



ORACLE

Move to the Cloud

Not only for Techies ...





Roy F. Swonger

Vice President
Database Upgrade,
Utilities & Patching

 [@RoyFSwonger](https://twitter.com/RoyFSwonger)



Mike Dietrich

Distinguished Product Manager
Database Upgrade
and Migrations

 <https://MikeDietrichDE.com>

 @MikeDietrichDE

 mikedietrich



Daniel Overby Hansen

Senior Principal Product Manager
Cloud Migration

 <https://dohdatabase.com>

 [@dohdatabase](https://twitter.com/dohdatabase)

 [dohdatabase](https://www.linkedin.com/company/dohdatabase)



William Beauregard

Senior Principal Product Manager
Data Pump and Utilities

 [william-beauregard-3053791](#)



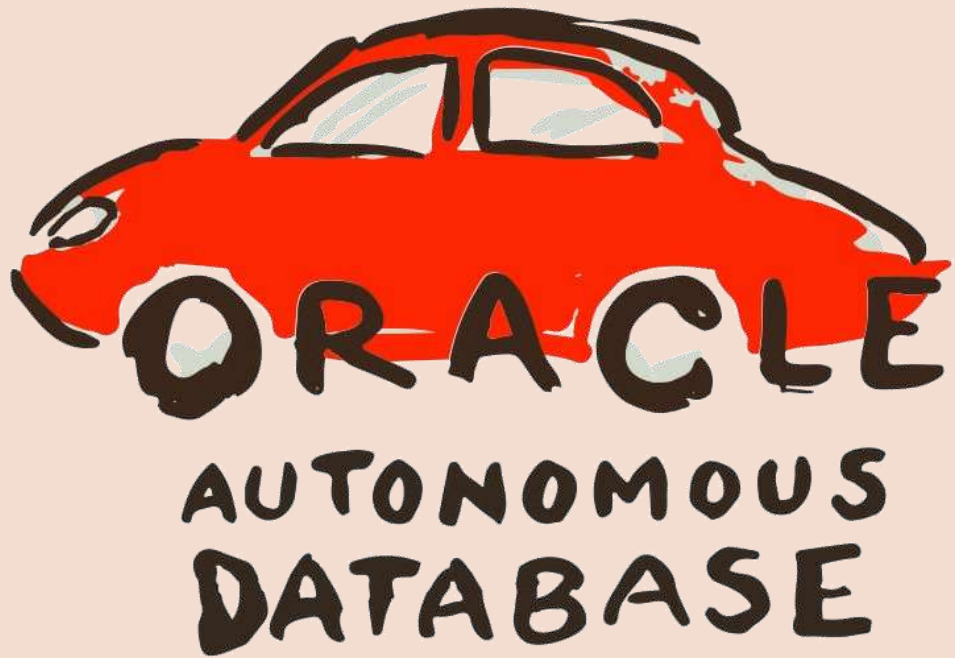
Rodrigo Jorge

Senior Principal Product Manager
Database Patching and Upgrade

 <https://dbarj.com.br/en>

 @rodrigojorgedba

 rodrigoaraujorge



Autonomous Database

How to Move Your Data into an
Autonomous Database

”An autonomous database is a **cloud database** that uses machine learning to **automate** database tuning, security, backups, updates, and other routine management tasks traditionally performed by DBAs.”

[What Is an Autonomous Database – Oracle.com](#)

Autonomous Database | Infrastructure



Shared ——— A simple and elastic choice. Oracle autonomously operates all aspects of the database life cycle from database placement to backup and updates.

Dedicated ——— A private cloud in public cloud choice. A completely dedicated compute, storage, network and database service for only a single tenant. Dedicated infrastructure provides for the highest levels of security isolation and governance. The customer has customizable operational policies to guide Autonomous Operations for workload placement, workload optimization, update scheduling, availability level, over provisioning and peak usage



Autonomous Database | Restrictions

Restrictions apply

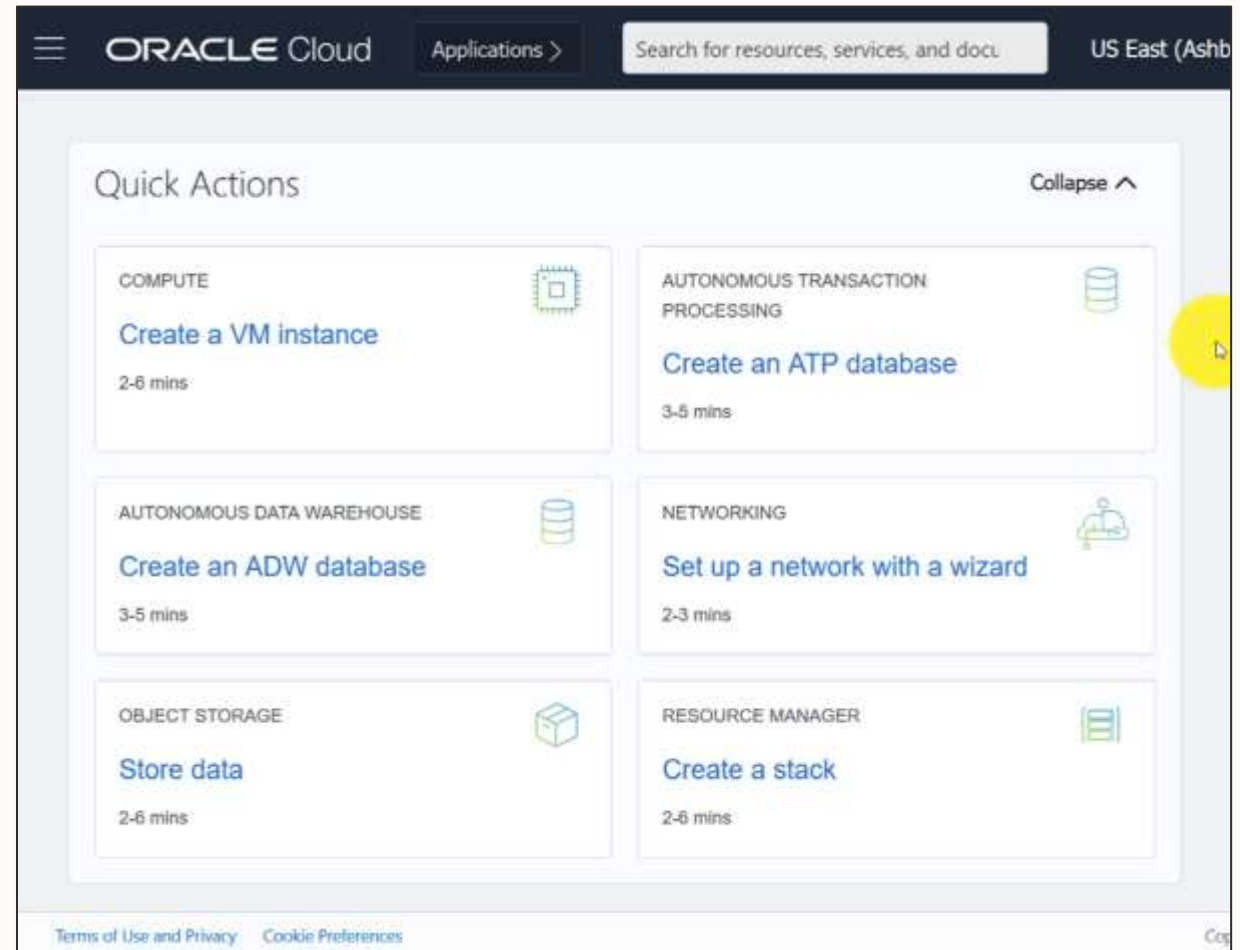
- Parameters
- SQL commands
- Data types
- Features

Pro tip: Cloud services evolve rapidly, check the [documentation](#) for up-to-date information



Autonomous Database | Always Free

- Free of charge
- Usage:
 - Small-scale apps
 - Development
 - Testing
- Quick start showcase
 - [MuShop](#)



[Watch on YouTube](#)



”When it comes to migration, Autonomous Database differs because you will be **migrating data**, not databases”

Autonomous Database | Migration Techniques

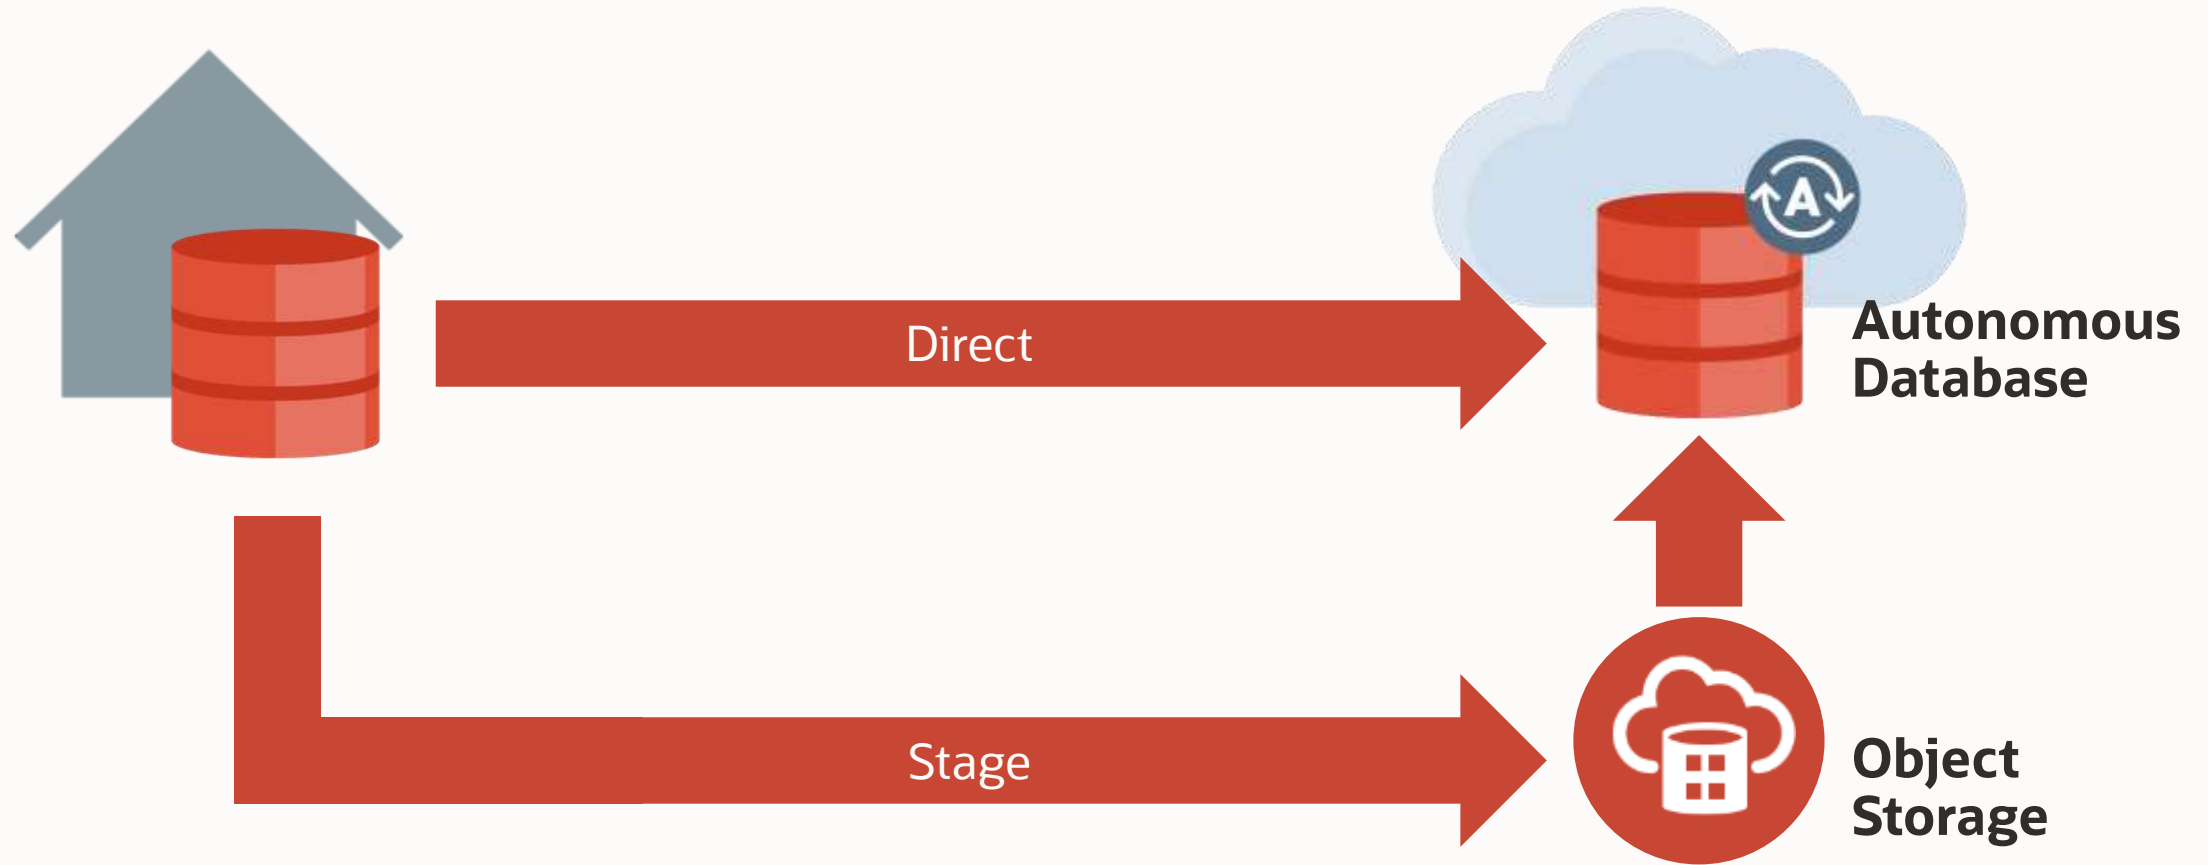




Photo by [Alejandro Morelos](#) on [Unsplash](#)

Importing

Death Metal

Demo | Data Set



The screenshot shows a Kaggle dataset page for 'Death Metal'. The title is 'Death Metal' with the subtitle 'Death metal bands and albums'. It was updated 4 years ago (Version 2) by user 'zjf'. The page features a navigation bar with 'Data', 'Tasks', 'Notebooks (21)', 'Discussion (2)', 'Activity', and 'Metadata'. There are buttons for 'Download (71 MB)' and 'New Notebook'. The 'Usability' score is 7.6. The 'Context' section describes the data source as 'Metal-Archives.com' and provides details about the search criteria for 'death metal' bands and albums. The description also mentions that the dataset banner is the cover art of the album 'Towards the Megalith' by the band Disma.

Downloading the Data Set

[Watch on YouTube](#)



Options | Overview



1.
SQL Developer Web

2.
SQL Developer

3.
SQL*Loader

4.
Data Pump

5.
DBMS_CLOUD

6.
MV2ADB

Options | **SQL Developer Web**

Easily accessible from OCI console

Quick and simple

Works on:

- CSV
- XML
- JSON
- XLS/XLSX
- Avro

[Quickstart lab](#)

[Jeff Smith blog post](#)



Options | SQL Developer Web



[Watch on YouTube](#)



Options | Overview



1.
SQL Developer Web

2.
SQL Developer

3.
SQL*Loader

4.
Data Pump

5.
DBMS_CLOUD

6.
MV2ADB

Options | **SQL Developer**

Local installation - [Download](#)

Quick and simple

Works on:

- CSV
- XML
- JSON
- XLS/XLSX
- Avro

Loads from:

- Local file
- OCI object stage

[Quickstart lab](#)



Options | SQL Developer



Loading Data with
SQL Developer

[Watch on YouTube](#)



Options | Overview



1.
SQL Developer Web

2.
SQL Developer

3.
SQL*Loader

4.
Data Pump

5.
DBMS_CLOUD

6.
MV2ADB

Options | **SQL*Loader**

Highly configurable

Can transform data

Loads from:


- Local file

Works on:

- CSV
- Text



Options | SQL*Loader



[Watch on YouTube](#)



Options | Overview



1.
SQL Developer Web

2.
SQL Developer

3.
SQL*Loader

4.
Data Pump

5.
DBMS_CLOUD

6.
MV2ADB

Options | Data Pump

Fast unload and load

All or selected data

Best option for big data

Can transform metadata

Loads from:

- Oracle Database (database link)
- OCI Object Storage

Pro tip: Use [SQL Developer](#)



Options | Data Pump

Parameter `MAX_DATAPUMP_PARALLEL_PER_JOB` is 50 -
and can't be changed

Network mode does not support
parallel meta data load



Options | Data Pump



DEMO

Upload into ADW with Data Pump

[Watch on YouTube](#)



Options | Overview



1.
SQL Developer Web

2.
SQL Developer

3.
SQL*Loader

4.
Data Pump

5.
DBMS_CLOUD

6.
MV2ADB

Options | **DBMS_CLOUD**

PL/SQL interface

Loads from:

- OCI Object Storage
- Amazon AWS S3
- Microsoft Azure Object Store

Works on:

- CSV
- Data Pump
- ORC
- Parquet
- Avro
- Zipped files

[Quickstart lab](#)



Options | **DBMS_CLOUD**

Troubleshooting

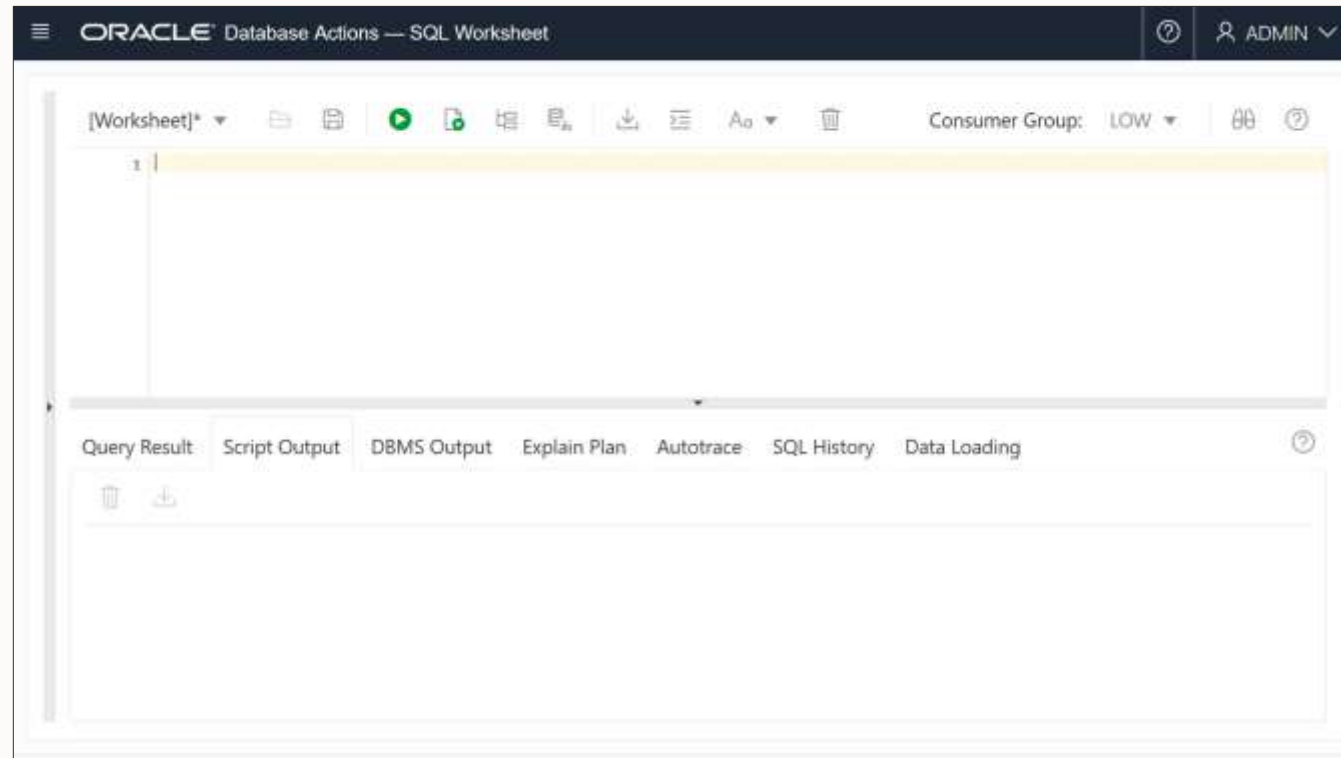
Coming to on-premises

Documentation

Use SQL Developer



Options | DBMS_CLOUD



[Watch on YouTube](#)



Options | Overview



1.
SQL Developer Web

2.
SQL Developer

3.
SQL*Loader

4.
Data Pump

5.
DBMS_CLOUD

6.
MV2ADB

Options | MV2ADB

"One button approach"

Uses Data Pump (schema mode)


Loads from:

- Oracle Database

Documentation: [Doc ID 2463574.1](#)

Support import over DB link (`--netlink`)

Runs on Linux / Solaris



Pro Tip:
Check log file to find out how MV2ADB uses
Data Pump

Options | MV2ADB



Enable Data Pump metrics (METRICS=Y and LOGTIME=ALL)

```
$ mv2adb.bin auto --conf mv2adb_metal.cfg --nosudo -dpdebug
```

Exclude statistics

```
EXTRA_EXPDP=EXCLUDE=STATISTICS
```

Enable Data Pump compression (license required)

```
COMPRESSION=ALL
```

Options | MV2ADB



```
oracle@hol:~$ more mv2adb_metal.cfg
```

[Watch on YouTube](#)



Options | Overview




	SQL Developer Web	SQL Developer	SQL*Loader	Data Pump	DBMS_CLOUD	MVZADB
DATA FORMAT						
CSV	x	x	x	x	x	
Text		x	x			
Text - advanced			x			
XML	x					
JSON	x					
Delimited	x	x	x			
Data Pump				x	x	x
Excel	x					
ORC					x	
Parquet					x	
Avro	x				x	

	SQL Developer Web	SQL Developer	SQL*Loader	Data Pump	DBMS_CLOUD	MVZADB
DATA SOURCE						
Local file	x	x	x	x	x	x
Object storage		x		x	x	x
Amazon S3					x	
Amazon S3 compatible source					x	
Azure BLOB					x	



Options | **Big Data**



Only one option: **Data Pump**

Follow recommendations and best practices

Options | Zero Downtime




ORACLE | GoldenGate

Only option: [Oracle GoldenGate](#)

ADB support from Oracle GoldenGate 12.3.0.1.2

Options | Zero Downtime




Zero Downtime Migration




Logical Online method using Data Pump + GoldenGate

GoldenGate | Explained

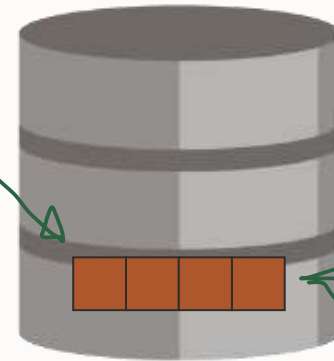


Statements captured



-  INSERT INTO ...
-  UPDATE ...
-  DELETE FROM ...

Statements replayed



Initial copy

-  INSERT INTO ...
-  UPDATE ...
-  DELETE FROM ...



GoldenGate | Cloud Native

New Cloud Native service: **OCI GoldenGate**

Runs GoldenGate 21c, **managed** by Oracle

Auto-scale: true cloud elasticity, low operations cost

Very **attractive** pricing

Supports:

- Oracle Database 11.2.0.4 and higher



Pro Tip: Watch a short intro on [YouTube](#)



GoldenGate | Cloud Native



ORACLE Cloud Applications > Germany Central (Frankfurt)

Get Started Dashboard

Quick Actions Collapse ^

- COMPUTE: Create a VM instance (2-6 mins)
- AUTONOMOUS TRANSACTION PROCESSING: Create an ATP database (3-5 mins)
- AUTONOMOUS DATA WAREHOUSE: Create an ADW database (3-5 mins)
- NETWORKING: Set up a network with a wizard (2-3 mins)
- RESOURCE MANAGER: Create a stack
- OBJECT STORAGE: Store data

Account Center

- User Management: Add a user to your tenancy
- Billing: Current billing cycle charges \$0.00, Days elapsed in billing cycle 29 / 31

What's New

Terms of Use and Privacy Cookie Preferences Copyright © 2021, Oracle and/or its affiliates. All rights reserved.

[Watch on YouTube](#)



GoldenGate | Certification Matrix

GoldenGate version: 19.1.0.0.200714

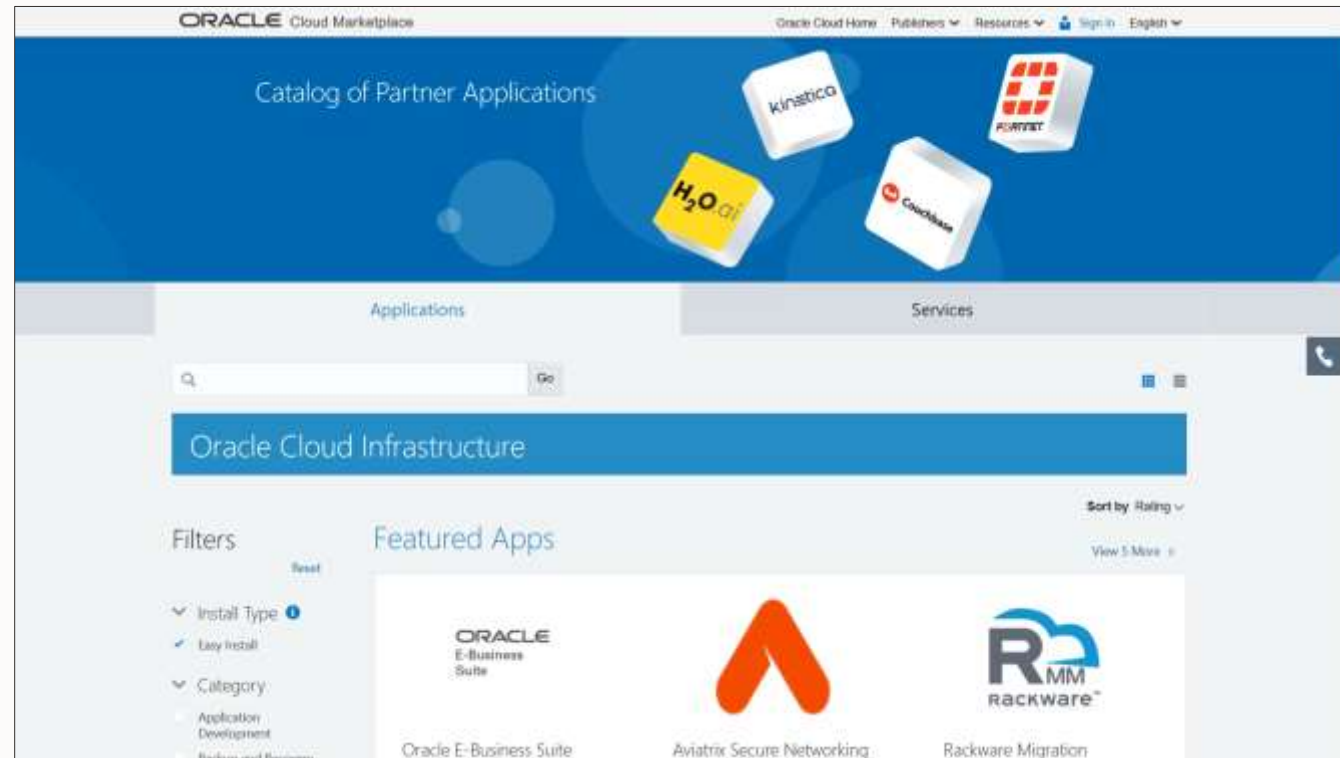
Processor Type	OS Version	OS Update Level	11.2.0.4	12.1.0.1	12.1.0.2	12.2	18	19
HP-UX Itanium (64-bit)	11,31	1409+	X	X	X	X	X	X
IBM AIX on POWER Systems (64-bit)	7,1	5.SP1+	X	X	X	X	X	X
IBM AIX on POWER Systems (64-bit)	7,2	2.SP1+	X		X	X	X	X
Linux on System Z (64-bit)	Red Hat Enterprise Linux 6	32+	X	X	X	X	X	
Linux on System Z (64-bit)	Red Hat Enterprise Linux 7		X		X	X	X	X
Linux on System Z (64-bit)	SLES 12	3+			X	X	X	X
Linux x86-64	Oracle Linux 6	4+	X	X	X	X	X	
Linux x86-64	Oracle Linux 7		X		X	X	X	X
Linux x86-64	Red Hat Enterprise Linux 6	4+	X	X	X	X	X	
Linux x86-64	Red Hat Enterprise Linux 7		X		X	X	X	X
Microsoft Windows x64 (64-bit)	2012		X	X	X	X	X	
Microsoft Windows x64 (64-bit)	2016					X	X	X
Microsoft Windows x64 (64-bit)	2019							X
Microsoft Windows x64 (64-bit)	2012 R2		X		X	X	X	X
Oracle Solaris on SPARC (64-bit)	11,3	SRU 3.31+	X	X	X	X	X	X
Oracle Solaris on SPARC (64-bit)	11,4	SRU 2+	X	X	X	X	X	X
Oracle Solaris on x86-64 (64-bit)	11,3	SRU 3.31+	X	X	X	X	X	X
Oracle Solaris on x86-64 (64-bit)	11,4	SRU 2+	X	X	X	X	X	X

Oracle GoldenGate 19.1 certification matrix



GoldenGate | Marketplace Image

<https://cloudmarketplace.oracle.com>



[Watch on YouTube](#)






Autonomous JSON Database

The New Kid on the Block

”

Oracle Autonomous JSON Database is a **cloud document database service** that makes it simple to develop JSON-centric applications. It features simple document APIs, serverless scaling, high performance ACID transactions, comprehensive security, and low pay-per-use pricing.

Autonomous JSON Database | Provision



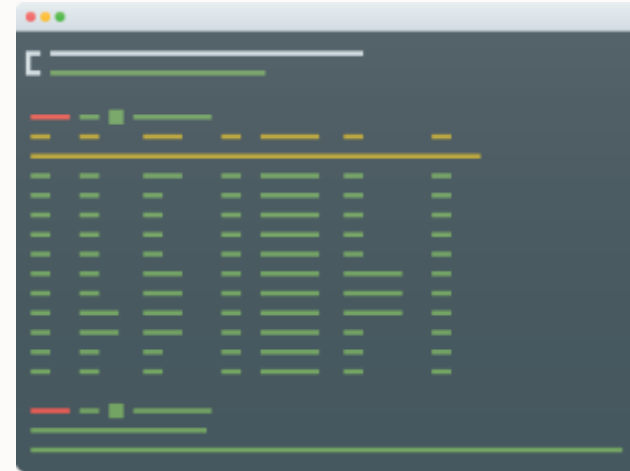
[Watch on YouTube](#)

Options | **SQLCI**




Uses SODA

Natively available in SQLCI



Options | SQLCI



[Watch on YouTube](#)



Options | **REST API**




Fast and easy


Via ORDS

ORDS already configured in ADB

Jeff Smith [blog post](#)



Options | REST API



[Watch on YouTube](#)



Autonomous JSON Database | Further Reading



[Introducing Oracle Autonomous JSON Database for application developers](#)

Julian Dontcheff: [How is Oracle Autonomous JSON Database different from Oracle ATP and MongoDB?](#)





Photo by [Jakob Boman](#) on [Unsplash](#)

Deep Dive

Recommendations | Schema Advisor

[Oracle Autonomous Database Schema Advisor \(Doc ID 2462677.1\)](#)

```
@install_adb_advisor.sql ADB_ADVISOR ADB_ADVISOR  
  
CONNECT ADB_ADVISOR/ADB_ADVISOR  
SET SERVEROUTPUT ON FORMAT WRAPPED  
SET LINES 3000  
EXEC ADB_ADVISOR.REPORT (schemas=>'METAL',adb_type=>'ADW');
```

Pro Tip:
Other migration targets are ATD, ATPD
and ADWD



Recommendations | Cloud Premigration Advisor Tool (CPAT)

Cloud Premigration Advisor Tool (CPAT) Analyzes Databases for Suitability of Cloud Migration (Doc ID 2758371.1)

- Successor to the Schema Advisor

```
=====  
Cloud Premigration Advisor Check Details List  
=====  
Report ~~~~~  
~~~~~ Check Name:   has_user_defined_objects_in_sys  
CPAT Check Result: BLOCKER  
CPAT ~~~~~  
Target Description:   User-defined objects in SYS and SYSTEM schemas will not migrate.  
Migration Failure Impact: Any applications relying on user-defined objects in SYS and SYSTEM will fail.  
Report Action:       User-defined objects were detected in SYS/SYSTEM schemas. Consider moving them  
Report out prior to the migration.  
  
Database Relevant Objects:  
~~~~~ OWNER  OBJECT_NAME                SUBOBJECT_NAME  OBJECT_TYPE  
-----  
Source SYS      MY_VERIFICATION_FUNCTION      FUNCTION  
Source SYSTEM  IDXSIZE_UQ2                   INDEX  
Source SYSTEM  IDXSIZE_UQ3                   INDEX  
SYSTEM ISEQ$$_83287              SEQUENCE
```



Recommendations | Cloud Premigration Advisor Tool (CPAT)

Does not install anything in your database

- Supports read only databases (e.g. physical standby)

JSON or Text output

```
appVersion: "22.1.4-1"
checkList:
  0:
    action: "User-defined objects were detected in SYS/SYSTEM schemas. Consider moving them out prior to the migration."
    description: "User-defined objects in SYS and SYSTEM schemas will not migrate."
    executedSQL: "SELECT owner, object_name, subobject_name, object_type FROM sys.dba_objects WHERE owner IN('SYS','SYSTEM') AND or
object_name NOT LIKE 'SYS%NT%' AND object_name NOT LIKE 'AQ%%$%' AND object_name NOT LIKE 'AQ%SRVNTFN%' AND object_nam
'!' AND object_name NOT LIKE 'AQ!_%' ESCAPE '!' AND object_name NOT LIKE 'SYS$SERVICE!_METRICS%' ESCAPE '!' AND object
'TTS_TB$$', 'TTS_USR$', 'V$LCH_AUDIT', 'XDB_INSTALLATION_TAB', 'PROXY$', 'SQLPLUS_PRODUCT_PROFIL
('UTL_RECOMP_SORT_IDX1', 'DAM_CONFIG_PARAM_UK1', 'SCHEDULER$_EVENT_QTAB_HIST', 'I_PROXY$', 'I_AU
object_name LIKE 'new!_values%!_T' ESCAPE '!' OR object_name LIKE 'old!_value%!_COLL' ESCAPE '!' OR
object_name LIKE 'SQLDATA%' OR object_name LIKE 'SQLDEF%' OR object_name LIKE 'SQLDEPENDENCY%' OR
'SQLERROR%' OR object_name LIKE 'SYSTP%' OR object_name IN ('ALERT_TYPE', 'DBMS_XS_PRIVID_LIST',
'SQLRANGEDEFINITIONSTRUCT', 'SQLVERSIONINFOSTRUCT', 'SQLUNMATCHEDINPUTSEXCEPTION', 'SQLSTRINGS
object_name LIKE 'AQ$AQ!_%' ESCAPE '!' OR object_name LIKE 'X!_$$KS' ESCAPE '!' OR
'V$FALSE_PING', 'V$PING', 'USER_SOURCE_TAB_COLUMNS', 'USER_ADVISOR_DIRECTIVE
LIKE 'OWNER!_MIGRATE!_UPDATE!_%' ESCAPE '!' OR object_name LIKE 'VALIDATE!_%' ESCAPE '!' OR
(object_type = 'TRIGGER' AND object_name IN ('ATP_DBFS_MOUNT', 'ATP_DBFS_UNMOUNT',
object_name LIKE 'ORA$AUTOTASK!_%' ESCAPE '!')) AND NOT (object_type IN ('JAVA RESOURCE', 'JAVA CLASS', 'JAVA SOURCE') A
'CW2!_OLAP!_%' ESCAPE '!' OR object_name IN ('SDO_RDF_EXP_IMP', 'SQLJUTL2', 'BLAST_CUR',

failureImpact:
  name: "has_user_defined_objects_in_sys"
relevantObjectsData:
  0:
    OBJECT_NAME: "MY_VERIFICATION_FUNCTION"
    OBJECT_TYPE: "FUNCTION"
```



Recommendations | Multipart Uploads

Recommended for files larger than 100 MB

Use [OCI CLI](#)

```
oci os object put \  
  --namespace ... -bn ... --file ... --name ... \  
  --part-size 1024 \  
  --parallel-upload-count 4
```

Max part size is 50 GB

OCI CLI [installation guide](#)

Recommendations | Bulk Uploads

Recommended for **many** files

Use [OCI CLI](#)

```
oci os object bulk-upload \  
  -ns ... -bn ... --src-dir ...
```

Does multipart and parallel uploads automatically

Optionally,

- Finetune uploads using `--parallel-upload-count` and `--part-size`
- Prefix all file names with `--object-prefix`
- Include or exclude files selectively using patterns and `--include` and `--exclude`

Recommendations | Checksum

- Avoid in-flight corruption

```
[oracle@hol]$ md5sum metal*.dmp

5edf66ed92086b4f69580fc27b75f662 metal_01.dmp
59eb25ff2a0f648c051a9212e0861979 metal_02.dmp
29951a56abe074d9151c27728d88e9eb metal_03.dmp
c8860e7a71e74f8013068240b598c116 metal_04.dmp
0d05d258e4b501c657cd9490b7e48715 metal_05.dmp
1e367394a31e2ce45d2aeb6a3d4f9507 metal_06.dmp
9c276aa580c0e57c0829f274d04d15de metal_07.dmp
0d560d0ce57c47425424e17604d8ec49 metal_08.dmp
```

```
SQL> SELECT object_name, checksum
        FROM DBMS_CLOUD.LIST_OBJECTS(
            '<credential_name>',
            '<location_uri>');

metal_01.dmp 5edf66ed92086b4f69580fc27b75f662
metal_02.dmp 59eb25ff2a0f648c051a9212e0861979
metal_03.dmp 29951a56abe074d9151c27728d88e9eb
metal_04.dmp c8860e7a71e74f8013068240b598c116
metal_05.dmp 0d05d258e4b501c657cd9490b7e48715
metal_06.dmp 1e367394a31e2ce45d2aeb6a3d4f9507
metal_07.dmp 9c276aa580c0e57c0829f274d04d15de
metal_08.dmp 0d560d0ce57c47425424e17604d8ec49
```

- Windows: `Get-FileHash *.dmp -Algorithm MD5`
- Manifests as ORA-31693 ORA-29913 ORA-29104



Recommendations | ADB Resources

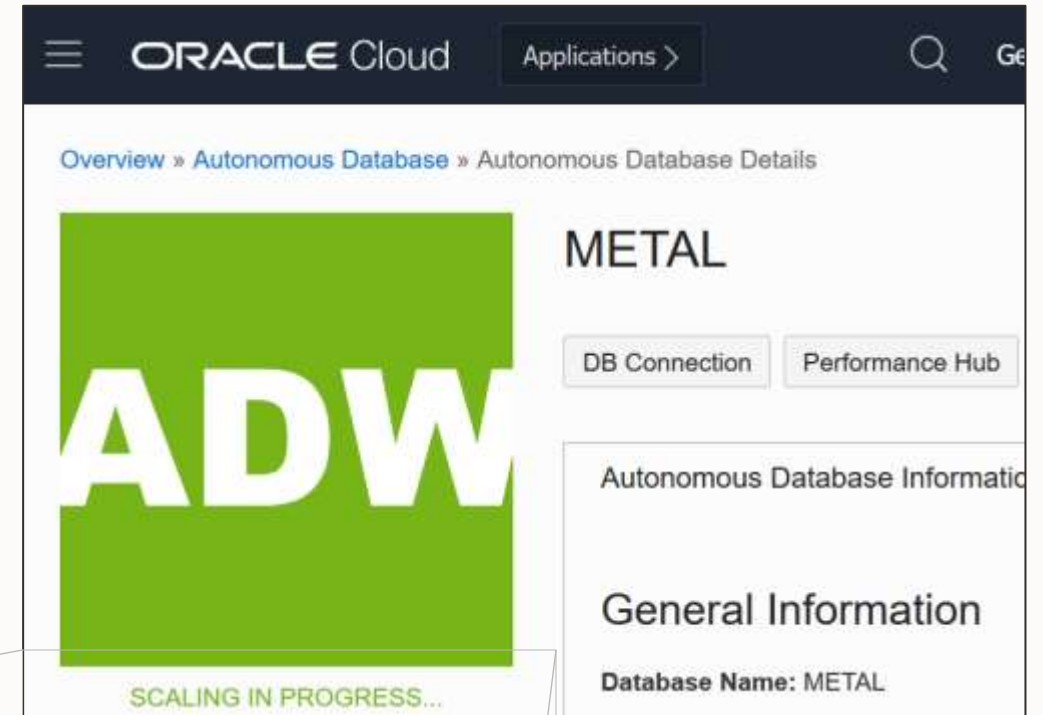
When loading database, **scale up** on CPUs

I/O scales linear with CPU

Connect to **high** service

Loading from the object store

typically faster than database link



Recommendations | ADB Resources

Allocate adequate storage **before** import

```
ALTER DATABASE DATAFILE 2556 RESIZE 100G;
```

Tablespaces are **bigfile** - 32 TB limit

For ADB dedicated:

- use `allow_rowid_column_type=true` to allow tables with rowid columns

Recommendations | Data Pump



Don't use `SYS AS SYSDBA`

Applies to export and import



Caution: Do not start Export as `SYSDBA`, except at the request of Oracle technical support. `SYSDBA` is used internally and has specialized functions; its behavior is not the same as for general users.



Recommendations | Data Pump



Always use a **parameter file**

Applies to export and import

```
$ more export.par

DIRECTORY=my_data_pump_dir
DUMPFILE=dumpfile%U.dmp
LOGFILE=logfile.log
SCHEMAS=HR
EXCLUDE=STATISTICS
LOGTIME=ALL
METRICS=YES
FLASHBACK_TIME=SYSTIMESTAMP
PARALLEL=4
FILESIZE=5G
TRANSFORM=OMIT_ENCRYPTION_CLAUSE:Y

$ expdp parfile=export.par
```



Recommendations | Data Pump



Always export to **multiple files**

DUMPFILe applies to export and import
FILESIZE applies only to export
More than 99 files, use DUMPFILe=dumpfile%L.dmp

```
$ more export.par

DIRECTORY=my_data_pump_dir
DUMPFILe=dumpfile%U.dmp
LOGFILE=logfile.log
SCHEMAS=HR
EXCLUDE=STATISTICS
LOGTIME=ALL
METRICS=YES
FLASHBACK_TIME=SYSTIMESTAMP
PARALLEL=4
FILESIZE=5G
TRANSFORM=OMIT_ENCRYPTION_CLAUSE:Y

$ expdp parfile=export.par
```



Recommendations | Data Pump



Always use **schema mode**

Applies to export

```
$ more export.par

DIRECTORY=my_data_pump_dir
DUMPFILE=dumpfile%U.dmp
LOGFILE=logfile.log
SCHEMAS=HR
EXCLUDE=STATISTICS
LOGTIME=ALL
METRICS=YES
FLASHBACK_TIME=SYSTIMESTAMP
PARALLEL=4
FILESIZE=5G
TRANSFORM=OMIT_ENCRYPTION_CLAUSE:Y

$ expdp parfile=export.par
```



Recommendations | Data Pump



Always exclude **statistics**

Applies to export and import

```
$ more export.par

DIRECTORY=my_data_pump_dir
DUMPFILE=dumpfile%U.dmp
LOGFILE=logfile.log
SCHEMAS=HR
EXCLUDE=STATISTICS
LOGTIME=ALL
METRICS=YES
FLASHBACK_TIME=SYSTIMESTAMP
PARALLEL=4
FILESIZE=5G
TRANSFORM=OMIT_ENCRYPTION_CLAUSE:Y

$ expdp parfile=export.par
```



Recommendations | Data Pump



Always include **diagnostics**

Applies to export and import
LOGTIME available from 12.1, METRICS from 11.2.0.4

```
$ more export.par

DIRECTORY=my_data_pump_dir
DUMPFILE=dumpfile%U.dmp
LOGFILE=logfile.log
SCHEMAS=HR
EXCLUDE=STATISTICS
LOGTIME=ALL
METRICS=YES
FLASHBACK_TIME=SYSTIMESTAMP
PARALLEL=4
FILESIZE=5G
TRANSFORM=OMIT_ENCRYPTION_CLAUSE:Y

$ expdp parfile=export.par
```



Recommendations | Data Pump - Diagnostics

No diagnostics

```
Processing object type SCHEMA_EXPORT/TABLE/TABLE
Processing object type SCHEMA_EXPORT/TABLE/TABLE_DATA
. . imported "METAL"."ALBUMS"          988.8 KB    28069 rows
. . imported "METAL"."BANDS"          3.444 MB    37723 rows
. . imported "METAL"."REVIEWS"        66.47 MB    21510 rows
```

All diagnostics

```
16-OCT-20 17:26:57.158: Processing object type SCHEMA_EXPORT/TABLE/TABLE
16-OCT-20 17:26:58.262: Startup took 1 seconds
16-OCT-20 17:26:58.264: Startup took 1 seconds
16-OCT-20 17:26:59.082:      Completed 3 TABLE objects in 1 seconds
16-OCT-20 17:26:59.082:      Completed by worker 1 1 TABLE objects in 1 seconds
16-OCT-20 17:26:59.082:      Completed by worker 2 1 TABLE objects in 0 seconds
16-OCT-20 17:26:59.082:      Completed by worker 3 1 TABLE objects in 0 seconds
16-OCT-20 17:26:59.313: Processing object type SCHEMA_EXPORT/TABLE/TABLE_DATA
16-OCT-20 17:27:01.943: . . imported "METAL"."ALBUMS"      988.8 KB    28069 rows in 2 seconds using external_table
16-OCT-20 17:27:03.778: . . imported "METAL"."BANDS"      3.444 MB    37723 rows in 2 seconds using external_table
16-OCT-20 17:27:12.644: . . imported "METAL"."REVIEWS"    66.47 MB    21510 rows in 13 seconds using external_table
```

Recommendations | Data Pump



Always make **consistent** exports

Applies to export only

```
$ more export.par

DIRECTORY=my_data_pump_dir
DUMPFILE=dumpfile%U.dmp
LOGFILE=logfile.log
SCHEMAS=HR
EXCLUDE=STATISTICS
LOGTIME=ALL
METRIC=YES
FLASHBACK_TIME=SYSTIMESTAMP
PARALLEL=4
FILESIZE=5G
TRANSFORM=OMIT_ENCRYPTION_CLAUSE:Y

$ expdp parfile=export.par
```



Recommendations | Data Pump



Always use **parallel**

Enterprise Edition only

Applies to export and import

OCI: Number of OCPUs
On-premises: Number physical cores x 2

```
$ more export.par

DIRECTORY=my_data_pump_dir
DUMPFILE=dumpfile%U.dmp
LOGFILE=logfile.log
SCHEMAS=HR
EXCLUDE=STATISTICS
LOGTIME=ALL
METRIC=YES
FLASHBACK_TIME=SYSTIMESTAMP
PARALLEL=4
FILESIZE=5G
TRANSFORM=OMIT_ENCRYPTION_CLAUSE:Y

$ expdp parfile=export.par
```



Recommendations | Data Pump



Always remove **column encryption**

Applies to import only

```
$ more export.par

DIRECTORY=my_data_pump_dir
DUMPFILE=dumpfile%U.dmp
LOGFILE=logfile.log
SCHEMAS=HR
EXCLUDE=STATISTICS
LOGTIME=ALL
METRIC=YES
FLASHBACK_TIME=SYSTIMESTAMP
PARALLEL=4
FILESIZE=5G
TRANSFORM=OMIT_ENCRYPTION_CLAUSE:Y

$ expdp parfile=export.par
```



Recommendations | Data Pump



Consider using **compression**

Advanced Compression Option license required

Applies to export only
Algorithms: BASIC, LOW, MEDIUM, HIGH

```
$ more export.par

DIRECTORY=my_data_pump_dir
DUMPFILE=dumpfile%U.dmp
LOGFILE=logfile.log
SCHEMAS=HR
EXCLUDE=STATISTICS
LOGTIME=ALL
METRIC=YES
FLASHBACK_TIME=SYSTIMESTAMP
PARALLEL=4
FILESIZE=5G
COMPRESSION=ALL
COMPRESSION_ALGORITHM=MEDIUM

$ expdp parfile=export.par
```



Recommendations | Data Pump - Compression Comparison



12.2 EBS Database - export

Example 1

	File Size MB	Ratio	Time
NONE	5500	1	4m 54s
ALL BASIC	622	8,9	4m 58s
ALL LOW	702	7,8	5m 24s
ALL MEDIUM	567	9,7	4m 55s
ALL HIGH	417	13,2	5m 13s

Example 2

	File Size MB	Ratio	Time
NONE	5800	1	2m 33s
ALL BASIC	705	8,2	3m 3s
ALL LOW	870	6,6	8m 11s
ALL MEDIUM	701	8,2	3m 1s
ALL HIGH	509	11,3	12m 16s



Tweaks | Data Pump



Use VIEWS_AS_TABLES to export
subset of data

```
SQL> CREATE VIEW v1 AS  
      SELECT *  
      FROM t1  
      WHERE origin='INTERNAL'  
  
$ expdp tables=v1
```

Applies to export only

Tweaks | Data Pump



Use **wildcards** for URL based file names

```
$ impdp dumpfile=http://...../exp01.dmp,  
          http://...../exp02.dmp,  
          http://...../expnn.dmp  
  
$ impdp dumpfile=http://...../exp%u.dmp
```

Applies to import only



ADB Compliance | Data Pump



Transform IOTs to tables

```
TRANSFORM=DWCS_CVT_IOTS:Y  
TRANSFORM=CONSTRAINT_USE_DEFAULT_INDEX:Y  
TRANSFORM=SEGMENT_ATTRIBUTES:N  
REMAP_TABLESPACE=%:DATA  
EXCLUDE=INDEX,CLUSTER,INDEXTYPE,  
        MATERIALIZED_VIEW,  
        MATERIALIZED_VIEW_LOG,  
        MATERIALIZED_ZONEMAP,DB_LINK  
DATA_OPTIONS=GROUP_PARTITION_TABLE_DATA  
PARTITION_OPTIONS=MERGE
```

Applies to import only



ADB Compliance | Data Pump



Enforce proper **naming standard** by using constraint name for PK and FK indexes

```
TRANSFORM=DWCS_CVT_IOTS:Y
TRANSFORM=CONSTRAINT_USE_DEFAULT_INDEX:Y
TRANSFORM=SEGMENT_ATTRIBUTES:N
REMAP_TABLESPACE=%:DATA
EXCLUDE=INDEX,CLUSTER,INDEXTYPE,
        MATERIALIZED_VIEW,
        MATERIALIZED_VIEW_LOG,
        MATERIALIZED_ZONEMAP,DB_LINK
DATA_OPTIONS=GROUP_PARTITION_TABLE_DATA
PARTITION_OPTIONS=MERGE
```

Applies to import only



ADB Compliance | Data Pump



Remove **segment customization**

```
TRANSFORM=DWCS_CVT_IOTS:Y
TRANSFORM=CONSTRAINT_USE_DEFAULT_INDEX:Y
TRANSFORM=SEGMENT_ATTRIBUTES:N
REMAP_TABLESPACE=%:DATA
EXCLUDE=INDEX,CLUSTER,INDEXTYPE,
        MATERIALIZED_VIEW,
        MATERIALIZED_VIEW_LOG,
        MATERIALIZED_ZONEMAP,DB_LINK
DATA_OPTIONS=GROUP_PARTITION_TABLE_DATA
PARTITION_OPTIONS=MERGE
```

Applies to import only



ADB Compliance | Data Pump



Remap any tablespace to
DATA tablespace

```
TRANSFORM=DWCS_CVT_IOTS:Y
TRANSFORM=CONSTRAINT_USE_DEFAULT_INDEX:Y
TRANSFORM=SEGMENT_ATTRIBUTES:N
REMAP_TABLESPACE=%:DATA
EXCLUDE=INDEX,CLUSTER,INDEXTYPE,
        MATERIALIZED_VIEW,
        MATERIALIZED_VIEW_LOG,
        MATERIALIZED_ZONEMAP,DB_LINK
DATA_OPTIONS=GROUP_PARTITION_TABLE_DATA
PARTITION_OPTIONS=MERGE
```

Applies to import only



ADB Compliance | Data Pump

For expdp:

```
EXCLUDE=INDEX,CLUSTER,INDEXTYPE,  
MATERIALIZED_VIEW,  
MATERIALIZED_VIEW_LOG,  
MATERIALIZED_ZONEMAP,DB_LINK  
DATA_OPTIONS=GROUP_PARTITION_TABLE_DATA
```

Follow [ADW best practice](#) and exclude these object types

For impdp:

```
TRANSFORM=DWCS_CVT_IOTS:Y  
TRANSFORM=CONSTRAINT_USE_DEFAULT_INDEX:Y  
TRANSFORM=SEGMENT_ATTRIBUTES:N  
REMAP_TABLESPACE=:DATA  
EXCLUDE=INDEX,CLUSTER,INDEXTYPE,  
MATERIALIZED_VIEW,  
MATERIALIZED_VIEW_LOG,  
MATERIALIZED_ZONEMAP,DB_LINK  
PARTITION_OPTIONS=MERGE
```

ADB Compliance | Data Pump



Follow ATP best practice and exclude these object types

For expdp:

EXCLUDE=CLUSTER,DB_LINK

For impdp:

TRANSFORM=DWCS_CVT_IOTS:Y

TRANSFORM=CONSTRAINT_USE_DEFAULT_INDEX:Y

TRANSFORM=SEGMENT_ATTRIBUTES:N

REMAP_TABLESPACE=%:DATA

EXCLUDE=CLUSTER,DB_LINK



ADB Compliance | Data Pump



Only **SecureFiles LOBs** are allowed

```
SQL> select value from v$parameter where name = 'db_securefile';  
  
ALWAYS
```

No need to transform

```
TRANSFORM=LOB_STORAGE:SECUREFILE
```

Applies to import



Data Pump | Troubleshooting

```
ORA-39001: invalid argument value  
ORA-39000: bad dump file specification  
ORA-31640: unable to open dump file for read  
ORA-27037: unable to obtain file status  
Linux-x86_64 Error: 2: No such file or directory
```

Validate Object Storage URI and Credentials
[MOS Doc ID 2468298.1](#)

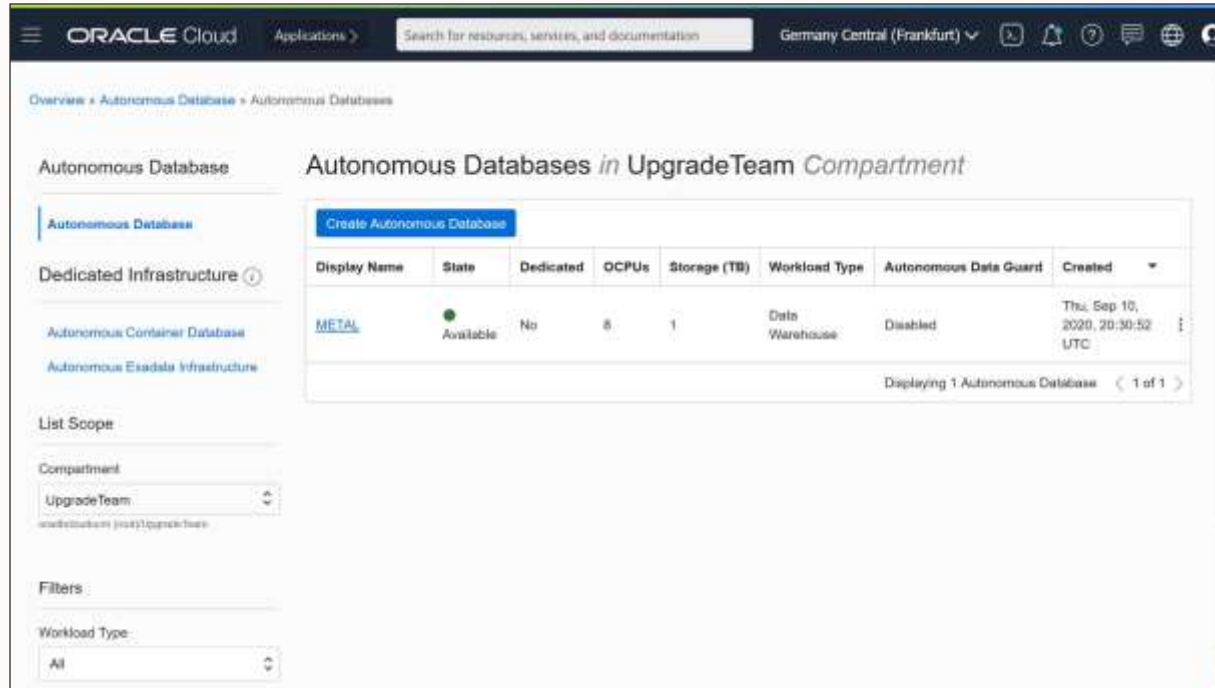




Photo by [Dave Hoefler](#) on [Unsplash](#)

Almost there...

Autonomous Database | SQLcl



[Watch on YouTube](#)

[Kris Rice: SQLcl and OCI Cloud Shell](#)

Pro Tip: It is a full client, so you can use `impdp` and `expdp` as well



Autonomous Database | Quickstart Workshop

Autonomous Database Quickstart Workshop

Provisioning Autonomous Database (ADW and ATP)

Introduction

Objectives

Prerequisites

Video Preview

STEP 1: Choosing ADW or ATP from the Services Menu

STEP 2: Creating the

Introduction

This lab walks you through the steps to get started using the Cloud (Autonomous Data Warehouse [ADW] and Autonomous Transaction Processing) Oracle Cloud. You will provision a new ADW instance and connect to it from Oracle SQL Developer Web.

Note: While this lab uses ADW, the steps are identical for creating an Autonomous Transaction Processing database.

Estimated time: 5 minutes

Objectives

Contents

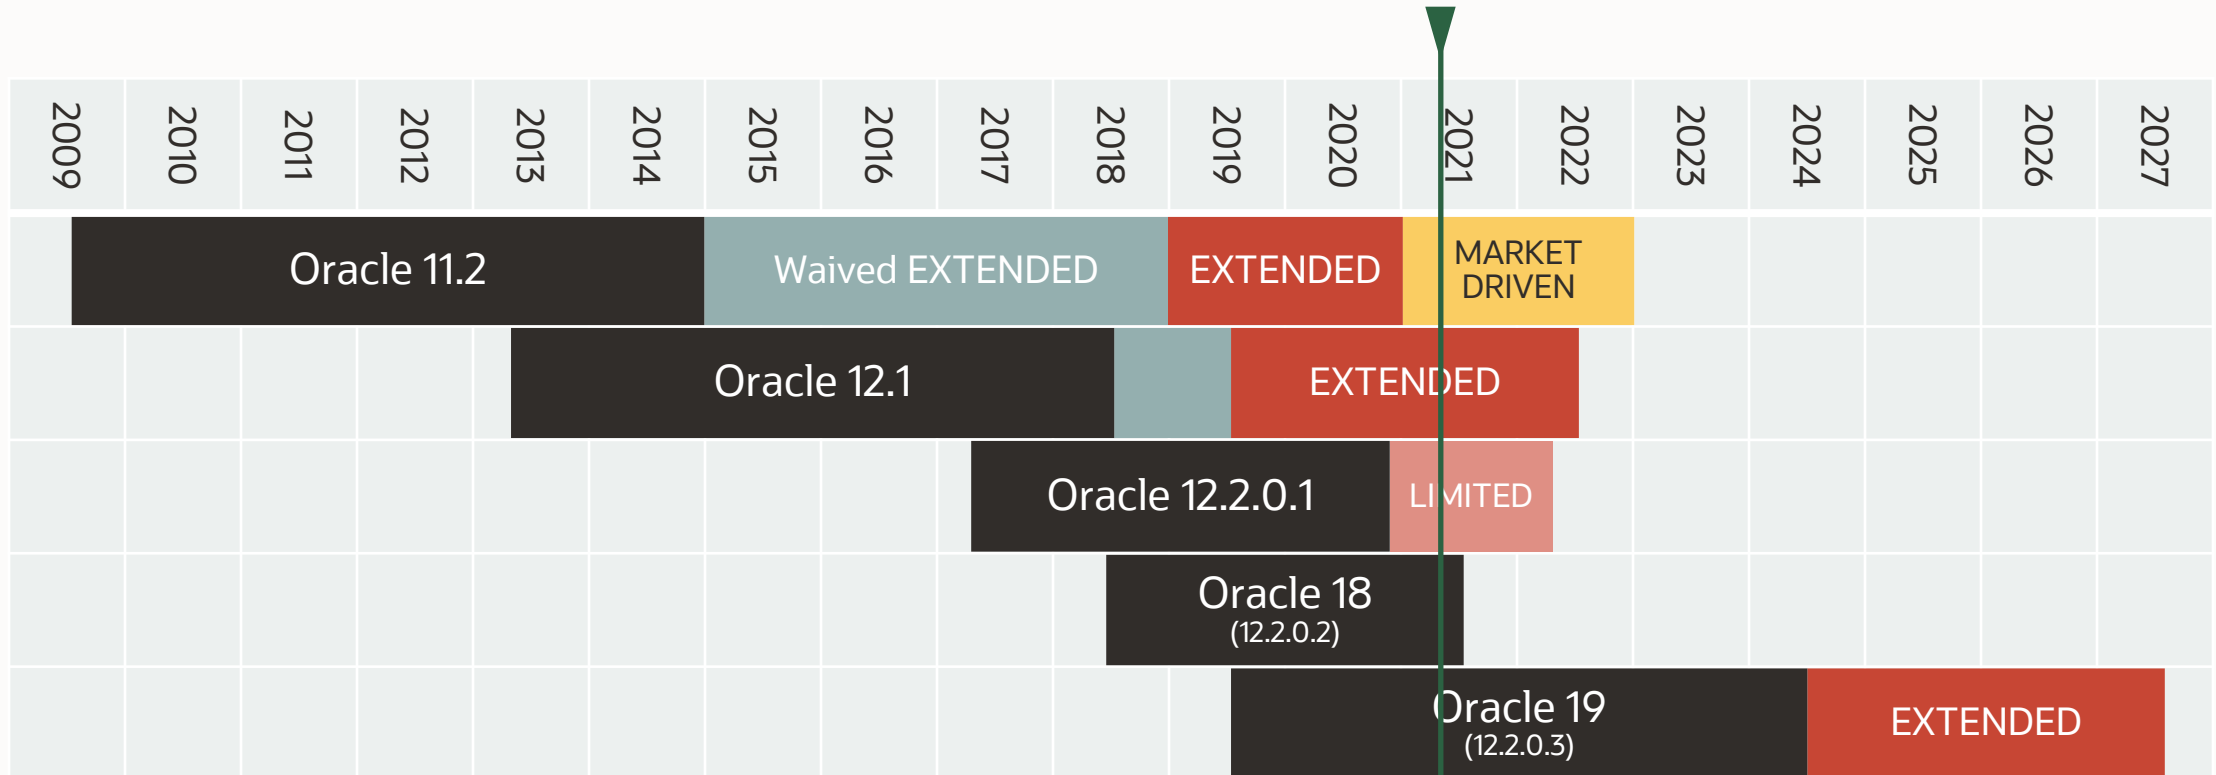
- Prerequisites
- Introduction
- Lab 1: Provision Autonomous Database
- Lab 2: Explore SQL Developer Web
- Lab 3: Loading Data
- Lab 4: Querying External Data
- Lab 5: Visualizing Your Data

[Autonomous Database Quickstart Workshop](#)



Database Cloud Service

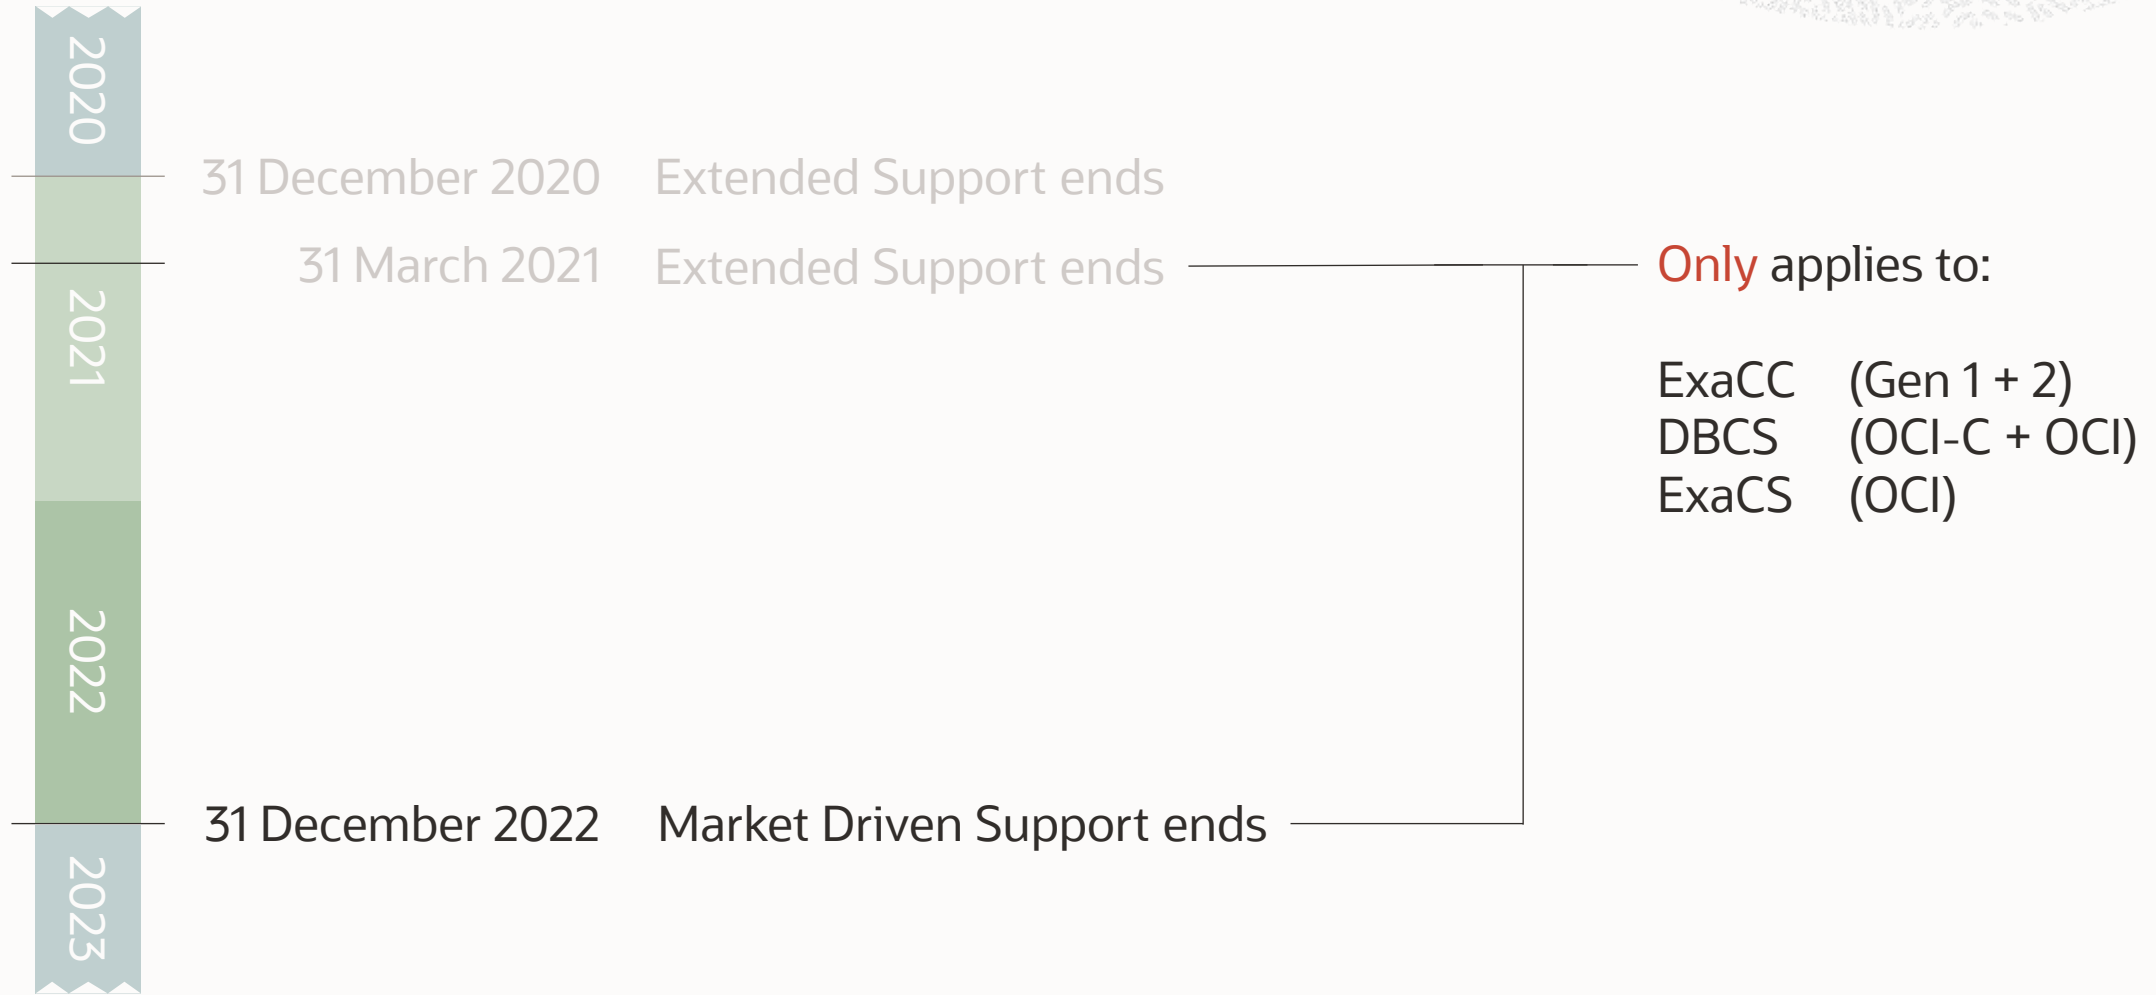
Support Timeline | Cloud 11.2.0.4



Premier Support
 Waived Extended Support
 Paid Extended Support
 Market Driven Support
 Limited Error Correction



Support Timeline | Cloud 11.2.0.4



Database Cloud Service | Virtual Machines



Entry-level, provision with GI or LVM (fast-provision)

Restrictions:

- Only **one CDB** - the pre-created one
- Can't upgrade OS or GI
- Can't install another Oracle Home
- COMPATIBLE is always default - unless
- Drop and create database **not** supported (drop and restore is)



Database Cloud Service | Bare Metal



Mid-level, provision with GI

Restrictions:

- As **many CDBs** as you want
- Can't upgrade OS or GI
- DATA disk group, max. 16 TB
- Only one database edition
- Only one database pr. Oracle Home

Database Cloud Service | Exadata



World's best database machine, provision with GI

Restrictions:

- As **many CDBs** as you want

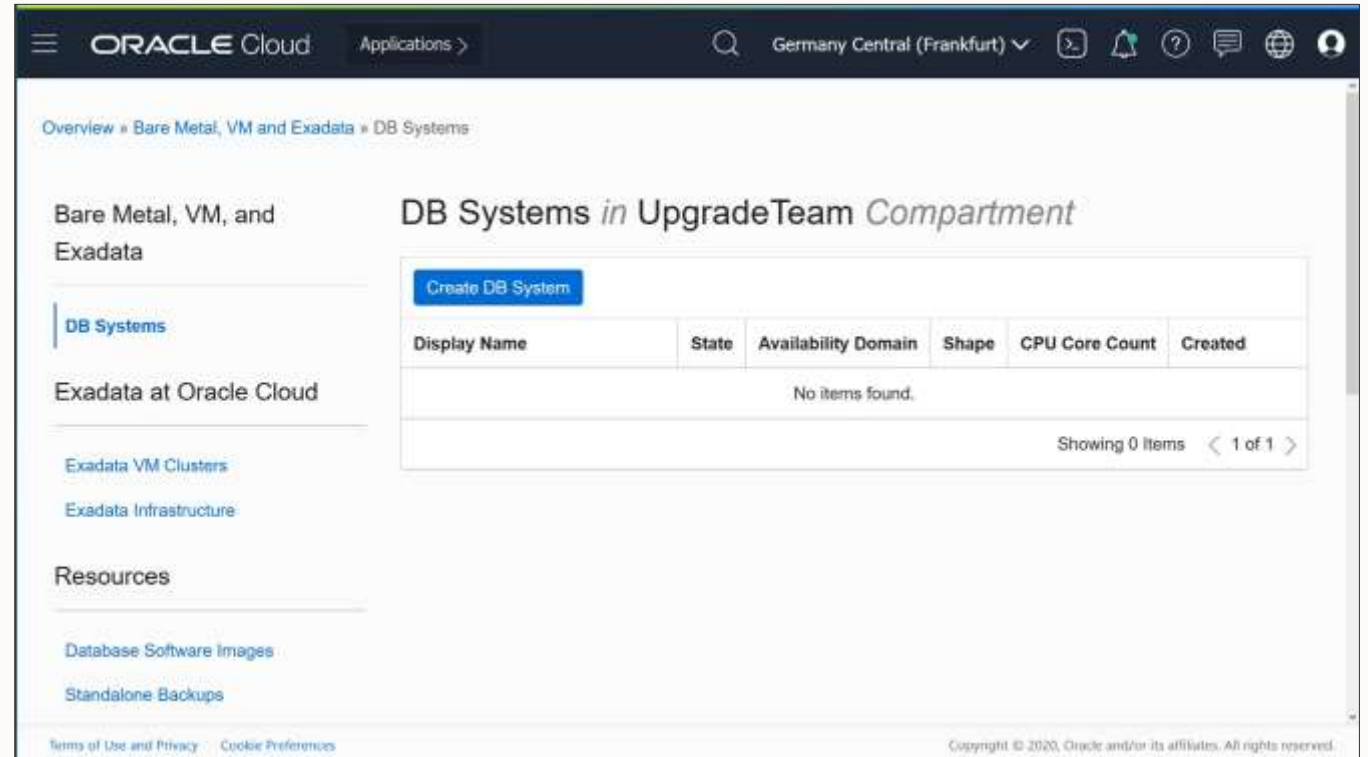


Database Cloud Service | Database Images

For VM and BM and Exadata

Customize your database software:

- Oracle Home
- Release Update
- One-offs




[Watch on YouTube](#)



Migration | **Scaling**

Typically, during migration, you need:



CPU




I/O
throughput



Network
throughput

Migration | **Scaling Virtual Machines**



Change **shape** up and down
Shape changes **offline**
X7 offers from 2 to 24 OCPUs



Scales online, but **up only** with amount of total storage
Allocated in quota between DATA and RECO
Storage is network attached - needs network bandwidth to read/write



Scales with number of OCPUs



Migration | **Scaling Virtual Machines**



[Blog post](#)

Network throughput


Shape	Throughput MB/s
VM.Standard.2.1	128
VM.Standard.2.2	256
VM.Standard.2.4	512
VM.Standard.2.8	1024
VM.Standard.2.16	2048
VM.Standard.2.24	3200

I/O throughput


Storage (GB)	Throughput MB/s
256	120
1024	480
2048	960
4096	1280
10240	1600
20480	3200



Migration | **Scaling Virtual Machines**



Migration | **Scaling Bare Metal**



Scales **up and down**
Scales **online**




Locally attached NVMe disks



25 Gbps network interface
Theoretically 3200 MB/s

Migration | **Scaling Exadata**



Scales **up and down**
Scales **online**



Exadata storage system



2 x 25 Gbps network interface
Theoretically 6400 MB/s

Migration | Transfer Speed

Use multipart and bulk uploads

Recommendations | Multipart Uploads

Recommended for files larger than 100 MB

Use [OCI CLI](#)

```
oci os object put \  
--namespace ... -bn ... --file ... --name ... \  
--part-size 1024 \  
--parallel-upload-count 4
```

Max part size is 50 GB

[OCI CLI installation guide](#)

© Copyright © 2021, Oracle and/or its affiliates



Recommendations | Bulk Uploads

Recommended for *many* files

Use [OCI CLI](#)

```
oci os object bulk-upload \  
-ns ... -bn ... --src-dir ...
```

Does multipart and parallel uploads automatically

[OCI CLI installation guide](#)

© Copyright © 2021, Oracle and/or its affiliates



Migration | Multitenant

Oracle Database 19c:	Max PDBs	3
	DBCS EE-HP or EE-EP or ExaCS	4096

Keep **within limits** - MAX_PDBS:

```
SQL> ALTER SYSTEM SET MAX_PDBS=3 SCOPE=BOTH;
```

DBCS tooling don't care about PDBs - create and drop as you want

Migration | TDE

TDE Tablespace Encryption is a **must-have** in the cloud

Databases in OCI are allowed to use Isolated Keystore mode

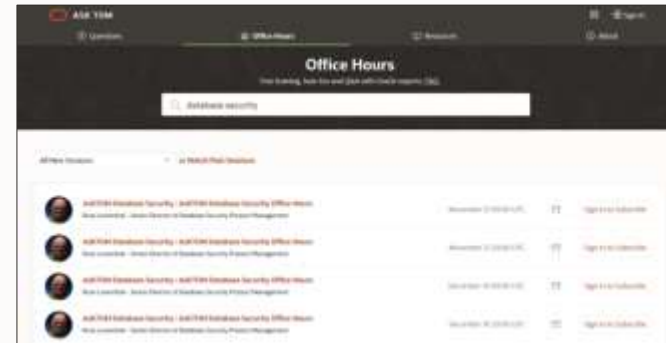
- But it is currently not supported by tooling

Learn encryption:

[Basic introduction](#)



[AskTOM Office Hours](#)



MIGRATION

methods


ZDM

DATA
GUARD


TRANS-
PORTABLE

DATA PUMP
GOLDENGATE

Introduction | **Birds-Eye**



Introduction | **Key Features**



- Simple
- Reliable
- Scalable
- Secure
- Proven
- Free

Prerequisites And Features | Location

Location

Release

Platform

Edition

Architecture

Encryption

Source database can be located

- On-premises
 - Commodity hardware
 - Exadata + ExaCC + ODA
- Oracle Cloud Infrastructure Classic (OCI-C)
- Oracle Cloud Infrastructure (OCI)
 - Cross-region / cross-location migration
 - System migration
- Other clouds

Prerequisites And Features | Release

Location

Release

Platform

Edition

Architecture

Encryption

Source databases

- 11.2.0.4
- And **anything newer**

Target databases

- Physical migrations: Database release **must** be the same
- Logical migrations: Database release can be the **same or higher**

Pro Tip: It is possible to migrate to a higher patch level. If needed, ZDM invokes `datapatch`



Prerequisites And Features | Platform

Location

Release

Platform

Edition

Architecture

Encryption

Supported source platform

- Linux

Supported target platforms

- Autonomous (dedicated and shared)
 - ATP
 - ADW
- DB Systems (VM, BM and Exadata)
- Exadata Cloud at Customer
- Exadata (on-prem)



Prerequisites And Features | Edition

Location

Release

Platform

Edition

Architecture

Encryption

Supported editions

- Enterprise Edition
- Standard Edition

Standard Edition restriction:

For zero downtime approach select **logical migration**

Migration between editions

- Restrictions apply

Pro Tip: True *Zero Downtime* requires an MAA compliant application



Prerequisites And Features | Architecture

Location

Release

Platform

Edition

Architecture

Encryption

Any architecture supported (non-CDB and CDB)

	PHYSICAL	LOGICAL
<u>NON-CDB</u>		
- Migrate directly into a PDB		X
- Keep it as a non-CDB	X	(X)
- Convert to PDB (downtime)	X	
<u>CDB</u>		
- Migrate entire CDB	X	
- Migrate one PDB only		X



Prerequisites And Features | Architecture

Location

Release

Platform


Edition

Architecture

Encryption

Any type (single instance, RAC One Node, RAC)

- Single instance **can** be migrated to RAC
- RAC One Node will be migrated to RAC (physical only)
- RAC will be migrated to RAC (physical only)



Prerequisites And Features | Architecture

Location

Release

Platform


Edition

Architecture

Encryption

Any type (single instance, RAC One Node, RAC)

- Single instance can be migrated to RAC
- RAC One Node **will** be migrated to RAC (physical only)
- RAC **will** be migrated to RAC (physical only)



Prerequisites And Features | Encryption

Location

Release

Platform

Edition

Architecture

Encryption

Source database

- Unencrypted
- Encrypted

Target database

- Is **always** encrypted
- Unencrypted databases gets encrypted **on-the-fly**

Network connectivity

- **Always** and **transparently** encrypted

Migration | Options



PHYSICAL


Online: Data Guard + switchover

Offline: Backup + restore

Standard Edition - offline only

Physical Migration | Overview

Users are connected to source database




Physical Migration | Overview

Provision target database in OCI



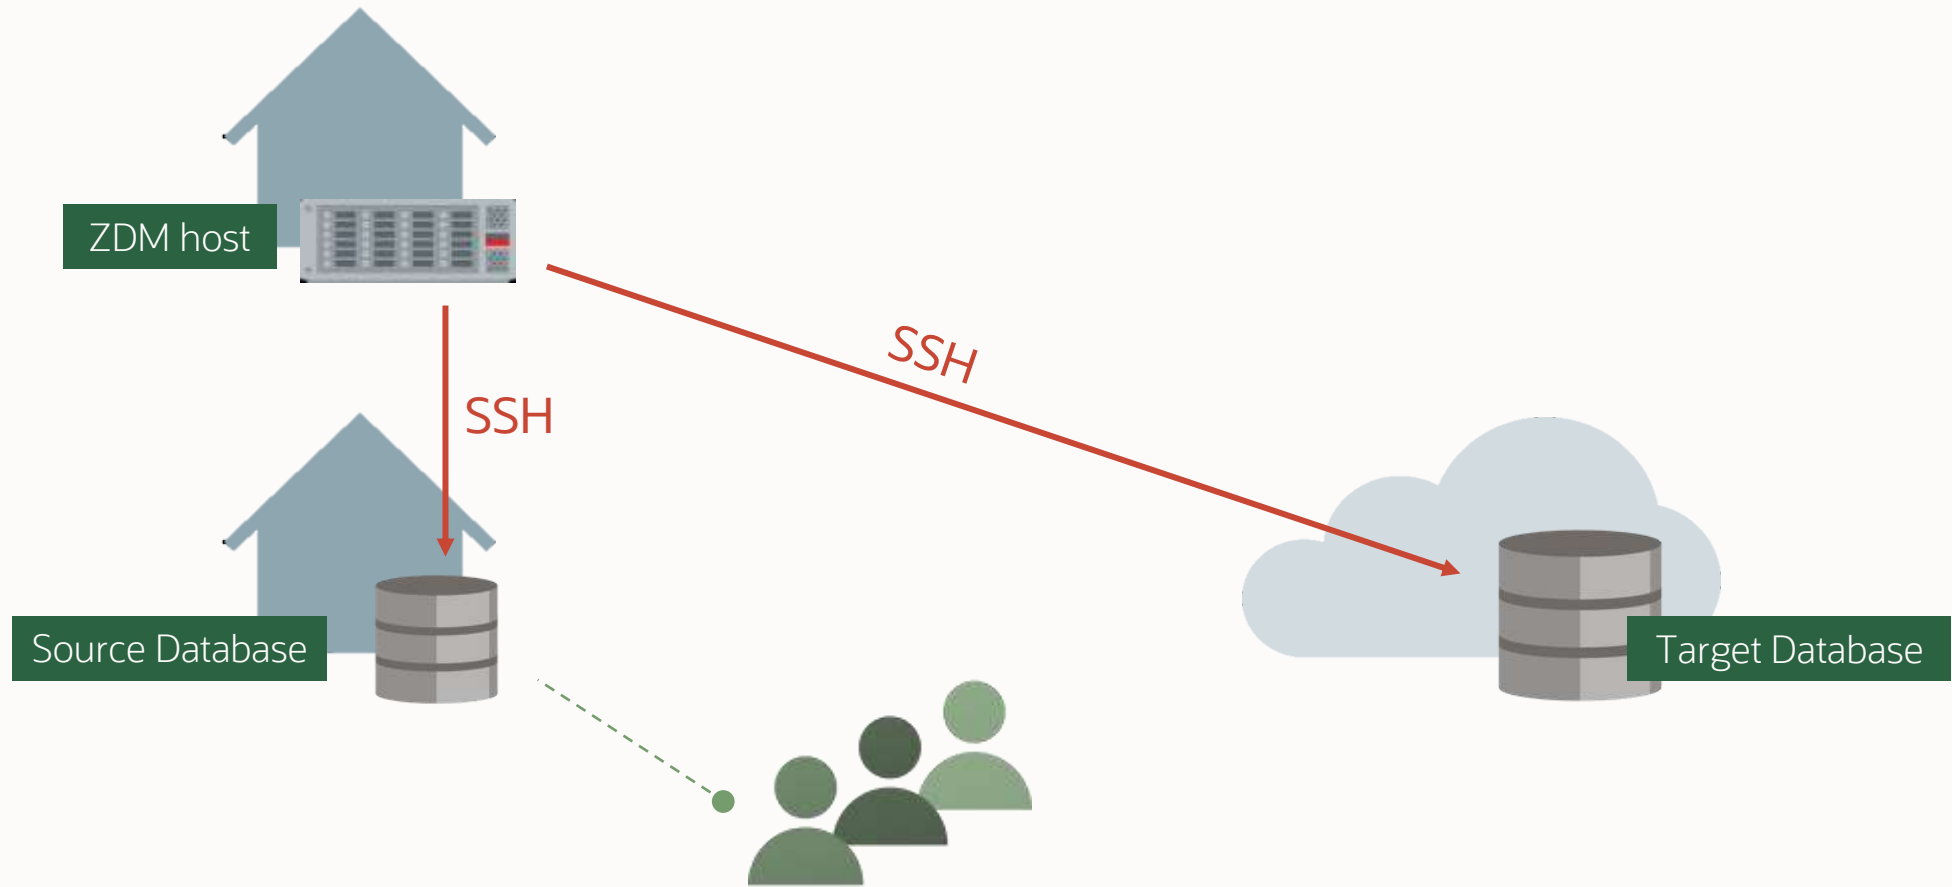
Physical Migration | Overview

Download and install ZDM



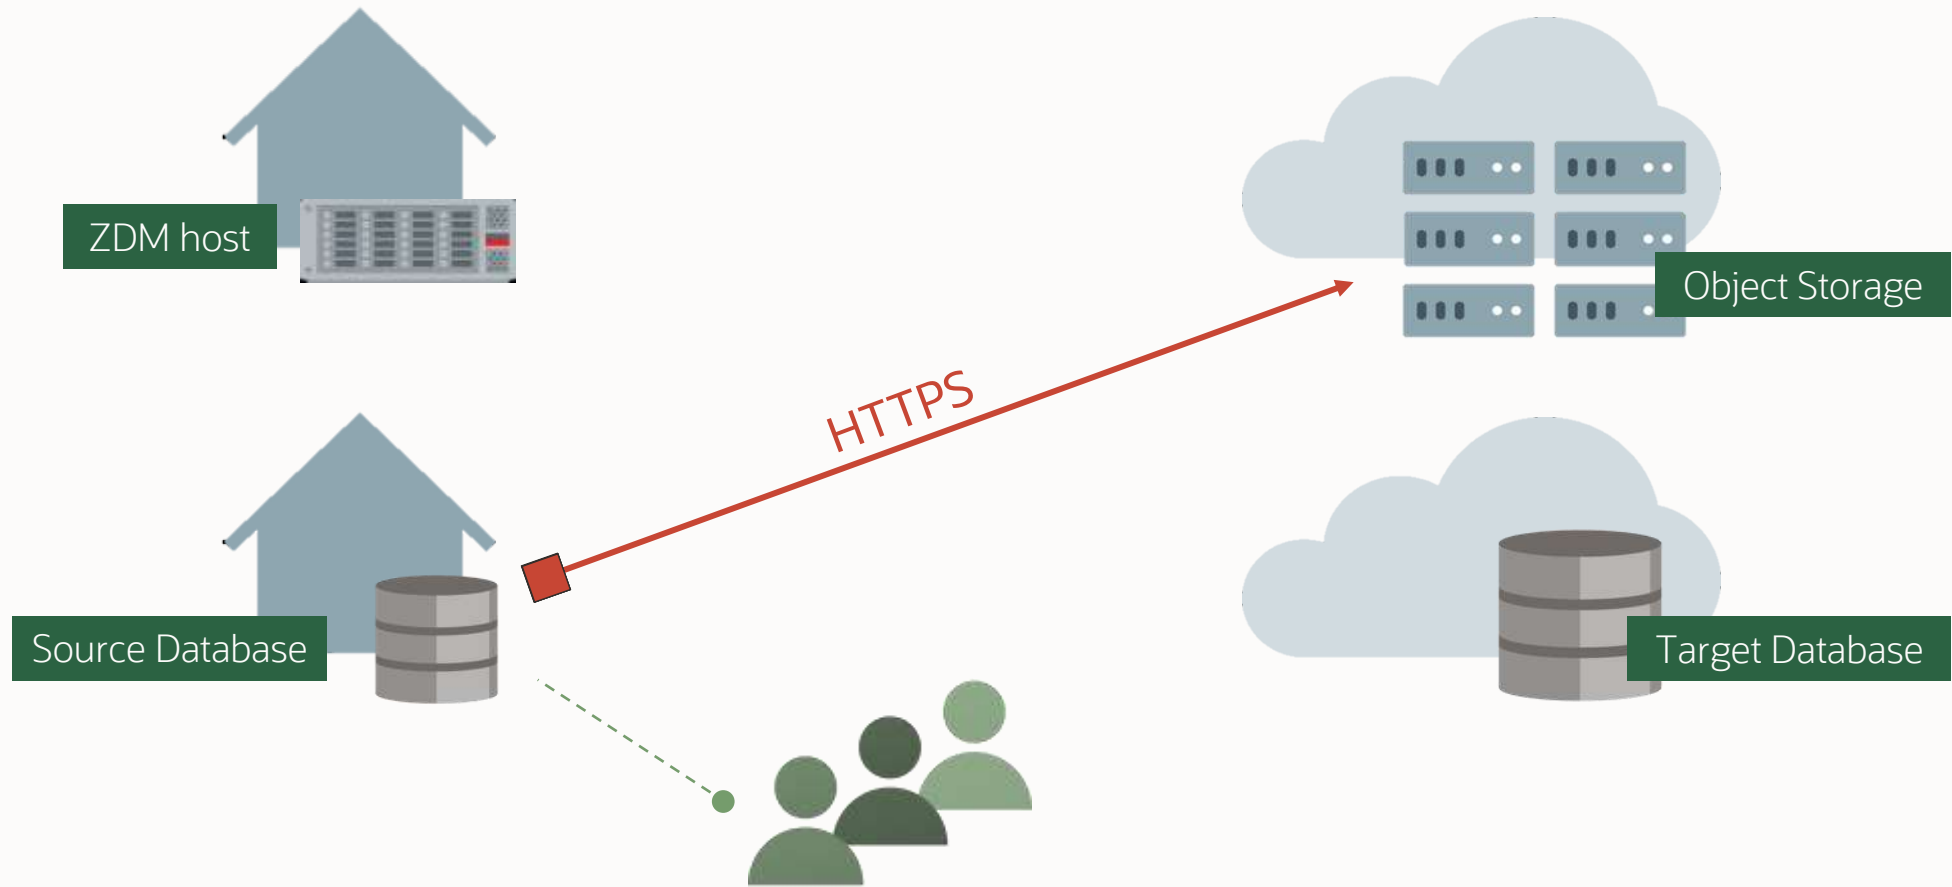
Physical Migration | Overview

ZDM connects to source and target database



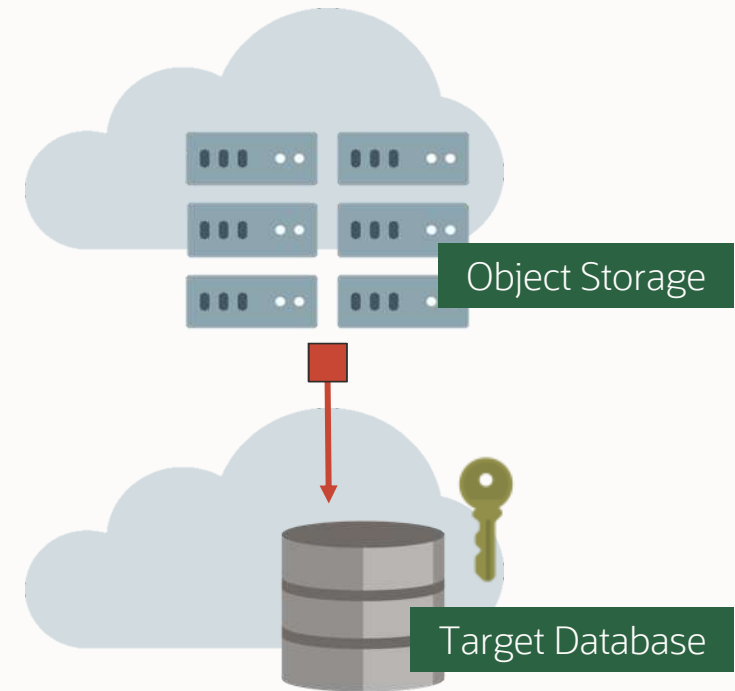
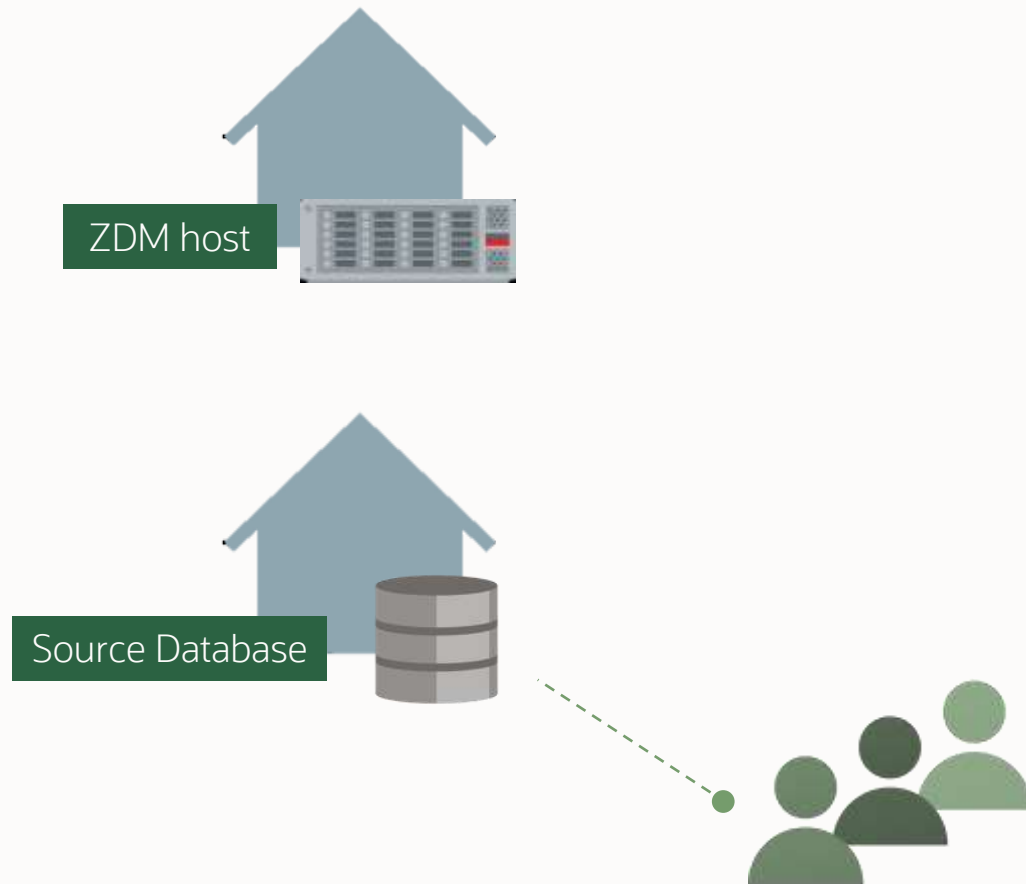
Physical Migration | Overview

Back up source database to object storage



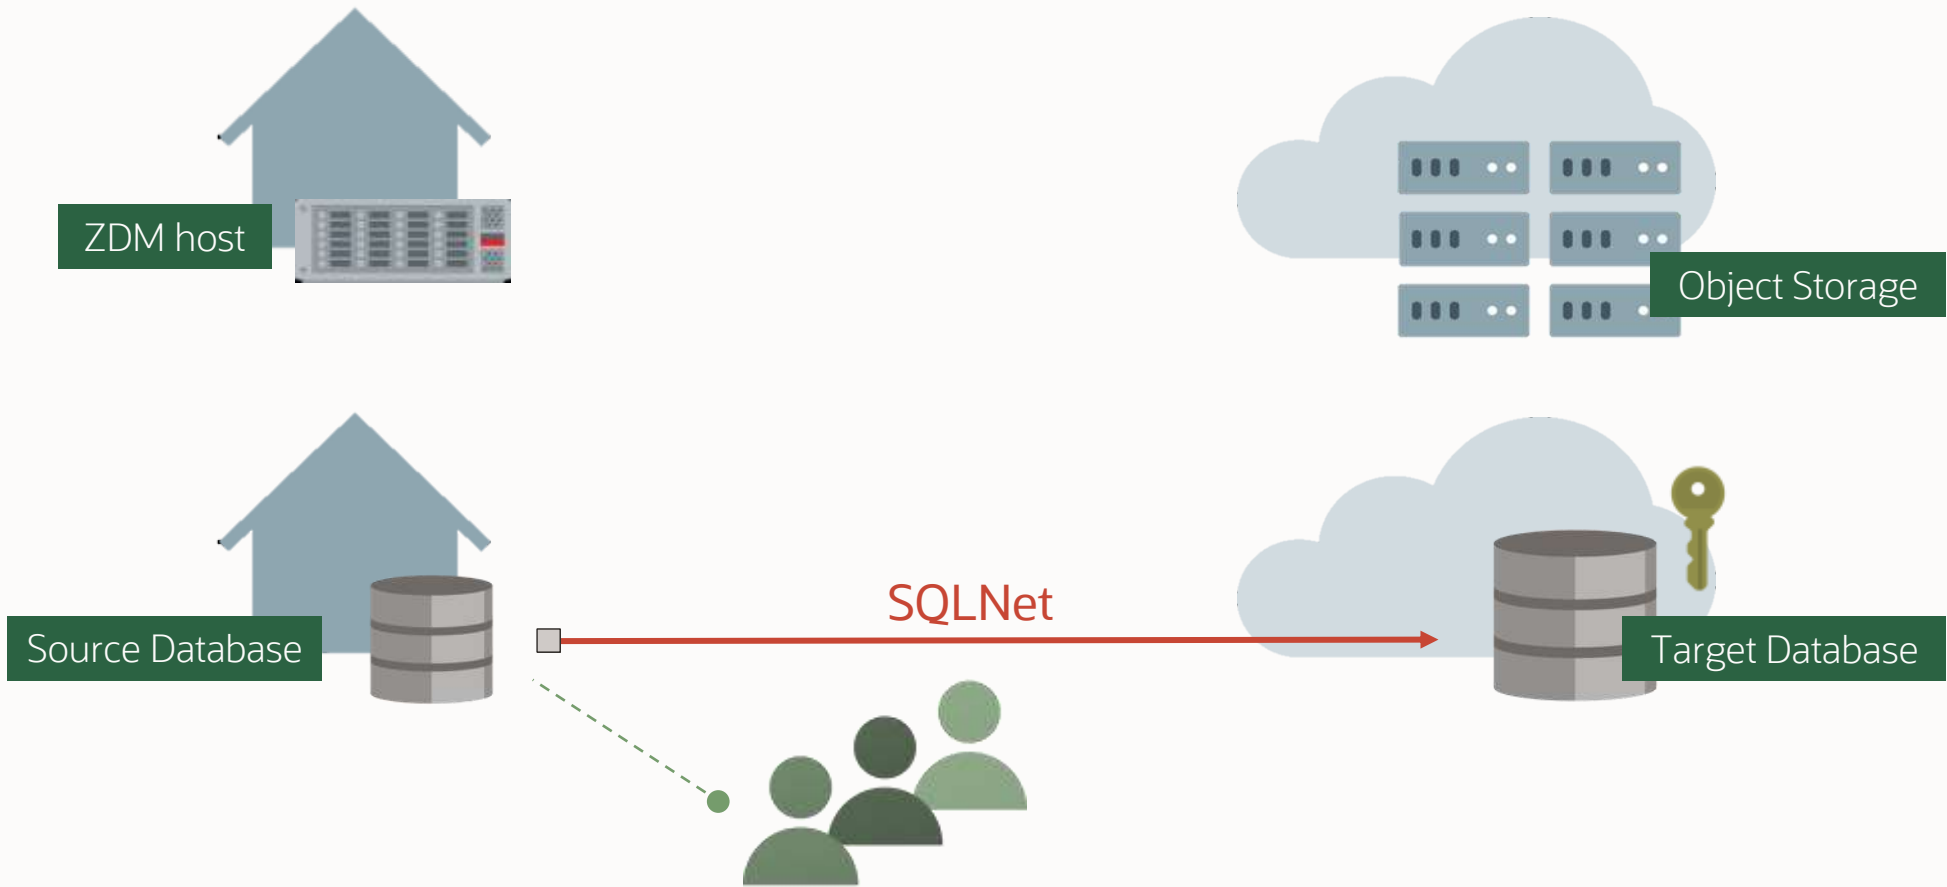
Physical Migration | Overview

Instantiate standby database from backup




Physical Migration | Overview

Synchronize via redo apply

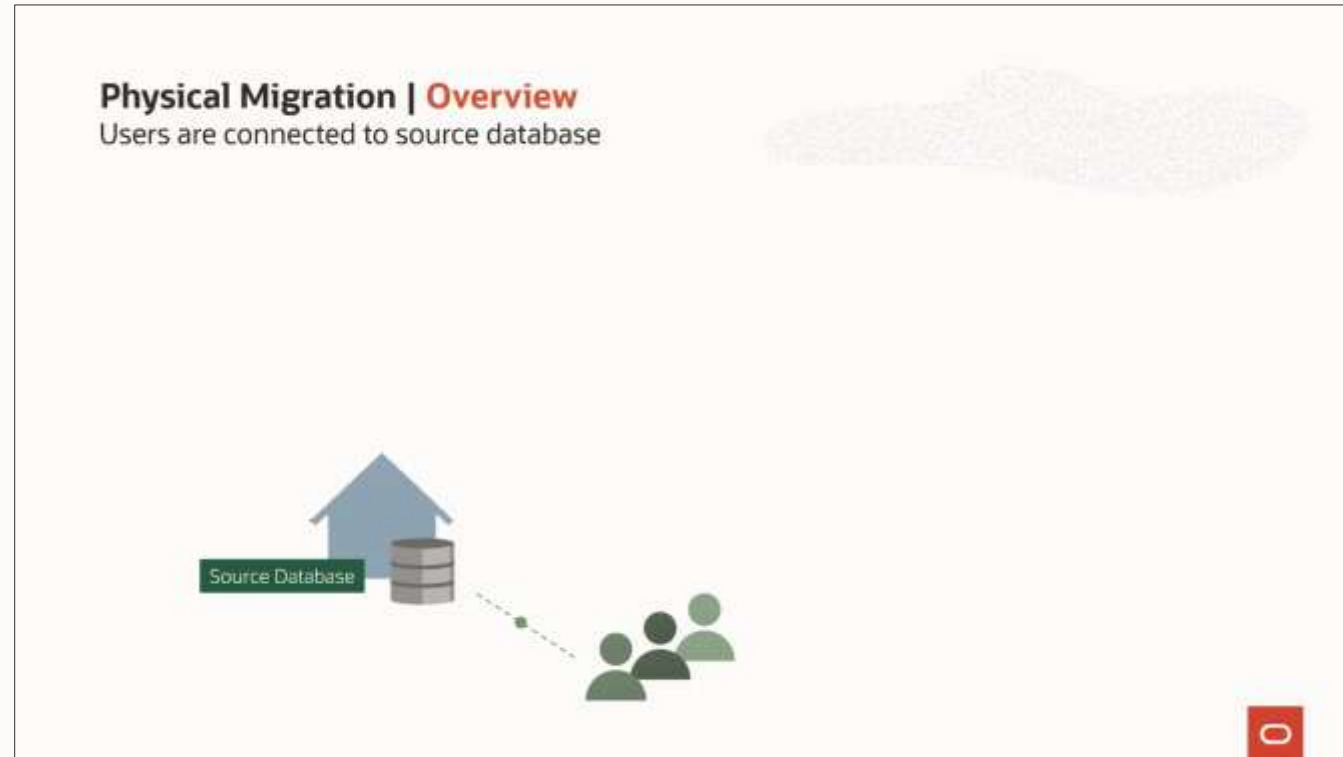


Physical Migration | Overview

At your will, switchover sessions



Physical Migration | Overview



Migration | Options



PHYSICAL

LOGICAL

Online: Data Guard + switchover

Offline: Backup + restore

Standard Edition - offline only

Online: Data Pump + GoldenGate

Offline: Data Pump


Via dump file or database link

Standard Edition - any approach



Logical Migration | Overview

Users are connected to source database



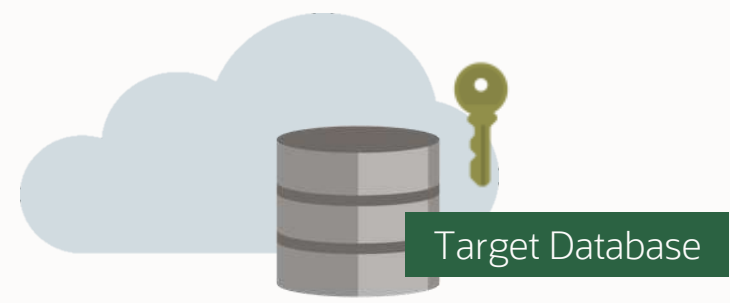
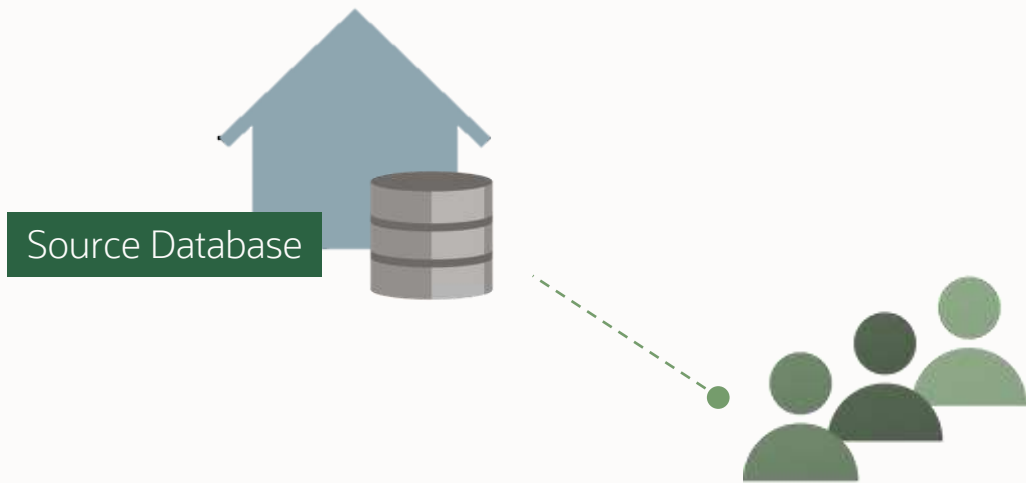
Logical Migration | Overview

Provision target database in OCI




Logical Migration | Overview

Download and install ZDM




Logical Migration | Overview

ZDM connects to source and target database



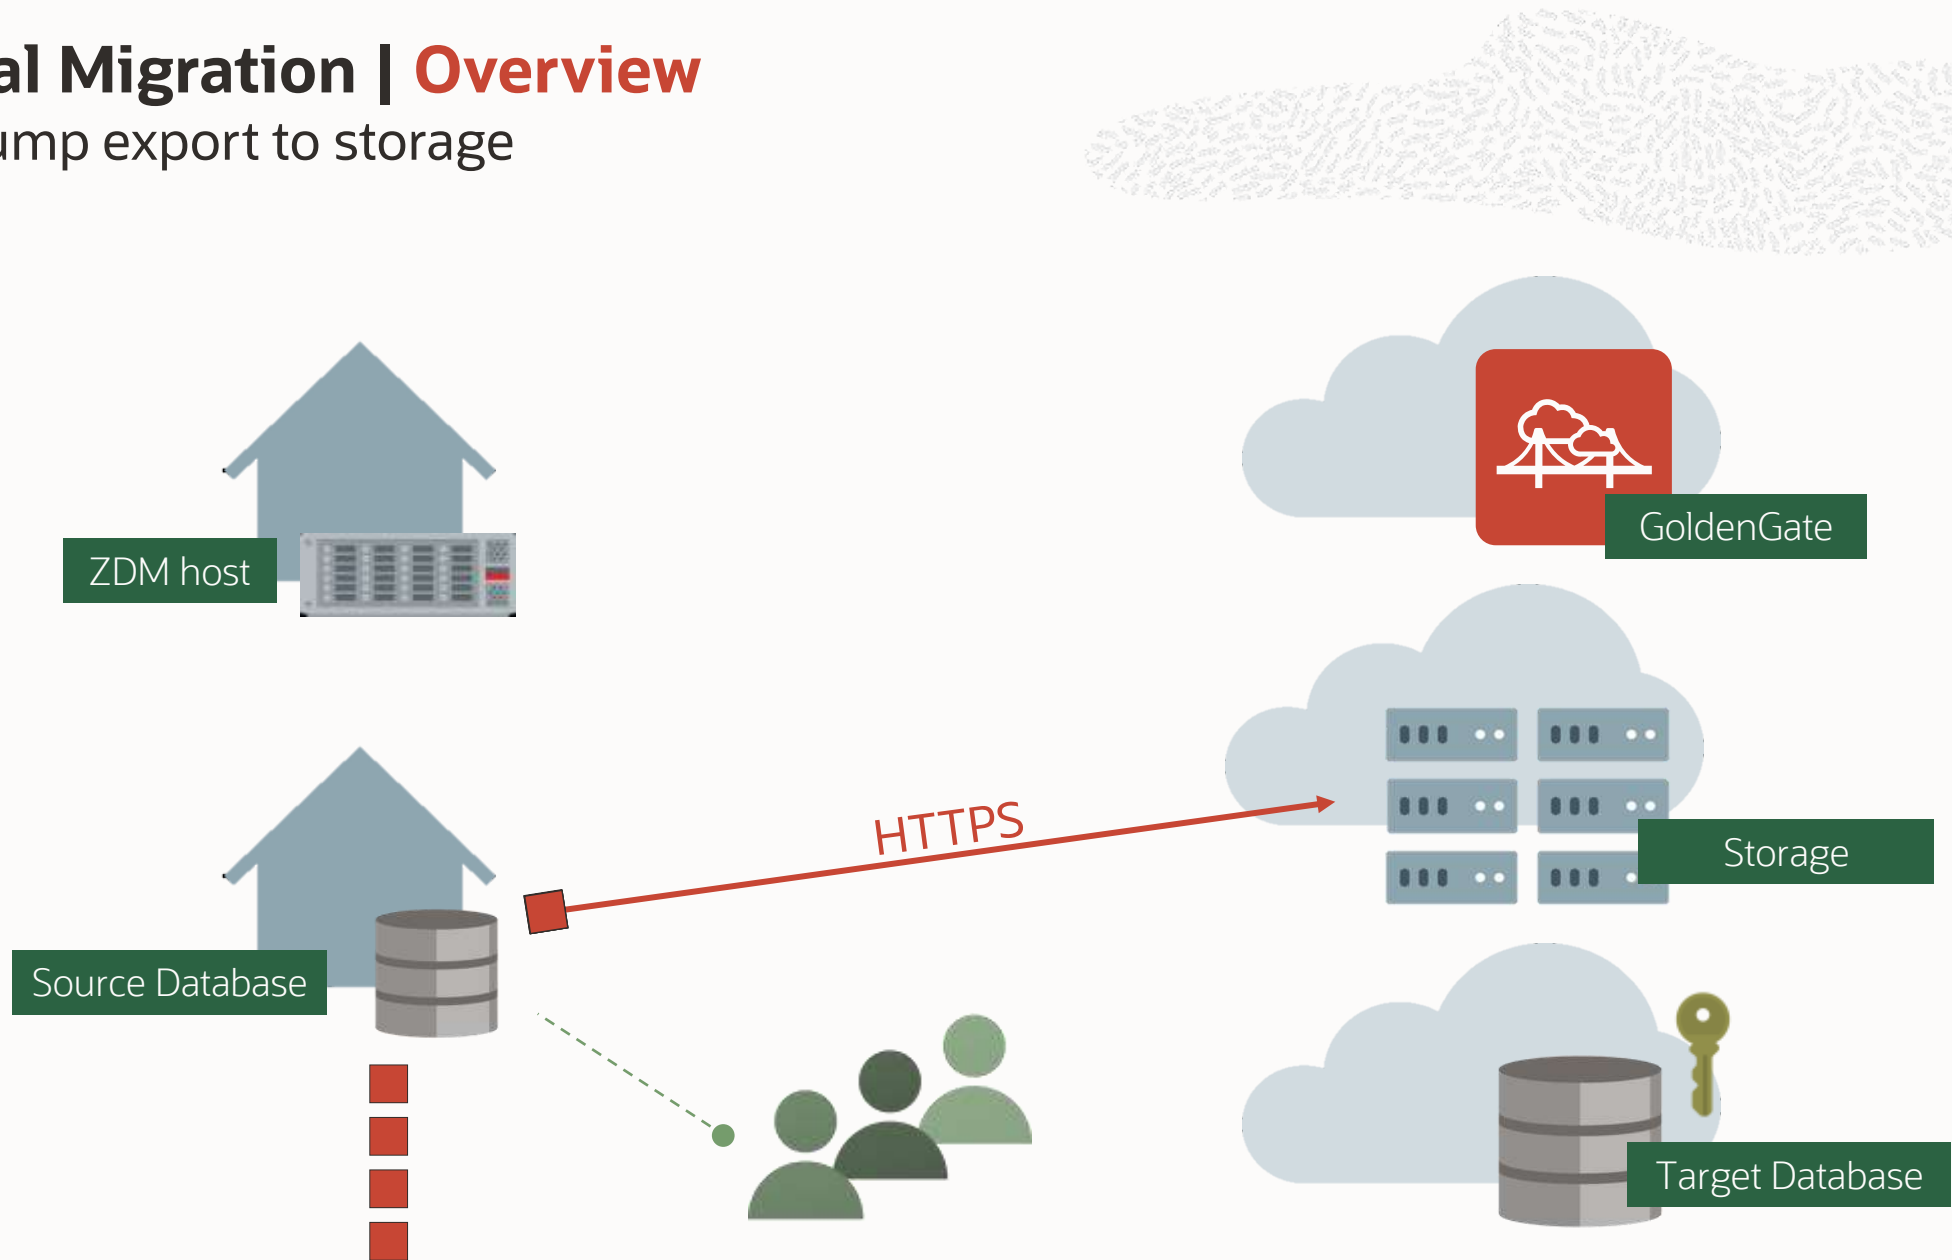
Logical Migration | Overview

Provision GoldenGate and capture on source



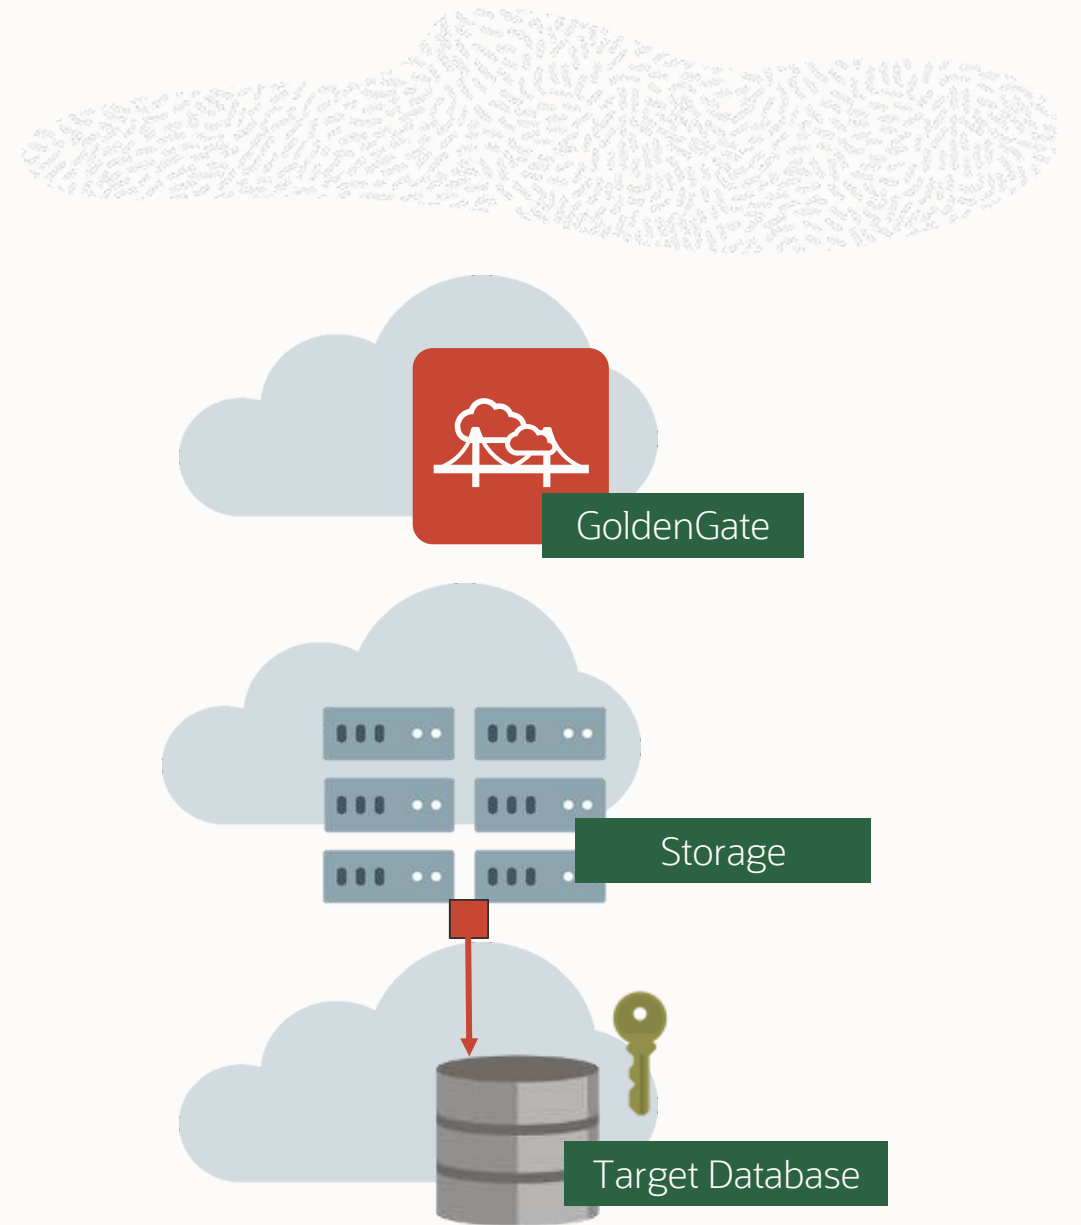

Logical Migration | Overview

Data Pump export to storage



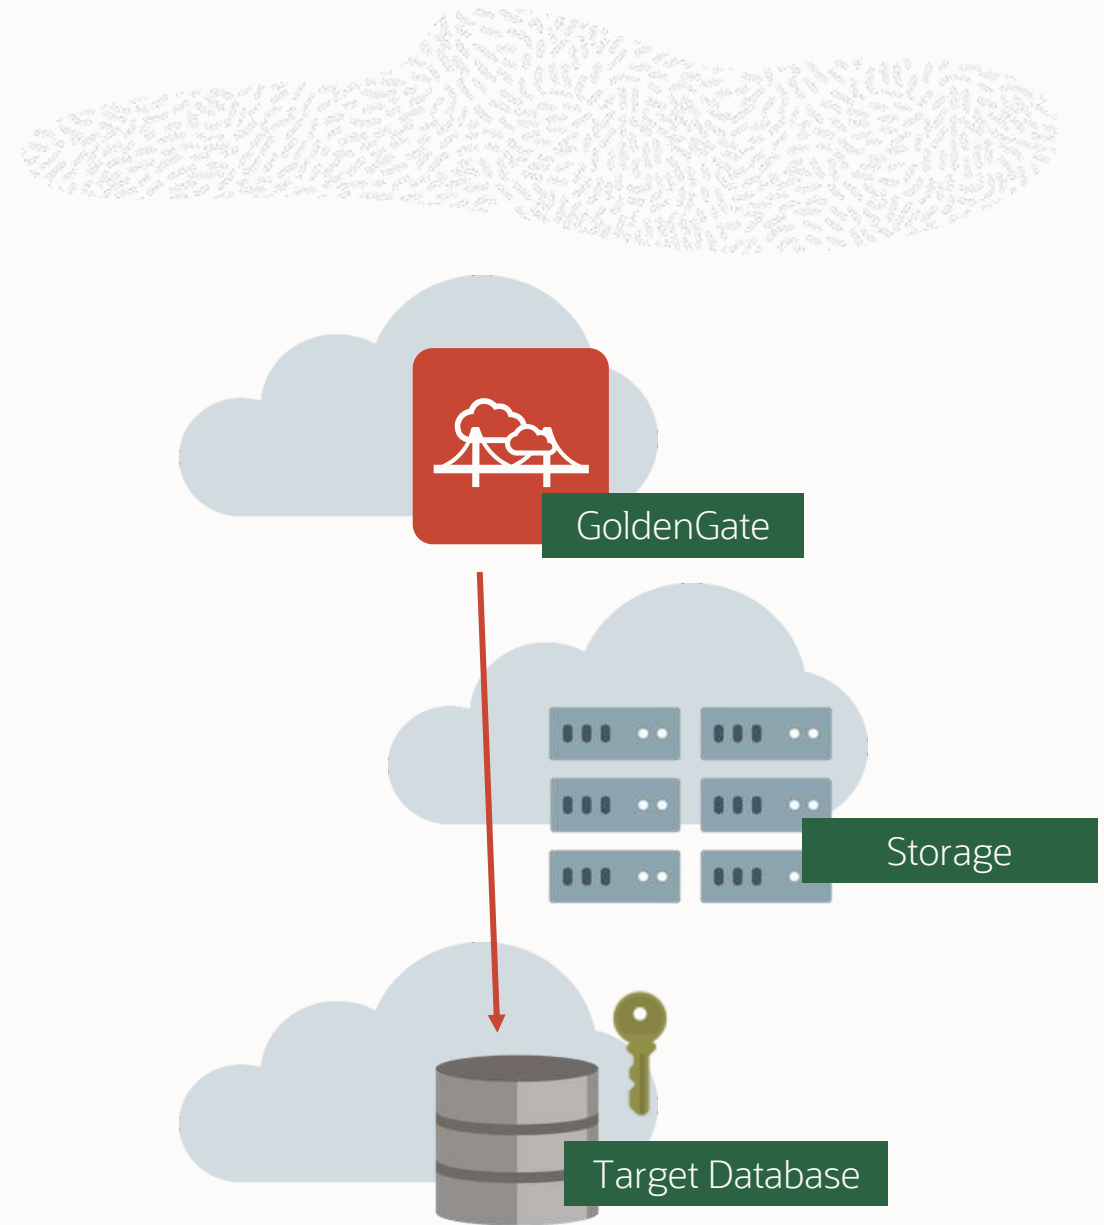
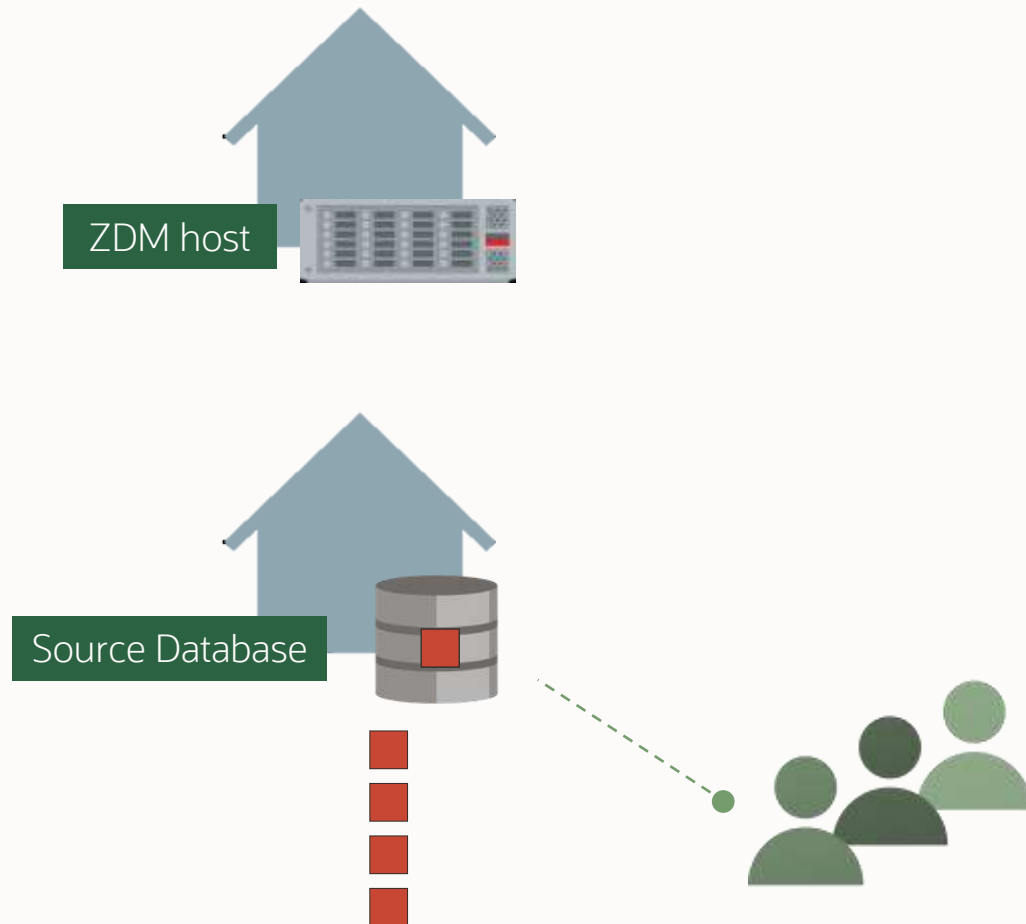
Logical Migration | Overview

Data Pump import from storage



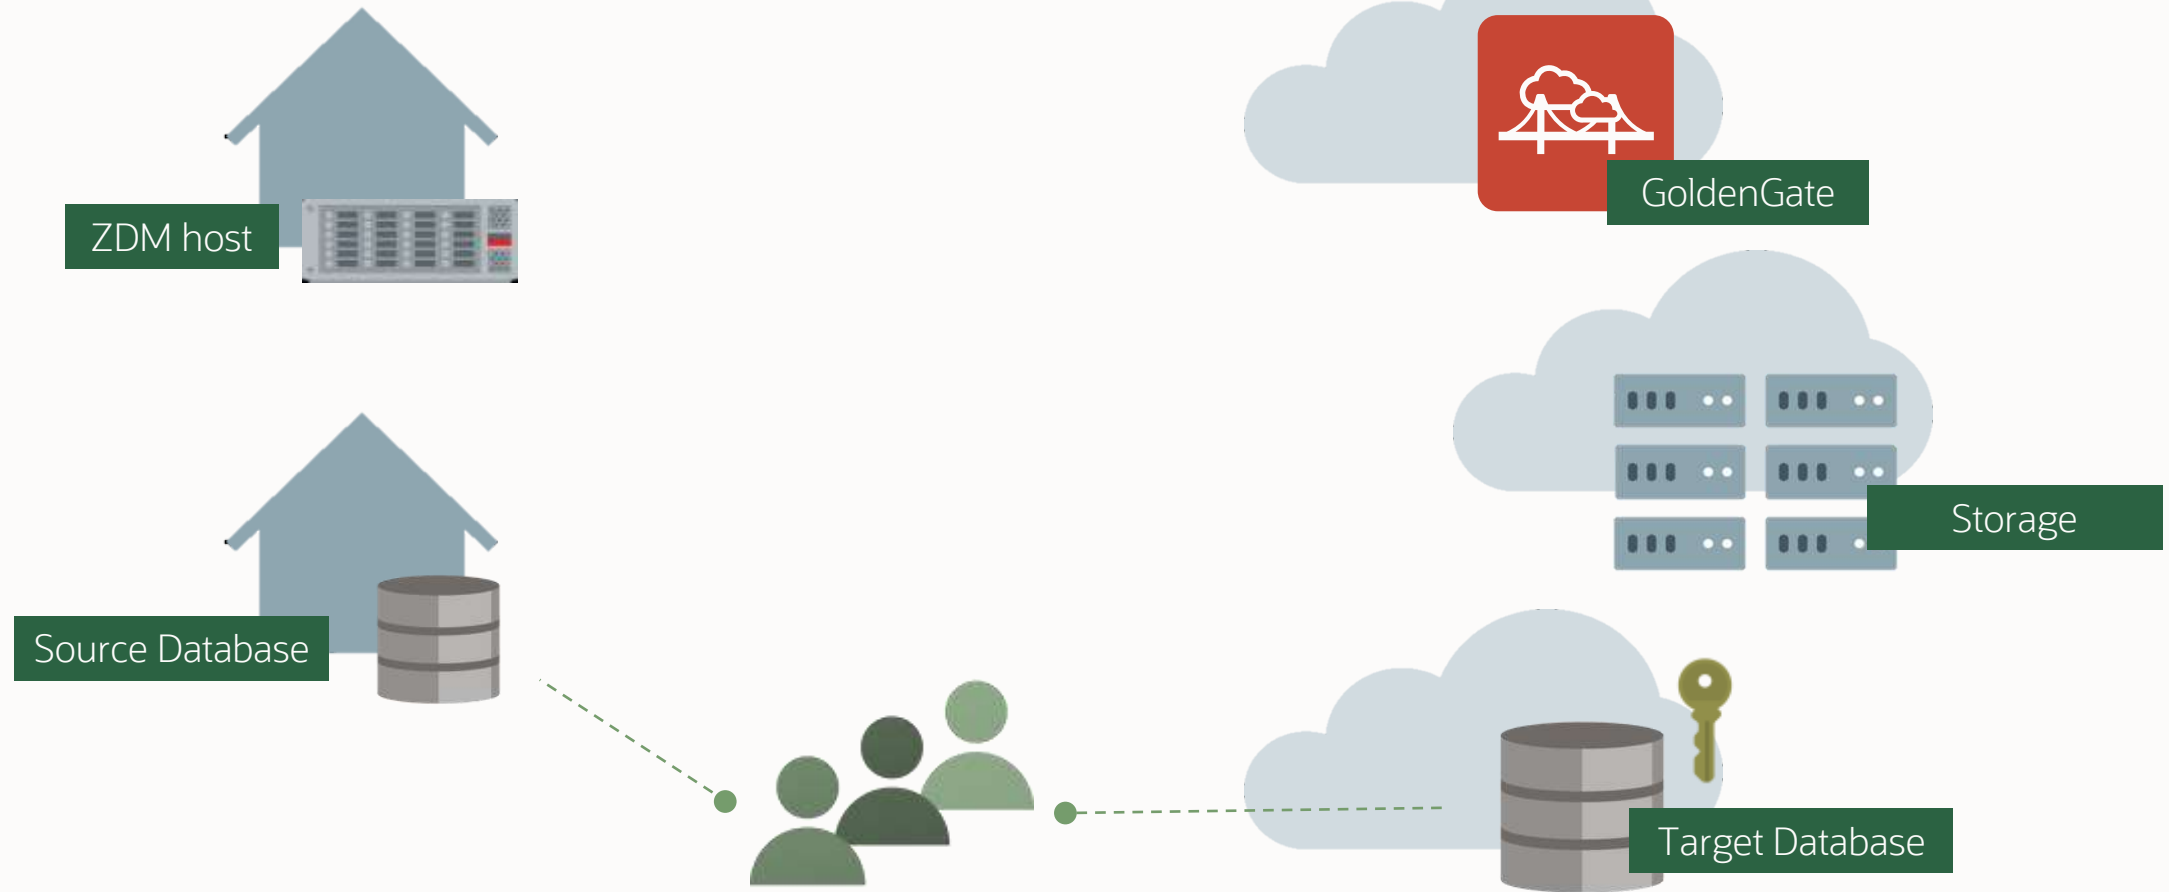
Logical Migration | Overview

Configure apply on target



Logical Migration | Overview

At your will, switchover sessions



Logical Migration | Overview

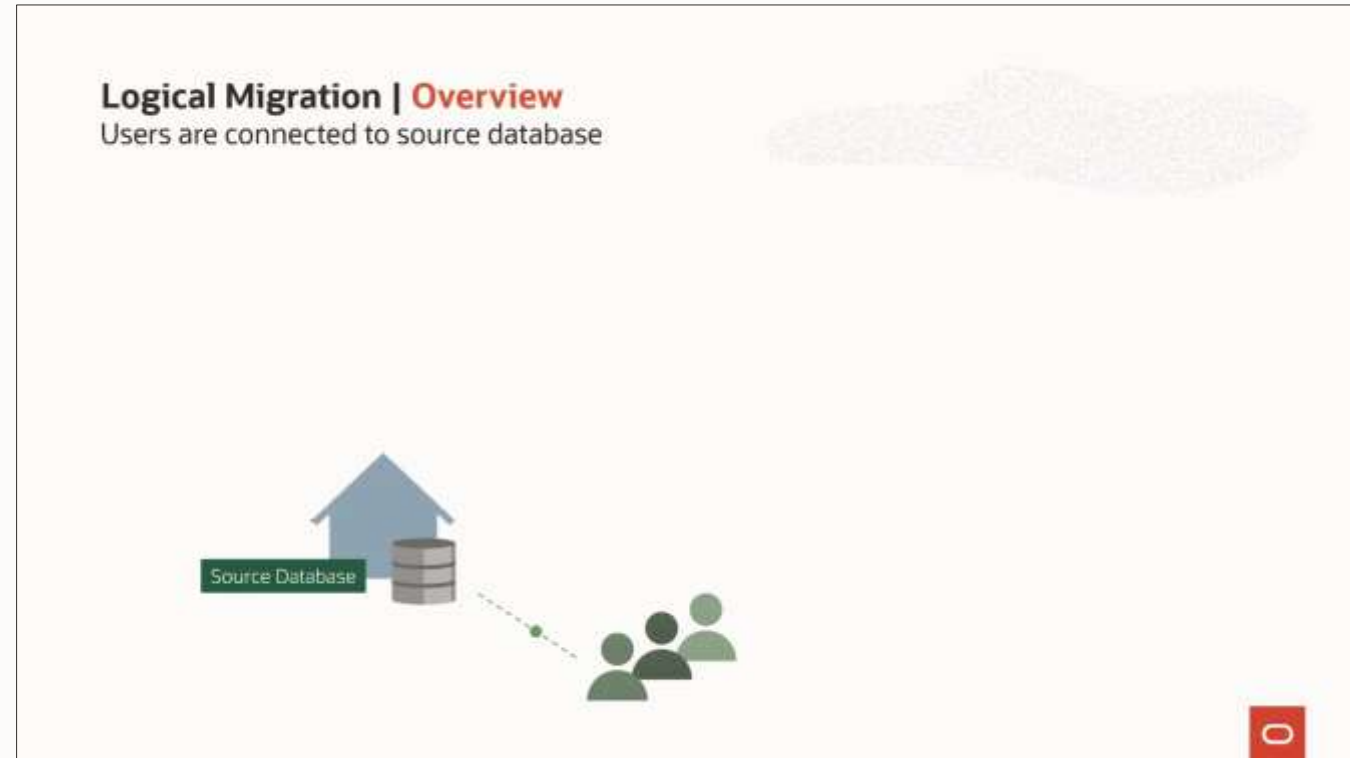




Photo by [Anne Nygård](#) on [Unsplash](#)

Install

ZDM Service Host

Service Host | Prerequisites

- Oracle Linux 7
- Install on dedicated server (recommended)
- Network connectivity to source and target database host
- Check the [documentation](#)

Service Host | SSH Key Format

SSH key format must be PEM format

```
$ ssh-keygen -t rsa -m PEM

$ cat ~/.ssh/id_rsa

-----BEGIN RSA PRIVATE KEY-----
MIIG5QIBAAKCAyEA0GPlWoSFfU8+6zgOymj47d9NTxRJYr5U9seFAcz3/aaWEP5k
ZT0FjipCIziBcnYzs0jKPLSrSoPnYGJxJuYbDj6pwMNH/f0SfhAibjHD3+Buj5cc
...
```

Service Host | **Installation**

Provision Oracle Linux Machine



[Watch on YouTube](#)





Migration

Physical

Migration | Options



PHYSICAL

LOGICAL

Online: Data Guard + switchover

Offline: Backup + restore

Standard Edition - offline only

Online: Data Pump + GoldenGate

Offline: Data Pump

Via dump file or database link

Standard Edition - any approach




Physical Migration | Demo



[Watch on YouTube](#)



Physical Migration | **Testing**



[Watch on YouTube](#)

Physical Migration | Benefits

- Well-known method
- Seamless switchover with properly configured application
- Excellent fallback
 - Requires license for ASO on-prem
 - 1-year limited license for ASO available
- Migrate entire database
 - AWR
 - SQL Plan Baseline
 - SQL Profiles
 - ...



Physical Migration | **Benefits**

- RMAN compression automatically applied
 - Advanced Compression Option not needed during migration

Physical Migration | Considerations

- Migrate to same version only
- Convert to PDB requires additional downtime
- Entire database is migrated
 - *Old baggage*
- Standard Edition is offline only (backup/restore)

Physical Migration | **Very Large Databases**

- Can co-exist with existing Data Guard
 - Switchovers supported
 - Must be disabled prior to OCI switchover
- Level 0 backup is required
 - Backup is streamed directly to OCI
 - The faster the connection to OCI, the faster the backup
 - Must suspend other backup activities
 - No disk space required for backup
 - Potentially large amount of archive logs must be stored on disk
 - Number of channels and compression algorithm configurable
 - Exadata on-prem and ExaCC can use existing backup or ZDLRA



Physical Migration | **Very Large Databases**

- Redo can be compressed
 - Requires Advanced Compression Option
 - [How To Calculate The Required Network Bandwidth Transfer Of Redo In Data Guard Environments \(Doc ID 736755.1\)](#)
- Automatic backup and Data Guard must be configured on OCI database **after** switchover
 - No support for cascading standby
 - Increases downtime



REDO APPLY

benchmark

Redo apply, TB/Day	11.2.0.4	12.1.0.2	12.2	MIRA 2x 12.2	MIRA 4x 12.2
Batch	57	57	57	115	226
OLTP	14	15	15	29	60

Source: [Redo Apply Best Practices – Oracle Data Guard and Active Data Guard](#)



redo

TRANSPORT AND APPLY

benchmark

Connection, Gbps	11.2.0.4	12.1.0.2	12.2	MIRA 2x 12.2	MIRA 4x 12.2
Batch	57 / 6	57 / 6	57 / 6	115 / 11	226 / 22
OLTP	14 / 2	15 / 2	15 / 2	29 / 3	60 / 6

Source: [Redo Apply Best Practices – Oracle Data Guard and Active Data Guard](#)



Physical Migration | Redo Apply Best Practices

Redo Apply Best Practices - Oracle Data Guard and Active Data Guard

Physical Migration | Different Patch Level

You can migrate to a higher patch level

- Example: 19.7.0 to 19.11.0

Procedure

- Switch over to OCI database
- ZDM executes `datapatch`
- Patch on-premises Oracle Home

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

Physical Migration | **Fallback**

Caution: Fallback requires license for Advanced Security Option on source database
Get 1-year limited license for Advanced Security Option when you migrate to OCI

Procedure ([MAA Practices for Cloud Migration Using ZDM \(Doc ID 2562063.1\)](#)):

```
$ $ZDM_HOME/bin/zdmcli abort job -jobid n

--OCI database
SQL> ALTER DATABASE COMMIT TO SWITCHOVER TO PHYSICAL STANDBY;

--On-premises database
srvctl modify database -d CDB1 -role primary -startoption open
SQL> ALTER DATABASE COMMIT TO SWITCHOVER TO PRIMARY;
SQL> ALTER DATABASE OPEN;
```

Physical Migration | Q&A

Q: Does ZDM support source database using Oracle Key Vault for TDE keys?

A: No



Photo by [eelias](#) on [Unsplash](#)

Migration

Logical

Logical Migration | Options



PHYSICAL

Online: Data Guard + switchover

Offline: Backup + restore

Standard Edition - offline only

LOGICAL

Online: Data Pump + GoldenGate

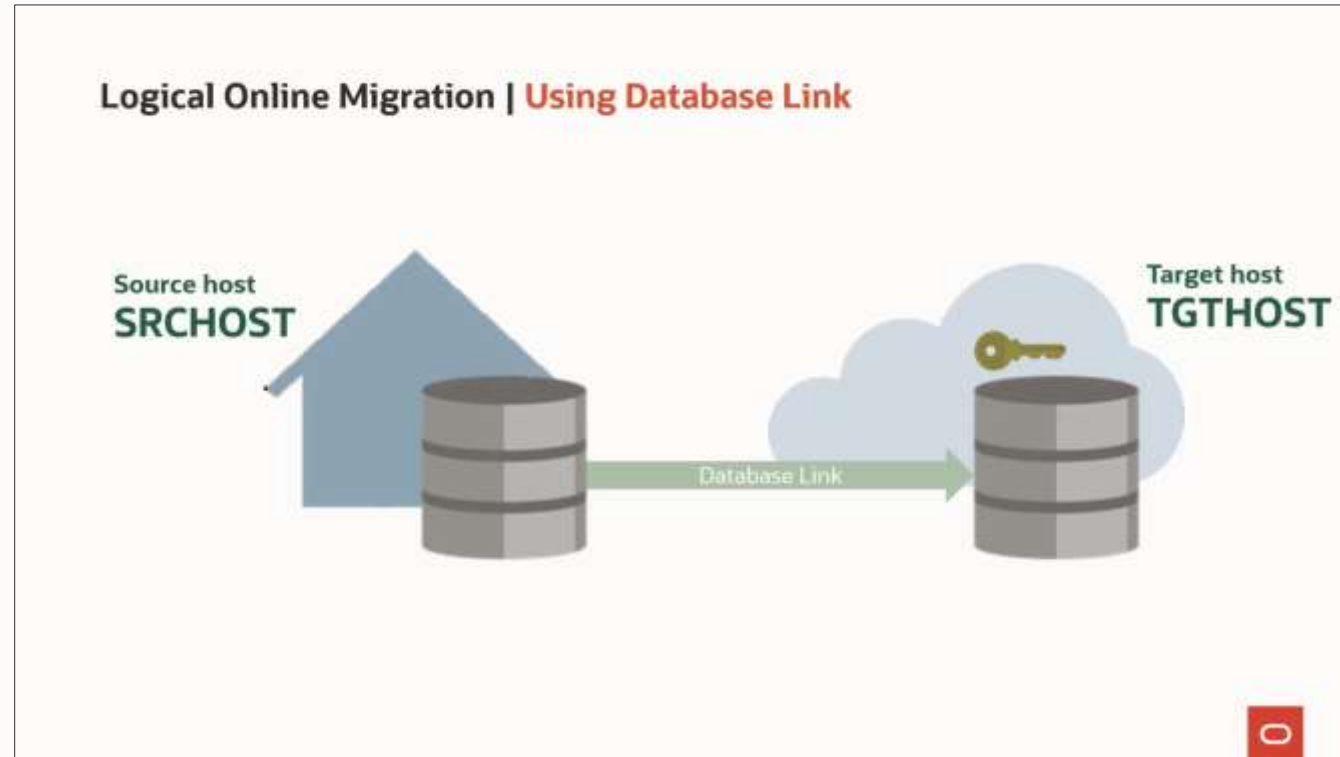
Offline: Data Pump

Via dump file or database link

Standard Edition - any approach



Logical Migration | Demo



[Watch on YouTube](#)

Logical Migration | GoldenGate

When I say Oracle GoldenGate

What do **you** say?

Logical Migration | GoldenGate



”

Oracle GoldenGate for Oracle – Database Migrations can be used for 183 days to perform migrations into Oracle databases located in Oracle Cloud Infrastructure using the following tools:

Oracle Zero Downtime Migration

Oracle Cloud Infrastructure Database Migration

[Cloud Marketplace: Oracle GoldenGate for Oracle – Database Migrations](#)

- Applies to the Cloud Marketplace image **only**
- Covers the source and target database
- For migration to OCI including ExaCC only

Logical Migration | GoldenGate




”

*The image also contains a docker image of Oracle GoldenGate 21c for Oracle Database for download; the docker image is used with Oracle Zero Downtime Migration to migrate to Oracle **Exadata Cloud@Customer** targets*

[Cloud Marketplace: Oracle GoldenGate for Oracle – Database Migrations](#)

- For ExaCC migrations avoid using GoldenGate in OCI
- [Deploy Docker image in your own data center](#)
- Get the same benefits as ExaCS migration

Logical Migration | GoldenGate



[Watch on YouTube](#)

Logical Migration | GoldenGate

Installation

Deploy from Cloud Marketplace

Configuraton

Handled by ZDM

Monitoring

Via GoldenGate Hub
(as seen in demo)



Logical Migration | Recommendations

Use Data Pump in schema mode

```
DATAPUMPSETTINGS_JOBMODE=SCHEMA  
INCLUDEOBJECTS-1=owner:SH  
INCLUDEOBJECTS-2=owner:OE
```

Use parallel option (on-prem = 2 x physical cores - OCI = number of OCPUs)

```
DATAPUMPSETTINGS_DATAPUMPPARAMETERS_EXPORTPARALLELSMDEGREE=n  
DATAPUMPSETTINGS_DATAPUMPPARAMETERS_IMPORTPARALLELSMDEGREE=n
```

Optionally, change the ignorable Data Pump errors

- Default: ORA-31684, ORA-39111, ORA-39082

Pro Tip: ZDM automatically adds compression and encryption to Data Pump exports



Logical Migration | Recommendations

Gather dictionary stats after the initial Data Pump load:

```
exec dbms_stats.gather_schema_stats('SYS');  
exec dbms_stats.gather_schema_stats('SYSTEM');  
  
--or  
  
exec dbms_stats.gather_dictionary_stats;
```

Logical Migration | CPAT

Get up-to-date recommendations on your migration

- You can run CPAT manually
- Replace CPAT in ZDM with newer version

[Cloud Premigration Advisor Tool \(CPAT\) Analyzes Databases for Suitability of Cloud Migration \(Doc ID 2758371.1\)](#)

Logical Migration | CPAT

Sample output

```
Premigration advisor output:  
Cloud Premigration Advisor Tool Version 21.0.0  
Cloud Premigration Advisor Tool completed with overall result: WARNING  
Cloud Premigration Advisor Tool generated report location:  
/u01/app/oracle/zdm/zdm_SALES_fra3wg_1/out/premigration_advisor_report.json  
RESULT: WARNING  
  
Schemas Analyzed (1): SH  
A total of 15 checks were performed  
There were 0 checks with FATAL results  
There were 0 checks with BLOCKER results  
There were 3 checks with WARNING results
```

Logical Migration | CPAT

Sample output

```
timezone_table_compatibility_higher
```

```
RESULT: WARNING
```

```
DESCRIPTION: The source database TZ_VERSION cannot be lower than the target  
TZ_VERSION.
```

```
ACTION: Request that Cloud Database Support change the Timezone Version on  
your target Database.
```

Logical Migration | Benefits

- Free features
 - Oracle GoldenGate (via Cloud Marketplace)
 - Data Pump Compression
 - Data Pump Encryption
- Optionally, remodel your schema and data
 - Migrate to SecureFile LOBs is default



Logical Migration | SecureFile



” *SecureFiles is the default storage mechanism for LOBs starting with Oracle Database 12c, and Oracle **strongly recommends SecureFiles** for storing and managing LOBs, rather than BasicFiles. BasicFiles will be deprecated in a future release.*

[Database SecureFiles and Large Objects Developer's Guide](#)

Always transform LOBs to SecureFiles LOBs

```
$ impdp ... TRANSFORM=LOB_STORAGE:SECUREFILE
```

Logical Migration | SecureFile



Importing as BasicFiles

```
10-OCT-20 21:43:21.848: W-3 . . imported "SCHEMA"."TABLE" 31.83 GB 681025 rows in 804 seconds using direct_path
```

Importing as SecureFiles

```
15-OCT-20 18:16:48.663: W-13 . . imported "SCHEMA"."TABLE" 31.83 GB 681025 rows in 261 seconds using external_table
```



Logical Migration | Benefits

- Migrate to higher release
- Migrate directly into a PDB
- Online option for SE2
- Configure backup of target database in advance
 - Test backup/restore functionality
- Configure Data Guard in advance
 - Protect database immediately after switch over to OCI

Logical Migration | Considerations

- Target database time zone file version must be equal to or higher than source

```
SQL> select * from v$timezone_file;
```

- Possibly patches are recommended on source database
 - [11g](#)
 - [12c and newer](#)
- GoldenGate [supported data types](#)

Logical Migration | Considerations

- Export or re-create public and other not exported objects
 - Synonyms
 - Database links
 - ...
- Diagnostic and tuning related information
 - AWR
 - SQL Plan Baselines
 - SQL Profiles
 - SQL Patches
 - ...



Logical Migration | Considerations

- Statistics are not exported by Data Pump
 - According to our best practices
 - `exclude=statistics`
- Gather at target after import
 - Requires time
 - Column usage information (`COL_USAGE$`) is not populated
 - Histograms are not created if `method_opt` is `SIZE AUTO`
 - Populate from source database via database link using [DBMS_STATS.MERGE COL USAGE](#)
 - Table, schema or database statistics preferences not present
 - `DBMS_STATS.SET_TABLE_PREFS`



Logical Migration | Considerations

- Or, use DBMS_STATS
 - See webinar [Performance Stability, Tips and Tricks and Underscores](#)
 - Remember to transfer database/schema/table statistics preferences as well
 - Like `method_opt`, `degree`, `stale_pct` etc.
 - `DBMS_STATS.EXPORT_TABLE_PREFS`

Logical Migration | Considerations

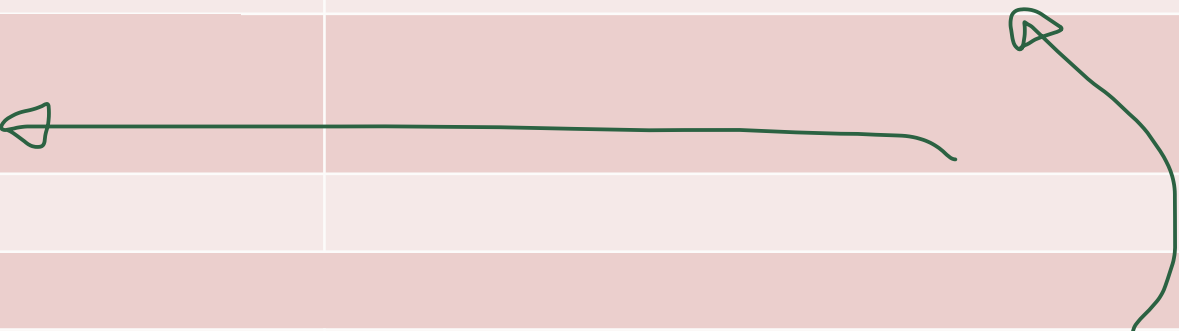
- Fallback
- DDL is not replicated
 - Unless you manually change the GoldenGate configuration
- Character set must match



Logical Migration | Considerations

- Sequences must be *forwarded* on switch-over

SOURCE	TARGET
Sequence <i>seq1.currval</i> = 100	



The diagram illustrates the concept of sequence forwarding during a logical migration. It features a table with two columns: 'SOURCE' and 'TARGET'. The first row of the table contains the text 'Sequence seq1.currval = 100' under the 'SOURCE' column. A green arrow originates from the 'TARGET' column, curves downwards and to the left, and then points back to the 'SOURCE' column, indicating that the sequence value is being forwarded from the target back to the source.



Logical Migration | Considerations

- Sequences must be *forwarded* on switch-over
 - Recreate sequences after switchover
 - Set sequence

```
alter sequence seq1 increment by 5473;  
select seq1.nextval from dual;  
alter sequence seq1 increment by 1;
```

Increment by source nextval - target nextval

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



Logical Migration | Testing

- Oracle GoldenGate supports Flashback Database
 - [Does Goldengate Support Oracle RDBMS Flashback Features? \(Doc ID 966212.1\)](#)
- Use backup/restore or cloning for
 - Autonomous Database
 - Standard Edition 2

Logical Migration | **Very Large Databases**

- Can co-exist with existing Data Guard
 - Switchover and failovers not supported
- Scaling up on CPUs is advantageous
- Data Pump export does **not** use `FLASHBACK_SCN` or `FLASHBACK_TIME`
- GoldenGate trail files typically
 - 30-40 % of redo
 - Compress at least 1:4, most likely up to 1:8
- Automatic backup and Data Guard can be configured on OCI database **before** switchover



Logical Migration | What if ...

I have an EBS database?

- EBS has not verified the use of GoldenGate for migrations. The certified way of migrating EBS databases is transportable tablespaces

I want to use Oracle GoldenGate Veridata?

- It is not part of the free offering of Oracle GoldenGate. You to need to license that separately. For basic verification consider using [DBMS_COMPARISON](#)

I have a table without a primary or unique key?

- Not a problem if all scalar columns together guarantee uniqueness. Replication on such tables will suffer performance-wise.
If uniqueness can't be guaranteed in any way, then table is not supported

Logical Migration | What if ...

I have a big-Endian database?

- Talk to us! Perhaps you can become a beta test for ZDM on other platforms. Otherwise check our webinars [Migration Strategies – Insights, Tips and Secrets](#) and [Move to the Cloud – Not only for techies](#)

My target is an ExaCC? I am not allowed to use OGG in OCI?

- For ExaCC migrations you can [get a Docker image containing Oracle GoldenGate that you can use for migrations to ExaCC](#). The Docker image can be deployed in your own data center and avoid remote OCI resources



Photo by [Alexander Andrews](#) on [Unsplash](#)

Details

Pro Tips | Troubleshooting

ZDM service host

- `$ZDM_BASE/chkbase/scheduled`
- `$ZDM_BASE/crsdata/hostname/rhp`

Source and target hosts

- `$ORACLE_BASE/zdm/zdm_<db_unique_name_<zdm job id>/zdm/log`

Clear ZDM logs for easier troubleshooting

```
$ $ZDM_HOME/bin/zdmservice stop
$ rm $ZDM_BASE/crsdata/*/rhp/rhpserver.log*
$ rm $ZDM_BASE/chkbase/scheduled/*
$ $ZDM_HOME/bin/zdmservice start
```

Pro Tip: You can abort a job using
`zdmcli abort job -jobid n`



Pro Tips | Troubleshooting

Other sources:

- Alert log
- Data Pump process trace file DM00
- Data Pump log file
 - Directory referenced by directory object
 - `$ORACLE_HOME/rdbms/log/<PDB GUID>`
- GoldenGate logs
 - `/u02/deployments/ogg_deployment_name/var/log`

Pro Tip: Before creating a Service Request: [SRDC - Data Collection For Database Migration Using Zero Downtime Migration \(ZDM\) \(Doc ID 2595205.1\)](#)



Pro Tips | **Troubleshooting**

Oracle Zero Downtime Migration 21 .1 Release Notes

- [Troubleshooting](#)
- [Known Issues](#)

Move to Oracle Cloud Using Zero Downtime Migration

- [Troubleshooting](#)

Pro Tips | ZDM Log File

Tailing migration log file:

```
$ $ZDM_HOME/bin/zdmcli migrate database \  
-rsp /home/zdmuser/std.rsp \  
...  
  
$ tail -n 50 -f "`ls -td /u01/app/oracle/chkbase/scheduled/* | head -1`"  
  
zdmhost: 2021-05-06T18:14:25.590Z : Starting zero downtime migrate operation ...  
zdmhost: 2021-05-06T18:14:25.625Z : Executing phase ZDM_VALIDATE_TGT  
zdmhost: 2021-05-06T18:14:25.634Z : Fetching details of user-managed OCI database "ocidl.database..."  
zdmhost: 2021-05-06T18:14:26.840Z : Lifecycle state of OCI database "ocidl.database...": "Available"  
zdmhost: 2021-05-06T18:14:29.365Z : Type of OCI database "ocidl.database...": "Virtual Machine Database System"  
zdmhost: 2021-05-06T18:14:29.466Z : Verifying configuration and status of target database "sales"  
zdmhost: 2021-05-06T18:14:33.889Z : Global database name: SALES.SUB02121342350.DANIEL.ORACLEVCN.COM  
zdmhost: 2021-05-06T18:14:33.890Z : Target PDB name : SALES  
zdmhost: 2021-05-06T18:14:33.891Z : Database major version : 19  
zdmhost: 2021-05-06T18:14:36.711Z : Database parameter ENABLE_GOLDENGATE_REPLICATION is set to true.  
zdmhost: 2021-05-06T18:14:36.713Z : Oracle GoldenGate database admin user "GGADMIN" has required privileges.  
zdmhost: 2021-05-06T18:14:36.714Z : Execution of phase ZDM_VALIDATE_TGT completed  
zdmhost: 2021-05-06T18:14:36.742Z : Executing phase ZDM_VALIDATE_SRC
```

Pro Tips | Custom Scripts

Run your own script before or after any phase

The script is executed on either source or target

- Autonomous DB only .sql scripts are possible

Relevant information is available as environment variables

- Database
- Oracle Home
- ZDM Phase
- ...

Pro Tip:

To list **all phases**: `zdmcli migrate database -rsp -listphases`



Pro Tips | GoldenGate Certificate

GoldenGate Hub provided by OCI Marketplace image comes with a self-signed certificate

Best:

- [Implement your own properly signed certificates](#)

Or, for test environments:

- [GoldenGate and self signed certificate? Zero Downtime Migration - GoldenGate Hub Certificate Known Issues \(Doc ID 2768483.1\)](#)

Pro Tips | GoldenGate Privileges

Replicat process connects a regular database user

Determine how to grant privileges:

- DBA / PDB_DBA role
- INSERT ANY, DELETE ANY, UPDATE ANY
- Connor McDonald's schema grant

In doubt? Do like in Autonomous Databases...

Pro Tips | GoldenGate Health Check

Generate report:

- Check prerequisites
- Database characteristics
- Find database objects of interest
- Extract/replicat statistics

Oracle GoldenGate Integrated Extract/Replicat Health Check Database - SALES SUB02121342350 DANIEL.ORACLEVCN.COM | Instance - CDB1

OVERVIEW DATABASE TOOLS REPORTMAP

MENU: OVERVIEW Expand All Collapse All

General Findings section shows the results of sanity checks. Questionable results are highlighted. The details are visible in the later sections.

General Findings

COMPONENT	TYPE	NAME	ALERT	REASON	STAT	INFO
DATABASE	Configuration	RAC	IMPO	Multitenant Database (CDB/PDB) in use MDDL		
DATABASE	Configuration	streams_pool_size	GREEN	Usage: 0 Threshold: 85	sp_etc	

Back to Top

The summary of Database, Extract and Replicat is showing some basic information of the System. It contains of a static and dynamic part. Dynamic information is gathered in a 10 sec interval by default and can be changed with the PL/SQL API dbms_hc.set_parameter.

Database, Extract and Replicat Summary

Database (Instance#)		Comments
CDB1 (1)		
Current SCN (Time)	3193580 (2021-05-07 05:36:03)	Current Scn and the time
Database Version	19.0.0.0.0	Database Software version. Note that the COMPATIBLE
Database Status	ACTIVE	
Shutdown Pending	NO	
Active State	NORMAL	
Blocked	NO	
Archives	STARTED	



Pro Tips | GoldenGate Health Check

Generate report by:

- Installing objects in database: `ogghc_install.sql`
- Execute health check: `ogghc_run.sql`
- Optionally, clean-up objects: `ogghc_uninstall.sql`

For GoldenGate MicroServices Architecture find the scripts:

`/u01/app/ogg/oraclenn/lib/sql/healthcheck`

Photo by [Carolina Pimenta](#) on [Unsplash](#)

Finally ...

It is time to wrap up

Migration | Comparison



PHYSICAL

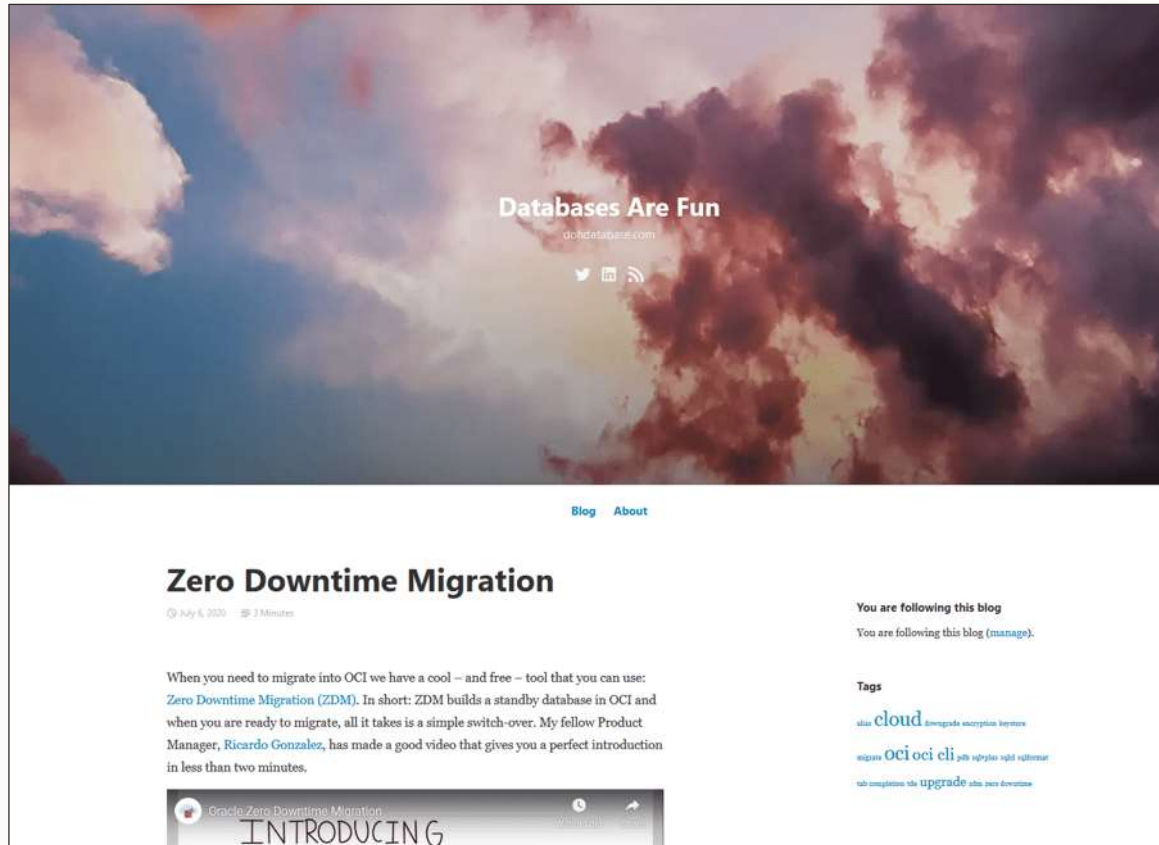
- Migrate entire database
- No data customization
- Same version / same architecture
- Well-known
- No online option for SE2
- Excellent fallback

LOGICAL

- Migrate schemas
- Remodel your data
- Cross-version / cross-architecture
- New skills
- Online option for SE2
- Complicated fallback



Wrapping Up | **Blog Posts**



A walkthrough with all the details

- *includes one on ExaCS*



Wrapping Up | YouTube



[YouTube Playlist](#)



Wrapping Up | Further Information



[Oracle Zero Downtime Migration Product Page](#)

[Oracle Zero Downtime Migration Documentation](#)

[Oracle Zero Downtime Migration Release Notes](#)

[Oracle Zero Downtime Migration Whitepaper](#)

[MAA Practices for Cloud Migration Using ZDM \(Doc ID 2562063.1\)](#)

[Hybrid Data Guard to Oracle Cloud Infrastructure](#)

MIGRATION

methods

ZDM

**DATA
GUARD**

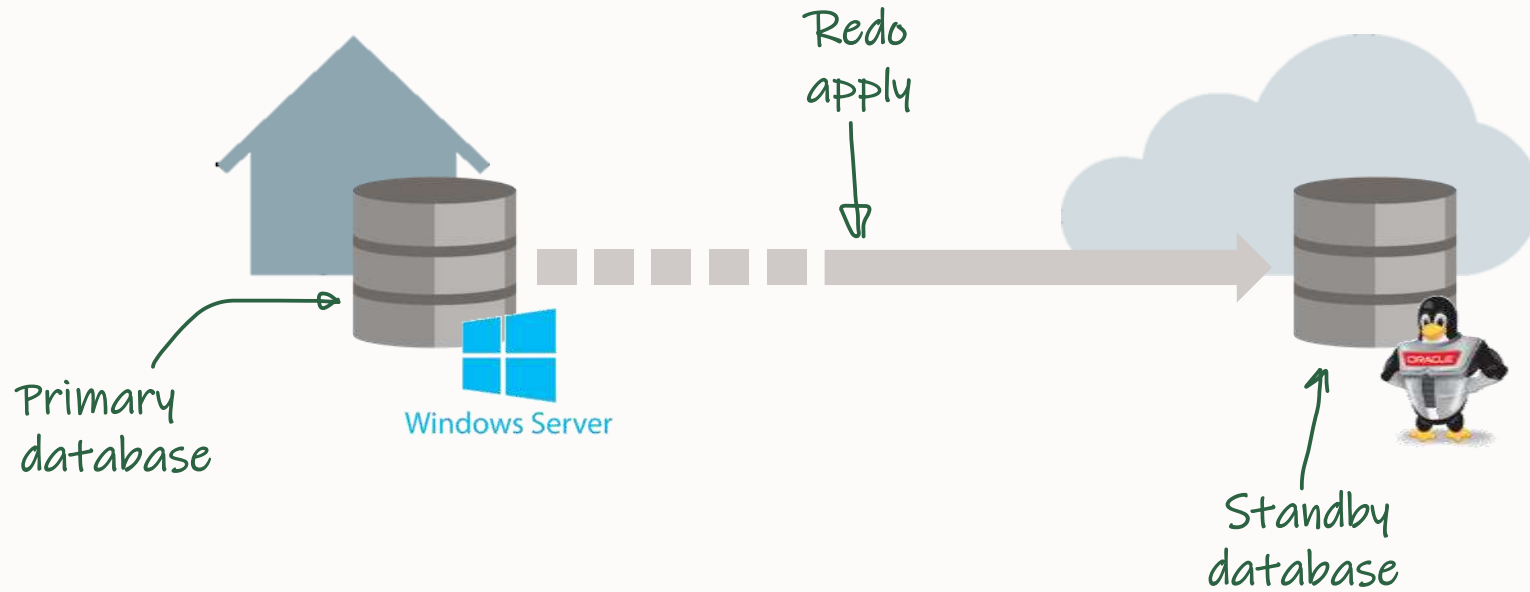
TRANS-
PORTABLE

DATA PUMP
GOLDENGATE



Heterogeneous Data Guard

Little Endian | Heterogeneous Data Guard



Little Endian | Linux Standby

★ Data Guard Support for Heterogeneous Primary and Physical Standbys in Same Data Guard Configuration (Doc ID 413484.1)

What differences are allowed between a Primary Database and a Data Guard Physical Standby Database (Redo Apply)?

This note is updated for Redo Apply and Oracle Data Guard 12c. It applies to all versions of Oracle Database 10g, 11g and Oracle Database 12c.

For information on supported configurations using Logical Standby (SQL Apply), see Support Note 1085687.1.

Scope and Application:

The simplest path when deploying Data Guard is to configure a homogeneous and symmetric primary/standby configuration. However, it is often useful to deploy a heterogeneous configuration either to utilize existing servers that happen to run different operating systems or to facilitate migrations from one platform to another with minimal downtime or risk. It is also reasonable for users to wish to reduce their disaster recovery investment by purposely configuring a standby system with less processing capacity than production, or by utilizing lower cost components than used for their primary system. Use the instructions and information provided in this support note to determine which platform combinations are supported within a single Data Guard configuration and any additional requirements or restrictions that may apply.

If a heterogeneous primary/standby configuration is under consideration, Oracle recommends that users conduct sufficient testing to be sure that required service levels will continue to be achieved following a switchover or failover to the standby system.

1. Determine the Platform ID for your primary and standby database.

You can find the PLATFORM_ID inside the database in the V\$DATABASE view using the query below:

```
SQL> select platform_id, platform_name from v$database;
```

PLATFORM_ID	PLATFORM_NAME
10	Linux IA (32-bit)

Differences between the primary server(s) and the standby server(s) are always supported as long as the Oracle software installed on all servers is of the same Oracle Platform as defined above, is certified to run on each server, and is the same Oracle Database Release and Patch Set. Examples of such differences that are supported include the following:

[Data Guard Support for Heterogeneous Primary and Physical Standbys in Same Data Guard Configuration \(Doc ID 413484.1\)](#)

Little Endian | Linux Standby



```
SQL> SELECT platform_name, endian_format
       FROM v$transportable_platform
       WHERE endian_format='Little';
```

PLATFORM_NAME	ENDIAN_FORMAT
Apple Mac OS (x86-64)	Little
HP IA Open VMS	Little
HP Open VMS	Little
HP Tru64 UNIX	Little
Linux IA (32-bit)	Little
Linux IA (64-bit)	Little
Linux x86 64-bit	Little
Microsoft Windows IA (32-bit)	Little
Microsoft Windows IA (64-bit)	Little
Microsoft Windows x86 64-bit	Little
Solaris Operating System (x86)	Little
Solaris Operating System (x86-64)	Little



Little Endian | Hybrid Data Guard



Hybrid Data Guard to Exadata Cloud Services

Production Database on Premises and Disaster Recovery with Oracle Cloud Infrastructure Exadata Cloud Service



MIGRATION

methods

ZDM

DATA
GUARD

TRANS-
PORTABLE

DATA PUMP
GOLDENGATE



Photo by [Florian Klauer](#) on [Unsplash](#)

Transport

Basics

Endianness | The Basics



Big-endian

increasing addresses →

...	4A _h	6F _h	68 _h	6E _h	...
...	'J'	'o'	'h'	'n'	...

Little-endian

increasing addresses →

...	6E _h	68 _h	6F _h	4A _h	...
...	'n'	'h'	'o'	'J'	...

Source: <https://en.wikipedia.org/wiki/Endianness>



Endianness Migration | Cloud is Linux



```
SQL> SELECT platform_name, endian_format
       FROM v$transportable_platform
       WHERE endian_format!='Little';
```

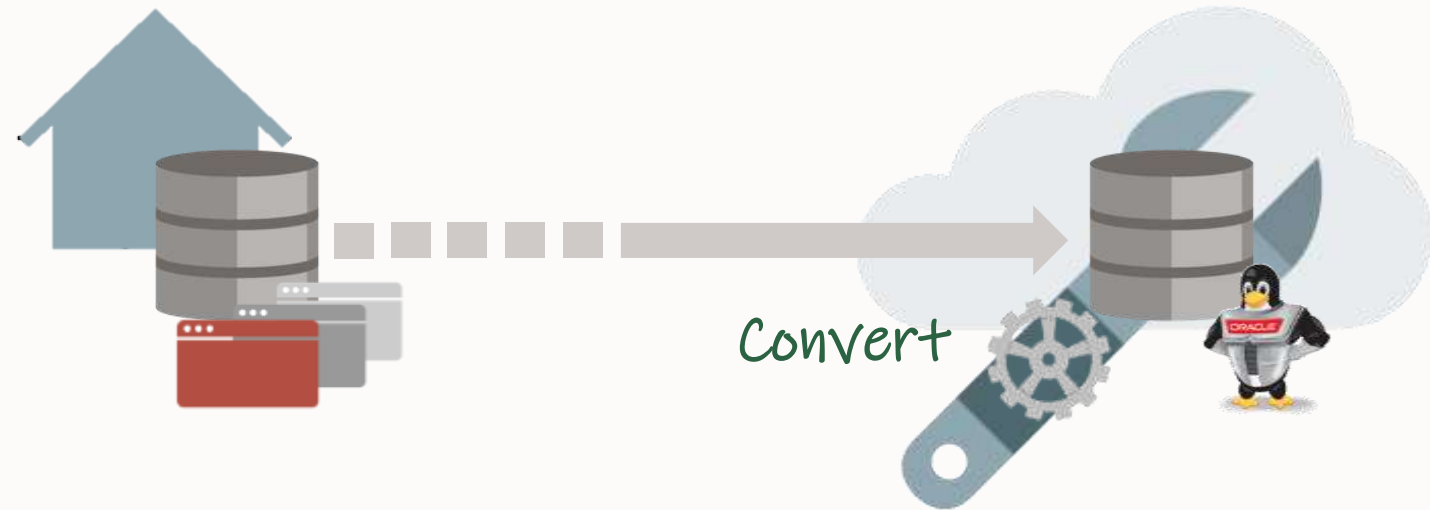
PLATFORM_NAME	ENDIAN_FORMAT
AIX-Based Systems (64-bit)	Big
Apple Mac OS	Big
HP-UX (64-bit)	Big
HP-UX IA (64-bit)	Big
IBM Power Based Linux	Big
IBM zSeries Based Linux	Big
Linux OS (S64)	Big
Solaris[tm] OE (32-bit)	Big
Solaris[tm] OE (64-bit)	Big



Endianness Migration | Transport

Big Endianness platforms

- HP-UX (64-bit)
- HP-UX IA (64-bit)
- AIX-Based Systems (64-bit)
- IBM zSeries Based Linux
- IBM Power Based Linux
- Solaris[tm] OE (32-bit)
- Solaris[tm] OE (64-bit)



Pro Tip: You can use Transportable Tablespace even for little Endian migrations

Endianness Migration | Convert



RMAN

Out-of-place conversion

2 x disk space needed

Supported in newest version of Perl scripts

DBMS_FILE_TRANSFER

In-flight conversion

1 x disk space needed

Not supported in Perl scripts version 4

Endianness Migration | Transportable Tablespace Concept

To move data we need two things:

META DATA

How is the following defined:

User/schema

Table

Columns

Indexes

Triggers

Grants

PL/SQL

...

DATA

The actual rows



Endianness Migration | Transportable Tablespace Concept

To move data we need two things:

META DATA

```
CREATE USER u1 ...  
CREATE TABLE t1 ( ...  
CREATE INDEX i1 ON t1.c1 ...  
CREATE TRIGGER trig1 ...  
GRANT SELECT ON t1 ...  
CREATE PROCEDURE p1 ...
```

Stored in
SYSTEM tablespace

DATA

```
INSERT INTO t1 (...) VALUES (...)  
INSERT INTO t1 (...) VALUES (...)  
INSERT INTO t1 (...) VALUES (...)  
INSERT INTO t1 (...) VALUES (...)
```

Stored in
user tablespaces

Endianness Migration | Transportable Tablespace Concept

SYSTEM			



Export with
Data Pump

DATA			

Copy
data files



Endianness Migration | Transportable Tablespace Concept



Tied to database version

Works **only in same** database version

Works **only in same** database architecture

Data Pump **works across:**

- Database version
- Database architecture (non-CDB / CDB)

Independent of database version

Works in **same or newer** database version

Works in **same or different** database architecture

Endianness Migration | Transportable Tablespace Concept

Transportable tablespaces works:

- To the **same or newer** database version
- For non-CDB to PDB conversion (and vice-versa)

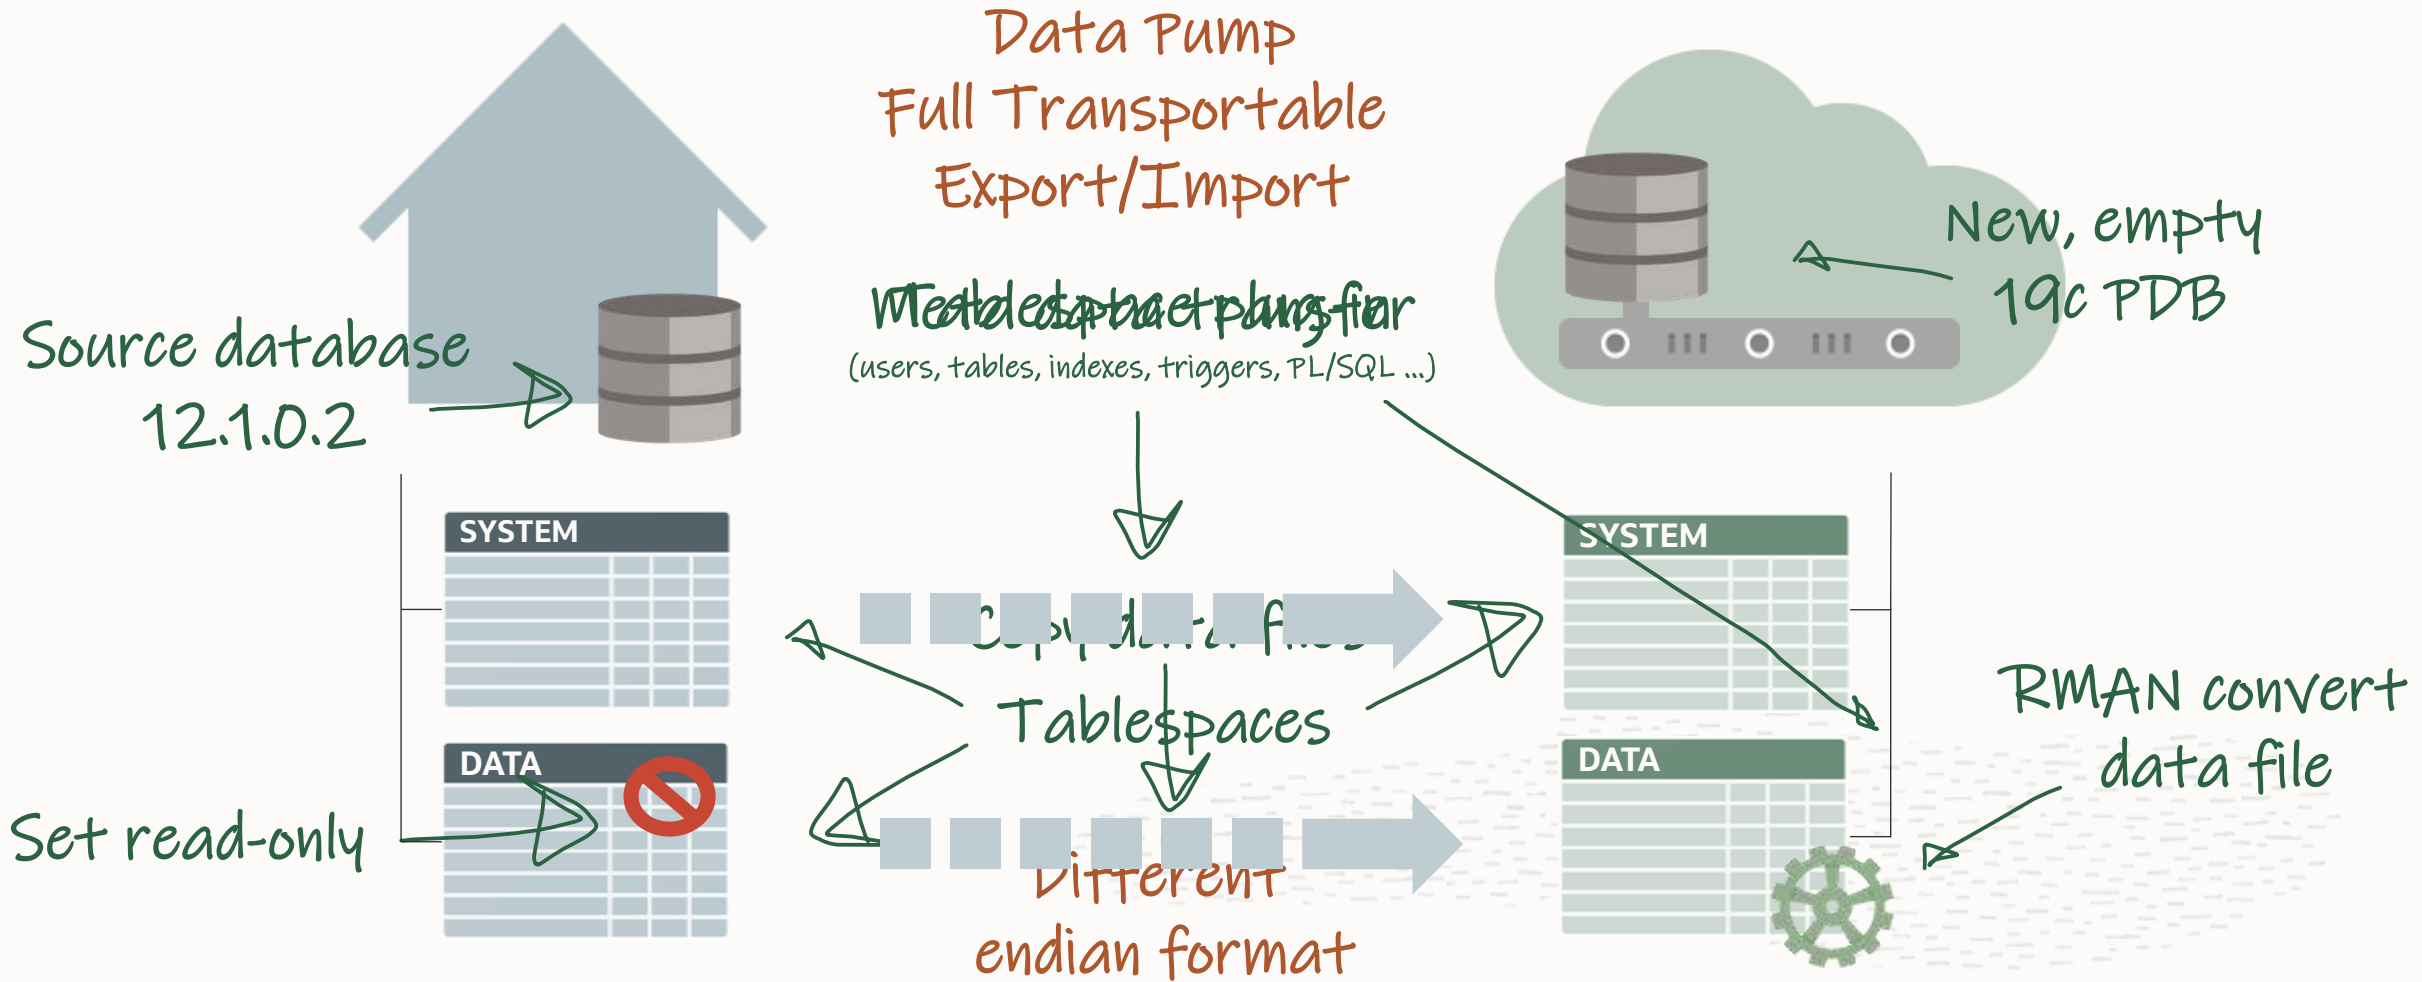


Photo by [Fahrul Azmi](#) on [Unsplash](#)

Transport

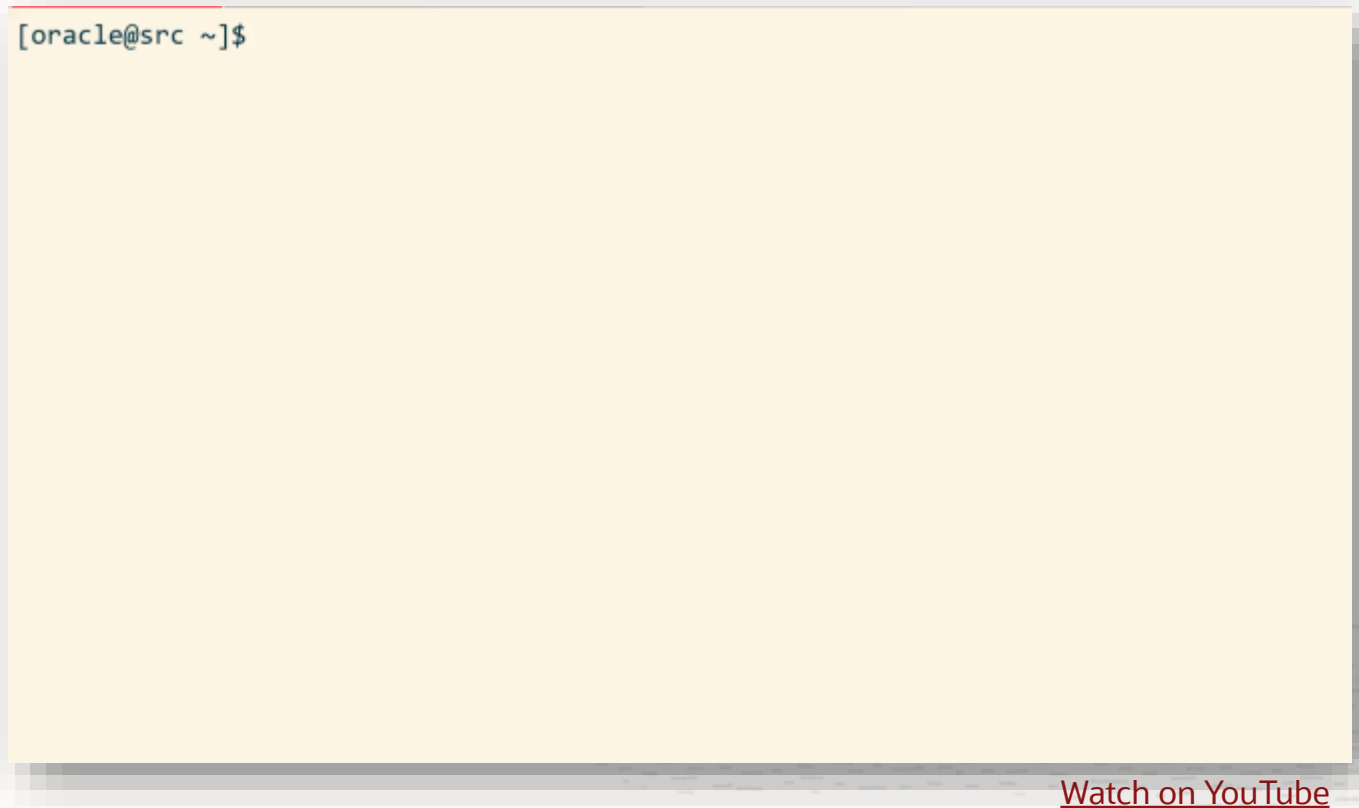
Methods

Endianness Migration | Full Transportable Export Import (FTEX)



Endianness Migration | Full Transportable Export Import (FTEX)

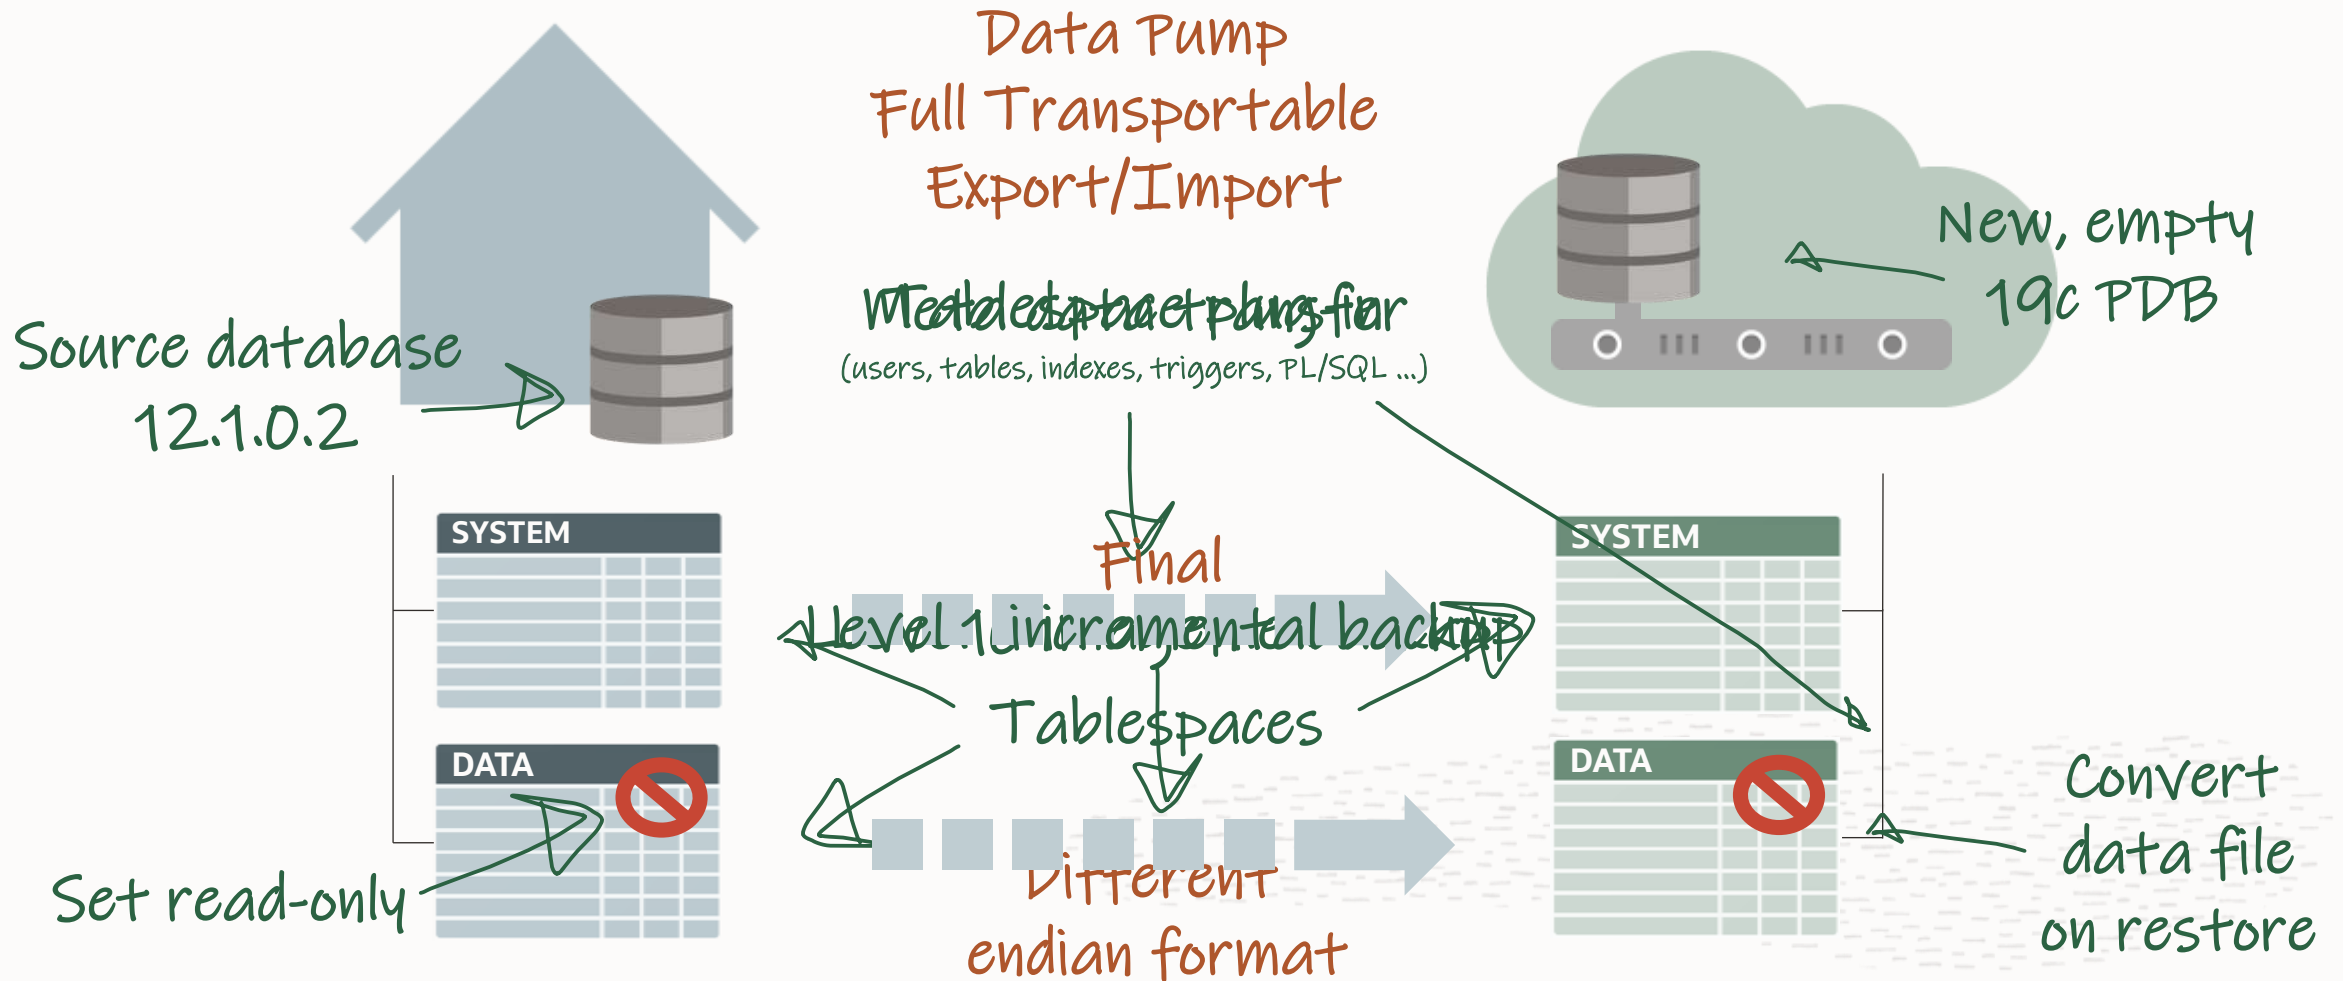
```
[oracle@src ~]$
```



[Watch on YouTube](#)



Endianness Migration | FTEX plus Incremental Backups

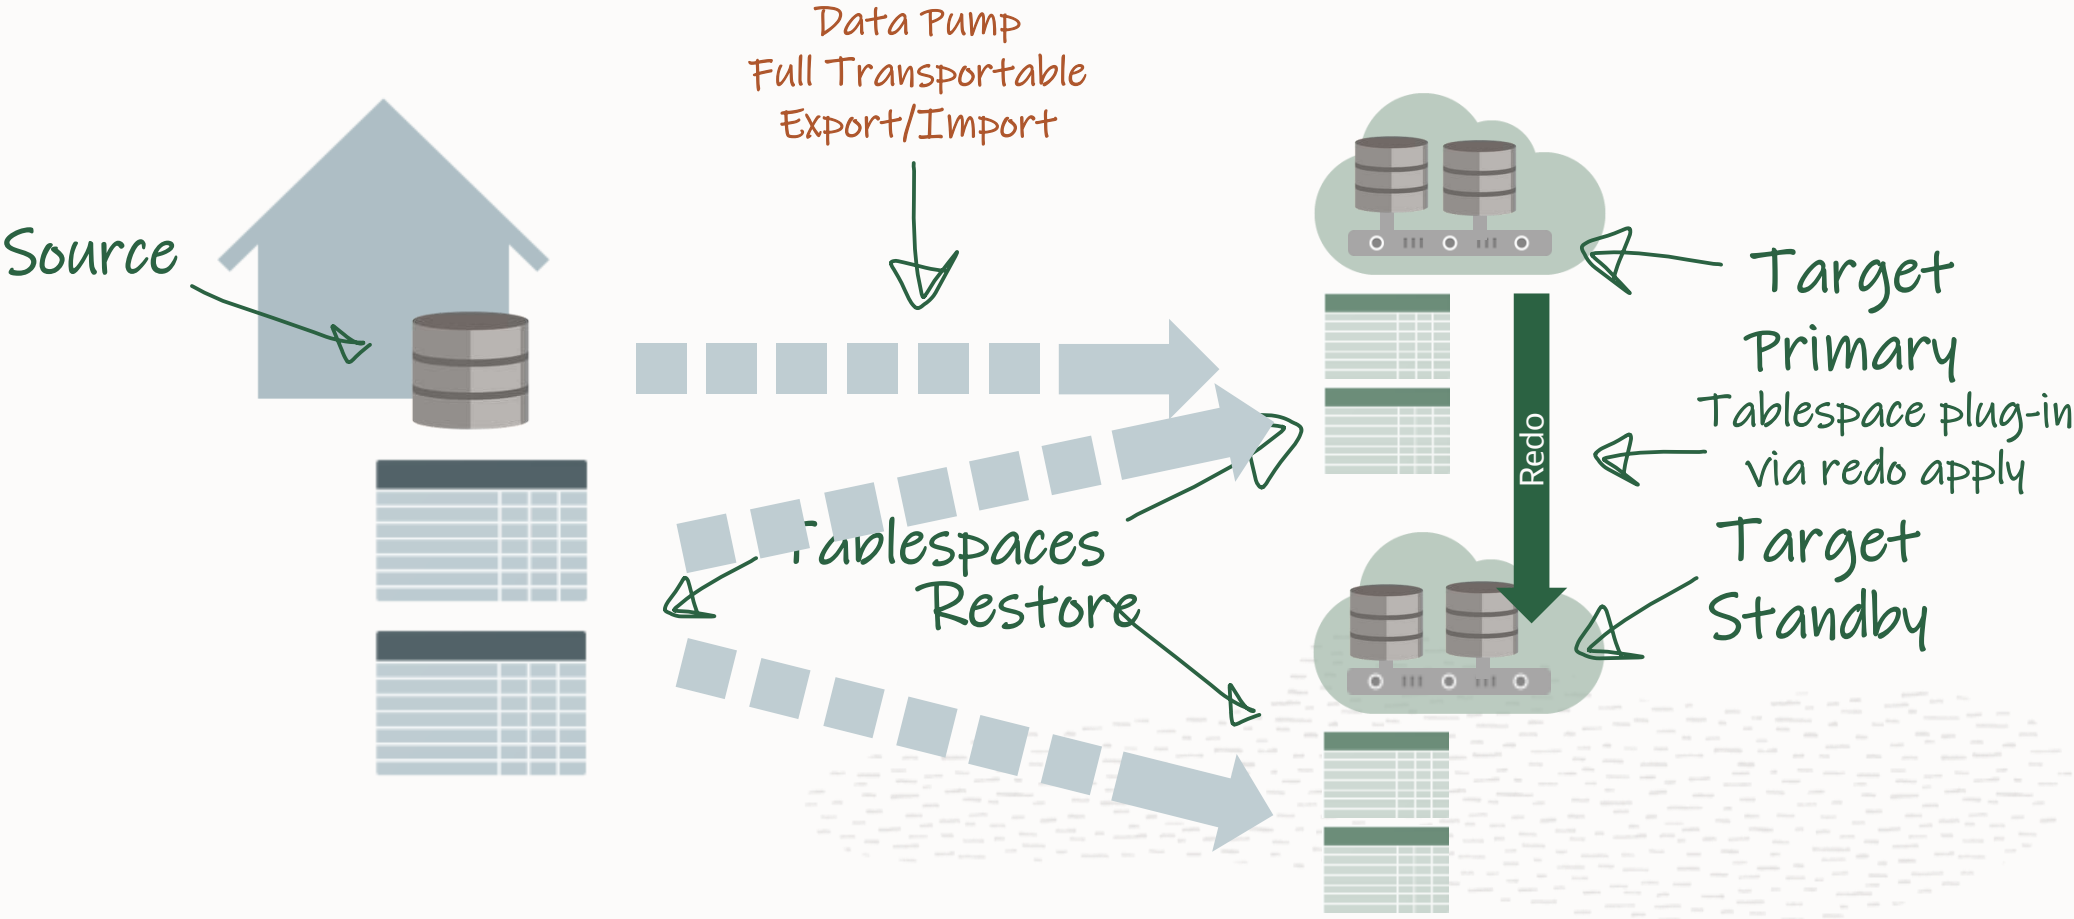


Endianness Migration | **FTEX** plus Incremental Backups

```
[oracle@src ~]$
```

[Watch on YouTube](#)

Data Guard | Transportable



Data Guard | Transportable



★ **Step by Step Process of Migrating non-CDBs and PDBs Using ASM for File Storage (Doc ID 1576755.1)**

In this Document

- [Purpose](#)
- [Scope](#)
- [Details](#)
- [Goal](#)
 - [Multitenant Plugin with OMF/ASM](#)
 - [Handling of Standby Databases during plugin.](#)
 - [12c Migrate non-CDB and Convert to a PDB using 12c RMAN Active Database Duplication](#)
 - [12c Unplug/Plug PDB to new CDB using 12c RMAN Active Database Duplication](#)
 - [Data Pump Full Transportable Using Incremental Backups to Migrate a pre-12c non-CDB to 12c PDB](#)
- [References](#)

APPLIES TO:

- Oracle Database Cloud Schema Service - Version N/A and later
- Oracle Database Exadata Cloud Machine - Version N/A and later



[Step by Step Process of Migrating non-CDBs and PDBs Using ASM for File Storage \(Doc ID 1576755.1\)](#)





Photo by [Sebastian Herrmann](#) on [Unsplash](#)

Transport

Checklist

Transportable | Starter Checklist

Database Creation

Backup / Recovery

TDE

PERL Scripts

Target database requirements

- `COMPATIBLE` must be the same or higher

Target database requirements with workarounds

- Identical character set
- Identical national character set
- Identical time zone (only with `TIMESTAMP WITH LOCAL TIME ZONE`)
- Identical time zone file version (only with `TIMESTAMP WITH TIME ZONE`)



Transportable | Starter Checklist

Database Creation

Backup / Recovery

TDE

PERL Scripts

[Blog post](#) on how to create a database in OCI with custom COMPATIBLE setting

Transportable | Starter Checklist

Database Creation

Backup / Recovery

TDE

PERL Scripts

To determine character set:

```
SQL> select * from nls_database_parameters;
```

Convert source database to Unicode with [DMU](#)

A few [Character set exceptions](#)

Transportable | Starter Checklist

Database Creation

Backup / Recovery

TDE

PERL Scripts

OCI databases are AL32UTF8

- Change in *Advanced Options*

Database 12.2 and higher

- PDB can use different character set

Recommendation

- Keep *production* CDB on AL32UTF8
- Provision *temporary* CDB with desired character set
- Create new empty PDB in *temporary* CDB
- Clone custom PDB to *production* CDB

Transportable | Starter Checklist

Database Creation

Backup / Recovery

TDE

PERL Scripts

OCI DB Systems are in UTC time zone

- Change in *Advanced Options*

Sets the OS time zone, which affects:

- SYSDATE
- SYSTIMESTAMP

[How to change the Time Zone in Oracle Database Hosted in OCI with an Example \(Doc ID 2459830.1\)](#)

Transportable | Starter Checklist

Database Creation

Backup / Recovery

TDE

PERL Scripts

To determine database time zone:

```
SQL> select dbtimezone from v$instance;
```

If source and target database time zone doesn't match

- Tables with TSLTZ are skipped
- Import using Data Pump afterwards

Database time zone is only relevant for columns of `TIMESTAMP WITH LOCAL TIME ZONE`

[Documentation](#)

Transportable | Starter Checklist

Database Creation

Backup / Recovery

TDE

PERL Scripts

OCI database time zone defaults to UTC

Change for CDB:

```
SQL> alter database cdb1 set time_zone = '+02:00';
```

PDBs can have different DB Time Zone:

```
SQL> alter pluggable database pdb1 set time_zone = '+04:00';
```

Only relevant for

- TSLTZ
- CURRENT_DATE
- CURRENT_TIMESTAMP
- LOCALTIMESTAMP

Transportable | Starter Checklist

Database Creation

Backup / Recovery

TDE

PERL Scripts

To determine database time zone file version:

```
SQL> select * from v$timezone_file;
```

If source and target database time zone file version doesn't match

- [Tables with TSTZ are skipped](#)
- Import using Data Pump afterwards

Database time zone file version is only relevant for columns of
TIMESTAMP WITH TIME ZONE

How to [create a database](#) with a non-default time zone file version

[Documentation](#)



Transportable | Starter Checklist

Database Creation

Backup / Recovery

TDE

PERL Scripts

Enable Block Change Tracking on source for incremental backups

```
SQL> SELECT status, filename FROM V$BLOCK_CHANGE_TRACKING;  
SQL> ALTER DATABASE ENABLE BLOCK CHANGE TRACKING;
```

- Conversion on destination is usually faster than on source
- PERL scripts will do the conversion
- Requires
 - Enterprise Edition (on-prem)
 - Enterprise Edition Extreme Performance (DBCS)
 - Exadata



Transportable | Starter Checklist

Database Creation
Backup / Recovery

TDE

PERL Scripts

TDE Encryption

- Not supported
- Only for same-Endianness migration

Transportable | Starter Checklist

Database Creation
Backup / Recovery
TDE

PERL Scripts

RMAN Incremental Backups

- [MOS Note: 2471245.1](#)
V4 PERL Scripts to reduce Transportable Tablespace Downtime using Cross Platform Incremental Backup
- Source: 10.2.0.3 or newer
- Target: 11.2.0.4 or newer

Transportable | General Best Practices



- Practice, practice, practice
 - Start on small database
 - Prove it works on production-size database
- Automate
 - To ensure consistency and avoid human error
- Save all logs and output
 - Data Pump, RMAN
- Clean-up procedure
 - In case of failure and rollback
 - To repeat tests
 - Ensure that source database gets offlined afterwards





Transport

Advanced

FTEX | Recommendations



Exclude SYS (at least in a PDB)

```
EXCLUDE=SYS_USER
```

Create TEMP tablespaces in advance, and exclude TEMP tablespaces

```
EXCLUDE=TABLESPACE:"IN('TEMP')"
```

Exclude statistics

```
EXCLUDE=TABLE_STATISTICS,INDEX_STATISTICS
```

Exclude Spatial users (removed in 19c)

```
EXCLUDE=SCHEMA:"IN('SPATIAL_CSW_ADMIN_USR','SPATIAL_WFS_ADMIN_USR')"
```



FTEX | Known Issues



Ignorable error: Package is removed in 12.2

```
Processing object type DATABASE_EXPORT/SYSTEM_PROCOBJECT/POST_SYSTEM_ACTIONS/PROCACT_SYSTEM
ORA-39083: Object type PROCACT_SYSTEM failed to create with error:ORA-04042: procedure, function,
package, or package body does not exist
```

Failing sql is:

```
BEGIN
SYS.DBMS_UTILITY.EXEC_DDL_STATEMENT('GRANT EXECUTE ON DBMS_DEFER_SYS TO "DBA"');COMMIT; END;
```

[ORA-39083 And ORA-04042 Errors On DBMS_DEFER_SYS When Importing Into 12.2 Database \(Doc ID 2335846.1\)](#)

FTEX | Known Issues



Ignorable error: Multimedia desupported in 19c, but code is still there

```
Processing object type DATABASE_EXPORT/NORMAL_OPTIONS/TABLE  
ORA-39342: Internal error - failed to import internal objects tagged with ORDIM due to ORA-00955:  
name is already used by an existing object.
```

Transportable | Further Reading



- [Master Note for Transportable Tablespaces \(TTS\) -- Common Questions and Issues \(Doc ID 1166564.1\)](#)
- [Transportable Tablespace \(TTS\) Restrictions and Limitations: Details, Reference, and Version Where Applicable \(Doc ID 1454872.1\)](#)
- [V4 PERL Scripts to reduce Transportable Tablespace Downtime using Cross Platform Incremental Backup \(Doc ID 2471245.1\)](#)
- [Known Issues for Cross Platform Transportable Tablespaces XTTS \(Doc ID 2311677.1\)](#)
- [Cross Platform Database Migration using ZDLRA \(Doc ID 2460552.1\)](#)
- [11G – Reduce Transportable Tablespace Downtime using Cross Platform Incremental Backup \(Doc ID 1389592.1\)](#)
- [12C – Reduce Transportable Tablespace Downtime using Cross Platform Incremental Backup \(Doc ID 2005729.1\)](#)
- Blog post: [What Is a Self-contained Transportable Tablespace Set](#)



MIGRATION

methods

ZDM

DATA
GUARD

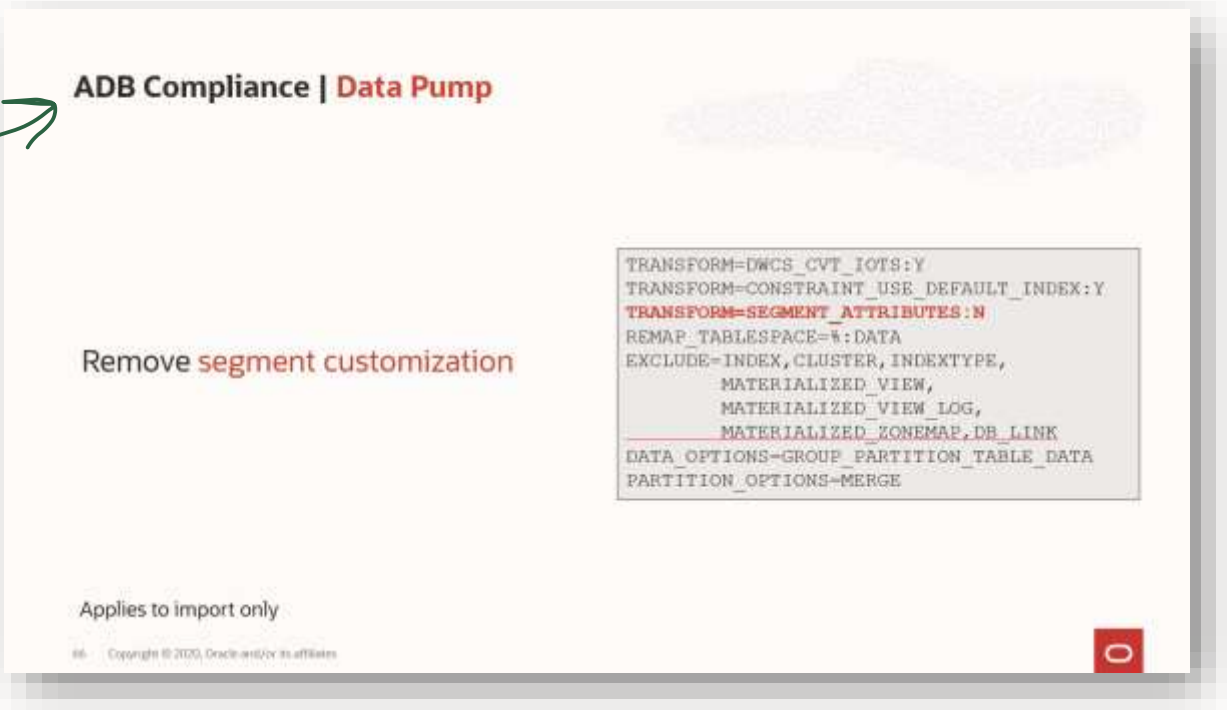
TRANS-
PORTABLE

DATA PUMP
GOLDENGATE

Data Pump

Use information from "Autonomous Database Deep Dive"

Except "ADB Compliance"



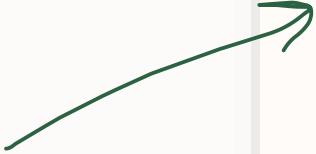
ADB Compliance | Data Pump

Remove segment customization

```
TRANSFORM=DWCS_CVT_IOTS:Y
TRANSFORM=CONSTRAINT_USE_DEFAULT_INDEX:Y
TRANSFORM=SEGMENT_ATTRIBUTES:N
REMAP_TABLESPACE=%:DATA
EXCLUDE=INDEX, CLUSTER, INDEXTYPE,
        MATERIALIZED_VIEW,
        MATERIALIZED_VIEW_LOG,
        MATERIALIZED_ZONEMAP, DB_LINK
DATA_OPTIONS=GROUP_PARTITION_TABLE_DATA
PARTITION_OPTIONS=MERGE
```

Applies to import only

66 Copyright © 2020, Oracle and/or its affiliates



Data Pump | SecureFiles



” *SecureFiles is the default storage mechanism for LOBs starting with Oracle Database 12c, and Oracle **strongly recommends SecureFiles** for storing and managing LOBs, rather than BasicFiles. BasicFiles will be deprecated in a future release.*

[Database SecureFiles and Large Objects Developer's Guide](#)

Always transform LOBs to SecureFiles LOBs

```
$ impdp ... TRANSFORM=LOB_STORAGE:SECUREFILE
```

Data Pump | SecureFiles



Importing as BasicFiles

```
10-OCT-20 21:43:21.848: W-3 . . imported "SCHEMA"."TABLE"      31.83 GB  681025 rows in 804 seconds using direct_path
```

Importing as SecureFiles

```
15-OCT-20 18:16:48.663: W-13 . . imported "SCHEMA"."TABLES"  31.83 GB  681025 rows in 261 seconds using external_table
```



Data Pump | ZDM



☆ (OCI) MV2OCI: move data to Oracle Cloud Database in "one-click" (Doc ID 2514026.1)

In the

DETAILS

[Abs](#)

[Hist](#)

M

MV



Zero Downtime Migration

Oracle Zero Downtime Migration

Simple, Automated, One button approach solution for moving your Oracle Databases into the Oracle Cloud

Download



Data Pump | MV2OCI




"One button approach"

Uses Data Pump (schema mode)

Documentation: [Doc ID 2514026.1](#)

Support import over DB link (`--netlink`)

Runs on Linux / Solaris



Pro Tip: Check log file to find out how MV2OCI uses Data Pump




Data Pump | **MV2OCI**



Dump files are moved directly to target host
- not staged on Object Store

Requires SSH and SQL*Net connectivity


Java executable in PATH



Pro Tip: For troubleshooting check
`/opt/mv2oci/out/log`



Data Pump | MV20CI



[Watch on YouTube](#)



Data Pump | MV2OCI



Enable Data Pump metrics (METRICS=Y and LOGTIME=ALL)

```
$ mv2oci.bin auto --conf mv2oci_metal.cfg --nosudo -dpdebug
```

Exclude statistics

```
EXTRA_EXPDP=EXCLUDE=STATISTICS
```

Enable Data Pump compression (license required)

```
EXTRA_EXPDP=COMPRESSION=ALL COMPRESSION_ALGORITHM=MEDIUM
```

Transform LOBs to SecureFiles

```
EXTRA_IMPDP=TRANSFORM=LOB_STORAGE:SECUREFILE
```

Data Pump | ZDM

ZDM automatically sets recommended defaults for Data Pump

- METRICS=Y
- LOGTIME=ALL
- COMPRESSION=ALL
- ENCRYPTION=ALL

Oracle Data Pump Defaults for Zero Downtime Migration

Note: use of compression and encryption by ZDM does not require additional licenses for the Advanced Security Option or Advanced Compression Option

GoldenGate



ORACLE® | GoldenGate

Use ZDM: Logical Online Migration

Alternatively: Use OCI GoldenGate

ORACLE



Performance Testing in the Cloud

Oracle Database 19c

Roy Swonger, Mike Dietrich & Daniel Overby Hansen

Oracle Database Upgrade, Utilities, Cloud Migration and Patching



The Usual Testing Challenges

Testing | Usual Challenges

- Only 10% of real data used
- Artificially created data sets
- Outdated data
- Tests done on a laptop
- No testing tools used
- No stale statistics refreshed

”

Testing?? Maybe next time ...



“ Help me - I have an upgrade problem ...”

In

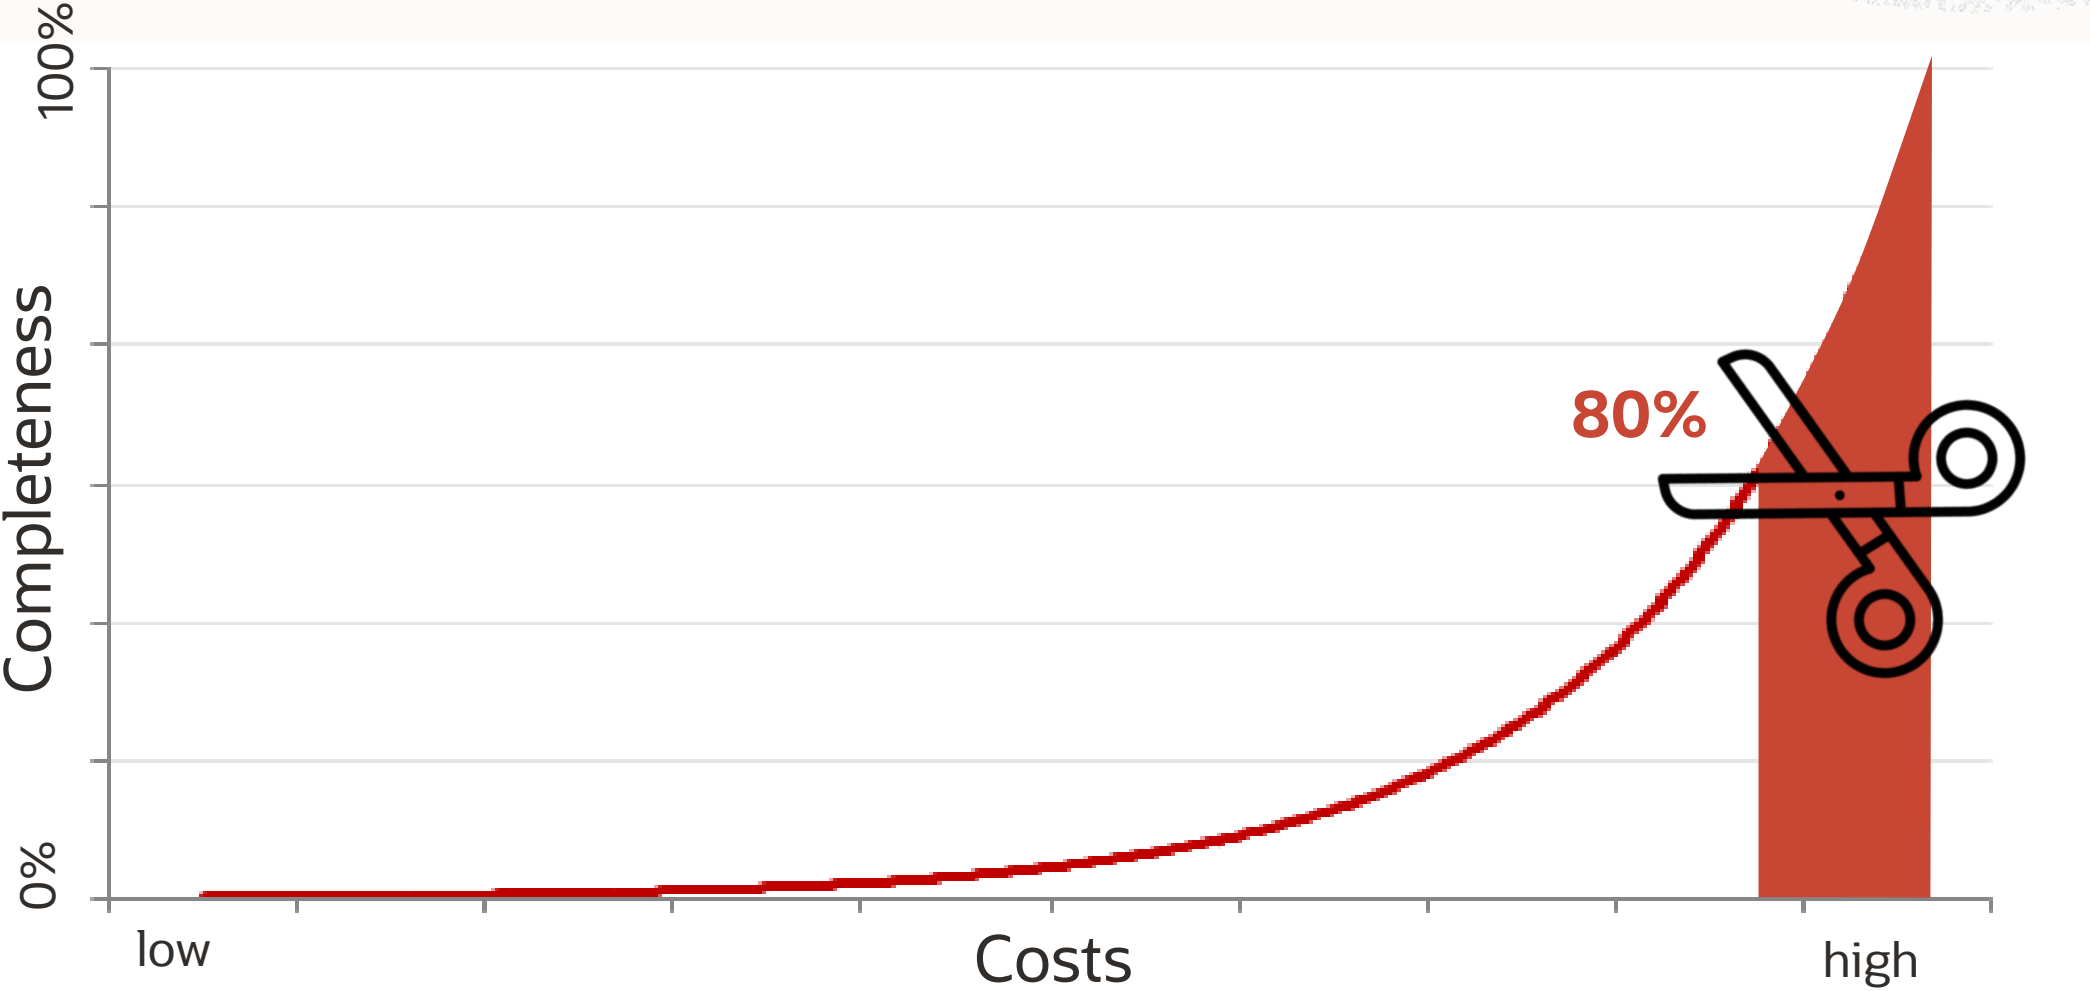
95%


of all cases, "upgrade problem" in **reality** is a **performance issue after upgrade**. Or not database related.

There is exactly one way to mitigate the risk.

TESTING!

Testing | Pareto Principle – 80/20 Rule





But what if
you get only old and outdated test hardware?

And no proper testing tools?



your key to

Successful Database Upgrades / Migrations

Step 1

Download and
install Oracle 19c

eDelivery.oracle.com

Step 2

Download and
install newest RU

MOS Note: 2118136.2

Step 3

Download and use
AutoUpgrade

MOS Note: 2485457.1

Step 4

Performance Stability
with SPM, STA and RAT





CAUTION

The following recommendations apply to creation of test environments
They do not necessarily apply to production environments

Challenge #1

Hardware

Why Use Cloud | License Included



Did you know?

When you provision a DBCS in OCI the license **includes** some of the Options and Management Packs?



Oracle Database

Offerings

and licenses included

	DBCS EE	DBCS EE-HP	DBCS EE-EP	ExaCS
Diagnostics Pack	x	x	x	x
Tuning Pack	x	x	x	x
Real Application Testing	x	x	x	x
Advanced Compression		x	x	x

Why Use Cloud | Use Cases



Diagnostics Pack

- Database Replay Compare Period Report
- AWR
- ASH
- Performance Hub



Tuning Pack

- SQL Access Advisor
- SQL Tuning Advisor
- SQL Profiles
- Real-time SQL and PL/SQL Monitoring

Why Use Cloud | Use Cases



Real Application Testing

- Database Replay
(if you want to capture on-prem, license needed there as well)
- SQL Performance Analyzer



Advanced Compression

- Compress your Data Pump export
(importing a compressed dump file is not licensed)
- Compress your backups of your OCI database using RMAN compression

Many more Options and Packs are included

Creating Database | **Multitenant**

Multitenant is the default in OCI

What if ... you are not there yet



Database Cloud Service | Virtual Machines



Entry-level, provision with GI or LVM (fast-provision)

Restrictions:

- Only **one CDB** - the pre-created one
- Can't install another Oracle Home
- COMPATIBLE is always default - unless
- Supports RAC
- Max. storage 40 TB

To create non-CDB

- Migrate your non-CDB using ZDM
- Restore a non-CDB on top of the pre-created CDB
 - `DB_NAME` and `DB_UNIQUE_NAME` must match exactly



Database Cloud Service | Bare Metal



Mid-level, provision with GI

Restrictions:

- As **many databases** as you want
- DATA disk group, max. 16 TB
- Only one database edition
- Only one database per Oracle Home
- No RAC

To create non-CDB

- It is not possible to create a non-CDB via the console - use [dbcli](#)
- Before using `dbcli` it is a good idea to update it: `cliadm update-dbcli`



Database Cloud Service | Bare Metal



```
[root]$ dbcli create-database \  
    --dbname DOH19 \  
    --databaseUniqueName DOH191 \  
    --no-cdb \  
    --dbconsole \  
    --adminpassword "<secret-password>" \  
    --dbshape odb4 \  
    --dbtype SI \  
    --version 19.10.0.0
```

Job details

```
-----  
                ID: 6d87fbf0-2b03-417c-aec9-09b69b0d6222  
Description: Database service creation with db name: DOH19  
    Status: Created  
    Created: May 9, 2021 8:33:11 AM UTC  
Progress: 0%  
Message:
```



Database Cloud Service | Bare Metal



```
[root]$ dbcli describe-job -i 6d87fbf0-2b03-417c-aec9-09b69b0d6222

Job details
-----
          ID: 6d87fbf0-2b03-417c-aec9-09b69b0d6222
Description: Database service creation with db name: DOH19
      Status: Running
      Created: May 9, 2021 8:33:11 AM UTC
      Progress: 0%
      Message:

Task Name                Start Time                End Time                Status
-----
Pre Database Creation Tasks May 9, 2021 8:33:11 AM UTC May 9, 2021 8:33:13 AM UTC Success
Database Home Creation    May 9, 2021 8:33:13 AM UTC May 9, 2021 8:33:13 AM UTC Running
```



Database Cloud Service | Exadata



World's best database machine, provision with GI

Restrictions:

- As **many databases** as you want
- Non-CDB are supported only for 12.1.0.2 and 19c
- Supports RAC
- Max. storage up to 598 TB
 - Depend on Exadata version and configuration
 - Additional storage can be allocated via additional storage servers

To create non-CDB

- It is not possible to create a non-CDB via the console - use `dbaasapi`
- [Creating non-CDB databases on the Service in OCI \(Doc ID 2528257.1\)](#)





Challenge #2

Tools

Remember?



NEW Episode 3
Performance Stability, Tips and Tricks and Underscores
120 minutes – Mar 4, 2021



Performance Stability Prescription

1.
Collect

3.
Analyze

5.
Manage

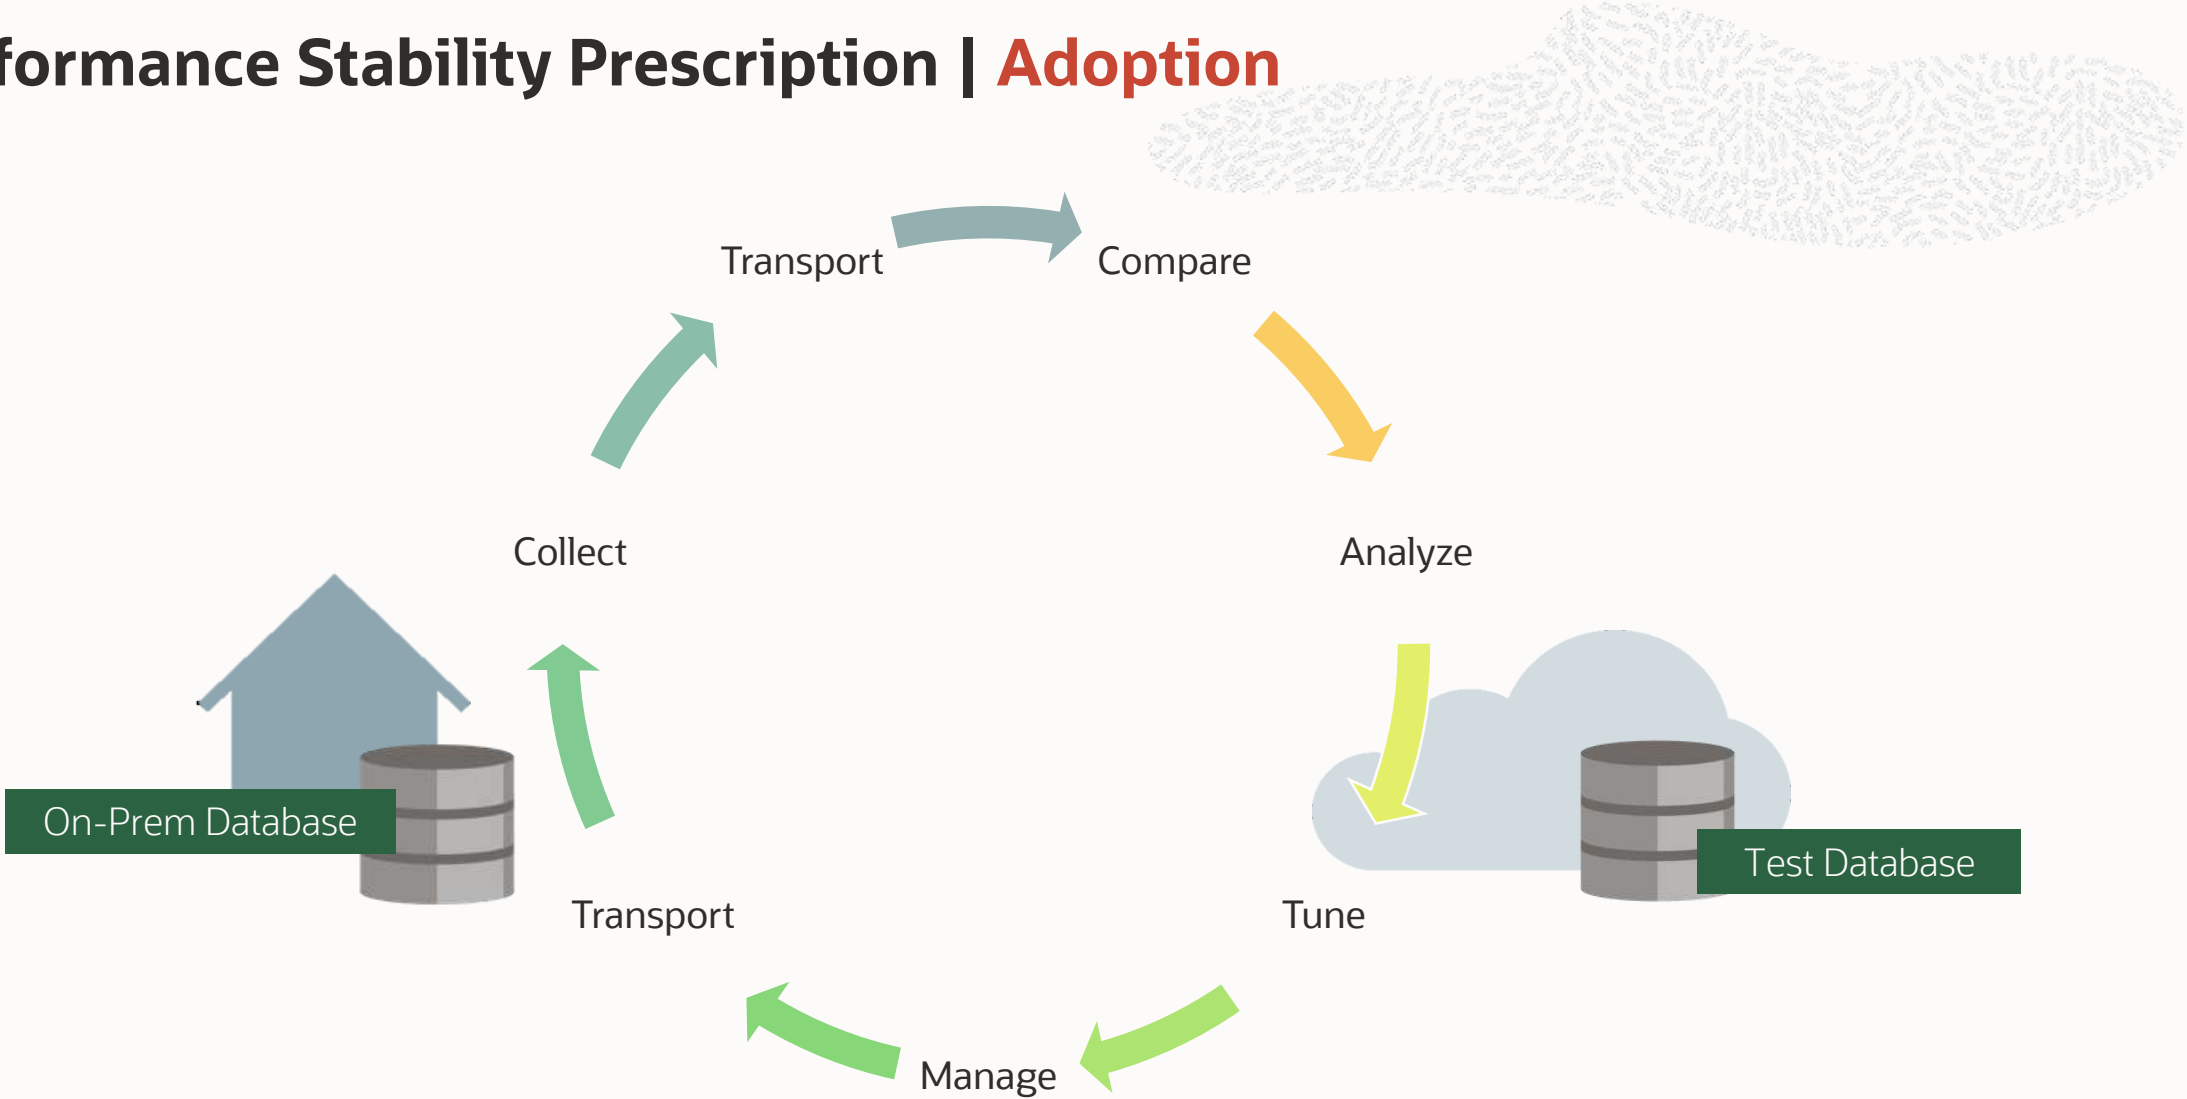
2.
Compare

4.
Tune

6.
Test



Performance Stability Prescription | Adoption



Performance Stability Prescription | **Toolbox**

AWR

Cursor Cache

SQL Tuning Set

Staging Tables

SQL Plan Management

SQL Tuning Advisor

SQL Performance Analyzer

Database Capture and Replay



Challenge #3

Migration

Migration | Move to the Cloud

<https://dohdatabase.com/webinars>



Migrate | Overview

When migrating for test

- Don't worry about downtime
- Make it simple
- Avoid affecting production database



Migrate | Data Pump

- Run at off-peak hours
- Use parallel at off-peak hours
- Remember to make entire export consistent - `flashback_time=systemtimestamp`
- Use Data Pump compression (license required)
- else use OS utilities for compressing
- Export from a snapshot standby database
- Database Replay - recover standby to desired SCN, then run Data Pump export

Migrate | Data Pump From Specific SCN

Export from standby database to:

- Offload production database
- Avoid ORA-01555: snapshot too old

1. Stop redo apply on standby database

```
DGMGRL> edit database 'stdby1' set state='apply-off';
```

2. Finish workload on primary database, e.g. prepare for Database Capture

3. Recover standby database to desired SCN

```
SQL> alter database recover managed standby database until change 16489354;
```

Migrate | Data Pump From Specific SCN

4. Convert to snapshot standby. It opens the database in *read write* mode which is required by Data Pump

```
DGMGRL> convert database 'stdby1' to snapshot standby;
```

5. Export your data with Data Pump. It is not necessary to use `flashback_scn` or `flashback_time` because no one is using the system

```
$ expdp system schemas=SH parallel=8 directory=mydir dumpfile=sh%U.dmp
```

6. Convert your standby back into a physical standby database

```
DGMGRL> convert database 'stdby1' to physical standby;
```

Pro tip: Snapshot Standby is part of Enterprise Edition



Migrate | Existing Backups

- Use existing backups
- If backup sets are not compressed, consider using OS utilities
- No effect on production database
- Allows you to test your recovery protocols
- Cloud is an attractive Disaster Recovery strategy - give it a try

Migrate | Multipart Uploads

Recommended for files larger than 100 MB

Use [OCI CLI](#)

```
oci os object put \  
  --namespace ... -bn ... --file ... --name ... \  
  --part-size 1024 \  
  --parallel-upload-count 4
```

Max part size is 50 GB

OCI CLI [installation guide](#)

Migrate | Bulk Uploads

Recommended for **many** files

Use [OCI CLI](#)

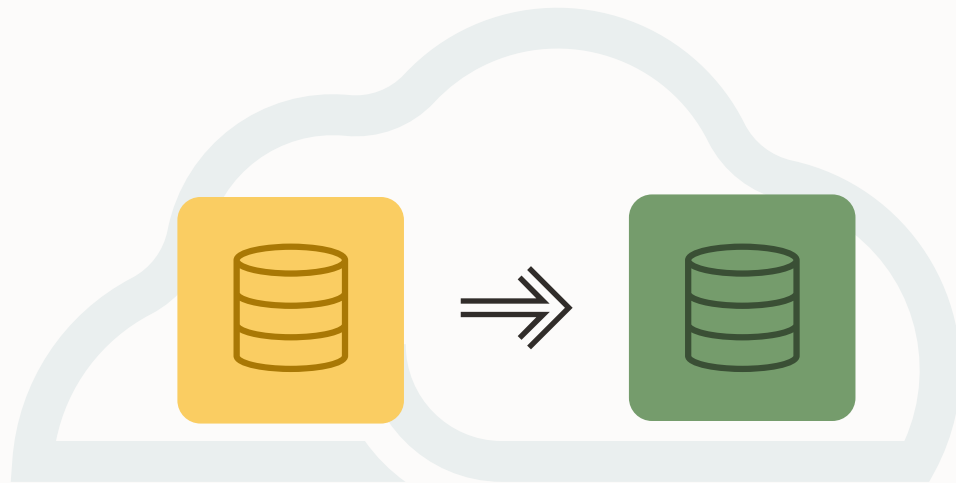
```
oci os object bulk-upload \  
  -ns ... -bn ... --src-dir ...
```

Does multipart and parallel uploads automatically

Optionally,

- Finetune uploads using `--parallel-upload-count` and `--part-size`
- Prefix all file names with `--object-prefix`
- Include or exclude files selectively using patterns and `--include` and `--exclude`

Upgrade | Overview



Upgrading in the Cloud

Upgrade | Overview



Cloud Tooling



AutoUpgrade



Data Pump



Upgrade | Overview




Cloud Tooling

- One-Button approach
- Requires fully supported system
- Slower than other options
- Not an option you can use for production upgrade



Upgrade | Overview




AutoUpgrade

- Breaks the cloud tooling
- Familiar option
- Likely what you will use for production upgrade
- Customizable
- Test your database during upgrade



Upgrade | Overview




Data Pump


- Migrate directly into new release
- Familiar option
- Probably not what you will use for production upgrade
- Customizable



Migrate | Your Data



BRING ALL YOUR DATA



Migrate | Storage Cap



ORACLE
Database
Cloud Service
Virtual
Machines

40 TB

ORACLE
Database
Cloud Service
Bare
Metal

16 TB

ORACLE
Database
Cloud Service
Exadata
Cloud Service

598 TB





ALL DATA

- Migrate structure and data
- Allows full testing - including the application
- Requires more resources
- Preferred method

STRUCTURE ONLY

- Migrate structure only
- Allows only partial testing
- Can use the smallest possible shape
- Alternative method



Migrate | Structure Only

Data Pump can export the structure but skip the actual data

```
$ expdp .... schemas=SH content=metadata_only
```

On import

- Statistics are maintained
If the source table had 100 rows, the statistics will reflect this
- Statistics are locked

Migrate | Structure Only

Alternatively, truncate tables and shrink data files

```
SQL> truncate table sh.sales;  
SQL> truncate table sh.costs cascade;  
...  
SQL> alter database datafile 10 resize 5m;
```

- Tablespace must be completely empty to resize the data files to a minimum
- Optionally, use the `cascade` keyword to truncate parent and child tables
- Statistics are cleared - import production stats afterwards
- **Don't do this in production!**



Migrate | **Subsetting**

Why not Data Subsetting? It is an option - but data subsetting is **hard!**

- Which data to remove?
 - By age
 - Random
- Integrity
 - Referential integrity
 - Logical integrity

Pro tip: Explore the options in [Data Masking and Subsetting Guide](#)



Migrate | **Subsetting**

Perhaps some data isn't needed for testing

- Audit
- Historical
- Logging
- Temporary

Data Pump can exclude tables

```
$ expdp .... schemas=SH exclude=table:in('COSTS')
```

If you need the tables, but want them empty:

```
$ expdp .... schemas=SH exclude=table:in('COSTS')  
$ expdp .... schemas=SH include=table:in('COSTS') content=metadata_only
```

Migrate | **Subsetting**

During restore or duplicate RMAN can exclude tablespaces

```
RMAN> restore database skip tablespace audit, historical, logging ...  
  
RMAN> duplicate ... skip tablespace audit, historical, logging ...
```

After restore/duplicate

- Tablespaces are offline
- Any tables stored in the tablespace are still defined in data dictionary
- Any tables stored in the tablespace can't be accessed
- They are **not** to be considered as empty tables

Challenge #4

Setup



**MAKE
CREATE
SLEEP
REPEAT**

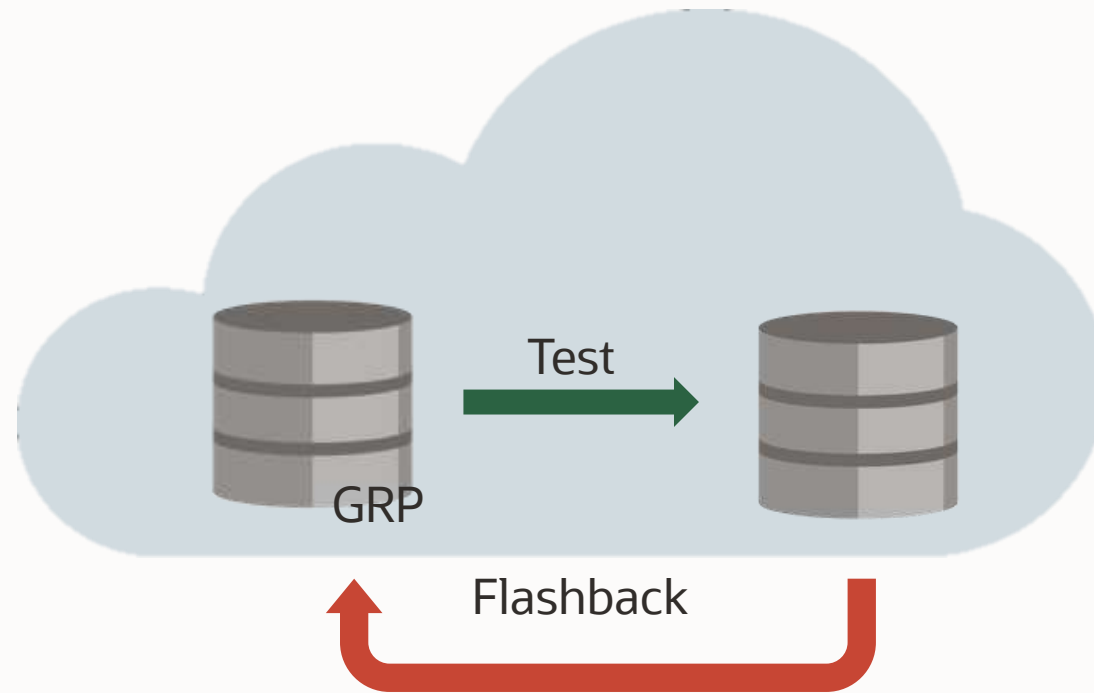
Photo by [Nick Fewings](#) on [Unsplash](#)

All tests need to be repeatable

Guaranteed Restore Points
Snapshot Standby

Setup | Guaranteed Restore Point


Set a Guaranteed Restore Point (GRP)
Flashback to the GRP



Setup | Guaranteed Restore Point

Set a Guaranteed Restore Point (GRP)

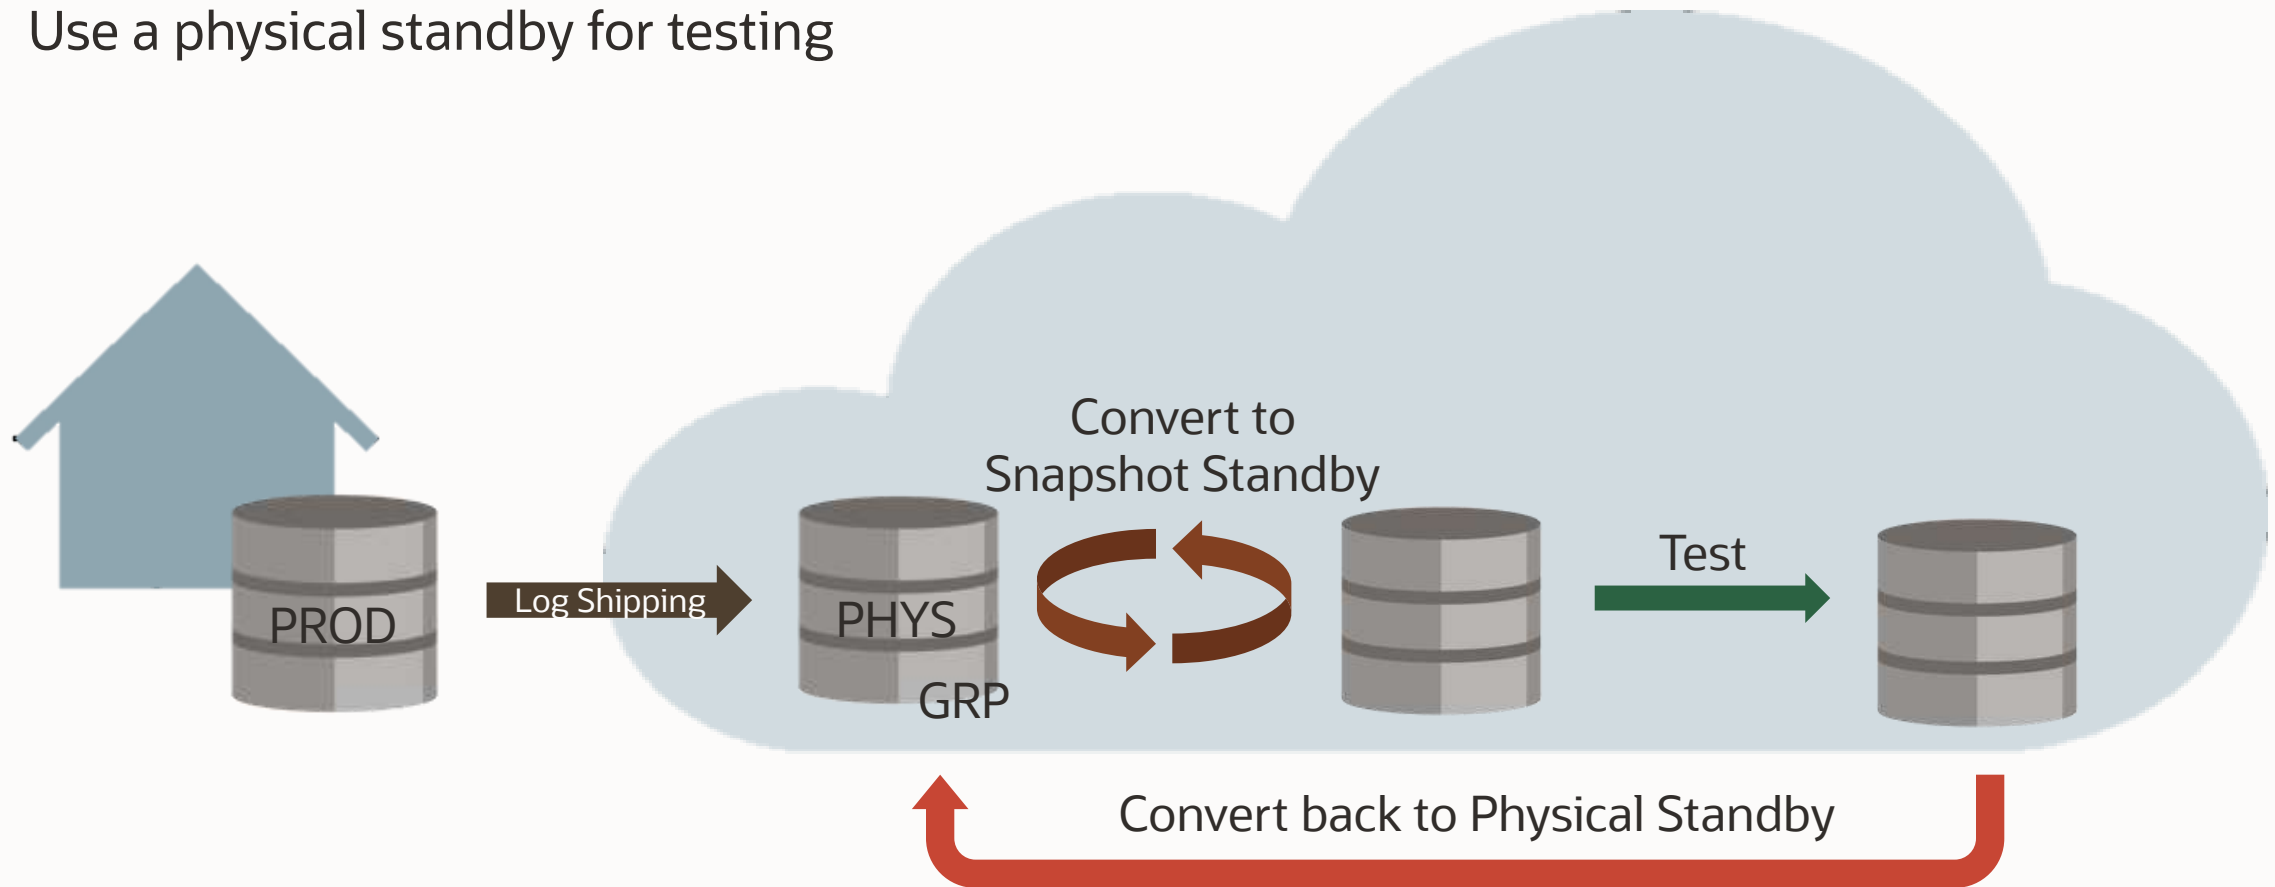
Flashback to the GRP

Before Test	After Test
<pre>CREATE RESTORE POINT grpt GUARANTEE FLASHBACK DATABASE;</pre>	
	
	<pre>SHUTDOWN IMMEDIATE</pre>
	<pre>STARTUP MOUNT;</pre>
	<pre>FLASHBACK DATABASE TO RESTORE POINT grpt;</pre>
	<pre>SHUTDOWN IMMEDIATE</pre>
<pre>STARTUP MOUNT;</pre>	
<pre>ALTER DATABASE OPEN RESETLOGS;</pre>	



Setup | Snapshot Standby

Use a physical standby for testing



Challenge #5

Statistics



Photo by [Veri Ivanova](#) on [Unsplash](#)

Statistics

Things to know

Statistics | Overview



Which statistics matter for our purpose?

- **System statistics**
 - Measure performance of system CPU and storage subsystem
- **Schema statistics**
 - User object statistics per schema
- **Database statistics?**
 - Includes schema stats, dictionary stats, fixed objects stats, system stats



System Statistics | Overview



”

The system statistics describe hardware characteristics such as I/O and CPU performance and utilization. System statistics enable the query optimizer to more accurately estimate I/O and CPU costs when choosing execution plans.

[Database 19c SQL Tuning Guide, chapter 10](#)

That **sounds** like a good idea

System Statistics | Recommendation



”

... in most cases you should *use the defaults* and not gather system statistics.

Databases supporting a *pure data warehouse workload* on an *Oracle Exadata Database Machine* can benefit from system statistics gathered using the *EXADATA* option

... if the workload is *mixed* or you are not in a position to test the effect of using *EXADATA* system statistics, then *stick to the defaults* even on this platform.

[Nigel Bayliss, Optimizer blog](#)

System Statistics | Reference

To delete system statistics (and revert to defaults)

```
SQL> EXEC DBMS_STATS.DELETE_SYSTEM_STATS
```

References:

- [Optimizer blog, Should You Gather System Statistics?](#)
- [SQL Tuning Guide, System Statistics](#)
- [SQL Tuning Guide, Guidelines for Gathering Optimizer Statistics Manually](#)
- [Database Performance Tuning Guide, Session and System Statistics](#)



Photo by [Clay Banks](#) on [Unsplash](#)

Statistics

Transporting

Transporting Statistics | Overview

”

When you transport optimizer statistics between databases, you must use `DBMS_STATS` to copy the statistics to and from a staging table, and tools to make the table contents accessible to the destination database.

[Database 19c SQL Tuning Guide, chapter 17](#)

You can transport the following statistics

- Schema
- Table
- *Database (rare)*
- *Dictionary and fixed objects (rare)*

Pro tip: You can read more about transporting statistics in the [SQL Tuning Guide](#)

Transporting Statistics | Use Cases



Restore complete database



- Best case
- All stats including COL_USAGE\$ are included

Data Pump

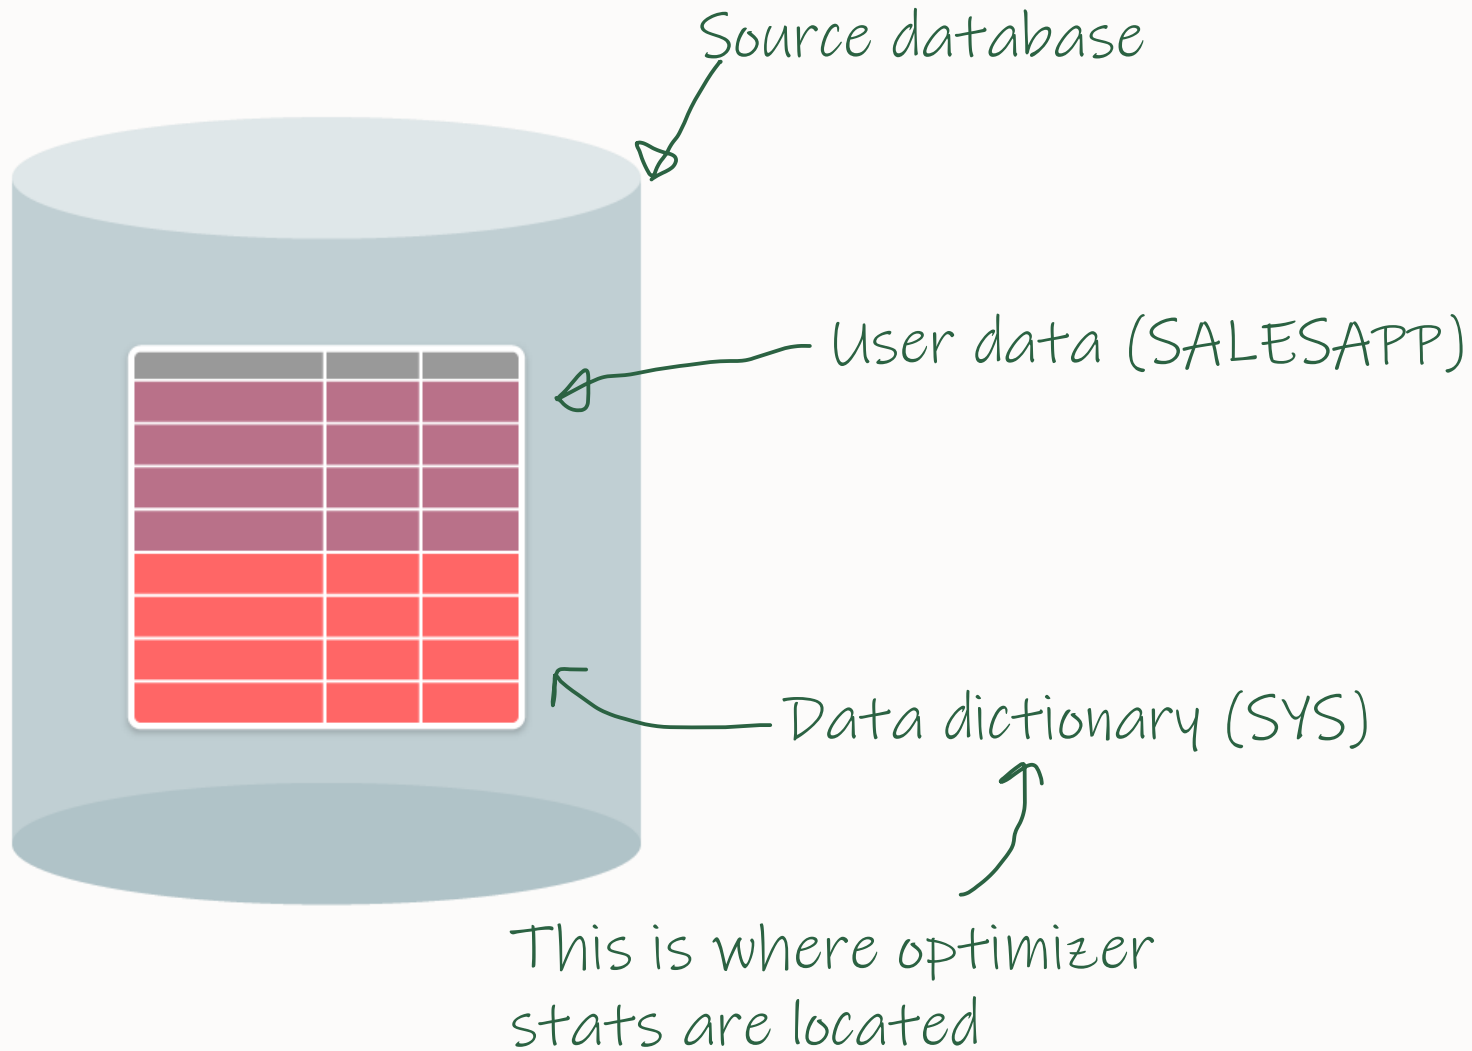
- Full database export
- Subset or schema only
- METADATA_ONLY



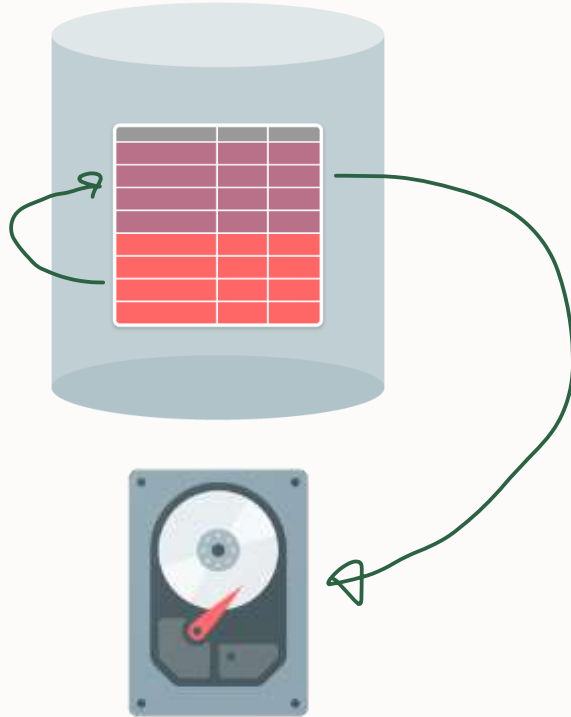
- Statistics need to be transported
- Especially COL_USAGE\$ will be missing
 - Consider METHOD_OPT=SKEWONLY



Transporting Statistics | Workflow



Transporting System Statistics | Workflow



Create staging table

```
SQL> EXEC DBMS_STATS.CREATE_STAT_TABLE (  
    ownname => 'SYSTEM',  
    statab => 'SYS_STATS_STG');
```

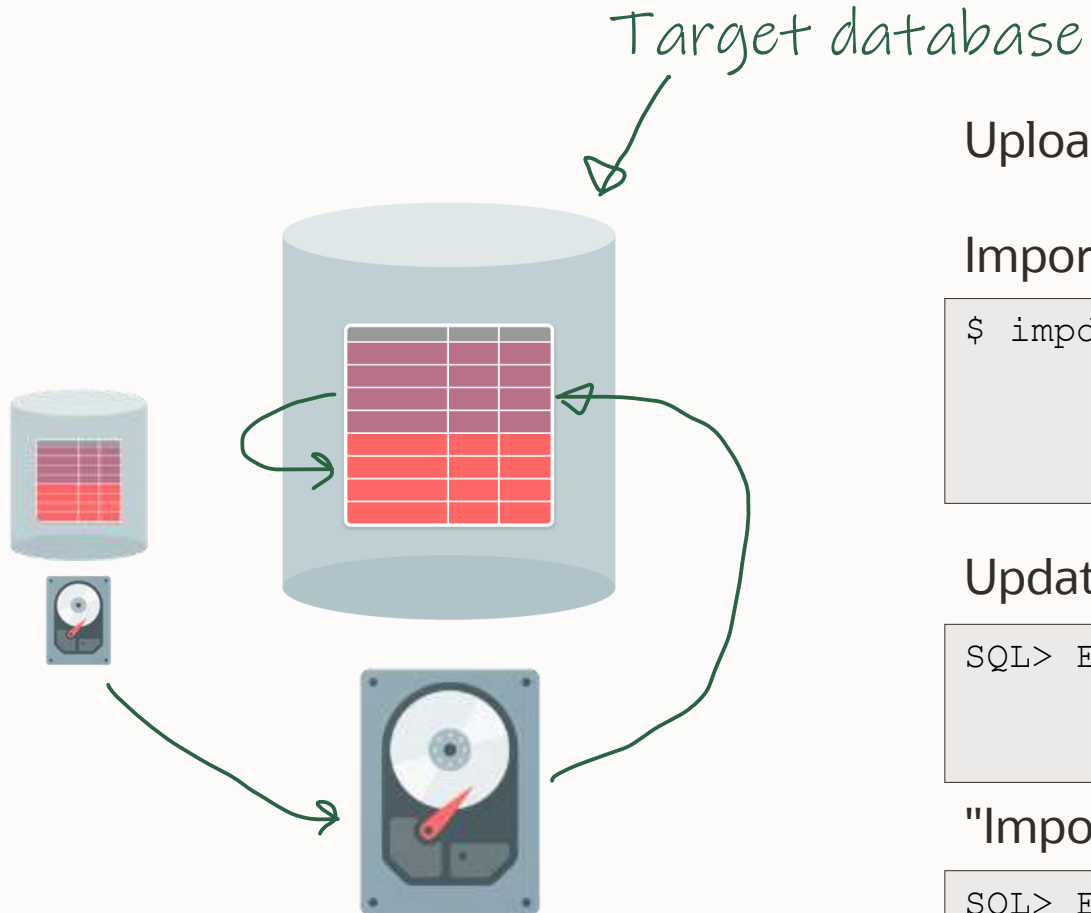
Export statistics to staging table

```
SQL> EXEC DBMS_STATS.EXPORT_SYSTEM_STATS (  
    statab => 'SYS_STATS_STG',  
    statown => 'SYSTEM');
```

Export staging table using Data Pump

```
$ expdp SYSTEM \  
    DIRECTORY=DATA_PUMP_DIR \  
    DUMPFILE=sys_stats_stg.dmp \  
    TABLES=SYS_STATS_STG
```

Transporting System Statistics | Workflow



Upload dump file to the object storage

Import staging table using Data Pump

```
$ impdp SYSTEM \  
  DIRECTORY=mydirectory \  
  DUMPFILE=sys_stats_stg.dmp \  
  TABLES=SYS_STATS_STG
```

Update statistics when you transport from a different version

```
SQL> EXEC dbms_stats.upgrade_stat_table (  
  ownname => 'SYSTEM',  
  statab => 'SYS_STATS_STG');
```

"Import" statistics from staging table

```
SQL> EXEC DBMS_STATS.IMPORT_SYSTEM_STATS (  
  statab => 'SYS_STATS_STG',  
  statown => 'SYSTEM');
```

You may repeat this for schema stats as well



Transporting Statistics | Demo

The screenshot shows a presentation slide with the title "Transporting Statistics | Demo". On the left side, there is a vertical flow diagram with two boxes: "Source" at the top and "Target" at the bottom, connected by a vertical line. To the right of this diagram is a list of six numbered steps:

1. Creating schema and adding data
2. Gathering and checking stats
3. Exporting stats
4. Creating empty schema and gathering stats
5. Importing statistics
6. Checking statistics

In the bottom right corner of the slide, there is a small red square icon containing a white circle, and below it, the text "Watch on YouTube" is displayed.

Transporting Statistics | Nice to Know

- The optimizer does not use statistics stored in a staging table - only from dictionary
- Importing statistics makes them current (i.e., not stale)
- You can transfer to a higher version - potentially the stats table must be upgraded

```
SQL> EXEC DBMS_STATS.IMPORT_SCHEMA_STATS ( ...  
  
ORA-20002: Version of statistics table "SALESAPP"."OPT_STATS_STG" is too old  
  
SQL> EXEC DBMS_STATS.UPGRADE_STAT_TABLE ('SALESAPP', 'OPT_STATS_STG');
```

- Incremental statistics: optionally export synopses as well

```
SQL> EXEC DBMS_STATS.EXPORT_SCHEMA_STATS (  
    ...  
    stat_category => 'OBJECT_STATS, REALTIME_STATS, SYNOPSIS');
```

Transporting Statistics | Nice to Know - 2

- Gather stats on the staging table after Data Pump import
 - Also, gather on staging table indexes
 - **Before** executing `DBMS_STATS.IMPORT_SCHEMA_STATS / IMPORT_TABLE_STATS`
- If enabled, imported statistics will be added as pending stats until you publish them

Statistics | Refresh?

Should you refresh object statistics when you upgrade or migrate to Oracle 19c?

- It depends
 - Don't do it if you **don't have** the exact amount of data to test with
 - Do it if you have the same amount of data
- Be cautious
 - When your source is 11.2, histograms can change
 - Avoid gradual change of plans when stats become stale
 - Workload is required to create correct statistics

Statistics | Disable Auto Tasks

In case you import your stats, better disable the Auto Task

```
SQL> EXEC DBMS_AUTO_TASK_ADMIN.DISABLE ( -  
      client_name => 'auto optimizer stats collection', -  
      operation   => NULL, -  
      window_name => NULL);
```

- Check:

```
SQL> SELECT client_name, status FROM dba_autotask_client;
```

CLIENT_NAME	STATUS
-----	-----
sql tuning advisor	ENABLED
auto optimizer stats collection	DISABLED
auto space advisor	ENABLED

Statistics | Gather Stats on Load

Since Oracle 12.1.0.2, statistics are gathered automatically on load

- CREATE TABLE AS SELECT (CTAS)
- INSERT AS SELECT (IAS)
- Feature can be turned off

```
_optimizer_gather_stats_on_load = FALSE
```

```
insert /*+append NO_GATHER_OPTIMIZER_STATISTICS*/ into MYTAB select ...
```



Photo by Jen Theodore on Unsplash

Should you lock statistics?

Locking Statistics | Use Cases

”

You can lock statistics to prevent them from changing.

[Database 19c SQL Tuning Guide, chapter 15](#)

- Certain static environments
- Highly volatile tables
- Enable use of dynamic statistics
- ... and all the exceptions

Locking Statistics | How to

Lock table statistics

```
SQL> EXEC DBMS_STATS.LOCK_TABLE_STATS(ownname=>'MYAPP', tabname=>'MY_VOLATILE_TAB1');
```

You can also lock at:

- Schema-level
- Partition-level

You can also unlock statistics

Pro tip: Locking table statistics also lock index and partition statistics



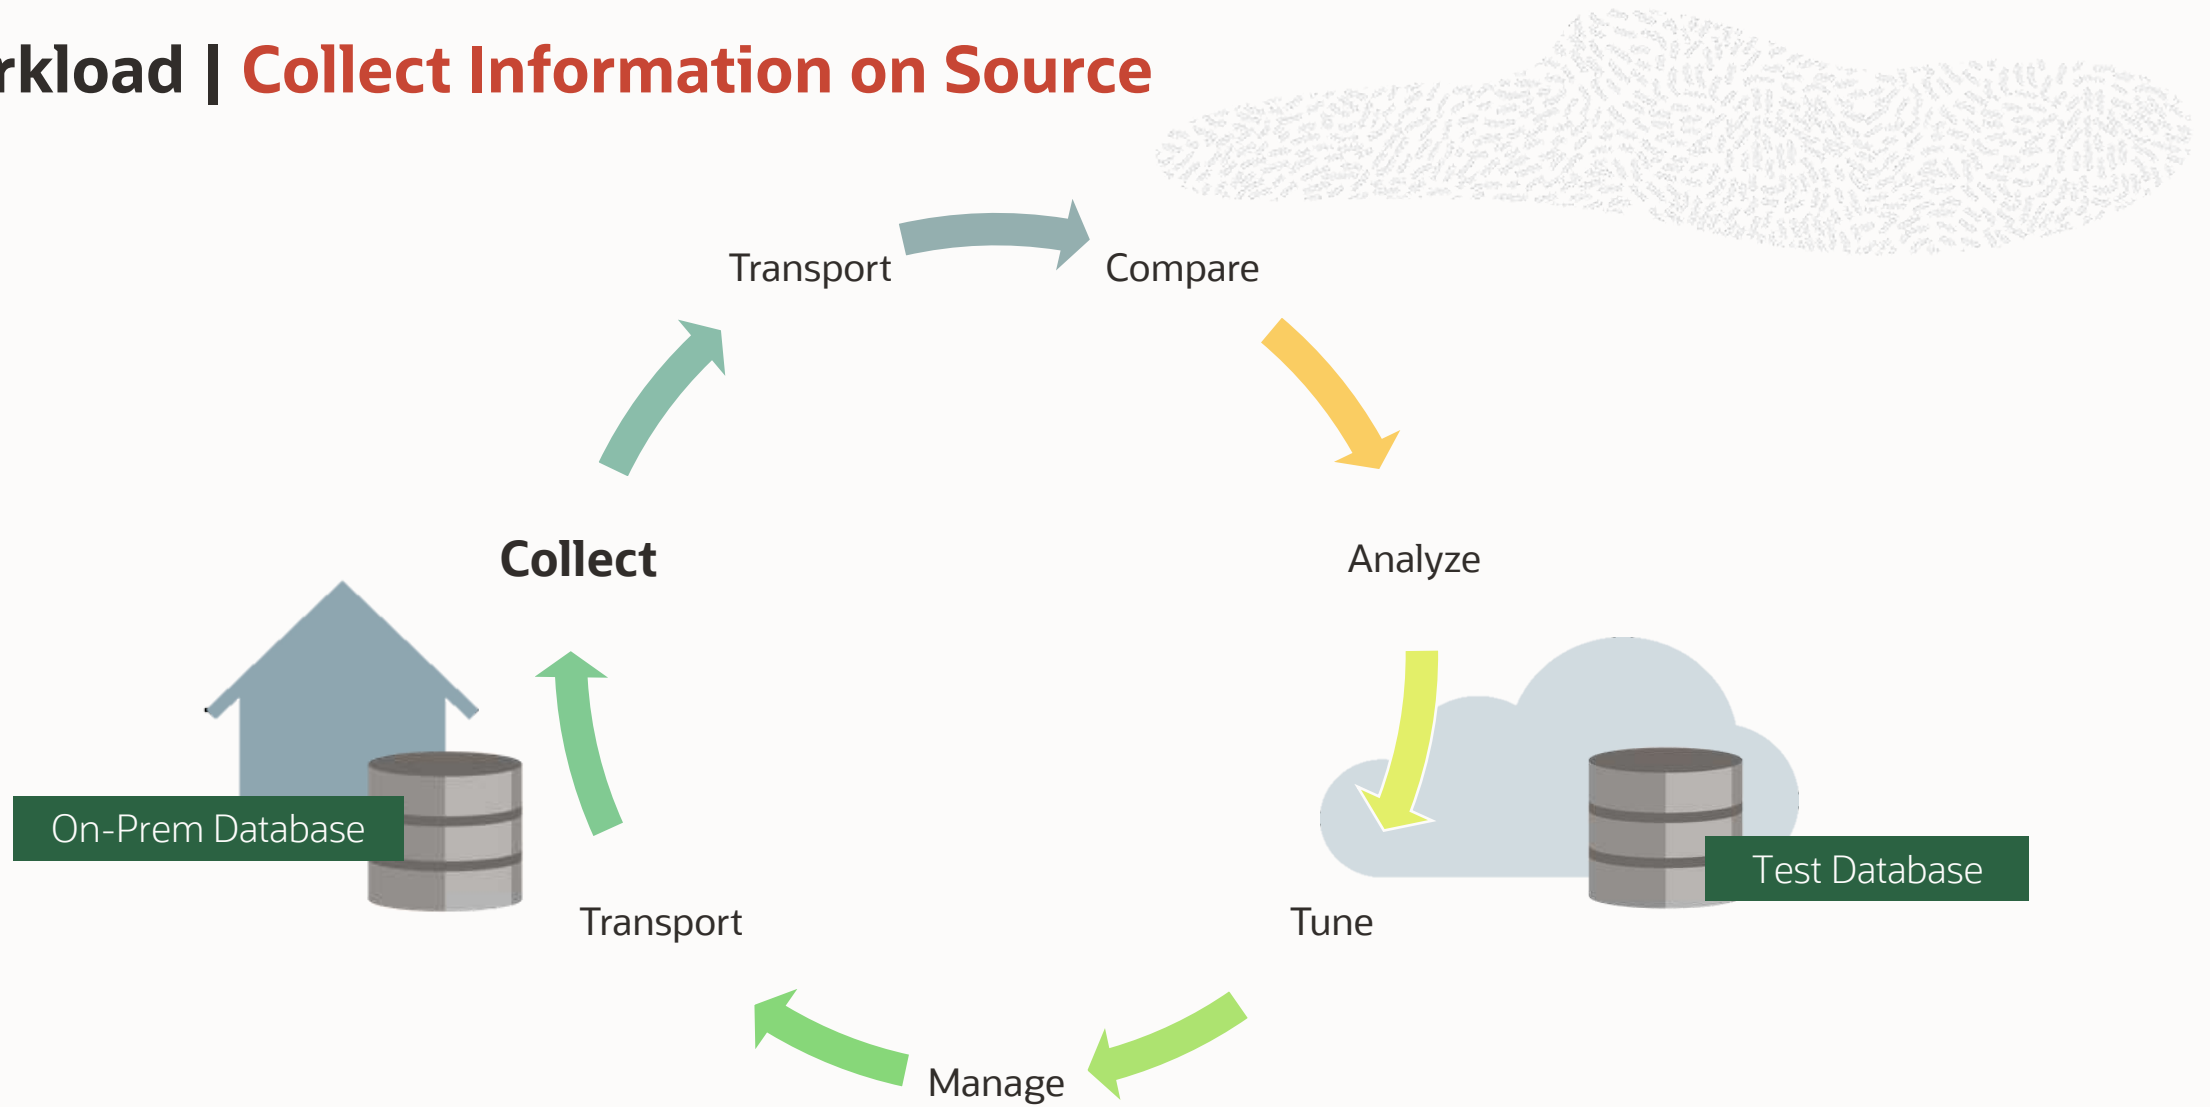
Locking Statistics | **Worth mentioning**

- Locking and unlocking statistics causes cursor invalidation
- To achieve plan stability, consider SQL Plan Management
- Statistics advisor will warn you about locked statistics
- Locking information is not exported

Challenge #6

Workload

Workload | **Collect Information on Source**



Workload | Objective



Collect workload information on source database



SQL Tuning Set

1. Create SQL Tuning Set (STS)
2. Capture workload from
 - a) AWR
 - b) Cursor cache



SQL Tuning Set | Definition



”

An SQL Tuning Set (STS) enables you to group SQL statements and related metadata in a single database object, which you can use to meet your tuning goals.

Specifically, SQL tuning sets achieve the following goals:

- *Providing input to the performance tuning advisors*
- *Transporting SQL between databases*

[Database 19c SQL Tuning Guide, chapter 23](#)


SQL Tuning Set | Definition



SQL statement



Context



Statistics



Plans



SQL Tuning Set | Create



First, create a SQL Tuning Set

```
SQL> BEGIN
  DBMS_SQLSET.CREATE_SQLSET (
    sqlset_name => 'UPG_STS_1',
    description => 'For upgrade - from source'
  );
END;
/
```



Pro tip: You can also use [DBMS_SQLTUNE](#) to create a SQL Tuning Set

SQL Tuning Set | Capture



Next, capture statements from AWR

```
SQL> DECLARE
  begin_id number;
  end_id number;
  cur sys_refcursor;
BEGIN
  SELECT min(snap_id), max(snap_id) INTO begin_id, end_id
  FROM dba_hist_snapshot;

  open cur for
  select value(p) from table(dbms_sqltune.select_workload_repository(
    begin_snap      => begin_id,
    end_snap        => end_id,
    basic_filter    => 'parsing_schema_name not in (''SYS'')',
    ranking_measure1 => 'elapsed_time',
    result_limit    => 5000,
    attribute_list  => 'ALL')) p;

  dbms_sqltune.load_sqlset('UPG_STS_1', cur);

close cur;

END;
/
```



Pro tip: Consider excluding other internal schemas like *DBSNMP*, *ORACLE_OCM*, *LBACSYS*, *WMSYS*, *XDB*, *SYSTEM*



SQL Tuning Set | Capture



Optionally, capture statements from cursor cache

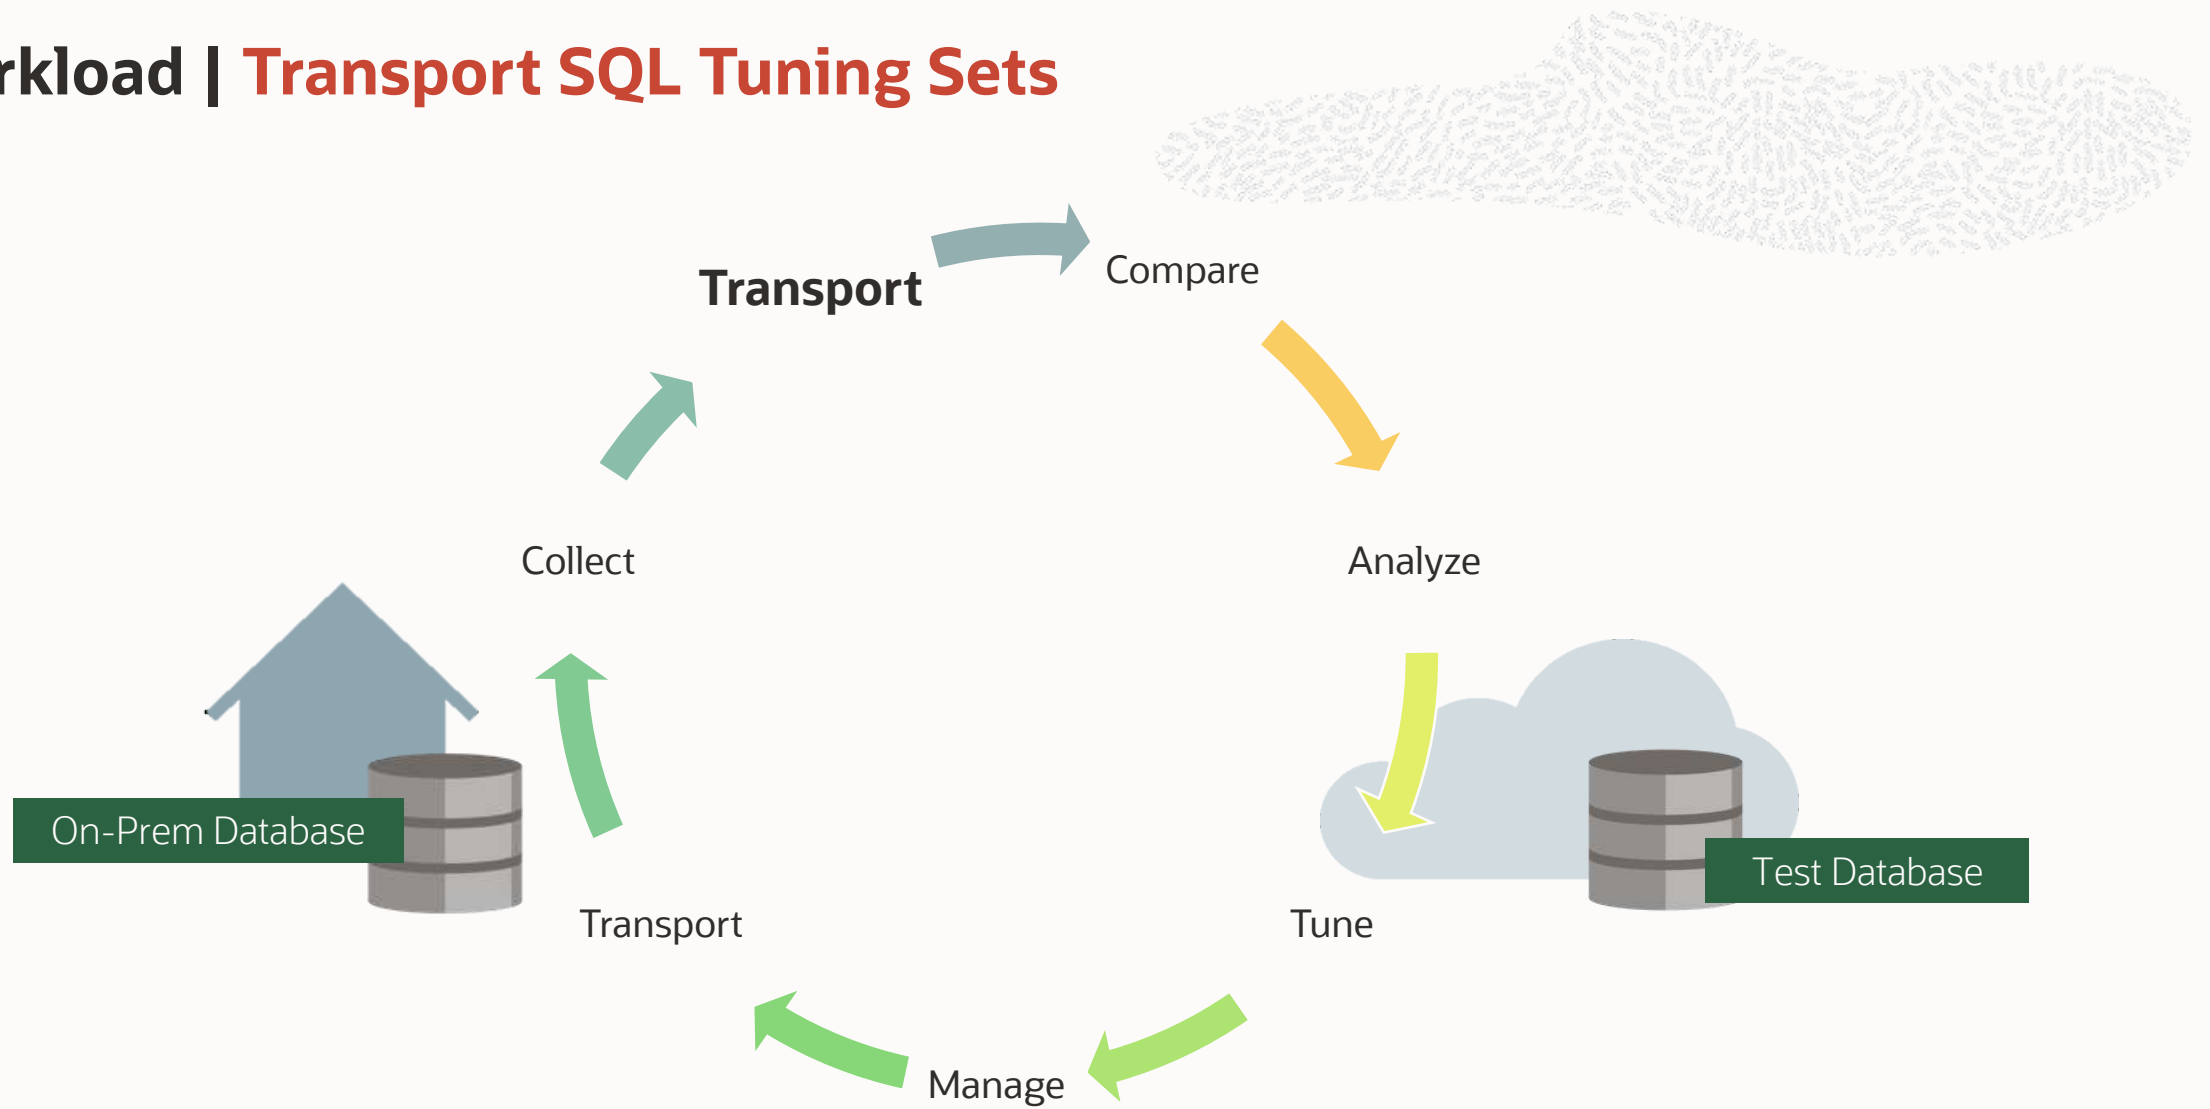
```
SQL> BEGIN
  DBMS_SQLSET.CAPTURE_CURSOR_CACHE_SQLSET(
    sqlset_name      => 'UPG_STS_1',
    time_limit       => 900,
    repeat_interval  => 60,
    capture_option   => 'MERGE',
    capture_mode     => DBMS_SQLTUNE.MODE_ACCUMULATE_STATS,
    basic_filter     => 'parsing_schema_name not in (''SYS'')',
    sqlset_owner     => NULL,
    recursive_sql    => 'HAS_RECURSIVE_SQL');
END;
/
```



Careful - puts load on your system

Pro tip: [SQL Tuning Guide](#) shows how to load all statements from a given schema

Workload | Transport SQL Tuning Sets



Workload | Objective



Transport SQL Tuning Set



SQL Tuning Set + Data Pump

1. Pack SQL Tuning Set into staging table
2. Export staging table
3. Upload to object storage
4. Import into test database in the cloud



SQL Tuning Set | Transport



Pack into staging table on **source** database

```
SQL> BEGIN
  DBMS_SQLTUNE.CREATE_STGTAB_SQLSET (
    table_name          => 'UPG_STGTAB_1');
  DBMS_SQLTUNE.PACK_STGTAB_SQLSET (
    sqlset_name         => 'UPG_STS_1',
    staging_table_name => 'UPG_STGTAB_1');
END;
```

Optionally, use `DBMS_SQLTUNE.REMAP_STGTAB_SQLSET` to remap between `CON_DBID`

Export with Data Pump

```
$ expdp user \
  directory=mydirectory
  dumpfile=upg_stgtab_1.dmp
  tables=UPG_STGTAB_1
```



SQL Tuning Set | **Transport**



Optionally, compress with Data Pump (license required)

```
$ expdp ... compression=all compression_algorithm=medium
```

Or, use OS utilities

```
$ zip upg_stgtab_1.zip upg_stgtab_1.dmp
```



SQL Tuning Set | **Transport**



1. Transfer dump file directly via `scp` or `rsync`

2. Or, use OCI CLI

```
$ oci os object put --namespace oradbclouducm \  
-bn zdm-oss \  
--file upg_stgtab_1.dmp \  
--part-size 100 \  
--parallel-upload-count 4
```

Create Pre-Authenticated Request (PAR) and download

```
$ wget https://objectstorage... /o/upg_stgtab_1.dmp
```



SQL Tuning Set | **Transport**



Import with Data Pump to **target** database

```
$ impdp user \  
    directory=mydirectory  
    dumpfile=upg_stgtab_1.dmp  
    tables=UPG_STGTAB_1
```

Unpack staging table

```
SQL> BEGIN  
    DBMS_SQLTUNE.UNPACK_STGTAB_SQLSET (  
        sqlset_name          => '%'  
        replace              => true  
        staging_table_name   => 'UPG_STGTAB_1'  
    );  
END;  
/
```



SQL Tuning Set | License



”

SQL Tuning Sets can also be accessed by way of database server APIs and command-line interfaces. Usage of any subprograms in the DBMS_SQLSET package to manage SQL Tuning Sets is part of the EE and EE-ES offerings.

In addition, the following subprograms, part of the DBMS_SQLTUNE package, provide an older interface to manage SQL Tuning Sets and are also part of the EE and EE-ES offerings:

*ADD_SQLSET_REFERENCE
CREATE_STGTAB_SQLSET
LOAD_SQLSET
SELECT_CURSOR_CACHE
UNPACK_STGTAB_SQLSET*

*CAPTURE_CURSOR_CACHE_SQLSET
DELETE_SQLSET
PACK_STGTAB_SQLSET
SELECT_SQLSET
UPDATE_SQLSET*

*CREATE_SQLSET
DROP_SQLSET
REMOVE_SQLSET_REFERENCE
SELECT_WORKLOAD_REPOSITORY*

[Database 19c Database Licensing Information User Manual](#)

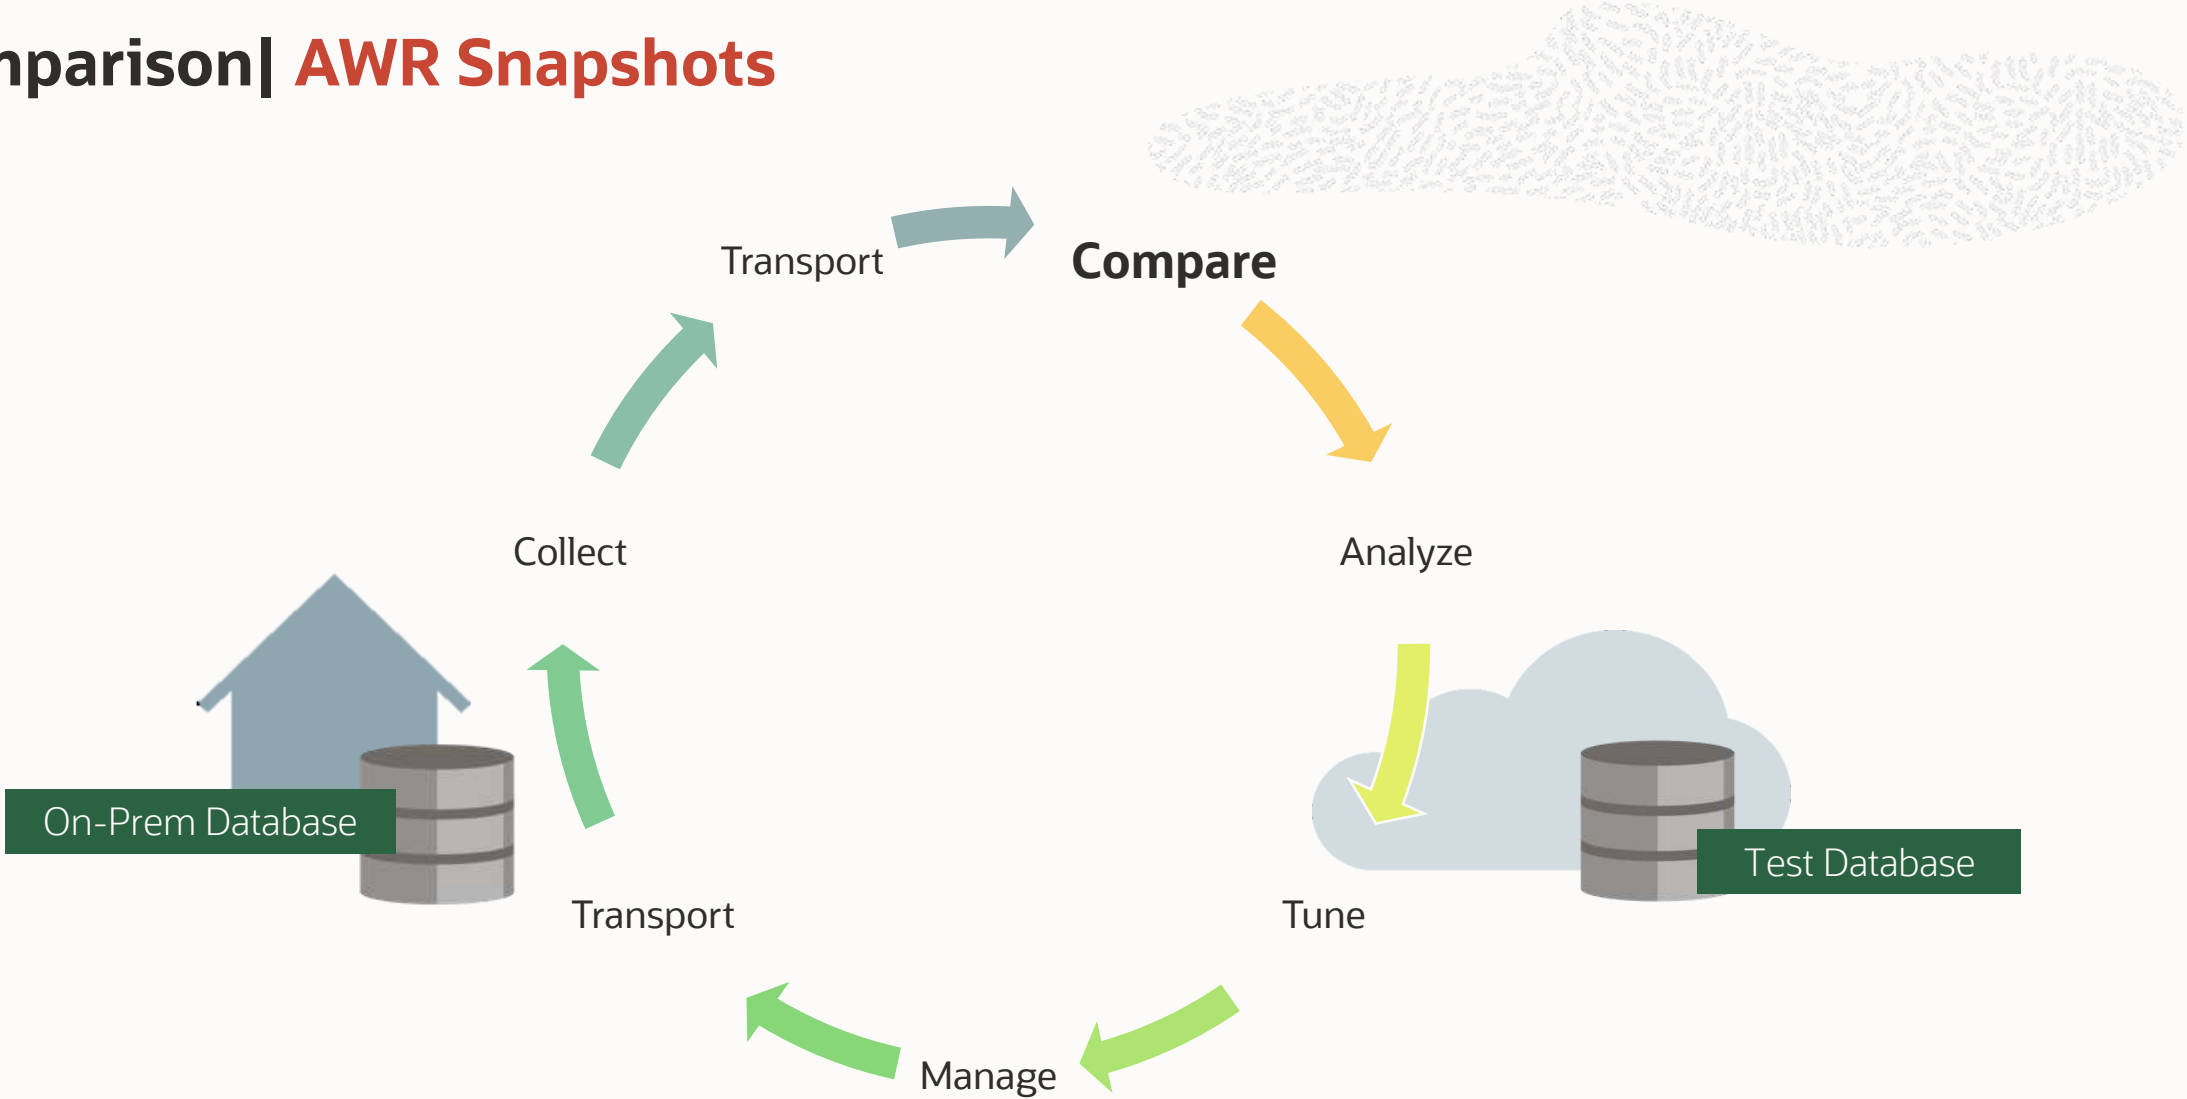
SQL Tuning Set | Recommendation

Always capture workload data into SQL Tuning Sets

Challenge #7

Comparison

Comparison | **AWR Snapshots**



Comparison | Objective



Compare workload data



AWR snapshots and diff report

1. Identify two snapshots for a representative workload
2. Export AWR
3. Upload to object storage
4. Import AWR into test database in the cloud
5. Run representative workload in the cloud
6. Compare AWR snapshots



Due to completely different specs and load scenarios, the result may be misleading or not useful



Comparison | AWR Snapshots



Comparison | AWR Diff Report



Use script awrddrpi.sql →

```
Instances in this Workload Repository schema
~~~~~
  DB Id          Inst Num          DB Name          Instance          Host
  -----
  72245725          1          UPGR          UPGR  hol.localdom
* 753780962          1          CDB2          CDB2  hol.localdom

Database Id and Instance Number for the First Pair of Snapshots
~~~~~
Enter value for dbid:
```

Top Timed Events

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

1st						2nd					
Event	Wait Class	Waits	Time(s)	Avg Time(ms)	%DB time	Event	Wait Class	Waits	Time(s)	Avg Time(ms)	%DB time
CPU time			68,289.05		43.73	db file sequential read	User I/O	22,193,998	114,919.21	5.18	23.17
db file sequential read	User I/O	6,686,953	37,737.81	5.64	24.17	enq: SS - contention	Configuration	3,913	98,997.90	25,299.74	19.96
gc buffer busy	Cluster	12,508,244	23,886.55	1.91	15.30	CPU time			73,786.55		14.88
TCP Socket (KGAS)	Network	680,629	12,514.65	18.39	8.01	row cache lock	Concurrency	73,940	48,472.30	655.56	9.77
db file scattered read	User I/O	1,572,296	4,271.68	2.72	2.74	reliable message	Other	41,148	47,600.87	1,156.82	9.60



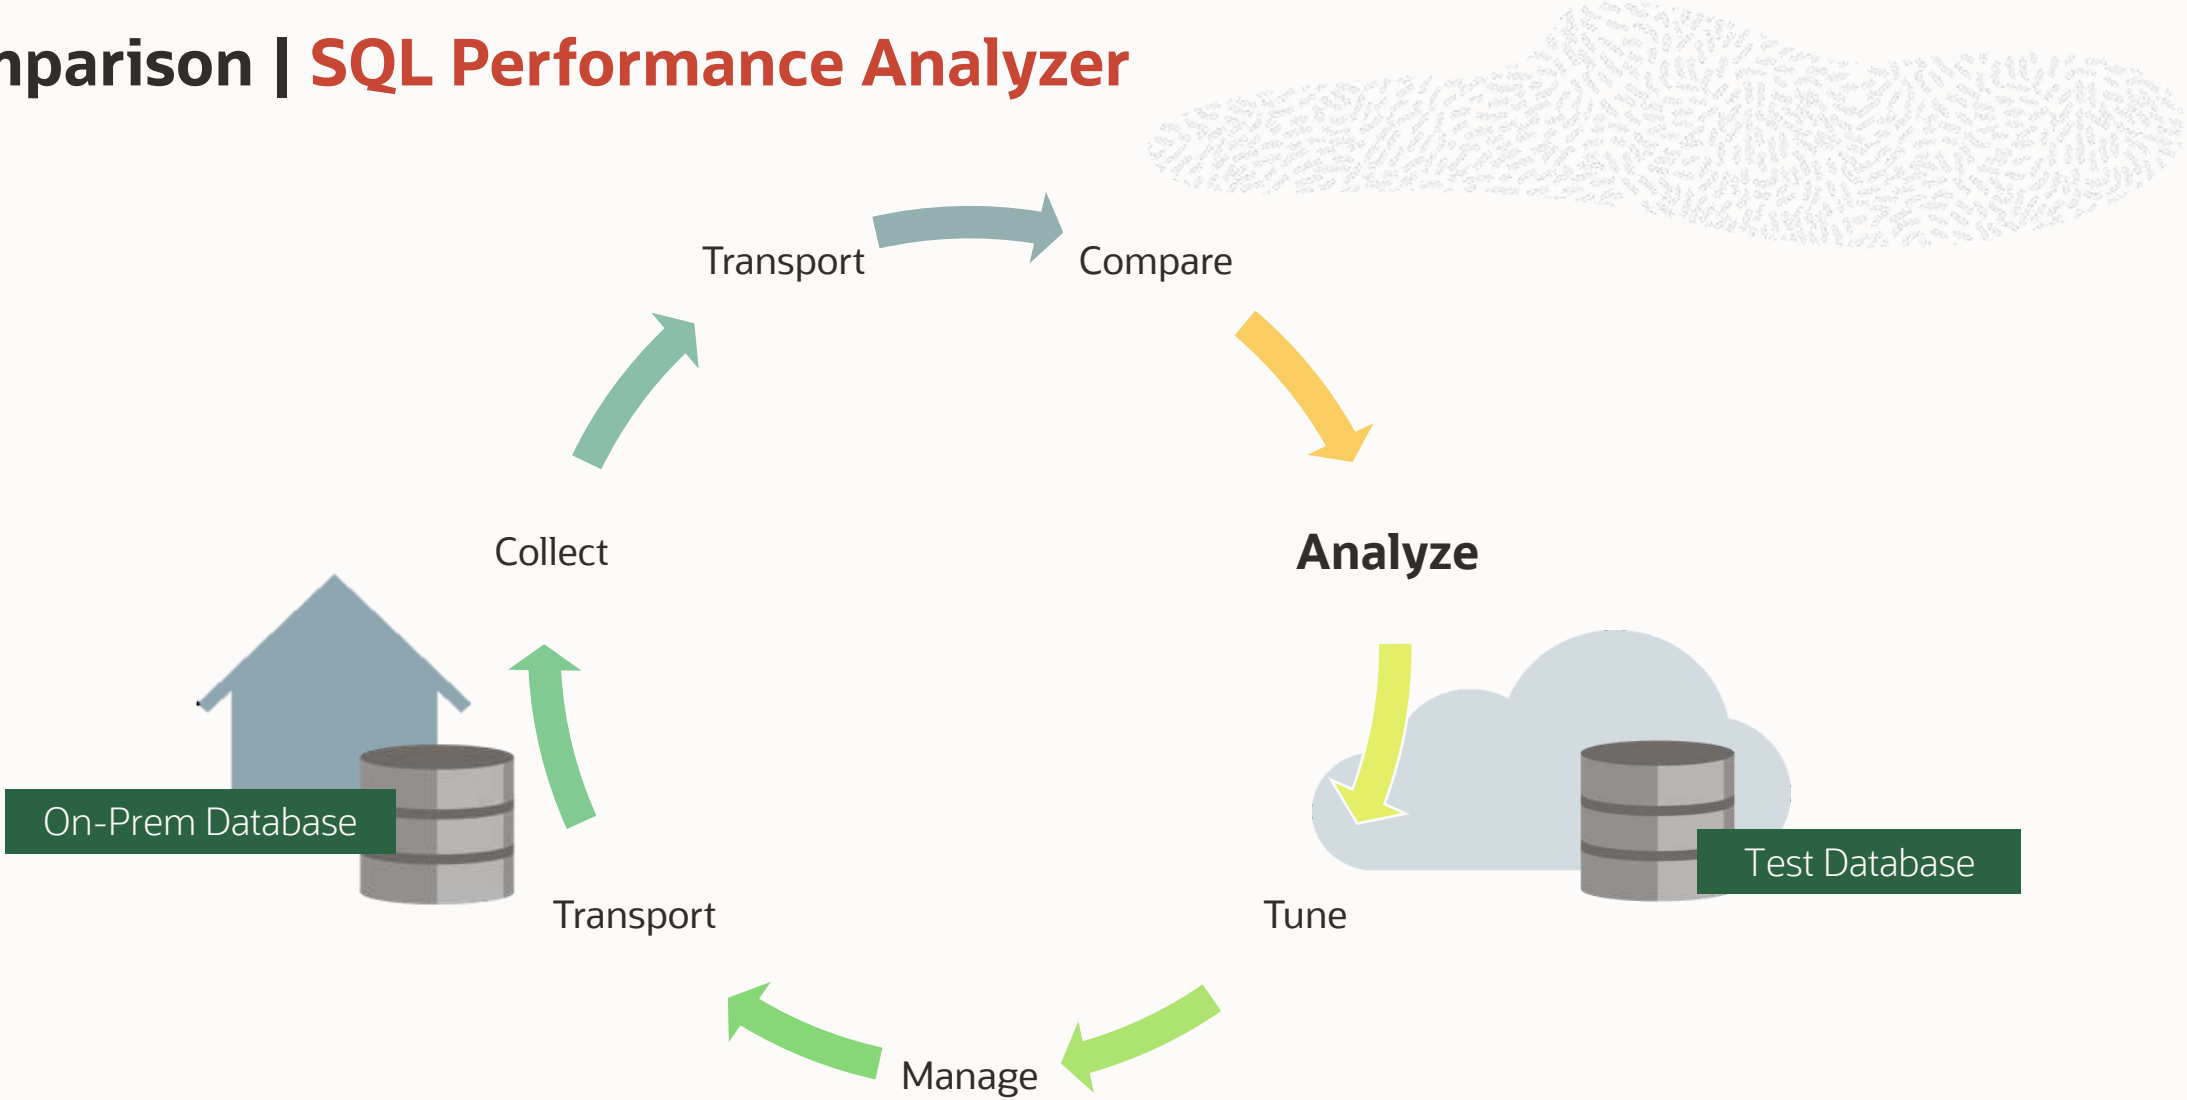
Comparison | **AWR Snapshots**



For AWR Diff Reports to be really useful

You should run *before* and *after* workload in cloud

Comparison | SQL Performance Analyzer



Comparison | Objective



Find bad plans by simulating SQL execution



SQL Performance Analyzer

1. Build ANALYSIS task
2. Simulate SQL statements captured on production
 - a) AWR
 - b) Cursor Cache
3. Compare results
4. Remedy issues
5. Repeat analysis



SQL Performance Analyzer | SPA



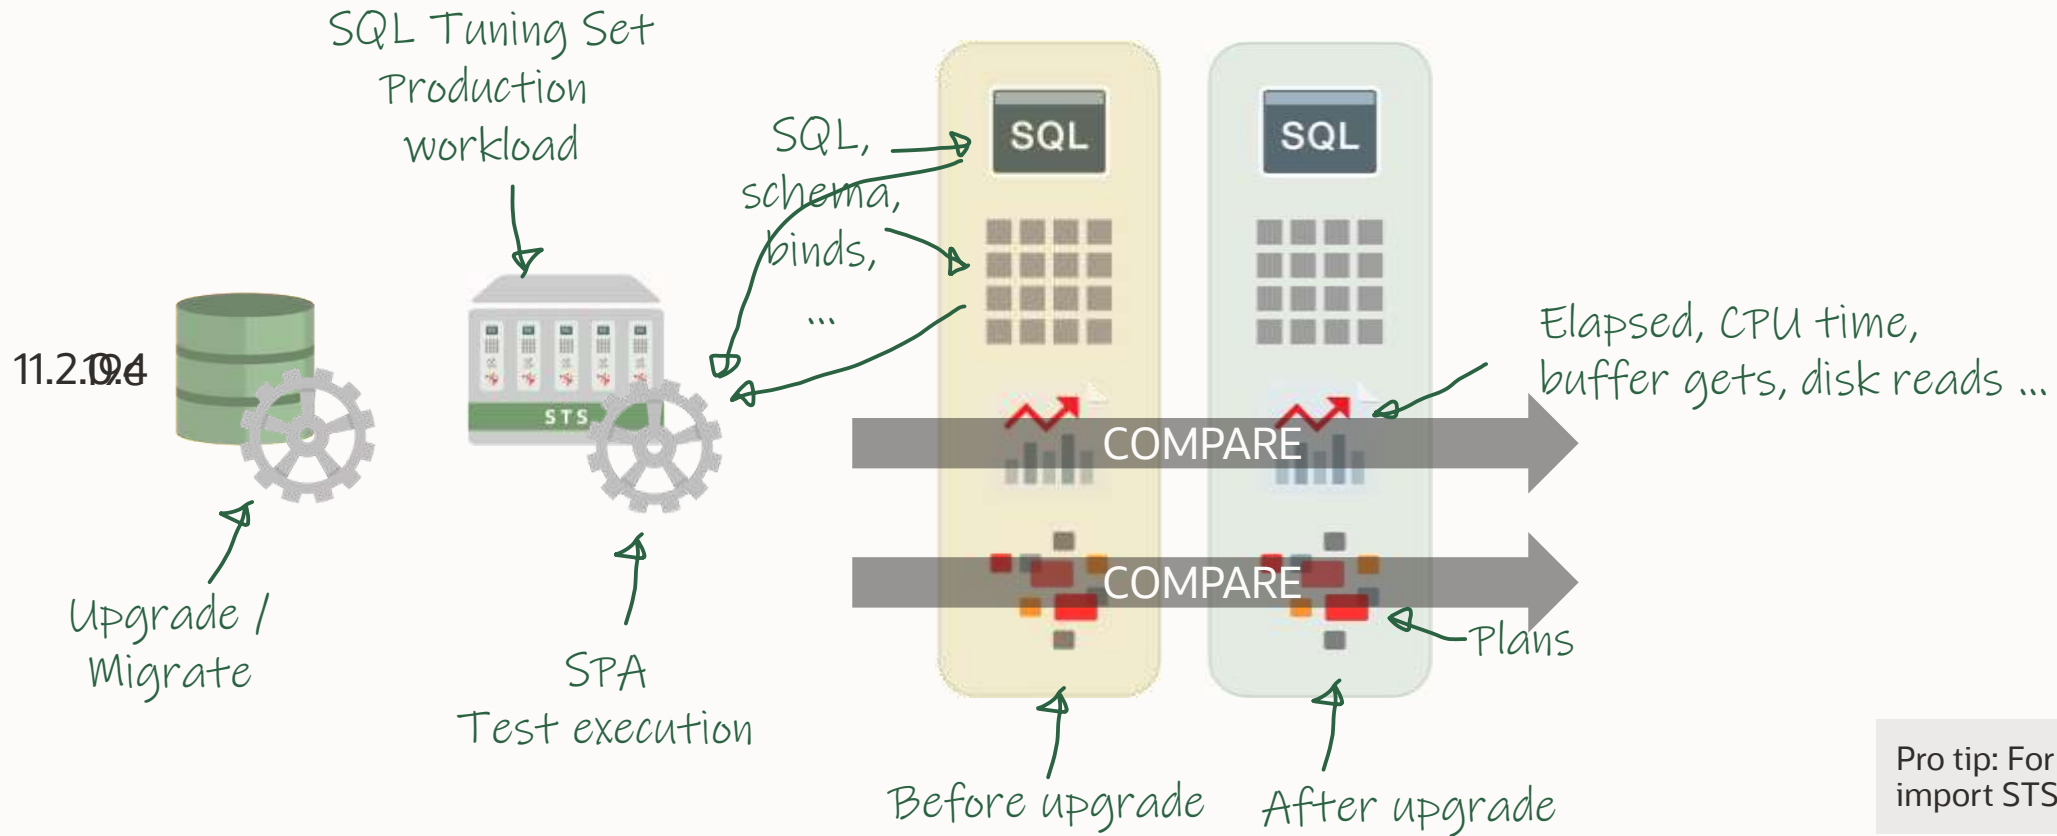
”

*SPA provides fine-grained assessment of environment changes on SQL **execution plans** and **statistics** by running the SQL statements both in isolation and serially manner in before-change and after-change environments.*

SPA functionality is well integrated with existing SQL Tuning Set (STS), SQL Tuning Advisor, and SQL Plan Management functionality.

[Oracle Database Real Application Testing Data Sheet](#)

SPA | Concept



Pro tip: For migrations, import STS into target database



SPA | Regressed Report



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

From production workload

From test execution

SPA | Regressed Report



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

SPA | Regressed Report



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

SQL Details: czzzubf8fjz96

Parsing Schema APPS

Execution Frequency 3

SQL Text

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

Single Execution Statistics

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



SPA | Regressed Report



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

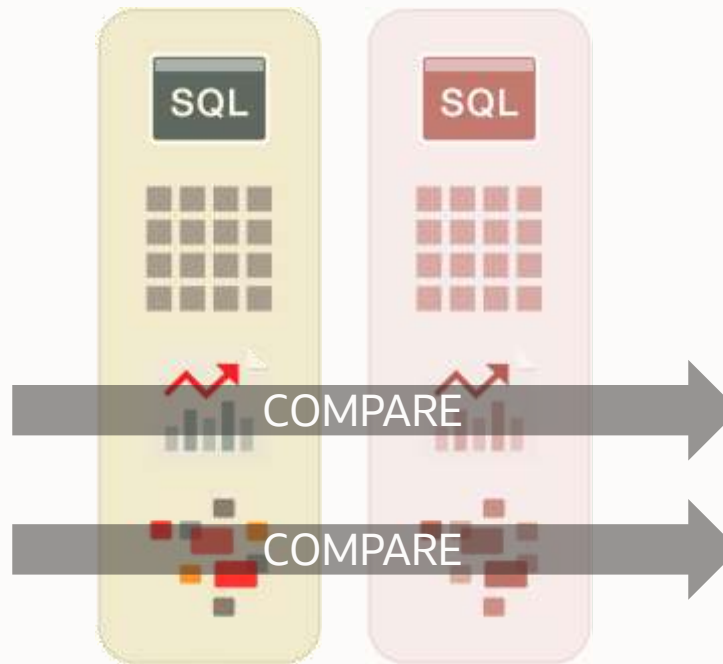


Plan Comparison
SQL_TRIAL_1353942463446
 Plan Hash Value 1165613724

Expand All | Collapse All

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



SPA | Continuous Improvement






Photo by [Andrew George](#) on [Unsplash](#)

Some ideas for testing with SPA

- init.ora Parameters
- New statistics
- Optimizer settings
- DBMS_OPTIM_BUNDLE_fix_control settings
- Histogram types
- New synopsis for incremental statistics
- Specific patches

SPA | Validate Rows and Values

NEW FEATURE since Oracle 18c

```
exec DBMS_SQLPA.SET_ANALYSIS_TASK_PARAMETER ( -  
    task_name => 'my_spa_task', -  
    parameter => 'COMPARE_RESULTSET', -  
    value      => 'TRUE');
```

SPA | Scripts

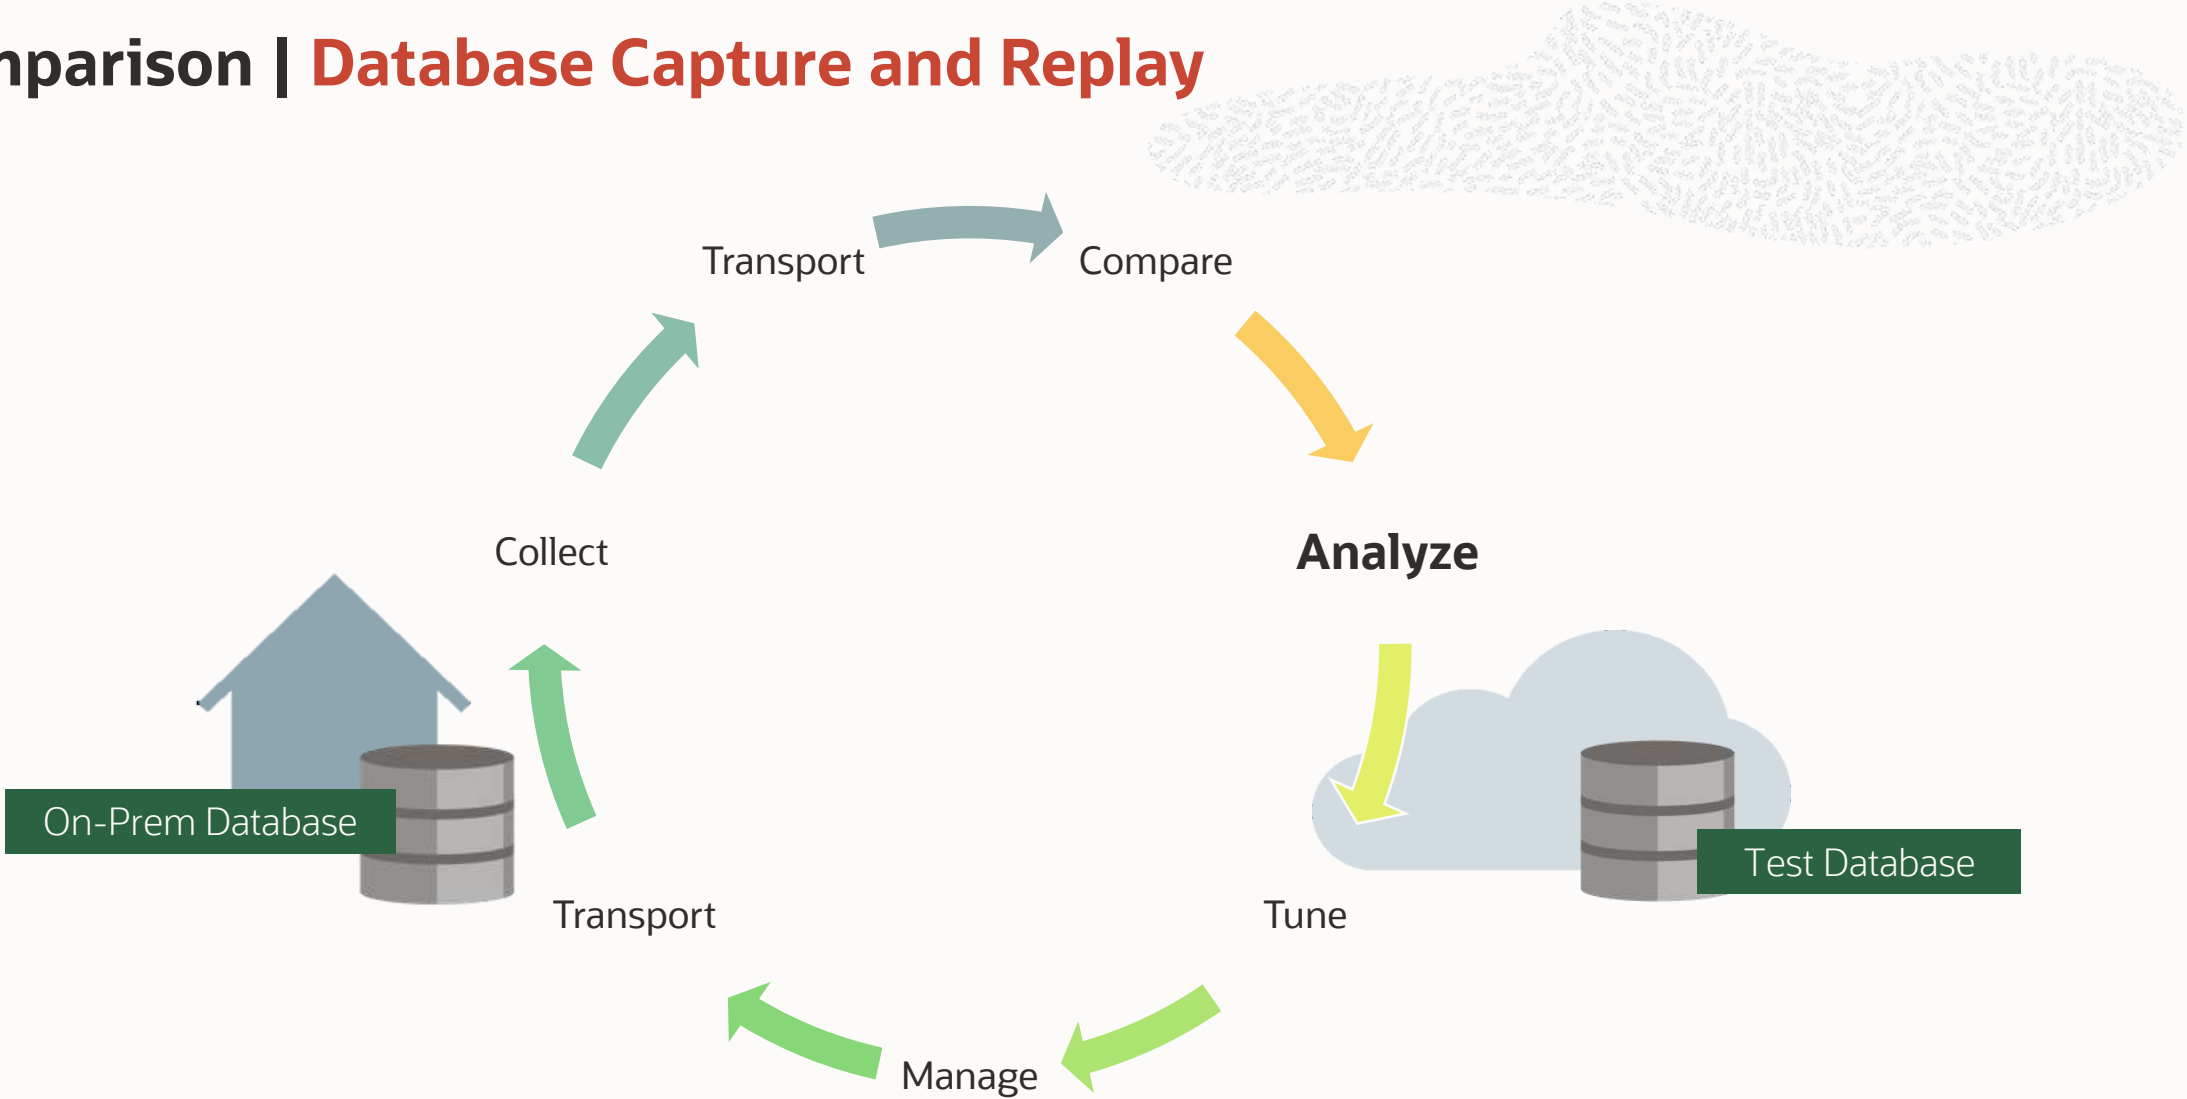
Find the scripts to try out SPA on:

<https://MikeDietrichDE.com/scripts/>

11.2, 12.1, 12.2, 18c, 19c	capture_awr.sql	Capture statements from AWR – uses the oldest and newest available snapshot by default and limits to 5000 statements by elapsed time.
11.2, 12.1, 12.2, 18c, 19c	capture_cc.sql	Capture statements from Cursor Cache – uses a conservative test setting for 30 seconds only polling the CC every 5 seconds – needs to be adjusted for live systems.
19c and below	spa_cpu.sql	Run a full SQL Performance Analyzer task on captured SQL Statements – compare on CPU
19c and below	spa_buffer.sql	Run a full SQL Performance Analyzer task on capture SQL Statements – compare on buffer gets
19c and below	spa_elapsed.sql	Run a full SQL Performance Analyzer task on capture SQL Statements – compare on elapsed time
19c and below	spa_report_cpu.sql	Generate a regressed report comparing two workloads based on CPU
19c and below	spa_report_buffer.sql	Generate a regressed report comparing two workloads based on buffer gets
19c and below	spa_report_elapsed.sql	Generate a regressed report comparing two workloads based on elapsed time



Comparison | Database Capture and Replay



Database Replay | Overview



”

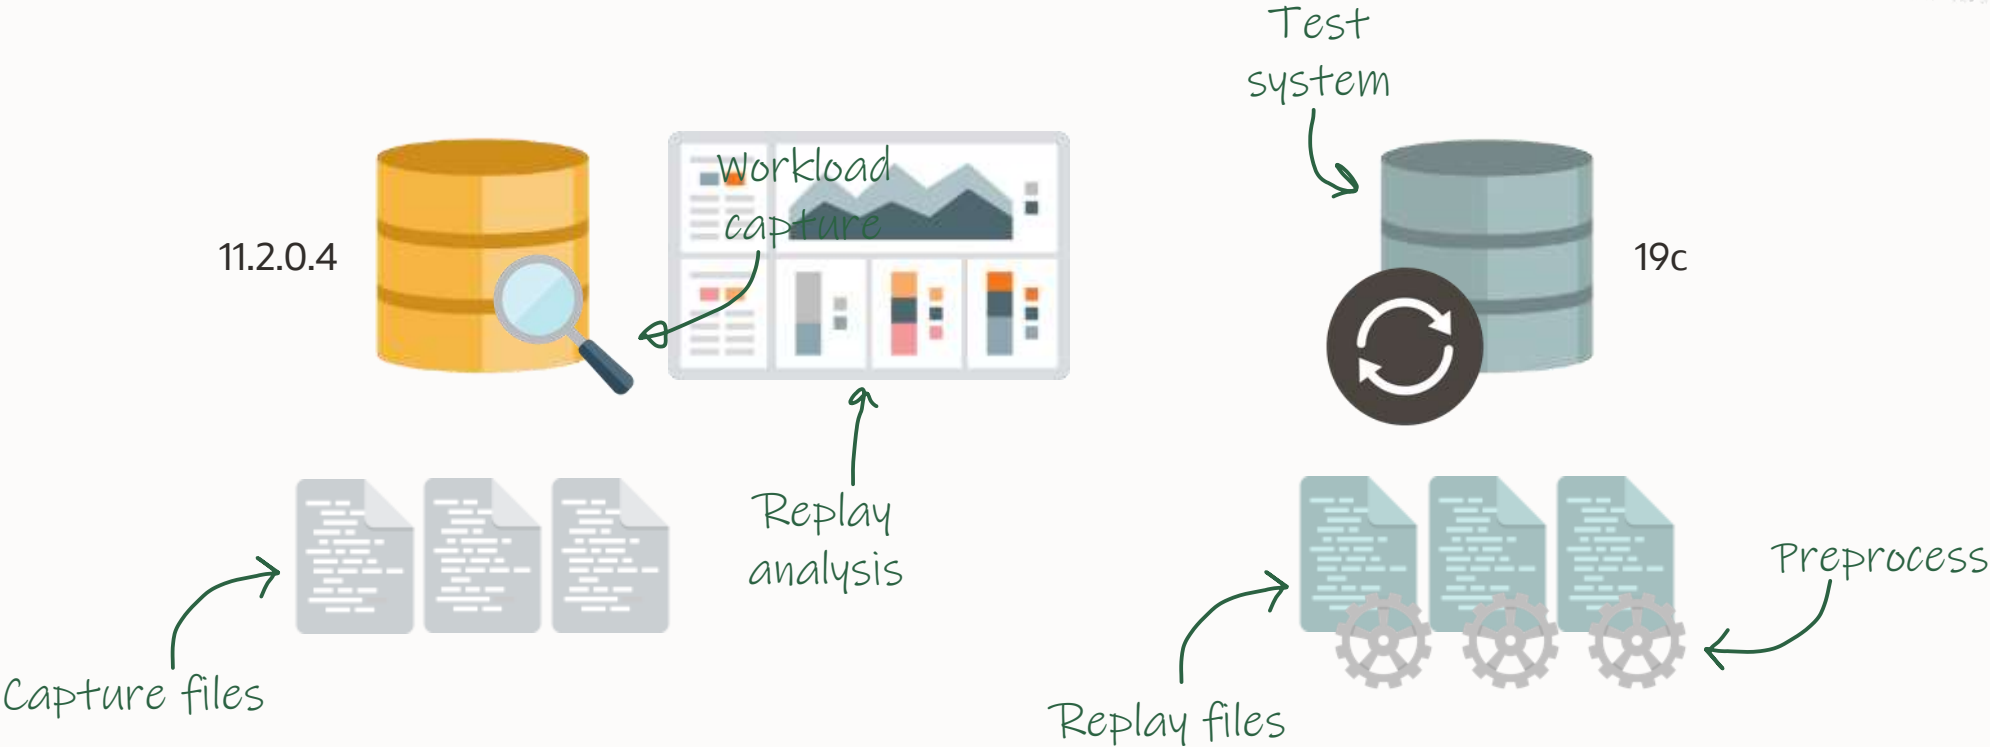
You can use Database Replay to capture a workload on the production system and replay it on a test system with the exact timing, concurrency, and transaction characteristics of the original workload.

This enables you to test the effects of a system change without affecting the production system.


[Database 19c Testing Guide, chapter 9](#)

Requires Real Application Testing on source

Database Replay | Overview



Database Replay | Facts

1. Platform independent
2. RAC compliant - optionally, change the number of nodes
3. Per-PDB capture/replay 
4. Capture and replay across database releases
5. Download scripts



Database Replay | To Consider

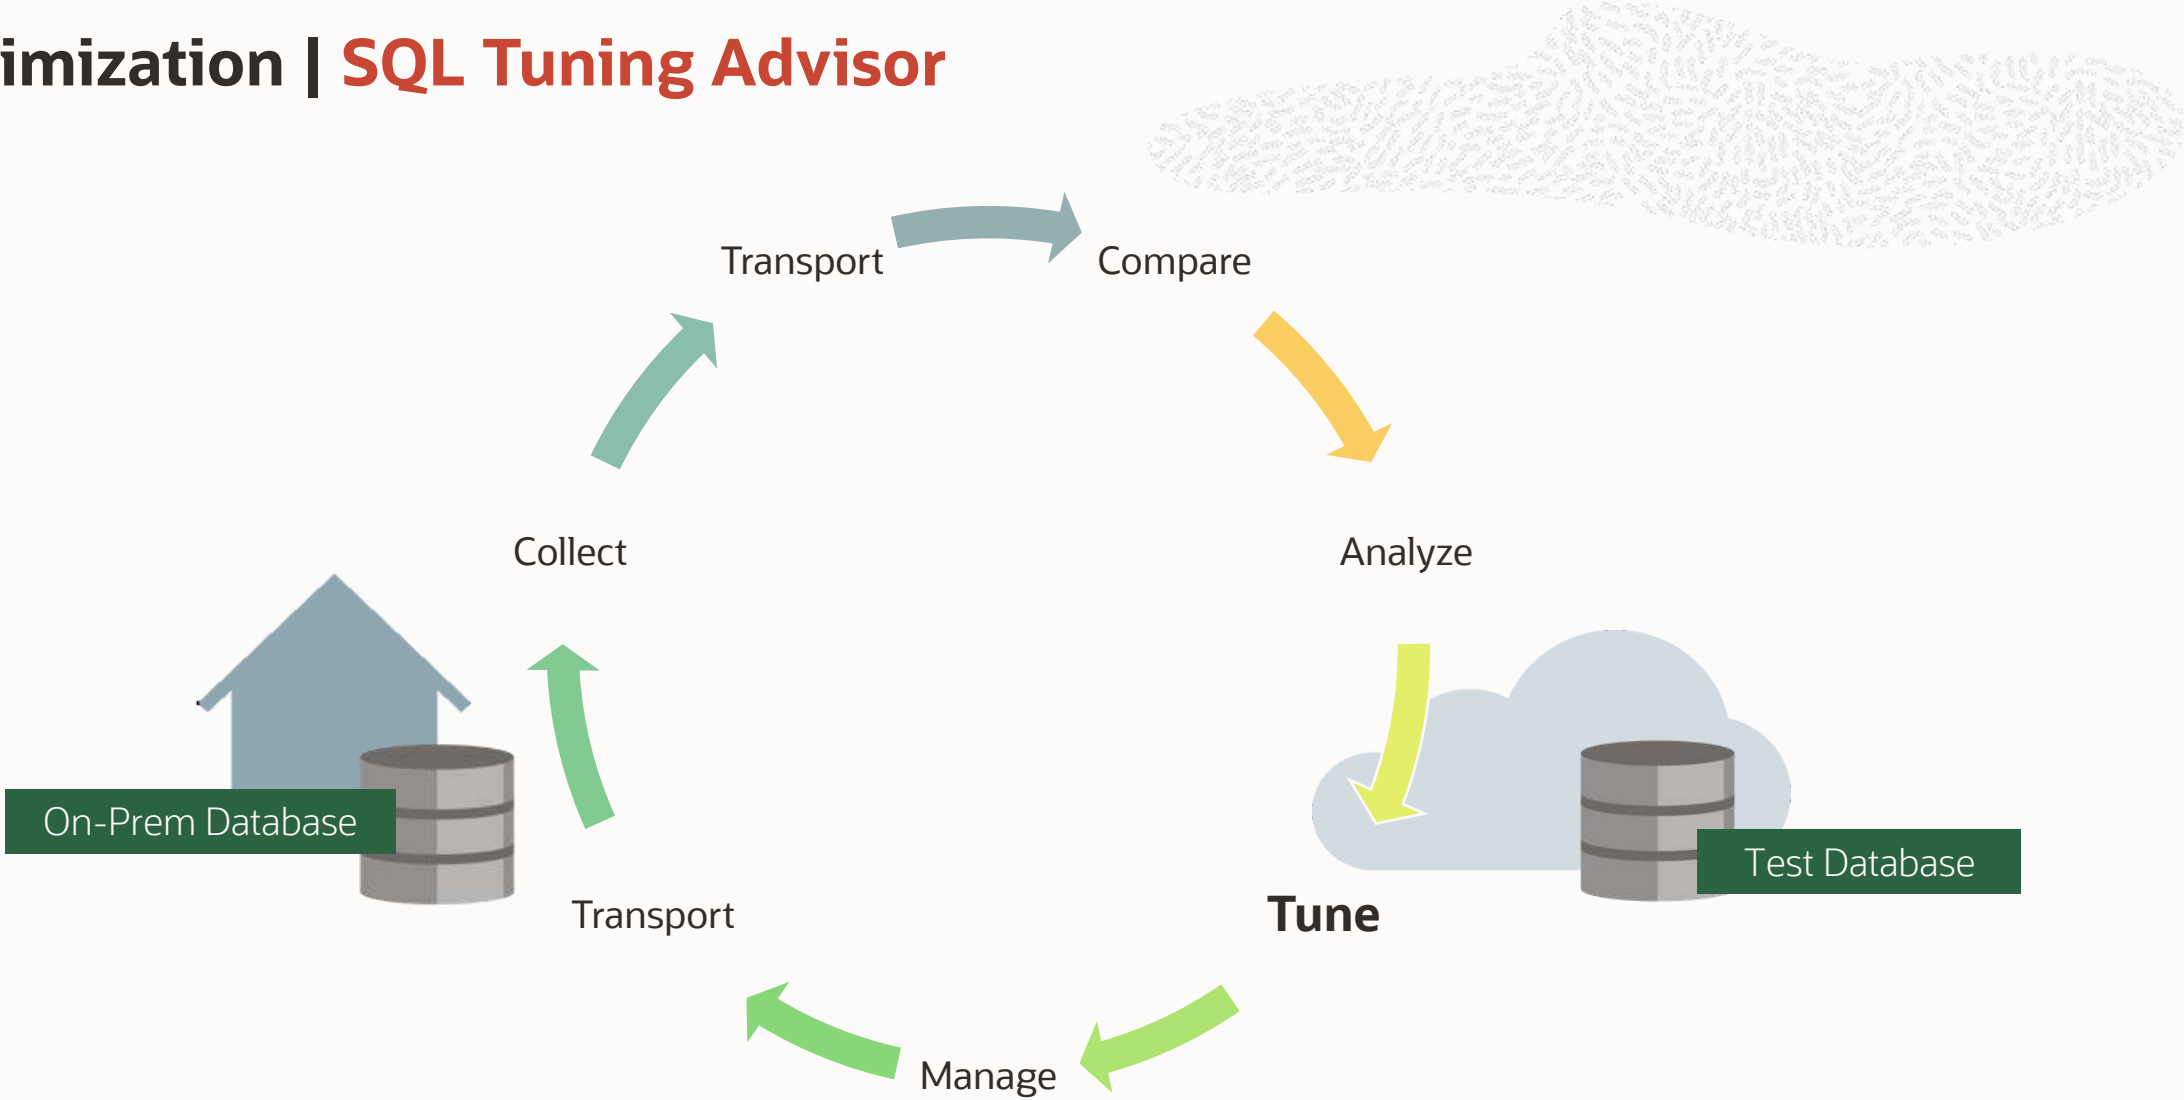
1. Workload capture restrictions
2. Not suitable with **external dependencies**
 - Database link, external tables, UTL_HTTP
3. Recommended to **restart** database before capture
 - Startup in restricted mode, capture automatically sets unrestricted mode
4. Work best from dedicated SCN
 - Data Pump FLASHBACK_SCN or restore to specific SCN



Challenge #8

Optimization

Optimization | SQL Tuning Advisor



SQL Tuning Advisor



”

SQL Tuning Advisor is SQL diagnostic software in the Oracle Database Tuning Pack.

...

SQL Tuning Advisor is a mechanism for resolving problems related to suboptimally performing SQL statements.

[Database 19c SQL Tuning Guide, chapter 24](#)

SQL Tuning Advisor | Findings

Types of findings:

1. Collection of object statistics
 2. Creation of indexes
 3. Rewriting SQL statements
 4. Creation of SQL profiles
- and more

Pro tip: SQL Developer has a good [interface](#) to SQL Tuning Advisor



SQL Profiles



”

A SQL profile is a database object that contains auxiliary statistics specific to a SQL statement.

...

The corrected statistics in a SQL profile can improve optimizer cardinality estimates, which in turn leads the optimizer to select better plans.

[Database 19c SQL Tuning Guide, chapter 26](#)



SQL Profiles | Facts

1. Part of Tuning Pack
 - Included in some cloud offerings
2. Stores a set of hints that causes the optimizer to select a plan
3. Affects one statement only



SQL Profiles | Facts

4. You can enable/disable a profile
5. Transparent to application
 - Does not require application changes
6. Persistent and transportable
 - [Documentation](#)
7. Useful with literals using `FORCE_MATCH=TRUE`

SQL Profile | Testing



1. Enable profile for selected environments only

```
SQL> exec :p_name := dbms_sqltune.accept_sql_profile(  
    task_name=>'TASK_21944',  
    name=>'XT_PROFILE',  
    category=>'TEST_ENV');
```

2. Verify the profile – it doesn't get used by the optimizer in the live environment

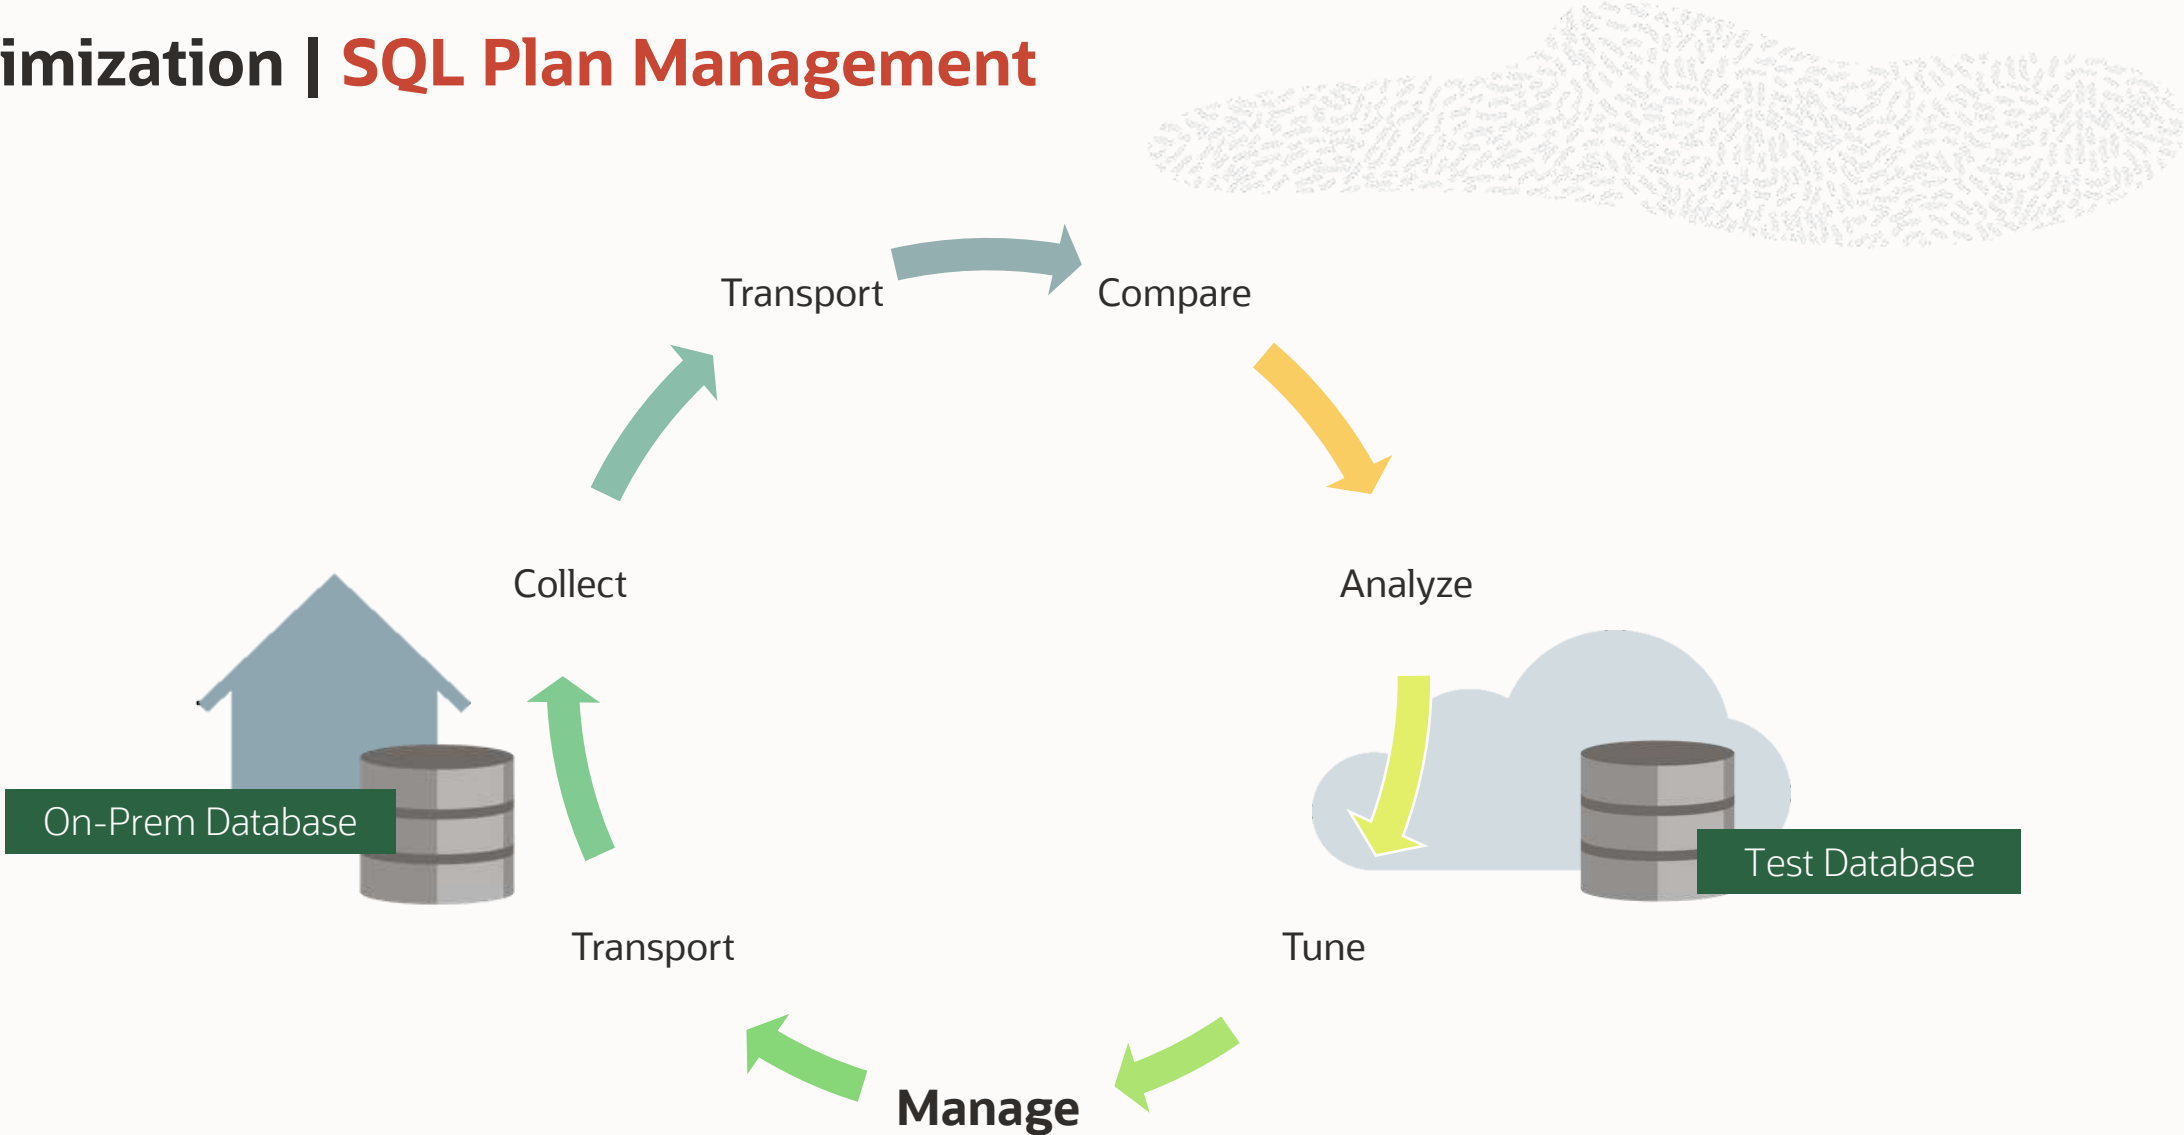
```
SQL> alter session set sqltune_category='TEST_ENV';
```

3. Accept and make visible to all sessions ('DEFAULT')

```
SQL> exec dbms_sqltune.alter_sql_profile(  
    name=>'XT_PROFILE',  
    attribute_name=>'CATEGORY',  
    value=>'DEFAULT');
```



Optimization | SQL Plan Management



SQL Plan Management | SPM



”

*SQL plan management uses a mechanism called a **SQL plan baseline**, which is a set of accepted plans that the optimizer is allowed to use for a SQL statement.*

...

SQL plan management prevents performance regressions caused by plan changes.

[Database 19c SQL Tuning Guide, chapter 27](#)

SPM | Concept



SQL SQL

Repeatable SQL



Filtering changed

- Server statistics
- SQL text
- SQL parameters
- Profile
- Action

```
OPTIMIZER_USE_SQL_PLAN_BASELINES=TRUE
```

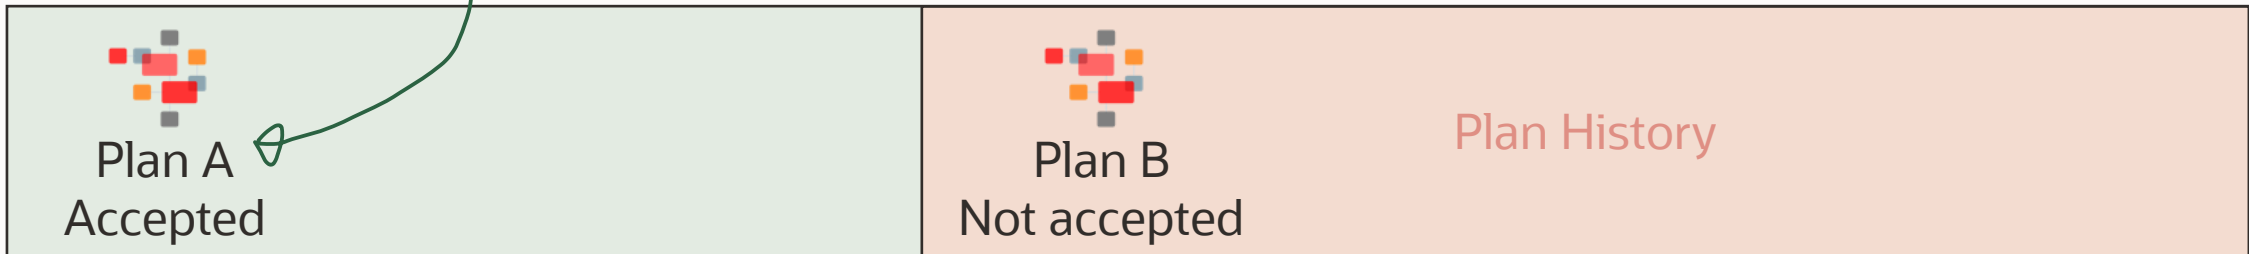
No plans in baseline

Plan A Filter

Plan A is used

Plan history

Dedicated part of SQL plan baseline



SPM | Plans



The plans in a SQL plan baseline can be:

- Enabled
- Accepted
- Fixed

To change status use `DBMS_SPM.ALTER_SQL_PLAN_BASELINE`

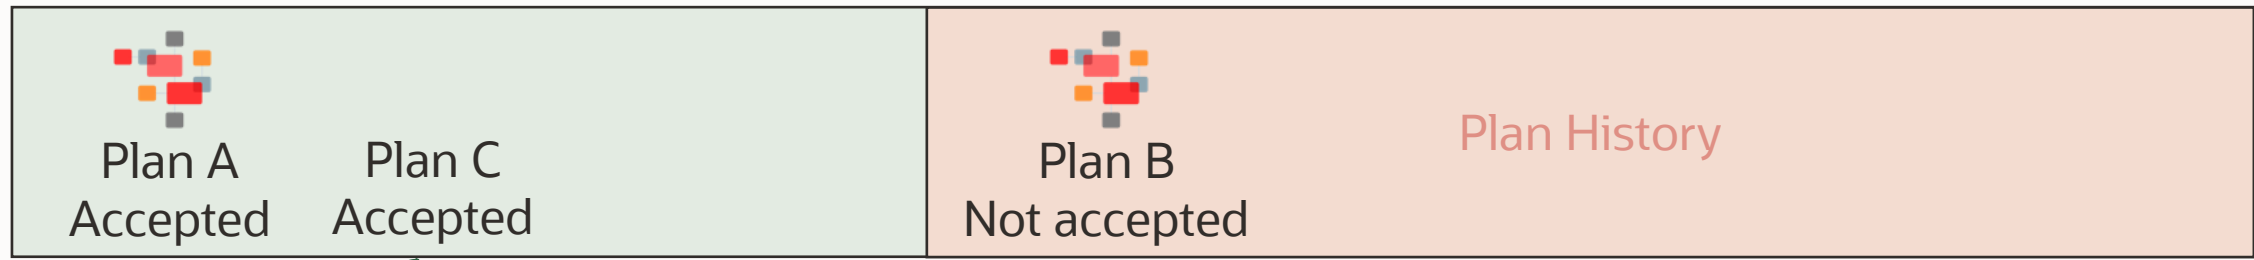
You can also prevent plans from getting purged by setting the `autopurge` property.

Pro tip: The *Accepted* attribute can only be set by a test execution

SPM | Load from STS



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



Automatically
accepted



SPM | What if ... literals



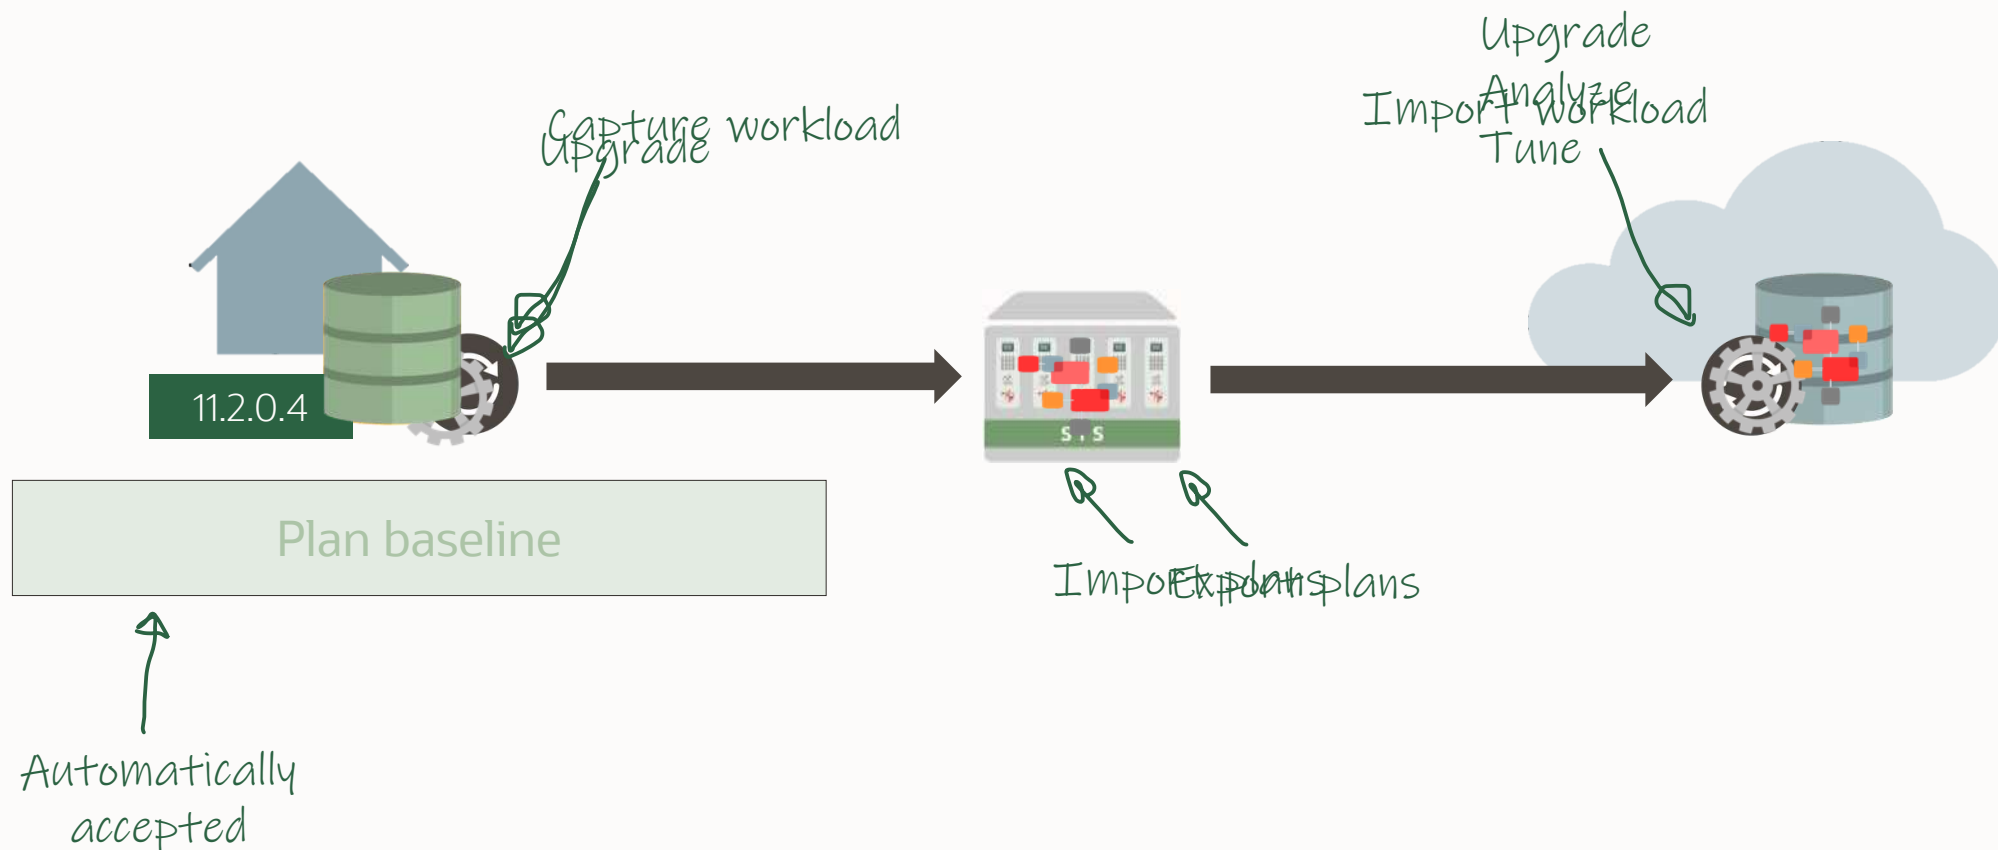
SQL Plan Management in a system with literals is not a good fit

- Many distinct statements
- CURSOR_SHARING = FORCE? No!
- SQL profiles can do force matching

Optimal solution: Change your application to use bind variables



SPM | Use Case



SPM | Use Case



Plan baseline

```
SQL> DECLARE
  plans_loaded NUMBER;
  filter VARCHAR2(255);
BEGIN
  filter := 'sql_id='czzzubf8fjz96' AND plan_hash_value='1165613724''';

  plans_loaded := DBMS_SPM.LOAD_PLANS_FROM_SQLSET (
    sqlset_name => 'UPG_STS_1',
    basic_filter => filter
  );
END;
/
```

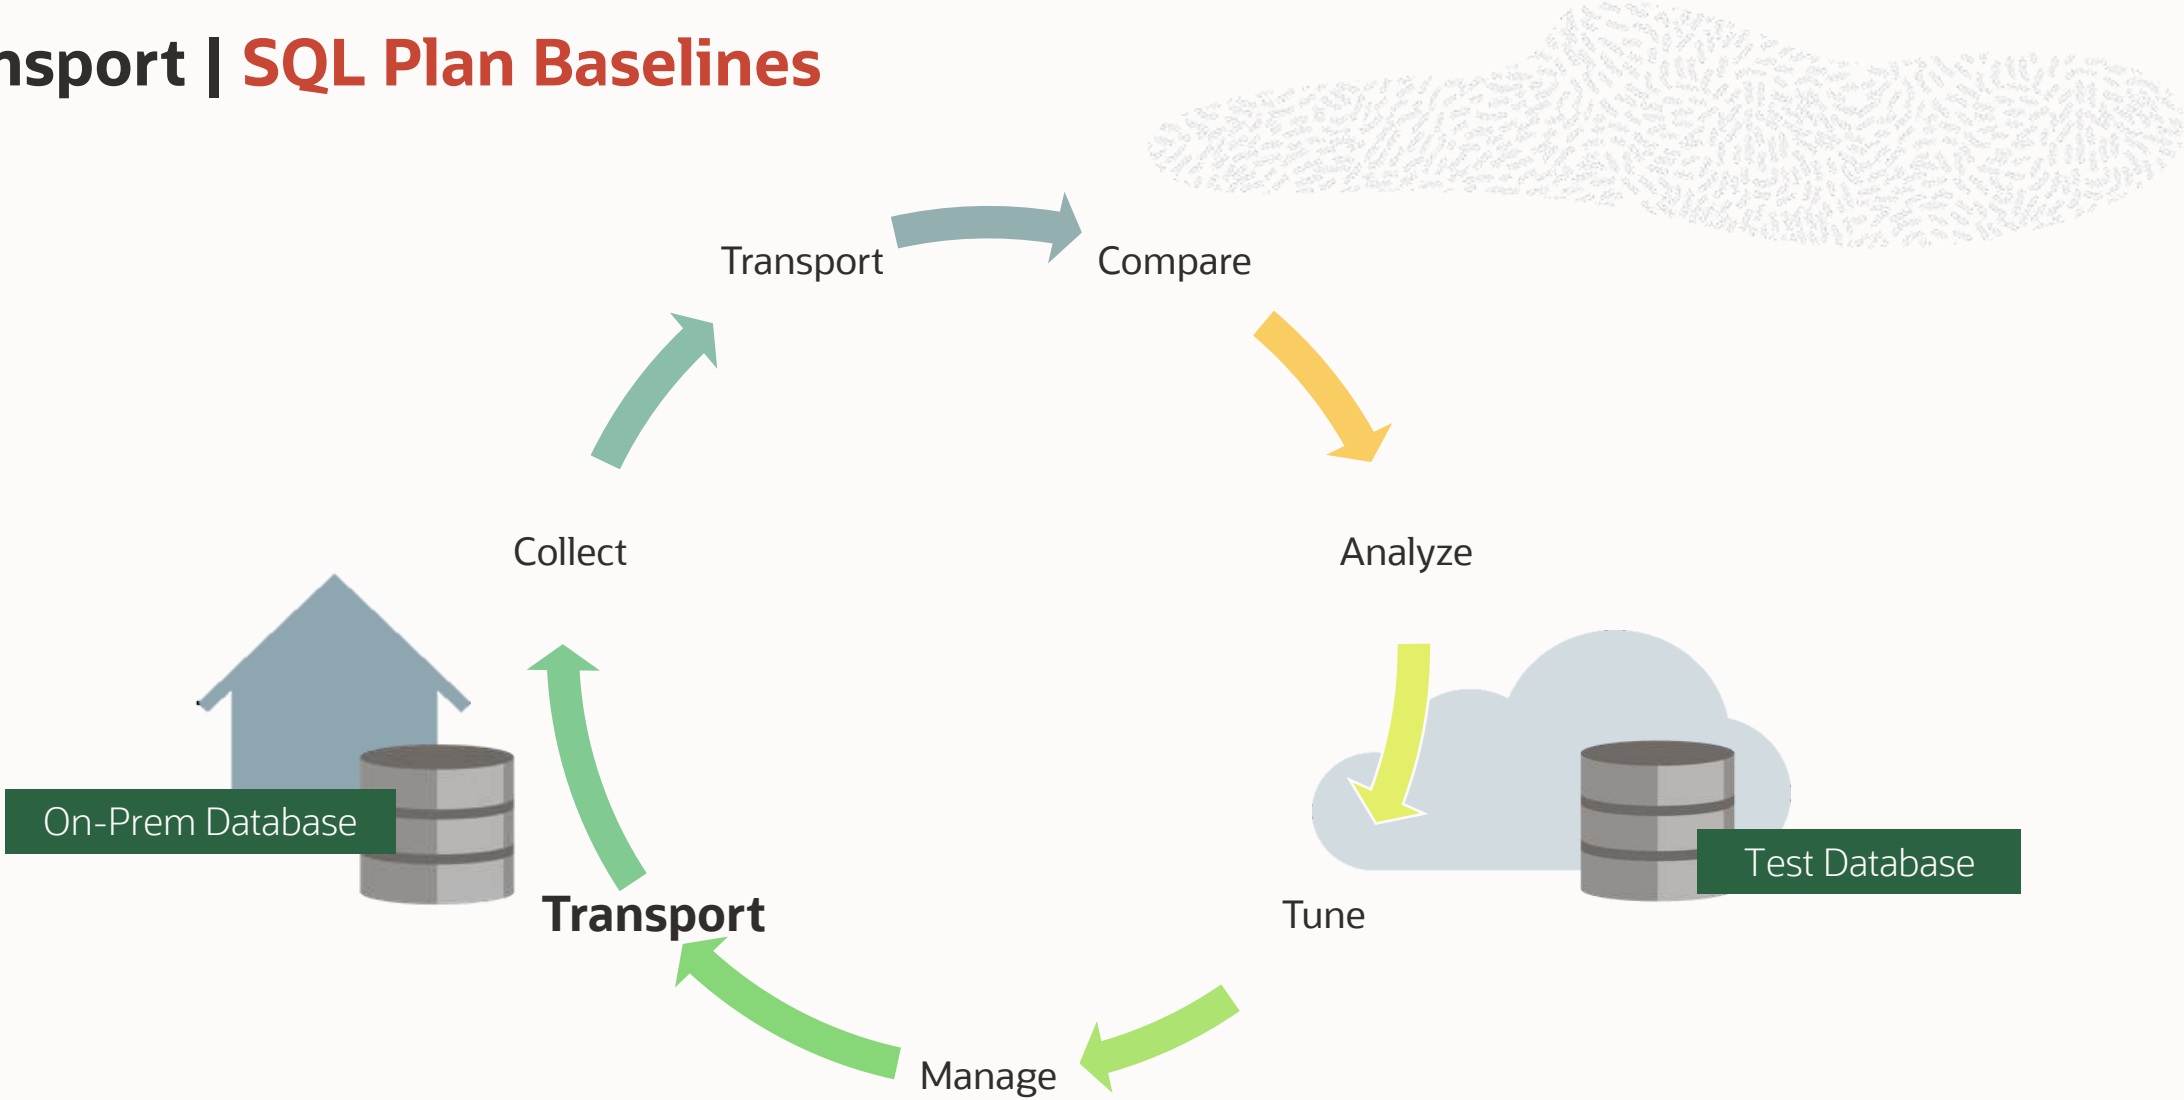
Pro tip: The function `LOAD_PLANS_FROM_SQLSET` can also *fix* the plans



Challenge #9

Transport

Transport | SQL Plan Baselines



Transporting | SQL Plan Baseline

Prepare

Extract

Transfer

Load

SQL Plan Baselines are stored in data dictionary

To transfer - information must be converted to a transportable format and stored in a **staging table**

```
SQL> BEGIN
      DBMS_SPM.CREATE_STGTAB_BASELINE (
        table_name    => 'SPB_STAGING',
        table_owner   => 'SPM');
      END;
/
```

Transporting | SQL Plan Baseline

Prepare

Extract

Transfer

Load

Select the baselines that you want to transfer

To extract the **fixed** and **accepted** plans

```
SQL> DECLARE
    l_count NUMBER;
BEGIN
    l_count := DBMS_SPM.PACK_STGTAB_BASELINE (
        table_name     => 'SPB_STAGING',
        table_owner    => 'SPM',
        enabled        => 'YES',
        fixed          => 'YES');
END;
/
```

Pro tip: You can also use
`dba_sql_plan_baselines` to find plans



Transporting | SQL Plan Baseline

Prepare

Extract

Transfer

Load

Use Data Pump to transfer that single table

```
SQL> CREATE DATABASE LINK cloud_link ... ;
```

```
$ impdp system network_link=cloud_link tables=SPM.SPB_STAGING  
...
```

Optionally,

- Export to dump file
- Transfer using `scp` or `rsync`
- Import from dump file

Transporting | SQL Plan Baseline

Prepare
Extract
Transfer

Load

Finally, load the baselines from the staging table into the data dictionary

```
SQL> DECLARE
    l_count NUMBER;
BEGIN
    l_count := DBMS_SPM.UNPACK_STGTAB_BASELINE (
        table_name => 'SPB_STAGING',
        table_owner => 'SPM');
END;
/
```

- This requires Enterprise Edition license on-prem

Pro tip: You can apply filters to limit the baselines to import



Transporting | **SQL Plan Baseline**

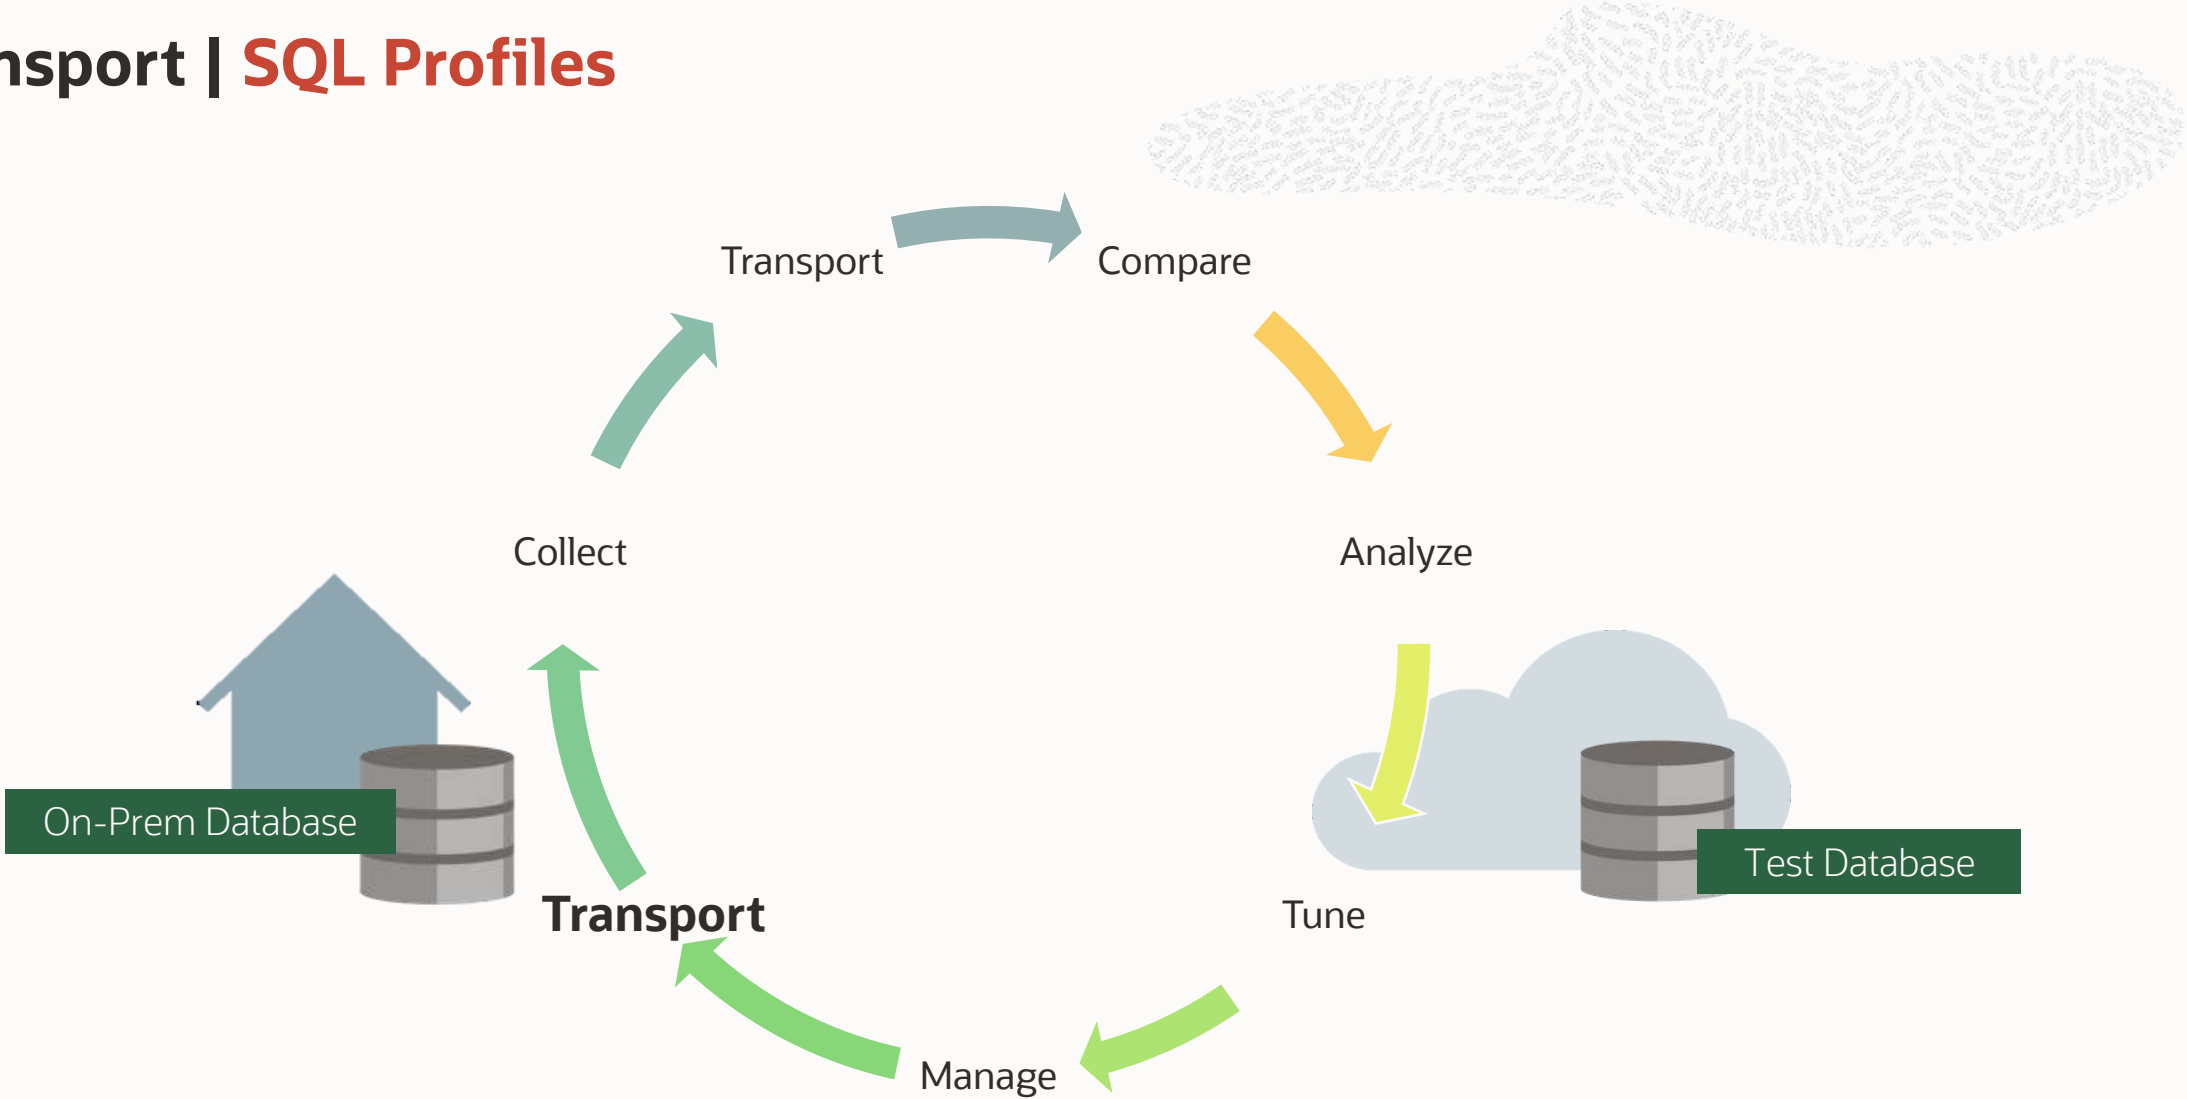
Additional resources

Documentation

[Technical Brief, SQL Plan Management in Oracle Database 19c](#)

[Database 19c, SQL Tuning Guide, Overview of SQL Plan Management](#)

Transport | SQL Profiles



Transporting | SQL Profiles

Prepare

Extract

Transfer

Load

SQL Profiles are stored in data dictionary

To transfer - profiles must be converted to a transportable format and stored in a **staging table**

```
SQL> BEGIN
      DBMS_SQLTUNE.CREATE_STGTAB_SQLPROF (
        table_name    => 'STAGING',
        table_owner   => 'SQLPROFILES');
      END;
/
```

Transporting | SQL Profiles

Prepare

Extract

Transfer

Load

Select the profiles that you want to transfer

To extract all profiles from DEFAULT category

```
SQL> BEGIN
      DBMS_SQLTUNE.PACK_STGTAB_SQLPROF (
        staging_table_name    => 'STAGING',
        staging_schema_owner => 'SQLPROFILES');
      END;
  /
```

Pro tip: You can filter on `profile_name` and `profile_category` as well



Transporting | SQL Profiles

Prepare

Extract

Transfer

Load

Use Data Pump to transfer that single table

```
SQL> CREATE DATABASE LINK cloud_link ... ;  
  
$ impdp system network_link=cloud_link \  
    tables=SQLPROFILES.STAGING ...
```

Transporting | SQL Profiles

Prepare
Extract
Transfer

Load

Finally, load the profiles from the staging table into the data dictionary

```
SQL> BEGIN
      DBMS_SQLTUNE.UNPACK_STGTAB_SQLPROF (
        staging_table_name    => 'STAGING',
        staging_schema_owner  => 'SQLPROFILES',
        replace                => TRUE);
      END;
/
```

- This requires a Tuning Pack license on-prem

Pro tip: You can load a SQL profile into the same or higher release



Transporting | SQL Profiles

Additional resources

[Documentation](#)

[Database 19c, SQL Tuning Guide, Transporting a SQL Profile](#)

[How to Move SQL Profiles from One Database to Another \(Including to Higher Versions\)
\(Doc ID 457531.1\)](#)





Photo by [Dušan veverkolog](#) on [Unsplash](#)

Want to Know More?

Webinar: Performance Stability

[Recording](#)

[Slides](#)



Photo by [Bill Jelen](#) on [Unsplash](#)

Upgrading in the Cloud

Upgrade | **Scaling**

Consider **scaling CPU** when upgrading

- Bare Metal and Exadata scales **online**
- Virtual Machine requires **downtime** (10-15 min)

Non-CDB: Default parallel 4 (max 8)





CDB\$ROOT: Default parallel 4 (runs first, serial)

PDBs: Default parallel 2 (run in parallel, as many as there are CPUs)

Pro Tip:
During upgrade I/O is of no importance



Upgrade | Shapes



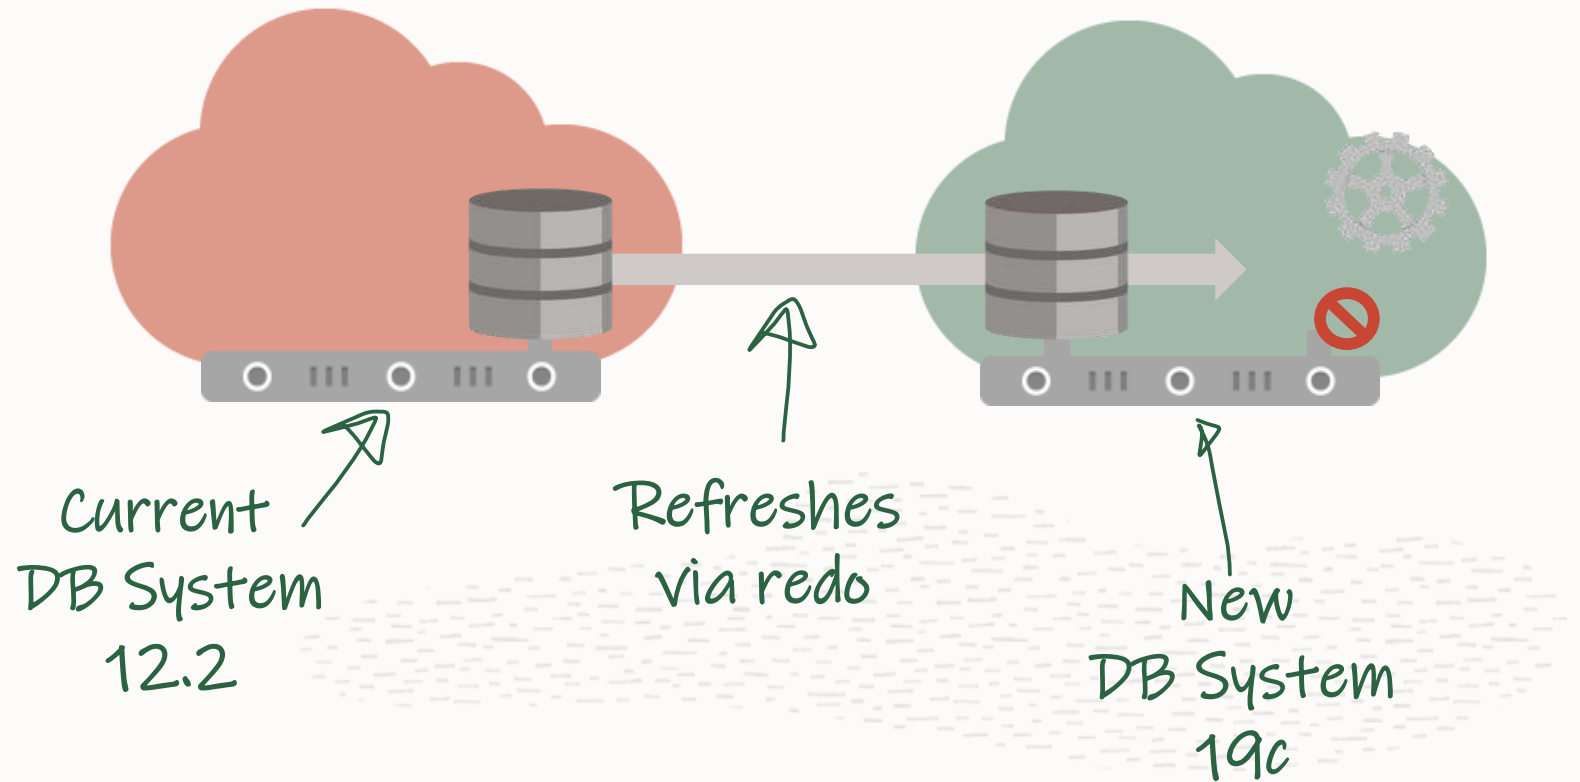
Virtual Machine | Tooling



No tooling yet, **coming soon**

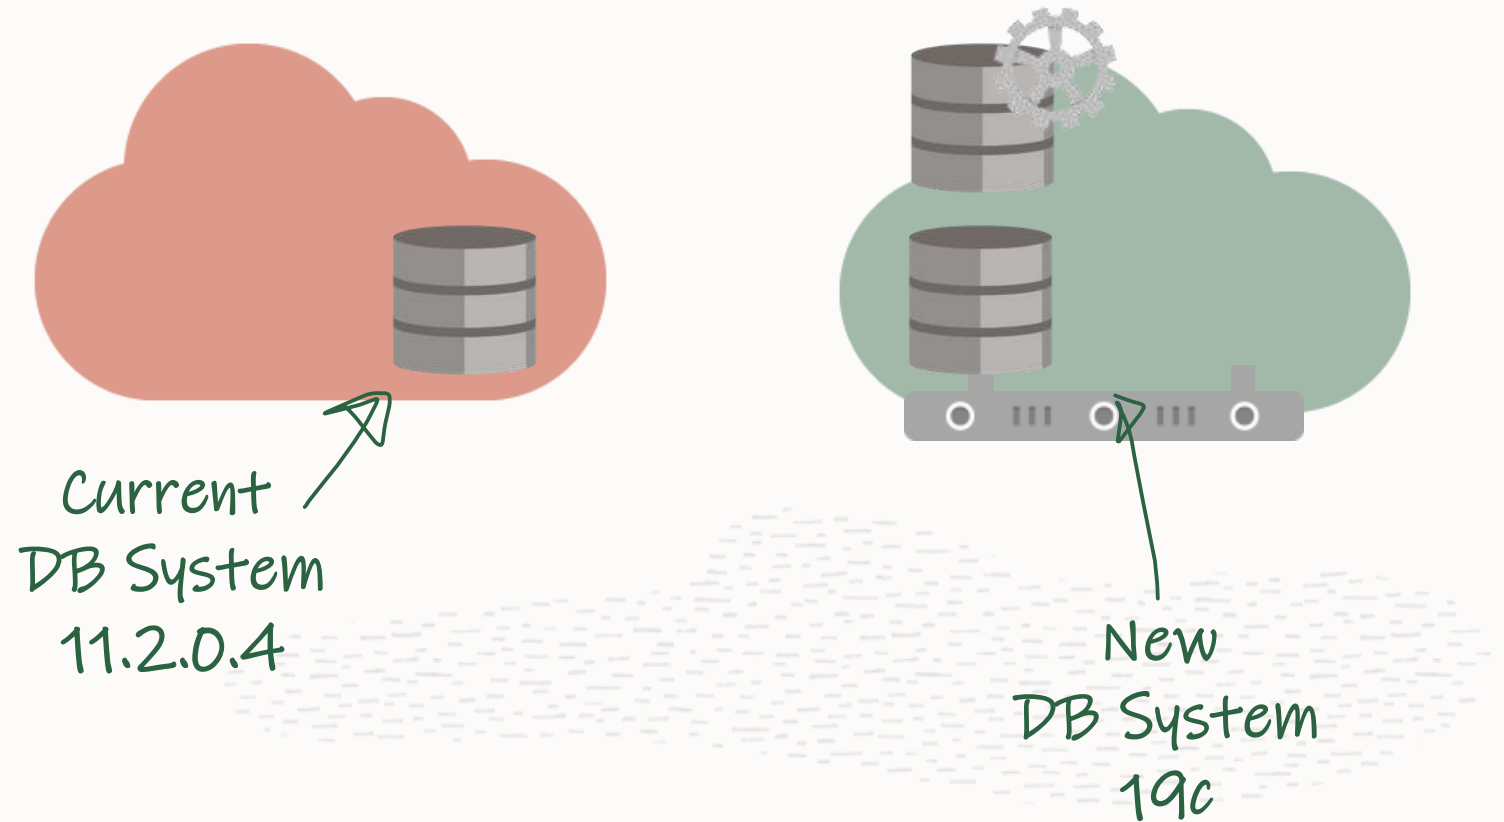
Virtual Machine | Upgrade PDB to 19c

1. Create refreshable PDB
2. Upgrade
3. Open



Virtual Machine | Upgrade Non-CDB to 19c

1. Move to new host
2. Upgrade
3. Plug in



Virtual Machine | **Blog Post Series**



Introduction

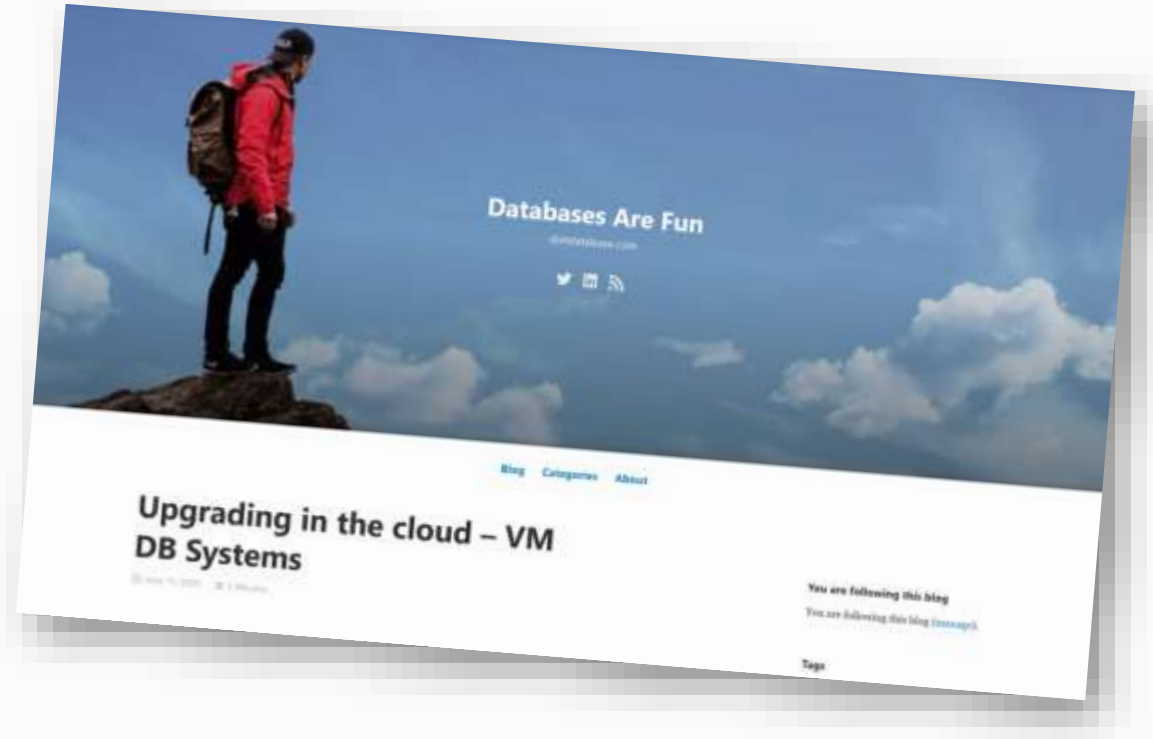
12.2.0.1 PDB to 19c

What about downgrade?

Transfer Speed

11.2.0.4 to 19c (simple)

11.2.0.4 to 19c (minimal downtime)



Bare Metal | Tooling



No tooling yet, **coming soon**



Bare Metal | Grid Infrastructure

Check your Grid Infrastructure version

```
$ crsctl query crs activeversion  
Oracle Clusterware active version on the cluster is [19.0.0.0.0]
```

Check [documentation](#) for database support:

” You can use *Oracle Database 19c*, *18c*, *Oracle Database 12c releases 1 and 2*, and *Oracle Database 11g release 2 (11.2.0.3 or later)* with *Oracle Grid Infrastructure 19c*.

No Grid Infrastructure support? **Order new system!**

- Not possible to upgrade Grid Infrastructure

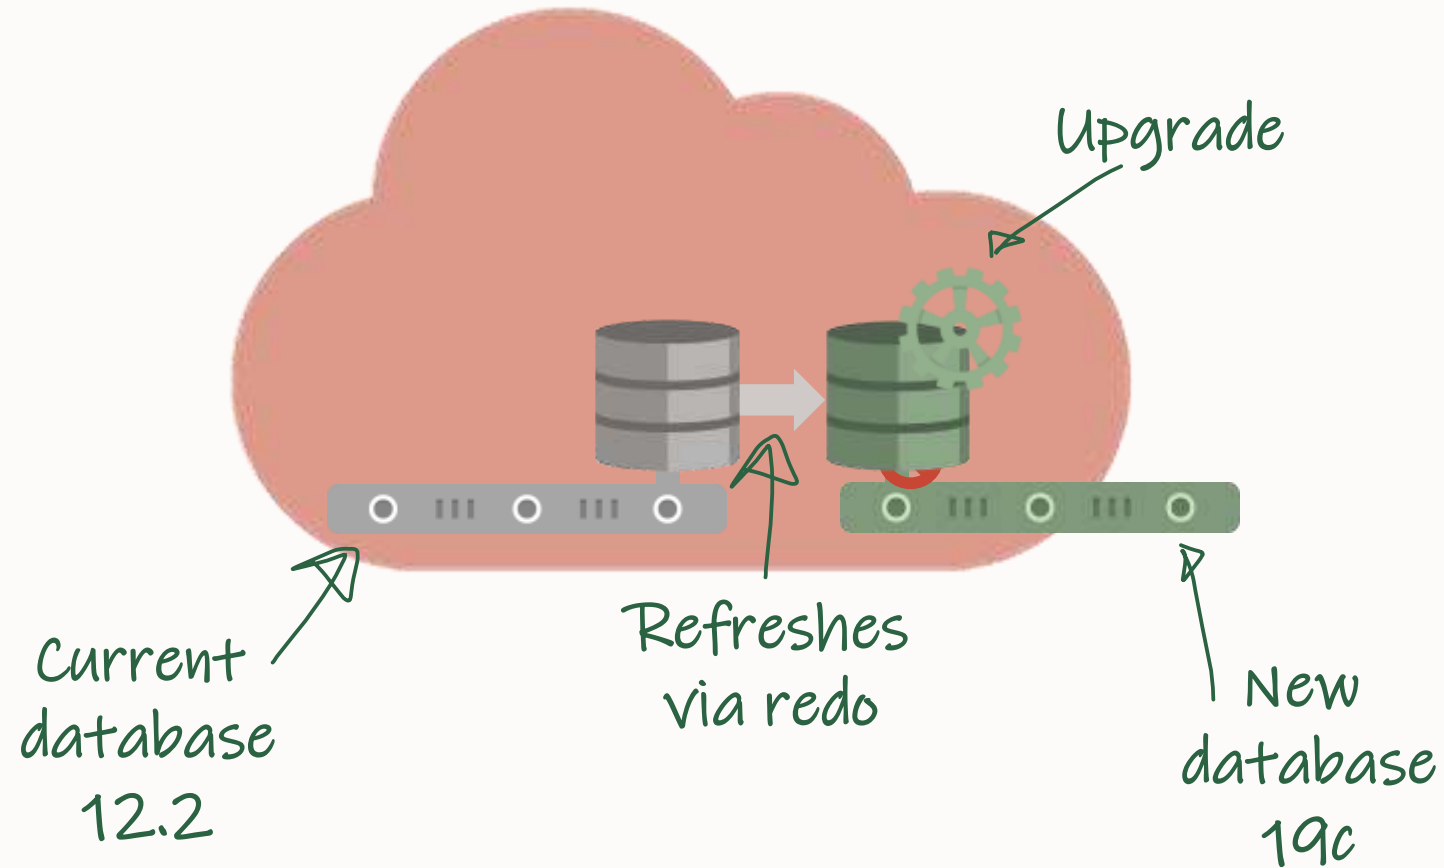


Bare Metal | Upgrade PDB to 19c

Create refreshable PDB

Upgrade

Open



Upgrade | Exadata Cloud Service



★ [Upgrading to 19c Oracle Database on Exadata Cloud Service \(Doc ID 2628228.1\)](#)

In this Document

[Goal](#)

[Solution](#)

[Overview](#)

[Database 19c Upgrade Prerequisites](#)

[Step 1.1 - Validate Upgrade Steps on Test](#)

[Step 1.2 - Golden Gate Supported Versions](#)

[Step 1.3 - Oracle Application Express\(APEX\) Supported Versions](#)

[Step 1.4 - Validate Minimum Software Requirements](#)

[Step 1.4.1 - Required Exadata Database Server Software version](#)

[Step 1.4.2 - Required Grid Infrastructure version](#)

[Step 1.5 - Install the latest Cloud Tooling](#)

[Step 1.6 - Database Prerequisites](#)

[Step 1.7 - Create Target Database 19c ORACLE HOME](#)

[Step 1.7.1 - Create an empty 19c ORACLE HOME](#)

[Step 1.7.2 - Apply any one-off or merge patches](#)

[Step 1.7.3 - Validate the new 19c ORACLE HOME version](#)

[Step 1.8 - Run Pre-Checks Information Tool on Primary DB](#)

[MOS Note: 2628228.1](#)
[Upgrading to 19c Oracle Database on Exadata Cloud Service](#)



Upgrade | Exadata Cloud Service

1. Check (and upgrade) Exadata Server software (DomU) to ≥ 19.2
2. Upgrade Cloud Tooling
3. Grid Infrastructure upgrade to GI 19c
4. Create empty 19c database home
5. Database upgrade to Oracle 19c
6. **Optional:** Plugin into a CDB



Upgrade | Exadata Cloud Service

Upgrade Grid Infrastructure to Oracle GI 19c

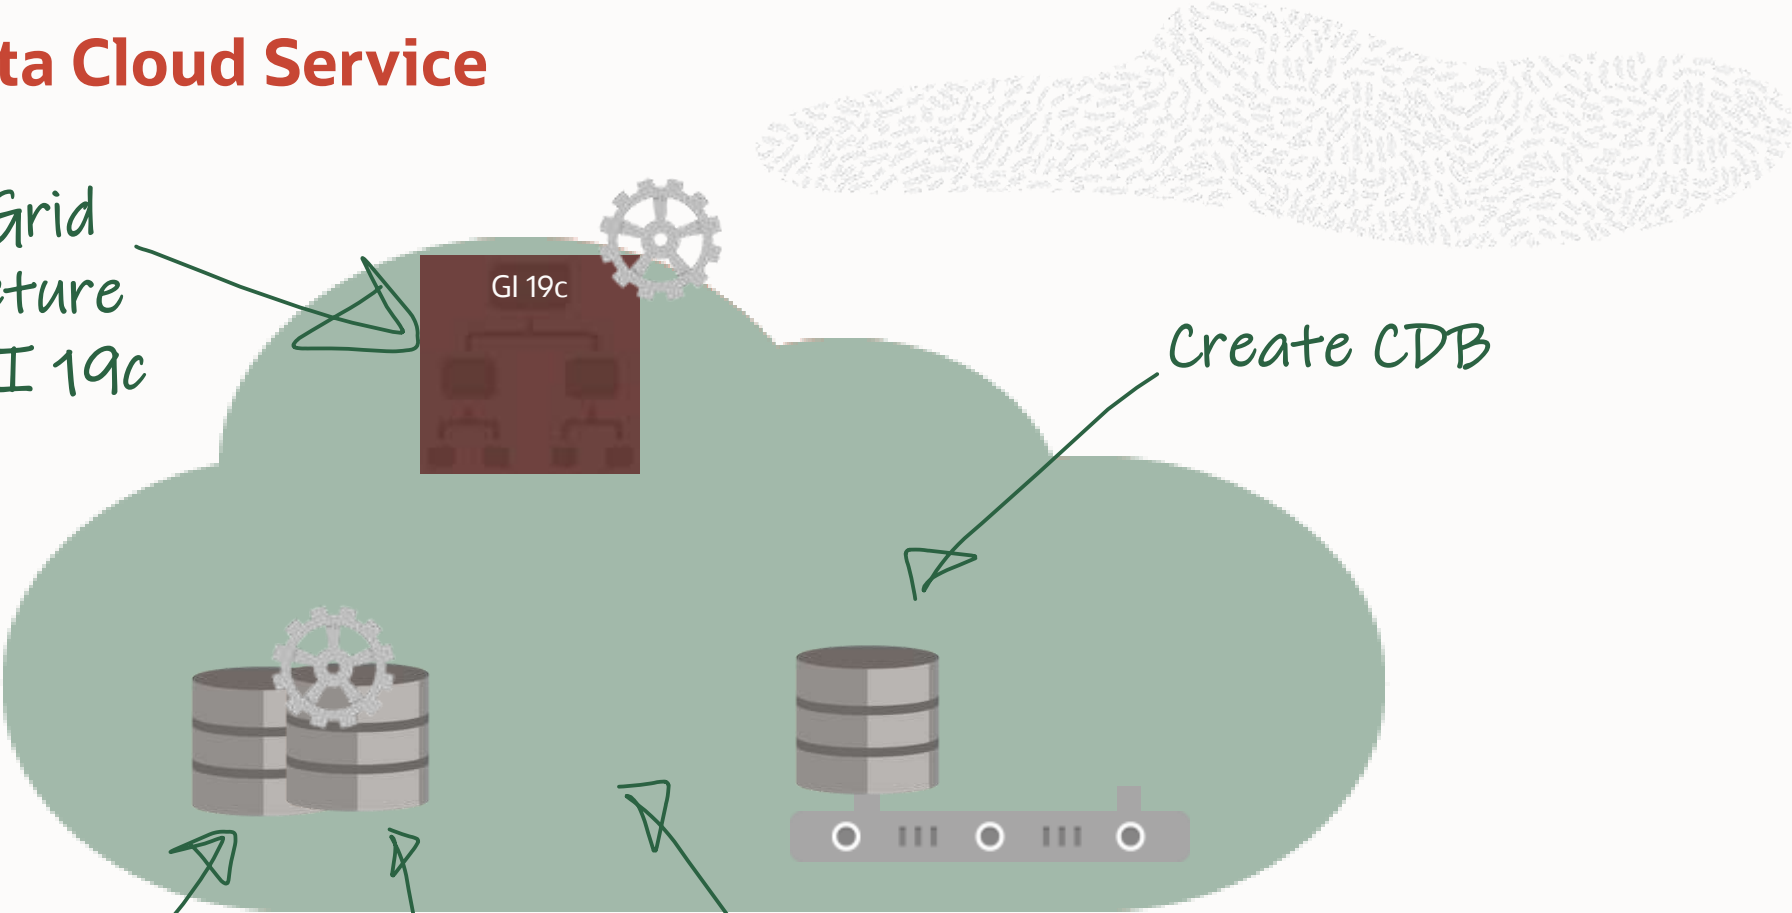


Create CDB

Current DB System 11.2.0.4 - 18c

DB upgrade to Oracle 19c

Plugin



Upgrade | ExaCS – Best Practices

[MOS Note: 2628228.1](#)

[Upgrading to 19c Oracle Database on Exadata Cloud Service](#)

- Supplement to the Oracle Database 19c Upgrade Guide

[MOS Note: 884522.1](#)

[Preupgrade.jar Download](#)

- Always download the newest preupgrade.jar beforehand
- Copy it to the new Oracle 19c home on ExaCS

```
cp preupgrade_19_cbuid_7_1f.zip $NEW_ORACLE_HOME/rdbms/admin/  
cd $NEW_ORACLE_HOME/rdbms/admin/  
unzip preupgrade_19_cbuid_7_1f.zip
```

Upgrade | ExaCS – Best Practices

Deploy the most recent Release Update

- List the available images for download:

```
[root@exacs-exa1 ~]# dbaascli  
DBAAS> cswlib list
```

- Download the image

```
DBAAS> cswlib download --version 19000 --bp JUL2020
```

- Activate the image

```
DBAAS> cswlib activateBP --version 19000 --bp JUL2020
```

- Deploy a new home with OCI Console or with OCI CLI

- <https://blog.alexplyth.io/2020/05/13/oracle-exacs-creating-oracle-homes-with-older-ru-bps/>

Upgrade | ExaCS – Best Practices

Pay attention to pre- and post-upgrade tasks

Tooling creates a GRP automatically

- Name: `BEFORE#UPGRADE#<integer>#`
- For the standby as well

Wallet must be set to **auto-login**

Remove all underscores during upgrade

Downgrades are supported by tooling

- Don't change `COMPATIBLE` if you'd like to preserve the downgrade option

Upgrade | ExaCS - Post-upgrade tasks

Observe the following after the upgrade:

- RMAN catalog upgrade
- Check for the `SQLNET.ALLOWED_LOGON_VERSION` parameter behavior
- Regathering fixed objects statistics with `DBMS_STATS` after the database has been warmed up
- Upgrade from Oracle 12.1.0.2 only:
 - Transfer Unified Audit records after upgrade
- Upgrading the Time Zone File version
- Check the post upgrade fixups log

Pro Tip:
Check [Chapter 6 Post-Upgrade Tasks for Oracle Database](#) for up-to-date information



Upgrade | ExaCS – Conversion non-CDB to PDB

Create the CDB at first

Run Plugin Compatibility check

```
SET SERVEROUTPUT ON

DECLARE
compatible CONSTANT VARCHAR2(3) := CASE DBMS_PDB.CHECK_PLUG_COMPATIBILITY(
pdb_descr_file => '/home/oracle/PDB12201.xml', pdb_name => 'PDB12201')
WHEN TRUE THEN 'YES'
ELSE 'NO' END;

BEGIN
DBMS_OUTPUT.PUT_LINE(compatible);
END;
/
```

```
SELECT time, name, cause, type, message, status, action
FROM pdb_plug_in_violations
WHERE name = '<pdb_name>' AND status not in ('RESOLVED');
```

Upgrade | ExaCS – Conversion non-CDB to PDB

Create the CDB at first

Run Plugin Compatibility check

Test the conversion:

```
[root@exacs-exa1 ~]# dbaascli  
DBAAS> database convert_to_pdb --dbname <database_name> --cdbname  
<db_unique_name> --precheck
```

Perform the conversion:

```
DBAAS> database convert_to_pdb --dbname <database_name> --cdbname  
<db_unique_name>
```

Check PDB status:

```
SELECT name, open_mode, restricted FROM v$pdb WHERE name='<pdb_name>';
```

Upgrade | ExaCS – Rename a PDB

Plugin will keep the database's name as PDB name

Rename a PDB in ExaCS:

```
ALTER PLUGGABLE DATABASE KARL121 CLOSE INSTANCES=ALL
ALTER PLUGGABLE DATABASE KARL121 UNPLUG INTO '/tmp/karl121.xml' ENCRYPT
USING <a-secret-passphrase>;
DROP PLUGGABLE DATABASE KARL121 KEEP DATAFILES;
CREATE PLUGGABLE DATABASE MARIE19 USING '/tmp/karl121.xml' NOCOPY KEYSTORE
IDENTIFIED BY <keystore_password> DECRYPT USING <a-secret-passphrase>;
ALTER PLUGGABLE DATABASE MARIE19 OPEN INSTANCES=ALL;
ALTER PLUGGABLE DATABASE MARIE19 SAVE STATE;
```


PLEASE NOTE: FILES ARE STILL STORED IN THE ORIGINAL LOCATION

Photo by [Carolina Pimenta](#) on [Unsplash](#)

Finally ...

It is time to wrap up

Migration | **Keep It Simple**



Wrapping Up | There Are More



- RMAN Restore + Incremental Backup
- RMAN Duplicate From Active Database
- Cross-platform Transportable Database
- Transportable PDB
- Cross-platform Transportable PDB
- Refreshable PDB
- PDB Relocate
- PDB Remote Clone
- PDB Unplug/Plug
- SQL Developer Database Copy
- Database Links and CTAS
- ...



Further Reading



oracle.com/goto/move

Database / Oracle Database Cloud Migration

Move to the Oracle Cloud

Move your Database to the Oracle Cloud.

Simple & Efficient

Oracle automated tools make it seamless to move your on-premises database to the Oracle Cloud with virtually no downtime. Using the same technology and standards on-premises and in the Oracle Cloud, you can facilitate the same products and skills to manage your cloud-based Oracle Databases as you would on any other platform.

Flexible

You can directly migrate your Oracle Database to the Oracle Cloud from various source databases into different target cloud deployments depending on your requirements and business needs. A well-defined set of tools gives you the flexibility to choose the method that best applies to your needs.

Cost Effective

The same flexibility that lets you directly migrate your Oracle Database to the Oracle Cloud is applied to finding the most cost effective solution for the purpose and duration of the migration. Even if the automated tools determine that an Oracle licensable product should be used to optimize your migration, Oracle will provide a cost neutral solution.

Highly Available & Scalable

The tight integration of all migration tools with the Oracle Database lets you maintain control and gain better efficiency when moving your databases to the Oracle Cloud, while the Maximum Availability Architecture (MAA)-approved tools as well as Zero Downtime Migration (ZDM)-based migrations ensure that your migration is handled as smoothly as possible.



Migration | **More Information**

[Watch the recordings](#)

[Get the slides](#)



Thank you!

