Migration to Oracle Autonomous Database

Part 3: Migrating

Photo by Jason Leung on Unsplash

Oracle

DBAs

run the world

MIKE DIETRICH
Vice President
Database Upgrade, Migrations & Patching

mikedietrich

@mikedietrichde.com

https://mikedietrichde.com





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

dohdatabase

@dohdatabase.com

https://dohdatabase.com





RODRIGO JORGE
Distinguished Product Manager
Database Upgrade, Migrations & Patching

rodrigoaraujorge

@dbarj.com.br

https://dbarj.com.br





ALEX ZABALLA
Distinguished Product Manager
Database Upgrade, Migrations & Patching

alexzaballa

@alexzaballa.bsky.social

https://alexzaballa.com





KLAUS GRONAU
Consulting Member of Technical Staff
Database Upgrade, Migrations & Patching

klaus-gronau-39a43aa9

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

Download the Slides

https://MikeDietrichDE.com/slides

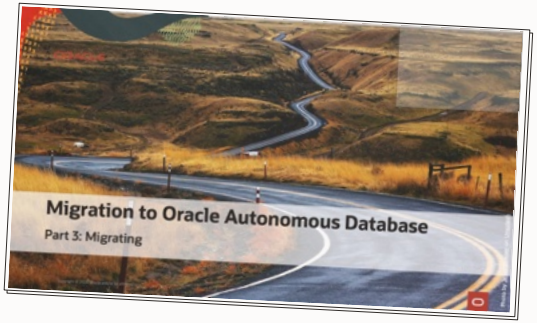

Virtual Classroom Seminars	Slides
Episode 16 (replaces Episode 1 from Feb 2021) Oracle Database Release and Patching Strategy for 19c and 23c *115 minutes – May 10, 2023*	
Episode 17 From SR to Patch – Insights into the Oracle Database Development process *55 minutes – June 22, 2023*	
Episode 18 Cross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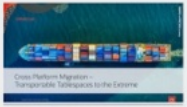
Episode 19 Move to Oracle Database 23ai – Everything you need to know about Multitenant PART 1 *145 min – May 16, 2024*	
Episode 20 Move to Oracle Database 23ai – Everything you need to know about Multitenant PART 2 *100 min – June 28, 2024*	

Recorded Web Seminars

https://MikeDietrichDE.com/videos

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

Let's Do This Together

1 PLANNING

Watch recording
Get slides

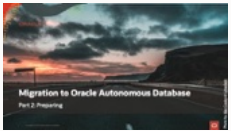

2 PREPARING

Watch recording
Get slides

3 MIGRATING

Watch **soon**
Get slides

4 OPERATING

July 10, 15:00 CET
Sign up

Recap

What happened in Part 2?

Migration to Autonomous Database is always a logical migration

- Move the data, not the database

Getting an Overview

1 Estate Explorer

2 Cloud Premigration Advisor Tool

3 Cloud Migration Advisor

Getting an Overview

1 Estate Explorer

2 Cloud Premigration Advisor Tool

3 Cloud Migration Advisor

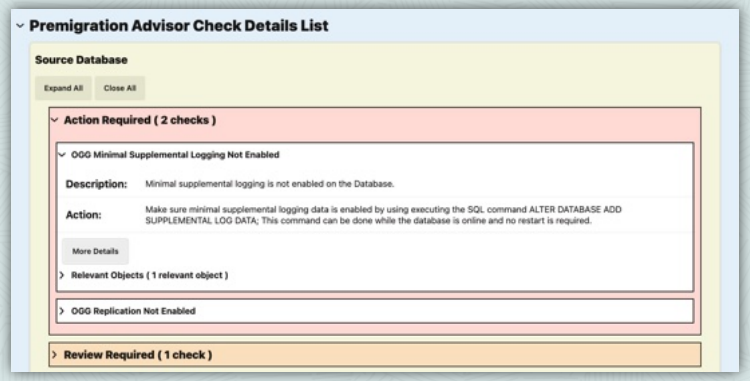

Getting an Overview

2

Cloud Premigration Advisor Tool

Evaluate an Oracle Database for compatibility with Autonomous Database

- Use Cloud Premigration Advisor Tool (CPAT)
- Download CPAT from MOS Note: 2758371.1

Databases - Recap

These are the databases we were going to migrate

Example Databases

In this series, we will use two databases:

- **The Simple Database**
- **The Complex Database**

Example Databases

In this series, we will use two databases:

- **The Simple Database**
  Based on standard Oracle schemas

HR: Human Resources
CO: Customer Orders
SH: Sales History

```
@hr_install.sql
@co_install.sql
@sh_install.sql
```

- **The Complex Database**

Example Databases

In this series, we will use two databases:

- **The Simple Database**

- **The Complex Database**
  Standard schemas and manually created objects

HR: Human Resources
CO: Customer Orders
SH: Sales History

+

External Tables
External Library
Tables with encrypted columns
Java Objects
Tables with XML columns
XML Schemas
Tables using Spatial
Profiles using custom password verification functions
Tables with ROWID columns
SQL Patches and SQL Plan Baselines
Table Clusters
Jobs using DBMS_JOB
Scheduler Jobs running external scripts
Procedure calling DBMS_SHARED_POOL + UTL_HTTP

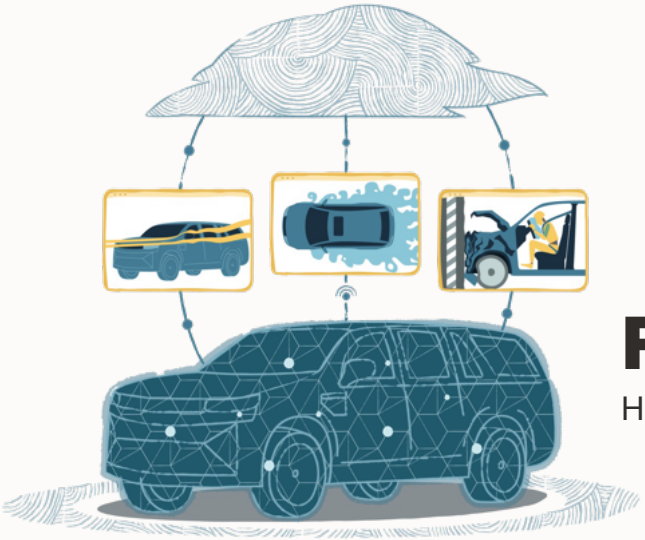

Fixing

How did we fix the CPAT findings?

CPAT | Simple Database

--migrationmethod **datapump**

Premigration Advisor Report Summary

Report Result	Review Required
Number of schemas analyzed:	3
List of schemas analyzed:	[PDBADMIN, CO, HR]

Report Results Overview

Source Database		Target Database		Migration Method		Additional Tasks	
Action Required	0	Action Required	0	Action Required	0	Action Required	0
Review Required	0	Review Required	1	Review Required	0	Review Required	0
Review Suggested	1	Review Suggested	2	Review Suggested	0	Review Suggested	3
Passed	12	Passed	27	Passed	1	Passed	10

Return to Table of Contents

Fixing Findings | Simple Database

- Scheduler Jobs
- Directory Objects
- Trusted Server Entries
- Auditing

Fixing Findings

Scheduler Jobs

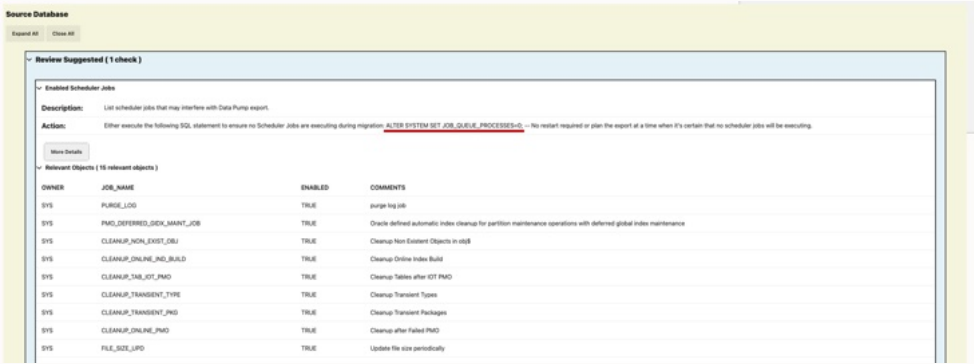

Fixing Findings

Directory Objects

Fixing Findings

Trusted Server Entries

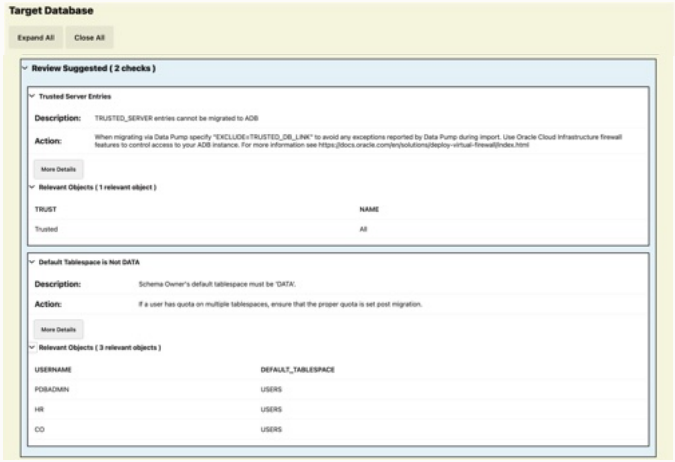

Fixing Findings

Auditing

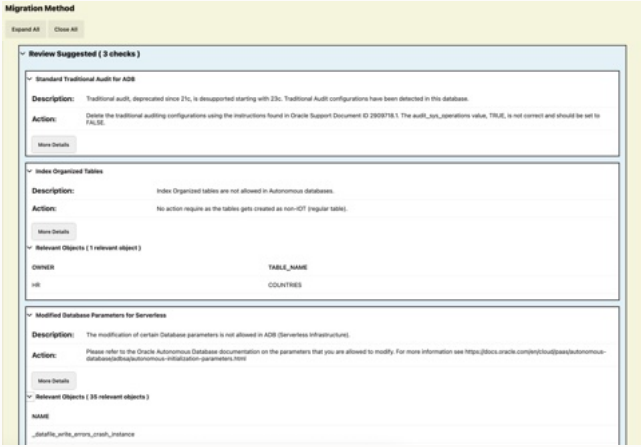

CPAT | Simple Database with Online

--migrationmethod **all**

Premigration Advisor Report Summary

Report Result	Action Required
Number of schemas analyzed:	4
List of schemas analyzed:	[PDBADMIN, CO, HR, C##DBLCMUSER]

Report Results Overview

Source Database		Target Database		Migration Method		Additional Tasks	
Action Required	2	Action Required	1	Action Required	1	Action Required	0
Review Required	1	Review Required	2	Review Required	0	Review Required	0
Review Suggested	1	Review Suggested	2	Review Suggested	1	Review Suggested	5
Passed	17	Passed	33	Passed	6	Passed	15

Return to Table of Contents

Fixing Findings | Simple Database with Online

- Scheduler Jobs
- Dictionary Objects
- Trusted Server Entries
- Auditing
- Supplemental Logging
- Streams Pool Size

Fixing Findings

Supplemental Logging

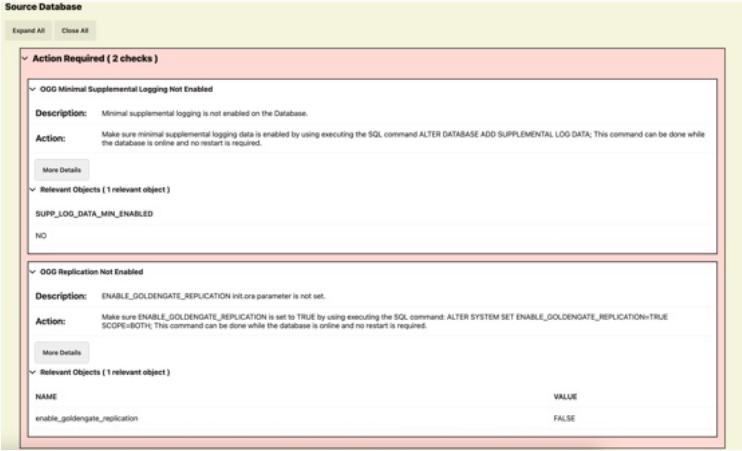

Fixing Findings

Streams Pool

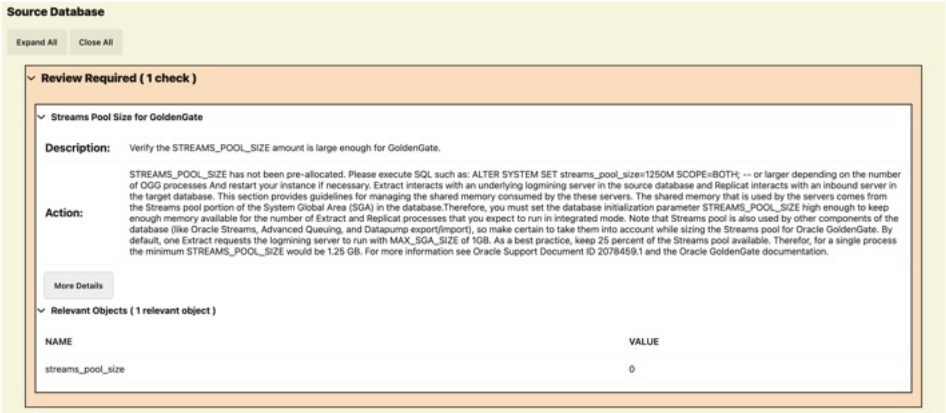

```
[oracle@db19c cpat]$
```

CPAT | Complex Database

--migrationmethod **datapump**

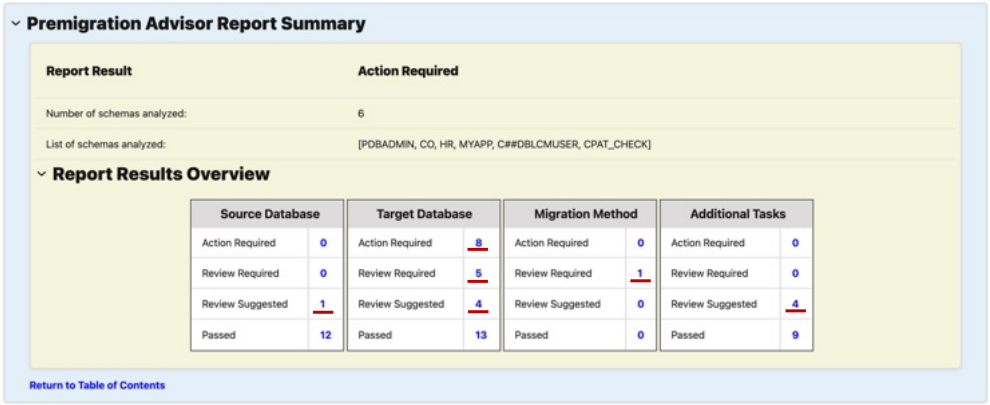

Premigration Advisor Report Summary

Report Result	Action Required
Number of schemas analyzed:	6
List of schemas analyzed:	[PDBADMIN, CO, HR, MYAPP, C##DBLCMUSER, CPAT_CHECK]

Report Results Overview

Source Database		Target Database		Migration Method		Additional Tasks	
Action Required	0	Action Required	8	Action Required	0	Action Required	0
Review Required	0	Review Required	5	Review Required	1	Review Required	0
Review Suggested	1	Review Suggested	4	Review Suggested	0	Review Suggested	4
Passed	12	Passed	13	Passed	0	Passed	9

Return to Table of Contents

Fixing Findings | Complex Database

- Previous Findings, plus ...
- Libraries
- System Privileges
- XML Objects, Tables, Columns
- Java Sources, Objects
- Common Objects
- Non-Exported Grants
- Scheduler Jobs
- External Tables
- Restricted Packages
- Directory Objects Accessibility

Libraries

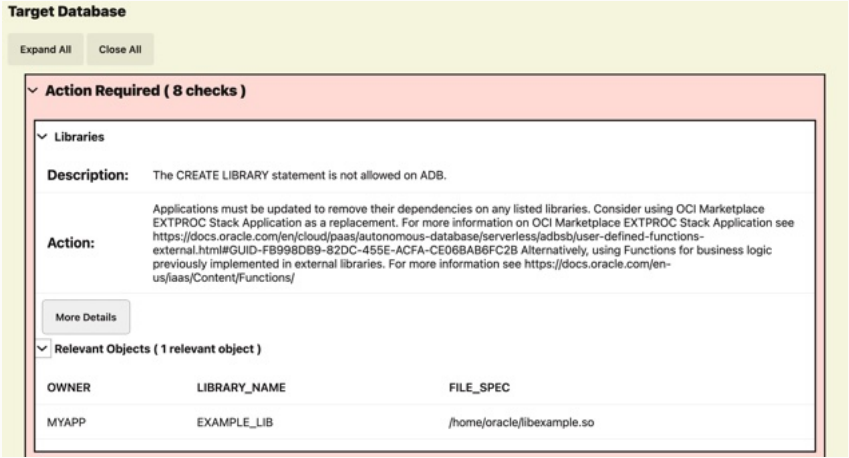

System Privileges

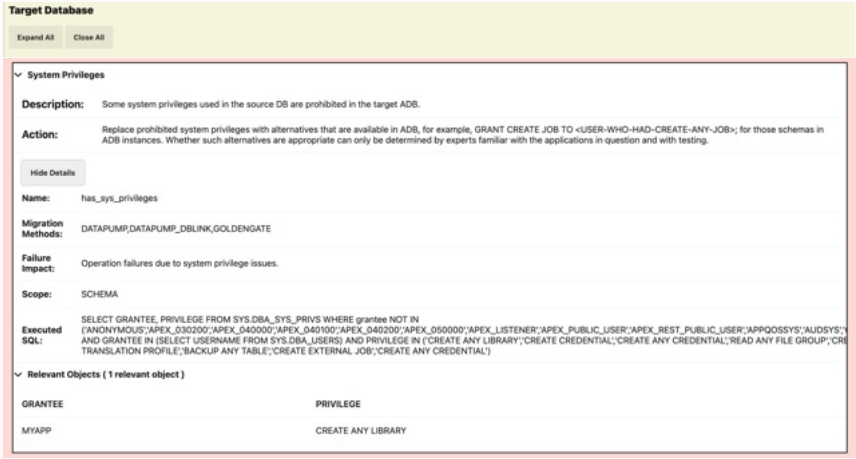

XML Objects

XML Tables

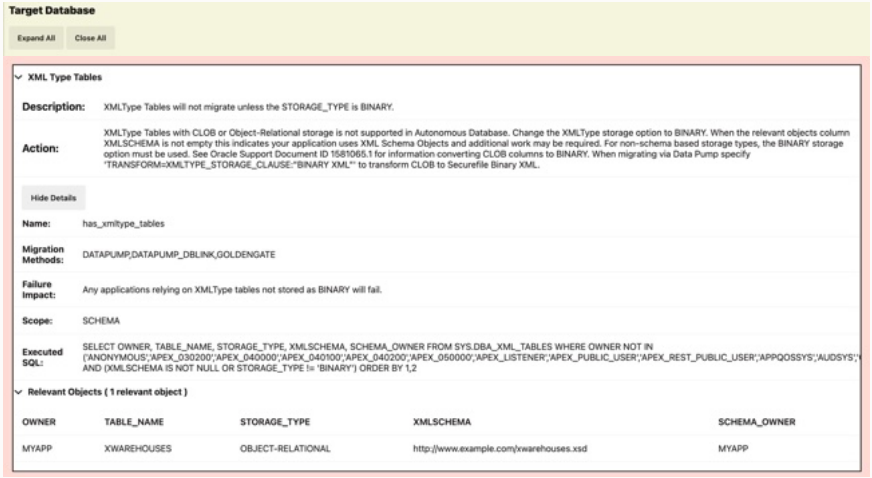

XML Columns

Java Sources

Java Objects

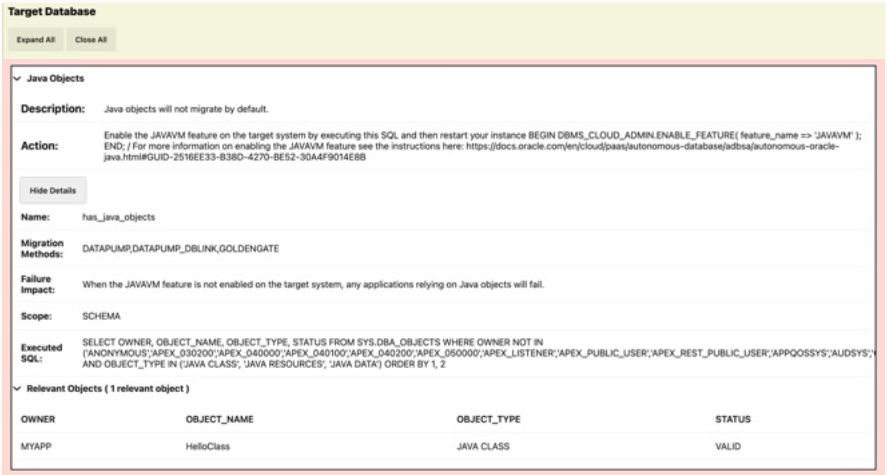

Common Objects

Non-Exported Objects Grants

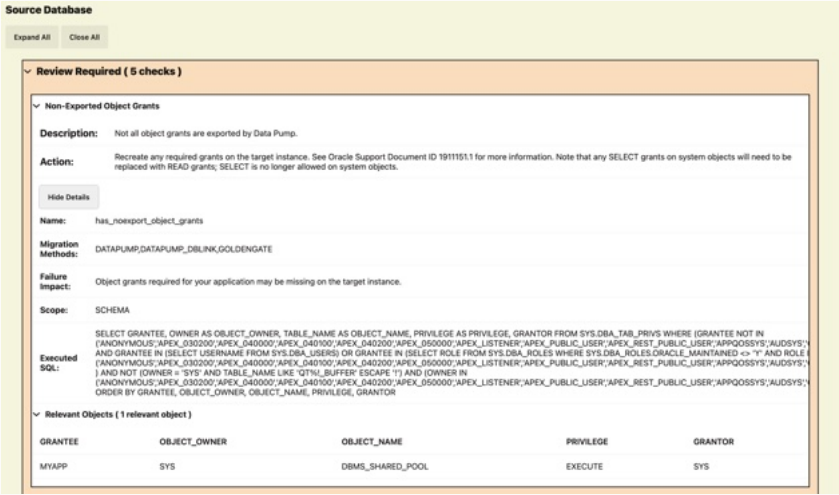

Scheduler Jobs

External Tables

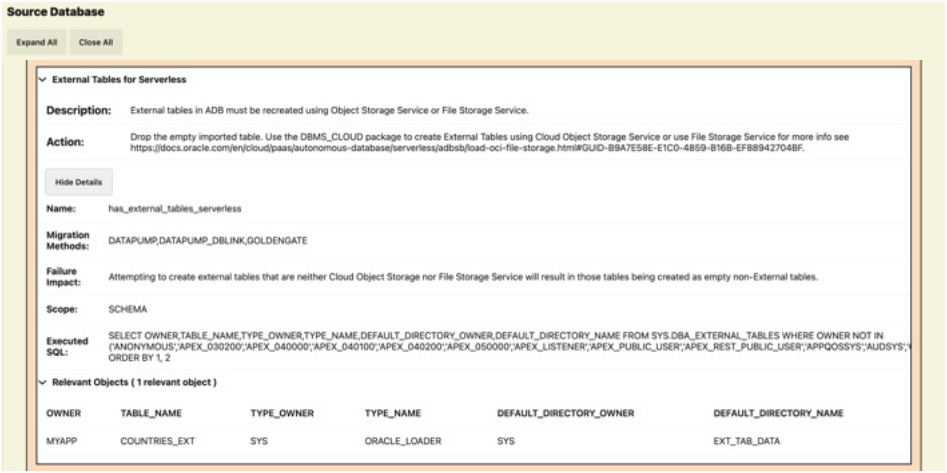

Restricted Packages

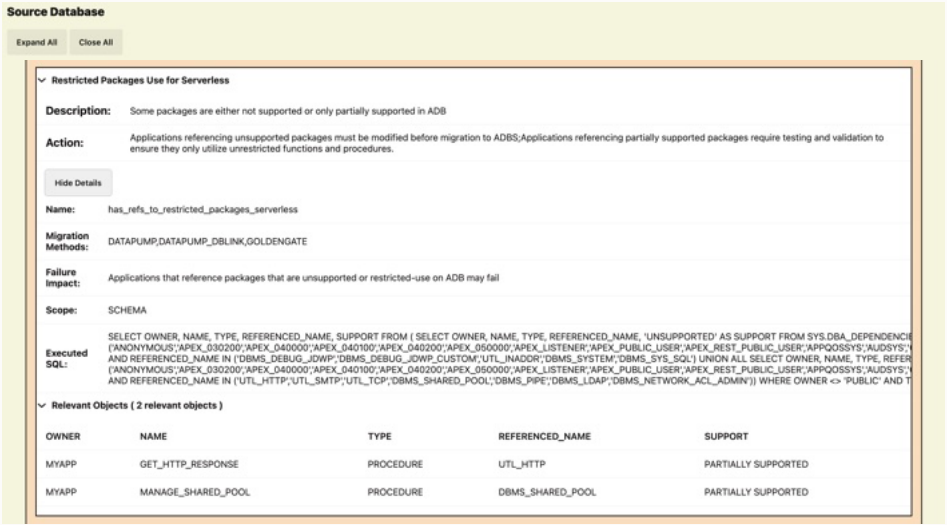

Directory Objects Accessibility

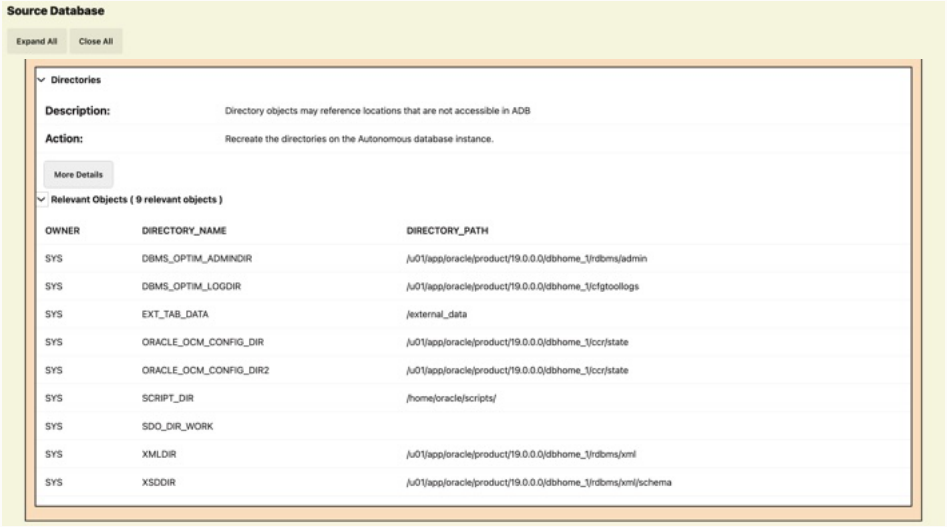

Connections

What's the best option?

Predefined Database Service Names

Autonomous DWH

dbname_high
dbname_medium
dbname_low

Autonomous TP

dbname_tpurgent
dbname_tp
dbname_high
dbname_medium
dbname_low

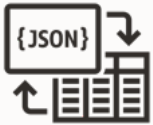

Autonomous JSON

dbname_tpurgent
dbname_tp
dbname_high
dbname_medium
dbname_low

Notes for Importing with Oracle Data Pump

https://docs.oracle.com/en/cloud/paas/autonomous-database/serverless/adbsb/load-data-data-pump-notes.html

“For the best import performance use the **HIGH** database service for your import connection and set the parallel parameter to one quarter the number of ECPUs (.25 x ECPU count).

If you are using OCPU compute model, set the parallel parameter to the number of OCPUs (1 x OCPU count).”

Connections

For migrations, clearly **TPURGENT** usually provides best results

- For Autonomous Data Warehouse (ADW), you can request this service by creating a Service Request (SR)

Concurrency

Database Service Name	Concurrent Statements with Compute Auto Scaling Disabled	Concurrent Statements with Compute Auto Scaling Enabled
tpurgent	75 × number of ECPUs	75 × number of ECPUs
tp	75 × number of ECPUs	75 × number of ECPUs
high	3	9
medium	0.25125 × number of ECPUs A decimal result is truncated.	0.75375 × number of ECPUs A decimal result is truncated.
low	75 × number of ECPUs	75 × number of ECPUs

https://docs.oracle.com/en/cloud/paas/autonomous-database/serverless/adbsb/manage-service-concurrency.html#GUID-6E4DCD27-CDAA-432D-A90B-485C19EF72B0

Concurrency Limits

Database Service Name	Concurrent Statements with OCPU Auto Scaling Disabled	Concurrent Statements with OCPU Auto Scaling Enabled
tpurgent	300 × number of OCPUs	300 × number of OCPUs
tp	300 × number of OCPUs	300 × number of OCPUs
high	3	9
medium	1.26 × number of OCPUs	3.78 × number of OCPUs
low	300 × number of OCPUs	300 × number of OCPUs

https://docs.oracle.com/en/cloud/paas/autonomous-database/serverless/adbsb/manage-service-concurrency.html#GUID-6E4DCD27-CDAA-432D-A90B-485C19EF72B0

Storage

What's the best option?

NFS vs Object Storage?

- It depends …

NFS vs Object Storage

Usually, NFS (or FSS) is recommended over object storage for the migration

- Easy to setup
- Performance seems to be better
- NFS requires "private endpoint access only"
- Object storage's advantage: pre-authenticated

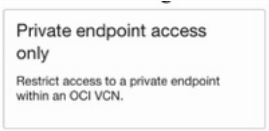

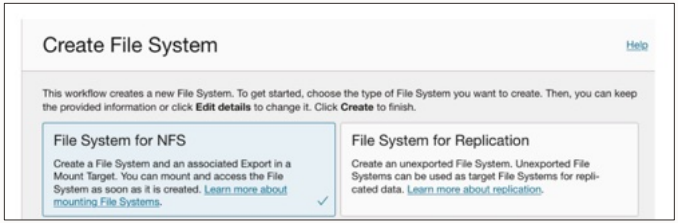

High Performance Storage

Use HPMT (High Performance Mount Target)

OPTION	THROUGHPUT	INCLUDED
HPMT-20	up to 20 Gbs	20 TB
HPMT-40	up to 40 Gbs	40 TB
HPMT-80	up to 80 Gbs	80 TB

- https://docs.oracle.com/en-us/iaas/releasenotes/filestorage/high-performance-mount-targets.htm
- https://docs.oracle.com/en-us/iaas/Content/Resources/Assets/whitepapers/file-storage-performance-guide.pdf

Use High Performance Mount Target during migration for better throughput

- 30 days minimum subscription

Latency

Have a closer look

Connection and latency test tool - adbping (Doc ID 2863450.1)

- Diagnose high-latency issues in customer workloads
- Validate if the ADB service is healthy and not the root cause of latency

Latency

```
[oracle@db19caz adbping]$ ./adbping -u admin -p                \
> -s atpaz_high -w \
> /home/oracle/wallet -j /u01/app/oracle/product/19.0.0/dbhome_1/javavm/jdk/jdk8 \
> -d 30
+++Test Summary+++
   Test Client: sqlplus
   Number of concurrent threads: 1
   Duration (secs): 30
   SQL executed: select 1 from dual;
   Pass: 27 Fail: 0
   Test start date: 2025-06-02 15:25:21.517810+00:00
   Test end date: 2025-06-02 15:25:52.126871+00:00
   SQL Execution Time(ms) : Min:0 Max:10 Avg:1.111 Median:0 Perc90:10 Perc95:10 Perc99:10
   Connect + SQL Execution Time(ms) : Min:1115.78 Max:1292.391 Avg:1131.607 Median:1123.914 Perc90:1155.317 Perc95:1155.977 Perc99:1292.391

Interpretation of the results
------------------------------

       1. Pass/Fail count: Indicates the total number of connections passed/failed in defined duration by the defined number of threads.

       2. SQL execution time: Time taken to just execute the SQL. Connection time not included.
          For sqlplus, this would be the elapsed time reported by sqlplus.

       3. Connect + SQL Execution Time: Time taken to connect and execute SQL.
          For sqlplus, this would be the time to connect and run the sql.
          For java, it would be time taken to getConnection() and execute the query.

       4. Java connection pool stats: Reports the time taken to setup the java connection pool and the initial and max size.
          All query executions do a getConnection() and execute the SQL.

       5. Perc90, Perc95, Perc99: This is the percentile value indicating 90%, 95% or 99% of the latencies are below the respective value.
```

Latency

```
[oracle@db19caz adbping]$ ./adbping -u admin -p                \
> -s atpaz_high -w \
> /home/oracle/wallet -j /u01/app/oracle/product/19.0.0/dbhome_1/javavm/jdk/jdk8 \
> -d 30
+++Test Summary+++
   Test Client: sqlplus
   Number of concurrent threads: 1
   Duration (secs): 30
   SQL executed: select 1 from dual;
   Pass: 27 Fail: 0
   Test start date: 2025-06-02 15:25:21.517810+00:00
   Test end date: 2025-06-02 15:25:52.126871+00:00
   SQL Execution Time(ms) : Min:0 Max:10 Avg:1.111 Median:0 Perc90:10 Perc95:10 Perc99:10
   Connect + SQL Execution Time(ms) : Min:1115.78 Max:1292.391 Avg:1131.607 Median:1123.914
Perc90:1155.317 Perc95:1155.977 Perc99:1292.391
```

Latency

```
[oracle@db19caz oci]$ ./connping11 -l admin/                @atpaz_high --period=5

RWP*Connect/OCIPing Release 3.2.1.0 Production on Mon, 02 Jun 2025 15:32:39 UTC
Connected default database with reconnect to:
Oracle Database 23ai Enterprise Edition Release 23.0.0.0.0 - for Oracle Cloud and Engineered Systems
connect:114.25 ms, ociping:0.959 ms, dualping:0.998 ms, sid=55917, inst#=8, time=1.1
connect:117.93 ms, ociping:1.004 ms, dualping:0.997 ms, sid=43883, inst#=8, time=2.1
connect:115.45 ms, ociping:1.120 ms, dualping:1.138 ms, sid=13832, inst#=8, time=3.1
connect:115.62 ms, ociping:1.206 ms, dualping:1.357 ms, sid=13832, inst#=8, time=4.1
connect mean=115.81, stddev=1.33, min=114.25, max=117.93
ociping mean=1.07, stddev=0.10, min=0.96, max=1.21
dualping mean=1.12, stddev=0.15, min=1.00, max=1.36
```

Maintenance

What else do you need to take care on?

Choose Maintenance Plan

Two maintenance plans are available

- You can adjust it later on if needed

Choose Maintenance Plan

Maintenance happens weekly

- Typically, on weekends
- Window will be always the same
- Assigned during creation

Take note for your migration!!

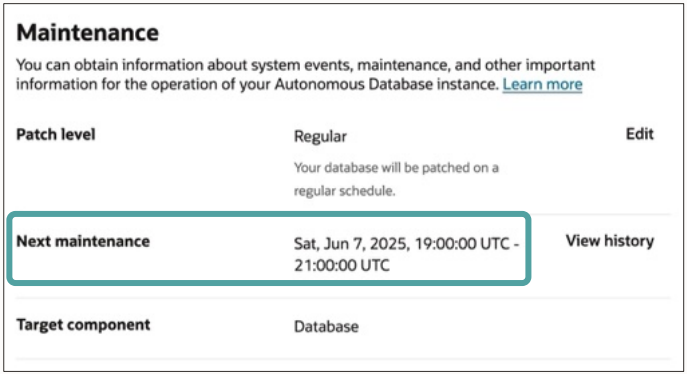

```
--Query your recent maintenance window on ADB Serverless
--Note: Exact timing for upcoming maintenance
--       window is available 24 hours before.

select actual_start_date, actual_end_date,
       maintenance_status
  from db_notifications
  where type = 'maintenance' and time > sysdate -7;

ACTUAL_START_DATE          ACTUAL_END_DATE            MAINTENANCE_STATUS
-------------------------  -------------------------  ---------------------
2025-05-24 19:00:02 GMT    2025-05-24 21:00:24 GMT    Completed
```

If you want to change the maintenance window, you must open an SR

- Only ±2 hours possible

Data Pump

The simple approach

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

- You may request one-off fixes via an SR

Allocate a sufficient number of ECPUs

- 32 should be the minimum when you import

Export: **PARALLEL** 2x of physical cores

Import: **PARALLEL**=ECPU/4, or higher

- Scale up to the maximum for migrations

Ensure CLUSTER_DATABASE=TRUE

- Allows Data Pump workers across nodes

```
NAME             TYPE    VALUE
---------------- ------- -----
cluster_database boolean TRUE
```

Most simple method: Data Pump

Datapump with Files

Datapump with DB Links

Datapump with Dump Files

- More control over parallelism
- Storage Overhead
- No source-target connection interoperability requirement
- Requires Object Storage / File Storage setup

Datapump with DB Links

- Network throughput and latency dependency
- Faster for smaller databases
- Requires DB Link setup
- And there is more ...

Be aware of network link import limitations

- May have a significant impact on performance

Network Link Imports - LOBs

- A network round trip is required for reach row with a LOB
- If you have millions of LOB rows and a high latency connection to ADB, this may have a significant negative impact
- Check your latency from source database to ADB instance using
  - How to measure network latency for Oracle Database applications in OCI (Doc ID 3008087.1)
  - Connection and latency test tool - adbping (Doc ID 2863450.1)

```
./ociping -l user/password@myadb_high --period=5
RWP*OCIPing Release 3.2.1.0 Production on Fri, 30 May 2025 15:23:35 UTC
Connected default database to:
Oracle Database 23ai Enterprise Edition Release 23.0.0.0.0 - for Oracle
Cloud and Engineered Systems
0.998 0.0
1.008 1.0
0.987 2.0
0.999 3.0
1.054 4.0
ociping (ms) mean=1.009, stddev=0.023, min=0.987, max=1.054
```

Network Link Imports - Metadata

- A network link import does not import metadata in parallel
- On complex schemas this may have a significant negative impact
- If a network link import is required,
  - Import metadata before migration - `CONTENTS=METADATA_ONLY`
  - Load rows only during the migration - `CONTENTS=DATA_ONLY`

ADB Migration Overview

Source: 19c IaaS

- Based on sample schemas 19.2
- HR, PM, IX, SH and PI
- Some additional objects

ATP Serverless
Oracle19.27

ATP Dedicated
Oracle 19.26

ADB Dedicated | Migration Example

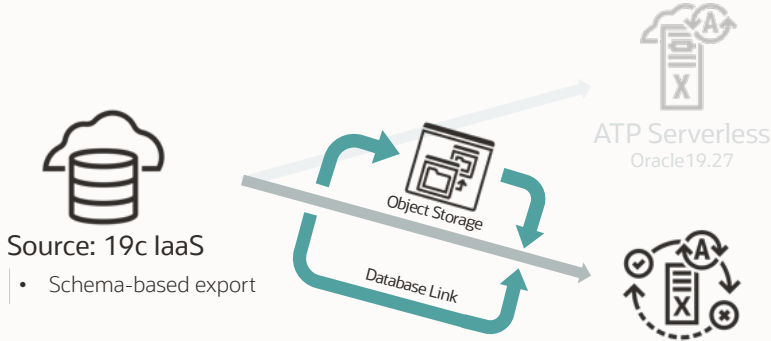

ADB Dedicated | Demo

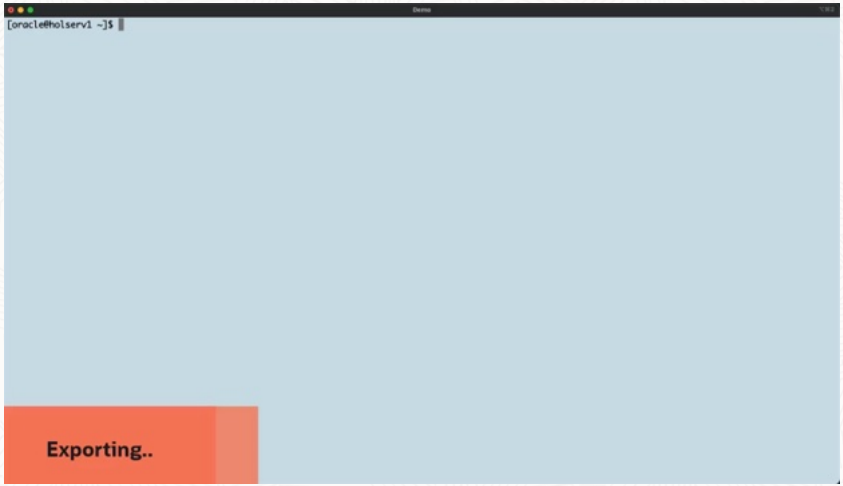

ADB Serverless | Migration Example

Source: 19c IaaS

- Schema-based export

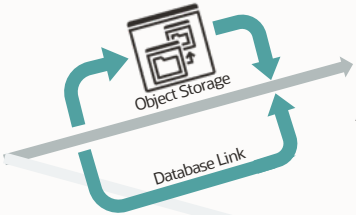

ATP Serverless

Oracle19.27

- Fixing some objects

ATP Dedicated

Oracle 19.26

ADB Serverless | Demo

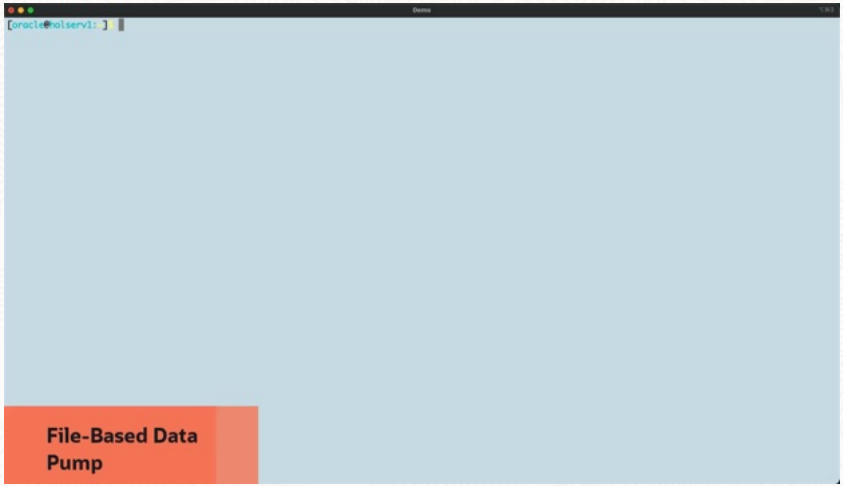

Automation

How AMA scripts ease migrations

What is AMA?

Autonomous Database Migration Automation (AMA)

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

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

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

You won't start with filming right away

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

Now you can start filming your scenes

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

AMA Workflow

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

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

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

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

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


Parameter File

Adjust;

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

Documentation

AMA Documentation is available at request

- Documents the entire flow and all options and parameters

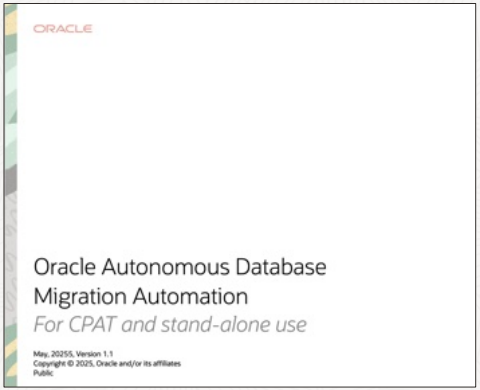

AMA Demonstration

Part 1 - Configuration

AMA Workflow

AMA | Planning Phase

On-Prem - Source

- Gather stats for SYS / SYSTEM
- Create Migration user
- Enable restricted session
- Set JOB_QUEUE_PROCESSES=0

ADB-S - Target

- Enable OLAP / JAVA in ADB-S

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

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

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

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

AMA Workflow

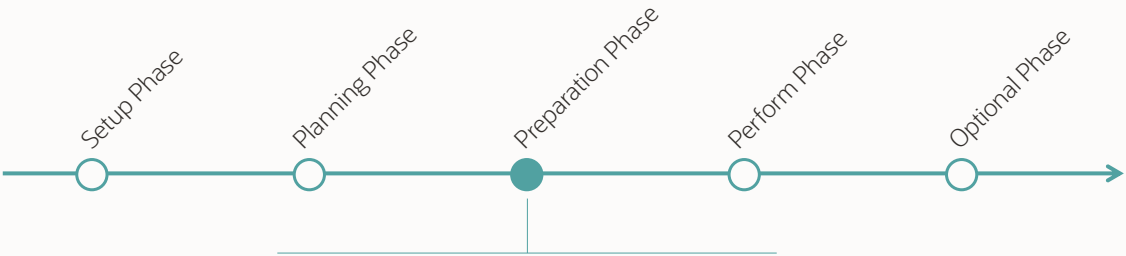

AMA | Preparation Phase

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

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

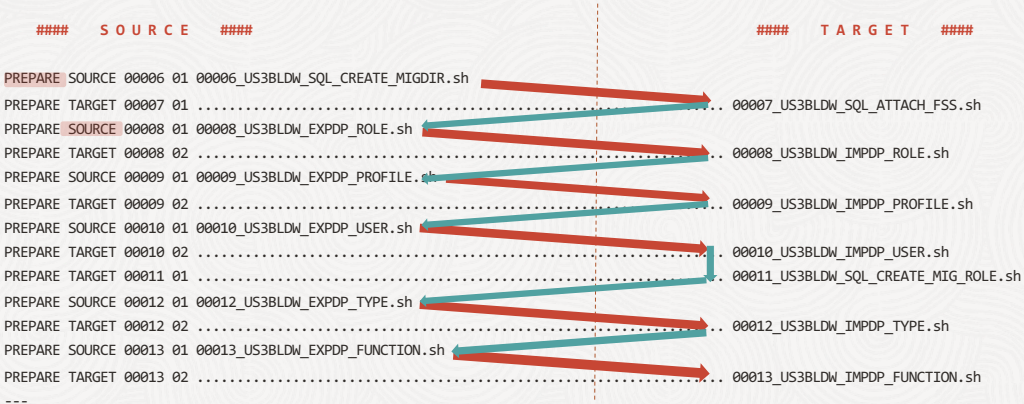


```
---
```

AMA Workflow

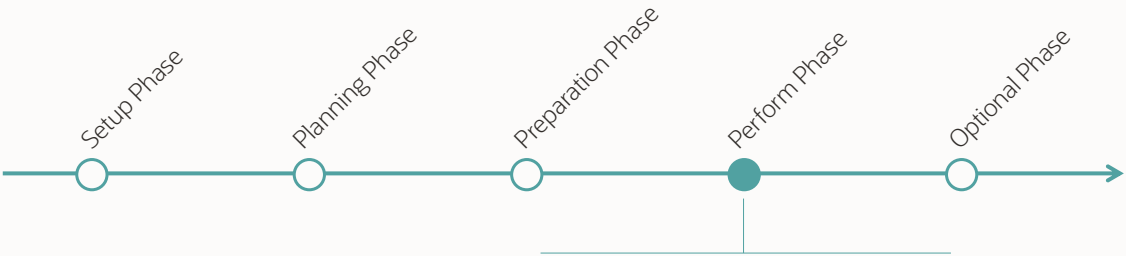

AMA | Perform Phase

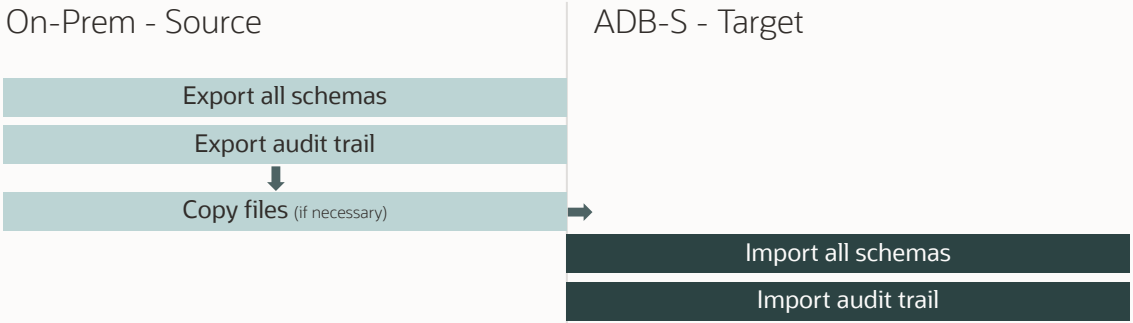

AMA Workflow

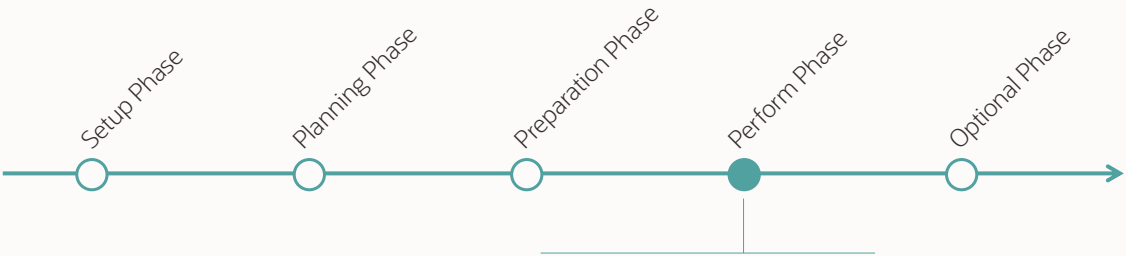

```
---
--- PERFORM PHASE ---
--- Commonly in this phase nothing depends on each other (except you for example have objects that depend on objects stored in other schema) ---
--- So export jobs can be started in parallel and imports once the export finished ---
---
---
```

#### S O U R C E ####	#### T A R G E T ####
PERFORM SOURCE 00014 01 00014_US3BLDW_EXPDP_AUDIT_TRAILS.sh	
PERFORM TARGET 00014 02	00014_US3BLDW_IMPDP_AUDIT_TRAILS.sh
PERFORM SOURCE 00015 01 00015_US3BLDW_SQL_GEN_SQL_PROFILE_STAGE_TAB.sh	
PERFORM SOURCE 00015 02 00015_US3BLDW_EXPDP_SQL_PROFILES.sh	
PERFORM TARGET 00015 03	00015_US3BLDW_IMPDP_SQL_PROFILES.sh
PERFORM TARGET 00015 04	00015_US3BLDW_SQL_APPL_SQL_PROFILE_STAGE_TAB.sh
PERFORM SOURCE 00016 01 00016_US3BLDW_SQL_GEN_SQL_PATCHES_STAGE_TAB.sh	
PERFORM SOURCE 00016 02 00016_US3BLDW_EXPDP_SQL_PATCHES.sh	
PERFORM TARGET 00016 03	00016_US3BLDW_IMPDP_SQL_PATCHES.sh
PERFORM TARGET 00016 04	00016_US3BLDW_SQL_APPL_SQL_PATCHES_STAGE_TAB.sh
PERFORM SOURCE 00017 01 00017_US3BLDW_EXPDP_SCHEMA_FUSION.sh	
PERFORM TARGET 00017 02	00017_US3BLDW_IMPDP_SCHEMA_FUSION.sh
PERFORM SOURCE 00018 01 00018_US3BLDW_EXPDP_SCHEMA_FUSION_OCSERVER11G.sh	
PERFORM TARGET 00018 02	00018_US3BLDW_IMPDP_SCHEMA_FUSION_OCSERVER11G.sh

```
...
```

AMA | Perform Phase

On-Prem - Source

ADB-S - Target

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

Export network ACLs

- Import network ACLS
- Enable Advanced Queues
- Recompilation

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

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

```
---
---
--- END OF MIGRATION ---
```

AMA Workflow

AMA Demonstration

Part 2 - Migration

Done!!

AMA can run a migration fully automated and completely unattended

Works with MS Windows as source database

Database links, directories, external tables, XML binary objects, APEX applications

- Work-in-progress

ZDM

Zero Downtime Migration

Oracle Zero Downtime Migration

Watch on YouTube

DMS

Data Migration Service

Oracle Zero Downtime Migration

Watch on YouTube

ADB@Azure

Using ZDM to migrate to Azure

Autonomous Database @ Azure

Watch on YouTube

Autonomous Database @ Azure

Zero Downtime Migration documentation

Exploring NFS Storage Options for Oracle ZDM Migrations to Oracle Database@Azure

Network topology and connectivity for Oracle Database@Azure - Migration connectivity design

Step-by-step Guide: Logical **Offline** Migration to ADB-S on Oracle Database@Azure
Step-by-step Guide: Logical **Online** Migration to ADB-S on Oracle Database@Azure

Success

Validating a migration

How can you proof that no data was lost during the migration?

Validating the Migration

It's a logical migration into a different database platform

- ADB performs a lot of transformations:
  IOTs to heap tables, external tables to heap tables
- Database links are different:
  Uses credentials and a different connection string
- Different kinds of storage:
  tablespace, removing table storage customization,
  changing to binary XML etc.
- And so on ...

Start by validating the Data Pump export and import

```
Import: Release 19.0.0.0.0 - Production on Mon Apr 28 08:49:43 2025
Version 19.27.0.0.0
Copyright (c) 1982, 2019, Oracle and/or its affiliates.  All rights reserved.

Connected to: Oracle Database 19c Enterprise Edition Release 19.0.0.0.0 - Production
28-APR-25 08:49:44.678: W-1 Startup took 0 seconds
.
.
.
.
28-APR-25 08:51:05.528: Job "DPUSER"."ADB_MIGR" successfully completed at Mon Apr 28 08:51:05 2025 elapsed 0 00:01:21
```

```
Import: Release 19.0.0.0.0 - Production on Mon Apr 28 08:49:43 2025
Version 19.27.0.0.0
Copyright (c) 1982, 2019, Oracle and/or its affiliates.  All rights reserved.

Connected to: Oracle Database 19c Enterprise Edition Release 19.0.0.0.0 - Production
28-APR-25 08:49:44.678: W-1 Startup took 0 seconds
.
.
.
.
28-APR-25 08:51:05.528: Job "DPUSER"."ADB_MIGR" successfully completed at Mon Apr 28 08:51:05 2025 elapsed 0 00:01:21
28-APR-25 08:51:05.528: Job "DPUSER"."ADB_MIGR" completed with 56 error(s) at Mon Apr 28 08:51:05 2025 elapsed 0 00:01:21
```

```
--Is this an ignorable error?

ORA-31684: Object type USER:"APPUSER" already exists
```


```
--Is this an ignorable error?

ORA-39082: Object type VIEW:"APPUSERS"."MyCaseSensitiveView"
created with compilation warnings
```


```
--Is this an ignorable error?

ORA-01653: unable to extend table APPUSER.T1 by 8192
ORA-39171: Job is experiencing a resumable wait
```


```
--Is this an ignorable error?

ORA-12899: value too large for column (actual: 3, maximum: 2)
```


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

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

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


```
$ python3 dpla.py import.log

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

...

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

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

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

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

...

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

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

Data Pump Log Analyzer

Table Details

Search for Table...

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

Data Pump Log Analyzer

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

Also, usable for diagnostics and performance tuning

Then, validate your database

Validate Your Database

1. Objects
2. Rows
3. Data

Validate Your Database

1

Objects

1. Recompile invalid objects

2. Compare number of objects

```
--Generate a list of objects in the source
--subtract the objects in target to find missing objects

select owner, object_type, object_name, status
from   dba_objects@sourcedb

minus

select owner, object_type, object_name, status
from   dba_objects;
```


```
--Constraints are not listed in DBA_OBJECTS

select   owner, table_name, count(table_name)
from     dba_constraints@sourcedb
where    constraint_name not like 'BIN%'
group by owner, table_name

minus

select   owner, table_name, count(table_name)
from     dba_constraints
where    constraint_name not like 'BIN%'
group by owner, table_name;
```


Using Advanced Queueing

- AQ creates some queue structures on demand only
- Blog post

Validate Your Database

Source database

Queue table

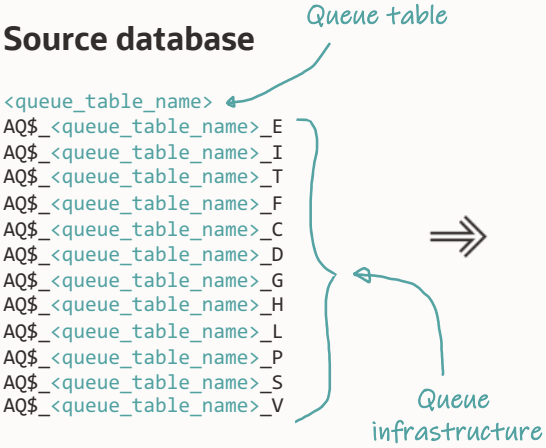

Target database

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

Take into account in comparing source and target database object count

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

Other objects also change
- like database links and directories

Validate Your Database

- Validate objects using
  - `@?/rdbms/admin/utlrp`, or
  - `@?/rdbms/admin/utlprp n`
- Use `DBA_ERRORS` to find cause of invalidation
- What Objects Are Created When Creating a Queue Table ? (Doc ID 224027.1)
- Things to Consider When Importing Advanced Queues using Oracle Data Pump

Validate Your Database

1 Objects

2 Rows

3 Data

Validate Your Database

2

Rows

1. Compare number of rows exported and imported

2. Count and compare number of rows

Data Pump keeps tracks of unloaded and loaded rows

```
cat export.log

;;;
Export: Release 19.0.0.0.0 - Production on Mon Jun 2 13:57:40 2025
Version 19.27.0.0.0

Copyright (c) 1982, 2019, Oracle and/or its affiliates.  All rights reserved.
Connected to: Oracle Database 19c Enterprise Edition Release 19.0.0.0.0 - Production
.
.
.
.
Processing object type SCHEMA_EXPORT/TABLE/CONSTRAINT/CONSTRAINT
. . exported "APPUSER"."T1"                              119.0 KB        105 rows
. . exported "APPUSER"."T2"                              13.29 GB   32411047 rows
. . exported "APPUSER"."T3"                              35.20 GB   78910231 rows
. . exported "APPUSER"."T4"                              19.02 MB      57174 rows
Master table "DPUSER"."SYS_EXPORT_SCHEMA_02" successfully loaded/unloaded
******************************************************************************
Dump file set for DPUSER.SYS_EXPORT_SCHEMA_02 is:
  /home/oracle/dpdir/appuser.dmp
Job "DPUSER"."SYS_EXPORT_SCHEMA_02" successfully completed at Mon Jun 2 13:58:05 2025 elapsed 0 00:00:23
```

```
grep -w exported export.log | grep -w rows | awk '{print $4,$7}' > exp.txt

"APPUSER"."T1" 105
"APPUSER"."T2" 32411047
"APPUSER"."T3" 78910231
"APPUSER"."T4" 57174
```


```
grep -w exported export.log | grep -w rows | awk '{print $4,$7}' > exp.txt

grep -w imported import.log | grep -w rows | awk '{print $4,$7}' > imp.txt

diff exp.txt imp.txt
```


How do you validate the row count when using Oracle GoldenGate?

```
spool count_source.log
select /*+ parallel */ 'APPUSER.T1 ' || count(1) from appuser.t1;
select /*+ parallel */ 'APPUSER.T2 ' || count(1) from appuser.t2;
select /*+ parallel */ 'APPUSER.T3 ' || count(1) from appuser.t3;
select /*+ parallel */ 'APPUSER.T4 ' || count(1) from appuser.t4;
```


Counting Rows

- Requires either a full table or index scan
- Counting rows is usually faster with an index on a NOT NULL column
- Don't use parallel query on small tables
- The bigger the database, the longer it takes

Validate Your Database

1 Objects

2 Rows

3 Data

Validate Your Database

3

Data

1. Ensure data matches on source and target

2. Different techniques
   - Oracle GoldenGate Veridata
   - `DBMS_COMPARISON`
   - `DBMS_CRYPTO` and `STANDARD_HASH`

Validate Your Database

- Oracle GoldenGate Veridata
- Details on DBMS_COMPARISON, DBMS_CRYPTO and STANDARD_HASH

Do you have time for a full verification in your migration window?

Full Verification

Test

Test

Final test

- Full verification
- Use refined runbook
- No further changes to runbook

Migration

- Partial verification
- Use same runbook

Partial Verification

You decide on the scope of the partial verification

- Tolerated and maximum downtime
- Data criticality
- Verify entire table or a sample
- Business requirements
- Audit requirements
- Regulations

Partial Verification

Build a plan that satisfies the requirements:

• APPUSER.T1	Business data	Sample 10 %
• APPUSER.T2	Insert-only table	No verification
• APPUSER.T3	Generated data	No verification
• APPUSER.T4	Regulated data	Full verification
• APPUSER.T5	Rarely updated	Verify last 3 months of data

Be sure to automate your verification

- Save the output and log files

Try it out, please!!

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

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

YouTube | @UpgradeNow

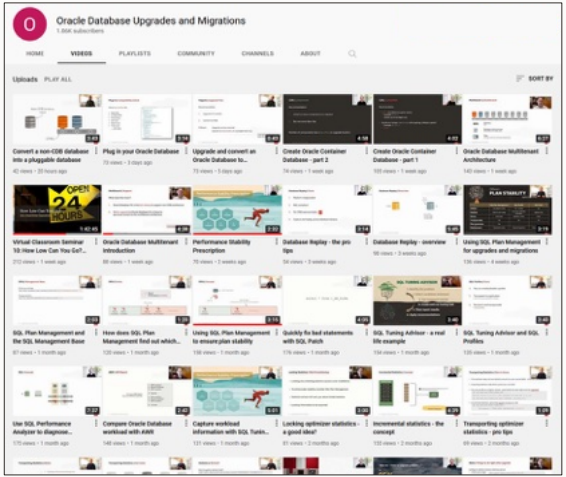

Link

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

SCAN ME

Thank You