ORACLE

Migrating Very Large Databases

Virtual Classroom Series

Photo by Rhys Moult on Unsplash

ROY SWONGER
Vice President
Database Upgrade, Utilities & Patching

royfswonger

@royfswonger




MIKE DIETRICH
Distinguished Product Manager
Database Upgrade and Migrations

mikedietrich

@mikedietrichde

https://mikedietrichde.com




DANIEL OVERBY HANSEN
Senior Principal Product Manager
Cloud Migrations

dohdatabase

@dohdatabase

https://dohdatabase.com




RODRIGO JORGE
Senior Principal Product Manager
Database Patching and Upgrade

rodrigoaraujorge

@rodrigojorgedba

https://dbarj.com.br/en

Webinar | Get The Slides

https://MikeDietrichDE.com/slides

Webinar | Upcoming

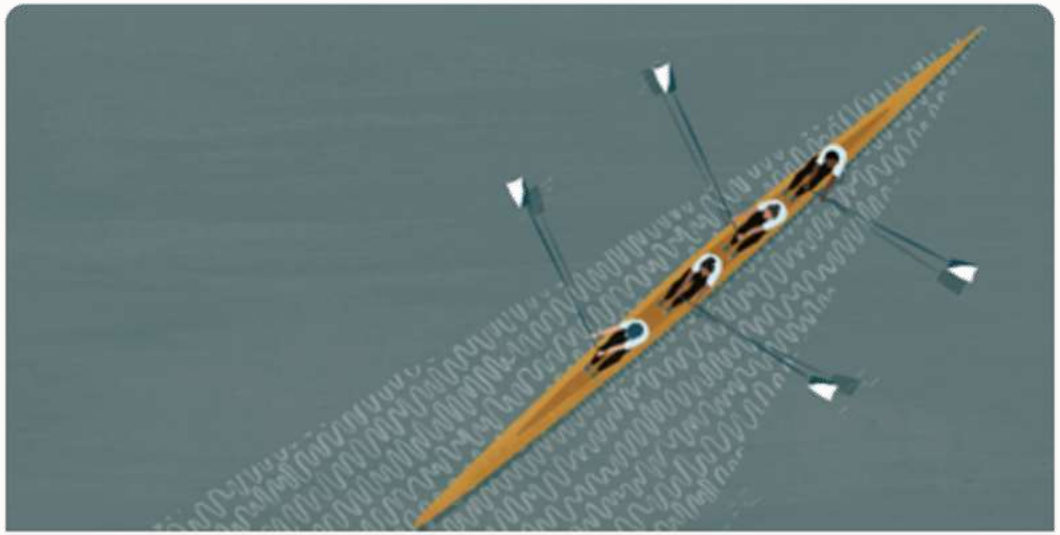

Data Pump Extreme - Deep Dive with Development

March 3, 2022 | 09:00 GMT / 10:00 CET / 11:00 EET / 13:00 GST

Duration: 120 mins

SCAN ME

https://go.oracle.com/LP=114938?elqCampaignId=302203

105 minutes – Feb 4, 2021

Episode 2

AutoUpgrade to Oracle Database 19c

115 minutes – Feb 20, 2021

Episode 3

Performance Stability, Tips and Tricks and Underscores

120 minutes – Mar 4, 2021

Episode 4

Migration to Oracle Multitenant

120 minutes – Mar 16, 2021

Episode 5

Migration Strategies – Insights, Tips and Secrets

120 minutes – Mar 25, 2021

Episode 6

Move to the Cloud – Not only for techies

115 minutes – Apr 8, 2021

Episode 7

Cool Features – Not only for DBAs

110 minutes – Jan 14, 2021

Episode 8

Database Upgrade Internals – and so much more

110 minutes – Feb 11, 2021

Episode 9

Performance Testing Using the Oracle Cloud for Upgrades and Migrations

90 minutes – May 19, 2021

NEW Episode 10

How Low Can You Go? Minimal Downtime Upgrade Strategies

100 minutes – Oct 26, 2021

Recorded Web Seminars

https://MikeDietrichDE.com/videos

INTRODUCTION

What is very large?

It depends …

Source and target system's Endianness?

Endianness | The Basics

Big-endian

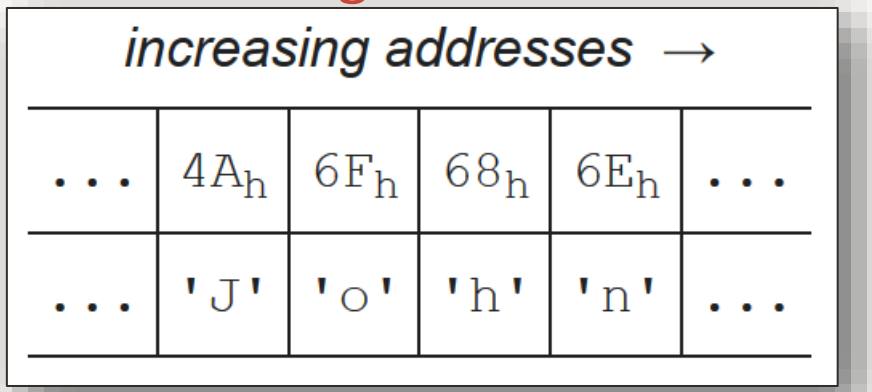

Little-endian

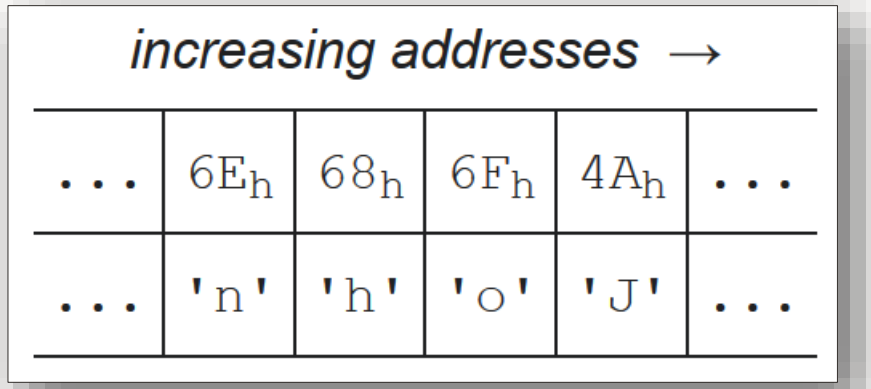

Source: https://en.wikipedia.org/wiki/Endianness

Endianness | Platforms

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

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

Endianness | Platform IDs

```
SELECT PLATFORM_ID, PLATFORM_NAME, ENDIAN_FORMAT
  FROM V$TRANSPORTABLE_PLATFORM
  ORDER BY PLATFORM_ID;


PLATFORM_ID PLATFORM_NAME                                ENDIAN_FORMAT
----------- -------------------------------------------- --------------
          1 Solaris[tm] OE (32-bit)                      Big
          2 Solaris[tm] OE (64-bit)                      Big
          3 HP-UX (64-bit)                               Big
          4 HP-UX IA (64-bit)                            Big
          5 HP Tru64 UNIX                                Little
          6 AIX-Based Systems (64-bit)                   Big
...
```

Endianness | Common Scenario

Big Endianness to little Endianness platforms

Upgrade included?

Downtime requirement?

Conversion to Multitenant?

Which one is the best technique?

It depends …

Migration | Techniques

Data Pump	Data Guard	FTEX
RMAN	TTS	OGG

same

ENDIAN

migration technique

- **RMAN**
- Trans-portable
- Data Guard
- Golden Gate

Incremental | Concept

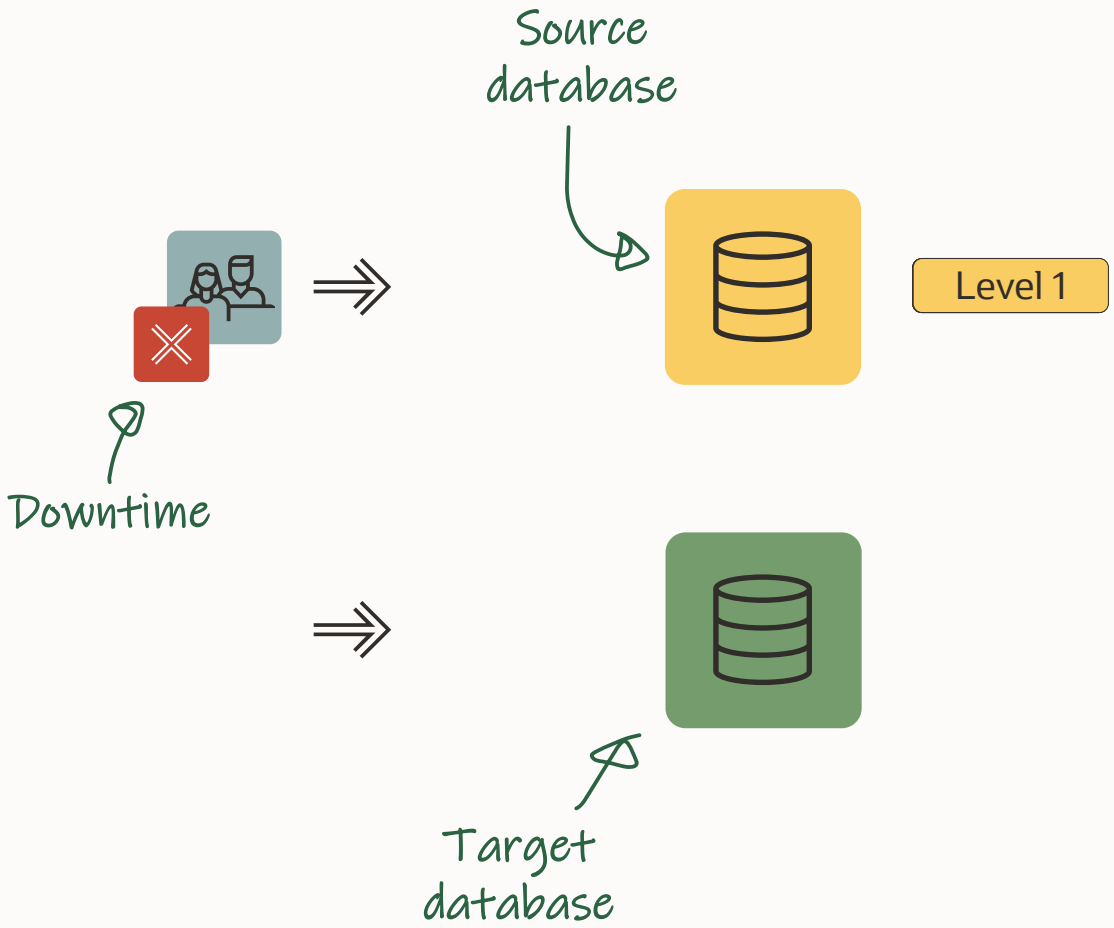


```
RMAN> BACKUP INCREMENTAL
      LEVEL 1 ...
```

```
RMAN> RECOVER DATABASE ...
```

Incremental | Benefits

- Simple and easy
- Well-known process
- Use existing backups
- Independent of file system, raw devices and ASM
- Some cross-platform capabilities

Incremental | Procedure

SOURCE DATABASE	TARGET DATABASE
`backup incremental level 0 database ... ;`	
	`restore database;`
`backup incremental level 1 database ... ;`	
	`recover database;`
DOWNTIME	
`backup incremental level 1 database ... ;`	
	`recover database;`
	`alter database open resetlogs;`

Incremental backups are useful when there is no SQL*Net connectivity between source and target

Incremental backups are useful when source database release can't be installed on target host

Pro tip: Any release of RMAN can restore and recover a previous release backup

Block Change Tracking is recommended to speed up incremental backups

Pro tip: BCT is an Enterprise Edition feature, but requires Active Data Guard if enabled on standby database

RMAN Compression can significantly reduce the size and duration of the backup

Pro tip: Most compression algorithms require Advanced Compression Option

Secure your RMAN backup with TDE Tablespace Encryption or RMAN Encryption

Pro tip: Requires Advanced Security Option

Using multisection backups is important in databases with bigfile tablespaces

Pro tip: The keyword `SECTION SIZE` controls the use of multisection backups

To recover the latest changes use an incremental backup or archive logs

What about upgrade?

RMAN can restore and recover from backups made in a previous release

Pro tip: The database must be opened immediately in `UPGRADE` mode and upgraded

Incremental | Upgrade

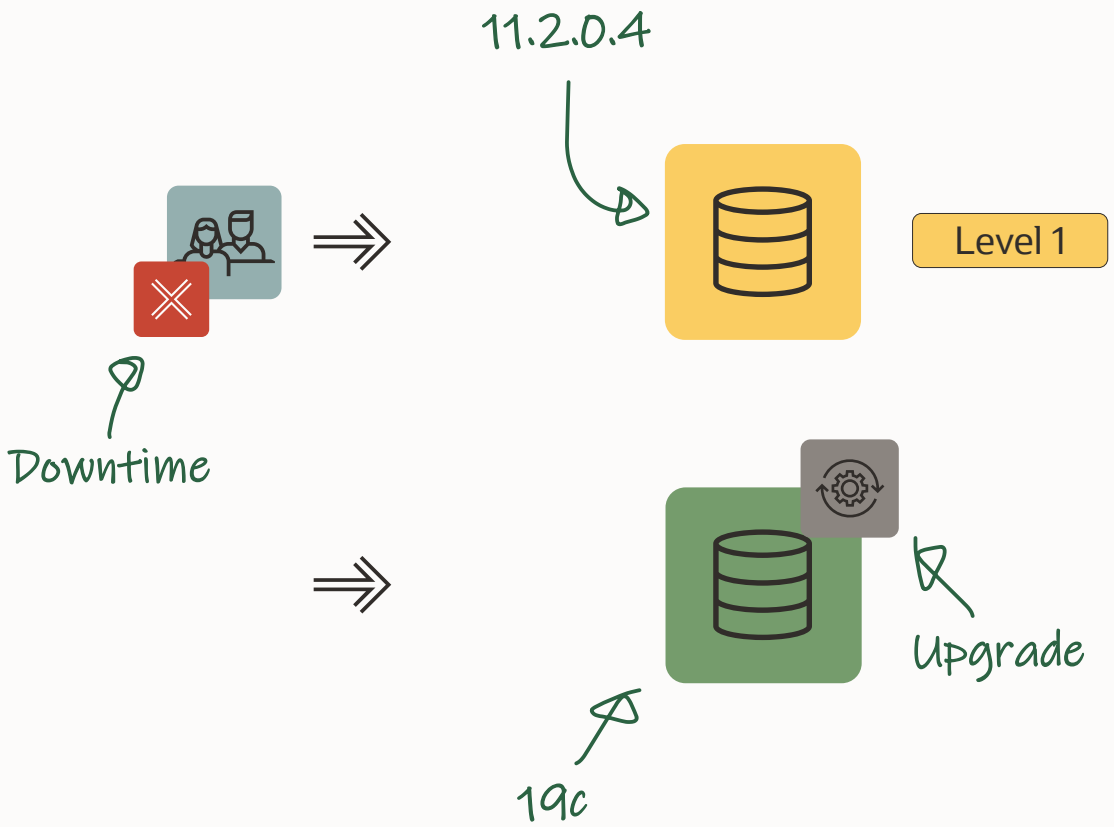


```
RMAN> BACKUP INCREMENTAL
      LEVEL 1 ...
```

```
RMAN> RECOVER DATABASE ...

SQL> ALTER DATABASE OPEN RESETLOGS UPGRADE;
```

Incremental | AutoUpgrade

SOURCE DATABASE	TARGET DATABASE
`backup incremental level 0 database ... ;`	
	`restore database;`
`java -jar autoupgrade.jar -mode analyze`	
DOWNTIME	
`java -jar autoupgrade.jar -mode analyze`	
`java -jar autoupgrade.jar -mode fixups`	
`backup incremental level 1 database ... ;`	
	`recover database;`
	`alter database open resetlogs upgrade;`
	`java -jar autoupgrade.jar -mode upgrade`

What about PDB conversion?

Convert to PDB after migration (and upgrade) using `noncdb_to_pdb.sql`

Cloning a non-CDB directly
into a CDB (NON$CDB cloning)
is not recommended for large databases

Can you offload the work from the source database?

Yes, you can perform the backups on a standby database

Or simly re-use any backups that are being taken already

Your target database must
be protected by Data Guard?

You can restore source data files on to future standby database in advance

Incremental | Data Guard

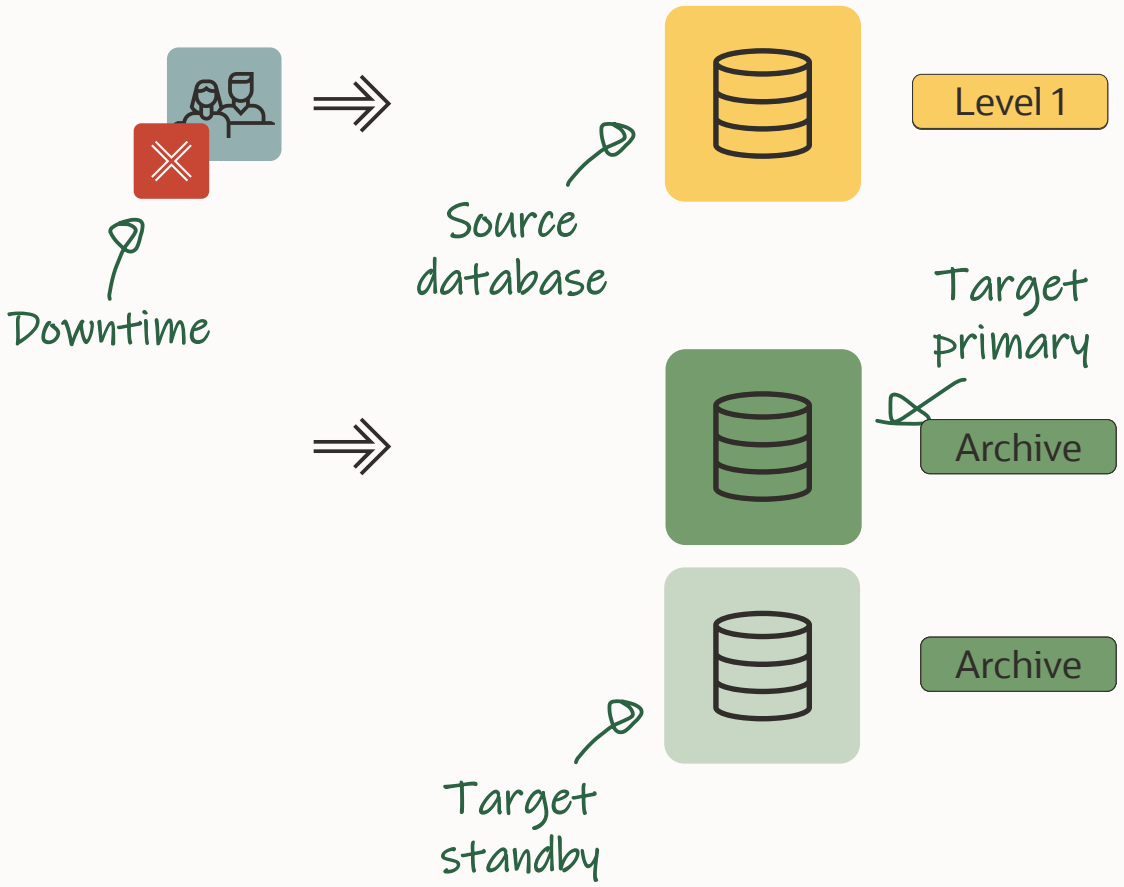


```
SQL> ALTER DATABASE OPEN RESETLOGS;
```

Configure:
- redo transport
- Redo apply

```
RMAN> RECOVER DATABASE ...
```

Your target database must
have a valid backup before go-live?

The backup pieces used by the migration, can be used for disaster recovery as well

Incremental | Backup

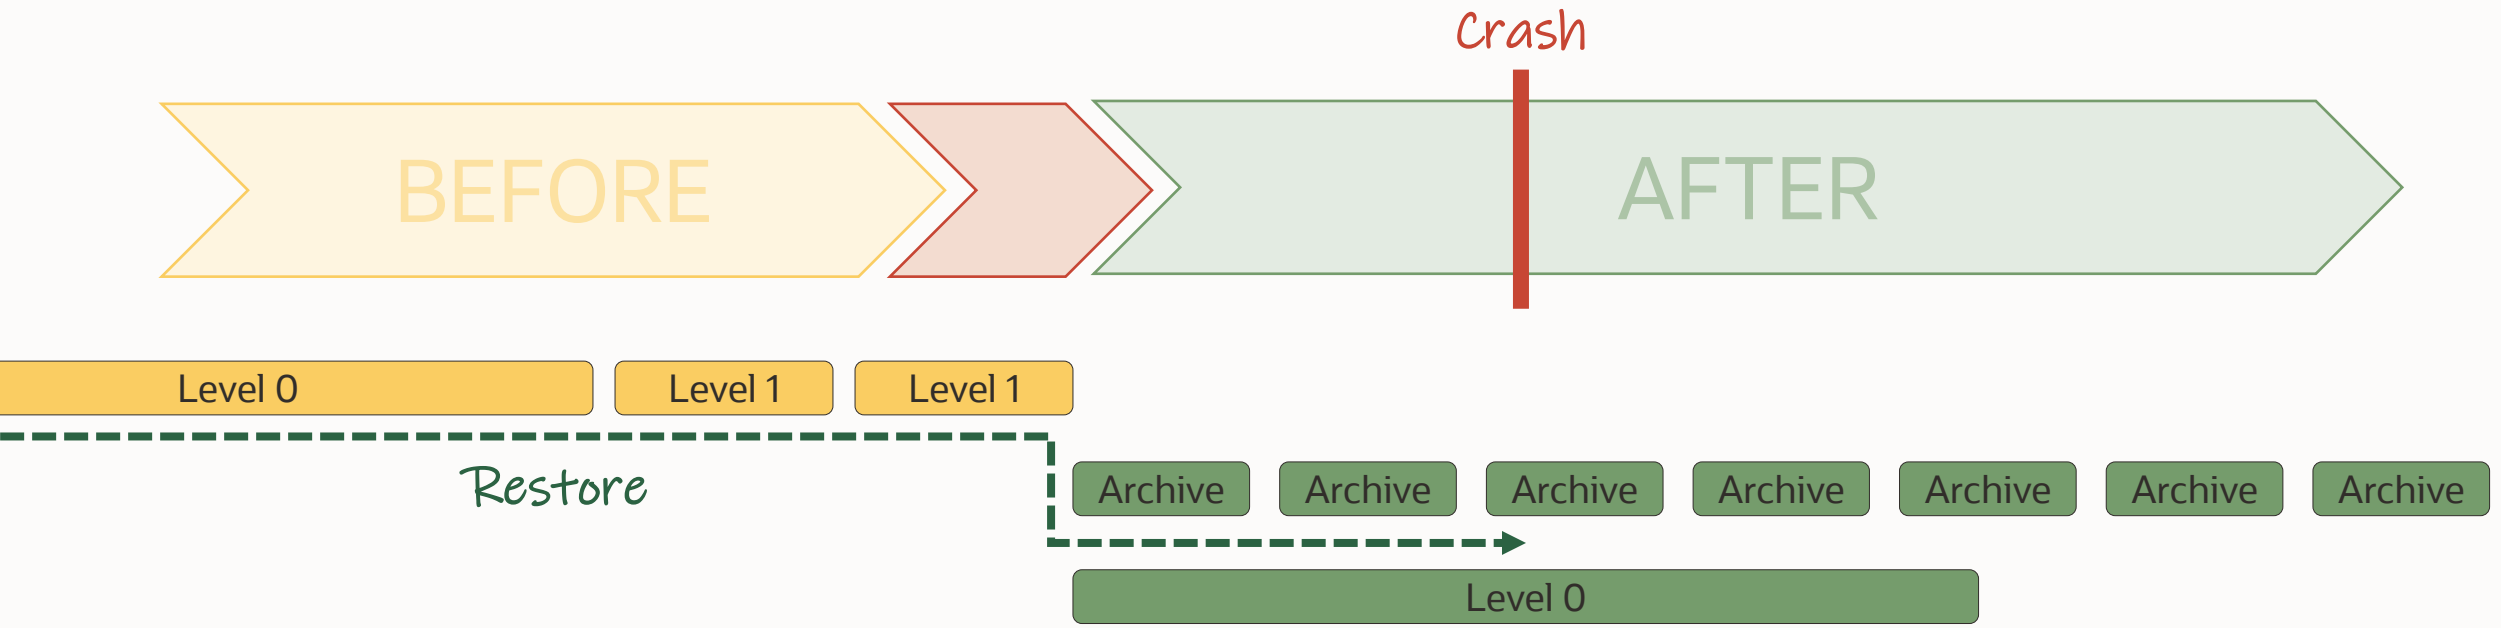

It works even if you upgraded the database

It does not work if you also converted to a PDB

Following a PDB conversion, new backups of the data files are required before go-live

Your target database must be RAC?

No problem, you can even restore a single instance to a RAC

To make recovery as easy as possible,
use shared storage as much as possible

Incremental | RAC

- Backups on shared storage enables multi-instance recovery
- Recovery is easier with SPFile and password file on shared storage
- For encrypted databases also place keystore on shared storage

How about your fallback plan?

Incremental | Fallback

- To roll back (before go live):
  Source environment is preserved
- To fall back (after go live):
  Redo process in reverse order
- Unless database was upgraded:
  Downgrade
- Unless database was converted:
  Data Pump and GoldenGate

same

ENDIAN

migration technique

- RMAN
- Trans-portable
- **Data Guard**
- Golden Gate

Data Guard | Concept

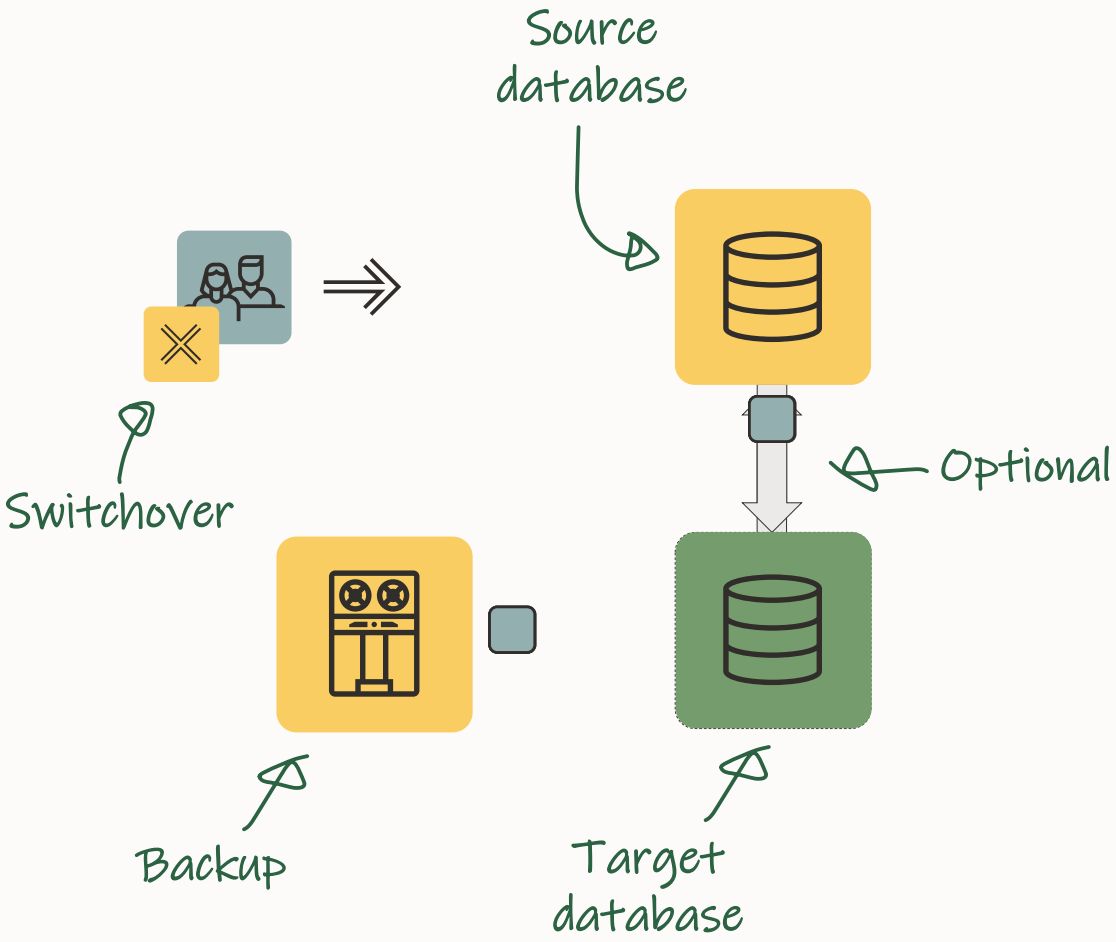

Configure:
- Redo transport
- Redo apply

```
RMAN> RESTORE STANDBY CONTROLFILE ... ;
RMAN> RESTORE DATABASE ... ;
RMAN> RECOVER DATABASE UNTIL ... ;
```

Data Guard | Benefits

- Preferred solution
- Well-known, simple and easy
- Seamless switchover
  with properly configured application
- Some cross-platform capabilities

Source Oracle Home
is needed on target host

Data Guard | Platform Certification

1. Migrate database to new hardware and upgrade from 11.2.0.4 to 19c

2. Target host must run Oracle Linux 8

3. To use Data Guard, you must install Oracle Database 11.2.0.4 on target host

4. Oracle Database 11.2.0.4 is not certified on Oracle Linux 8

Data Guard not possbile

Do you need the same platform on source and target host?

Data Guard | Heterogeneous

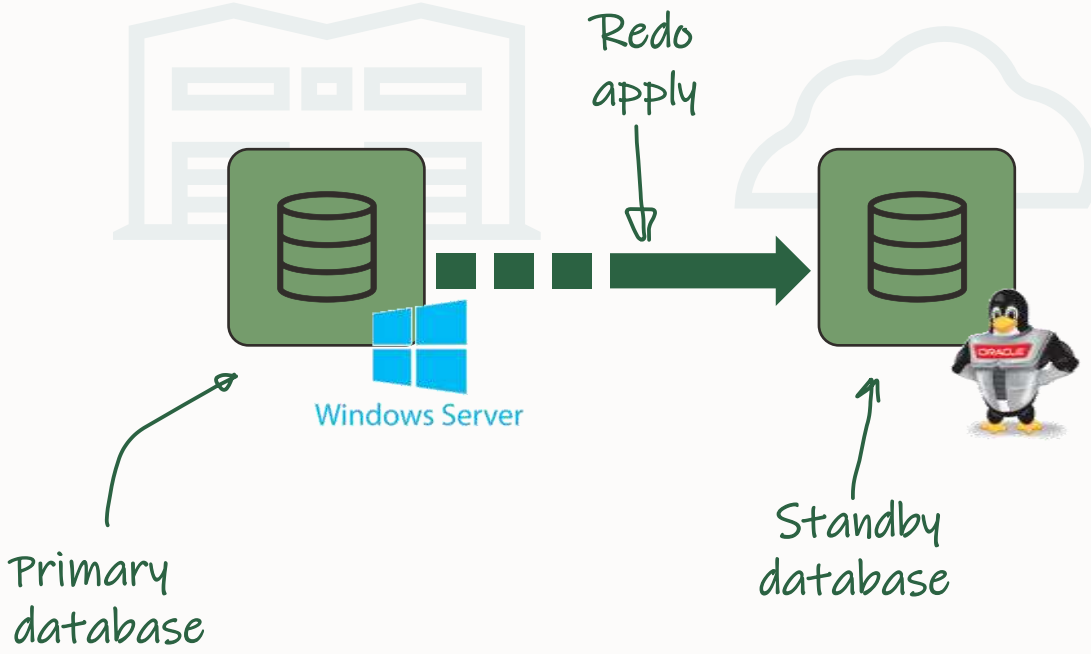

Little Endian | Linux Standby

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
```

SQL*Net connectivity is required between source and target database

RMAN Compression can reduce the size and duration of the backup significantly

Pro tip: Most compression algorithms require Advanced Compression Option

Secure your RMAN backup with TDE Tablespace Encryption or RMAN Encryption

Pro tip: Requires Advanced Security Option

Using multisection backups is important in databases with bigfile tablespaces

Pro tip: The keyword `SECTION SIZE` controls the use of multisection backups

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

How To Calculate The Required Network Bandwidth
Transfer Of Redo In Data Guard Environments (Doc ID 736755.1)

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

Consider compressing redo when using very slow connections

Pro tip: Requires Advanced Compression Option

What about upgrade?

Downtime is required to upgrade database after switchover

Data Guard | Upgrade

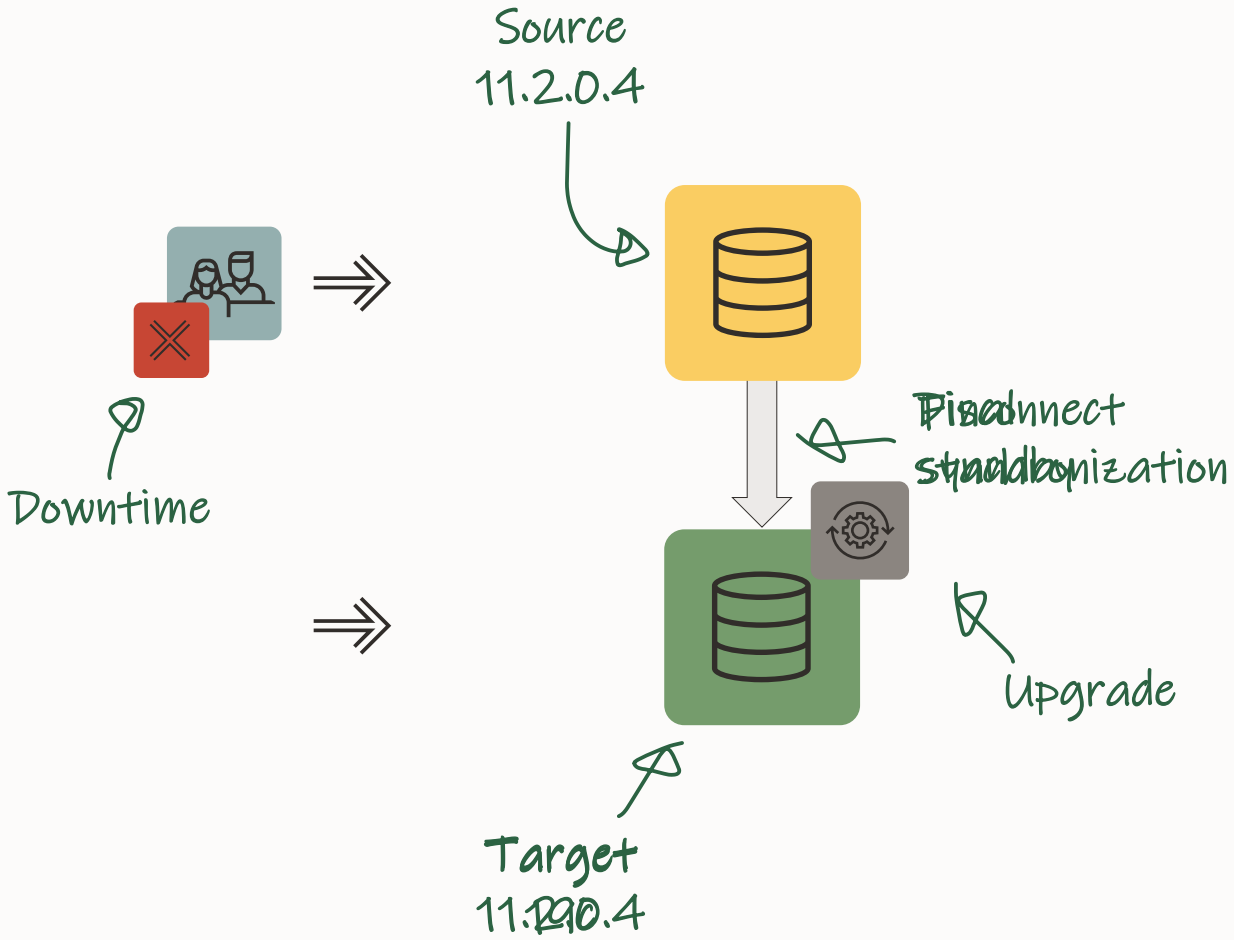

Data Guard | AutoUpgrade

SOURCE DATABASE	TARGET DATABASE
`java -jar autoupgrade.jar -mode analyze`	
DOWNTIME	
`alter system flush redo to ... confirm apply;`	
	`recover managed standby database cancel;`
	`recover standby database;`
	`alter database recover managed standby database finish;`
	`alter database activate physical standby database;`
	`alter database open;`
	`java -jar autoupgrade.jar -mode deploy`

Or use a Transient Logical Standby database for rolling upgrade using DBMS_ROLLING

Pro tip: Watch How Low Can You Go? Zero Downtime Operations for details

What about PDB conversion?

Convert to PDB after migration (and upgrade) using `noncdb_to_pdb.sql`

Can you offload the work from the source database?

Yes, you can. Instantiate the standby database from a backup

Your target database must
be protected by Data Guard?

You can build the future
standby database in advance
and connect it as a cascading standby

Data Guard | Cascading Standby

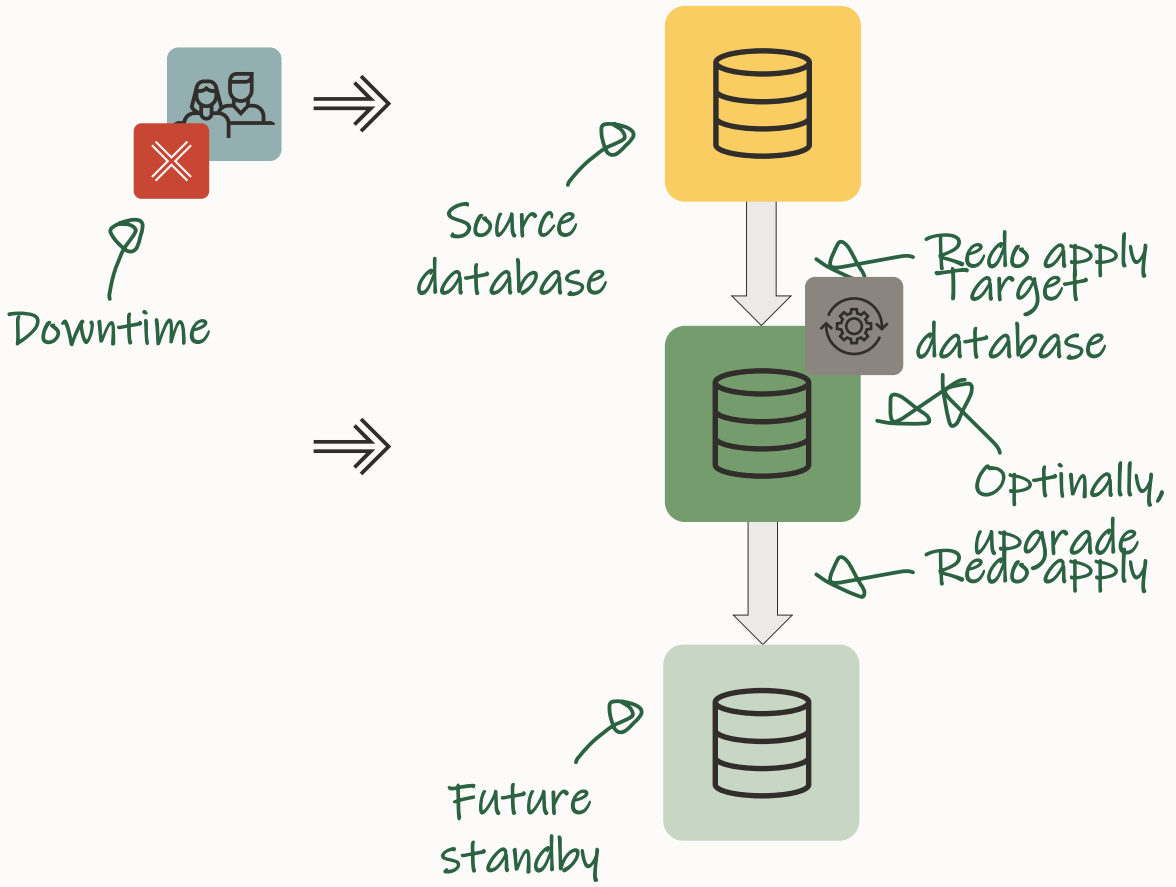

Your target database must
have a valid backup before go-live?

Before go-live perform level 0 backup of the target database

It works even if you upgrade the database

It does not work if you also convert to PDB

Following a PDB conversion new backups of the data files are required before go-live

Your target database must be RAC?

No problem, your standby database can be configured as a RAC database

How about your fallback plan?

Data Guard | Fallback

- To roll back (before go live):
  Source database is untouched
- To fall back (after go live):
  Switchover
- Unless database was upgraded:
  Downgrade
- Unless database was converted:
  Data Pump and GoldenGate

cross

ENDIAN

migration technique

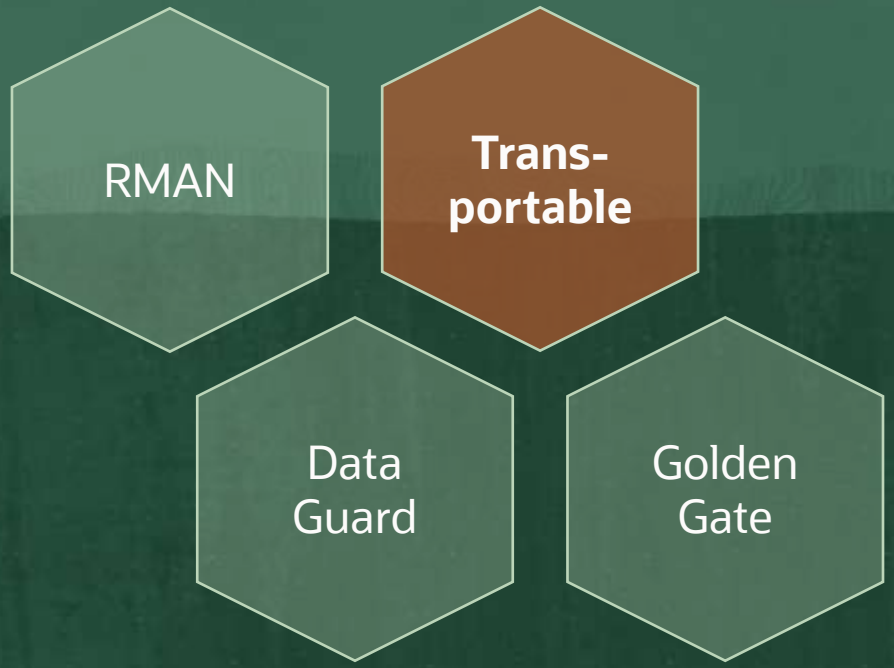

Transport | Concept

Unplug and plugin data tablespaces

Migrate meta information from dictionary

Transport | Concept

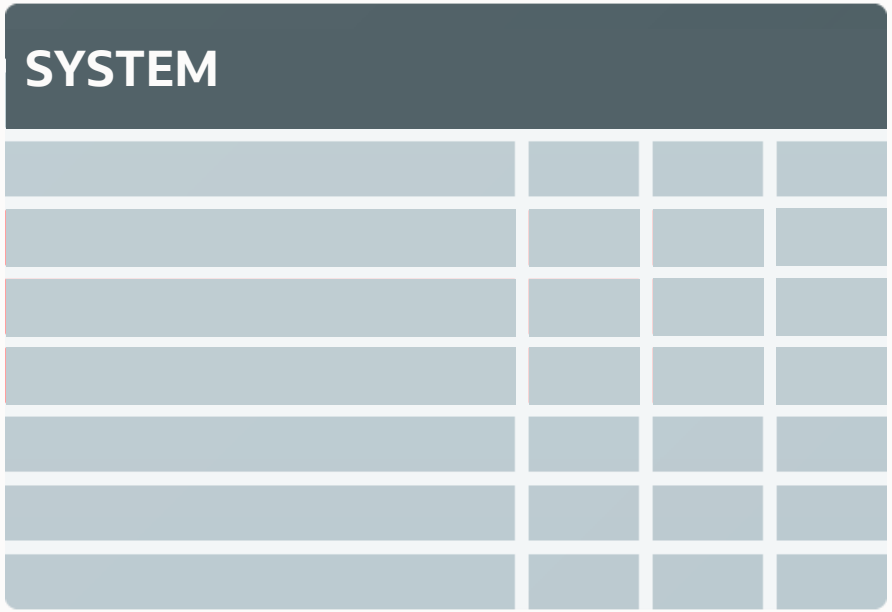

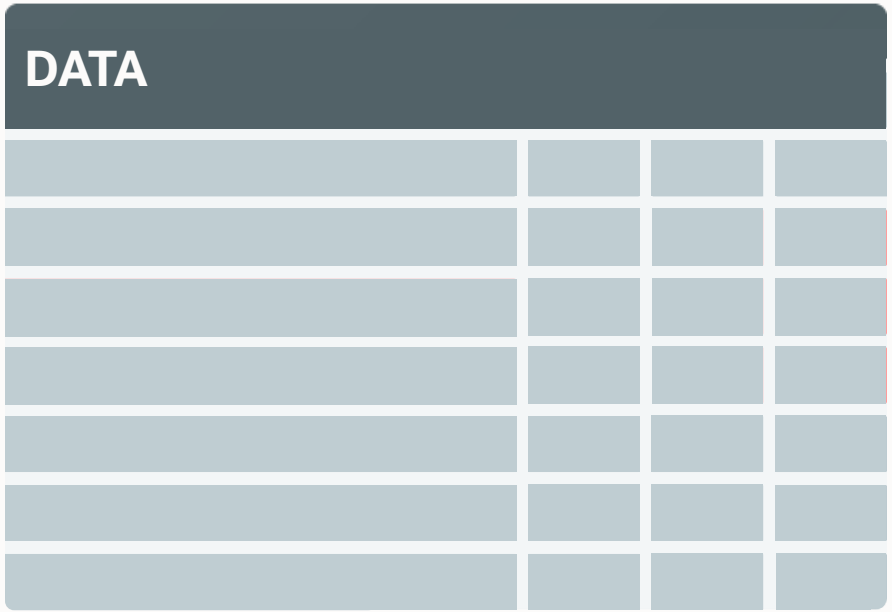

Export with Data Pump

Copy data files

Transport | Benefits

Endianness independent

Direct migration to same or higher version

Direct migration into CDB architecture

Transport | Convert

RMAN	DBMS_FILE_TRANSFER
Out-of-place conversion	In-flight conversion
2 x disk space needed	1 x disk space needed
Supported in newest version of Perl scripts	Not supported in Perl scripts version 4

Photo by Fahrul Azmi on Unsplash

Transport

Methods

Migration | FTEX plus Incremental Backups

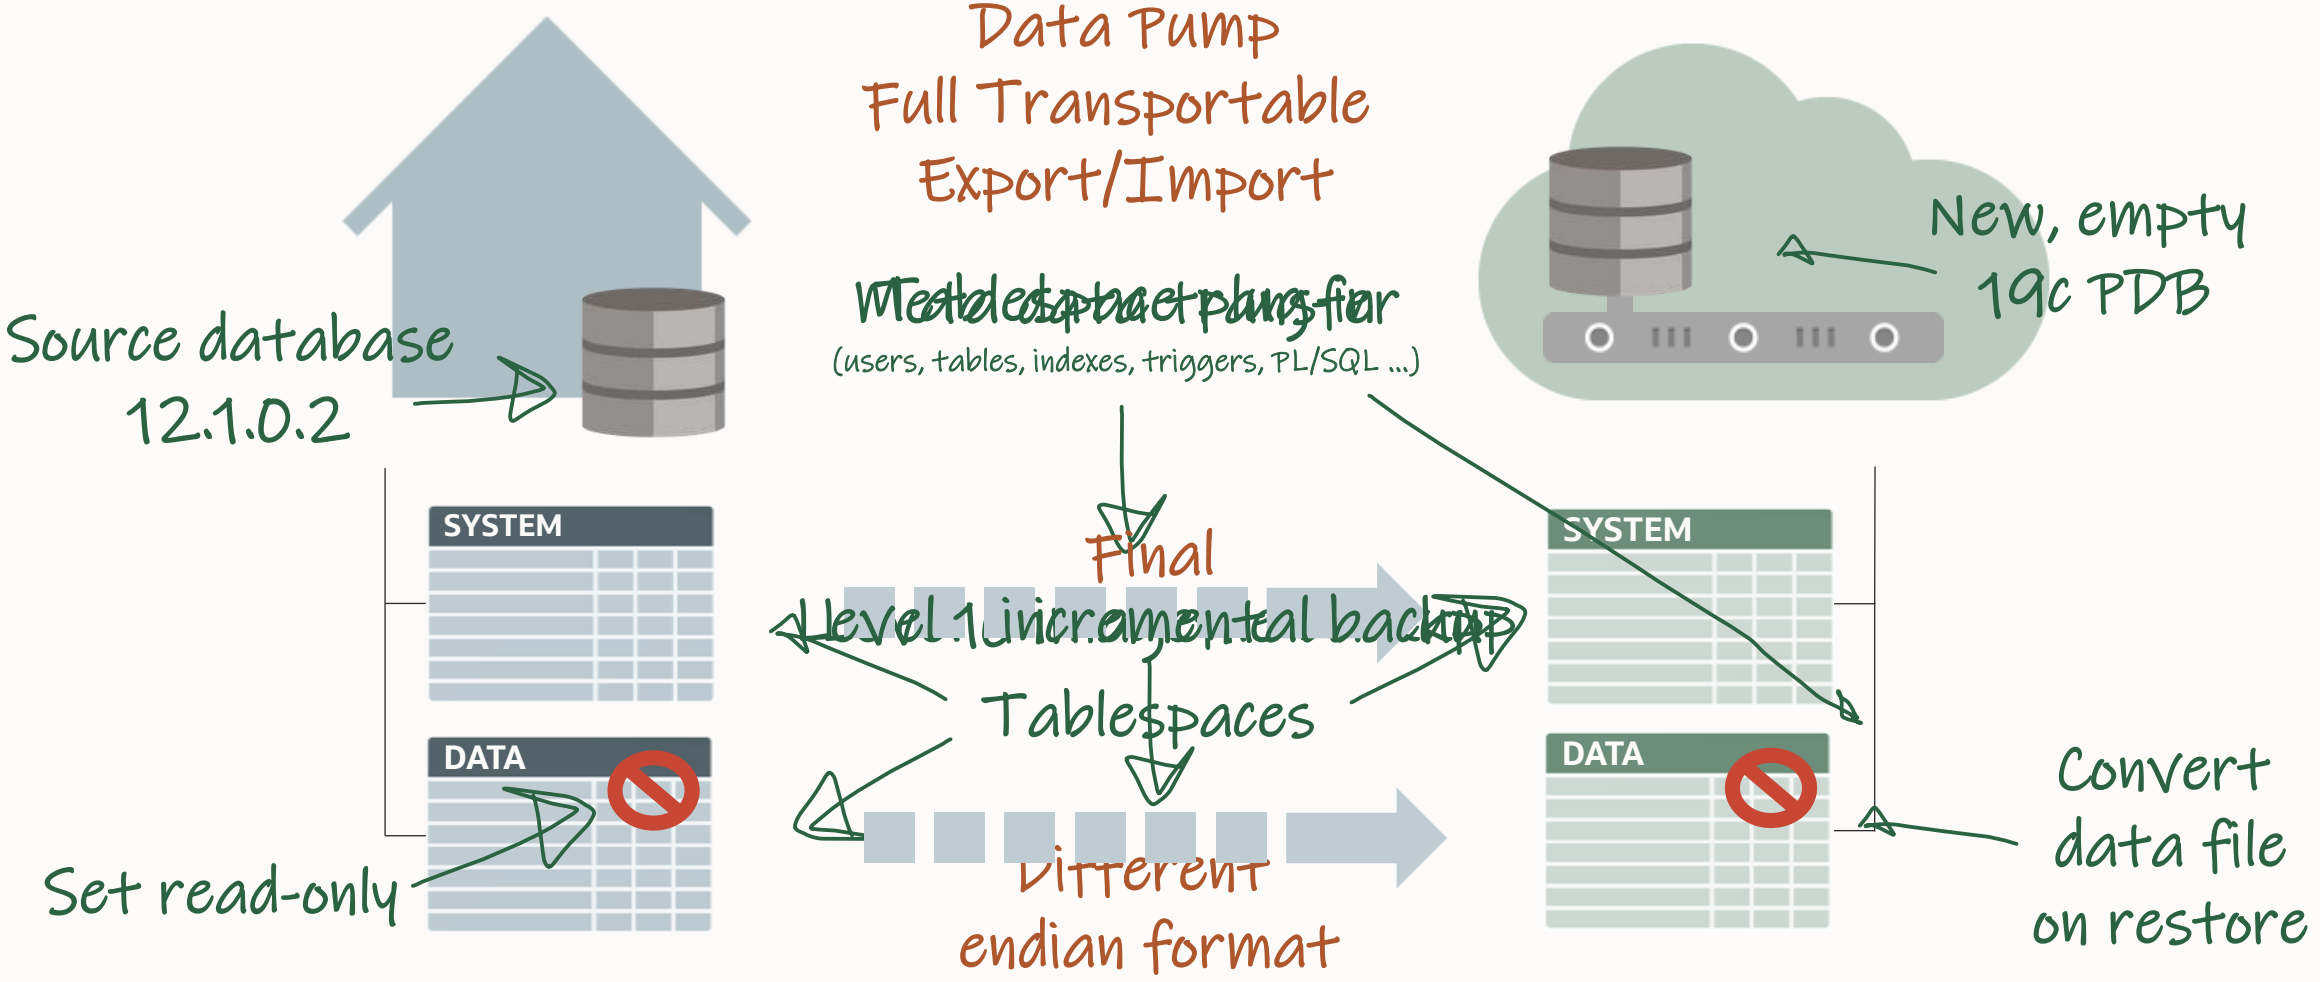

Migration | Oracle E-Business Suite

For EBS cross platform migrations to Oracle 19c, please see:

- MOS 2674405.1
  Using Transportable Tablespaces to Migrate Oracle E-Business Suite Release 12.2 Using Oracle Database 19c Enterprise Edition On a Multitenant Environment)
- Blog post

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

To determine character set:

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

Convert source database to Unicode with DMU

Other Character set workarounds

Transportable | Starter Checklist

Database Creation

Backup / Recovery

TDE

PERL Scripts

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

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

Workaround

- Decrypt before migration
- Reminder: Online decryption available since Oracle 12.2.0.1

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 | 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
- Offline source database afterwards

How about your fallback plan?

Transport | Fallback

- To roll back (before go live):
  Source environment is preserved
- To fall back (after go live):
  Reverse process back to source
- Unless database was upgraded:
  Data Pump and GoldenGate
- Unless database was converted:
  Data Pump and GoldenGate

Photo by asoggetti on Unsplash

Full Transportable Export Import

Step-by-step with PERL scripts

Transport with Incremental Backups | 6 Phases

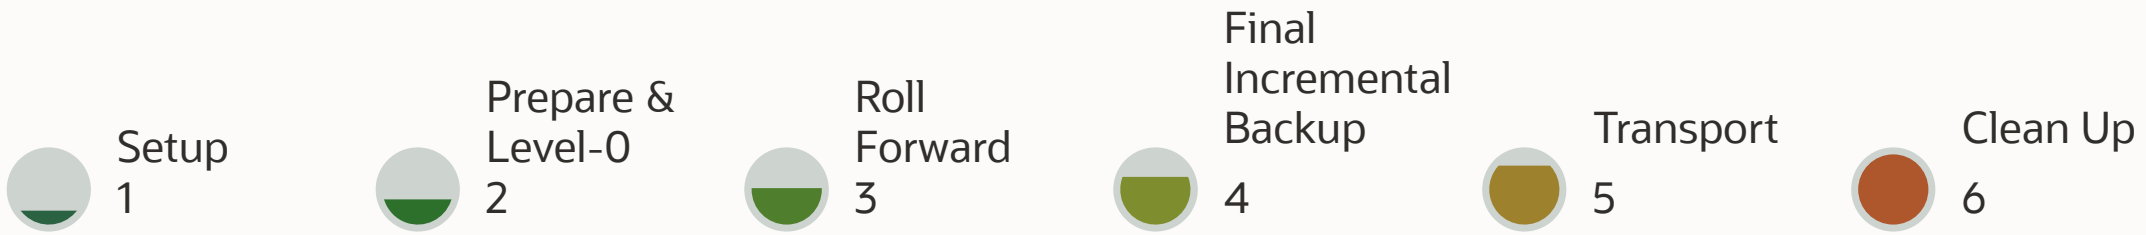

Phase 1 - Setup | Database Creation

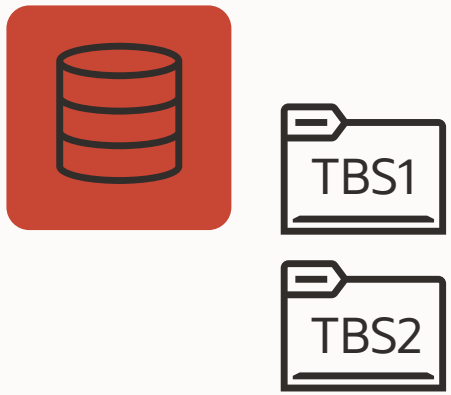

Create database or PDB

`COMPATIBLE` equal or higher
Identical database character sets
Identical national character sets
Identical time zone versions
Identical database time zone setting

Phase 1 - Setup | Download PERL Scripts

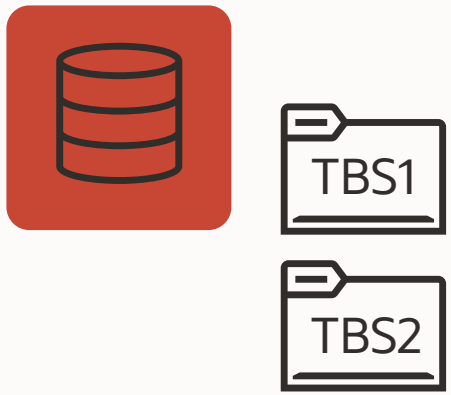

MOS Note: 2471245.1 – V4 PERL Scripts

Name	Kind	Date Modified
xtt.newproperties	Document	19. Feb 2019 at 06:53
xttprep.tmpl	Document	24. May 2017 at 15:57
xtt.properties	Java pr...ties file	5. Mar 2021 at 07:35
xttdriver.pl	Perl script	7. Jul 2019 at 18:21
xttcnvrtbkupdest.sql	SQL source	24. May 2017 at 15:57
xttdbopen.sql	SQL source	24. May 2017 at 15:57
xttstartupnomount.sql	SQL source	24. May 2017 at 15:57

Phase 1 - Setup | Standard Method

RMAN backup and convert

```
xttdriver.pl --backup
xttdriver.pl --restore
```

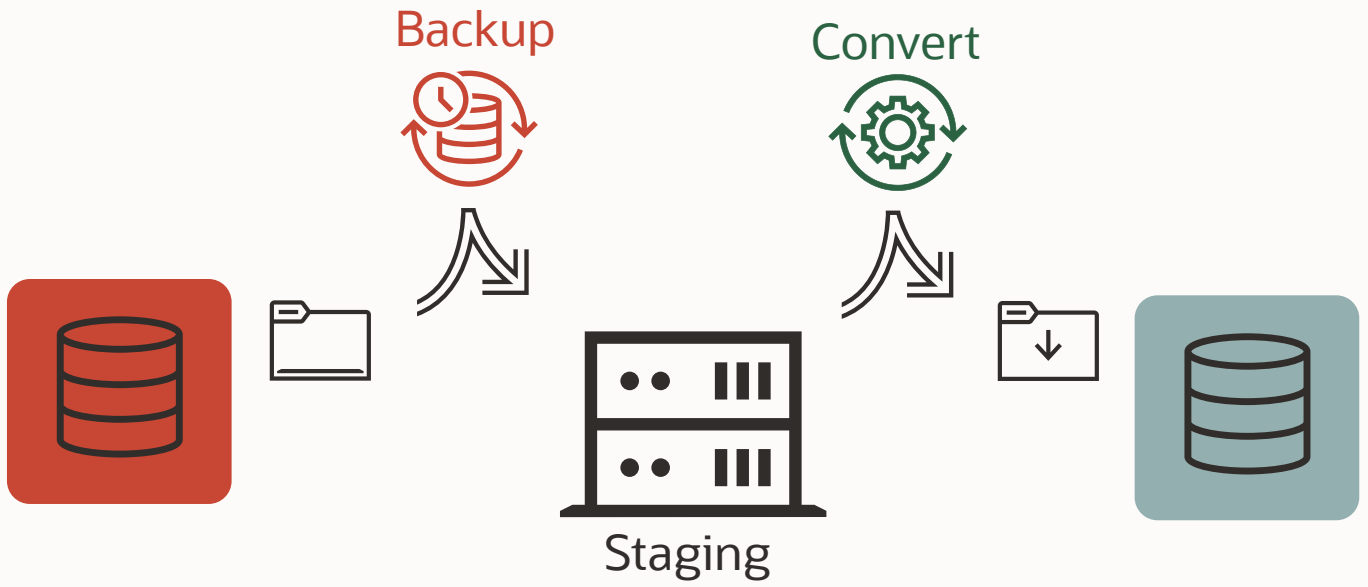

Phase 1 - Setup | Non-Supported Method

DBMS_FILE_TRANSFER

```
xttdriver.pl -S
xttdriver.pl -G
```

No support with V4 PERL scripts
2 TB limitation

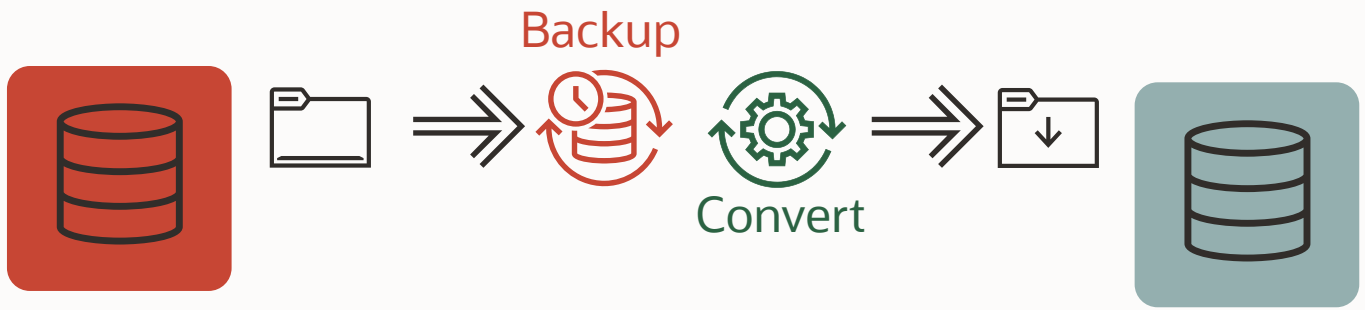

Phase 1 - Setup | Configure

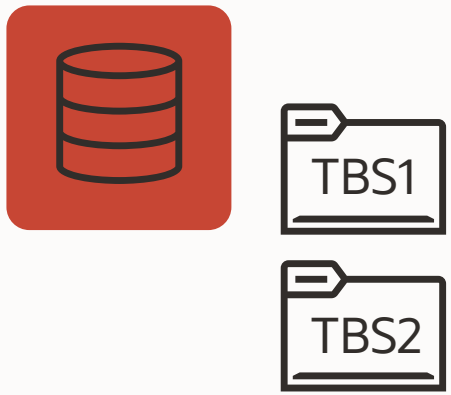

xtt.properties

```
tablespaces=TBS1,TBS2

platformid=13

src_scratch_location=/NFS_backups/
dest_scratch_location=/NFS_backups/

dest_datafile_location=+DATA

asm_home=/u01/app/19/grid
asm_sid=+ASM1

parallel=16
rollparallel=2
```

Phase 2 - Prepare | Level 0 Backup

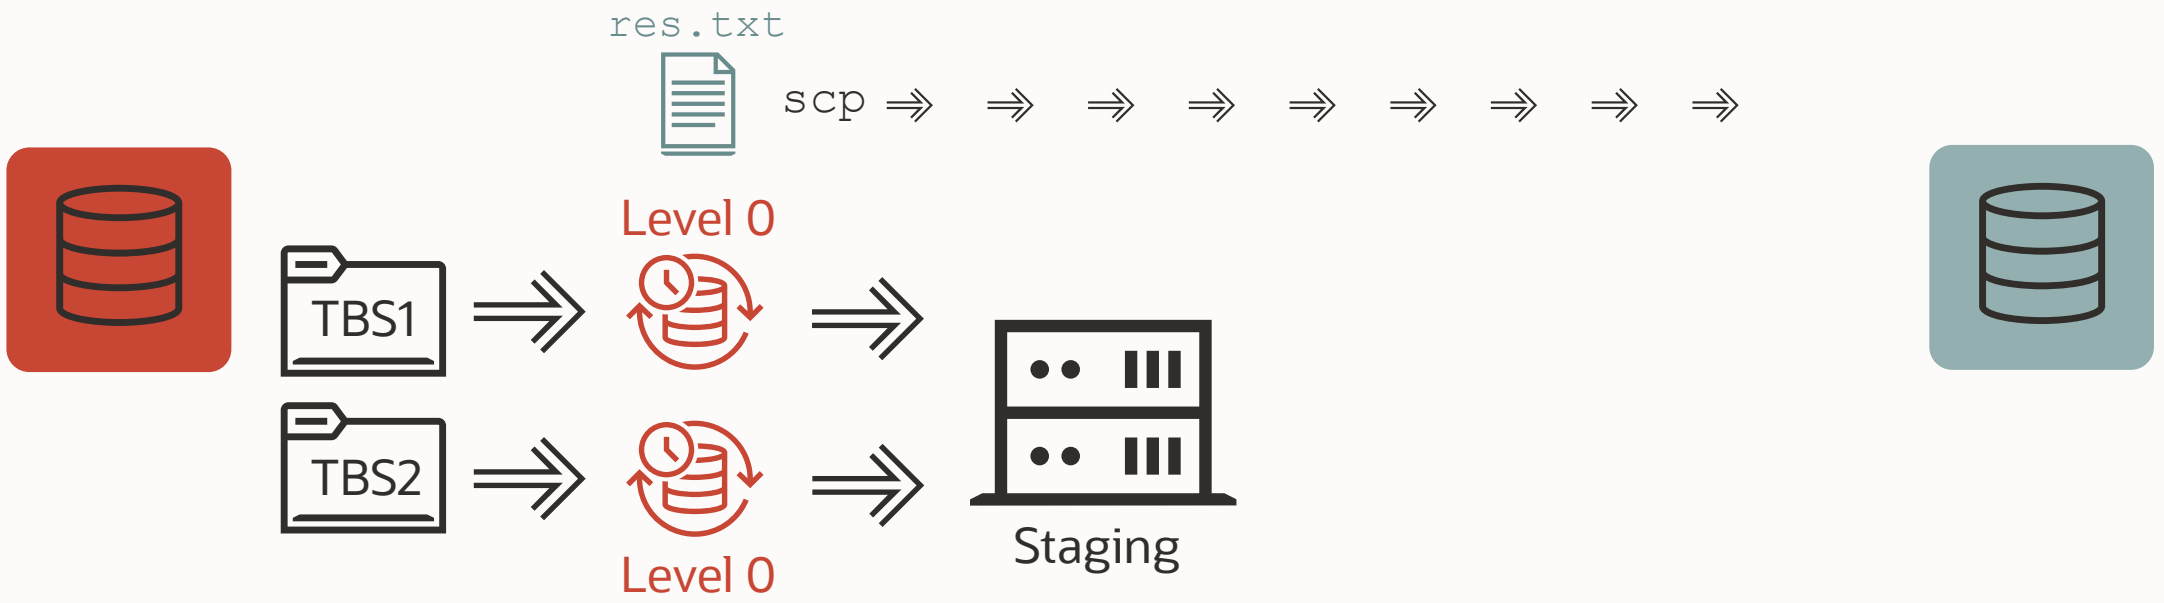


```
xttdriver.pl --backup
```

Phase 2 – Convert/Restore | Level 0 Backup

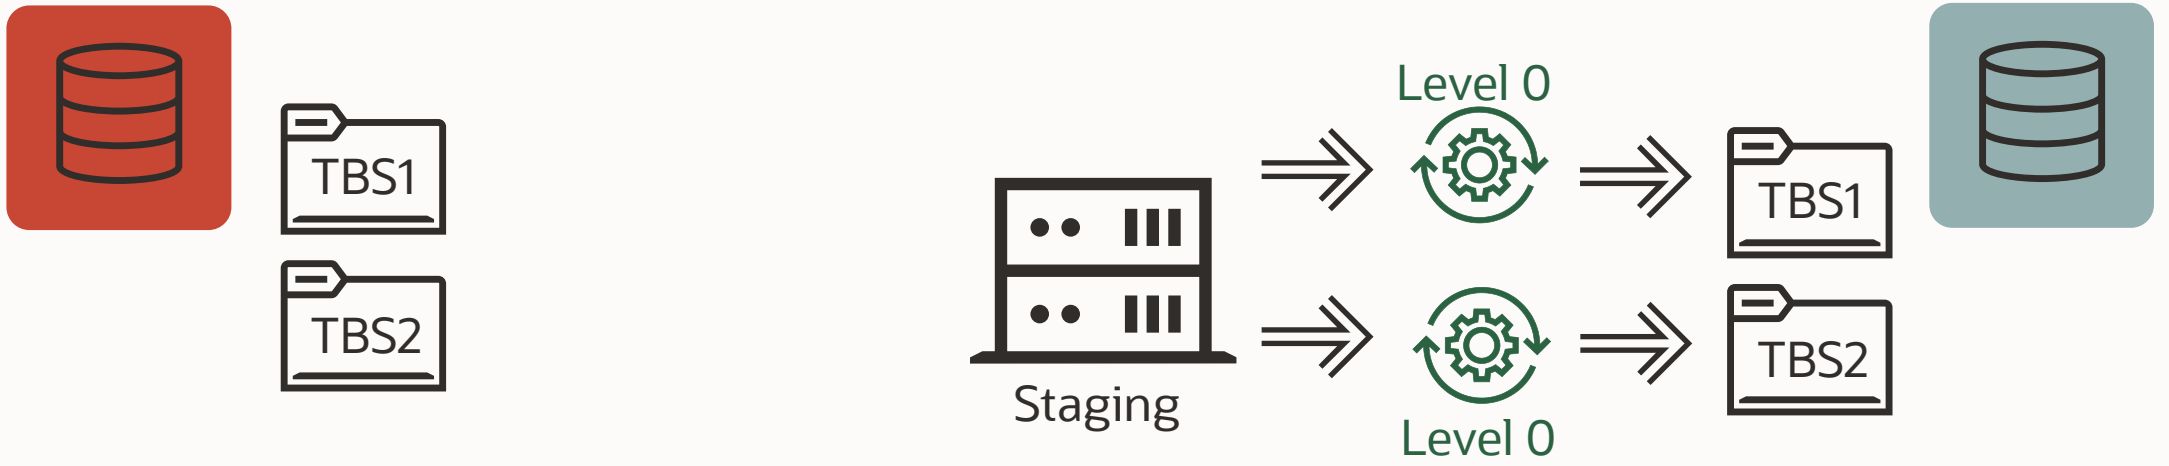


```
xttdriver.pl --restore
```

Phase 3 – Roll Forward | Level 1 Backup

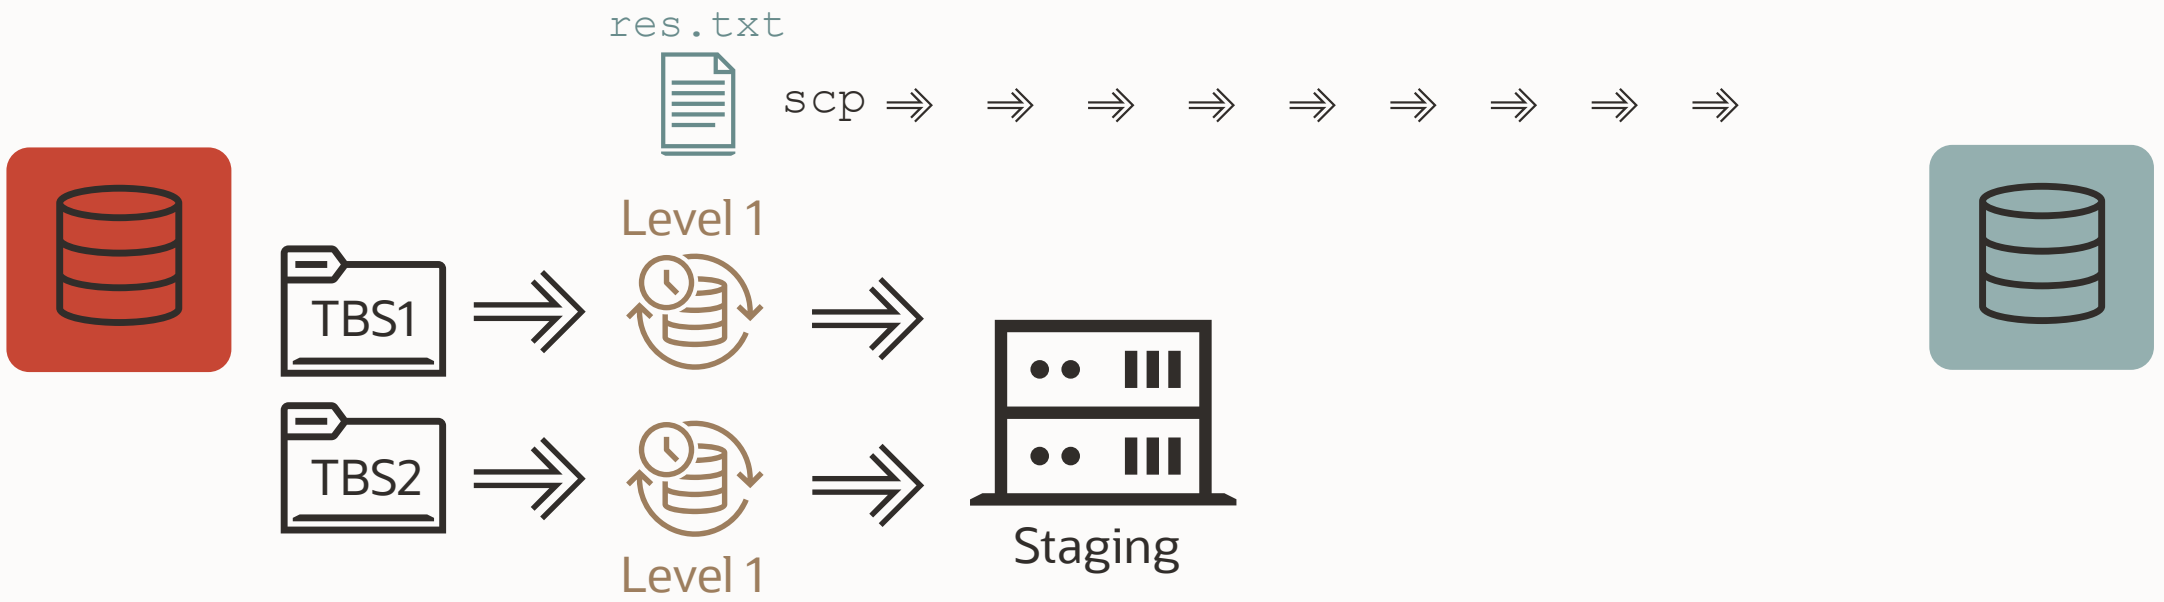


```
xttdriver.pl --backup
```

Phase 3 – Convert/Restore/Merge | Level 1 Backup

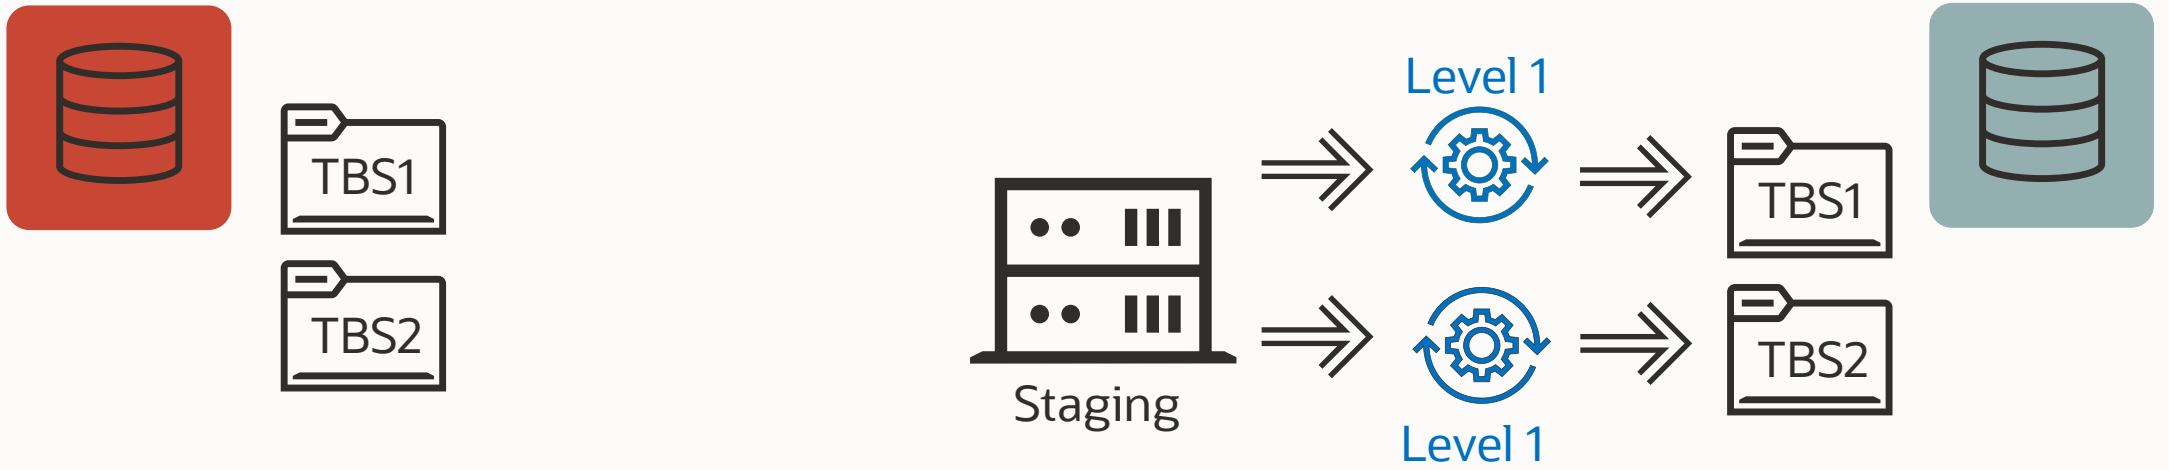


```
xttdriver.pl --restore
```

Phase 3 – Repeat | Level 1 Backup/Convert/Restore/Merge

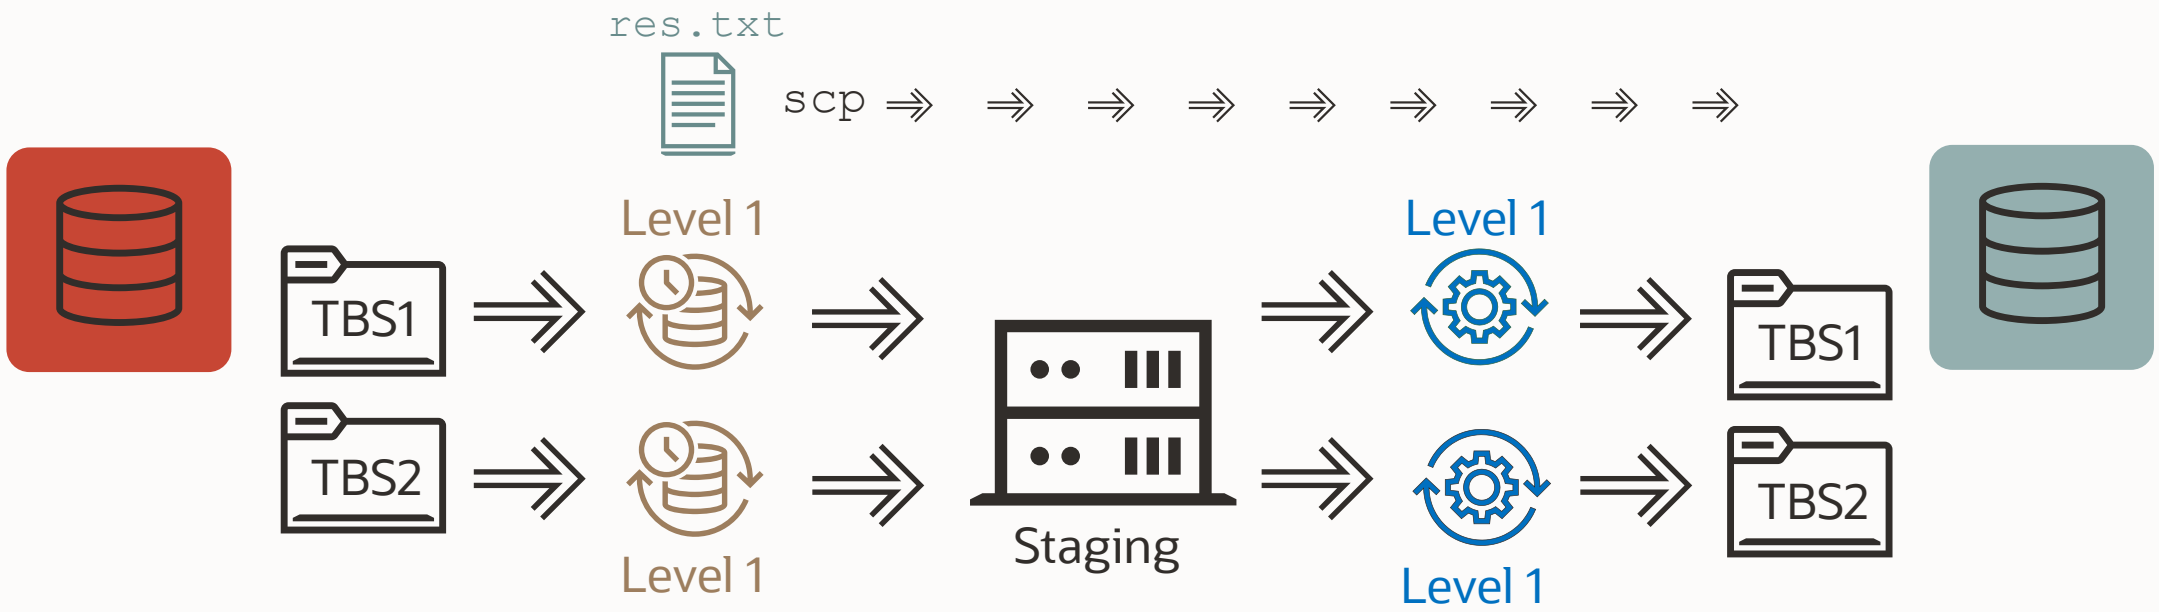

xttdriver.pl --backup

xttdriver.pl --restore

Phase 4 – Final Level 1 | Read Only

Phase 4 – Final Level 1 | Backup/Convert/Restore/Merge

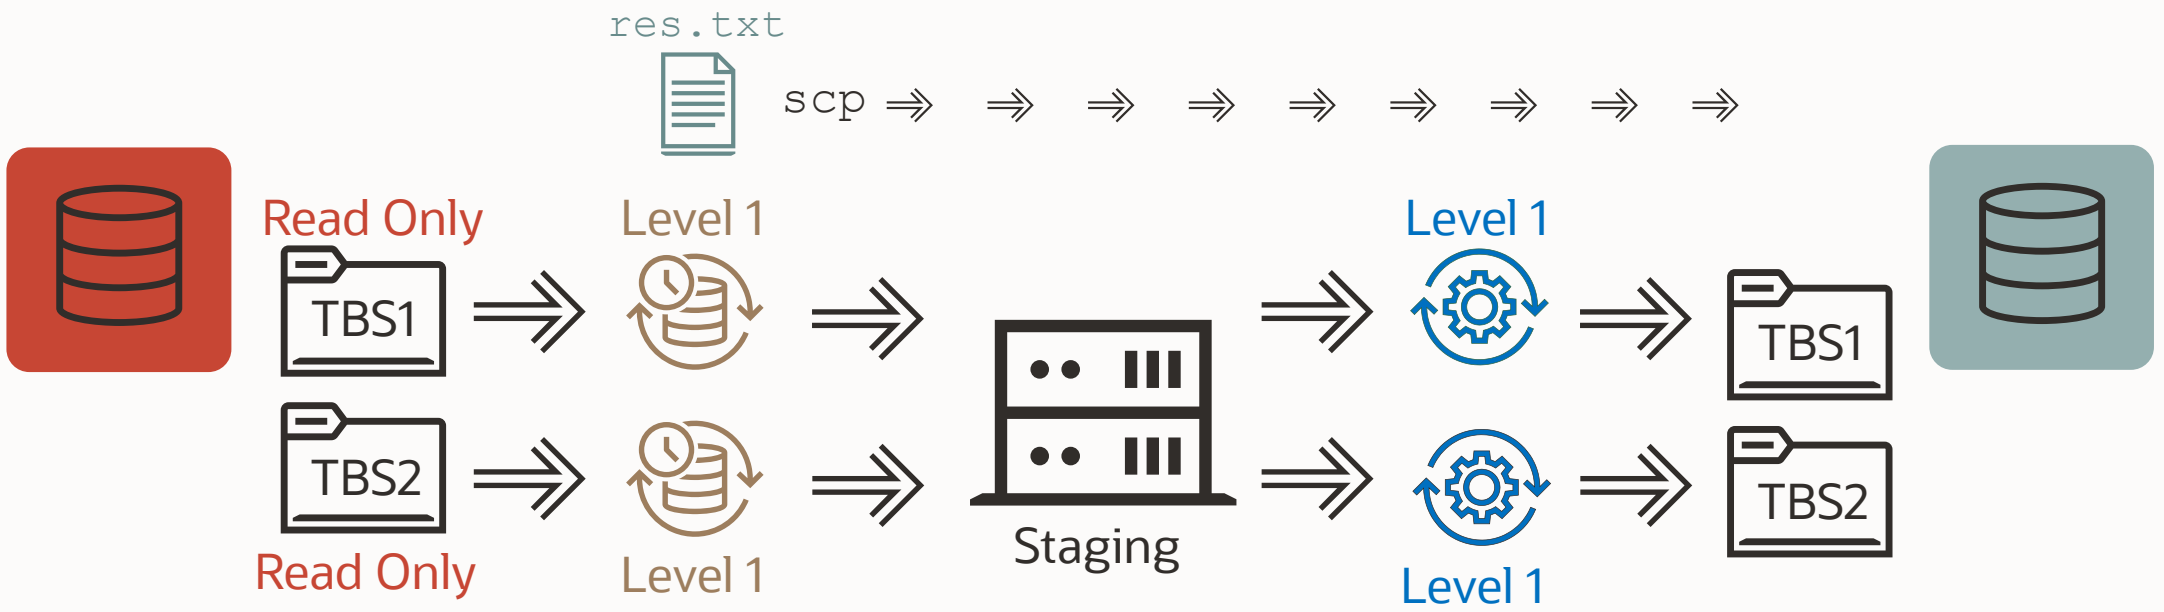

xttdriver.pl --backup

xttdriver.pl --restore

Phase 5 – Transport | Full Transportable Export Import

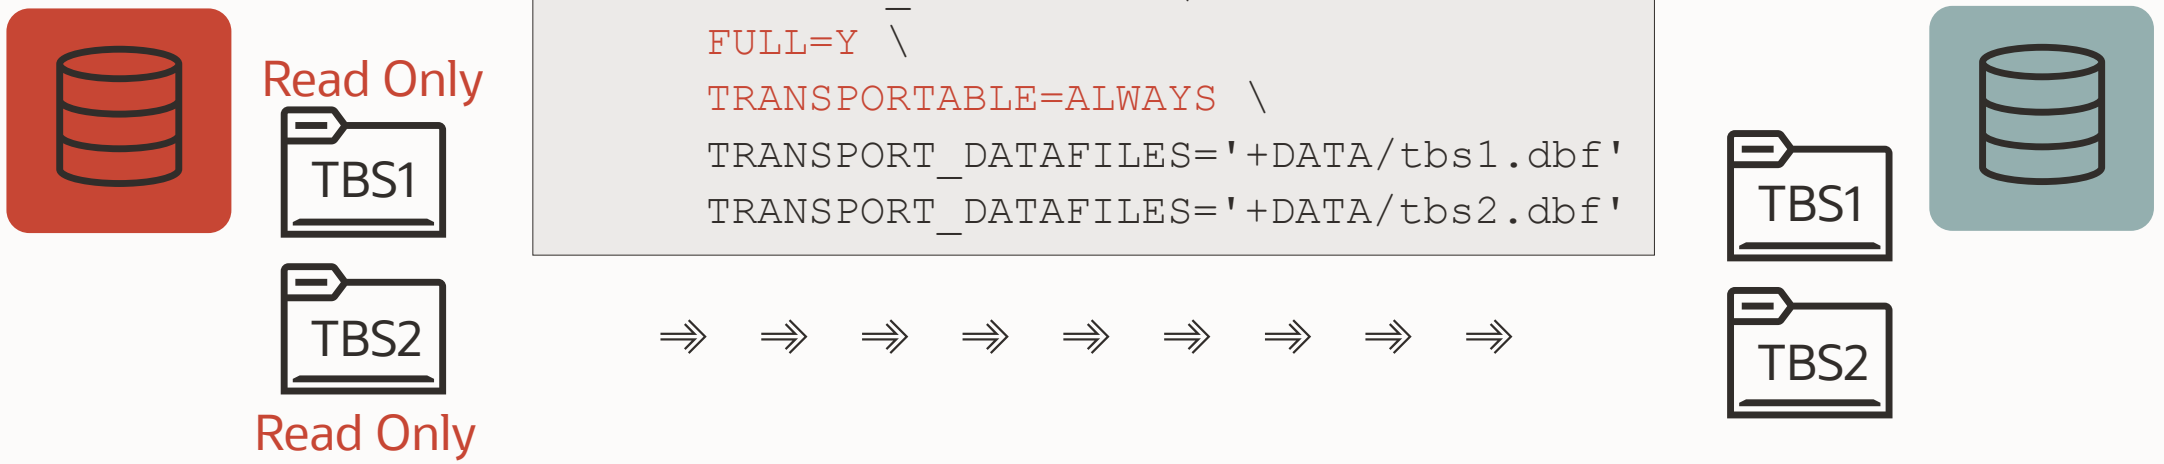

Phase 6 - Cleanup | Validation

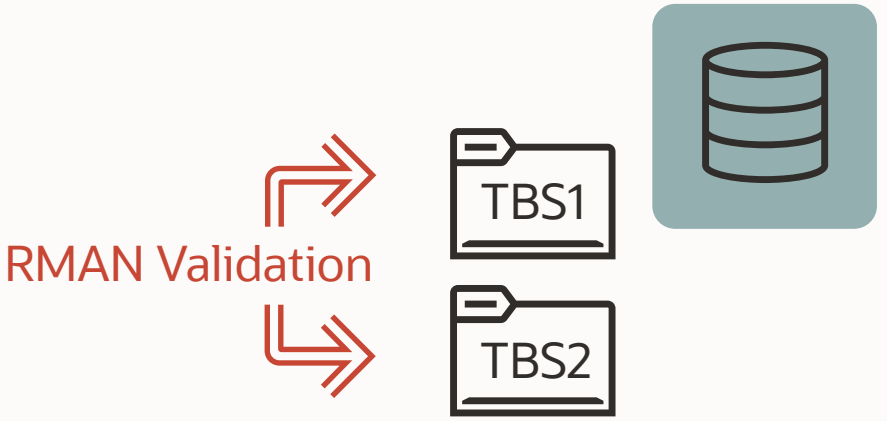

Transport with Incremental Backups | Demo

Watch on YouTube

Data Guard | Transportable

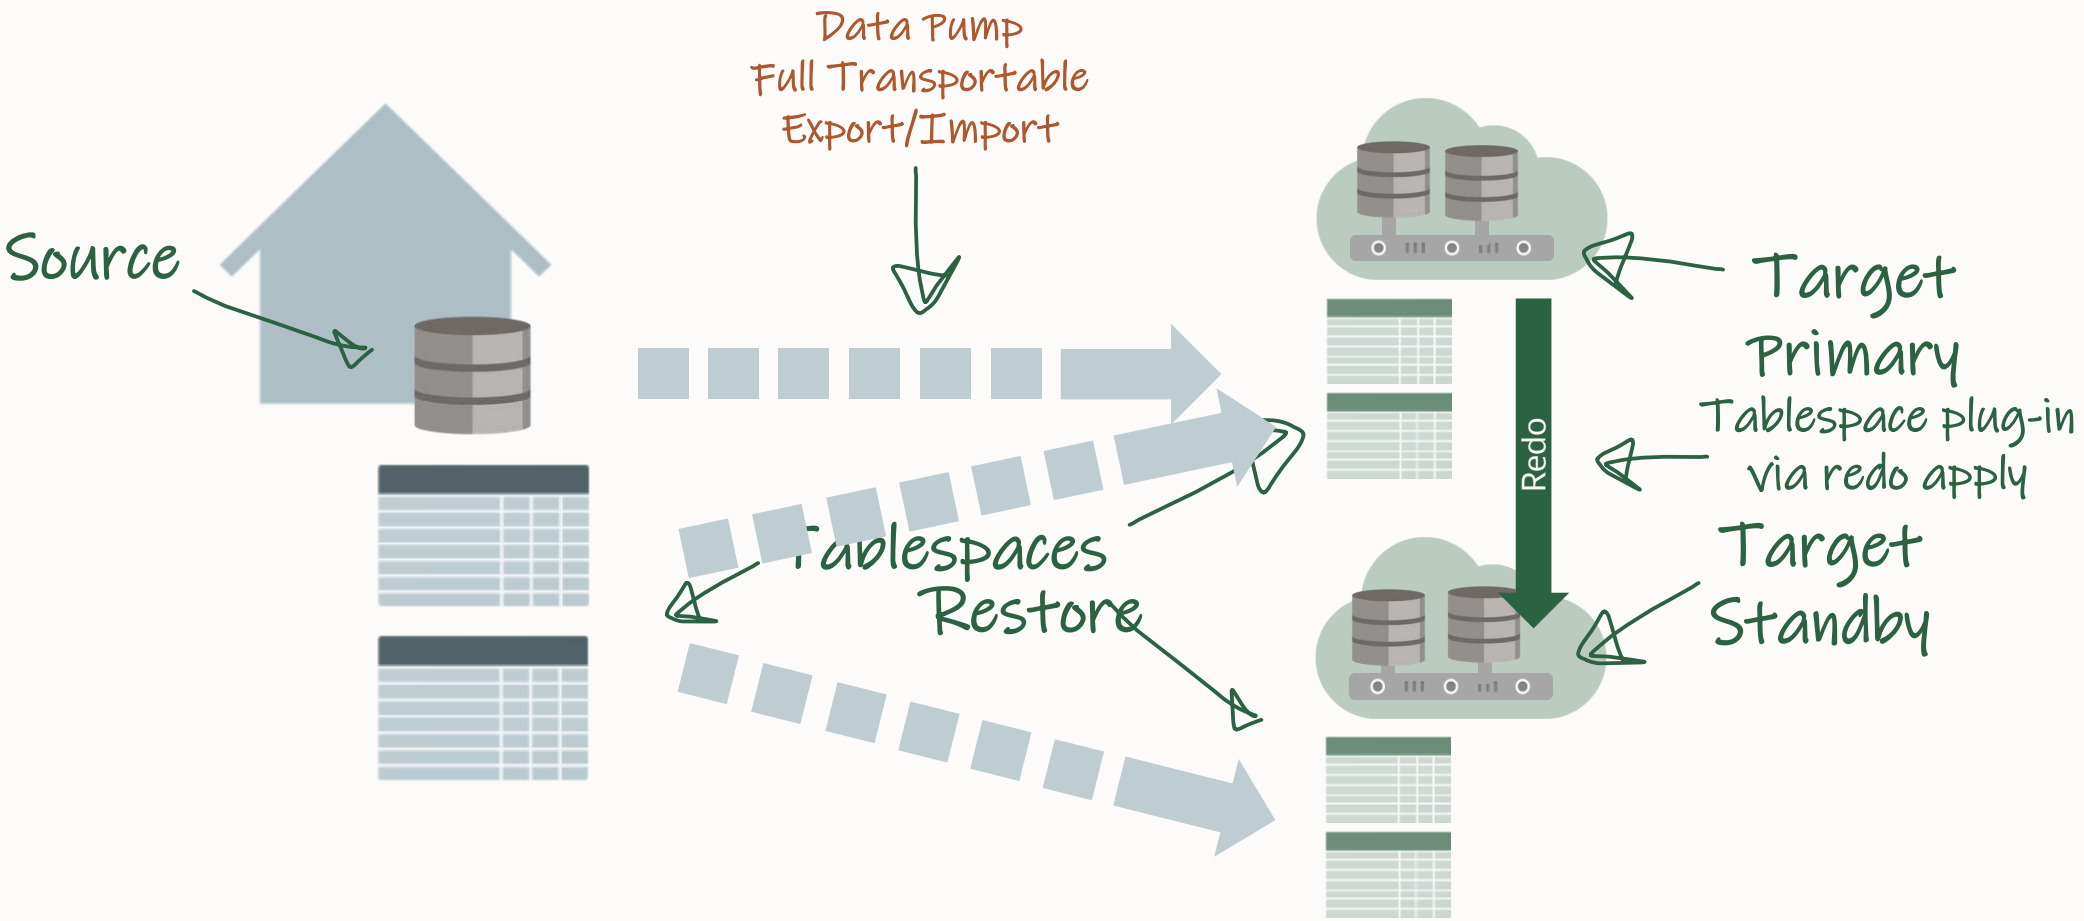

Data Guard | Transportable

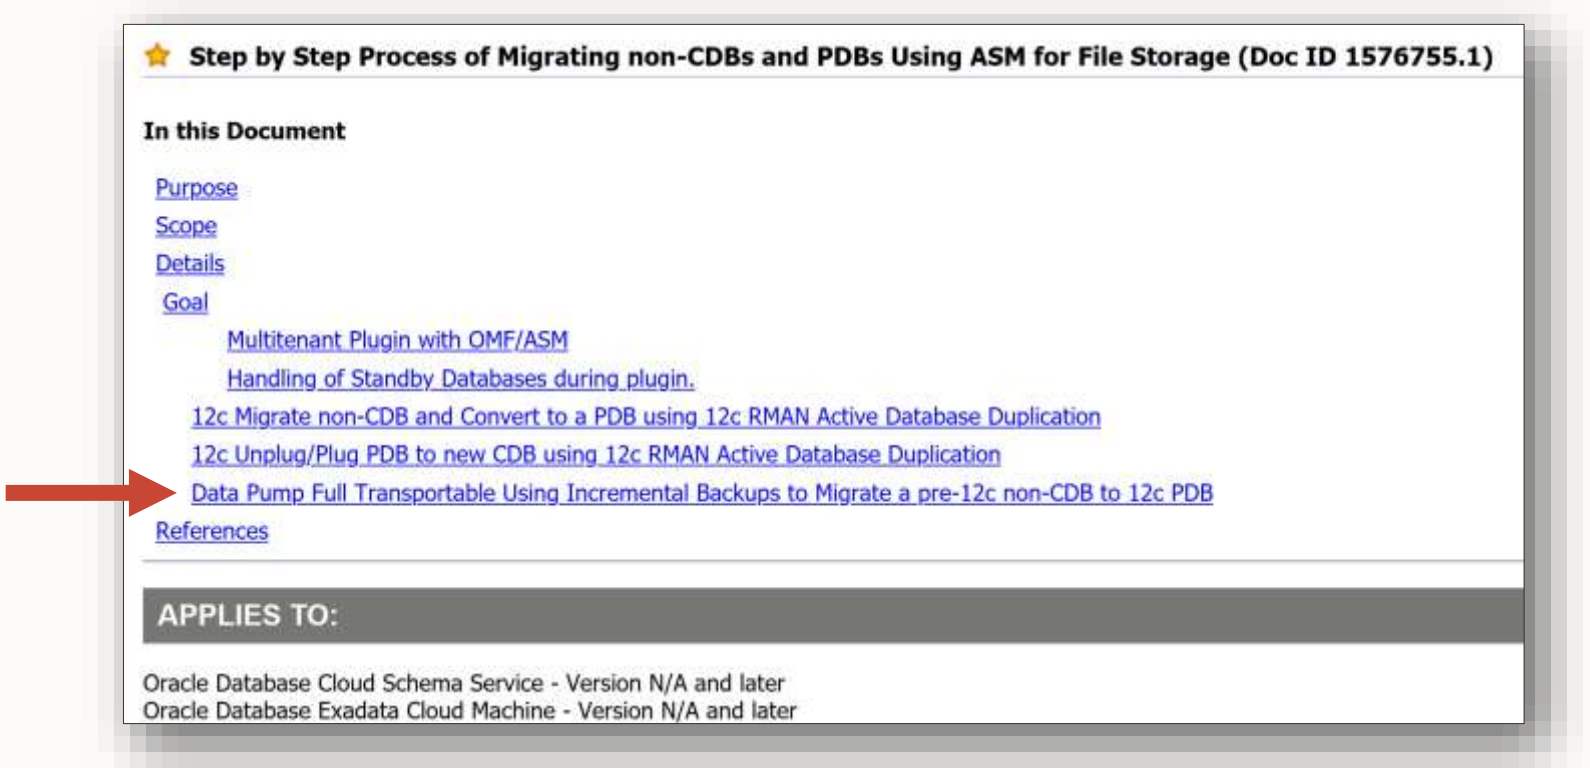

Step by Step Process of Migrating non-CDBs and PDBs Using ASM for File Storage (Doc ID 1576755.1)

230 Terabyte in less than 24 hours

PERL scripts in action at a top health care insurance

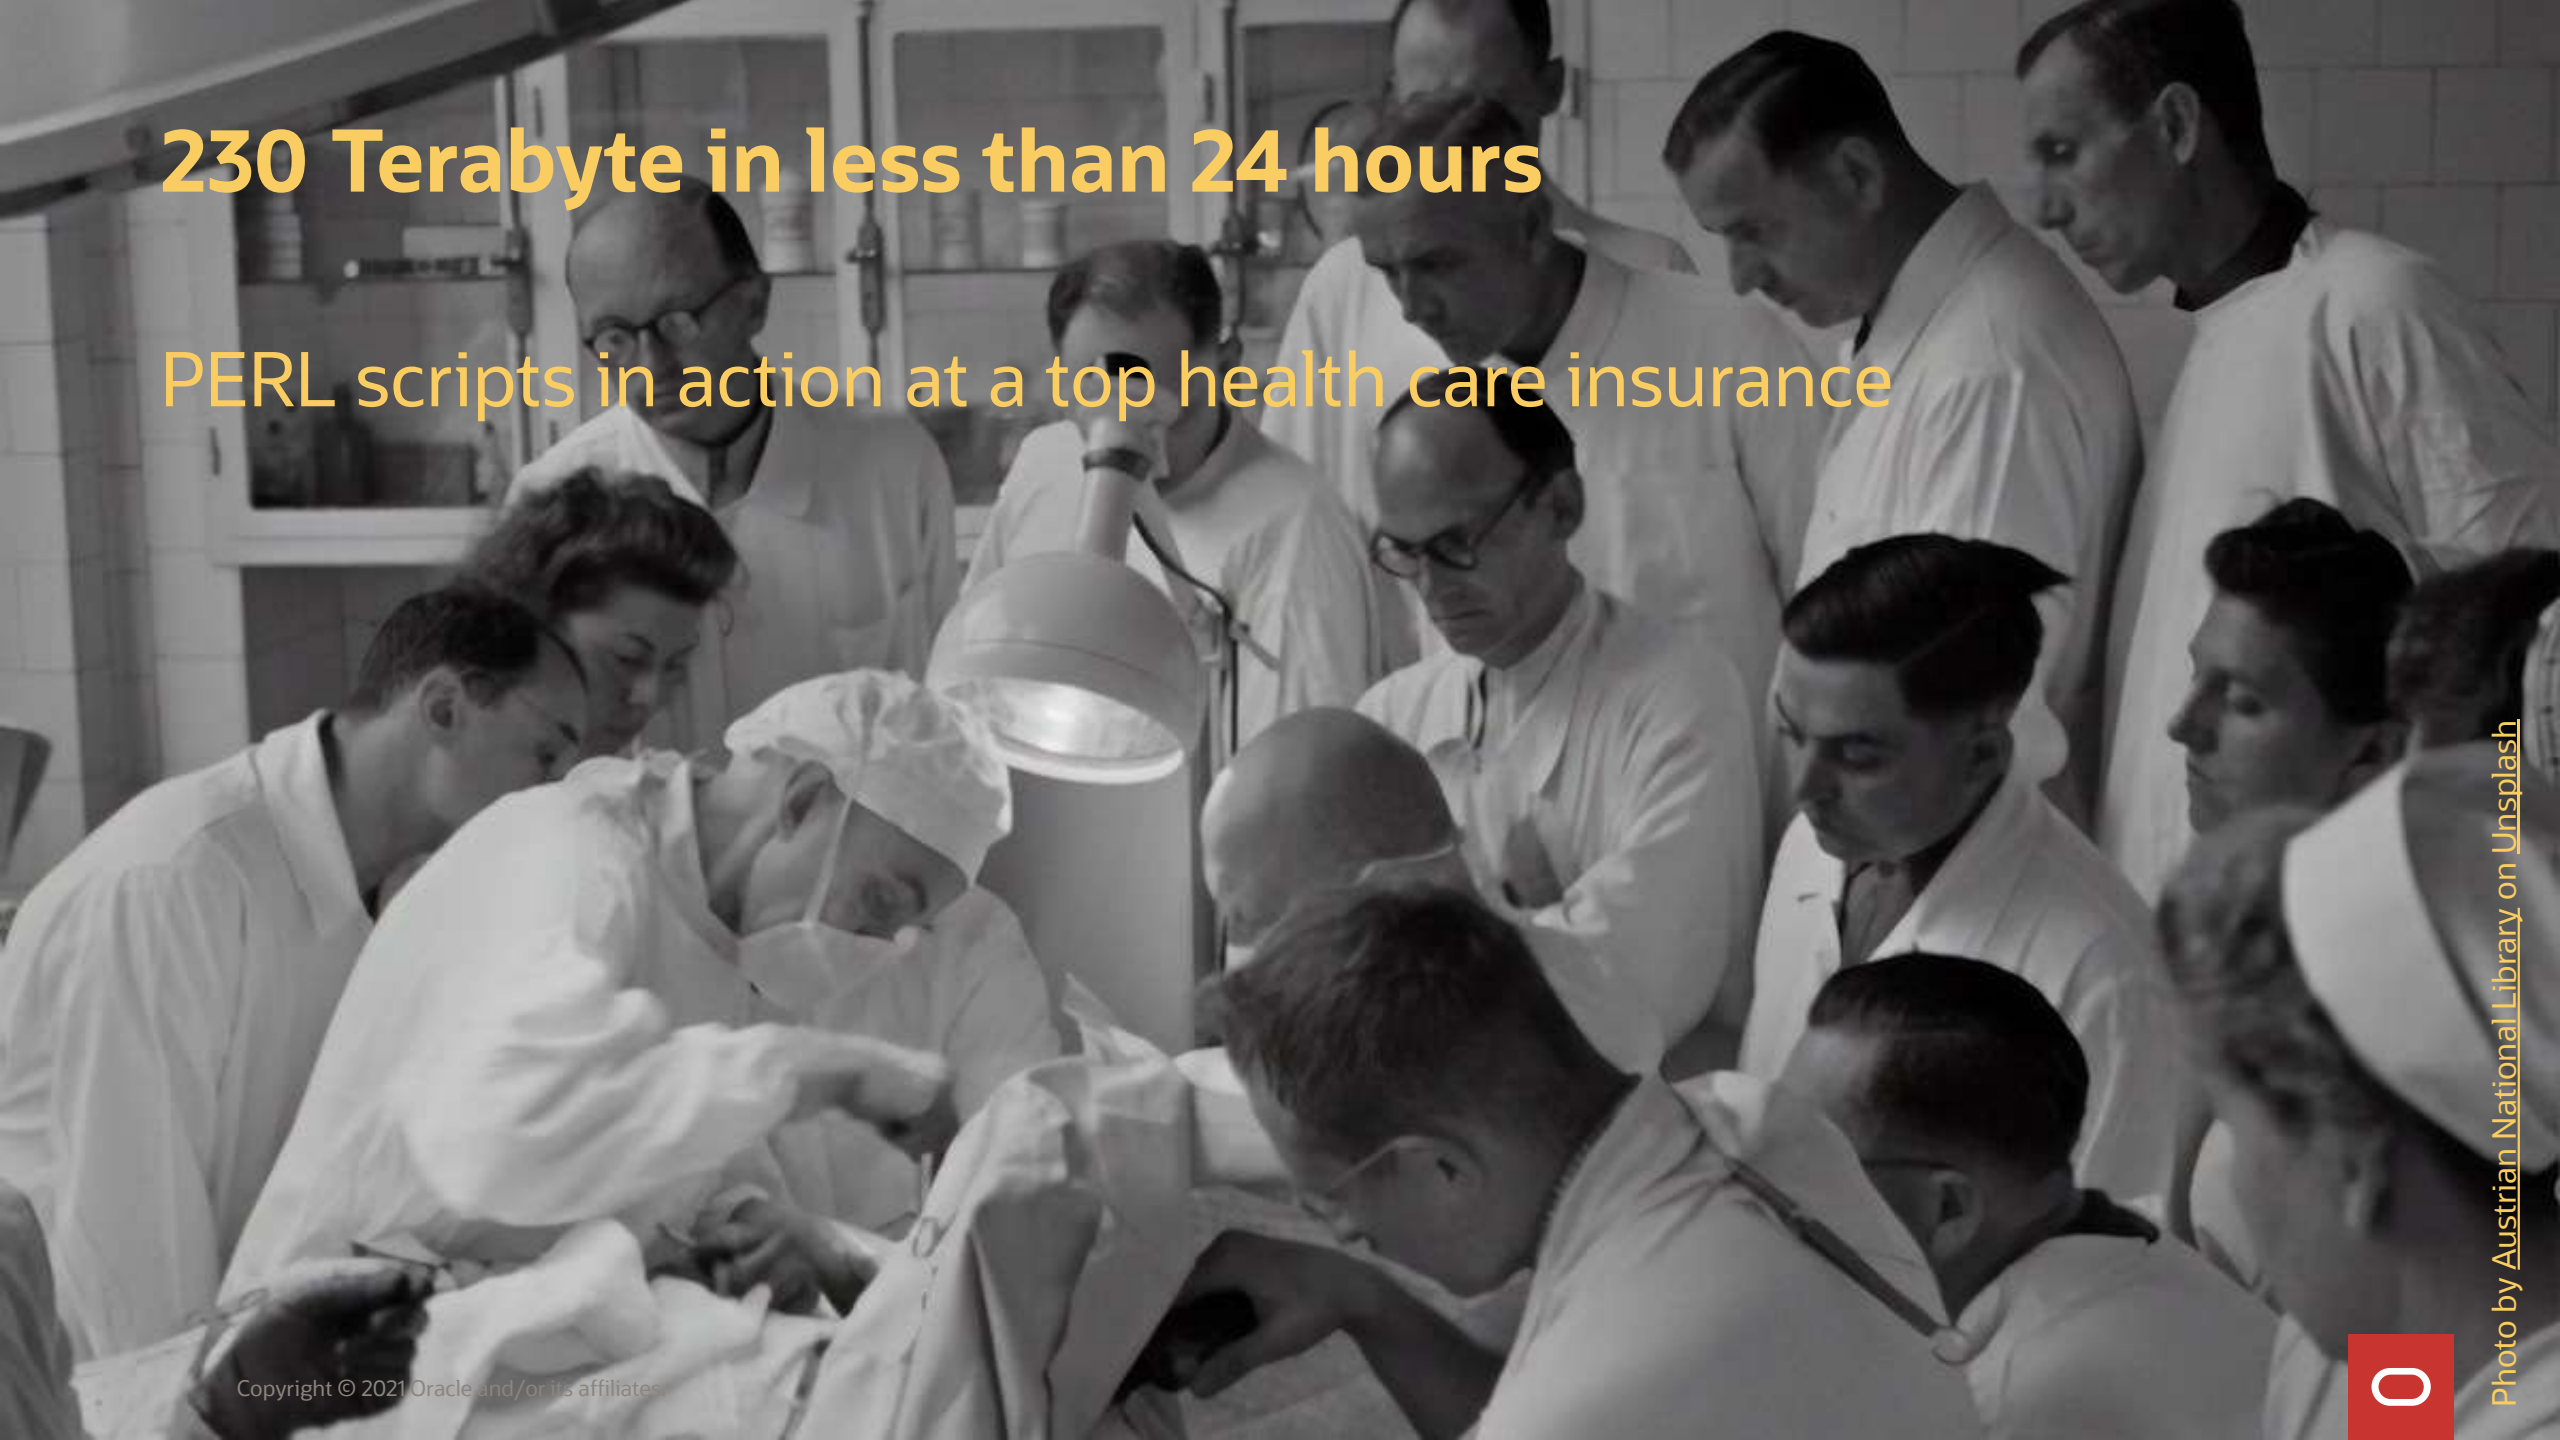

Photo by Austrian National Library on Unsplash

Customer Case | Health Care

Customer

Project 2017

Constraints

Preparation

Migration

Success?

Remarks

Customer

- One of the top healthcare insurance providers in the United States
  - Over 50,000 employees, over $50 BILLION annual revenue

Oracle Partner

- Centric Consulting, a management and technology consulting company
- Oracle Platinum Partner

Customer Case | Health Care

Customer

Project 2017

Constraints

Preparation

Migration

Success?

Remarks

Source

- AIX 5.3, Oracle Database 11.2.0.3, DB on filesystem

Target

- Exadata running Oracle Linux, Database 12.1.0.2, RAC/ASM

Enterprise data warehouse & operational data store

- Critical for day-to-day operations
- Minimizing downtime is critical
- Data Guard for DR

Customer Case | Health Care

Customer

Project 2017

Constraints

Preparation

Migration

Success?

Remarks

Huge, active database

- 230+ TB (and growing!)
- Generates ~1.2TB redo *per hour*

Initial attempts using Oracle GoldenGate were unsuccessful

- Could not keep up with massive redo generation

Earlier version (V.2) of PERL migrations scripts

- Did not handle addition of tablespaces during migration
- Single-threaded file transfer

Customer Case | Health Care

Customer

Project 2017

Constraints

Preparation

Migration

Success?

Remarks

Single-threaded file transfer

- Use `PARALLEL` in `xtt.properties`
- Number of data files to be processed in parallel

Transfer was too slow during initial tests

- 100 MB/sec throughput
- For 230 TB: almost 27 days (!) just for the Prepare phase

Customer Case | Health Care

Customer

Project 2017

Constraints

Preparation

Migration

Success?

Remarks

Workarounds

- 40 identical directories, each held a complete XTTS PERL script installation
- Distribute >530 tablespaces into 40 tablespace groups
- Run 40 jobs concurrently with PARALLEL=2, or 80 files at a time

Result: ~800 MB/sec throughput

- Reduced prepare phase from 27 days to 6 days

Customer Case | Health Care

Customer

Project 2017

Constraints

Preparation

Migration

Success?

Remarks

Additional customizations

- Cross-check scripts to ensure all tablespaces were being migrated
- Custom scripts
  - Automate 40 parallel script executions
  - Data Pump import par file for the plug-in step
- Load balanced RMAN CONVERT
  - Conversion on all four nodes in the Exadata
  - Result: Over 230 TBs converted in under 10 hours

Customer Case | Health Care

- Customer
- Project 2017
- Constraints
- **Preparation**
- Migration
- Success?
- Remarks

Environment

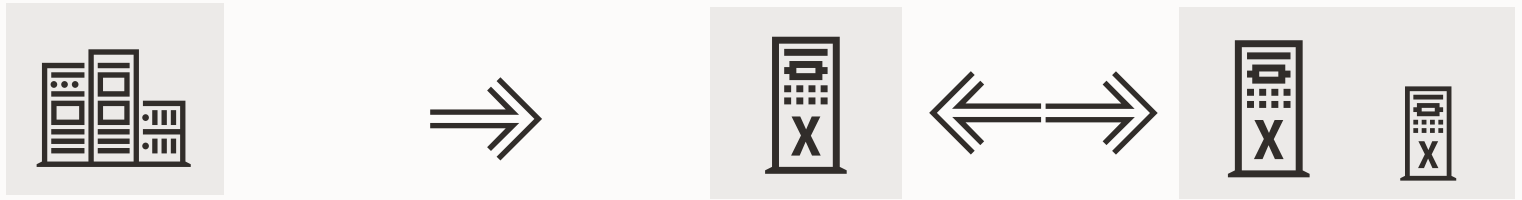

Customer Case | Health Care

Customer

Project 2017

Constraints

Preparation

Migration

Success?

Remarks

Migration and upgrade completed in one phase

- AIX to Linux
- 11.2.0.3 to 12.1.0.2
- Single instance to RAC
- File system to ASM
- 230+ TB

Customer Case | Health Care

Customer

Project 2017

Constraints

Preparation

Migration

Success?

Remarks

Completed in an 18-hour READ ONLY window!

	Phase	Location	Duration
0	Prepare Phase (RMAN Level 0)	Source	6 Days
1	Convert Image Copies	Target	10 Hrs
2	Incremental Backup (Pass #1)	Source	24 Hrs
3	Convert / Roll-Forward	Target	12 Hrs
4	Increment SCN	Source	1 Min
5	Incremental Backup (Pass #2)	Source	12 Hrs
6	Convert / Roll-Forward	Target	6 Hrs
7	Increment SCN	Source	1 Min
8	Incremental Backup (Pass #3)	Source	6 Hrs
9	Convert / Roll-Forward	Target	3 Hrs
10	Increment SCN	Source	1 Min
	*** READ-ONLY Window		18 Hrs
11	Place Tablespaces in READ-ONLY Mode	Source	30 Mins
12	Incremental Backup (Final)	Source	3 Hrs
13	Full Meta-Data Export	Source	3 Hrs
14	Create Dynamic Import Plug-In Par File	Source	1 Hr
15	Convert / Roll-Forward	Target	1.5 Hrs
16	Import Plug-In Step	Target	11 Hrs
17	Full Meta-Data Import	Target	2 Hrs
	*** Database Migrated ***		

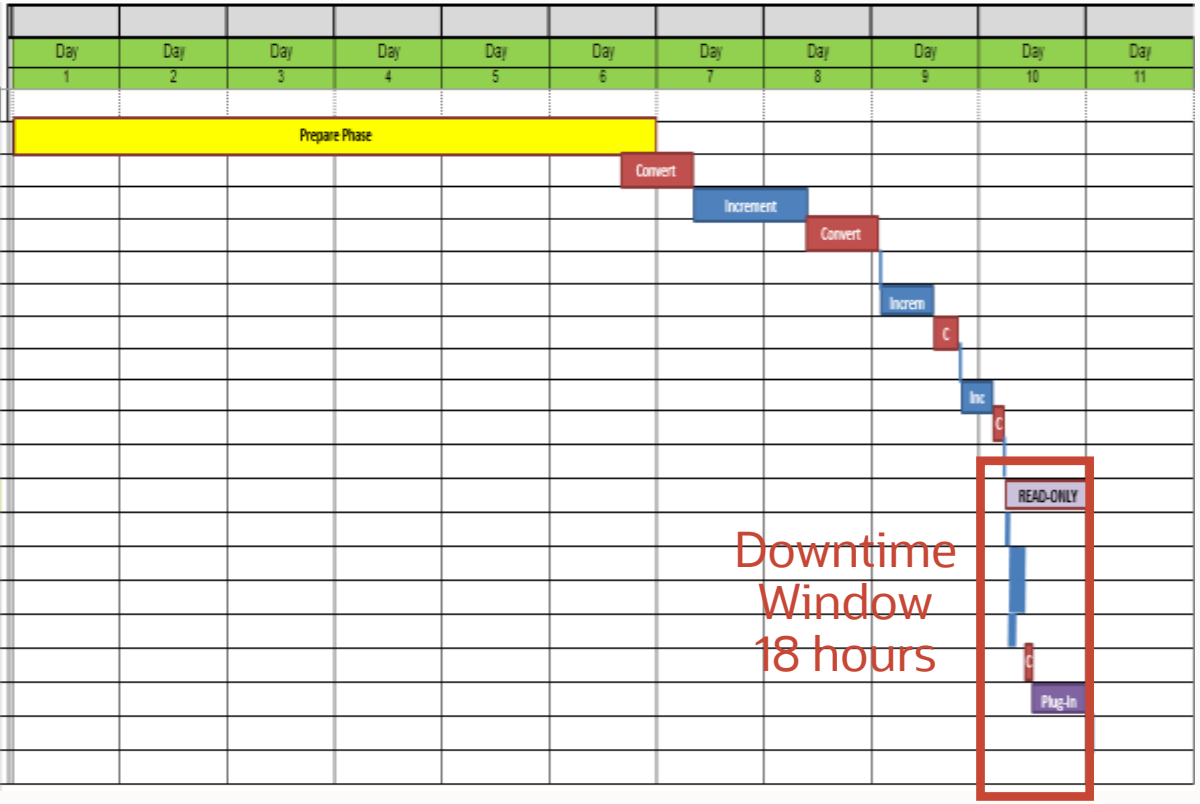

Customer Case | Health Care

Customer

Project 2017

Constraints

Preparation

Migration

Success?

Remarks

Get the latest version of the PERL scripts

Plan for the "unexpected"

Customize the process for VLDBs

- Otherwise, the Prepare Phase may take very long

ExaCC migration

PUSHING THE LIMITS

1500 databases with an ZDLRA

Photo by Patrick Hendry on Unsplash

ExaCC Migration | Strategic Overview

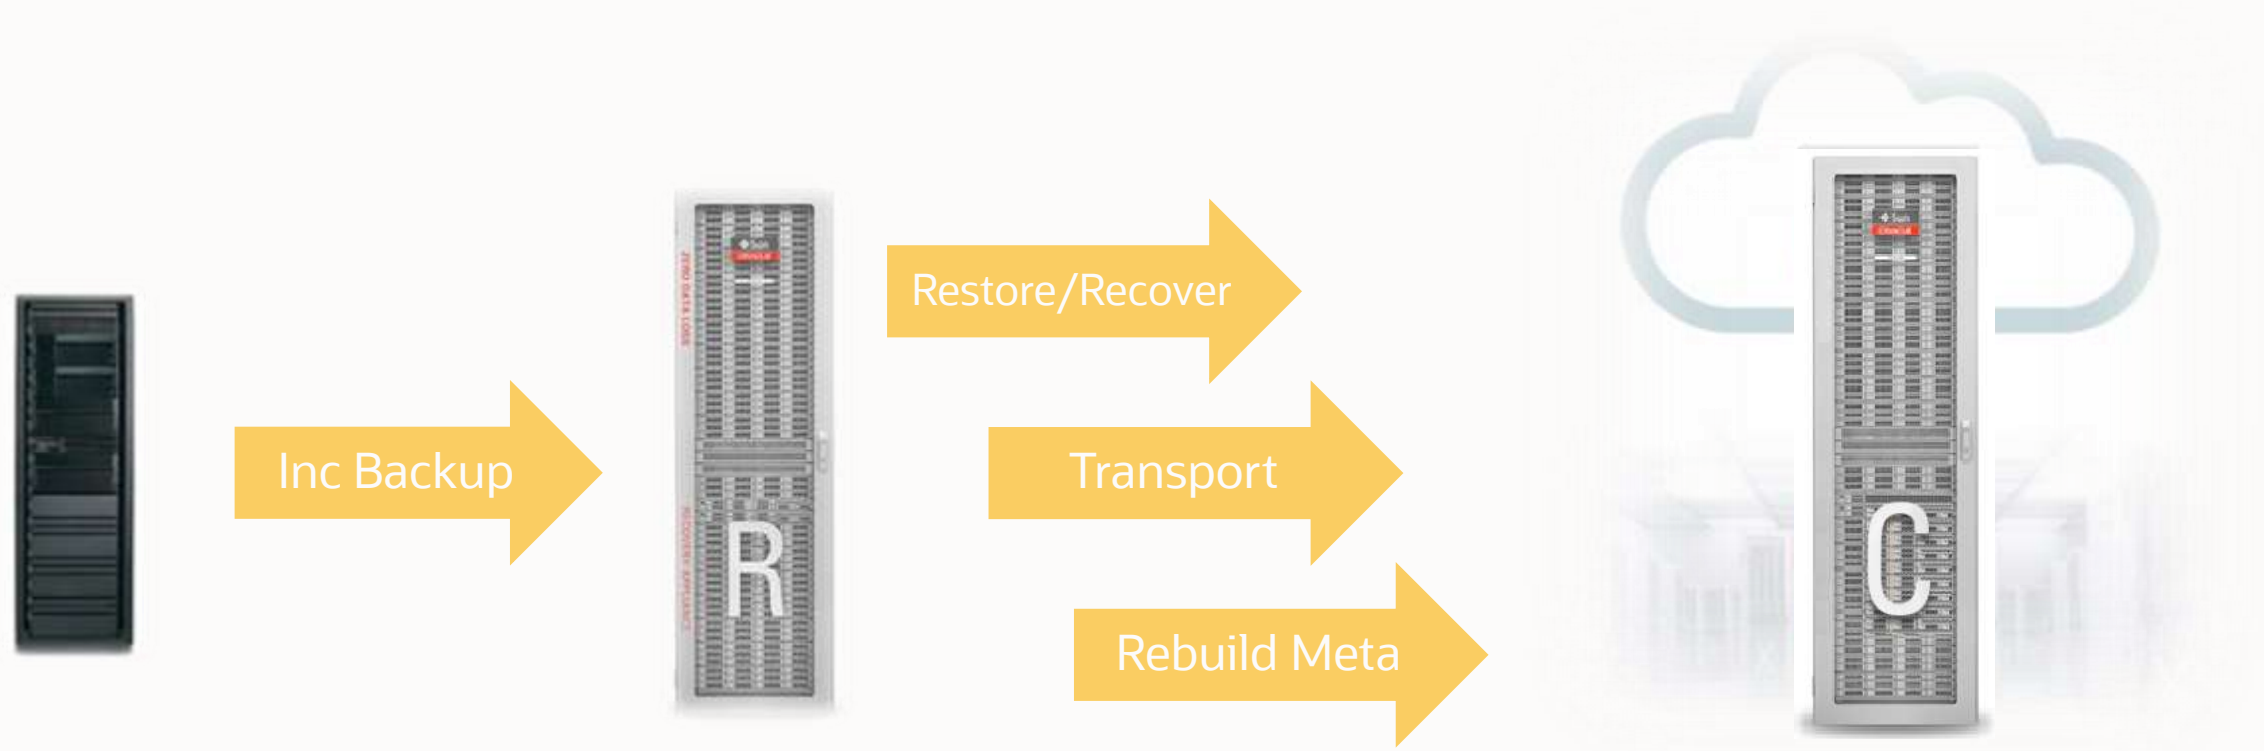

ExaCC Migration | Timeline

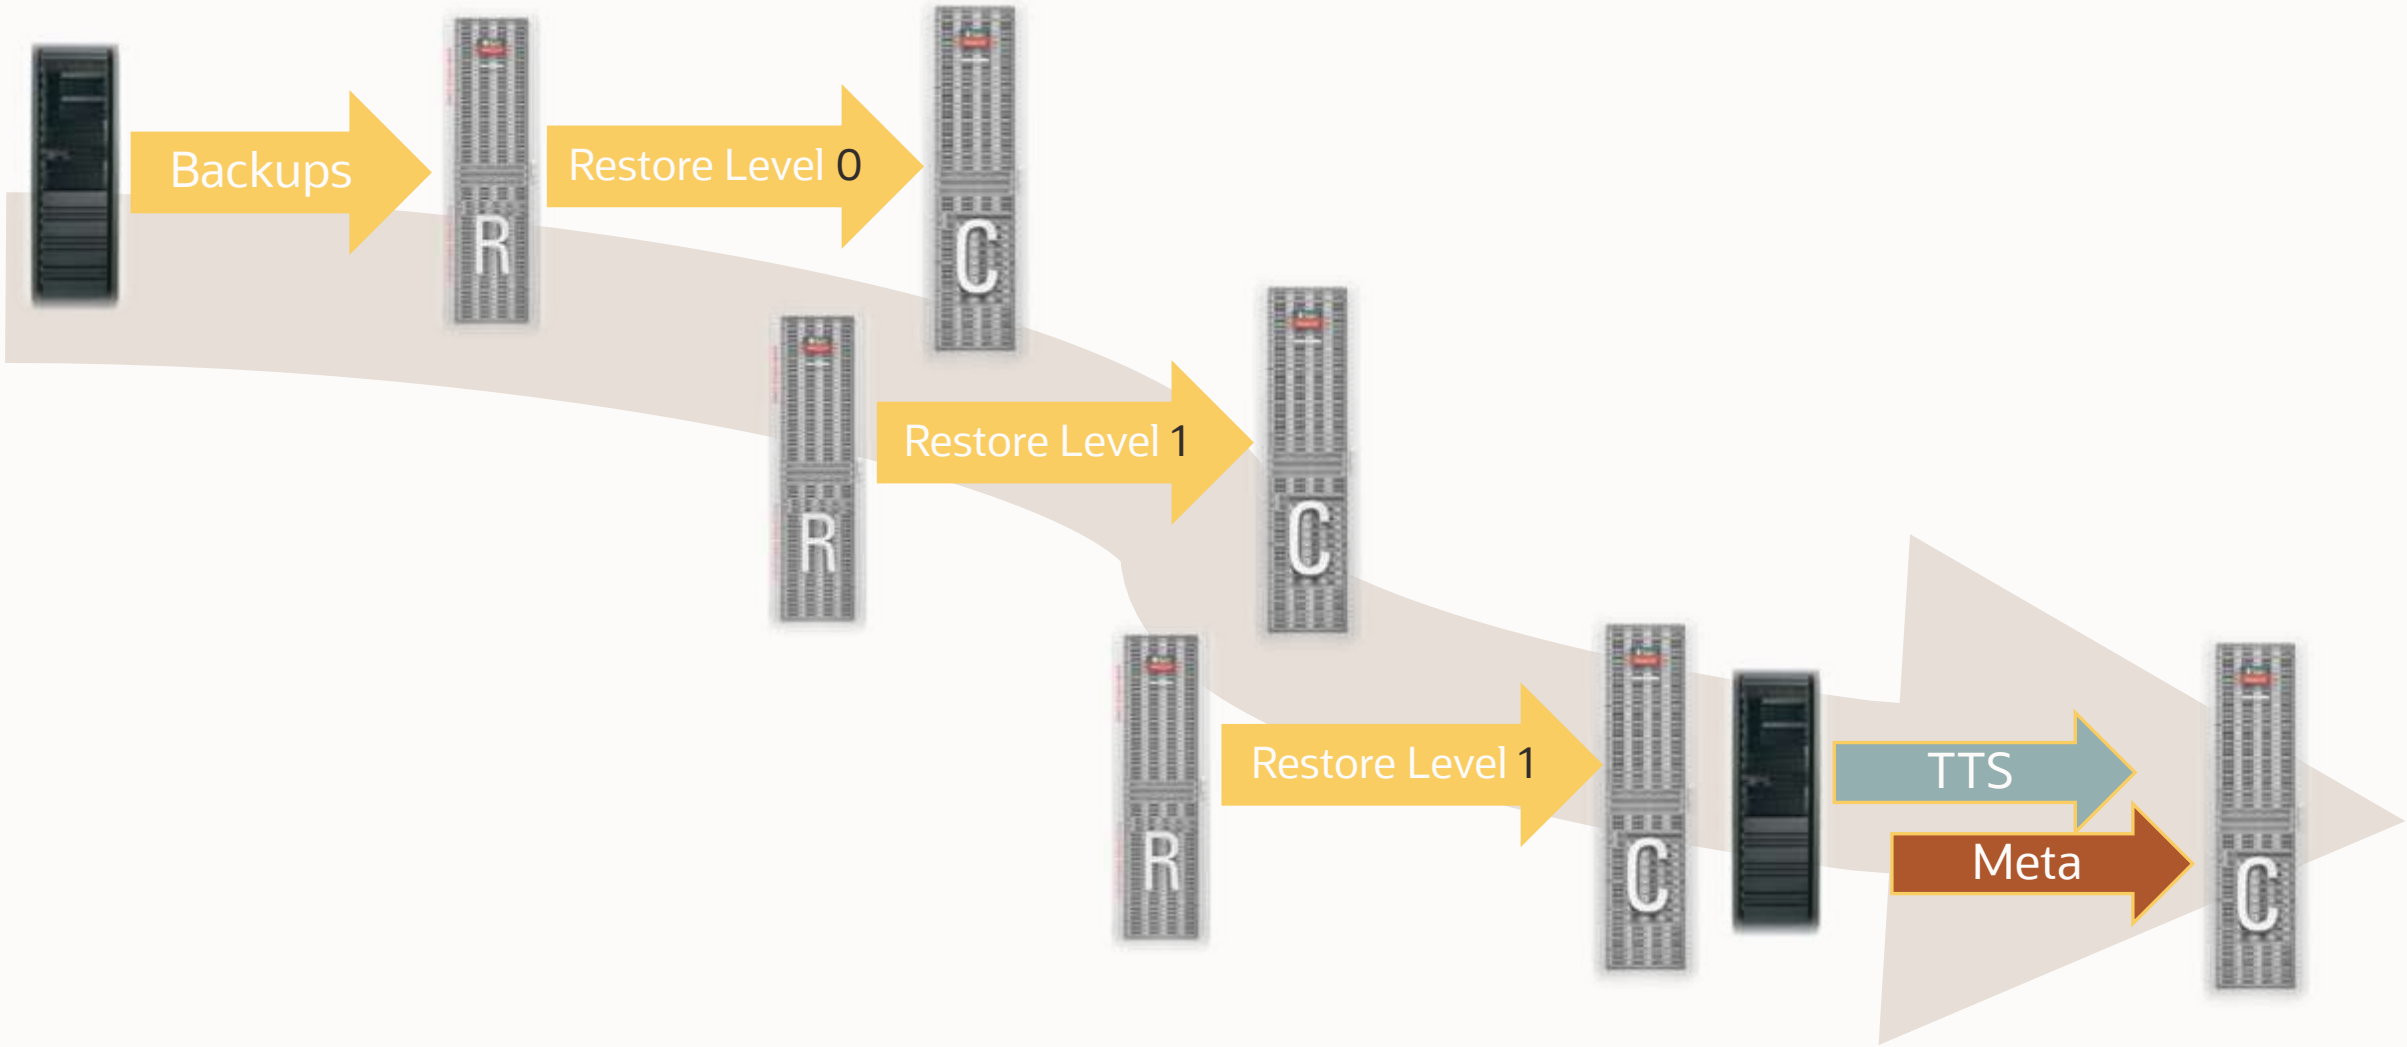

Migration Details

Part 1: The ZDLRA

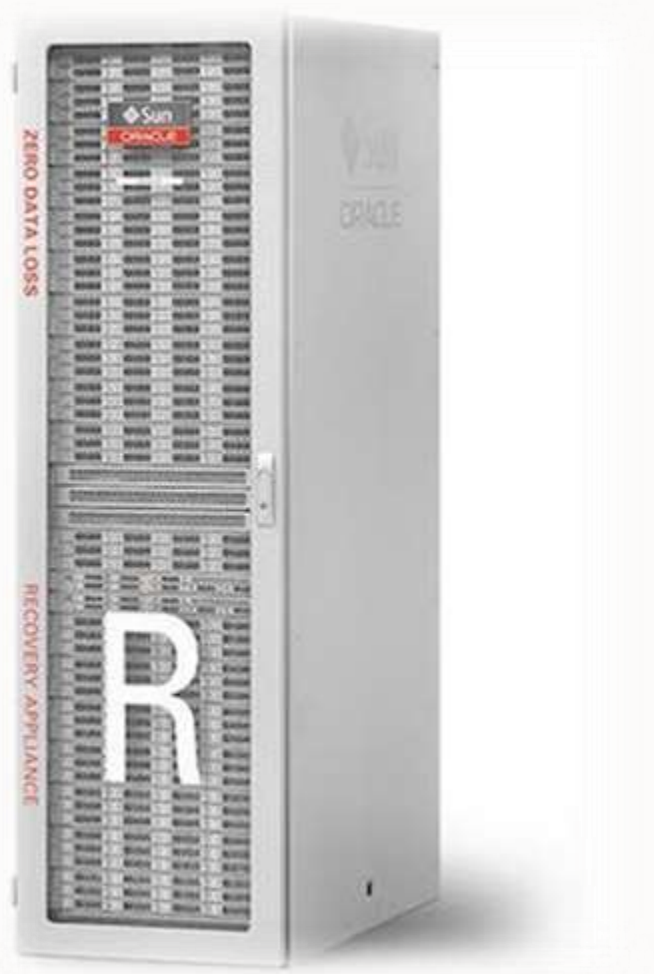

ExaCC Migration | PERL Scripts

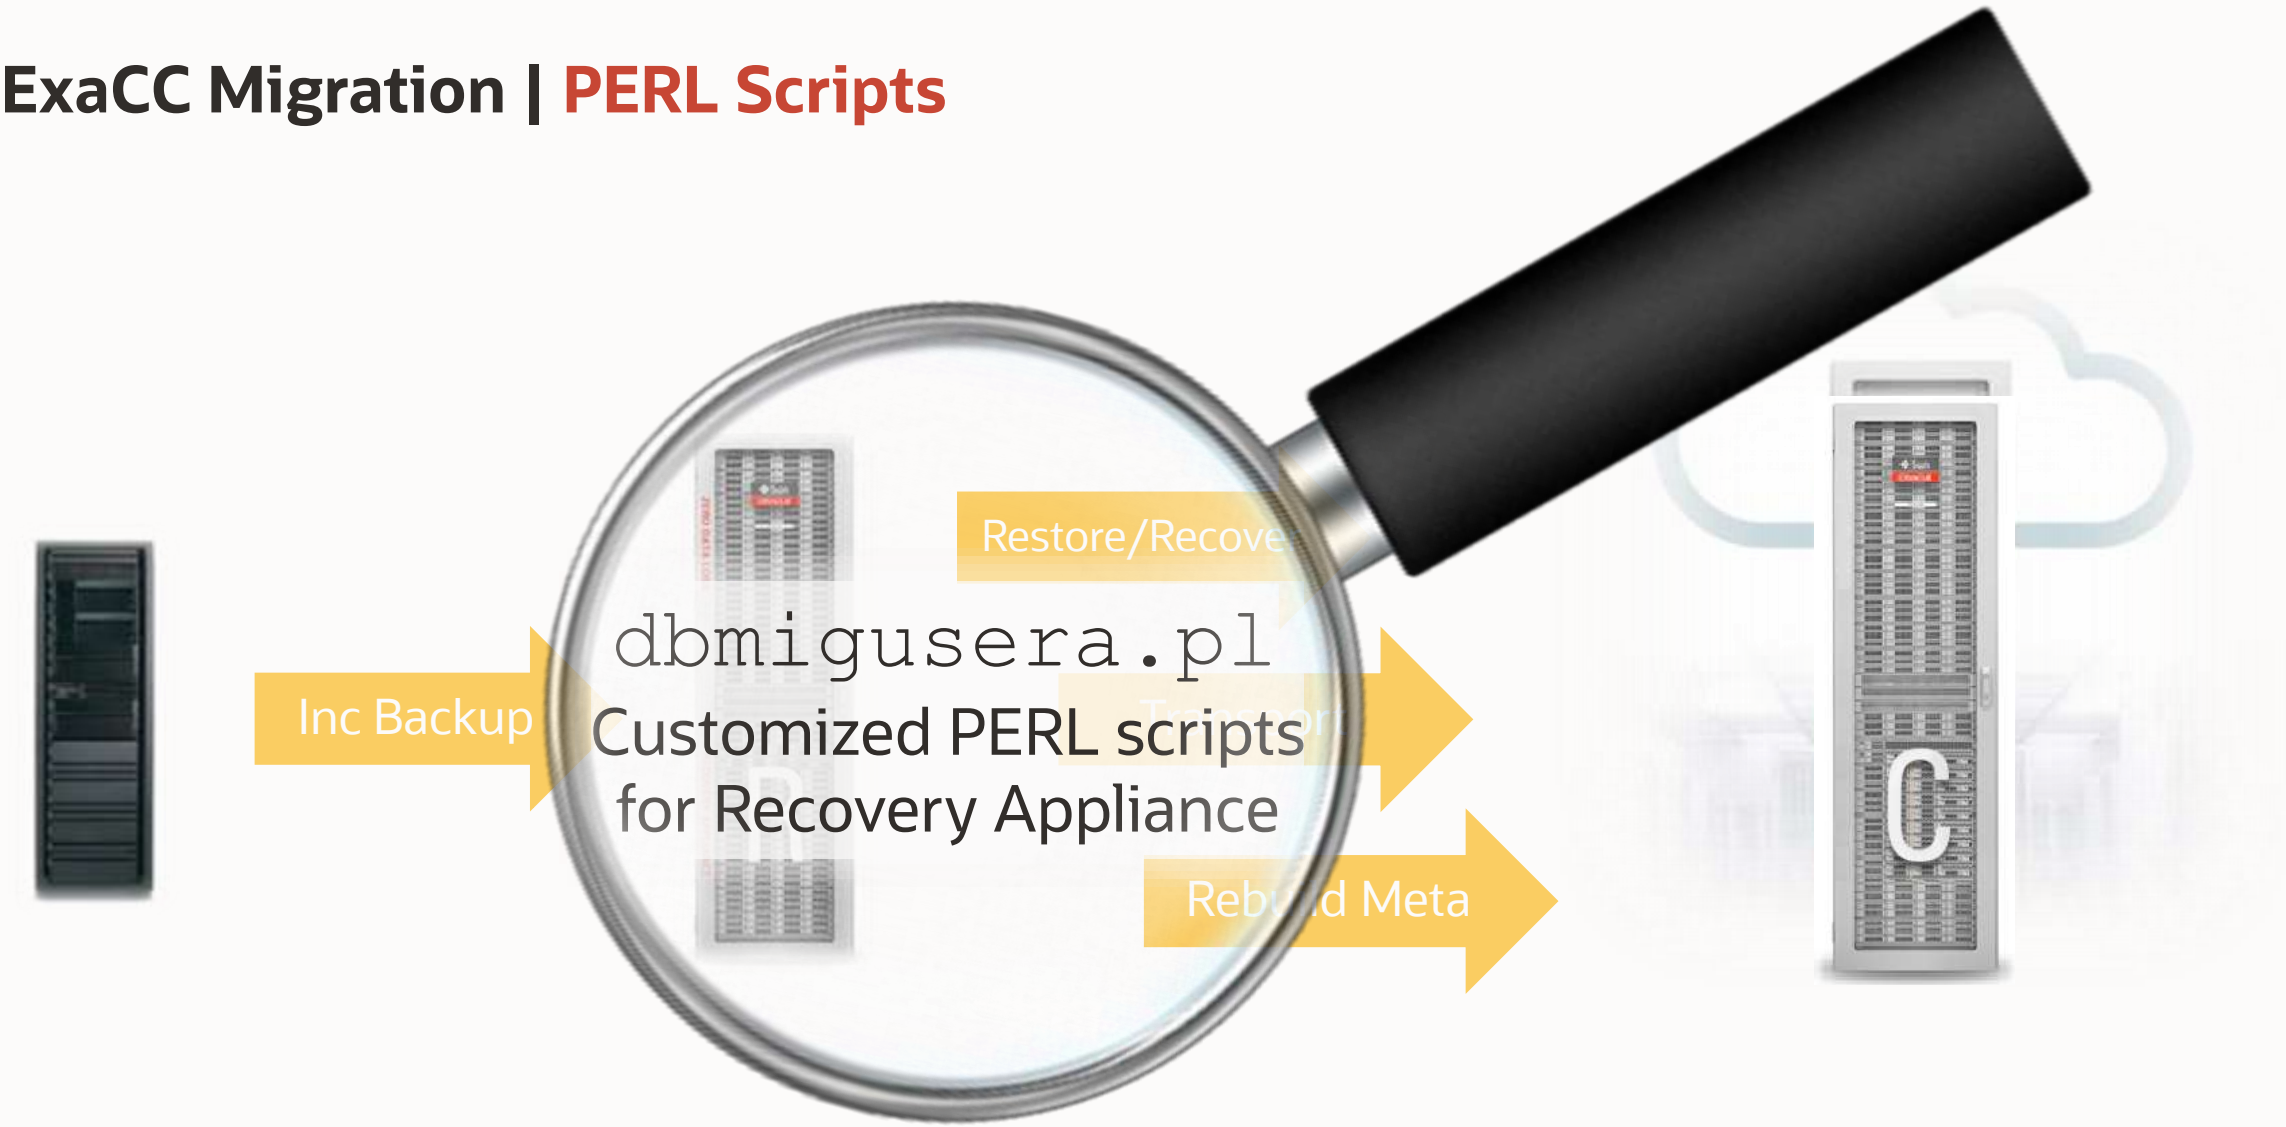

ExaCC Migration | libra.so

Install most recent `libra.so`

- MOS Note: 2219812.1
  ZDLRA: Download new sbt library

Recovery Appliance sbt library download locations:
==========
RA HPUX-IA64
RA Linux64
RA AIX-PPC64
RA Solaris-Sparc64
RA ZLinux64
RA Windows64
RA SolarisX64

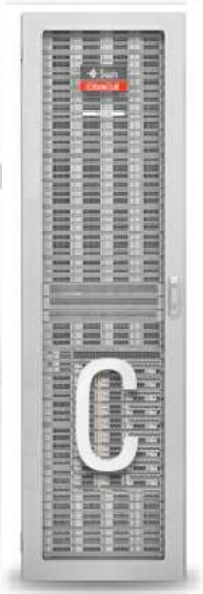

ExaCC Migration | dbmigusera.pl

Download package `dbmigusera.pl` incl. `xtt.properties`

- MOS Note: 2460552.1 - Cross Platform Database Migration using ZDLRA
- Deploy package

ExaCC Migration | xtt.properties

Customize `xtt.properties`

```
# SBT parameter configuration to be used for restore and recover operations
sbtlibparms="SBT_LIBRARY=/u01/app/oracle/product/12.2.0.1/dbhome_1/lib/libra.so, …')"

# The number of datafiles that will be restored / recovered in parallel.
resparallel=16

# The list of tablespaces to be migrated
ttsnames=T14_1,T14_2,T14_3,T14_4,T14_5,T14_6,T14_7,T14_8,T14_9,T14_10,T14_11

# The directory to which the datafiles will be restored
storageondest=+DATAC1/dbmig/datafile

sourceplatid=2
dbid=4173218531
retrycount=2
```

ExaCC Migration | Level 0 Backup

Level 0 Backup

Restore/Recover

```
$ perl dbmigusera.pl --restore …
```

ExaCC Migration | Level 1 Backup

Inc Backup

Restore/Recover

```
$ perl dbmiguser.pl --recover
```

Migration Details

Part 2: Transportable Tablespaces to ExaCC

ExaCC Migration | Transportable Tablespaces

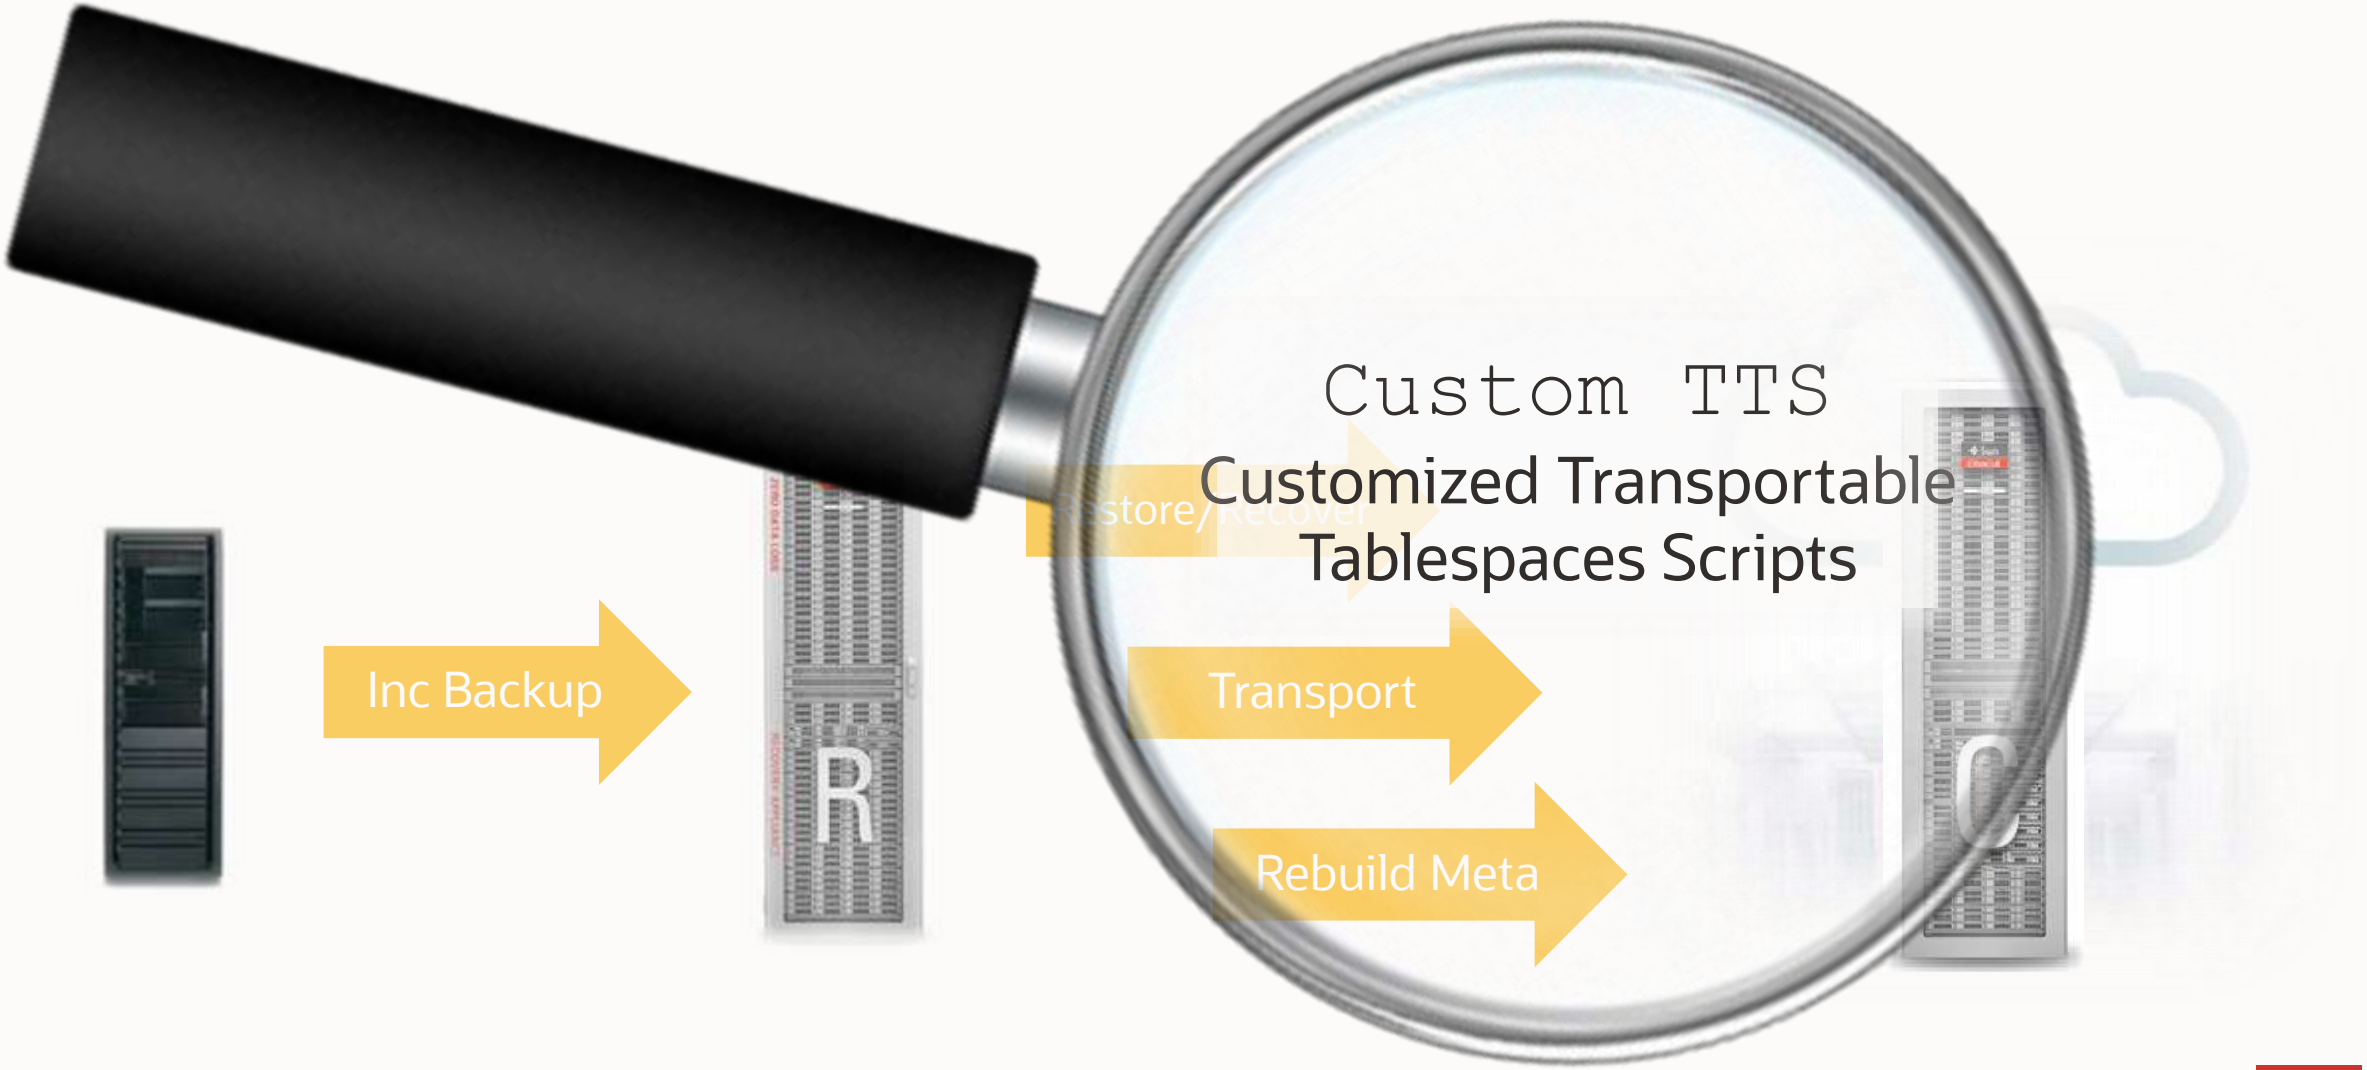

Photo by Alexander Andrews on Unsplash

Transport

New and advanced features

Data Pump | Test Mode for Transportable Tablespaces

Test TTS/FTEX export without "read only"

- TTS_CLOSURE_CHECK = ON | OFF | FULL | TEST_MODE

```
DIRECTORY=DP_DIR
DUMPFILE=tts.dmp
LOGFILE=logfile.log
TTS_CLOSURE_CHECK=TEST_MODE
TRANSPORT_TABLESPACES=(TTS)
```

NEW IN 19c

Data Pump | Keep Tablespaces Read-Only for TTS

Allow "read only" tablespaces for TTS import

- TRANSPORTABLE=NEVER|ALWAYS|KEEP_READ_ONLY|NO_BITMAP_REBUILD

```
DIRECTORY=DP_DIR
DUMPFILE=tts.dmp
LOGFILE=logfile.log
TRANSPORT_DATAFILES='/CDB2/pdb1/tts.dbf'
TRANSPORTABLE=KEEP_READ_ONLY
```

```
TABLESPACE_NAME    STATUS
------------------ ----------
SYSTEM             ONLINE
SYSAUX             ONLINE
UNDOTBS1           ONLINE
TEMP               ONLINE
TTS                READ ONLY
```

NEW IN 19c

Migration | ZDRLA

Learn the latest techniques using Recovery Appliance to simplify and speed up cross-platform database migration activities, including moving from non-Multitenant to modern Multitenant Database architectures.

MOS note
Cross Platform Database Migration using ZDLRA (Doc ID 2460552.1)

AskTOM Office Hours

Accelerate Cross-Platform Database Migration with Recovery Appliance

Transportable | Important MOS Notes

- 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)

zero downtime

MIGRATION

technique

RMAN

Trans-portable

Data Guard

Golden Gate

GoldenGate | Concept

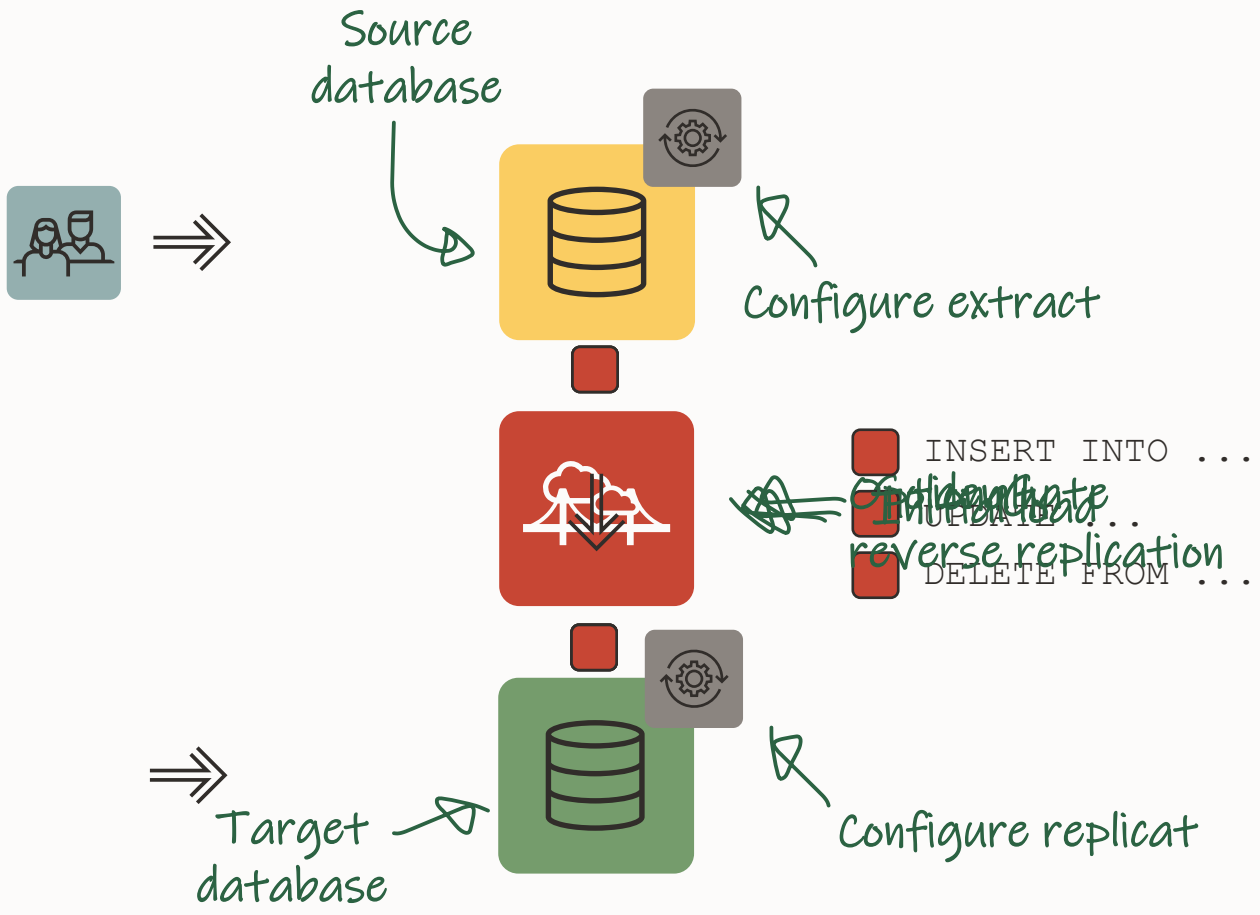

GoldenGate | Benefits

- True zero downtime
- Extremely flexible
- Cross-version and cross-architecture
- Cross-platform and cross-endian
- Test before go-live using Flashback Database (Doc ID 966212.1)

Pro tip: Active Data Guard included in GoldenGate license

GoldenGate | Architecture

- SQL*Net connection between databases
  - Alternatively, a distribution path between two GoldenGate instances
- Compress trail files to reduce network load
- Example: Database generating 10 TB redo
  - Trail files typically 30-40 % of redo
  - Compress at least 1:4, most likely up to 1:8
  - Result: 400 GB to 1000 GB trail files

GoldenGate | Initial load

If your initial load is a Data Pump import:

- 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
  - …

To strengthen security, you can encrypt the GoldenGate trail files

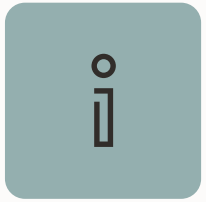

GoldenGate requires
database minimal supplemental logging
which does not impose a significant overhead

GoldenGate | 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 to support GoldenGate
  - 11g
  - 12c and newer

- DDL replication
  - Truncate
  - Sequences

Can you use GoldenGate on your database?

GoldenGate | Database Readiness

```
SQL> select * from dba_goldengate_support_mode;

OWNER OBJECT_NAME                    SUPPORT_MODE
CO    CUSTOMERS                      ID KEY
CO    ORDERS                         ID KEY
CO    ORDER_ITEMS                    FULL
CO    PRODUCTS                       ID KEY
CO    STORES                         ID KEY
```

GoldenGate | Database Readiness

SUPPORT_MODE

FULL

ID KEY
PLSQL

INTERNAL
NONE

Pro tip: Visit the documentation for more details

GoldenGate | Database Readiness

SUPPORT_MODE	FULL	ID KEY PLSQL	INTERNAL NONE

Pro tip: Visit the documentation for more details

GoldenGate | Database Readiness

SUPPORT_MODE	FULL	ID KEY PLSQL	INTERNAL NONE

Pro tip: Visit the documentation for more details

GoldenGate | Database Readiness

What's wrong in this Oracle Database running 12.2?

Identify columns supported as of Oracle Database 18c …

```
SQL> select * from dba_goldengate_support_mode;

OWNER OBJECT_NAME          SUPPORT_MODE
CO    CUSTOMERS            ID KEY
CO    ORDERS               ID KEY
CO    ORDER_ITEMS          FULL
CO    PRODUCTS             ID KEY
CO    STORES               ID KEY
```

GoldenGate | Database Readiness

Oracle Database 21c New Feature

```
SQL> select * from dba_goldengate_support_mode;

OWNER OBJECT_NAME          SUPPORT_MODE       DESCRIPTION
CO    CUSTOMERS            ID KEY             A very good explanation
CO    ORDERS               ID KEY             Another good explanation
CO    ORDER_ITEMS          FULL
CO    PRODUCTS             ID KEY             A third explanation
CO    STORES               ID KEY             Good explanation comes in
                                              abundance these days
```

GoldenGate | Database Readiness

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

GoldenGate | Database Readiness

```
SQL> select * from dba_goldengate_not_unique;

OWNER  TABLE_NAME                          BAD_COLUMN
IX     AQ$_ORDERS_QUEUETABLE_L             N
IX     AQ$_STREAMS_QUEUE_TABLE_L           N
SH     SALES                               N
SH     COSTS                               N
SH     SUPPLEMENTARY_DEMOGRAPHICS          N
SH     CAL_MONTH_SALES_MV                  N
SH     FWEEK_PSCAT_SALES_MV                N
```


GoldenGate | Database Readiness

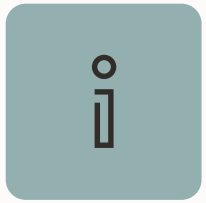

If the application maintains uniqueness,
but it is not enforced on the database,
use a `KEYCOLS` clause to let GoldenGate use it

Pro tip: For further information, read Ensuring Row Uniqueness in Source and Target Tables

GoldenGate | Health Check

Generate report:

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

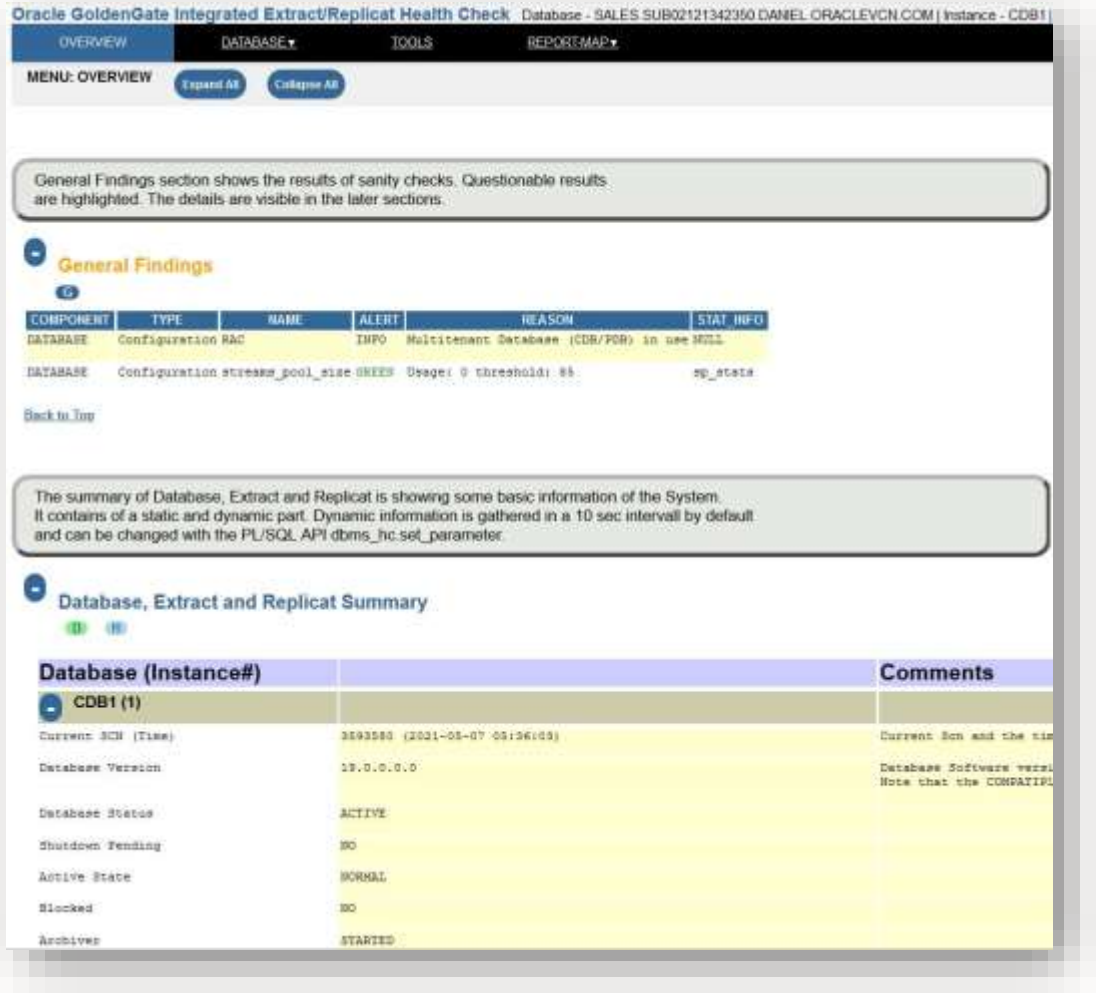

GoldenGate | Health Check

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`

How about upgrading?

GoldenGate can extract from one release, and replicate into another

Pro tip: You can even migrate from very old releases using multiple instances of GoldenGate

How about PDB conversion?

GoldenGate can replicate
from non-CDB directly into a PDB

Can I offload the work
from the source database?

Yes, you can extract from a *downstream* database

GoldenGate | Downstream

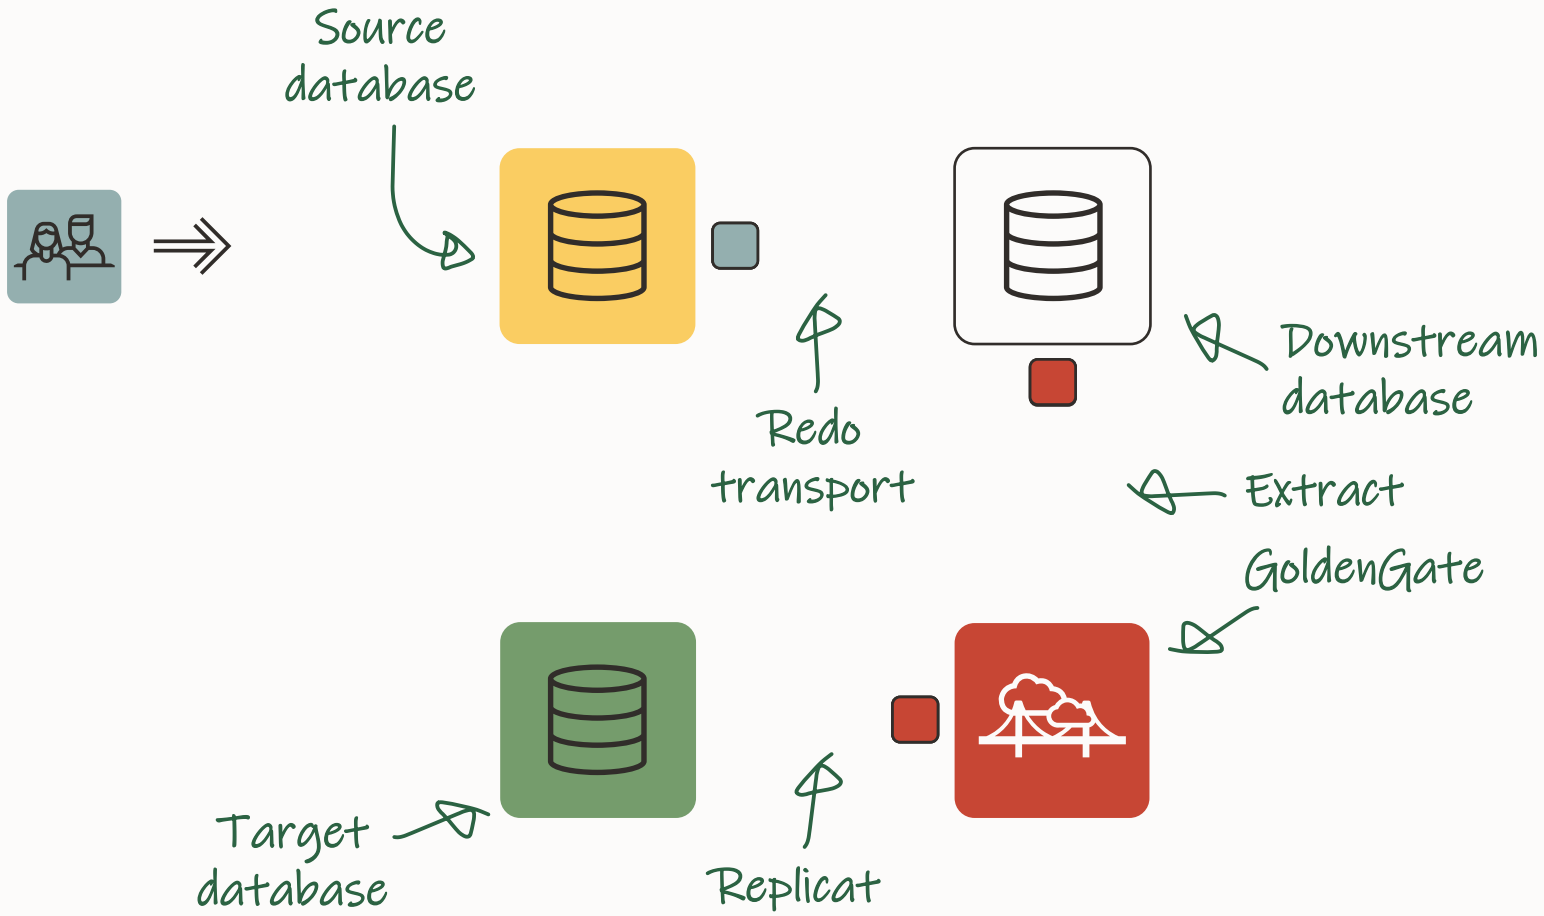

Your target database must be protected by Data Guard?

After the initial load on the target database, start building your Data Guard

GoldenGate | Data Guard

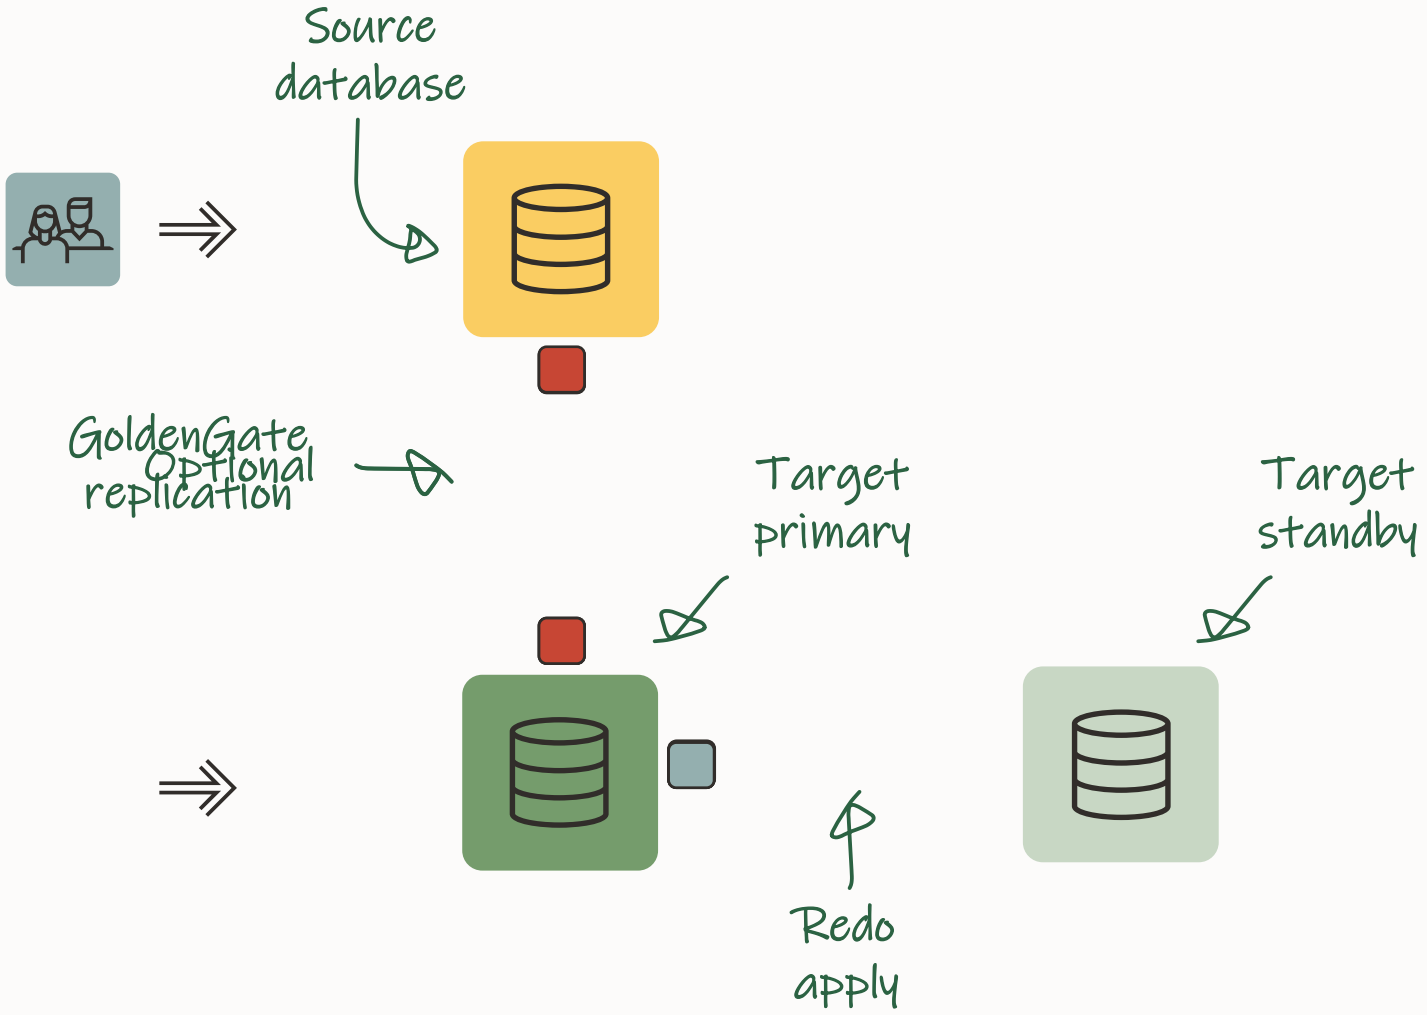

Your target database must
have a valid backup before go-live?

Perform and verify your backups
after initial load, but before switchover

GoldenGate | Backup

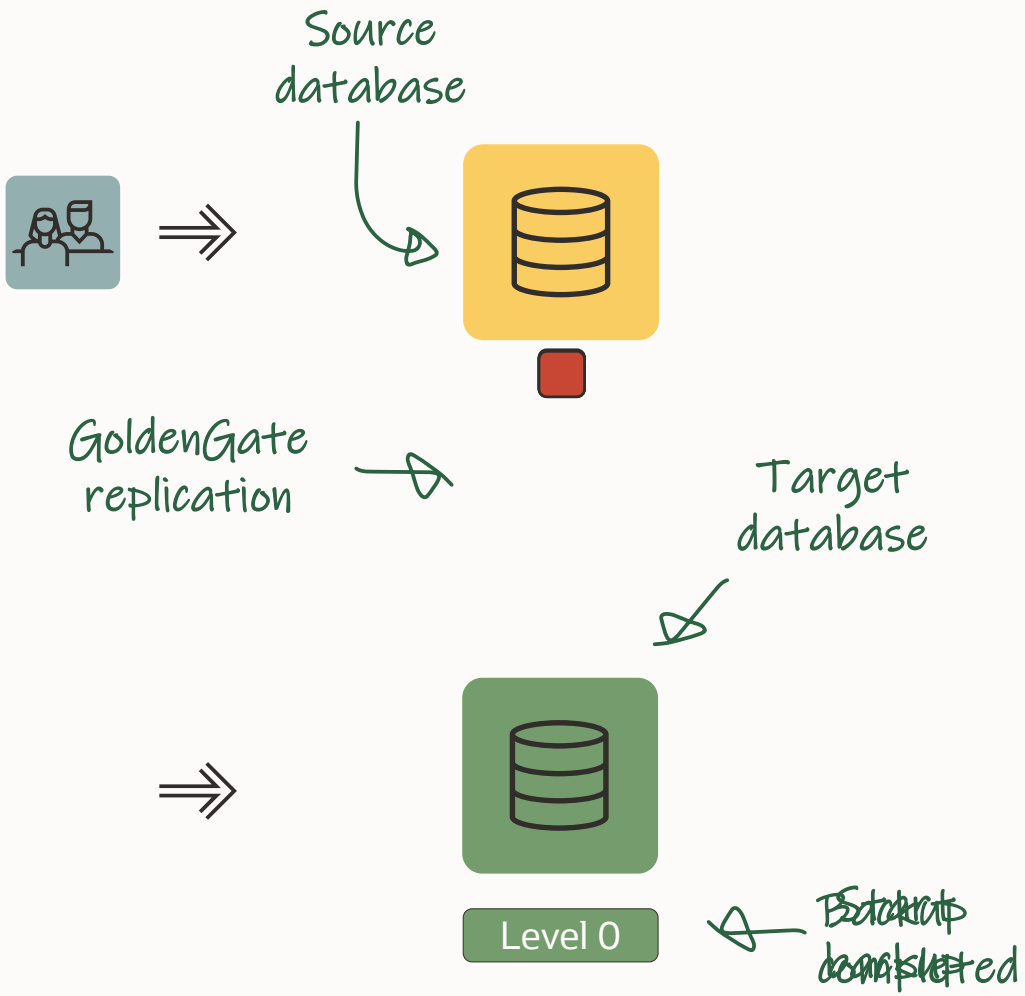

Your target database must be RAC?

No problem, GoldenGate can extract from and replicate to a RAC database

How about your fallback plan?

GoldenGate | Fallback

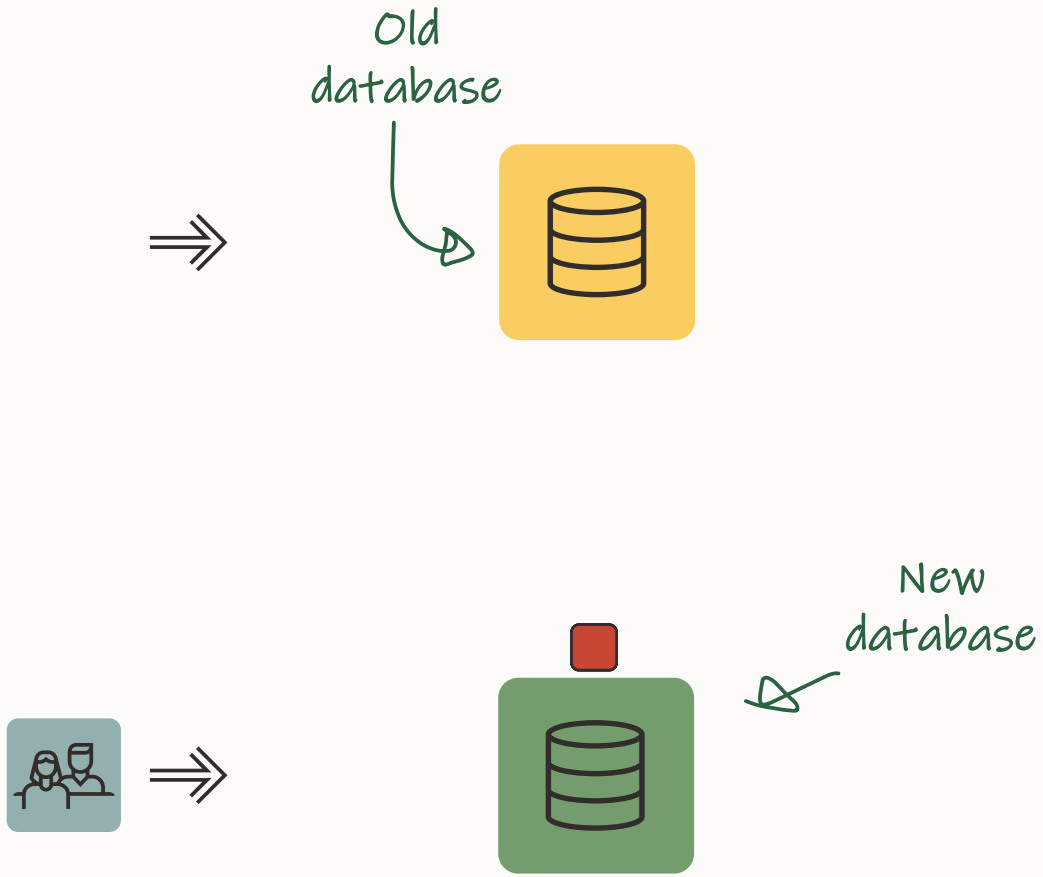

GoldenGate | Fallback

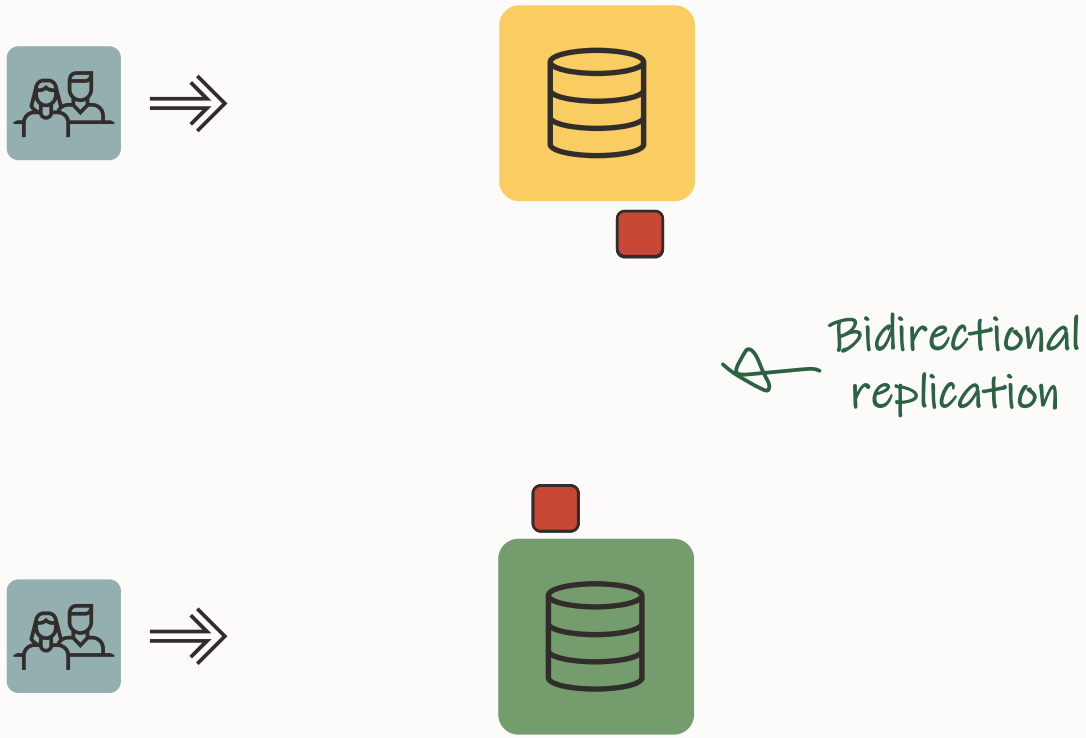

GoldenGate | Additional Resources

Certifications

GoldenGate 19.1: Using Oracle GoldenGate on Oracle Cloud Marketplace

OCI Marketplace: Oracle GoldenGate for Oracle

Oracle GoldenGate Best Practices: Instantiation from an Oracle Source Database (Doc ID 1276058.1)

GoldenGate | Technical Briefs

Oracle Database Migration with an Oracle GoldenGate Hub Configuration

Zero Downtime Database Upgrade Using Oracle GoldenGate

Oracle GoldenGate with Oracle RAC Configuration Best Practices

Transparent Role Transitions With Oracle Data Guard and Oracle GoldenGate

SUMMARY

Which one is the best technique?

It still depends …

Migration | Techniques

Data Guard

FTEX

RMAN

TTS

OGG

Webinar | Upcoming

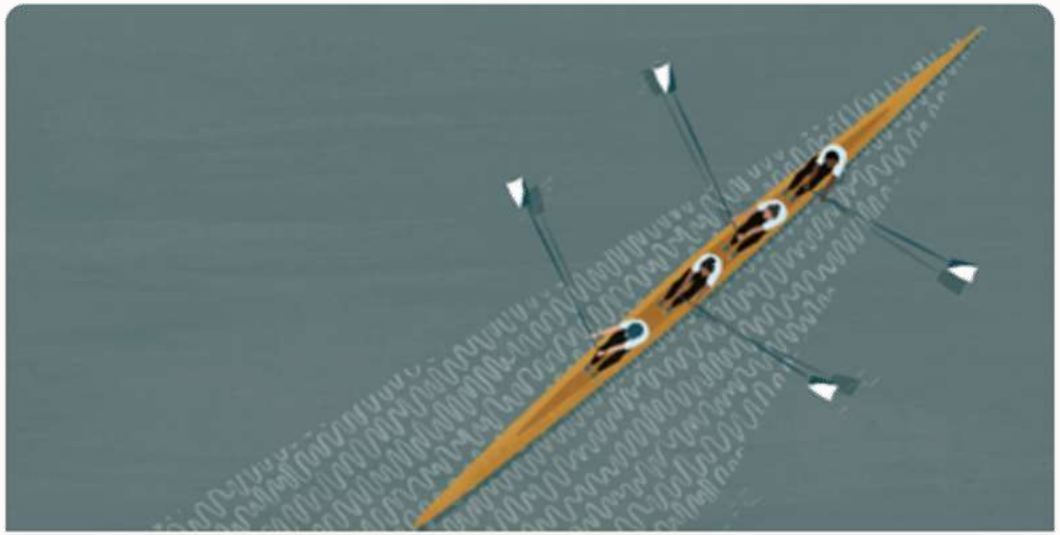

Data Pump Extreme - Deep Dive with Development

March 3, 2022 | 09:00 GMT / 10:00 CET / 11:00 EET / 13:00 GST

Duration: 120 mins

SCAN ME

https://go.oracle.com/LP=114938?elqCampaignId=302203

105 minutes – Feb 4, 2021

Episode 2

AutoUpgrade to Oracle Database 19c

115 minutes – Feb 20, 2021

Episode 3

Performance Stability, Tips and Tricks and Underscores

120 minutes – Mar 4, 2021

Episode 4

Migration to Oracle Multitenant

120 minutes – Mar 16, 2021

Episode 5

Migration Strategies – Insights, Tips and Secrets

120 minutes – Mar 25, 2021

Episode 6

Move to the Cloud – Not only for techies

115 minutes – Apr 8, 2021

Episode 7

Cool Features – Not only for DBAs

110 minutes – Jan 14, 2021

Episode 8

Database Upgrade Internals – and so much more

110 minutes – Feb 11, 2021

Episode 9

Performance Testing Using the Oracle Cloud for Upgrades and Migrations

90 minutes – May 19, 2021

NEW Episode 10

How Low Can You Go? Minimal Downtime Upgrade Strategies

100 minutes – Oct 26, 2021

Recorded Web Seminars

https://MikeDietrichDE.com/videos/

https://dohdatabase.com/webinars/

YouTube | Oracle Database Upgrades and Migrations

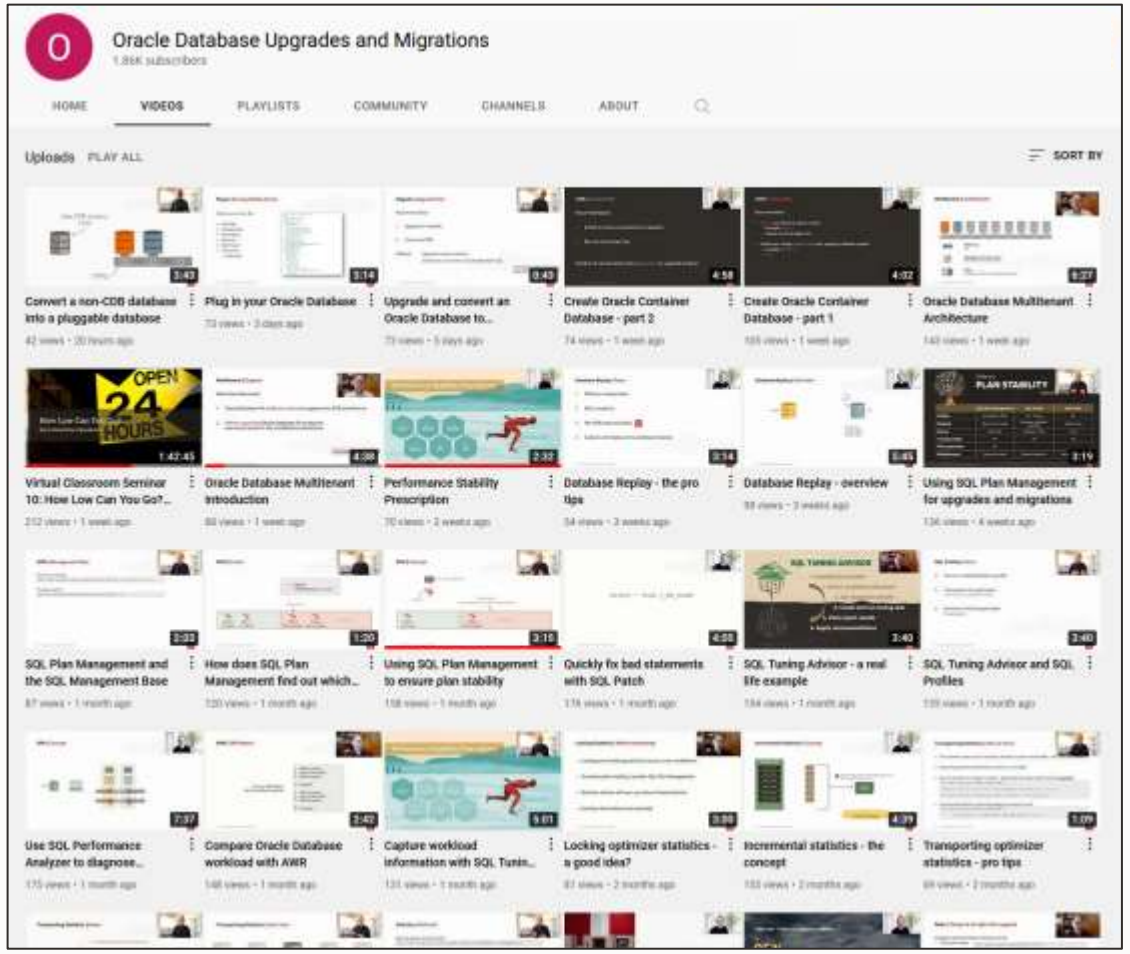

Link

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

SCAN ME

THANK YOU

Visit our blogs:

https://MikeDietrichDE.com

https://DOHdatabase.com

https://www.dbarj.com.br/en

THANK YOU

Webinars:

https://MikeDietrichDE.com/videos

YouTube channel:

OracleDatabaseUpgradesandMigrations

THANK YOU

DATA PUMP EXTREME

Deep Dive with Development

March 3, 2022 – 10:00h CET

THANK YOU