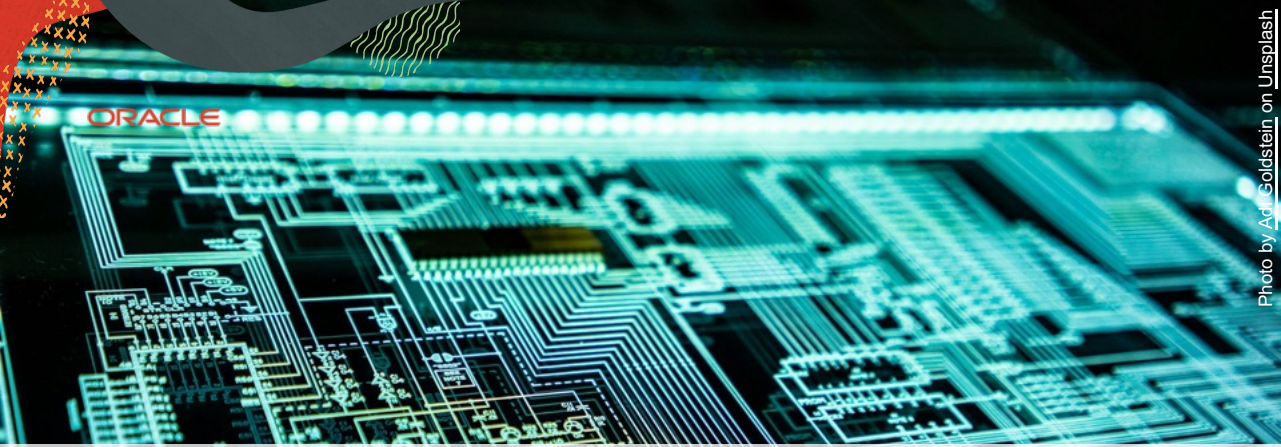

Photo by Adi Goldstein on Unsplash

The Data Pump Whisperer

DOAG 2025 Konferenz + Ausstellung, November 2025

Mike Dietrich

Vice President

mikedietrich

@mikedietrichde.com

https://mikedietrichde.com

Daniel Overby Hansen

Distinguished Product Manager

dohdatabase

@dohdatabase.com

https://dohdatabase.com

Best Practices

Get the most out of Data Pump

Always use the Data Pump Bundle Patch

More than 230 functional and performance fixes

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

A global provider of financial services

Apply the Data Pump Bundle Patch without downtime

DEMO

Apply Data Pump Bundle Patch

Watch on YouTube

Data Pump Bundle Patch

- Data Pump Recommended Proactive Patches For 19.10 and Above (Doc ID 2819284.1)
- The patch is not RAC rolling installable
  - The patch is non-binary online installable
  - Apply while the database instance is running
  - Don't use Data Pump or `DBMS_METADATA`

Ensure dictionary and fixed objects statistics are accurate

- Before export
- Before import
- Immediately after import

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

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

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

Use parallel and multiple dump files

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


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

Oracle Cloud Infrastructure

Number of ECPUs / 4

Number of OCPUs

On-prem (x86-64)

2 x physical cores

On-prem (other)

Depends

Use at least 32 ECPUs when importing into Autonomous Database

- Use the elasticity in Autonomous Database and scale even higher

Parallel Architecture

```
expdp ... parallel=4
```

Parallel Architecture

Parallel Architecture

```
expdp ... parallel=4
```

Control process

Worker processes

```
select /*+ parallel(2) */ * from t1
select * from t2
select * from t3
select * from t4
```

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


```
--Split dump files into minor files for easier transport
expdp ... parallel=8 dumpfile=exp%L.dmp filesize=10000M
```

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

expdp ... checksum=yes

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


For best protection against dump file tampering, use encrypted dump files

- Requires Advanced Security Option

```
-- Protect your dump files from alteration by using encryption
-- Creating an encrypted dump file requires Advanced Security Option

expdp ... encryption=all encryption_algorithm=AES256
```

```
-- Protect your dump files from alteration by using encryption
-- Creating an encrypted dump file requires Advanced Security Option

expdp ... encryption=all encryption_algorithm=AES256
```

DEMO

Encrypted exports

Watch on YouTube

Transportable jobs can use parallel in Oracle Database 21c

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

expdp ... full=y transportable=always parallel=16
expdp ... tablespace=<list> parallel=16

impdp ... parallel=16
```

Parallel Transportable Benchmark

Oracle E-Business Suite database
600.000+ objects

Export parallel 1	2h 2m
Import parallel 1	6h 44m
Total	**8h 46m**

Export parallel 16	1h 8m
Import parallel 16	1h 23m
Total	**2h 31m**

New Features

Take your exports and imports to the next level

Speed up imports by using NOVALIDATE constraints

A Constraint Can Be

VALIDATED

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

NOT VALIDATED

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

Most constraints are VALIDATED

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

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

create table sales ( .... );

insert into sales as select ... ;

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

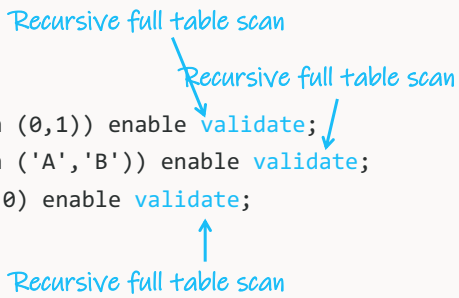

```
-- Add constraints with NOVALIDATE keyword regardless of state in source database
-- Significantly speeds up add constraints for larger tables

impdp ... transform=constraint_novalidate:y
```

```
--Transforming constraints to NOVALIDATE to speed up import

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

```
--Transforming constraints to NOVALIDATE to speed up import

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

Database validates new rows

Benchmark, 1 billion rows

Importing VALIDATE constraints

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

Importing NOVALIDATE constraints

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

Real-World Example, Importing 180G Schema

Starting point	44m 50s
Adding Data Pump Bundle Patch	36m 53s
Using NOVALIDATE constraints	7m 25s

NOVALIDATE constraints prevent the optimizer from certain query rewrites

Validate constraints after import, or even after go-live

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

```
alter table sales add constraint c_sales_1 check (c1 in (0,1)) enable novalidate;

-----------------------
------- GO LIVE -------
-----------------------

-- Validate constraints
-- Optionally, use parallel query

alter session force parallel query;
alter table sales modify constraint c_sales_1 enable validate;
```

Also available in Oracle Database 19c via 19.27 Data Pump Bundle Patch

Also available in Oracle Autonomous Database 19c

```
--Add the following to your ZDM response file to set the parameter during
--ADB migrations using ZDM

DATAPUMPSETTINGS_METADATATRANSFORMS-1=name:CONSTR_NOVALIDATE,value:1
```

NOVALIDATE Constraints

- 19.26 and earlier also requires patch 37280692
- Allow optimizer to perform query rewrites even with NOVALIDATE constraints using QUERY_REWRITE_INTEGRITY
- Data Pump always validates certain constraints:
  - On DEFAULT ON NULL columns
  - Used by a reference partitioned table
  - Used by a reference partitioned child table
  - Table with Primary key OID
  - Used as clustering key on a clustered table

NOVALIDATE Constraints

- Available in Database Migration Service (DMS)
- In the OCI console, Advanced parameters> DATAPUMPSETTINGS: Metadata Transforms

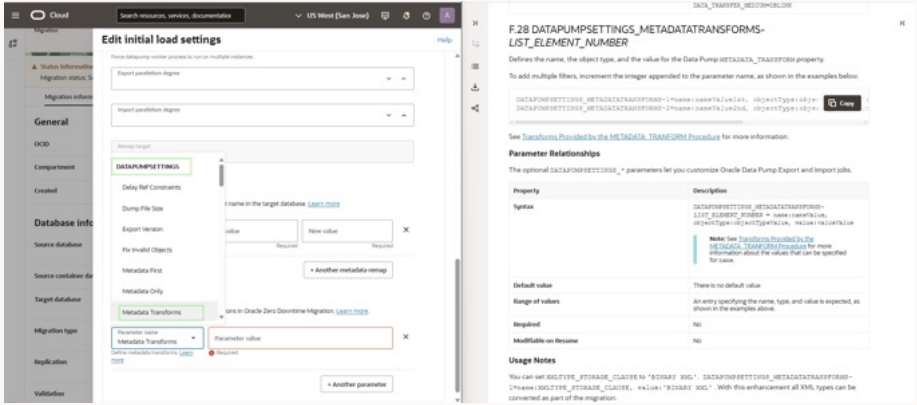

Use index size to determine parallel degree on index creation

Index Creation

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

Really good for few big indexes

Before 12.1

Worker 1 `CREATE INDEX .... PARALLEL 16`

Index Creation

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

From 12.1

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

Index Creation

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

From 23

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

How Data Pump Create Indexes

1. Calculate the optimal parallel degree
2. Create indexes

How Data Pump Create Indexes

1 Calculate the optimal parallel degree

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

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

Explained.
```


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

Explained.

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

...

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

How Data Pump Create Indexes

2 Create indexes

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

How Data Pump Create Indexes

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

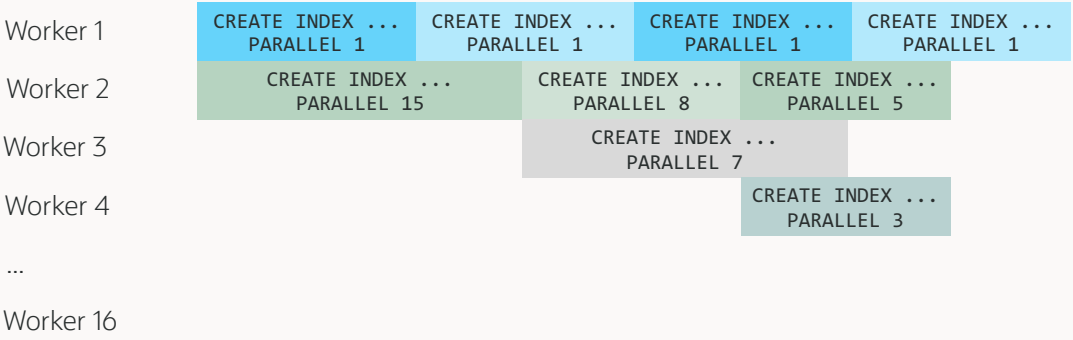

Importing with former index method

```
10-MAY-25 16:18:55.130: W-12 Processing object type SCHEMA_EXPORT/TABLE/INDEX/INDEX
10-MAY-25 16:36:46.902: W-30      Completed 480 INDEX objects in 1071 seconds
```

Importing with new index method

```
10-MAY-25 16:47:50.267: W-4 Processing object type SCHEMA_EXPORT/TABLE/INDEX/INDEX
10-MAY-25 16:59:17.006: W-3      Completed 480 INDEX objects in 686 seconds
```

Index Creation

Requires:

- Oracle AI Database 26ai
- Oracle Database 23.8 plus Data Pump Bundle Patch
- Oracle Database 19.26 plus Data Pump Bundle Patch

Also available in Oracle Database 19c via 19.26 Data Pump Bundle Patch

Convert time zone data on import

Time Zone File Version Check

Source
Version 43

Target
Version 42

```
create table t1 (
   ...
   c1 timestamp with timezone
   ...
)
```

```
Import: Release 19.0.0.0.0 - Production on Sun Sep 30 06:17:06 2025
Version 19.27.0.0.0

ORA-39002: invalid operation
ORA-39405: Oracle Data Pump does not support importing from a source database
with TSTZ version 43 into a target database with TSTZ version 42.
```

DEMO

Convert time zone data on import

Watch on YouTube

```
INSERT /*+ APPEND ENABLE_PARALLEL_DML ... */ INTO RELATIONAL( .... )
  SELECT ORA_DST_CONVERT("C1")
   FROM "TESTUSER"."ET$01145D870001" KU$
```

Time Zone Data Convert

- If needed, Data Pump converts `TIMESTAMP WITH TIMEZONE` to the target time zone file version.
- This works if the target time zone file is higher or lower than the source.
- Expect a small overhead for the conversion.
- Requires 19.27 including Data Pump Bundle Patch.

Also available in Oracle Database 19c via 19.27 Data Pump Bundle Patch

Copy the dump files before the export completes

- Oracle Database 23ai

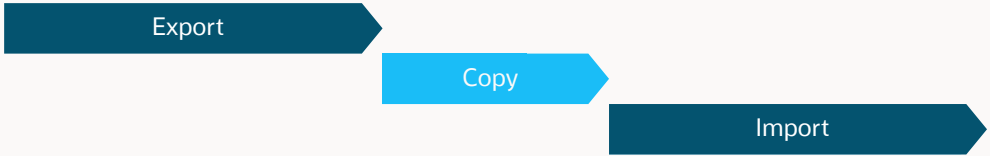

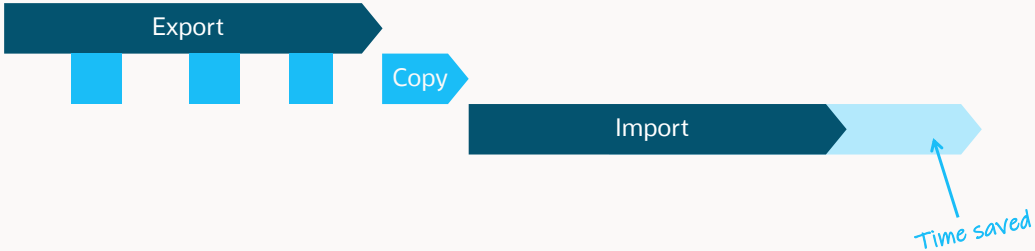

```
dbms_datapump.get_status(l_handle, l_stsmask, NULL, l_job_state, l_status);
```

```
dbms_datapump.get_status(l_handle, l_stsmask, NULL, l_job_state, l_status);

for i in l_status.job_status.files.first..l_status.job_status.files.last() loop




end loop;
```

```
dbms_datapump.get_status(l_handle, l_stsmask, NULL, l_job_state, l_status);

for i in l_status.job_status.files.first..l_status.job_status.files.last() loop

   l_dump_file := l_status.job_status.files(i);

   if (l_dump_file.file_bytes_written = l_dump_file.file_size) then
      dbms_output.put_line('DONE: ' || l_dump_file.file_name);
   end if;

end loop;
```

Copy Dump Files

- Blog post: Copy Data Pump Files Before the End of the Export

Use diagnostics views to assist in troubleshooting

- Oracle Database 23ai

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

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

Diagnostic Views

- New views available in Oracle Database 23ai:
  - `v$datapump_process_info`
  - `v$datapump_sessionwait_info`
  - `v$datapump_processwait_info`

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

Oracle LiveLabs:
Supercharge data movement with Data Pump

Key Learnings

1. Use Data Pump Bundle Patch
2. Benefit from faster index creation
3. Use NOVALIDATE constraints

YouTube | Oracle Database Upgrades and Migrations

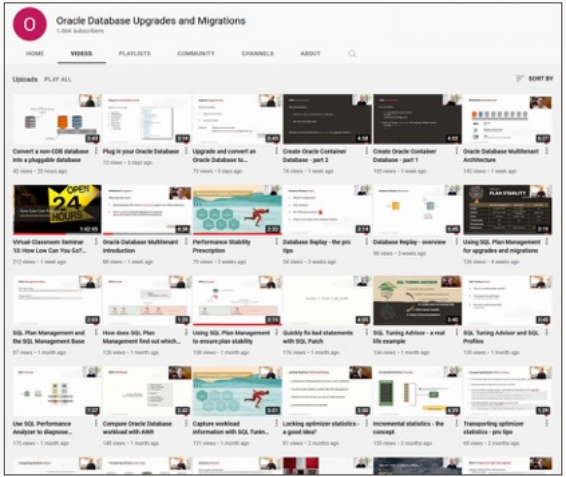

https://www.youtube.com/@upgradenow

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

SCAN ME

Find Slides and Much More on Our Blogs

MikeDietrichDE.com

Mike.Dietrich@oracle.com

dohdatabase.com

Daniel.Overby.Hansen@oracle.com

DBArj.com.br

Rodrigo.R.Jorge@oracle.com

AlexZaballa.com

Alex.Zaballa@oracle.com

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

Recorded Web Seminars

https://MikeDietrichDE.com/videos

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


ORACLE