:00)
Required

* Start Workshop Now?

I consent to receive emails from LiveLabs for my reservation and I agree that I will not upload sensitive personal or company information to Oracle Cloud Infrastructure *

Preview Sandbox Instructions

Submit Reservation

My Reservations

All your current workshop reservations are shown below. You can edit active or pending reservations, view workshop details, attend an available workshop, or delete a reservation.

To access this page again click the user dropdown in the top right corner and select **My Reservations**

Note: The status of your reservations will be emailed to you. Check your mail for any status updates.

Hitchhiker's Guide for Upgrading to Oracle Database 23ai

Wednesday October 8th, 12:59am (00:59) PST

Pending creation

Refresh the page a few times
until a lab environment is assigned

My Reservations

All your current workshop reservations are shown below. You can edit active or pending reservations, view workshop details, attend an available workshop, or delete a reservation.

To access this page again click the user dropdown in the top right corner and select **My Reservations**

Note: The status of your reservations will be emailed to you. Check your mail for any status updates.

Hitchhiker's Guide for Upgrading to Oracle Database 23ai

Wednesday October 8th, 12:59am (00:59) PST

Launch Workshop

Details Delete

Introduction

About this Workshop

Oracle Database 23ai is a *Long Term Support Release*. This lab combines several upgrade and migration methods, and equips you with performance features and tools to ensure stability when you move to any new Oracle Database release. By upgrading to Oracle Database 23ai, you will have Premier Support until the end of 2031 and Extended Support for a period thereafter. There is a direct upgrade path to Oracle Database 23ai from Oracle Database 19c and 21c, regardless of the Release Update applied.

Estimated Workshop Time: 120 minutes

Objectives

In this workshop, you will:

- Upgrade databases
- Use Performance Stability Prescription to ensure performance stability
- Convert to multitenant architecture
- Migrate databases using Data Pump and transportable tablespaces

Expand All Tasks

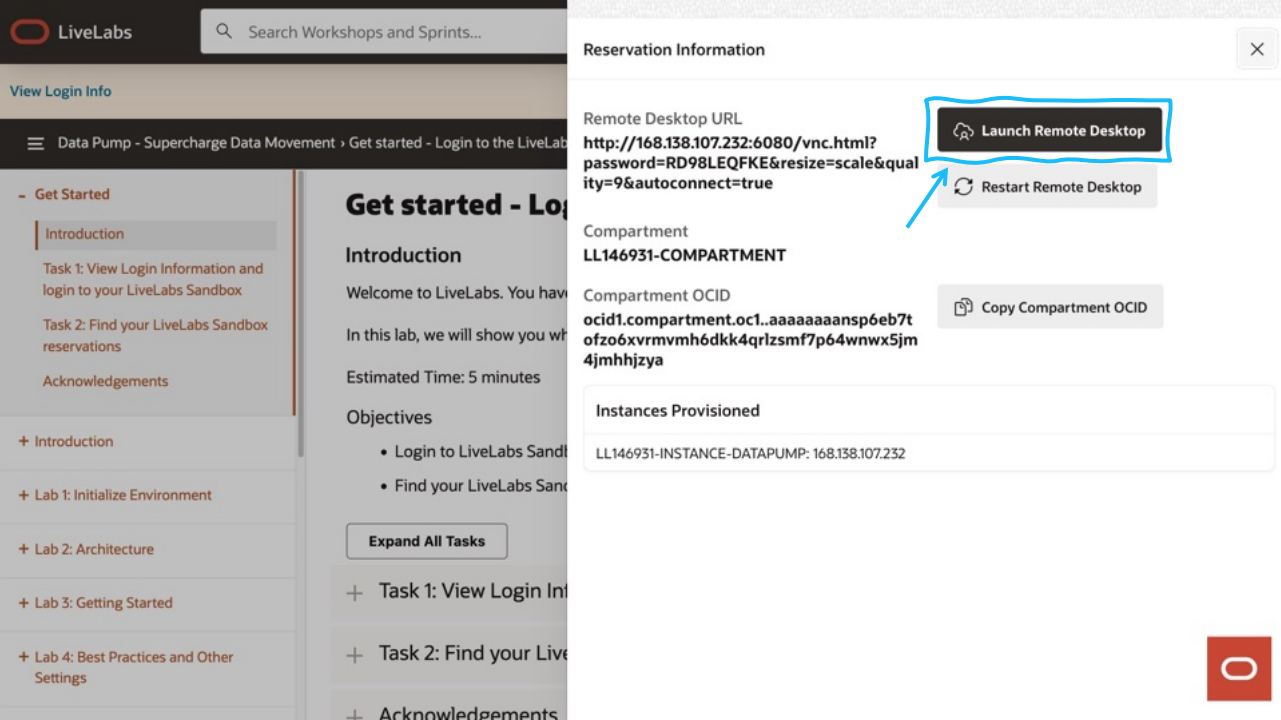
LiveLabs
Search Workshops and Sprints...
View Login Info
Data Pump - Supercharge Data Movement › Get started - Login to the LiveLab
Get Started
Introduction
Task 1: View Login Information and login to your LiveLabs Sandbox
Task 2: Find your LiveLabs Sandbox reservations
Acknowledgements
Introduction
Lab 1: Initialize Environment
Lab 2: Architecture
Lab 3: Getting Started
Lab 4: Best Practices and Other Settings
Get started - Log
Introduction
Welcome to LiveLabs. You have
In this lab, we will show you wh
Estimated Time: 5 minutes
Objectives
Login to LiveLabs Sandb
Find your LiveLabs Sand
Expand All Tasks
Task 1: View Login Inf
Task 2: Find your Live
Acknowledgements
Reservation Information
Remote Desktop URL
http://168.138.107.232:6080/vnc.html?password=RD98LEQFKE&resize=scale&quality=9&autoconnect=true
Launch Remote Desktop
Restart Remote Desktop
Compartment
LL146931-COMPARTMENT
Compartment OCID
ocid1.compartment.oc1..aaaaaaaansp6eb7tofzo6xvrmvmh6dkk4qrlzsmf7p64wnwx5jm4jmhhjzya
Copy Compartment OCID
Instances Provisioned
LL146931-INSTANCE-DATAPUMP: 168.138.107.232

HTTPS-Only Mode Alert

Secure Site Not Available

You've enabled HTTPS-Only Mode for enhanced security, and a HTTPS version of **168.138.107.232:6080** is not available.

Learn More...

What could be causing this?

- Most likely, the website simply does not support HTTPS.
- It's also possible that an attacker is involved. If you decide to visit the website, you should not enter any sensitive information like passwords, emails, or credit card details.

If you continue, HTTPS-Only Mode will be turned off temporarily for this site.

Continue to HTTP Site | Go Back

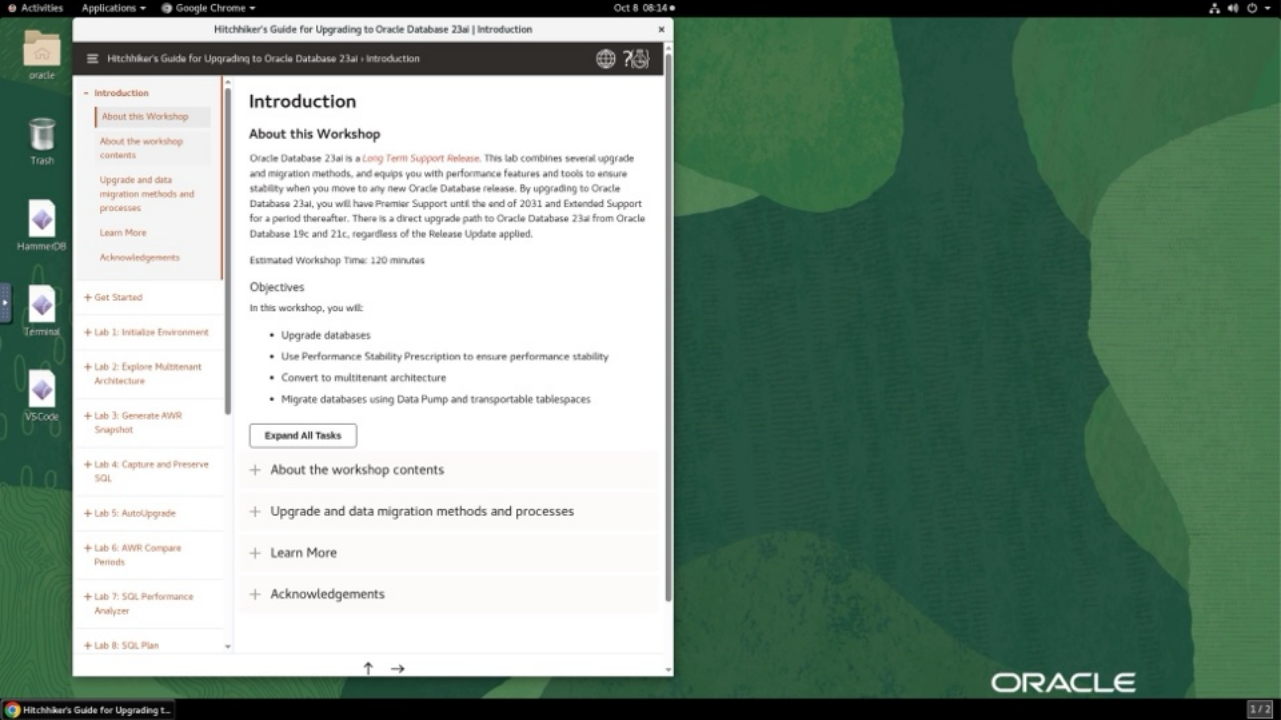
Activities
Applications
Google Chrome
Oct 8 08:14
Hitchhiker's Guide for Upgrading to Oracle Database 23ai | Introduction
Hitchhiker's Guide for Upgrading to Oracle Database 23ai › Introduction
Introduction
About this Workshop
About the workshop contents
Upgrade and data migration methods and processes
Learn More
Acknowledgements
Get Started
Lab 1: Initialize Environment
Lab 2: Explore Multitenant Architecture
Lab 3: Generate AWR Snapshot
Lab 4: Capture and Preserve SQL
Lab 5: AutoUpgrade
Lab 6: AWR Compare Periods
Lab 7: SQL Performance Analyzer
Lab 8: SQL Plan
Introduction
About this Workshop
Oracle Database 23ai is a Long Term Support Release. This lab combines several upgrade and migration methods, and equips you with performance features and tools to ensure stability when you move to any new Oracle Database release. By upgrading to Oracle Database 23ai, you will have Premier Support until the end of 2031 and Extended Support for a period thereafter. There is a direct upgrade path to Oracle Database 23ai from Oracle Database 19c and 21c, regardless of the Release Update applied.
Estimated Workshop Time: 120 minutes
Objectives
In this workshop, you will:
Upgrade databases
Use Performance Stability Prescription to ensure performance stability
Convert to multitenant architecture
Migrate databases using Data Pump and transportable tablespaces
Expand All Tasks
About the workshop contents
Upgrade and data migration methods and processes
Learn More
Acknowledgements
oracle
Trash
HammerDB
Terminal
VSCode
ORACLE
Hitchhiker's Guide for Upgrading t...
1 / 2

You can copy/paste from the instructions

- SHIFT + CTRL + C
- SHIFT + CTRL + V

Introduction

Fast Track: Upgrade to Oracle AI Database 26ai

Overview

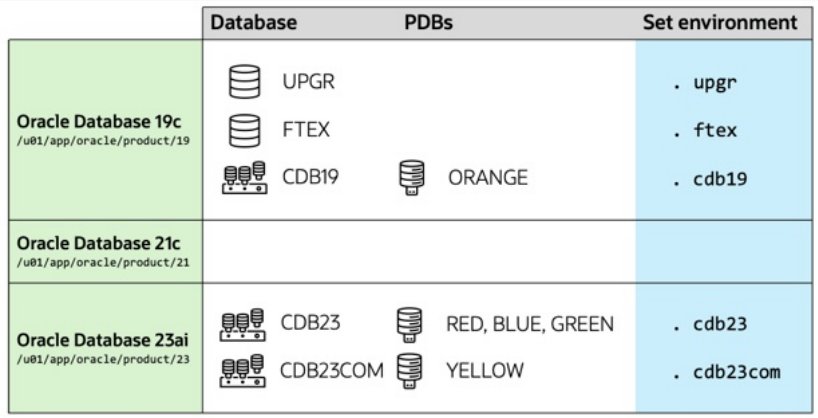

Overview

Labs

1. Initialize the environment

2. Explore multitenant architecture

Overview

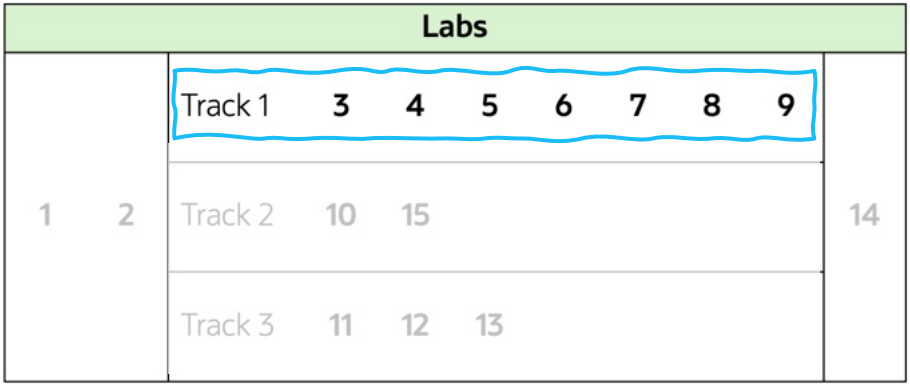

Labs

3. Generate AWR snapshot
   - Run a load generator
   - Capture SQL from cursor cache

4. Capture and preserve SQL
   - Capture SQL from AWR

Labs

Take a break while it upgrades

5. Upgrade to Oracle Database 23ai
6. Compare AWR diff report
   - Run load generator again
7. Analyze with SQL Performance Analyzer
8. Fix with SQL Plan Management
9. Fix with SQL Tuning Advisor

Start Your Engines

Let's Get Started

https://livelabs.oracle.com

08789-XSBU-QTBE

Multitenant Migration

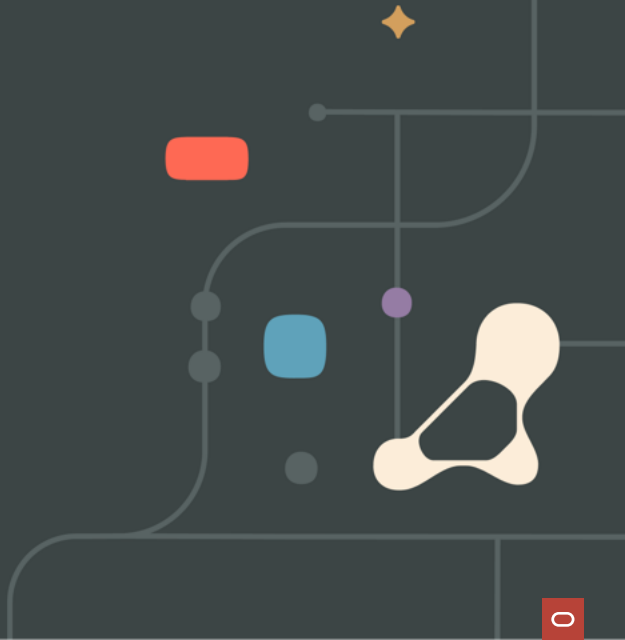

Non-CDB to PDB conversion is irreversible

What are your rollback options?

ROLLBACK

Backup / restore

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

Copy data files

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

ROLLBACK

Backup / restore

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

Copy data files

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

Refreshable clone PDB

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

Requires Oracle Database 12.2 or newer

Upgrade via Refreshable Clone PDB

CREATE

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

REFRESH

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

OUTAGE

Disconnect users and refresh PDB for the last time

UPGRADE

Upgrade PDB to Oracle AI Database 26ai

Upgrade via Refreshable Clone PDB

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

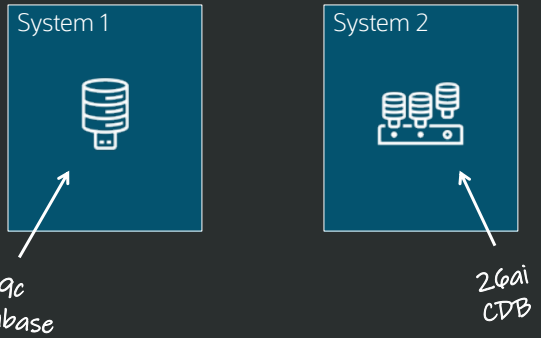

Upgrade via Refreshable Clone PDB

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

Could be same system as well

Upgrade via Refreshable Clone PDB

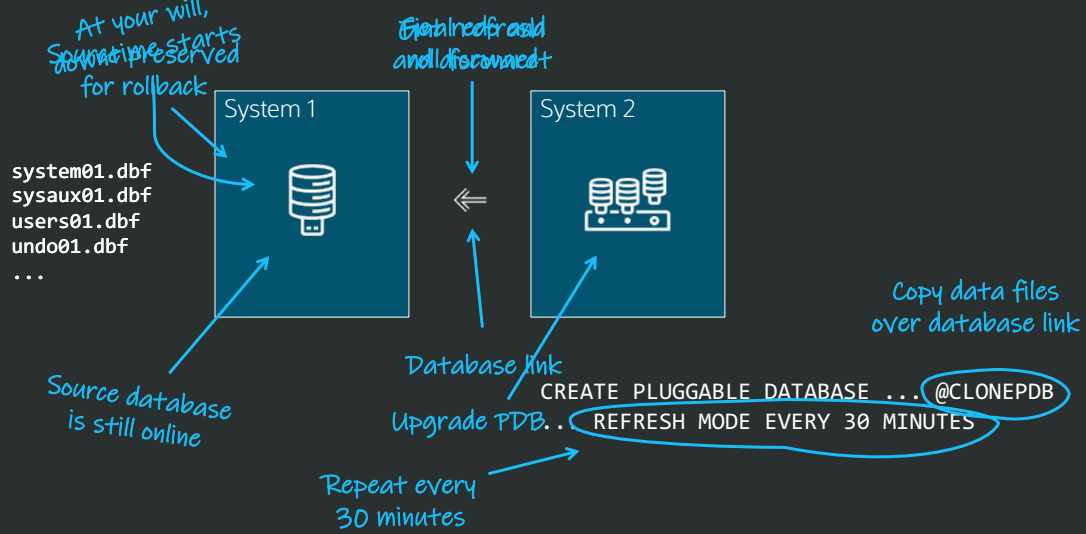

Refreshable Clone

Source non-CDB	Target CDB

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

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

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

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

You can drop user and database link after migration

Refreshable Clone

Source non-CDB	Target CDB

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

Refreshable Clone

Source non-CDB	Target CDB

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

Refreshable Clone

Source non-CDB	Target CDB

```
upg1.source_home=/u01/app/oracle/product/19
upg1.target_home=/u01/app/oracle/product/23
upg1.sid=NONCDB1
upg1.target_cdb=CDB1
upg1.source_dblink.NONCDB1=CLONEPDB 1800
upg1.target_pdb_name.NONCDB1=PDB1
```

Refreshable Clone

Source non-CDB	Target CDB

```
upg1.source_home=/u01/app/oracle/product/19
upg1.target_home=/u01/app/oracle/product/23
upg1.sid=NONCDB1
upg1.target_cdb=CDB1
upg1.source_dblink.NONCDB1=CLONEPDB 1800
upg1.target_pdb_name.NONCDB1=PDB1
upg1.start_time=23/11/2025 02:00:00
--Specify relative start time
--upg1.start_time=+1h30m
```

Refreshable Clone

Source non-CDB	Target CDB

```
upg1.source_home=/u01/app/oracle/product/19
upg1.target_home=/u01/app/oracle/product/23
upg1.sid=NONCDB1
upg1.target_cdb=CDB1
upg1.source_dblink.NONCDB1=CLONEPDB 300
upg1.target_pdb_name.NONCDB1=PDB1
upg1.start_time=25/01/2025 02:00:00
upg1.parallel_pdb_creation_clause=4
```

Help us! The cloning led to a massive network overload causing an outage.

Anonymous user that didn't set the parameter

Refreshable Clone

1

Run on source

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

2

Run on target

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

Refreshable Clone


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

```
upg1.start_time=23/11/2025 02:00:00
```

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

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

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

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

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

--Or reschedule it
upg> proceed -job 101 -newstarttime 30/11/2025 06:30:00
```

Works for unplug-plug upgrades as well

Works everywhere

Works everywhere

- Base Database Service
- Exadata Database Service
- Exadata Cloud@Customer
- On-prem and other clouds
- Exascale migration
- Non-CDB to PDB migration
- PDB upgrade
- On-prem to cloud migration

Key Benefits of Upgrade via Refreshable Clone PDB

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

Upgrading in the cloud

- Blog post: Upgrade Pluggable Database to Oracle Database 23ai
- Blog post: Upgrade Oracle Base Database Service to Oracle Database 23ai
- Blog post: Upgrade from 19c to 23ai using AutoUpgrade –ExaDB-D or ExaC@C – Part 1
- Blog post: When A Refreshable Clone Takes Over The Service
- Blog post: Upgrade Pluggable Database to Oracle Database 23ai
- Documentation: Proceed command

Refreshable clone works only with deferred recovery on standby database

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

Data Guard

Plug-in on primary propagates to standby database via redo

1 Enabled recovery

2 Deferred recovery

Enabled Recovery

1

Enabled recovery

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

Standby records PDB creation

Standby locates data files

MRP applies redo to PDB

PDB is immediately protected

Default

2

Deferred recovery

Enabled Recovery

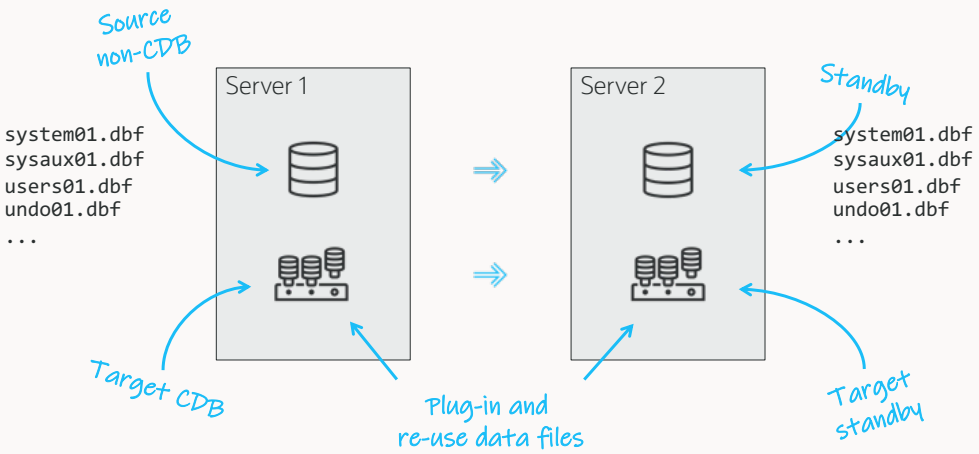

All data files on primary and standby must be at the same SCN

Enabled Recovery

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

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

Enabled Recovery

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

1 Regular files

2 OMF in regular file system

3 ASM

Enabled Recovery

1 Regular files

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

Enabled Recovery

2 OMF in regular file system

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

Enabled Recovery

3 ASM

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

Enabled Recovery | ASM

Primary

Standby

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

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

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

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

Same file,
but different name

Enabled Recovery | ASM

19c Non-CDB Primary

19c Non-CDB Standby

The manifest file contains

- File path on primary database only
- Not standby database

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

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

Enabled Recovery | ASM

Target primary

23ai
CDB
Primary

```
SQL> create pluggable database PDB1 using '/tmp/manifest_DB.xml' ... ;
```

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

23ai
CDB
Standby

Target standby

Enabled Recovery | ASM

23ai
CDB
Primary

+DATA/DB_BOSTON/DATAFILE/users.273.1103046827

Redo record says:
Plug in this data file

No good, data file
has a different name

23ai
CDB
Standby

+DATA/DB_CHICAGO/DATAFILE/users.269.1103050009

Enabled Recovery | ASM

23ai
CDB
Primary

+DATA/DB_BOSTON/DATAFILE/users.273.1103046827

OK, let's check the OMF directory

23ai
CDB
Standby

+DATA/DB_CHICAGO/DATAFILE/users.269.1103050009

+DATA/CDB1_CHICAGO/<PDB_GUID>/DATAFILE

It's empty

Enabled Recovery | ASM

23ai
CDB
Primary

+DATA/DB_BOSTON/DATAFILE/users.273.1103046827

OK, let's check the OMF directory

23ai
CDB
Standby

+DATA/DB_CHICAGO/DATAFILE/users.269.1103050009

+DATA/CDB1_CHICAGO/<PDB_GUID>/DATAFILE

It's empty

I'll just move the file in ASM

There's no move command in ASM.
How about copying?

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

ASMCMD-8016: copy source '+DATA/DB_BOSTON/.../users.269.1103050009' and target
'+DATA/DB_CHICAGO/.../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

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

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

Data Guard | Re-use Data Files

Looking for file like on primary

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

Follows alias and finds the real file

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

Enabled Recovery | Missing Data Files

What if a standby database fails to find data files?

- If Active Data Guard and PDB Recovery Isolation is turned on
  - New feature in Oracle Database 21c
  - Recovery disabled for PDB
  - Recovery proceeds in the entire CDB, except in specific PDB
  - Standby automatically restores data files from primary and re-enables recovery afterward
  - PDB protected after auto-restore

- If not, recovery halts in the **entire** CDB
  - **This is a critical situation**

Deferred Recovery

1

Enabled recovery

`create pluggable database ... standbys=all`

Standby records PDB creation

Standby locates data files

MRP applies redo to PDB

PDB is immediately protected

2

Deferred recovery

`create pluggable database ... standbys=none`

Standby records PDB creation

Standby ignores data files

MRP skips redo

PDB protected after restore

Deferred Recovery

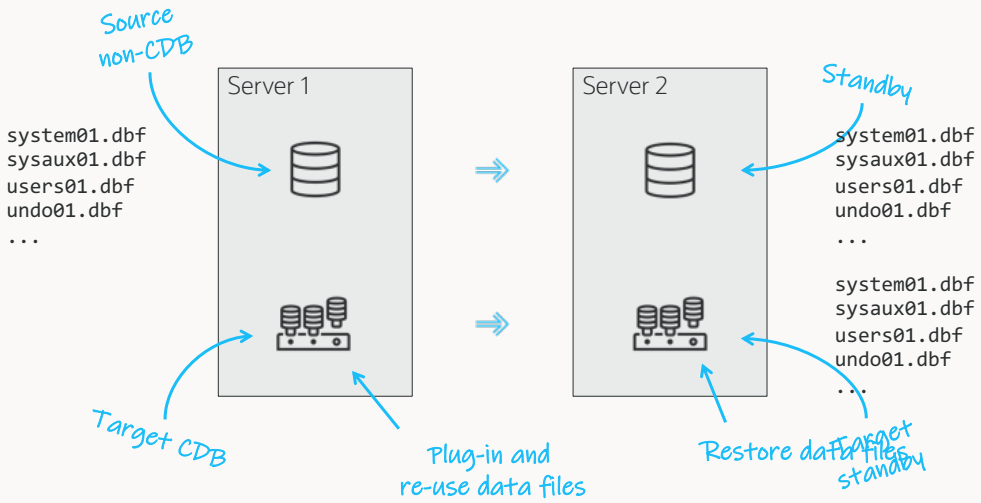

Deferred Recovery

Source
Non-CDB

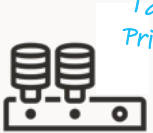

Target
Primary

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

Target
Standby

Deferred Recovery

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

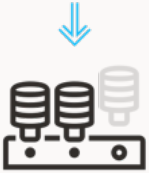

PDB created
Data files missing

Deferred Recovery

```
SQL> show pdbs

CON_NAME     OPEN MODE
PDB1         READ WRITE
```

```
SQL> show pdbs

CON_NAME     OPEN MODE
PDB1         MOUNTED
```

Deferred Recovery

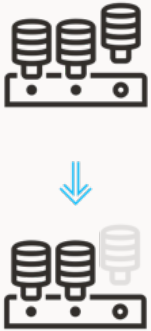

```
SQL> select name, recovery_status
     from v$pdbs;

NAME    RECOVERY_STATUS
PDB1    DISABLED
```

Deferred Recovery

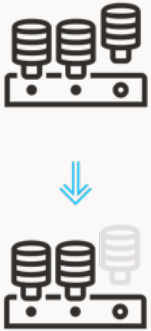

```
RMAN> restore pluggable database
      ... from service ... ;

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

Deferred Recovery

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

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

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

Data Guard and Multitenant Conversion

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

Don't jeopardize your Data Guard

- Test the procedure and verify your environment

Real World Customer Case

Techniker Krankenkasse

Move hundreds of Oracle Database19c PDBs to Oracle Database 23ai with ExaScale-only

Customer Case | Techniker Krankenkasse

Customer

Project

Constraints

Preparation

Upgrade

Success?

Remarks

Large German Public Health Insurance

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

Customer Case | Techniker Krankenkasse

Customer

Project

Constraints

Preparation

Upgrade

Success?

Remarks

Move to Oracle Database 23ai with ExaScale

- No ASM, no Oracle Database 19c anymore

Customer Case | Techniker Krankenkasse

Customer
Project
Constraints
Preparation
Upgrade
Success?
Remarks

`COMPATIBLE=23.0.0`

No ASM, no Oracle Database 19c anymore

Data Pump as rollback

Customer Case | Techniker Krankenkasse

Customer

Project

Constraints

Preparation

Upgrade

Success?

Remarks

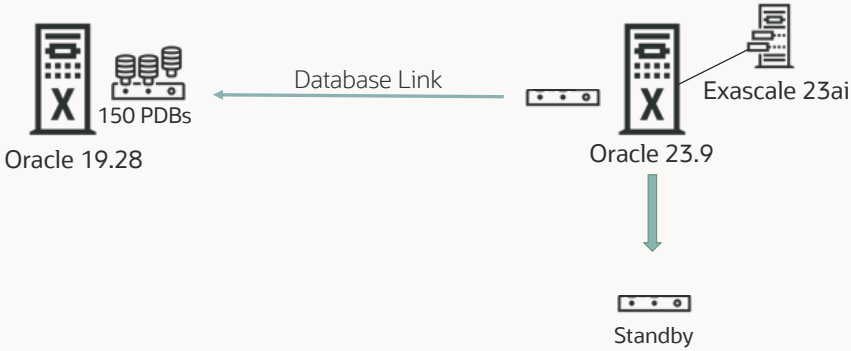

Customer Case | Techniker Krankenkasse

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

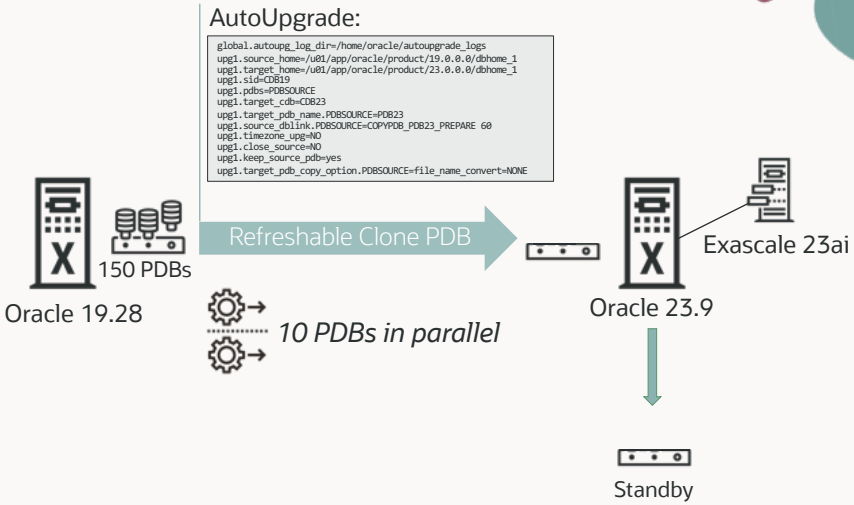

Customer Case | Techniker Krankenkasse

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

Oracle 19.28

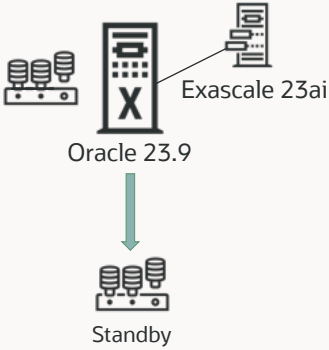

Customer Case | Techniker Krankenkasse

Customer

Project

Constraints

Preparation

Upgrade

Success?

Remarks

Yes!!

- Migrated 150 PDBs from TEST & DEV within less than 4 weeks
- Migrated larger environments already
- ExaScale snapshots are super-fast

Customer Case | Techniker Krankenkasse

Customer

Project

Constraints

Preparation

Upgrade

Success?

Remarks

Standby building

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

A number of SRs had to be opened for various areas

Optimizer works very well

Lab Exercise

Overview

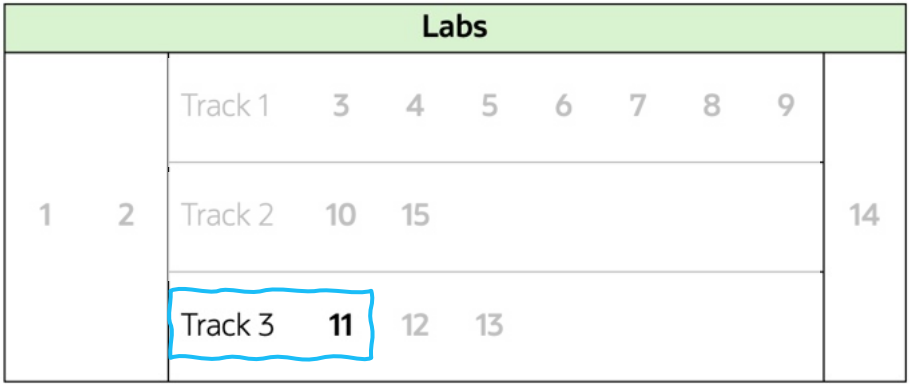

Labs

11. Upgrade non-CDB using refreshable clone PDB

12. Migrate data using Data Pump

13. Migrate data using Full Transportable Export/Import

Let's Get Started

https://livelabs.oracle.com

08789-XSBU-QTBE

Minimal Downtime

Reduce upgrade downtime to the time it takes to perform a switchover

- Rolling upgrades using a Transient Logical Standby

Rolling Upgrade | Standby Types

PHYSICAL	LOGICAL
Redo apply	SQL apply
Updated by changing data block	Updated by executing SQLs
Exact copy - block-by-block	Copy - data is the same

Rolling Upgrade | Concept

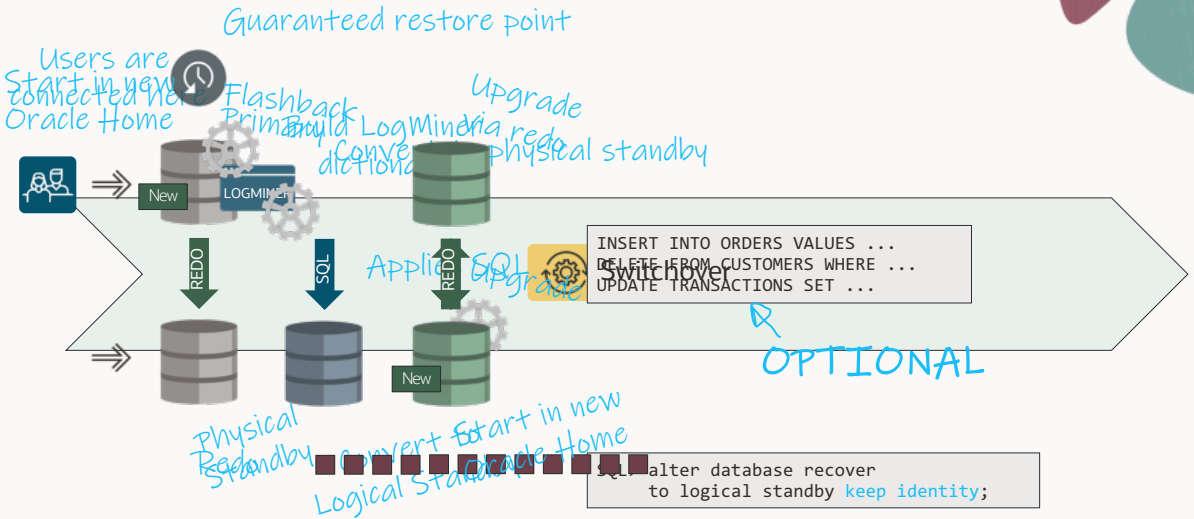

Rolling Upgrade | Options

MANUAL	DBMS_ROLLING
Part of Enterprise Edition	Requires Active Data Guard
Source must be 11.1.0.7	Source must be 12.1.0.2 or newer
Manual approach	Automated
Data Guard broker must be disabled	Data Guard broker can be enabled
	Recommended

Rolling Upgrade | Manual

MOS Note: 949322.1

Oracle11g Data Guard: Database Rolling Upgrade Shell Script

- Potentially not adjusted for Oracle 12c and newer
- Requires source is 11.2.0.3 or newer
- Does not work with Multitenant
- Not supported in 19c

Rolling Upgrade | Options

MANUAL	DBMS_ROLLING
Part of Enterprise Edition	Requires Active Data Guard
Source must be 11.1.0.7	Source must be 12.1.0.2 or newer
Manual approach	Automated
Data Guard broker must be disabled	Data Guard broker can be enabled
	Recommended

Rolling Upgrade | DBMS_ROLLING

```
SQL> exec dbms_rolling.init_plan;
SQL> exec dbms_rolling.build_plan;
SQL> exec dbms_rolling.start_plan;
```

6 SIMPLE STEPS

Upgrade database

```
SQL> exec dbms_rolling.switchover;
SQL> exec dbms_rolling.finish_plan;
```

```
...
Get current redo branch of the primary database
Wait until recovery is active on the primary's redo
branch
Reduce to a single instance if database is a RAC
Verify only a single instance is active if future
primary is RAC
Stop media recovery
Execute dbms_logstdby.build
Convert into a transient logical standby
Open database including instance-peers if RAC
Verify logical standby is open read/write
Get redo branch of transient logical standby
Get reset scn of transient logical redo branch
Configure logical standby parameters
Start logical standby apply
```

86 INSTRUCTIONS OR CHECKS

```
Start logical standby apply
Wait until apply lag has fallen below 600 seconds
Notify Data Guard broker that switchover to logical
standby database is starting
Log post-switchover instructions to events table
Switch database to a logical standby
Notify Data Guard broker that switchover to logical
standby database has completed
Wait until end-of-redo has been applied
...
```

After converting to logical standby database, take a level 0 backup

Also useful for
other maintenance activities

Can I use it on my database?

Determine database readiness

Do not create the logical standby on the same server as the primary

Supplemental logging is enabled automatically

- Introduces an overhead
- Increases amount of redo generated

When supplemental logging is enabled all DML cursors are invalidated

- Introduces an overhead
- Increases amount of redo generated

Not all data types and partitioning types are supported

- Introduces an overhead
- Increases amount of redo generated

Rolling Upgrade | Multitenant

- Rolling upgrade on container databases is fully supported
- Upgrade happens on CDB level - when you switchover - the entire CDB switches over
- The Transient Logical Standby can have a subset of the PDBs
- Adding new PDBs in primary after instantiating logical standby is possible, but cumbersome

Tips and tricks to ease your migration

For optimal performance all tables should have primary keys or unique keys

Use `ALTER DATABASE GUARD` to prevent accidental changes on logical standby

It is recommended to use three standbys for maximum protection

Upgrade Grid Infrastructure to new release before you start the process

Before starting rolling maintenance, test your Data Guard config

Plan your switchover during an off-peak period

Rolling Upgrade | Additional Information - 1

Documentation:

- Oracle Database Rolling Upgrades Using a Data Guard Physical Standby Database
- Oracle 19c Data Guard Concepts and Administration

MOS Notes:

- Transient Rolling Upgrade Using DBMS_ROLLING - Beginners Guide
- Rolling upgrade using DBMS_ROLLING - Complete Reference (Doc ID 2086512.1)
- MAA Whitepaper: SQL Apply Best Practices (Doc ID 1672310.1)
- Step by Step How to Do Swithcover/Failover on Logical Standby Environment (Doc ID 2535950.1)
- How To Skip A Complete Schema From Application on Logical Standby Database (Doc ID 741325.1)
- How to monitor the progress of the logical standby (Doc ID 1296954.1)
- How To Reduce The Performance Impact Of LogMiner Usage On A Production Database (Doc ID 1629300.1)
- Exadata Cloud Database 19c Rolling Upgrade With DBMS_ROLLING (Doc ID 2832235.1)

Rolling Upgrade | Additional Information - 2

MOS Notes:

- Handling ORA-1403 ora-12801 on logical standby apply (Doc ID 1178284.1)
- Troubleshooting Example - Rolling Upgrade using DBMS_ROLLING (Doc ID 2535940.1)
- DBMS Rolling Upgrade Switchover Fails with ORA-45427: Logical Standby Redo Apply Process Was Not Running (Doc ID 2696017.1)
- SRDC - Collect Logical Standby Database Information (Doc ID 1910065.1)
- MRP fails with ORA-19906 after Flashback of Transient Logical Standby used for Rolling Upgrade (Doc ID 2069325.1)
- What Causes High Redo When Supplemental Logging is Enabled (Doc ID 1349037.1)
- Logical Standby SQL APPLY Tuning Tips (Doc ID 2674154.1)

Bugs:

- BUG 22541208 - REPLICATION FAILS WITH ORA-02149 DROPING PARTITION WITH SYSTEM GENERATED NAME (fixed in 12.2 backport available for 12.1)
- BUG 31412209 - TRANSIENT LOGICAL STANDBY UPGRADE FAILING WITH ORA-600[KRVXSAU_122_12202_LCR_OP] (fixed in 12.2 backport available for lower versions)

Data Pump

Use the interactive console

```
-- Use the job name parameter to give your job a meaning name
-- Makes it easier to identify details of a specific job

expdp ... job_name=APP_EXPORT

impdp ... job_name=APP_IMPORT
```

```
$ impdp ... job_name=APP_IMPORT


Import: Release 19.0.0.0.0 - Production on Wed Apr 30 17:09:10 2025
Version 19.27.0.0.0

Copyright (c) 1982, 2019, Oracle and/or its affiliates.  All rights reserved.

Connected to: Oracle Database 19c Enterprise Edition Release 19.0.0.0.0 - Production
30-APR-25 17:09:13.051: W-1 Startup took 1 seconds
30-APR-25 17:09:13.779: W-1 Master table "DPUSER"."APP_IMPORT" successfully loaded/unloaded
30-APR-25 17:09:13.892: Starting "DPUSER"."APP_IMPORT":  dpuser/******** parfile=...
30-APR-25 17:09:13.901: W-1 Processing object type SCHEMA_EXPORT/USER
30-APR-25 17:09:13.974: W-1      Completed 1 USER objects in 0 seconds
30-APR-25 17:09:13.974: W-1      Completed by worker 1 1 USER objects in 0 seconds
```

```
$ impdp ... attach=APP_IMPORT
```

```
Job: APP_IMPORT
  Operation: IMPORT
  Mode: FULL
  State: EXECUTING
  Bytes Processed: 0
  Current Parallelism: 4
  Job Error Count: 0
  Job heartbeat: 2
  Dump File: /home/oracle/dpdir/faster-import-constraints.dmp

Worker 1 Status:
  Instance ID: 1
  Instance name: FTEX
  Host name: holserv1.livelabs.oraclevcn.com
  Object start time: Wednesday, 30 April, 2025 17:09:14
  Object status at: Wednesday, 30 April, 2025 17:09:14

(output truncated)
```

```
--Get status
status
status=120

--Enable tracing
trace=<nnn>

--Change the number of workers
parallel=<n>

--List all commands
help
```

You can also enter the interactive console by hitting CTRL+C on your import

- Hit it just once - otherwise - you kill the process

DEMO

Interactive Command Mode

Watch on YouTube

The control table

```
$ impdp ... job_name=APP_IMPORT


Import: Release 19.0.0.0.0 - Production on Wed Apr 30 17:09:10 2025
Version 19.27.0.0.0

Copyright (c) 1982, 2019, Oracle and/or its affiliates.  All rights reserved.

Connected to: Oracle Database 19c Enterprise Edition Release 19.0.0.0.0 - Production
30-APR-25 17:09:13.051: W-1 Startup took 1 seconds
30-APR-25 17:09:13.779: W-1 Master table "DPUSER"."APP_IMPORT" successfully loaded/unloaded
30-APR-25 17:09:13.892: Starting "DPUSER"."APP_IMPORT":  dpuser/******** parfile=...
30-APR-25 17:09:13.901: W-1 Processing object type SCHEMA_EXPORT/USER
30-APR-25 17:09:13.974: W-1      Completed 1 USER objects in 0 seconds
30-APR-25 17:09:13.974: W-1      Completed by worker 1 1 USER objects in 0 seconds
```

```
$ impdp ... job_name=APP_IMPORT KEEP_MASTER=Y


Import: Release 19.0.0.0.0 - Production on Wed Apr 30 17:09:10 2025
Version 19.27.0.0.0

Copyright (c) 1982, 2019, Oracle and/or its affiliates.  All rights reserved.

Connected to: Oracle Database 19c Enterprise Edition Release 19.0.0.0.0 - Production
30-APR-25 17:09:13.051: W-1 Startup took 1 seconds
30-APR-25 17:09:13.779: W-1 Master table "DPUSER"."APP_IMPORT" successfully loaded/unloaded
30-APR-25 17:09:13.892: Starting "DPUSER"."APP_IMPORT":  dpuser/******** parfile=...
30-APR-25 17:09:13.901: W-1 Processing object type SCHEMA_EXPORT/USER
30-APR-25 17:09:13.974: W-1      Completed 1 USER objects in 0 seconds
30-APR-25 17:09:13.974: W-1      Completed by worker 1 1 USER objects in 0 seconds
```

```
$ impdp ... job_name=APP_IMPORT MASTER_ONLY=Y


Import: Release 19.0.0.0.0 - Production on Wed Apr 30 17:09:10 2025
Version 19.27.0.0.0

Copyright (c) 1982, 2019, Oracle and/or its affiliates.  All rights reserved.

Connected to: Oracle Database 19c Enterprise Edition Release 19.0.0.0.0 - Production
30-APR-25 17:09:13.051: W-1 Startup took 1 seconds
30-APR-25 17:09:13.779: W-1 Master table "DPUSER"."APP_IMPORT" successfully loaded/unloaded
30-APR-25 17:09:13.892: Starting "DPUSER"."APP_IMPORT":  dpuser/******** parfile=...
30-APR-25 17:09:13.901: W-1 Processing object type SCHEMA_EXPORT/USER
30-APR-25 17:09:13.974: W-1      Completed 1 USER objects in 0 seconds
30-APR-25 17:09:13.974: W-1      Completed by worker 1 1 USER objects in 0 seconds
```

Use of the control table is not documented

```
SQL> select object_name, object_type, process_order
     from dpuser.monitoring where process_order > 0 order by process_order;

OBJECT_NAME                    OBJECT_TYPE                    PROCESS_ORDER
------------------------------ ------------------------------ -------------
CONSTR_VALIDATE                USER                                   1
UNLIMITED TABLESPACE           SYSTEM_GRANT                           2
DBA                            ROLE_GRANT                             3
                               DEFAULT_ROLE                           4
                               PROCACT_SCHEMA                         5
T1                             TABLE                                  6
T2                             TABLE                                  9
T1                             TABLE_DATA                            10
                                                                     10
T2                             TABLE_DATA                            11
                                                                     11
C_TAB1_C01                     CONSTRAINT                            14

(output truncated)

C_TAB2_C12                     CONSTRAINT                           104
C_TAB2_C02                     CONSTRAINT                           105
```

```
$ impdp ... abort_step=10
```

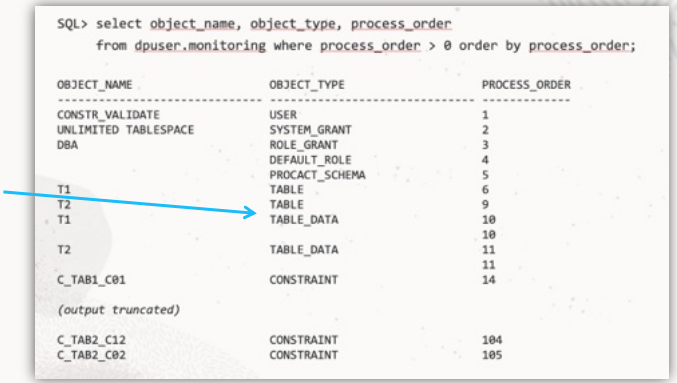

```
SQL> select object_name, object_type, process_order
     from dpuser.monitoring where process_order > 0 order by process_order;

OBJECT_NAME                    OBJECT_TYPE                    PROCESS_ORDER
------------------------------ ------------------------------ -------------
CONSTR_VALIDATE                USER                                       1
UNLIMITED TABLESPACE           SYSTEM_GRANT                               2
DBA                            ROLE_GRANT                                 3
                               DEFAULT_ROLE                               4
                               PROCACT_SCHEMA                             5
T1                             TABLE                                      6
T2                             TABLE                                      9
T1                             TABLE_DATA                                10
                                                                         10
T2                             TABLE_DATA                                11
                                                                         11
C_TAB1_C01                     CONSTRAINT                                14

(output truncated)

C_TAB2_C12                     CONSTRAINT                               104
C_TAB2_C02                     CONSTRAINT                               105
```

```
Import> start_job
```

Troubleshooting

In root and PDB

```
-- Change AWR snap interval to 15 minutes and create snapshot
exec dbms_workload_repository.modify_snapshot_settings(null, 15);
exec dbms_workload_repository.create_snapshot;
```

```
-- Change AWR snap interval to 15 minutes and create snapshot
exec dbms_workload_repository.modify_snapshot_settings(null, 15);
exec dbms_workload_repository.create_snapshot;

-- Optionally, enable SQL trace for Data Pump processes or specific SQL ID
alter system set events 'sql_trace {process: pname = dw | process: pname = dm} level=8';
alter system set events 'sql_trace[SQL: 03g1bnw08m4ds]';
```

```
-- Change AWR snap interval to 15 minutes and create snapshot
exec dbms_workload_repository.modify_snapshot_settings(null, 15);
exec dbms_workload_repository.create_snapshot;

-- Optionally, enable SQL trace for Data Pump processes or specific SQL ID
alter system set events 'sql_trace {process: pname = dw | process: pname = dm} level=8';
alter system set events 'sql_trace[SQL: 03g1bnw08m4ds]';

-- Run Data Pump job with trace (Doc ID 286496.1)
expdp ... metrics=yes logtime=all trace=1FF0300
impdp ... metrics=yes logtime=all trace=1FF0300
```

CONTROL PROCESS

Typically one: `dm00`

```
DB19_dm00_17468.trc
```

WORKERS

Typically many: `dwnn`

```
DB19_dw00_17469.trc
DB19_dw01_17470.trc
DB19_dw02_17471.trc
DB19_dw03_17472.trc
```

Tracing may generate a small overhead

- Up to 2-3 %

```
-- Change AWR snap interval to 15 minutes and create snapshot
exec dbms_workload_repository.modify_snapshot_settings(null, 15);
exec dbms_workload_repository.create_snapshot;

-- Optionally, enable SQL trace for Data Pump processes or specific SQL ID
alter system set events 'sql_trace {process: pname = dw | process: pname = dm} level=8';
alter system set events 'sql_trace[SQL: 03g1bnw08m4ds]';

-- Run Data Pump job with trace (Doc ID 286496.1)
expdp ... metrics=yes logtime=all trace=1FF0300
impdp ... metrics=yes logtime=all trace=1FF0300
```

```
Processing object type DATABASE_EXPORT/FINAL_POST_INSTANCE_IMPCALLOUT/MARKER
. . exported "SYS"."KU$_USER_MAPPING_VIEW"                5.890 KB      25 rows
. . exported "SYSTEM"."REDO_DB"                           25.59 KB       1 rows
```

```
02-NOV-21 19:43:59.380: W-1 Processing object type DATABASE_EXPORT/POST_SYSTEM_IMPCALLOUT/MARKER
02-NOV-21 19:43:59.387: W-1      Completed 1 MARKER objects in 0 seconds
02-NOV-21 19:43:59.830: W-1 . . exported "SYS"."KU$_USER    5.890 KB  25 rows in 0 seconds using external_table
02-NOV-21 19:43:59.923: W-1 . . exported "SYSTEM"."REDO_DB" 25.59 KB   1 rows in 0 seconds using direct_path
```

Enabling diagnostic information does not generate overhead

```
-- Change AWR snap interval to 15 minutes and create snapshot
exec dbms_workload_repository.modify_snapshot_settings(null, 15);
exec dbms_workload_repository.create_snapshot;

-- Optionally, enable SQL trace for Data Pump processes or specific SQL ID
alter system set events 'sql_trace {process: pname = dw | process: pname = dm} level=8';
alter system set events 'sql_trace[SQL: 03g1bnw08m4ds]';

-- Run Data Pump job with trace (Doc ID 286496.1)
expdp ... metrics=yes logtime=all trace=1FF0300
impdp ... metrics=yes logtime=all trace=1FF0300

-- Create AWR snapshot and produce AWR report
exec dbms_workload_repository.modify_snapshot_settings(null, <original-value>);
exec dbms_workload_repository.create_snapshot;
@?/rdbms/admin/awrrpt
```

In root and PDB

Troubleshooting

Collect:

- Data Pump log file
- AWR report - CDB and PDB level
- Data Pump trace files
  - Stored in the database trace directory
  - Control process file name: `*dm*`
  - Worker process file names: `*dw*`

New Data Pump diagnostic views

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

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

```
select * from v$datapump_process_info;

select * from v$datapump_processwait_info;
```

Analyzing log files

```
-- How do you deal with large Data Pump import log files?
-- In this example, the Data Pump import log file has almost 200.000 lines

$ du -h import.log
 29M   import.log

$ wc -l import.log
  189931 import.log
```

```
$ python3 dpla.py import.log

========================
 Data Pump Log Analyzer
========================

...

Operation Details
~~~~~~~~~~~~~~~~~
 Operation:                  Import
 Data Pump Version:          19.22.0.0.0
 DB Info:                    Oracle Database 19c EE Extreme Perf Release 19.0.0.0.0
 Job Name:                   DPJOB1
 Status:                     COMPLETED
   Processing:               -
 Errors:                     1267
   ORA- Messages:            1267
 Start Time:                 2024-04-11 09:30:55
 End Time:                   2024-04-12 10:33:01
 Runtime:                    25:03:06

Data Processing
~~~~~~~~~~~~~~~
 Parallel Workers:           128
 Schemas:                    27
 Objects:                    224755
 Data Objects:               188084
 Overall Size:               13.16 TB
```

```
$ python3 dpla.py import.log -e

========================
 Data Pump Log Analyzer
========================

...

ORA- MESSAGES DETAILS
~~~~~~~~~~~~~~~~~~~~~
(sorted by count):
```

Message	Count
ORA-39346: data loss in character set conversion for object COMMENT	919
ORA-39082: Object type PACKAGE BODY created with compilation warnings	136
ORA-39346: data loss in character set conversion for object PACKAGE_BODY	54
ORA-39082: Object type TRIGGER created with compilation warnings	36
ORA-39082: Object type PROCEDURE created with compilation warnings	29
ORA-31684: Object type USER already exists	27
ORA-39111: Dependent object type PASSWORD_HISTORY skipped, base object type USER already exists	27
ORA-39346: data loss in character set conversion for object PACKAGE	18
ORA-39082: Object type PACKAGE created with compilation warnings	10
ORA-39082: Object type VIEW created with compilation warnings	7
ORA-39346: data loss in character set conversion for object PROCEDURE	2
ORA-39082: Object type FUNCTION created with compilation warnings	2
Total	1267

```
$ python3 dpla.py import.log -o

=========================
 Data Pump Log Analyzer
=========================

...
```

Object	Count	Seconds	Workers	Duration
SCHEMA_EXPORT/TABLE/TABLE_DATA	188296	6759219	128	6759219
CONSTRAINT	767	37253	1	37253
TABLE	2112	3225	51	156
COMMENT	26442	639	128	18
PACKAGE_BODY	197	125	128	5
OBJECT_GRANT	5279	25	1	25
TYPE	270	6	1	6
ALTER_PROCEDURE	149	5	2	3
ALTER_PACKAGE_SPEC	208	4	3	2
PACKAGE	208	3	3	1
PROCEDURE	149	2	2	1

...

Total	224755	6800515	128	6796697

How about NOVALIDATE constraints?

Data Pump Log Analyzer

Table Details

Search for Table...

Table	Rows	Size	Seconds	Part	Subpart
SALES.ORDERS	118914251151	1.73 TB	878854	278	4448
SALES.INVOICES	115668171592	4.33 TB	805901	588	9408
SALES.TRANSACTIONS	115720037994	3.61 TB	611891	451	7216
FINANCE.EXPENSES	35091517646	258.14 GB	112962	367	0
MARKETING.CAMPAIGNS	11621627768	458.93 GB	82801	16	0
HR.EMPLOYEES	19433932893	296.19 GB	66156	2254	0
SALES.DOCUMENTS	4743542596	345.97 GB	48117	589	9424
SALES.REPORTS	4744610748	263.63 GB	42904	440	7040
INVENTORY.EQUIPMENT	9824954344	51.01 GB	33290	130	0
HR.PARTNERS	3983265247	83.62 GB	16388	3046	0

- Free to use
- Download from GitHub
- Not an official Oracle tool
- Created by Marcus Doeringer
  Our migration superstar

Autonomous AI Database

A migration approach

What's the story?

Flashback to October 2017

The idea?

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

KEYNOTE PRESENTATION

Oracle OpenWorld San Francisco 2017


OOW 2017
ORACLE

Photo by LARAM on Unsplash

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

Tim Hall

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

Photo by Leo on Unsplash

Autonomous AI Database – Where?

Public cloud

Autonomous AI Database

Cloud@Customer

Autonomous AI Database in a VM environment

One Autonomous AI Database – Two Deployment Choices

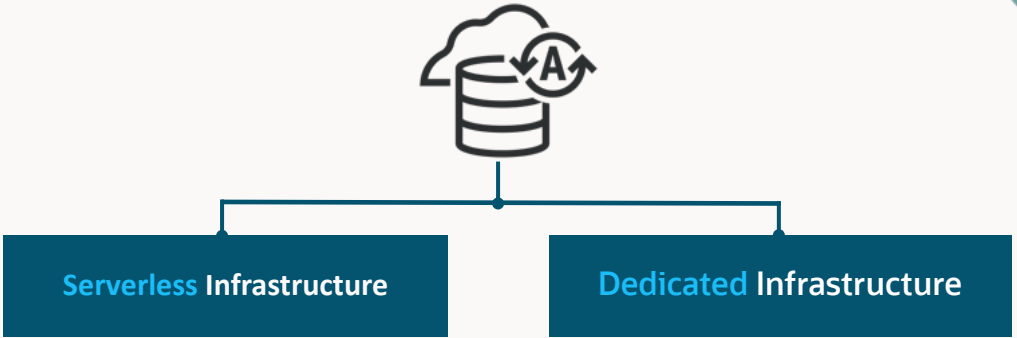

One Autonomous AI Database – Workload Choices

Autonomous AI Lakehouse (ADW)

- Data Warehouse, Data Mart
- Data Lake, Machine Learning

Autonomous Transaction Processing (ATP)

- Transactions, Batch, Reporting, IoT
- Application Dev, Machine Learning

Migration Planning

No migration without a proper runbook

Estate Modernization

But not every database is a great candidate for ADB

Migration to Autonomous AI Database is always a logical migration

- Move the data, not the database

Tools out-of-the-box

SQL Developer Web

Web-based Function rich, low code development env

No client software needed

Oracle REST Data Services

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

APEX

Execute SQL and PL/SQL

Build Data Models, generate DDL statements

Monitor and manage the DB

Essentials

Outside the database,
check OCI Documentation

Example: Deploy, start, stop, scale

Essentials

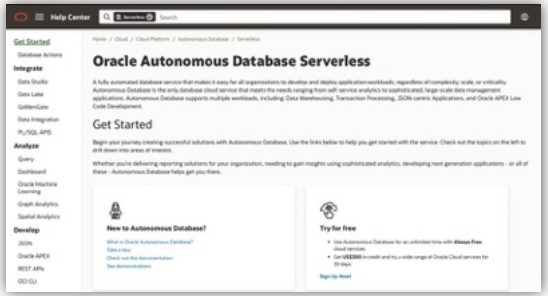

Inside the database,
check Database Documentation

Example: Schema, capabilities, connecting

Planning

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

- And which ones are good candidates?

Getting an Overview

1
Estate Explorer

2
Cloud Premigration Advisor Tool

3
Cloud Migration Advisor

Getting an Overview

1
Estate Explorer

2
Cloud Premigration Advisor Tool

3
Cloud Migration Advisor

Getting an Overview

1 Estate Explorer → **2** Cloud Premigration Advisor Tool → **3** Cloud Migration Advisor

Analyze 1000's of databases in just a few hours

Provide a detailed TCO to compare on-premises and cloud

View innovative visualizations and detailed reports

Optimize your Autonomous AI Databases using Elastic Pools

Getting an Overview

1 Estate Explorer

2 Cloud Premigration Advisor Tool

3 Cloud Migration Advisor

Getting an Overview

1 Estate Explorer

2 Cloud Premigration Advisor Tool

3 Cloud Migration Advisor

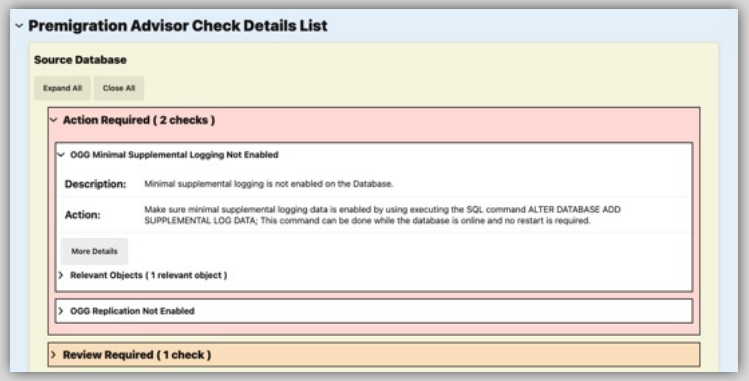

Getting an Overview

1 Estate Explorer

2 Cloud Premigration Advisor Tool

3 Cloud Migration Advisor

Getting an Overview

1 Estate Explorer

2 Cloud Premigration Advisor Tool

3 Cloud Migration Advisor

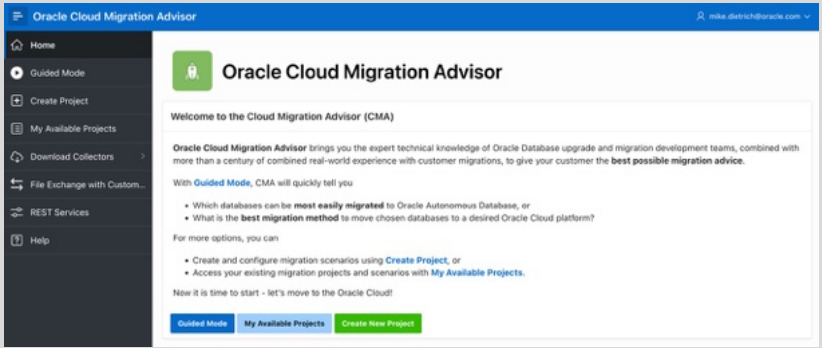

Getting an Overview

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

Cloud Premigration Advisor Tool

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

- Use Cloud Premigration Advisor Tool (CPAT)

Overview

Connects

Checks

Reports

Fixes

(optional)

Download CPAT from MOS Note: 2758371.1

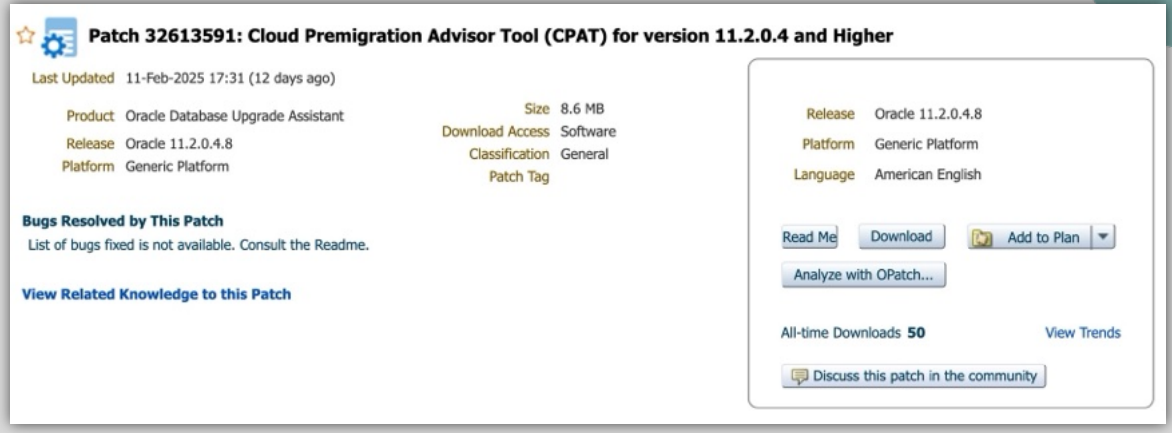

```
# One or more report formats separated by spaces
# json html text

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

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

CPAT COMPOSER

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

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

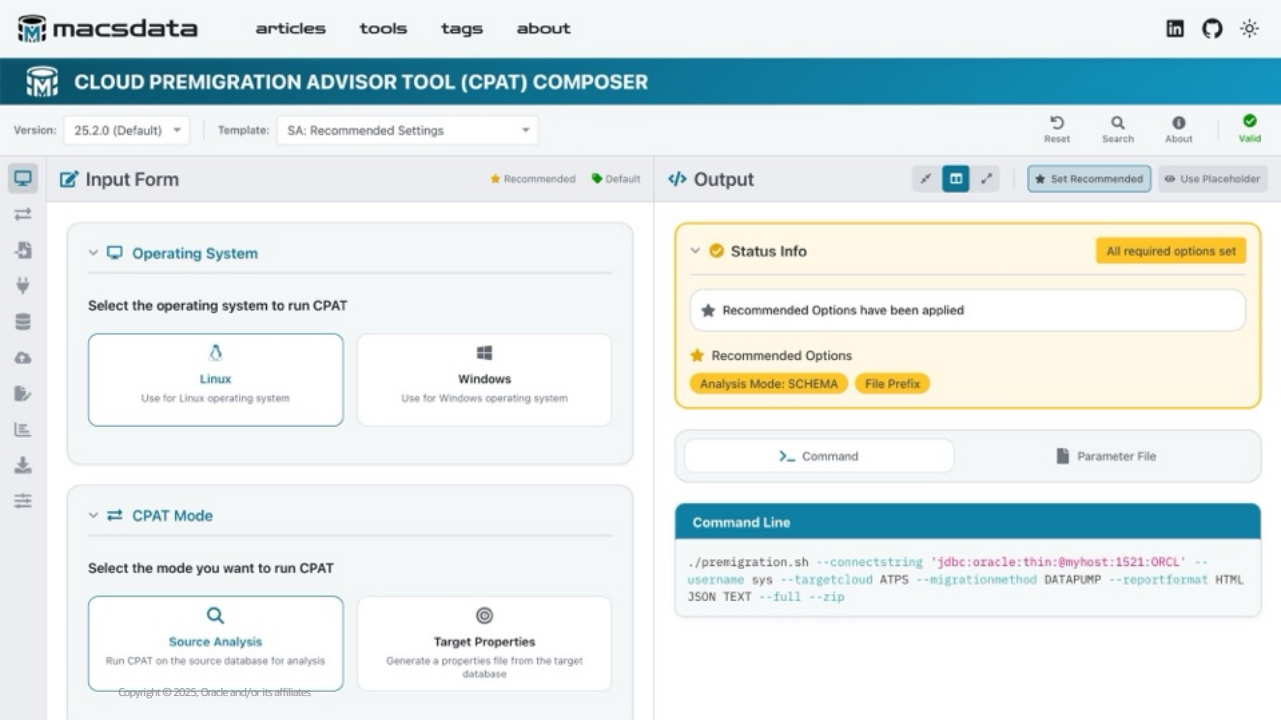
macsdata
articles
tools
tags
about
CLOUD PREMIGRATION ADVISOR TOOL (CPAT) COMPOSER
Version:
25.2.0 (Default)
Template:
SA: Recommended Settings
Reset
Search
About
Valid
Input Form
Recommended
Default
Operating System
Select the operating system to run CPAT
Linux
Use for Linux operating system
Windows
Use for Windows operating system
CPAT Mode
Select the mode you want to run CPAT
Source Analysis
Run CPAT on the source database for analysis
Target Properties
Generate a properties file from the target database
Output
Set Recommended
Use Placeholder
Status Info
All required options set
Recommended Options have been applied
Recommended Options
Analysis Mode: SCHEMA
File Prefix
Command
Parameter File
Command Line
./premigration.sh --connectstring 'jdbc:oracle:thin:@myhost:1521:ORCL' --username sys --targetcloud ATPS --migrationmethod DATAPUMP --reportformat HTML JSON TEXT --full --zip
Copyright © 2025, Oracle and/or its affiliates

ORACLE

Sat Feb 22 20:59:55 UTC 2025

Cloud Premigration Advisor Tool (CPAT) Report

CPAT Version: 25.2.1-1
Version Date: Feb 17, 2025
Days Since Last CPAT Update: 38 days

Table of Contents

Expand All | Close All

Premigration Advisor Report Summary

Report Result	Action Required
Number of schemas analyzed:	4

Expand All | Close All

Premigration Advisor Report Summary

Report Result	Action Required
Number of schemas analyzed:	4
List of schemas analyzed:	[PDBUSER, CO, HR, MYAPP]

Report Results Overview

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

Return to Table of Contents

Report Details

Return to Table of Contents

Report Details

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

More Details

Return to Table of Contents

Report Analysis Notes

More Details

Return to Table of Contents

Report Analysis Notes

Source Database Details

Source Database Version Information

Source Database Patch Information

Source Database Redo Information

Source Database Supplemental Information

Source Database Schema Summary Information

Premigration Advisor Check Details List

Source Database

Source Database Details

Source Cloud Vendor:	Oracle Cloud Infrastructure (Database)
Source Database Host Name:	dbsystemaz
Source Oracle SID:	ORCL
Source Database Created Date:	Fri Jan 24 22:23:51 UTC 2025
Source Database DBID:	1719058167
Source Database Unique Name:	ORCL_5tr_iad
Source Instance Name:	ORCL
Source Database Name:	ORCL
Source Database Username:	SYS
Source Database Port String:	x86_64/Linux 2.4.xx
Source Database Platform ID:	13
Source Database Container Name:	PDB_COMPLEX
Source DB Block Size in KB:	8
Source DB Combined Size of DATA, TEMP, LOG, and CONTROL File Usage in GB:	5.044
Source DB Size of DATA File Usage in GB:	1.856
Source DB Size of TEMP File Usage in GB:	0.17

More Details

Return to Table of Contents

Report Analysis Notes

Source Database Details

Source Database Version Information

Source Database Patch Information

Source Database Redo Information

Source Database Supplemental Information

Source Database Schema Summary Information

Premigration Advisor Check Details List

Source Database

Action Required (2 checks)

OGG Minimal Supplemental Logging Not Enabled

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

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

More Details

Relevant Objects (1 relevant object)

OGG Replication Not Enabled

Description: ENABLE_GOLDENGATE_REPLICATION init.ora parameter is not set.

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

More Details

Relevant Objects (1 relevant object)

The documentation has additional information on each CPAT check

Utilities Guide, Oracle Database 23ai

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

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

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

CPAT integration

Generate CPAT report

ZERO DOWNTIME MIGRATION
Run as part of its migration assessment

OCI DATABASE MIGRATION SERVICE
Run as part of its migration assessment

ENTERPRISE MANAGER MIGRATION WORKBENCH
Run as part of its migration assessment

SQL DEVELOPER / SQLcl
Through the MIGRATEADVISOR command

Data Pump

The simple approach

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

- You may request one-off fixes via an SR

Allocate a sufficient number of ECPUs

- 32 should be the minimum when you import

Export: PARALLEL 2x of physical cores

Import: PARALLEL=ECPU/4, or higher

- Scale up to the maximum for migrations

Most simple method: Data Pump

Datapump with Files

Datapump with DB Links

Automation

How AMA scripts ease migrations

What is AMA?

Autonomous AI Database Migration Automation (AMA)

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

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

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

You won't start with filming right away

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

Now you can start filming your scenes

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

AMA Workflow

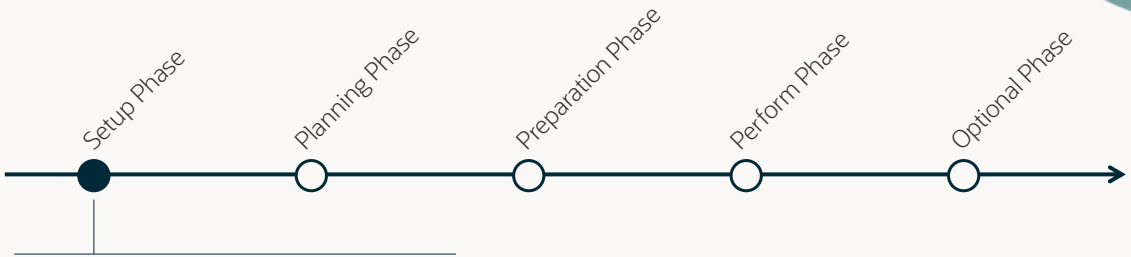

- Examination of source database (CPAT)
- Create migration directories
- Configure AMA parameter file

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

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

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

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

Parameter File

Adjust:

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

Documentation

AMA Documentation is available at request

- Documents the entire flow and all options and parameters

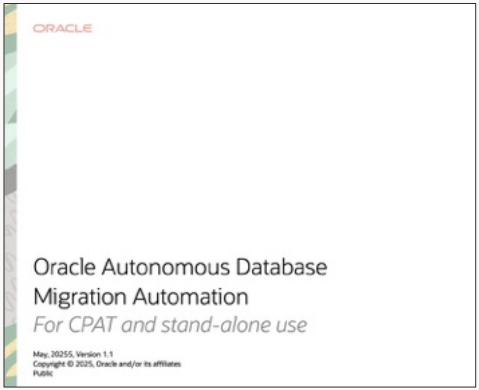

AMA Workflow

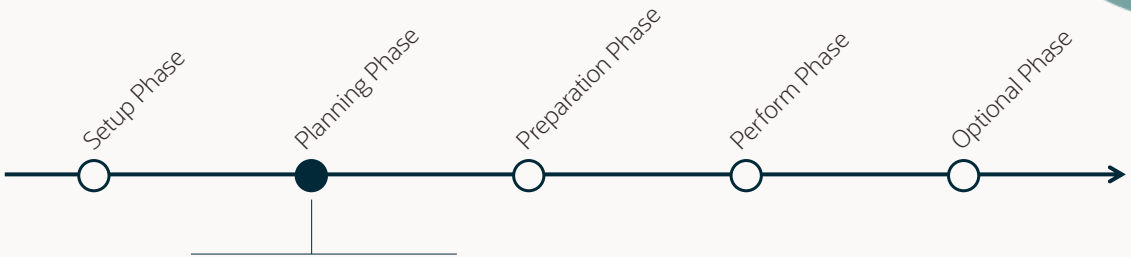

AMA | Planning Phase

On-Prem - Source

Gather stats for SYS / SYSTEM

Create Migration user

Enable restricted session

Set JOB_QUEUE_PROCESSES=0

ADB-S - Target

Enable OLAP / JAVA in ADB-S

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

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

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

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

AMA Workflow

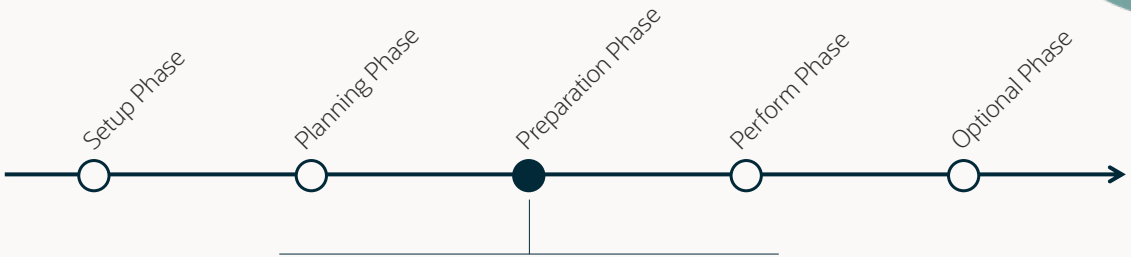

AMA | Preparation Phase

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

```
---
--- PREPARATION PHASE ---
--- All steps in this phase will prepare the source and target database ---
--- The scripts depend on each other, so execute in this phase one script after the other ---
---
---
```

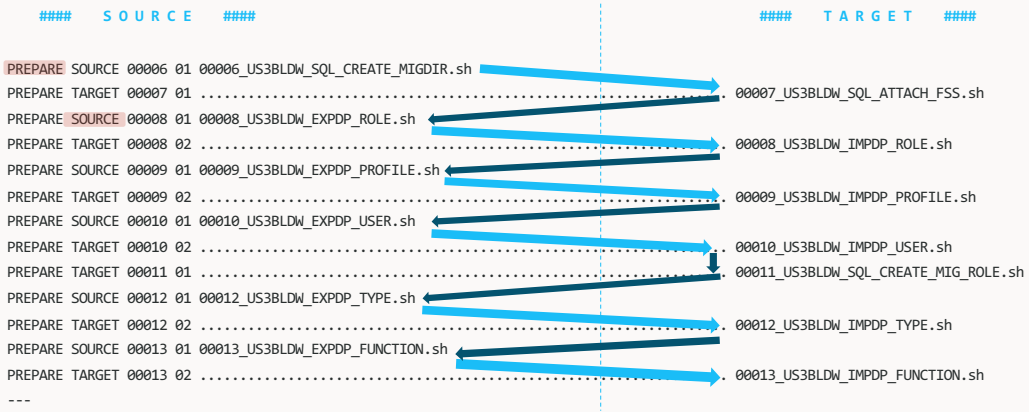


```
---
```

AMA Workflow

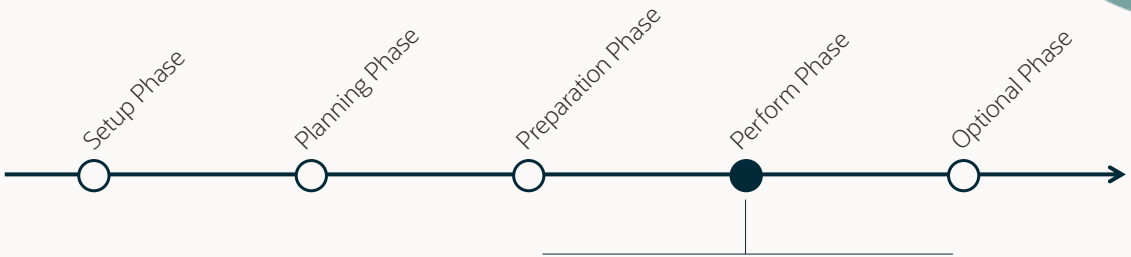

AMA | Perform Phase

On-Prem - Source

Export all schemas

Export audit trail

Copy files (if necessary)

ADB-S - Target

Import all schemas

Import audit trail

AMA Workflow

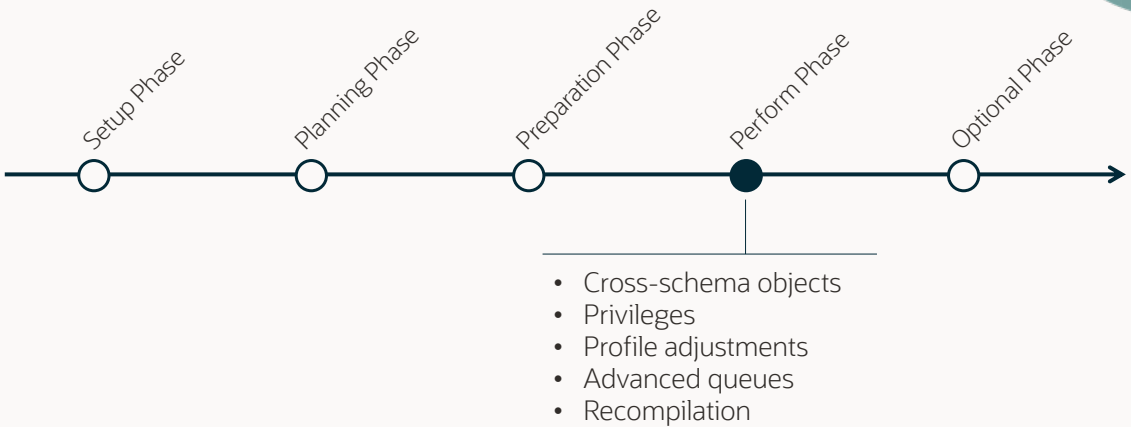

```
---
--- PERFORM PHASE ---
--- Commonly in this phase nothing depends on each other (except you for example have objects that depend on objects stored in other schema) ---
--- So export jobs can be started in parallel and imports once the export finished ---
---
---

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

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

AMA | Perform Phase

On-Prem - Source

ADB-S - Target

- `FOREIGN KEYS` cross-schemas
- `INDEXES` cross-schemas
- `FUNCTIONAL INDEXES` enableing
- `REVOKE` transition privileges
- `GRANT` privs SYS, SYSTEM, CTXSYS, objects
- Restore final profiles
- Set tablespace quotas

Export network ACLs

- Import network ACLS
- Enable Advanced Queues
- Recompilation

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

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

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

AMA Workflow

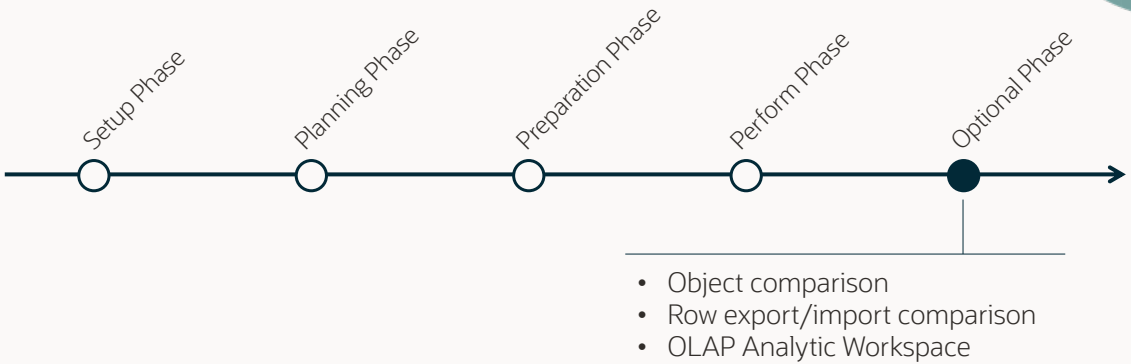

Done!!

AMA can run a migration fully automated and completely unattended

Works with Windows as source database

Database links, external tables, APEX applications

- Work-in-progress

Key Learnings

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

Virtual Classroom Seminar Series #22 – #25

1 PLANNING

Watch recording
Get slides

2 PREPARING

Watch recording
Get slides

3 MIGRATING

Watch recording
Get slides

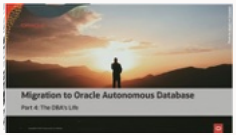

4 OPERATING

Wach recording
Get slides

Try it out, please!!

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

YouTube Channel

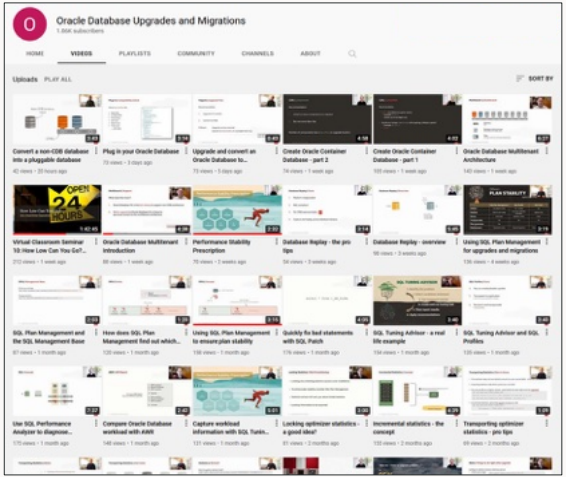

https://www.youtube.com/@upgradenow

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

SCAN ME


ORACLE