



ORACLE

Oracle Database Upgrade and Migration 19c & 23ai

Warsaw, February 2025

Oracle

DBAs

run the world





MIKE DIETRICH

Vice President
Database Upgrade, Migrations & Patching



mikedietch



@mikedietchde.com



<https://mikedietchde.com>





DANIEL OVERBY HANSEN

Distinguished Product Manager
Database Upgrade, Migrations & Patching



dohdatabase



@dohdatabase.com



<https://dohdatabase.com>

Get the Slides

<https://dohdatabase.com/slides>



Find Slides and Much More on Our Blogs



MikeDietrichDE.com

Mike.Dietrich@oracle.com



dohdatabase.com

Daniel.Overby.Hansen@oracle.com



DBArj.com.br

Rodrigo.R.Jorge@oracle.com



AlexZaballa.com

Alex.Zaballa@oracle.com

Web Seminar

Episode 16

(replaces Episode 1 from Feb 2021)

Oracle Database Release and Patching Strategy for 19c and 23c

115 minutes – May 10, 2023

Slides



Episode 17

From SR to Patch – Insights into the Oracle Database Development process


55 minutes – June 22, 2023



NEW Episode 18

Cross Platform Migration – Transportable Tablespaces to the Extreme

145 min – February 22, 2024



Episode 2

AutoUpgrade to Oracle Database 19c


115 minutes – Feb 20, 2021



Episode 3

Performance Stability, Tips and Tricks and Underscores

120 minutes – Mar 4, 2021



Episode 4

Migration to Oracle Multitenant



Recorded Web Seminars

<https://MikeDietrichDE.com/videos>

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



Webinar | Autonomous Database



Scan me to sign up

Prepare for Oracle Autonomous Database

– Everything you need to know about the migration

March 13, 15:00 CET

[Sign up](#)

AGENDA

09:30

Welcome
Release Strategy
Patching

11:15

AutoUpgrade
Multitenant

13:30

Data Pump
Autonomous Database

15:15

Migrating
"The Beast"

11:00

Coffee break

12:45

Lunch


15:00

Coffee break




Releases

Lifetime Support Policy



Lifetime Support Policy



Oracle Database 23ai

Production

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

Development / Testing

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

Patching

Question

How often do you patch your Oracle environments?

☐

Every quarter

☐

Twice per year


☐

Once per year

☐

Never

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 12k fixes with 19.26.0
- More than 300 security fixes

Just as easy as patching your smart phone

AutoUpgrade's mission for patching Oracle Database

AutoUpgrade and Patching



Download

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



Install

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



Patch

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


```
$ cat just_patch.cfg
```

```
patch1.source_home=/u01/app/oracle/product/19/dbhome_19_25_0  
patch1.target_home=/u01/app/oracle/product/19/dbhome_19_26_0  
patch1.sid=DB19
```

AutoUpgrade and Patching



Download

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



Install

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



Patch

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

```
$ cat do_it_all.cfg
```

```
global.keystore=/home/oracle/autoupgrade-patching/keystore  
patch1.source_home=/u01/app/oracle/product/19/dbhome_19_25_0  
patch1.target_home=/u01/app/oracle/product/19/dbhome_19_26_0  
patch1.sid=DB19  
patch1.folder=/home/oracle/autoupgrade-patching/patch  
patch1.patch=RU,OPATCH,OJVM,DPBP
```

```
$ cat do_it_all.cfg
```

```
global.keystore=/home/oracle/autoupgrade-patching/keystore  
patch1.source_home=/u01/app/oracle/product/19/dbhome_19_25_0  
patch1.target_home=/u01/app/oracle/product/19/dbhome_19_26_0  
patch1.sid=DB19  
patch1.folder=/home/oracle/autoupgrade-patching/patch  
patch1.patch=RU,OPATCH,OJVM,DPBP
```


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

Demo


One-Button Patching

- Download patches
- Install Oracle home
- Patch database


Watch on [YouTube](#)



AutoUpgrade Patching




AutoUpgrade Patching




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

AutoUpgrade Patching




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

AutoUpgrade Patching




- Gathers dictionary statistics if needed
- Recompiles Oracle-maintained objects if needed
- Executes checks - see Doc ID [2380601.1](#)

AutoUpgrade Patching




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

AutoUpgrade Patching




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

AutoUpgrade Patching



Patching Modes





Requirements

Current Requirements



Single instance



Oracle Database 19c



Always download
the latest version of AutoUpgrade

- My Oracle Support Doc ID [24854571](#)

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

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



Patching

In-Place Patching

Oracle Home, 19.24.0



`$ORACLE_HOME/OPatch/datapatch`

Out-of-Place Patching

Oracle Home, 19.24.0



```
SQL> SHUTDOWN IMMEDIATE
```



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

New Oracle Home
Oracle Home, 19.25.0



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.

In-Place

55.4%

Out-Of-Place

44.6%



AutoUpgrade patches out-of-place



Which patches should you install?

```
global.keystore=/home/oracle/autoupgrade-patching/keystore  
patch1.source_home=/u01/app/oracle/product/19/dbhome_19_25_0  
patch1.target_home=/u01/app/oracle/product/19/dbhome_19_26_0  
patch1.sid=DB19  
patch1.folder=/home/oracle/autoupgrade-patching/patch  
patch1.patch=RECOMMENDED
```




Recommended Patches

patch1.patch=RECOMMENDED

OPATCH

The latest OPatch

RU

The latest Release Update

OJVM

OJVM bundle matching Release Update

DPBP

Data Pump bundle patch matching RU



You can also add MRPs

- Get patch number from Primary Note for Database Quarterly Release Updates (Doc ID [888.1](#))



Primary Note for Database Quarterly Release Updates (Doc ID 888.1)

APPLIES TO:

Oracle Database - Standard Edition - Version 11.2.0.4 and later
Oracle Cloud Infrastructure - Exadata Cloud Service
Gen 1 Exadata Cloud at Customer (Oracle Exadata Database Cloud Machine)
Gen 2 Exadata Cloud at Customer
Oracle Database Cloud Exadata Service
Information in this document applies to any platform.

PURPOSE

The purpose of this document is to list Database patches for both proactive and reactive maintenance.

SCOPE

This document will be of interest to the DBA and to those responsible for proactive and reactive maintenance of the Oracle Database.

DETAILS

Please note that this document is maintained outside of the standard KM authoring system. Changes made directly to this document may be inadvertently overwritten when the document is next refreshed from outside of the standard KM authoring system. Rather than implementing changes to this document using the standard KM authoring system, please enter a comment as usual, and notify the document owner.[This section is not visible to customers.]

```
$ cat DB19.cfg
```

```
global.keystore=/home/oracle/autoupgrade-patching/keystore  
patch1.source_home=/u01/app/oracle/product/19/dbhome_19_25_0  
patch1.target_home=/u01/app/oracle/product/19/dbhome_19_26_0  
patch1.sid=DB19  
patch1.folder=/home/oracle/autoupgrade-patching/patch  
patch1.patch=RECOMMENDED,37546431
```

```
$ cat DB19.cfg
```

```
global.keystore=/home/oracle/autoupgrade-patching/keystore  
patch1.source_home=/u01/app/oracle/product/19/dbhome_19_25_0  
patch1.target_home=/u01/app/oracle/product/19/dbhome_19_26_0  
patch1.sid=DB19  
patch1.folder=/home/oracle/autoupgrade-patching/patch  
patch1.patch=RECOMMENDED,37546431  
patch1.patch=RECOMMENDED,MRP
```



You can also add specific one-off fixes

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

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

NOTE: This document will be updated once a week on Tuesdays at 9AM EST if any new patch is identified.

Through our review of service requests, we often find that issues encountered are the result of customers being on an older Release Update (RU) or Release Update Revision (RUR). Many issues will have been fixed in the latest updates, which we always recommend.

As noted in the update-specific tables below, fixes for known issues are targeted for inclusion in the first available RU or RUR. You can always find the latest RUs, RURs, other patches, lists of fixed bugs and known issues in [Primary Note for Database Proactive Patch Program\(Doc ID 888.1\)](#). Identifying and installing the latest updates (patches) helps ensure you are using the most current content for security, functional, regression and bug fixes, as well as minor enhancements and any emergency one-offs.

Note:

Beginning with the October 2022 patching cycle, 19c RURs will no longer be provided for 19.17.0 and above. No additional RURs will be delivered on any platform after the delivery of Oracle Database 19c RUR 19.16.2 in January, 2023.

Refer to [Sunsetting of 19c RURs and FAQ \(Doc ID 2898381.1\)](#) for further details.

To provide customers more frequent access to recommended and well-tested collections of patches, Oracle is pleased to introduce Monthly Recommended Patches (MRPs) starting Nov 2022. MRPs are supported only on Linux x86-64 platform.

Refer to [Introducing Monthly Recommended Patches \(MRPs\) and FAQ \(Doc ID 2898740.1\)](#) for further details.

In addition to the relevant patches listed below, you should apply patches based on the specific RU after reviewing the following My Oracle Support knowledge documents:

1. [Database PSU/BP/Update/Revision - Known Issues Primary Note\(Doc ID 1227443.1\)](#)
2. [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.

NOTE: Bug 32781163 is no longer recommended and this note will be updated when a replacement becomes available in the future

NOTE: Bug 31061145 was previously listed in this note but the fix was included in 19.21 in a disabled state. To enable the fix see Note 31061145.8 for instructions.

NOTE: If you are an Applications Unlimited (for example, EBS) customer, you should follow any directions given in your product-specific documentation on applicability of Release Updates (RUs) and Release Update Revisions (RURs).

The information below lists any additional patches (both rolling and non-rolling) that are recommended for installation on top of each RU. Click the relevant link for details.

```
$ cat DB19.cfg
```

```
global.keystore=/home/oracle/autoupgrade-patching/keystore  
patch1.source_home=/u01/app/oracle/product/19/dbhome_19_25_0  
patch1.target_home=/u01/app/oracle/product/19/dbhome_19_26_0  
patch1.sid=DB19  
patch1.folder=/home/oracle/autoupgrade-patching/patch  
patch1.patch=RECOMMENDED,36006910,36908826,35398148,36916250,  
36273767,34672698,34774667,29213893
```




How to download patches

Network Connectivity

The download stage requires:

1. Internet access
2. My Oracle Support credentials

Internet Access

AutoUpgrade connects to:

- `https://updates.oracle.com`
- `https://login-ext.identity.oraclecloud.com`
- `https://aru-akam.oracle.com`

URLs are part of a CDN, so expect changing IP addresses

- Use DNS names instead of IP addresses in your firewall

For connections via proxy use environment variables

- `https_proxy`

MOS Credentials

You must have:

- A valid My Oracle Support credential
- Connected to a Customer Support Identifier (CSI)
- Privilege to download patches using that CSI

Failure to meet the requirements:

- `*Connection Failed - You entered an incorrect user name or password.*`

MOS Credentials

For *download patches privilege* contact an MOS administrator in your organization.

- In *User Details*
- Select the appropriate *Support Identifier*
- Set the *Patches* drop-down to *Download*

Keystore

AutoUpgrade stores MOS credentials in a keystore

- Config file parameter: `global.keystore`
- Governs directory of AutoUpgrade keystore
- Password protected software keystore
- Optionally, an auto-open keystore
- No additional license needed

```
$ java -jar autoupgrade.jar ... -patch -load_password
```

```
$ java -jar autoupgrade.jar ... -patch -load_password
```

...

```
MOS> add -user <MOS username>
```


```
Enter your secret/Password:
```

```
Re-enter your secret/Password:
```




Your database host
doesn't have internet access?

Using Download Mode





You can download patches manually and place them in the *patch* folder

- Use a shared folder accessible to all hosts

In The Plans



We're working to support

- Oracle Database 23ai
- Cloning Oracle homes
- Gold images

Hands-on Lab

Patch Me If You Can

It's better to fail in our lab, than in production



[Access lab on Oracle Live Labs](#)

Datapatch

What is Datapatch?

`$ORACLE_HOME/OPatch/datapatch`

↳ `$ORACLE_HOME/sqlpatch/sqlpatch`

↳ `$ORACLE_HOME/sqlpatch/sqlpatch.pl`

```
use strict;  
use Getopt::Long;  
  
use sqlpatch.pm
```



Datapatch uses OPatch to determine patches installed in Oracle home

- From the database via **DBMS_QOPATCH**

ORA-20001: Latest xml inventory is not loaded into table

--Avoid the use of Queryable Inventory by using OPatch directly
--Safe for single instance databases. In RAC, it's vital you ensure
--all nodes are patched and the inventories on each node is similar.

./datapatch -noqi

./datapatch -local_inventory

Patch Apply Sequence

datapatch



1



Java patches

2



Bundle patches

3



One-off patches

Patch Rollback and Apply Queue

Binary Registry after opatch:

Patch 444 – Java Patch

Patch 555 – Bundle Patch

Patch 666 – One-off Patch

SQL Registry before datapatch:

Patch 111 – Java Patch

Patch 222 – Bundle Patch

Patch 333 – One-off Patch

```
$ ./datapatch
```

datapatch queue

Rollback:

Apply:

Rollback:

Cumulative:

Patch 222 to 555 – Bundle Patch

Apply:



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

Rollback Scripts



Apply/rollback scripts

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

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



Rollback scripts (zipped as BLOB):

One-off patches:

```
SELECT PATCH_DIRECTORY  
FROM   REGISTRY$SQLPATCH
```

Release Updates:

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

PATCH_ID	RU_VERSION	RU_BUILD_TS	LOB_SIZE_MB
29517242	19.3.0.0.0	04/10/2019 12:27:20	4
30125133	19.5.0.0.0	09/09/2019 18:05:49	18
30557433	19.6.0.0.0	12/17/2019 15:50:04	24
30869156	19.7.0.0.0	04/04/2020 03:50:18	30
31281355	19.8.0.0.0	07/03/2020 03:15:01	37
31771877	19.9.0.0.0	09/30/2020 18:32:49	43
32218454	19.10.0.0.0	01/08/2021 18:50:17	51
32545013	19.11.0.0.0	04/13/2021 00:40:09	59
32904851	19.12.0.0.0	07/16/2021 14:18:10	67
33192793	19.13.0.0.0	10/04/2021 16:50:50	75
33515361	19.14.0.0.0	12/25/2021 12:21:23	83
34133642	19.16.0.0.0	07/03/2022 02:22:23	99
34419443	19.17.0.0.0	09/24/2022 22:40:51	108
34765931	19.18.0.0.0	01/11/2023 17:17:38	116
35643107	19.21.0.0.0	09/30/2023 15:19:51	141

15 rows selected.

CON_ID	LOB_SIZE_MB
1	953
2	953
3	953
...	...
4098	953


Cleanup

Bug 37352486

PURGE OLD PATCH METADATA ZIP BLOB FROM DATAPATCH INVENTORY, IF NOT NEEDED

- Apply one-off patch or use 19.26.0 or newer
- Run `./datapatch -purge_old_metadata`
- This will execute the cleanup
- Be aware that the cleanup may take a bit in large CDB environments which have been patched frequently

Multitenant



- Datapatch starts with CDB\$ROOT
- Datapatch sorts PDBs by *priority* and *con_id*
 - Set priority using **ALTER PLUGGABLE DATABASE ... PRIORITY**
- Datapatch determines parallel degree based on **CPU_COUNT**



Datapatch patches *PDB\$SEED* automatically

- New PDBs are ready to go
- No need to execute Datapatch on new PDBs



Datapatch only patches open PDBs

- READ WRITE, READ ONLY, or UPGRADE



Unpatched PDBs will open in **RESTRICTED** mode

- Applies to failed or incomplete Datapatch runs
- Reported as plug-in violation

--Prevent PDBs from opening in restricted mode after a failed/incomplete
--Datapatch run. Quickly resolve the patching issue and re-run Datapatch.
--Use with caution.

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

Advanced Patching



Reduce downtime to the time it takes to perform a switchover

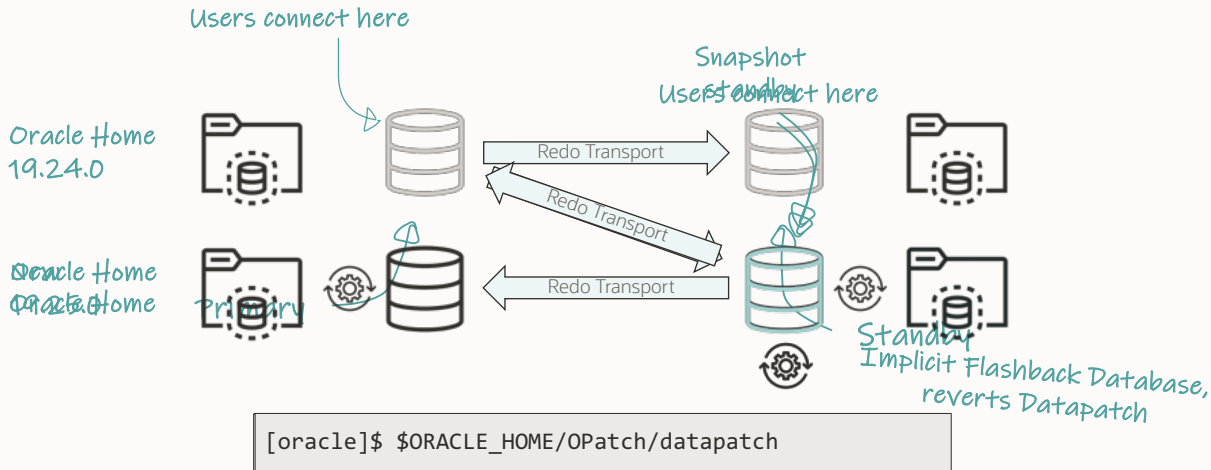
- [Data Guard Standby-First Patch Apply \(Doc ID 1265700.1\)](#)



Safely test and verify patches with Standby-First Patch Apply

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

Standby-First Patching





Patch must Standby-First installable

- Check the patch readme




Execute Datapatch on the primary database

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



Avoid database downtime with
RAC Rolling Patch Apply

RAC Rolling Patching



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



Keep GI and DB patch levels in sync

- This is what we test and run in our Cloud



GI and DB Patch Levels

Method	Status	Example
In sync	Recommended, best practice	GI 19.26 + DB 19.26
Within two Release Updates	Supported, good practice	GI 19.26 + DB 19.24 GI 19.24 + DB 19.26
Within three or more Release Updates	Supported, not recommended	GI 19.26 + DB 19.23 GI 19.23 + DB 19.26
Different releases	Supported, use only for upgrades	GI 23.7 + DB 19.26



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

- [RAC: Frequently Asked Questions \(Doc ID 220970.1\)](#)



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

- 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

Release updates are **always**:



Standby-First installable



RAC Rolling installable

Key Learnings



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

Break

We start again at 11:20



AutoUpgrade



Overview and new features





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

- DBUA is **desupported**


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

```
build.version 25.1.250212
```

```
build.date 2025/02/12 18:51:30 +0000
```

```
build.hash 24fd3b4f8
```

```
build.hash_date 2025/02/07 21:18:23 +0000
```

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

```
build.type production
```

```
build.label (HEAD, tag: v25.1)
```

```
build.MOS_NOTE 2485457.1
```

```
build.MOS_LINK https://support.oracle.com/.../?id=2485457.1
```

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

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

New Features

--Upgrade RMAN catalog after upgrade

--<https://dohdatabase.com/autoupgrade-new-features-upgrade-rman-catalog-schema>

```
$ cat CDB1.cfg
```

```
upg1.source_home=/u01/app/oracle/product/19/dbhome_19_26_0
```

```
upg1.target_home=/u01/app/oracle/product/13/dbhome_23_7_0
```

```
upg1.sid=CDB1
```

```
upg1.rman_catalog_connect_string=catalogdb
```

--Update OEM configuration after upgrade

--<https://dohdatabase.com/autoupgrade-new-features-update-enterprise-manager-configuration>

```
$ cat CDB1.cfg
```

```
upg1.source_home=/u01/app/oracle/product/19/dbhome_19_26_0
```

```
upg1.target_home=/u01/app/oracle/product/13/dbhome_23_7_0
```

```
upg1.sid=CDB1
```

```
upg1.upg1.emcli_path=/u01/app/oracle/oem
```

```
upg1.em_target_name=ORCL_myhost.domain.int
```



Upgrading and patching on Windows

Windows

- The Oracle Database instance runs as a Windows service
- Upgrading and out-of-place patching requires recreation of the Windows service using `oradim.exe`
- Windows supports passing the credentials via a *credential file*


```
$ cat patch_on_win.cfg
```

```
global.keystore=c:\oracle\autoupgrade\keystore  
patch1.source_home=c:\oracle\product\dbhome_19_25_0  
patch1.target_home=c:\oracle\product\dbhome_19_26_0  
patch1.sid=DB19  
patch1.folder=c:\oracle\patches  
patch1.patch=RECOMMENDED  
patch1.wincredential=c:\oracle\autoupgrade\credential
```




```
C:\> java -jar autoupgrade.jar  
-config patch_on_win.cfg  
-patch  
-loadwincredential "patch1"
```

```
C:\> java -jar autoupgrade.jar  
-config patch_on_win.cfg  
-patch  
-loadwincredential "patch1"
```

```
C:\> java -jar autoupgrade.jar  
-config patch_on_win.cfg  
-patch  
-mode deploy
```

Windows



--AutoUpgrade deletes the credential file after use for security reasons
--For automation it is desirable to keep the credential file

patch1.delete_credential_file=false



AutoUpgrade fully supports upgrading
Oracle RAC Database on Windows

Refreshable Clone PDBs

Non-CDB to PDB conversion is irreversible

What are your rollback options?

ROLLBACK



Backup / restore

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



Copy data files

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

ROLLBACK



Backup / restore

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



Copy data files

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



Refreshable clone PDB

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

Requires Oracle Database 12.2 or newer

Refreshable Clone PDB



CREATE

Create PDB from non-CDB over a database link



REFRESH

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



OUTAGE

Disconnect users and refresh PDB for the last time




CONVERT


To become a proper PDB, it must be converted

Refreshable Clone PDB

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



Source
non-CDB



Target
CDB

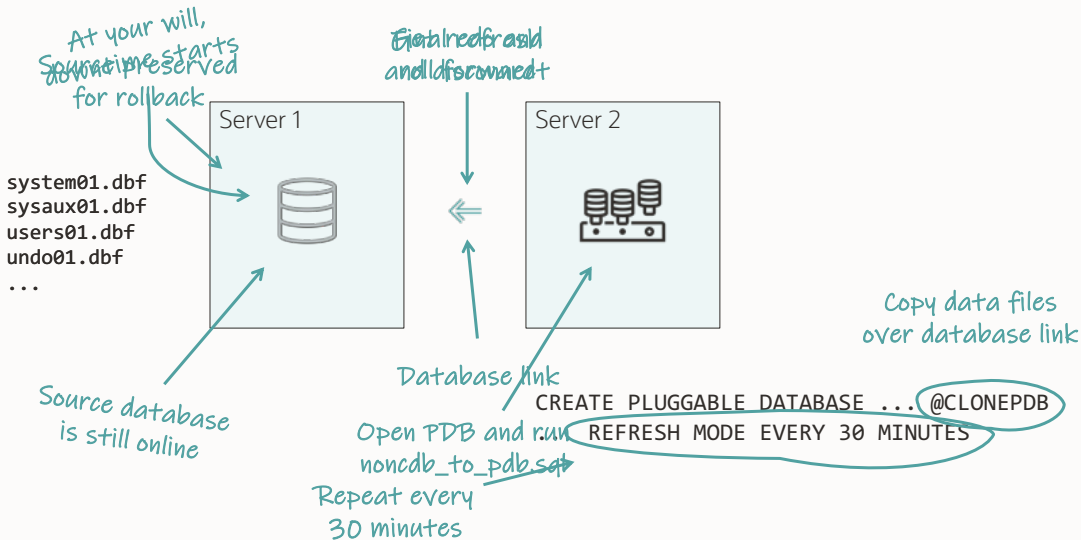
Refreshable Clone PDB

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



*Could be same
server as well*

Refreshable Clone



Refreshable Clone PDB

Source non-CDB

Target CDB



```
CREATE USER dblinkuser  
  IDENTIFIED BY ... ;  
  
GRANT CREATE SESSION,  
  CREATE PLUGGABLE DATABASE,  
  SELECT_CATALOG_ROLE TO dblinkuser;  
  
GRANT READ ON sys.enc$ TO dblinkuser;
```

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

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
```



You can drop user and database link after migration

- Config file parameter **drop_dblink**

Refreshable Clone PDB

Source non-CDB

Target CDB



```
upg1.source_home=/u01/app/oracle/product/12.2.0.1
upg1.target_home=/u01/app/oracle/product/19
upg1.sid=NONCDB1
upg1.target_cdb=CDB1
upg1.source_dblink.NONCDB1=CLONEPDB 300
upg1.target_pdb_name.NONCDB1=PDB1
```

Refreshable Clone

Source non-CDB

Target CDB



```
upg1.source_home=/u01/app/oracle/product/12.2.0.1
upg1.target_home=/u01/app/oracle/product/19
upg1.sid=NONCDB1
upg1.target_cdb=CDB1
upg1.source_dblink.NONCDB1=CLONEPDB 300
upg1.target_pdb_name.NONCDB1=PDB1
upg1.start_time=25/01/2025 02:00:00
--Specify relative start time
--upg1.start_time=+1h30m
```

Refreshable Clone PDB

Source non-CDB

Target CDB



```
upg1.source_home=/u01/app/oracle/product/12.2.0.1
upg1.target_home=/u01/app/oracle/product/19
upg1.sid=NONCDB1
upg1.target_cdb=CDB1
upg1.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
```



Refreshable Clone

1

Run on source

```
autoupgrade.jar ... -mode analyze
```

```
autoupgrade.jar ... -mode fixups
```

2

Run on target

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



Refreshable Clone PDB

1.

PDB
is created

2.

Data files
are copied

3.

Redo is
applied

4.

Final refresh

5.

Disconnect
and convert

autoupgrade.jar ... -mode deploy

upg1.start_time=25/01/2025 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 02/03/2025 02:30:00



Works for unplug-plug upgrades as well



The source non-CDB stays intact
to allow rollback



Zürcher
Kantonalbank

Customer Case | Zürcher Kantonalbank

Customer

Project

Constraints

Preparation

Migration

Success?

Remarks

A reliable partner for over 150 years

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



Customer Case | Zürcher Kantonalbank

Customer

Project

Constraints

Preparation

Migration

Success?

Remarks

Current situation

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

Planned solution: AutoUpgrade

Customer Case | Zürcher Kantonalbank

Customer

Project

Constraints

Preparation

Migration

Success?

Remarks

Test setup

- 3 non-CDB databases of different size

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

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

Customer Case | Zürcher Kantonalbank

Customer

Project

Constraints

Preparation

Migration

Success?

Remarks

Cloning user

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

Permissions

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

Database link

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

Customer Case | Zürcher Kantonalbank

Customer

Project

Constraints

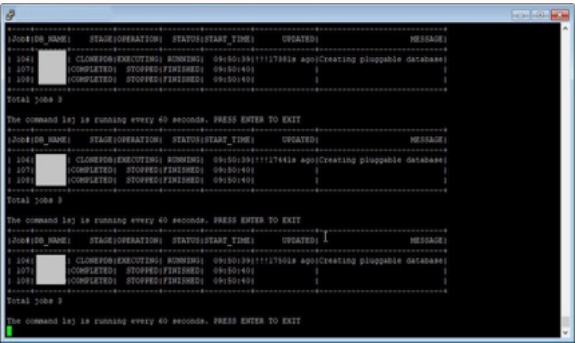
Preparation

Migration

Success?

Remarks

Migration in progress



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



Customer Case | Zürcher Kantonalbank

Customer

Project

Constraints

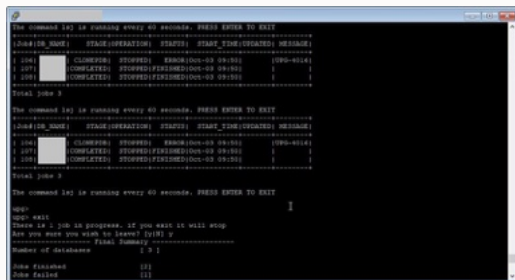
Preparation

Migration

Success?

Remarks

Migration completed



```
The command is running every 40 seconds. PRESS ENTER TO EXIT
=====
(Schema) STAGE(OPERATION) STATUS START TIME(UPDATED) MESSAGE
=====
[ 106 ] [ CLONED ] STOPPED ERROR(Oct-03 09:00) (999-0010)
[ 107 ] [ COMPLETED ] STOPPED(FINISHED) Oct-03 09:00 [ ]
[ 108 ] [ COMPLETED ] STOPPED(FINISHED) Oct-03 09:00 [ ]
=====
Total jobs 3

The command is running every 40 seconds. PRESS ENTER TO EXIT
=====
(Schema) STAGE(OPERATION) STATUS START TIME(UPDATED) MESSAGE
=====
[ 106 ] [ CLONED ] STOPPED ERROR(Oct-03 09:00) (999-0010)
[ 107 ] [ COMPLETED ] STOPPED(FINISHED) Oct-03 09:00 [ ]
[ 108 ] [ COMPLETED ] STOPPED(FINISHED) Oct-03 09:00 [ ]
=====
Total jobs 3

The command is running every 40 seconds. PRESS ENTER TO EXIT

^C^C
^C^C 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 ]
Jobs 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 archive logs from backup

User profile with IDLE_TIME lead to kill of the session

- Solution: assign a different profile to the clone user

Summary

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



Christian Marquardt
Database Specialist Expert & Solution Architect



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

- See [blog post](#) for details



Refreshable clone works only with deferred recovery on standby database

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

Multitenant

Create Container Database



1 Character set

2 Components

3 COMPATIBLE


Create Container Database

1 Character set

- Always choose AL32UTF8
- Allows PDBs with any character set

2 Components

3 COMPATIBLE




Create Container Database

1 Character set

2 Components

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

3 COMPATIBLE




Create Container Database

1 Character set


2 Components

3 COMPATIBLE

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

All initialization parameters			
 Update the initialization parameters only when it is required. Refer to the Oracle documentation to learn more about each initialization parameter and its valid set of values.			
(Storage related parameter(s) value is shown in MB) <input type="checkbox"/> Show advanced parameters			
Name	Value	Include in spfile	Category
undo_tablespace	UNDOTBS1	<input checked="" type="checkbox"/>	Cluster Database
sga_target	2379	<input checked="" type="checkbox"/>	SGA Memory
db_block_size (bytes)	8192	<input checked="" type="checkbox"/>	Cache and I/O
nls_language	AMERICAN	<input checked="" type="checkbox"/>	NLS
control_files	{~{ORACLE_BASE}}/oradata/...	<input checked="" type="checkbox"/>	File Configuration
remote_login_passwordfile	EXCLUSIVE	<input checked="" type="checkbox"/>	Security and Auditing
processes	320	<input checked="" type="checkbox"/>	Processes and Sessions
pga_aggregate_target	793	<input checked="" type="checkbox"/>	Sort, Hash Joins, Bitmap Indexes
nls_territory	AMERICA	<input checked="" type="checkbox"/>	NLS
open_cursors	300	<input checked="" type="checkbox"/>	Cursors and Library Cache
compatible	23.0.0	<input checked="" type="checkbox"/>	Miscellaneous
db_name	orcl	<input checked="" type="checkbox"/>	Database Identification
cluster_database	FALSE	<input type="checkbox"/>	Cluster Database
Description:			
compatible: Allows you to use a new release, while at the same time guaranteeing backward compatibility with an earlier release. Range of Values: Default to current release. Default Value: Release dependent			

Compatible





How does a standby database handles a plug-in operation

- Take care of your Data Guard during migration

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

```
SQL> select dbms_pdb.describe('/tmp/manifest_DB.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/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



23ai
CDB
Standby



+DATA/DB_BOSTON/DATAFILE/users.273.1103046827

Redo record says:
Plug in this data file

No good, data file
has a different name

+DATA/DB_CHICAGO/DATAFILE/users.269.1103050009

Enabled Recovery | ASM

23ai
CDB
Primary



+DATA/DB_BOSTON/DATAFILE/users.273.1103046827



23ai
CDB
Standby



OK, let's check the OMF directory

+DATA/DB_CHICAGO/DATAFILE/users.269.1103050009

+DATA/CDB1_CHICAGO/<PDB_GUID>/DATAFILE

It's empty

Enabled Recovery | ASM

23ai
CDB
Primary



+DATA/DB_BOSTON/DATAFILE/users.273.1103046827



23ai
CDB
Standby



OK, let's check the OMF directory

+DATA/DB_CHICAGO/DATAFILE/users.269.1103050009

+DATA/CDB1_CHICAGO/<PDB_GUID>/DATAFILE

It's empty



I'll just move the file in ASM



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

- Similar to file system soft links

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

Non-CDB standby data file



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



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



Standby PDB OMF location

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



Name does not matter.
Standby scans all files in OMF directory

Data Guard | Re-use Data Files

23ai
CDB
Primary



23ai
CDB
Standby



- Standby database scans its own OMF directory for data files
- Standby ignores file names and look at file headers
- Standby will find aliases and find the real file locations

Data Guard | Re-use Data Files

Primary



Standby



Looking for file like on primary



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



Follows alias and finds the real file



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**



What about AutoUpgrade
and enabled recovery?

Enabled Recovery | AutoUpgrade

The current version (24.1) does not support plugging in with enabled recovery

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



Deferred Recovery

1

Enabled recovery

`create pluggable database ... standbys=all`

Standby records PDB creation

Standby locates data files

MRP applies redo to PDB

PDB is immediately protected

2

Deferred recovery

`create pluggable database ... standbys=none`


Standby records PDB creation

Standby ignores data files

MRP skips redo

PDB protected after restore

Deferred Recovery



Deferred Recovery



Source
Non-CDB



Target
Primary

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



Target
Standby

Deferred Recovery



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



PDB created
Data files missing

Deferred Recovery



```
SQL> show pdbs
```

CON_NAME	OPEN	MODE
PDB1	READ	WRITE

```
SQL> show pdbs
```

CON_NAME	OPEN	MODE
PDB1	MOUNTED	

Deferred Recovery



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

NAME	RECOVERY_STATUS
PDB1	DISABLED

Deferred Recovery



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

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

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

Deferred Recovery



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

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

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

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

Further Details

Webinar:

[Move to Oracle Database 23ai –
Everything you need to know about Oracle Multitenant – Part 1](#)



Don't jeopardize your Data Guard

- Test the procedure and verify your environment



Operations



Share resources between PDBs

Resource Consolidation

Non-CDB
database



Memory



Background processes




Files



Resource Consolidation

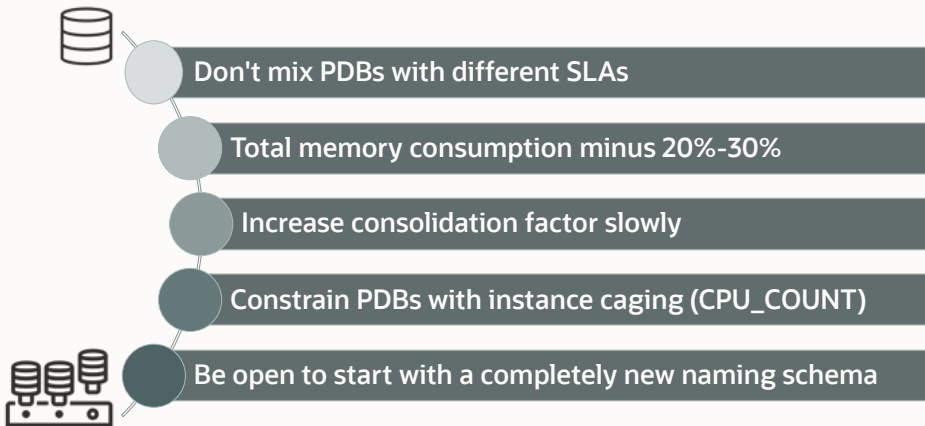


Resource Consolidation



Consolidation Strategies?

There is no "*best*" strategy



Using a Swingbench benchmark,
a single-core machine could host **nine non-CDBs**
before reaching 75 % CPU utilization

By going multitenant the number of databases reached **123 PDBs**

A US Health Care provider managed to

- Reduce the number of database instances by 7x
- Reduce the number of physical servers by 50 %



You can run multiple CDBs on the same host and out of the same Oracle home

Consolidation



Schema consolidation



Virtual Private Database



PDB consolidation

- Less complexity
- Better isolation
- Operational benefits
- Easier cloning

A global provider of financial services states

*The multitenant architecture gives us **complete client separation out of the box**, without having to maintain a Virtual Private Database setup.*

We went away from Virtual Private Database and consolidated our different clients in individual PDBs.

*This reduced the complexity of our database implementation and **made operations much easier**.*



Avoid noisy neighbors

- Allow sharing resources
but everyone must get a fair share

Method 1



Instance caging

- Most simple
- Define `CPU_COUNT` for each PDB
- Hard limit

Method 1



8 CPUs



CPU_COUNT=3



CPU_COUNT=2



CPU_COUNT=2



All non-CDBs
share 7 CPU

Method 1



8 CPUs



CPU_COUNT=7

CDB never uses more than 7 CPUs,
despite sum of PDBs
At peak, use more resources,
but never deplete the CDB



CPU_COUNT=4



CPU_COUNT=4



CPU_COUNT=4

PDBs might fight over CPUs,
but each process gets a fair share

Method 2



Memory allocation

- Simple
- Define **SGA_TARGET** for each PDB
- Hard limit

Method 2



8 GB memory



SGA_TARGET=7G



SGA_TARGET=4G



SGA_TARGET=4G



SGA_TARGET=4G

PDB may never use more than 4G of shared memory



If all PDBs are active, cache management comes into play



Requires use of Automatic Shared Memory Management

- Both in CDB and PDB



Optionally, allocate minimum shared pool and buffer cache for a PDB

- Use `SHARED_POOL_SIZE` and `DB_CACHE_SIZE`



You can combine method 1 and 2

- Instance caging and memory allocation

Method 3



Simple Resource Manager

- Elaborate, yet simple to implement
- Enable CDB resource manager
- Allocate minimum shares instead of hard limits
- For advanced use cases

Method 3



8 CPUs



CPU_COUNT=7



CPU_MIN_COUNT=2



CPU_MIN_COUNT=1



CPU_MIN_COUNT=1

At peak, may use up to 5 CPUs



4 CPUs are reserved,
3 are free for all

Method 3



8 GB memory



SGA_TARGET=7G



SGA_MIN_SIZE=2G



SGA_MIN_SIZE=1G



SGA_MIN_SIZE=1G

*At peak, may use
up to 4G shared memory*





Requires Resource Manager at root level

```
alter session set container=cdb$root;
```

```
-- Create an empty resource manager plan with no directives
```

```
exec dbms_resource_manager.clear_pending_area;
```

```
exec dbms_resource_manager.create_pending_area;
```

```
exec dbms_resource_manager.create_cdb_plan('CDB_PLAN');
```

```
exec dbms_resource_manager.validate_pending_area;
```

```
exec dbms_resource_manager.submit_pending_area;
```

```
-- Make plan active in root to enable CDB resource manager
```

```
alter system set resource_manager_plan=CDB_PLAN;
```

Method 4

Advanced Resource Manager



- Requires additional configuration, but much greater control
- Use directives instead of shares



You can still control resources inside a PDB
with Resource Manager



What about I/O?

- Exadata I/O Resource Management
- Or, **MAX_MBPS** and **MAX_IOPS**



You can run multiple CDBs on the same host and out of the same Oracle home

Inter-instance Resource Management

Shares resources like with non-CDBs:

- CPU_COUNT
- SGA_MAX_SIZE

Inter-instance CPU resource manager:


- Controls resource sharing using Linux c-groups
- Check [RESOURCE MANAGER CPU SCOPE](#)
- Exadata Database Machine and Autonomous Database



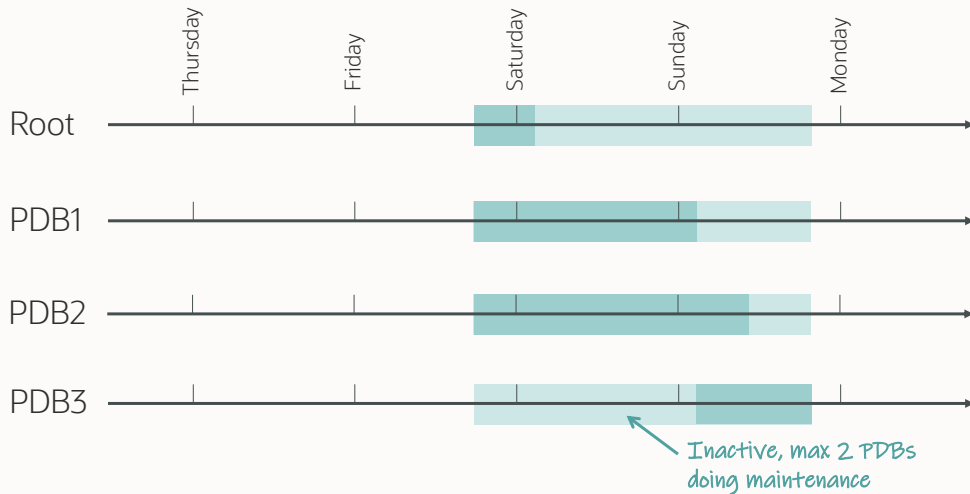


A word about automated maintenance tasks

Automated Maintenance Tasks




Automated Maintenance Tasks



--Change the amount of PDBs that can run maintenance tasks at the same time
--Default value 2


```
alter system set autotask_max_active_pdb=3;
```

Automated Maintenance Tasks



```
--Selectively disable maintenance tasks in a PDB  
--For instance, test databases or databases that are rebuilt frequently  
  
alter session set container=PDB2;  
alter system set enable_automatic_maintenance_pdb=false;
```

Automated Maintenance Tasks







Shift maintenance windows

- Optionally, shorten maintenance windows

Automated Maintenance Tasks



Automated Maintenance Tasks





Selectively disable individual maintenance tasks using **DBMS_AUTO_TASK_ADMIN**

- Does a test database need Automatic Segment Advisor?
- Or Evolve Advisor?



Resource Manager prevents maintenance tasks from *stealing* resources from users

- Consumer group `ORA$AUTOTASK`

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

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

Key Learnings



- 1 Get started with Multitenant
- 2 Understand plug-in on standby databases
- 3 Finetune your container database

Lunch Break

We start again at 13:30

Data Pump Top Tips

The background features a pattern of light gray concentric circles. In the top right corner, there is a blue circle with horizontal lines and a dark gray circle. In the bottom right corner, there is a red circle with a grid pattern and a green circle. A small red square with a white 'O' logo is located in the bottom right corner.

Supercharge data loading/unloading



Always use
Data Pump Bundle Patch





More than 200 functional and performance fixes

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

*Importing a complete application with data drops from almost 2.5 hours to **48 minutes** – by just applying the Data Pump bundle patch*

A global provider of financial services



Always ensure dictionary
and fixed objects statistics are accurate





Ensure dictionary and fixed objects statistics are accurate

- Before export
- Before import
- Immediately after import

```
begin
```

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

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

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

```
dbms_stats.gather_fixed_objects_stats;
```

```
end;
```

```
/
```

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

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



Use parallel and
multiple dump files



--Apply parallelism by simply specifying a degree
expdp ... parallel=8

--Use different parallel degree on import
impdp ... parallel=32



Oracle Cloud Infrastructure

Number of OCPUs

Number of ECPUs / 4



On-prem (x86-64)

2 x physical cores



On-prem (other)

Depends

--Use %L to allow multiple dump files
expdp ... parallel=8 dumpfile=exp%L.dmp

--Split dump files into minor files for easier transport
expdp ... parallel=8 dumpfile=exp%L.dmp **filesize=10G**

- After export, store a checksum in the dump file.
- Detects in-flight corruption or alteration.
- Specify other algorithms using checksum_algorithm parameter.

```
expdp ... checksum=yes
```

```
impdp ... verify_checksum=yes  
         verify_only=yes
```



Transportable jobs can use parallel
in Oracle Database 21c


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

```
expdp ... full=y transportable=always parallel=16
```

```
expdp ... tablespace=<list> parallel=16
```

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

Parallel Transportable | Benchmark

Oracle E-Business Suite database

600.000+ objects

Export parallel 1 2h 2m

Import parallel 1 6h 44m

Total 8h 46m

Export parallel 16 1h 8m

Import parallel 16 1h 23m

Total 2h 31m



Always convert to
SecureFile LOBs





You get the fastest LOB operations
with **SecureFile** LOBs

2007

Oracle Database 11g Release 1

-- Do you still have any old BasicFile LOBs in your database?

```
select * from dba_lobs where securefile='NO';
```



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

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

*By applying the Data Pump Bundle Patch
our 4.3 TB export with huge LOBs went
from over **21 hours to 3 hours 22 minutes***

A European government agency

*... Plus, by increasing parallel
from 4 to 12 the export dropped
to 1 hour 51 minutes*

A European government agency

*... Finally, we moved
to faster ASM based storage
bringing it to 1 hour 7 minutes*

A European government agency



Do you still have BasicFile LOBs?

- Use [DIY parallelism](#) during export

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

`impdp ... transform=lob_storage:securefile`

Importing as BasicFile LOBs

```
... imported "SCHEMA"."TABLE" 31.83 GB 681025 rows in 804 seconds using direct_path
```

Importing as SecureFile LOBs

```
... imported "SCHEMA"."TABLE" 31.83 GB 681025 rows in 261 seconds using external_table
```



Do you still have **LONG** and **LONG RAW**?

- Deprecated since Oracle8i

```
-- Convert LONG to CLOB, and LONG RAW to BLOB on import  
-- Be sure to change your application as well,  
-- PL/SQL interface for accessing LOBs and LONGs are not the same
```

```
impdp ... transform=long_to_lob:y
```



Speed up imports by using NOVALIDATE constraints





A Constraint Can Be

VALIDATED

All data in the table obeys the constraint.
The database guarantees that data is good.

NOT VALIDATED

All data in the table **may** obey the constraint.
The database **does not know** if data is good.



Most constraints are **VALIDATED**



On import, Data Pump creates constraints
in the same state as in the source

--Example of which commands Data Pump import might execute as part of an import

```
create table sales ( .... );
```

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

```
alter table sales add constraint c_sales_1 check (c1 in (0,1)) enable validate;  
alter table sales add constraint c_sales_2 check (c2 in ('A','B')) enable validate;  
alter table sales add constraint c_sales_3 check (c3 > 0) enable validate;
```

Recursive full table scan

Recursive full table scan

Recursive full table scan

```
-- Add constraints with NOVALIDATE keyword regardless of state in source database  
-- Significantly speeds up add constraints for larger tables
```

```
impdp ... transform=constr_novalidate
```

--Transforming constraints to NOVALIDATE to speed up import

```
alter table sales add constraint c_sales_1 check (c1 in (0,1)) enable novalidate;  
alter table sales add constraint c_sales_2 check (c2 in ('A','B')) enable novalidate;  
alter table sales add constraint c_sales_3 check (c3 > 0) enable novalidate;
```


No full table scan

--Transforming constraints to NOVALIDATE to speed up import

```
alter table sales add constraint c_sales_1 check (c1 in (0,1)) enable novalidate;  
alter table sales add constraint c_sales_2 check (c2 in ('A','B')) enable novalidate;  
alter table sales add constraint c_sales_3 check (c3 > 0) enable novalidate;
```

Database validates new rows

Benchmark, 1 billion rows

Importing VALIDATE constraints

```
10-AUG-24 00:32:28.716: W-1 Processing object type TABLE_EXPORT/TABLE/TABLE_DATA
10-AUG-24 00:36:42.762: W-1 . . imported "FUSION"."hwr_topic_t1" 151.2 GB 1044625000 rows in 254 seconds using external_table
10-AUG-24 00:45:41.226: W-1 Processing object type TABLE_EXPORT/TABLE/CONSTRAINT/CONSTRAINT
10-AUG-24 00:55:35.787: W-1      Completed 7 CONSTRAINT objects in 594 seconds
```

Importing NOVALIDATE constraints

```
10-AUG-24 00:14:56.050: W-1 Processing object type TABLE_EXPORT/TABLE/TABLE_DATA
10-AUG-24 00:19:10.311: W-1 . . imported "FUSION"."hwr_topic_t1" 151.2 GB 1044625000 rows in 254 seconds using external_table
10-AUG-24 00:29:20.841: W-1 Processing object type TABLE_EXPORT/TABLE/CONSTRAINT/CONSTRAINT
10-AUG-24 00:29:21.101: W-1      Completed 7 CONSTRAINT objects in 1 seconds
```



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

- Check QUERY REWRITE INTEGRITY



Validate constraints after import, or even **after go-live**

- Still requires a full scan of the table
- But can use parallel query
- And **no** table lock!

Exceptions

Data Pump always validates certain constraints:

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



Use with care if
you are transforming data on import



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

- Plus patch 37280692 - or be on 19.27

Even faster index imports





Use index size to determine parallel degree on index creation

- Coming in future 23ai Data Pump Bundle Patch

Index Creation

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

Before 12.1

Worker 1 `CREATE INDEX PARALLEL 16`

Really good for few big indexes



Index Creation

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

From 12.1

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

Really good for many small indexes



Index Creation

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

From 23

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

The best of both worlds





How Data Pump Create Indexes

- 1 Calculate the optimal parallel degree
- 2 Create indexes



How Data Pump Create Indexes

1 Calculate the optimal parallel degree

- Always parallel 1 when a table is less than 150 MB
- Customizable via `INDEX_THRESHOLD`
- Get optimal parallel degree using `EXPLAIN PLAN`

```
SQL> explain plan for create index i1 on t1(c1) parallel;
```

Explained.

```
SQL> explain plan for create index i1 on t1(c1) parallel;
```

Explained.

```
SQL> select * from table(dbms_xplan.display(format => 'ALL'));
```

...

Note

- automatic DOP: Computed Degree of Parallelism is 4 because of degree limit
- estimated index size: 655K bytes




How Data Pump Creates Indexes

2 Create indexes

- One worker creates small indexes (parallel 1) in large batches
- The next worker starts with the biggest index (measured by optimal parallel degree)

How Data Pump Creates Indexes

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





Benchmark, 1 billion rows

Importing with 19c settings constraints

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

Importing NOVALIDATE constraints + new index method

10-AUG-24 01:48:38.844: Job "SYSTEM"."SYS_IMPORT_TABLE_01" successfully completed at Sat Aug 10 01:48:38 2024 elapsed 0 00:10:40



We expect much better result
with more complex schemas



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

- Planned for future Data Pump Bundle Patch

Key Learnings



- 1 Use Data Pump Bundle Patch
- 2 Use parallel
- 3 Gather stats

Autonomous Database



A migration approach



The idea?

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

by **LESDN**

OOW 2017


ORACLE®





Will the DBA's be fired?

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



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

Tim Hall

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

Autonomous Database – Where?

Public cloud

Autonomous Database

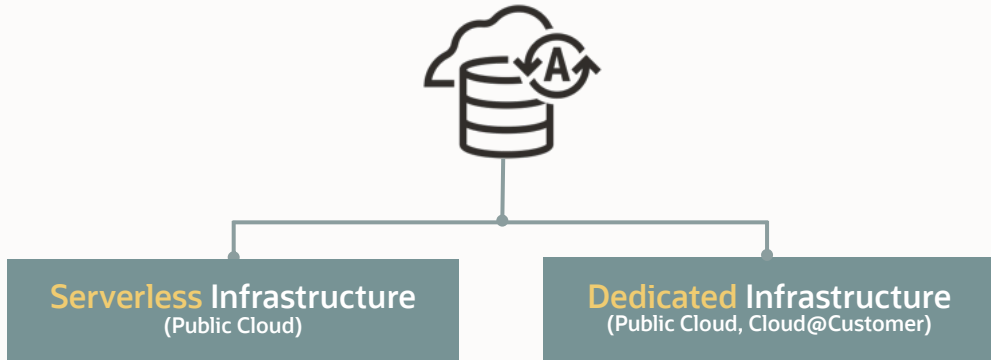


Cloud@Customer

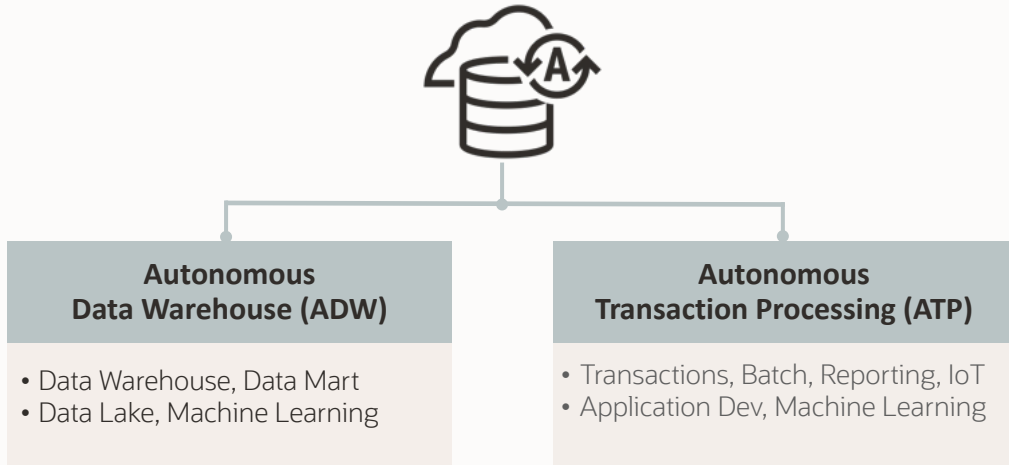
Autonomous Database in a VM environment



One Autonomous Database – Two Deployment Choices

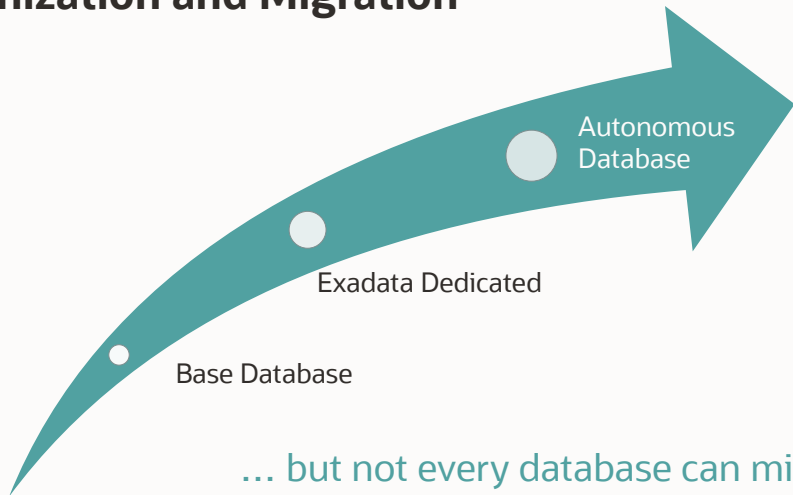


One Autonomous Database – Workload Choices



Migration Planning

Modernization and Migration





Migration to Autonomous Database is always a **logical** migration

- Move the data, not the database
- Spoiler: There may be another option soon

<https://www.oracle.com/database/upgrades/#adb-migration>

Oracle Database
AutoUpgrade utility

Cloud Migration Advisor
(CMA)

**Autonomous Database
Migration**

Database migration
resources

Load data into Oracle
Database

Hands-On Lab Training

Move to Autonomous Database

1. Estate Explorer

Identify and map potential targets for the migration to Oracle ADB.

> [Learn more about Estate Explorer](#)

2. Cloud Pre-Migration Advisor Tool (CPAT)

CPAT analyzes your source Oracle Database metadata for compatibility with your intended Oracle Autonomous Database target and migration method. It suggests a course of action to resolve each potential incompatibility it finds in the source.

> [Learn more about CPAT](#)

3. Cloud Migration Advisor (CMA)

CMA analyzes the CPAT report from your source database and provides detailed technical advice about likely migration targets and methods.

> [Learn more about CMA](#)

The tool gives you technical advice on the various migration options:

- Database Migration Service (DMS)
- Zero Downtime Migration (ZDM)
- Oracle Autonomous DB Migration Automation (In beta status currently, contact us).

Move to Autonomous Database



Estate Explorer

Identify and map potential targets for the migration to Oracle ADB.

[Click for more information](#)




Source Database	
Action Required	4
Review Required	0
Review Suggested	2
Passed	7

CPAT

Cloud Pre-Migration Advisor Tool – download from MOS Note: [2758371.1](#)

[Click for more information](#)



CMA

Cloud Migration Advisor, GUI landscape overview.

The tool gives you technical advice on the various migration options.

[Click for more information](#)



DMS

Database Migration Service

[Click for more information](#)

ZDM

Zero Downtime Migration

[Click for more information](#)

AMA

Oracle **Autonomous** DB Migration Automation

In beta status currently.

[Contact me](#)

Oracle Estate Explorer



Analyze 1000's of databases in just a few hours



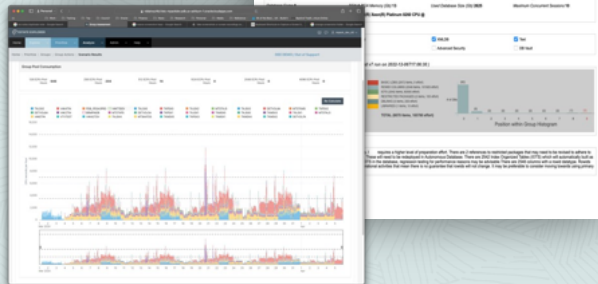

View innovative visualizations and detailed reports



Provide a detailed TCO to compare on-premises and cloud



Optimize your Autonomous Databases using Elastic Pools



<https://www.oracle.com/database/cloud-migration/estate-explorer/>

CPAT – Cloud Premigration Advisor Tool



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

Download CPAT from MOS Note: 2758371.1



Patch 32613591: Cloud Premigration Advisor Tool (CPAT) for version 11.2.0.4 and Higher

Last Updated 11-Feb-2025 17:31 (12 days ago)

Product Oracle Database Upgrade Assistant

Release Oracle 11.2.0.4.8

Platform Generic Platform

Size 8.6 MB

Download Access Software

Classification General

Patch Tag

Release Oracle 11.2.0.4.8

Platform Generic Platform

Language American English

Bugs Resolved by This Patch

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

View Related Knowledge to this Patch

Read Me

Download



Add to Plan



Analyze with OPatch...

All-time Downloads **50**

[View Trends](#)



Discuss this patch in the community

CPAT Howto

Links to the Complete Blog Post Series about CPAT (alexzaballa.com)

- [Part 1](#) – Starting with CPAT
- [Part 2](#) – Options for Running CPAT
- [Part 3](#) – Running CPAT
- [Part 4](#) – Troubleshooting CPAT
- [Part 5](#) – CPAT Checks
- [Part 6](#) – CPAT New Features – December – 2024
- [FAQ](#)

HTML Report

ORACLE

Thu Jan 30 14:03:27 UTC 2025

Cloud Premigration Advisor Tool (CPAT) Report

CPAT Version: 24.11.1

Version Date: Nov 19, 2024

Days Since Last CPAT Update: 96 days

Table of Contents

- [Premigration Advisor Report Summary](#)
- [Report Details](#)
- [Report Analysis Notes](#)
- [Source Database Details](#)
- [Source Database Version Information](#)
- [Source Database Patch Information](#)
- [Source Database Redo Information](#)
- [Premigration Advisor Check Details List](#)
- [Report Legend](#)

Report Results Overview

Source Database		Target Database		Migration Method		Additional Tasks	
Action Required	4	Action Required	6	Action Required	0	Action Required	0
Review Required	1	Review Required	5	Review Required	1	Review Required	1
Review Suggested	3	Review Suggested	5	Review Suggested	0	Review Suggested	4
Passed	5	Passed	14	Passed	0	Passed	8

[Return to Table of Contents](#)

HTML Report

▼ Report Details

CPAT Application Version:	24.11.1
Report Generated On:	Thu Jan 30 13:41:09 UTC 2025
Analysis Mode:	FULL
Target Cloud Type:	ATPS
Migration Method(s):	[DATAPUMP]
Command Line Options:	--connectstring jdbc:oracle:oci:@ --targetcloud ATPS --sysdba --pdbname GSI --outdir /home/oracle/cpat/cpat_24_11_1/GSI --logginglevel FINE --outfileprefix GSI --migrationmethod DATAPUMP --reportformat HTML TEXT --resultlevel R0

More Details

[Return to Table of Contents](#)

HTML Report

> Report Analysis Notes

> Source Database Details

> Source Database Version Information

> Source Database Patch Information

> Source Database Redo Information

Source Database Details

Source Cloud Vendor:	Oracle Cloud Infrastructure (Database)
Source Database Host Name:	onprem
Source Oracle SID:	CD81
Source Database Created Date:	Wed Jan 29 09:42:27 UTC 2025
Source Database DBID:	1165381187
Source Database Unique Name:	CD81_vr2_fra
Source Instance Name:	CD81
Source Database Name:	CD81
Source Database Username:	SYS
Source Database Port String:	x86_64Linux 2.4.xx
Source Database Platform ID:	13
Source Database Container Name:	OSI
Source DB Block Size in KB:	8
Source DB Combined Size of DATA, TEMP, LOG, and CONTROL File Usage in GB:	212.771
Source DB Size of DATA File Usage in GB:	209.085
Source DB Size of TEMP File Usage in GB:	0.668
Source DB Size of LOG File Usage in GB:	3.0
Source DB Size of CONTROL File Usage in GB:	0.017
Source Database Applications:	EBS (Owner: APPLSYS), FINANCIALS (Owner: AR)

[Return to Table of Contents](#)

HTML Report

▼ Premigration Advisor Check Details List

Source Database

Expand All Close All

> Action Required (4 checks)

> Review Required (1 check)

> Review Suggested (3 checks)

> Passed (5 checks)

Target Database

Expand All Close All

> Action Required (6 checks)

> Review Required (5 checks)

> Review Suggested (5 checks)

> Passed (14 checks)

Migration Method

Expand All Close All

> Review Required (1 check)

Additional Tasks

Expand All Close All

> Review Required (1 check)

> Review Suggested (4 checks)

> Passed (8 checks)

[Return to Table of Contents](#)

Source Database

Expand All Close All

▼ Action Required (4 checks)

> Users with Objects but No Quota

> User Defined Objects in SYS

> User Defined Objects in SYSTEM

> References to User Objects in SYS

▼ Review Required (1 check)

> 10G Password Version

▼ Review Suggested (3 checks)

> Unavailable Options for Serverless

> Streams Pool Size for Data Pump

> Enabled Scheduler Jobs

> Passed (5 checks)

HTML Report

Source Database

[Expand All](#)[Close All](#)

▼ Action Required (4 checks)

[Users with Objects but No Quota](#)[User Defined Objects in SYS](#)[User Defined Objects in SYSTEM](#)[References to User Objects in SYS](#)

▼ Review Required (1 check)

[10G Password Version](#)

▼ Review Suggested (3 checks)

[Unavailable Options for Serverless](#)[Streams Pool Size for Data Pump](#)[Enabled Scheduler Jobs](#)

> Passed (5 checks)

▼ Action Required (4 checks)

▼ Users with Objects but No Quota

Description:

Objects in the source database belonging to users without relevant tablespace quota (or who have not been granted UNLIMITED TABLESPACE) will not be migrated to the target environment.

Action:

Assign an appropriate quota to all listed users (or grant those users UNLIMITED TABLESPACE) before initiating the migration to complete transfer of their data to the target environment. Please refer to the Report Analysis Notes for additional information related to this check.

[More Details](#)

▼ Relevant Objects (39 relevant objects)

USERNAME	UNLIMITED_TABLESPACE	TABLESPACE_NAME	MAX_BYTES	ANALYZE_SEG_SIZE
AALOMARI	NO	USERS_T_S_01	524288000	YES
ABISWAL	NO	USERS_T_S_01	104857600	YES
AGILEMIG	NO	SHARED_T_M_01	5368709120	YES
AMITSING	NO	USERS_T_S_01	594804736	YES
DEVISINGH	NO	USERS_T_S_01	104857600	YES
FMSA	NO	USERS_T_S_01	157786200	YES

Optional: Load data into CMA

Oracle Cloud Migration Advisor

mike.dietrich@oracle.com

Home

Guided Mode

Create Project

My Available Projects

Download Collectors

File Exchange with Custom...

REST Services

Help

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?

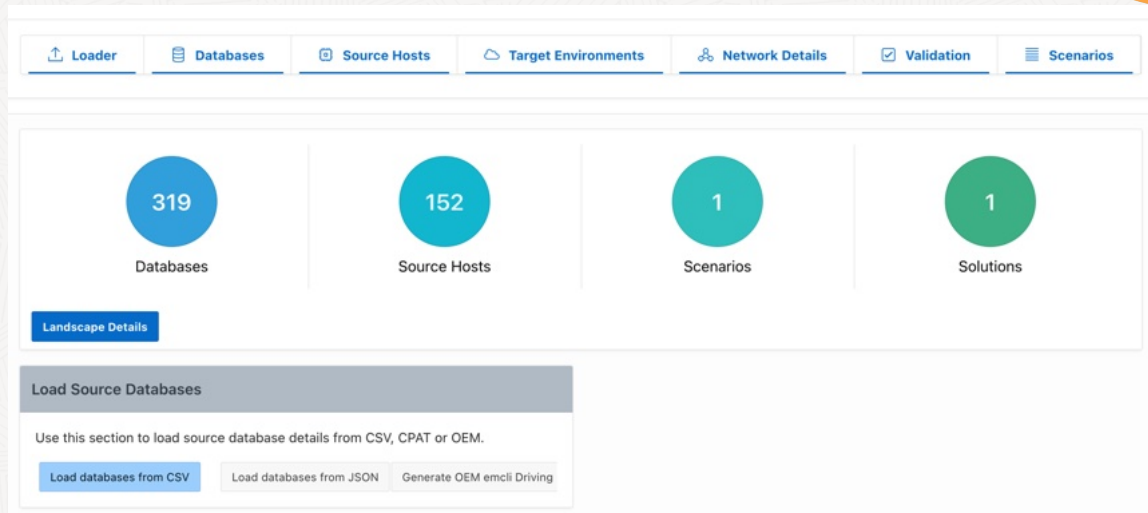
For more options, you can

- Create and configure migration scenarios using **Create Project**, or
- Access 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**

Project Overview



Project Overview

Oracle Database Version



Project Overview

Scenario Configuration

Select the **"Target Types"** and the **"Migration Methods"** you want to consider in this scenario.

Server types to advise

ADBD - Autonomous Database Dedicated	>>	ADBS - Autonomous Database Serverless
ExaDB C@C - Exadata on Cloud@Customer	>	Oracle Autonomous Database@Azure
ExaDB Dedicated - Exadata on Dedicated Infrastructure	<	
ExaDB XS - Exadata on Exascale Infrastructure	<<	
Oracle Autonomous Database@GCP		

Note: If there are more than two tied solutions, the server type with the higher priority (top) on this list will be selected.

Allowed migration technologies

☒ Enterprise Manager ☒ Golden Gate ☒ OCI Services

If you don't want one of the technologies above to be considered to perform the migrations, unselect it.

Migration methods to include

Classic Export and Import utilities	>>	ZDM Logical Online
Data Pump Conventional Export/Import	>	ZDM Logical Offline
Data Pump Full Transportable	<	ZDM Physical Offline
Data Pump Full Transportable + RMAN Convert	<<	ZDM Physical Online
Data Pump Import with DBLink		

Note: If there are more than two tied solutions, the migration method with the higher priority (top) on this list will be selected.

Project Overview

Databases

Target Environment

Migration Methods

Checks Result

Q Search: All Text Columns Go Actions

Reset

Target Name	Final Score (the higher the better)	Complexity Score (higher value = easier to migrate)	Target Accepted	Auto-Patching	Auto-Security	Administrator
ADBS - Autonomous Database Serverless	9.96	9.61	Yes	★★★★★	★★★★★	★☆☆☆☆
Oracle Autonomous Database@Azure	9.95	9.61	Yes	★★★★★	★★★★★	★☆☆☆☆


1 rows selectedTotal 2

Note: The **green row** is the currently defined target environment. You can change it by **selecting another method** from another target.

Migration Methods

This section shows the possible migration methods for the selected **target type** and **database**. You can change the defined migration method and target by clicking on the **"Select"** button (✓).


You can also check the instructions or simply click on the card and read more details about it on the "Decision Log".

**ZDM Logical Online**
Complexity: Easy | Downtime: Low
Migration Score: 9.96

The Logical Online Migration Workflow in ZDM leverages Data Pump for initial load and on Oracle GoldenGate for synchronization purposes and to preserve the online portion of the migration process. Oracle ZDM will run on a separate node and connect to both Source and Target to perform the migration.

✓ Select

Read Instructions

**ZDM Logical Offline**
Complexity: Easy | Downtime: Medium
Migration Score: 8.97

The Logical Offline Migration Workflow in ZDM leverages Data Pump for data transfer and target instantiation. Some Zero Downtime Migration logical migration work flows involve placing Oracle Data Pump dump files on storage media for transfer to the target database.

✓ Select

Read Instructions

CMA – How to get the tool?

Send us an email – it is an APEX app

- Vagrant build
- Zip file ready to be installed into ADB-free
- Oracle internal: cma.oraclecorp.com

Migration Execution

Logical Migrations

Move your data, not your database

How do you approach it?

- Full database export?
- Schema export?

We have tools

- ZDM – Zero Downtime Migration
- DMS – Database Migration Service
- AMA – Autonomous Migration Automation

Why no Full Database Export?

Oracle Autonomous Database is different


- Different admin user
- No DBA role
- Different key role names
- You can't put anything in SYS or SYSTEM
- XML Binary schema objects are not supported
- Libraries are not really supported
- Database Links and External Tables require extra work
- ...

Be aware that certain objects aren't part of schema exports!

AMA – Autonomous Migration Automation

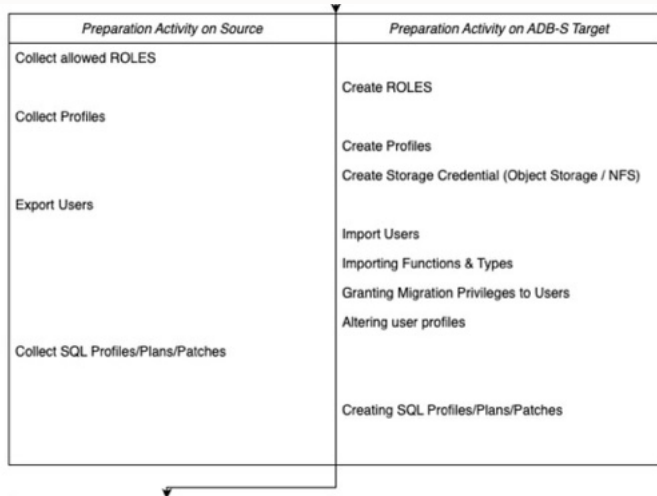
Before you start ...

<i>Before you start On-Prem</i>	<i>Before you start ADB-S Target</i>
<p>Gater sys/system stats</p> <p>create migration user</p> <p>enable restricted session</p> <p>set job_queue_processes to 0</p>	<p>Enable OLAP/Java in ADB</p>



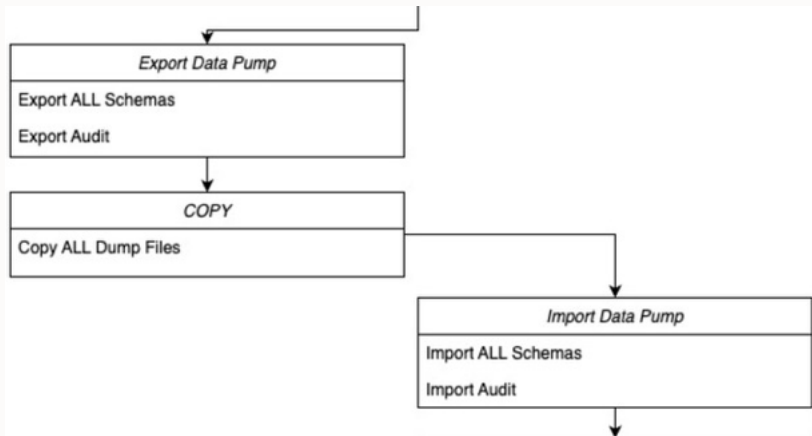
AMA – Autonomous Migration Automation

Preparation:



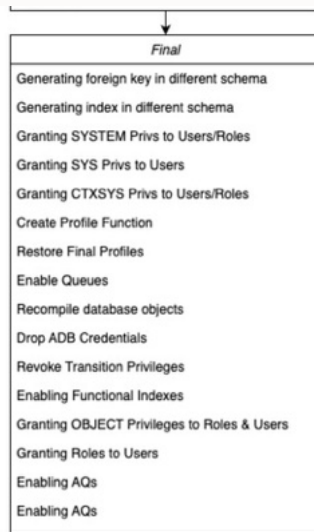
AMA – Autonomous Migration Automation

Export and import:



AMA – Autonomous Migration Automation

Final tasks:



Important Things To Consider

NFS vs Object Storage

We recommend to use NFS instead of object storage for the migration

- Easier to setup
- Performance seems to be better
- NFS requires "private endpoint access only"

Private endpoint access only

Restrict access to a private endpoint within an OCI VCN.

Create File System

[Help](#)

This workflow creates a new File System. To get started, choose the type of File System you want to create. Then, you can keep the provided information or click **Edit details** to change it. Click **Create** to finish.

File System for NFS

Create a File System and an associated Export in a Mount Target. You can mount and access the File System as soon as it is created. [Learn more about mounting File Systems.](#)



File System for Replication

Create an unexported File System. Unexported File Systems can be used as target File Systems for replicated data. [Learn more about replication.](#)

Pre-Defined Services to connect to ADB

Service	Parallel by Default	Resource Shares (when 100%)	Purpose
TPURGENT	Manual	12	Highest priority short-running queries
TP	N	8	short-running queries and transactions
For regular workloads, use TPURGENT, TP or LOW			
LOW	N	1	short-running queries and transactions

Tools out-of-the-box



SQL Developer Web

Web-based Function rich,
low code development env
No client software needed



Oracle REST Data Services

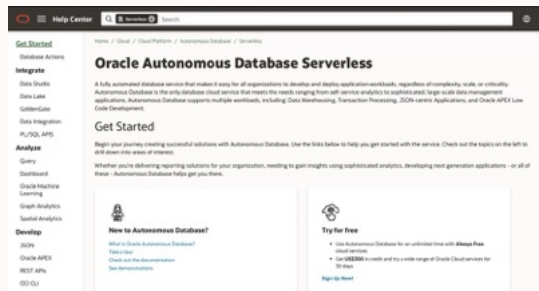
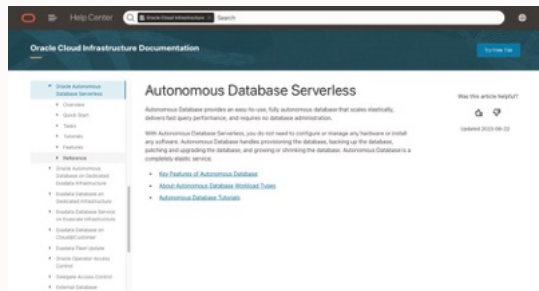
Execute SQL and PL/SQL
Build Data Models,
generate DDL statements
Monitor and manage the DB



APEX

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

Autonomous Database Essentials



OCI Documentation for everything
“outside” the database

i.e. deploy, start-stop, manage through the console

<https://docs.oracle.com/en-us/iaas/autonomous-database-serverless/index.html>

Database Documentation for everything
“inside” the database

i.e. schema, capabilities, using the database through a connection

<https://docs.oracle.com/en/cloud/paas/autonomous-database/serverless/index.html>



Do you plan to migrate to Autonomous Database?

Talk to us – we will guide you!

Key Learnings



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

Break

We start again at 15:15

Cross-Platform Migrations



Pushing the limits



The Beast

Introduction

Who is who?



ANDREAS GROETZ

Oracle DBA Tech Lead

Entain Services Austria GmbH

Entain is one of the world's largest sports betting and gaming groups. Leveraging the power of the Entain Platform, they bring moments of excitement into their customers lives through more than 30 iconic brands such as bwin, Coral, Ladbrokes and many more.

Entain operates on over 140 licenses across 40+ territories and employs over 29,000 talented workforce. Entain is listed on the London Stock Exchange and is a constituent of the FTSE 100 Index.

Σntain



Ladbrokes



sportingbet

CORAL 
GET CLOSER TO THE ACTION

bwin 

EUROBET 

SuperSport

party ker

Foxy BINGO

Challenges

What is special, what makes it so complex?

Migration Challenges



SPARC SuperCluster



ZDLRA



Exadata X9M Extreme Flash

Migration Challenges

180TB
size



SPARC SuperCluster



ZDLRA



Exadata X9M Extreme Flash

Migration Challenges

15TB
redo/day



SPARC SuperCluster



ZDLRA



Exadata X9M Extreme Flash

Migration Challenges



SPARC SuperCluster



Exadata X9M Extreme Flash



5 Physical Standby DBs

Local, and in different region, 2500km away

Constraints

Limiting factors, and other things to know



Photo by Mihaly Koles on Unsplash

Up to 15TB redo/day is beyond what Oracle GoldenGate will be able to synch


The system is highly active 24 x 7 x 365



Photo by Dave Hoefler on Unsplash

Very large database,
very constrained downtime

- 180+ TB database size
- 5-6 TB growth/month
- Every minute of downtime costs \$\$\$



Migrating tablespaces upfront or separately **definitely** not an option

- Way too many cross-dependencies
- Tablespaces aren't isolated



Photo by Linh Ha on Unsplash

Every complex Oracle data type you can imagine is used

- XML binary types
- Nested partitioned tables
- Evolved object types



Photo by [Masaki Komori](#) on [Unsplash](#)

Tight downtime window

- Dry run: 2 hours outage approved
 - Tablespace read-only
 - Full Transportable Export
- Live migration: 13 hours approved



Photo by Moreanella Breton on Unsplash

The available **Oracle V4 PERL** migration scripts would have been worked technically, but ...

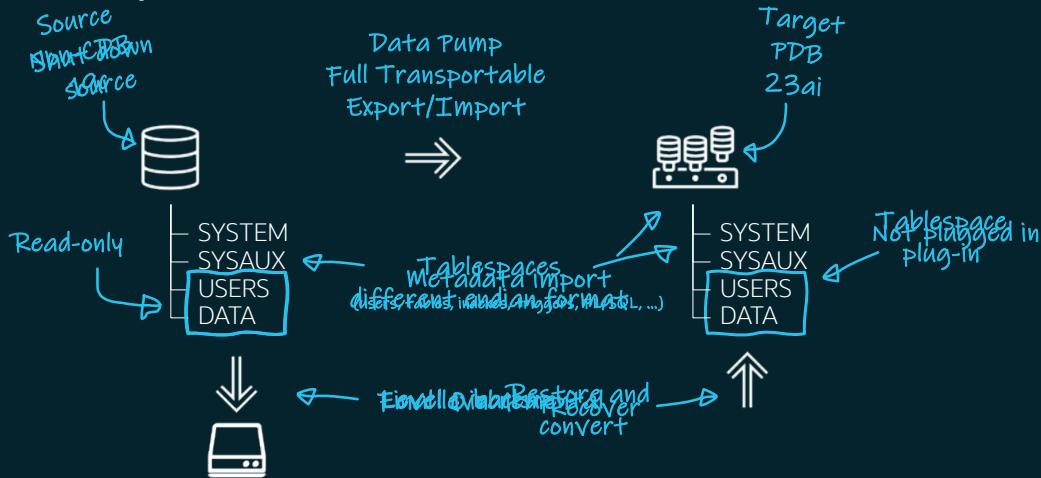
- No section-size backup support
- No standby backup support
- No selective PDB migration support

Migration




Typically, we use **Full Transportable Export/Import** for large cross-endian migrations

Concept



Scripts for Incremental Backup Automation

Backup / restore / recover



- No multisection backups
- No encrypted tablespaces
- Inefficient parallelism
- Incomplete multitenant support

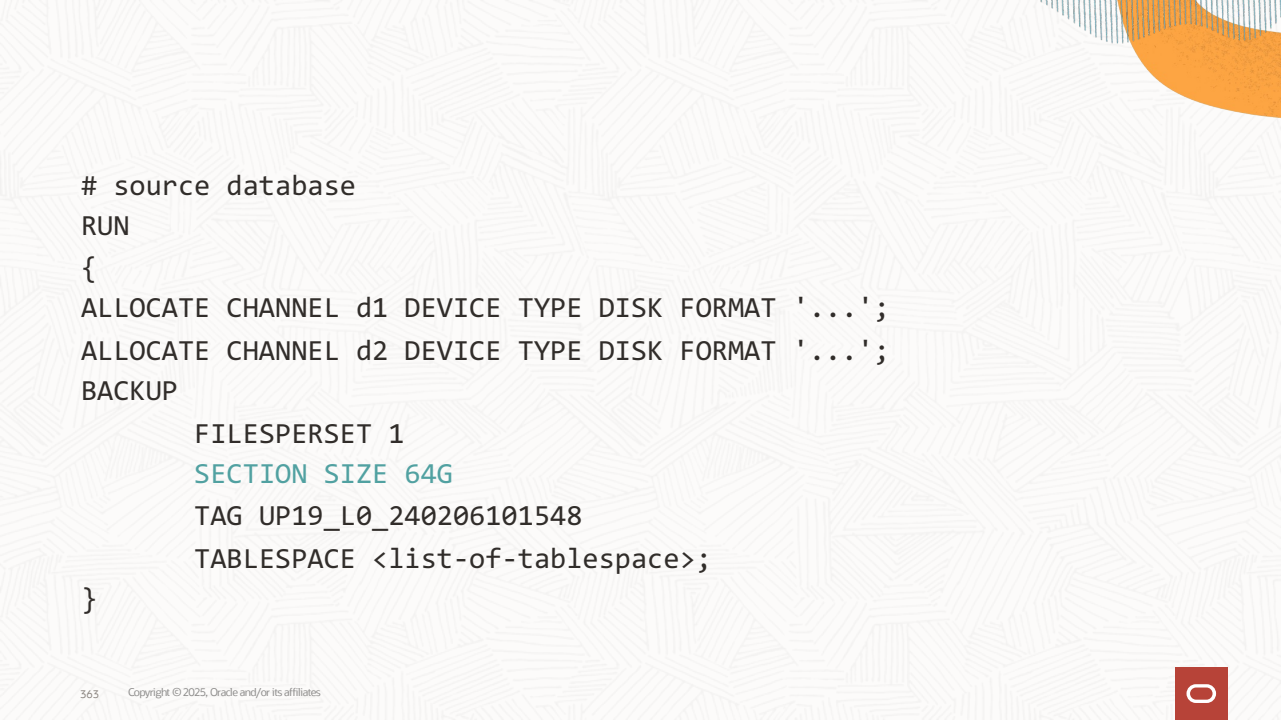


M5 is the **next-generation** cross-platform transportable tablespace procedure

- New RMAN functionality combined with Full Transportable Export/Import
- Doc ID [29991571](#)

M5 Migration Script

The new migrations scripts superseding the V4 PERL scripts



```
# source database
RUN
{
  ALLOCATE CHANNEL d1 DEVICE TYPE DISK FORMAT '...';
  ALLOCATE CHANNEL d2 DEVICE TYPE DISK FORMAT '...';
  BACKUP

    FILESPERSET 1
    SECTION SIZE 64G
    TAG UP19_L0_240206101548
    TABLESPACE <list-of-tablespace>;

}
```

```
# source database
RUN
{
  ALLOCATE CHANNEL d1 DEVICE TYPE DISK FORMAT '...';
  ALLOCATE CHANNEL d2 DEVICE TYPE DISK FORMAT '...';
  BACKUP
    FILESPERSET 1
    SECTION SIZE 64G
    TAG UP19_L0_240206101548
    TABLESPACE <list-of-tablespace>;
}
```

```
# target database
RUN
{
  ALLOCATE CHANNEL DISK1 DEVICE TYPE DISK FORMAT '...';
  ALLOCATE CHANNEL DISK2 DEVICE TYPE DISK FORMAT '...';
  RESTORE ALL FOREIGN DATAFILES TO NEW FROM BACKUPSET
    '<backup-set-1>',
    '<backup-set-2>',
    ...
    '<backup-set-n>'
};
```


Benefits

M5 procedure supports:

- Encrypted tablespaces
- Multisection backups
- Migrating multiple databases into the same CDB simultaneously
- Compressed backup sets
- Better parallelism

Requirements

- Source and target database must
 - be 19.18.0 or higher
 - use Data Pump Bundle Patch



Always use the latest version of M5 script

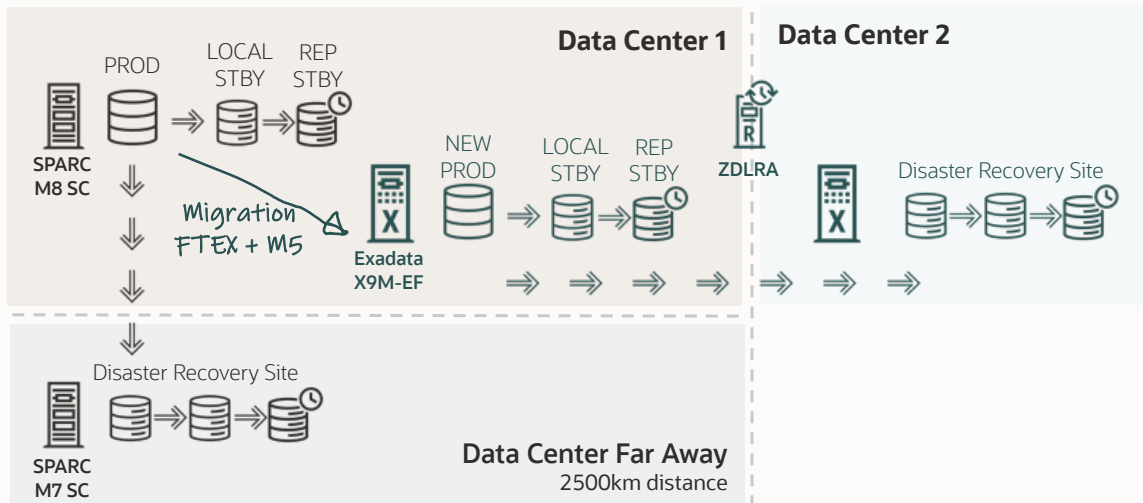
- Download from Doc ID [29991571](#)



Use Block Change Tracking for faster incremental backups

- Check the License Guide for details

Migration Plan



M5 Workflow

Configure

Level 0

Level 1

Outage

Final
Backup

Final
restore

Export
Import



SPARC
SuperCluster



Local standby




ZDLRA



Exadata
X9M Extreme Flash




M5 Workflow



- Download M5 script from Doc ID [2999157.1](#)
- Configure shared NFS
- Edit `dbmig_ts_list.txt`
- Edit `dbmig_driver.properties`
- Create new, empty target database

M5 Workflow



SPARC
SuperCluster



ZDLRA



Exadata
X9M Extreme Flash

M5 Workflow



- Start initial level 0 backup
 - Use driver script `dbmig_driver_m5.sh` `L0`
- Driver script creates a restore script
 - Restore using `restore_L0_<source_sid>_<timestamp>.cmd`
- Check logs

M5 Workflow

Configure

Level 0

Level 1

Outage

Final
Backup

Final
restore

Export
Import



SPARC
SuperCluster




ZDLRA



Exadata
X9M Extreme Flash

M5 Workflow



- Start level 1 incremental backup
 - Use driver script `dbmig_driver_m5.sh` [L1](#)
- Driver script creates a restore script
 - Restore using `restore_L1_<source_sid>_<timestamp>.cmd`
- Check logs
- Repeat as often as desired

M5 Workflow

Configure

Level 0

Level 1

Outage

Final
Backup

Final
restore

Export
Import



SPARC
SuperCluster




ZDLRA



Exadata
X9M Extreme Flash

M5 Workflow



- Maintenance window begins
- Read-only sessions can still use the database

M5 Workflow

Configure

Level 0

Level 1

Outage

Final
Backup

Final
restore

Export
Import



SPARC
SuperCluster




ZDLRA



Exadata
X9M Extreme Flash

M5 Workflow




- Start final level 1 incremental backup
 - Use driver script `dbmig_driver_m5.sh` **L1F**
 - Sets tablespaces read-only
 - Performs level 1 incremental backup
 - Starts Data Pump full transportable export
- Optionally, shuts down source database

M5 Workflow



- Driver script created a restore script
 - Restore using `restore_L1F_<source_sid>_<timestamp>.cmd`
- Check logs

M5 Workflow



SPARC
SuperCluster




ZDLRA



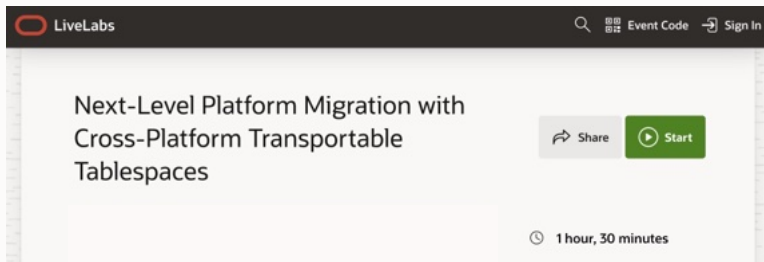
Exadata
X9M Extreme Flash

M5 Workflow



- Copy Data Pump dump file to *DATA_PUMP_DIR*
- Use import driver script in test mode
 - Start `impdp.sh <dump_file> <restore_log> test`
- Check generated parameter file
 - Use `impdp.sh <dump_file> <restore_log> run`
- Check Data Pump log file

Wanna try it out?



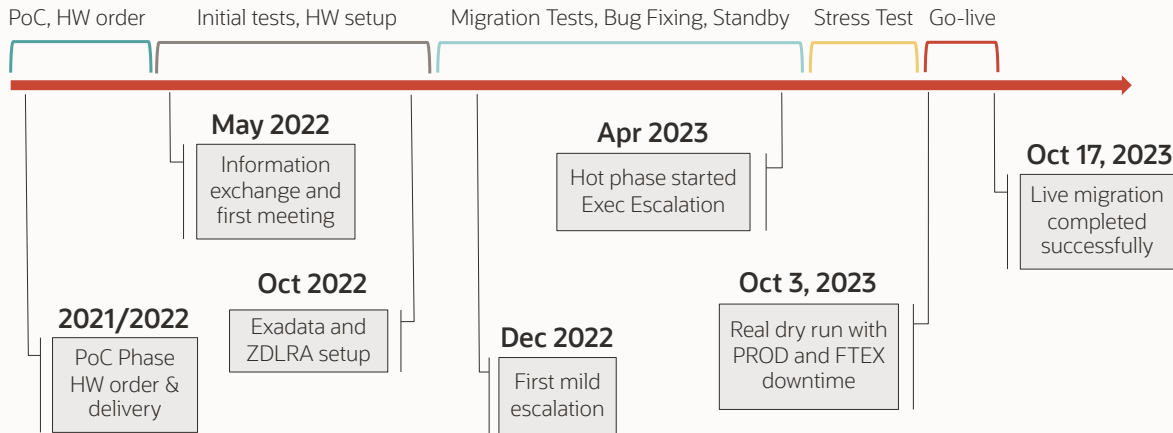
Oracle LiveLabs – Run the lab just inside your browser!

The Project Plan



Timelines and the Run Book

Overall Project Timeline



Key to Success: Runbook

Complex projects absolutely require a **detailed runbook**

ID	Task	Status	Responsible Primary Person	Responsible Secondary Person	Predecessor	Start Time (CEST)	Duration (hh:mm)	End Time (CEST)	Start Time (IST)	End Time (IST)	Actual Start Time (CEST)	Actual Duration	Actual End Time (CEST)	Comments - Blocker
----	------	--------	----------------------------------	------------------------------------	-------------	----------------------	---------------------	--------------------	------------------------	-------------------	--------------------------------	--------------------	------------------------------	-----------------------

- This run book covered over 200 individual tasks

A screenshot of a complex project runbook table. The table has many columns, including ID, Task, Status, Responsible Primary Person, Responsible Secondary Person, Predecessor, Start Time (CEST), Duration (hh:mm), End Time (CEST), Start Time (IST), End Time (IST), Actual Start Time (CEST), Actual Duration, Actual End Time (CEST), and Comments - Blocker. The rows are color-coded in green, blue, and yellow, representing different task statuses or categories. The table is densely populated with data, showing a detailed schedule for a project with over 200 tasks.

387 Copyright © 2025, Oracle and/or its affiliates

How many people were involved?



6 DBAs



Managers



2 data center
engineers



4 product
engineering



4 DB developer



War Room



2-3 infrastructure



2 monitoring,
communication



4 compliance



2 network



10 testing

Migration Issues



Some of them...

Where we started ...



PoC first FTEX **export**:

01-OCT-21 05:32:36.275: Job "SYSTEM"."SYS_EXPORT_FULL_01" successfully completed at
Fri Oct 1 05:32:36 2021 elapsed **0 04:25:22**

PoC first FTEX **import**:

05-OCT-21 01:48:59.534: Job "SYSTEM"."SYS_IMPORT_FULL_01" completed with **103000**
error(s) at Tue Oct 5 01:48:59 2021 elapsed **3 18:34:09**



The way forward ...

80 SRs opened and solved in various areas

Migration
ZDLRA
Standby
Infrastructure

18 one-off patches

5 merges

Daily calls with Oracle, countless evening / night / weekend hours

Many areas required special attention....

- Optimizer Statistics
- Scheduler jobs
- Resource Manager
- Cross Schema objects
- AQ
- Evolved Types/partitioned nested tables
- Binary XML
- Standby DBs
- ...



Issue 1 | Long Running Metadata Import

Fix applied to remedy export errors

- BUG 34201281 - MERGE ON DATABASE RU 19.12.0.0.0 OF 33963454 34052641

Result:

- Now Full Transportable import alone took over 6 days (!!)

```
08-JUN-22 16:21:17.887: W-1 Processing object type DATABASE_EXPORT/.../PROCACT_INSTANCE
14-JUN-22 18:56:58.813: W-1      Completed 108 PROCACT_INSTANCE objects in 527737 seconds
...
14-JUN-22 19:15:50.016: Job "SYSTEM"."SYS_IMPORT_TRANSPORTABLE_01" completed with 316 error(s)
                        at Tue Jun 14 19:15:49 2022 elapsed 6 06:45:38
```

Issue 1 | Long Running Meta Import

Long running action identified via tracing:

```
UPDATE "POSTMAN"."T_MAIL_LOG"  
SET "C_CVAR"=SYS_REMAP_XMLTYPE("C_CVAR")
```

- 300+ million rows

Issue in internal package DBMS_CSX_INT

- Fast merge of XMLTYPE is not happening as expected
 - *Reason:* Incorrect internal check query
- Tokens between source and target are not identical
 - *Reason:* Different Endianness

Issue 1 | Long Running Meta Import

Solution:

- Use workaround from MOS Note: 2309649.1 in UPGRADE mode
 - [MOS Note: 2309649.1 - How to Migrate Large Amount of Binary XML Data between Databases](#)

```
25-JUL-22 12:28:40.813: Job "SYSTEM"."SYS_IMPORT_TRANSPORTABLE_01" completed  
with 317 error(s) at Mon Jul 25 12:28:40 2022 elapsed 0 04:33:03
```

Issue 2 | Metadata API and Nested Tables

Full transportable import errors out for a **nested partitioned table**

```
PLS-00172: string literal too long  
ORA-39151: Table "DBA_XY"."X_GAMES" exists.  
All dependent metadata and data will be skipped due to table_exists_action
```


Root cause was a string overflow in the Metadata API

- Data Pump creates the index, and then alters it – here the overflow happened
- Side effect was a misleading error message

Issue 3 | Data Pump Bundle Patch

Many of the TTS and Metadata fixes got included into Data Pump Bundle

- [Data Pump Recommended Proactive Patches For 19.10 and Above \(Doc ID 2819284.1\)](#)



Issue 4 | Evolved Object Types

Evolved TYPES can lead to Data Pump errors during transportable import:

```
ORA-39083: Object type TABLE:"APPUSER"."CARS" failed to create with error:  
ORA-39218: type check on object type "APPUSER"."CAR_TYPE" failed  
ORA-39216: object type "APPUSER"."CAR_TYPE" hashcode or version number mismatch
```

Further Information:

- [Blog Post: Understand why Data Pump errors with evolved types](#)



Using evolved types in table definitions

--Create a new type. The type is now version 1

--Use the type in a table

```
CREATE TYPE CAR_INFO_TYPE IS OBJECT (model VARCHAR2(40));
```

```
CREATE TABLE CARS (id number, car_info car_info_type);
```

```
INSERT INTO CARS VALUES (1, car_info_type('Volvo V90'));
```

The type is now evolving



--Make a change to the type. The type is now version 2

```
ALTER TYPE CAR_INFO_TYPE ADD ATTRIBUTE horsepower NUMBER CASCADE NOT INCLUDING TABLE DATA;
```

```
INSERT INTO CARS VALUES (2, car_info_type('BMW 530i', 250));
```

Existing data is not updated



--Make another change to the type. The type is now version 3

```
ALTER TYPE CAR_INFO_TYPE ADD ATTRIBUTE color VARCHAR2(20) CASCADE NOT INCLUDING TABLE DATA;
```

```
INSERT INTO CARS VALUES (3, car_info_type('Hyundai Sonata', 160, 'Black'));
```

Evolved Types

```
SELECT * FROM CARS
```



CARS	
1	car_info_type v1: Volvo V90
2	car_info_type v2: BMW 530i, 250
3	car_info_type v3: Hyundai Sonata, 160, Black



DICTIONARY	
car_info_type v1	model
car_info_type v2	model, horsepower
car_info_type v3	model, horsepower, color



Data Pump recreates types during
Full Transportable Export/Import

Evolved Types

- To avoid **data corruption**,
Data Pump must recreate the exact same type evolution in target database
- Due to **implementation restrictions**,
it is not always possible to recreate the exact same type evolution
- In such situations, to avoid corruption,
Data Pump reports ORA-39218 or ORA-39216 on **import**



Evolved Types | Possible Solutions

- 1 Conventional Data Pump export
- 2 Manually recreate `type` in target database with matching evolution
- 3 Recreate `type` without evolution before export

[Blog post](#) with details

Issue 5 | Advanced Queueing

Source database

`<queue_table_name>`
`AQ$_<queue_table_name>_E`
`AQ$_<queue_table_name>_I`
`AQ$_<queue_table_name>_T`
`AQ$_<queue_table_name>_F`
`AQ$_<queue_table_name>_C`
`AQ$_<queue_table_name>_D`
`AQ$_<queue_table_name>_G`
`AQ$_<queue_table_name>_H`
`AQ$_<queue_table_name>_L`
`AQ$_<queue_table_name>_P`
`AQ$_<queue_table_name>_S`
`AQ$_<queue_table_name>_V`

Queue table

Queue
infrastructure

Target database

`<queue_table_name>`
`AQ$_<queue_table_name>_E`
`AQ$_<queue_table_name>_I`
`AQ$_<queue_table_name>_T`
`AQ$_<queue_table_name>_F`

Issue 5 | Advanced Queueing

Queue tables and underlying objects may change during import

- COMPATIBLE **during creation** of queue tables matters
- COMPATIBLE **during import** matters as well
- [MOS Note: 2291530.1 - Understanding How AQ Objects Are Exported And Imported](#)
- [Blog post: Changing data types in queue tables during import](#)

Options:

- Recreate the queue tables with "old" COMPATIBLE setting
- Benefit from new COMPATIBLE setting and test the application



Take into account when comparing source and target databases' object count

- Understanding How Advanced Queueing (AQ) Objects Are Exported And Imported (Doc ID [2291530.1](#))



Data Pump does not start queues

- Manually start queues after migration
- Use `DBMS_AQADM.START_QUEUE`

Issue 6 | Default Tablespaces

Due to a security fix export wants to write into the user's default tablespace

- Bug 27692190
- But default tablespaces are read-only while full transportable export runs

Workaround:

- Change default tablespace for all users to SYSTEM
- Full Transportable Export
- Full Transportable Import
- Revert default tablespaces back to original in target (and source)

Issue 7 | Exporting Statistics

Exporting statistics is slow using
`DBMS_STATS.EXPORT_SCHEMA_STATS`

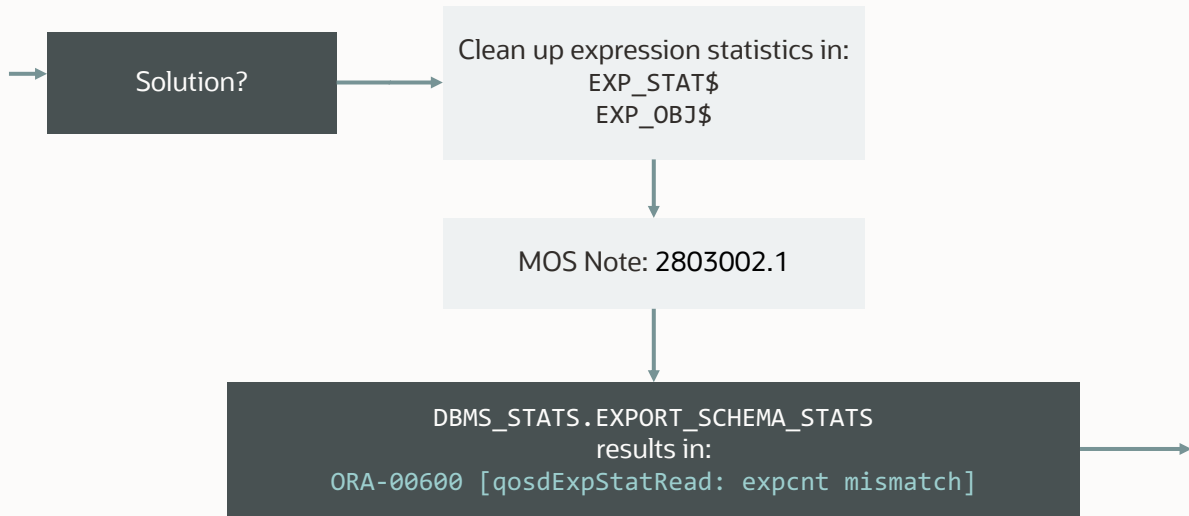


10046 trace reveals long runtime on:
`EXP_STAT$`
`EXP_OBJ$`

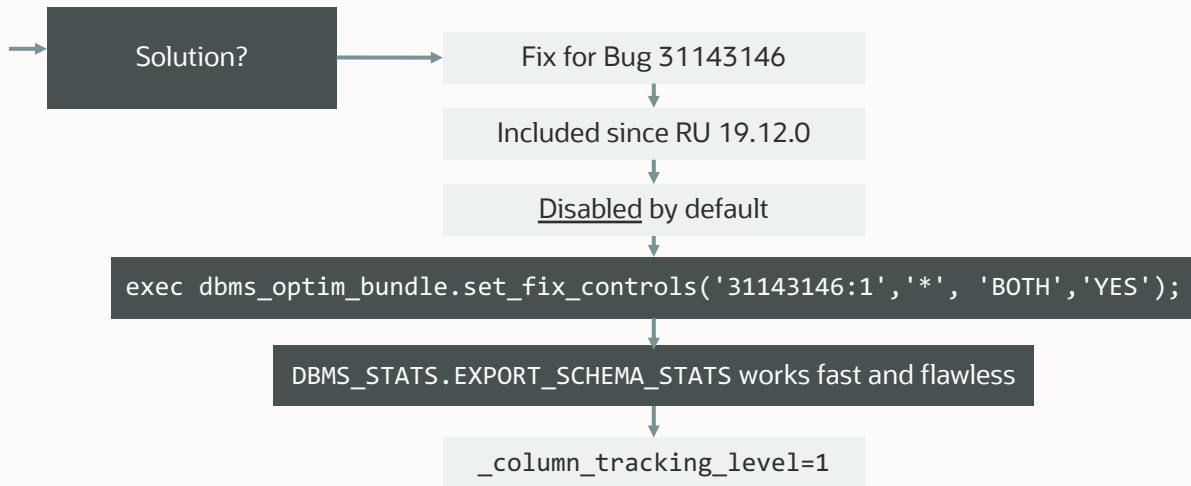
Expression Statistics for Auto-Indexing

Default: `_column_tracking_level=53;`

Issue 7 | Exporting Statistics



Issue 7 | Exporting Statistics



Issue 8 | Auditing

Tablespaces containing auditing tables **can't be set read-only**

Data Pump always unloads the audit records into the dump file

- Huge audit trail will lead to a huge dump file and longer outage

Options:

- Export audit records, and eventually import them afterwards
- Archive audit records, purge the audit trail

Key Learnings



- 1 Practice, practice, practice
- 2 Create a detailed runbook
- 3 Work together

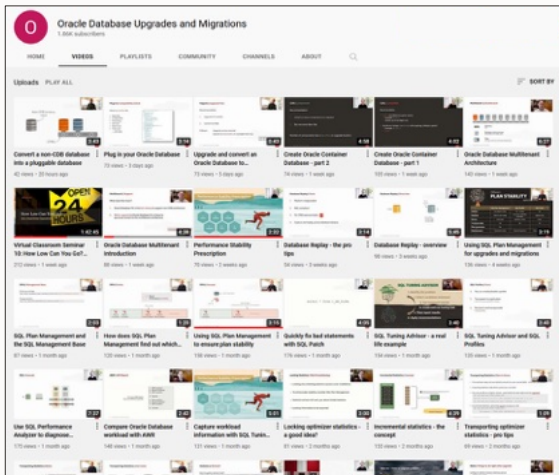
Oracle

DBAs

run the world



YouTube | Oracle Database Upgrades and Migrations



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

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



Thank You

