



# WHAT'S NEW IN ORACLE DATA PUMP?

**ROY F SWONGER**

**VICE PRESIDENT, DATABASE UPGRADE, UTILITIES AND PATCHING  
ORACLE**

## MICHIGAN ORACLE USERS SUMMIT 2020

MONDAY OCTOBER 26 – THURSDAY OCTOBER 29, 2020  
VIRTUAL EVENT





**Roy F. Swonger**

Vice President  
Database Upgrade, Utilities & Patching



@RoyFSwonger



# What's New in Data Pump 20c/21c?



Help Center Search Account

Oracle Database 20c is available only for preview. It is not available for production use. Upgrades to or from Oracle Database 20c are not supported.

Home / Database / Oracle / Oracle Database / Release 20

**Oracle Database 20c**

Get Started  
What's New

Recap highlights of Releases 19c, 18c and 12.2

## What's New in Data Pump 20c?

1. Validate a dumpfile by checksum
2. Export Cloud dumpfiles into an object store
3. Include and exclude objects in the same export or import operation
4. Control index compression during import
5. Parallelize TTS metadata operations
6. Resume stopped transportable tablespace jobs
7. Export and import native JSON datatype



## Validate a Dumpfile with a Checksum



- Export `CHECKSUM_ALGORITHM` or `CHECKSUM` parameters generate SHA or CRC
- Import `VERIFY_CHECKSUM` parameter uses the checksum to validate dumpfile(s)

- Examples:

```
expdp system hr DIRECTORY=dpump_dir1 DUMPFILE=hr.dmp  
CHECKSUM_ALGORITHM=SHA384
```

```
impdp system/oracle hr DIRECTORY=dpump_dir1 DUMPFILE=hr.dmp  
VERIFY_CHECKSUM=NO
```

## Validate a Dumpfile with a Checksum



- Details:

- `CHECKSUM_ALGORITHM = [CRC32 | SHA256 | SHA384 | SHA512]`
- Alternatively, specify `CHECKSUM=YES`, defaults to `SHA256`
- Must have `COMPATIBLE = 20.0` or higher
- `impdp` verifies checksum by default, can use `VERIFY_CHECKSUM=NO`
- `VERIFY_CHECKSUM` and `VERIFY_ONLY` parameters are mutually exclusive
- Overhead depends on file size

Documentation: Utilities Guide > Export > [Export Parameters](#) / [Import Parameters](#)



## Export Dumpfiles Written to Object Store



- Export/Import from Cloud and on-premises directly to/from cloud object store
- CREDENTIAL parameter specifies the object store permissions
  - Supports Oracle Cloud, Amazon S3, Microsoft Azure
- DUMPFILE parameter specifies the object store path

- Example:

```
expdp system hr DIRECTORY=dpump_dir1:file_name CREDENTIAL=user-  
credential DUMPFILE=https://objectstorage.us-phoenix-  
1.oraclecloud.com/exp%u.dmp...
```

Documentation: Utilities Guide > Export > [Export Parameters](#)

## Include & Exclude objects in Same Job



- Include and exclude objects within the same export or import job
- `INCLUDE` parameter processed first. It includes all objects identified by the parameter
- `EXCLUDE` parameter(s) processed second. It removes any objects in the list of include objects
- The parameters can appear in any order and achieve the same result
- Example:

```
expdp system hr DIRECTORY=dpump_dir1 DUMPFILE=exp%u.dmp  
SCHEMAS=hr,oe include=table exclude=statistics
```

Documentation: Utilities Guide > [Filtering During Data Pump Operations](#)



## Include & Exclude objects in Same Job



```
expdp system hr DIRECTORY=dpump_dir1 DUMPFILE=exp%u.dmp  
SCHEMAS=hr,oe include=table exclude=statistics
```

...

Processing object type SCHEMA\_EXPORT/TABLE/TABLE\_DATA

Total estimation using BLOCKS method: 448 KB

Processing object type SCHEMA\_EXPORT/TABLE/TABLE

Processing object type SCHEMA\_EXPORT/TABLE/PRE\_TABLE\_ACTION

Processing object type SCHEMA\_EXPORT/TABLE/GRANT/OWNER\_GRANT/OBJECT\_GRANT

Processing object type SCHEMA\_EXPORT/TABLE/COMMENT

Processing object type SCHEMA\_EXPORT/TABLE/INDEX/INDEX

Processing object type SCHEMA\_EXPORT/TABLE/CONSTRAINT/CONSTRAINT

Processing object type

~~SCHEMA\_EXPORT/TABLE/INDEX/STATISTICS/INDEX\_STATISTICS~~

Processing object type SCHEMA\_EXPORT/TABLE/CONSTRAINT/REF\_CONSTRAINT

Processing object type SCHEMA\_EXPORT/TABLE/TRIGGER

~~Processing object type SCHEMA\_EXPORT/TABLE/STATISTICS/TABLE\_STATISTICS~~

Processing object type SCHEMA\_EXPORT/TABLE/POST\_TABLE\_ACTION

## Control Import Index Compression



- Compress indexes during import with `INDEX_COMPRESSION_CLAUSE` transform

- Example:

```
impdp system hr DIRECTORY=dpump_dir1 DUMPFILE=hr.dmp SCHEMAS=hr  
TRANSFORM=INDEX_COMPRESSION_CLAUSE:COMPRESS ADVANCED LOW;
```

- Details:

- `COMPRESS ADVANCED LOW` recommended for ADB imports

Documentation: Utilities Guide > Import > [Transform](#)

## Parallel Metadata Operations with TTS



- Now transportable tablespace exports CAN have a degree of parallelism > 1
- Enhances performance for full transportable, and for databases like E-Business suite that have a lot of metadata

- Example:

```
expdp system hr DIRECTORY=dpump_dir1 DUMPFILE=tts.dmp  
TRANSPORT_TABLESPACES=tbs_1 TRANSPORT_FULL_CHECK=YES  
LOGFILE=tts.log PARALLEL=4
```

Documentation: Utilities Guide > Export > [Transportable Tablespace Mode](#)

## Transportable Jobs are Restartable



- Now Transportable tablespace `TRANSPORTABLE` jobs ARE restartable
- Matches non-transportable jobs, which are always restartable
- Enhances system availability and saves time because the source database is generally `READ ONLY` during the TTS export.

- Example

```
expdp system hr ATTACH=hr.export_job  
Export> START_JOB
```

Documentation: Utilities Guide > Export > [START\\_JOB](#)

## Export & Import Native JSON Datatype



- Data Pump supports the new binary JSON datatype
- SQL\*Loader can load JSON data
  - with Direct and Conventional Path load methods
  - from data files and LOBFILES (unstructured ASCII & binary)

Documentation:

Utilities Guide > Export > [Include, Exclude](#) Utilities Guide > Import > [Include, Exclude](#)

Utilities Guide > SQL\*Loader > [Loading Objects, Collections, and LOBs](#)

## SQL\*Loader: User-Defined Credentials



- Oracle SQL\*Loader can now read data from files in an object store
- Step 1: Create the sqlldr credential in the client wallet:

```
mkstore -wrl /usr/wallet -createEntry  
oracle.sqlldr.credential.obm_scott.username <username>
```

```
mkstore -wrl /usr/wallet -createEntry  
oracle.sqlldr.credential.obm_scott.password <password>
```

## SQL\*Loader: User-Defined Credentials



- Step 2: SQL\*Loader control file - dept.ctl:

```
LOAD DATA
INFILE 'https://objectstorage.us-phoenix-
1.oraclecloud.com/exp%u.dmp/myfiles/dept1.csv'
truncate
INTO TABLE dept
FIELDS TERMINATED BY ',' OPTIONALLY ENCLOSED BY '"'
(DEPTNO, DNAME, LOC)
```

- Step 3: Run SQL\*Loader

```
sqlldr scott dept.ctl credential=obm_scott log=dept.log
```

Documentation: Utilities Guide > SQL\*Loader > [Command-Line Parameters for SQL\\*Loader](#)



## What's New in Data Pump 19c?

1. Suppress Encrypted Columns Clause
2. Set Max Data Pump Jobs & Parallelism
3. Explicitly Enable Authenticated Roles
4. Use Any Object Store Credentials
5. Wildcards in Object Store Dumpfile Name
6. Transportable Tablespaces Test Mode
7. Transportable Tablespaces Import Read-Only Tablespaces



## Suppress Encrypted Columns Clause



- Migrate to a database having TDE encrypted tablespaces
  - TDE does not support encrypted columns (e.g., Oracle Cloud)
- Example:

```
impdp system hr DIRECTORY=dpump_dir1 DUMPFILE=hr.dmp SCHEMAS=hr  
TRANSFORM=OMIT_ENCRYPTION_CLAUSE:Y
```

- Details:
  - Default: N - column encryption clauses in `CREATE TABLE` are enabled
  - Valid for `TABLE` object types

**Documentation:** Utilities Guide > Import CLI > [TRANSFORM](#)

## Set Max Data Pump Jobs & Parallelism



- DBA can more easily govern Oracle Data Pump resource utilization
- `MAX_DATAPUMP_JOBS_PER_PDB` database parameter (**Changed**)
  - Default: 100, Range: 0 to **250** or **Auto** - 50% of `SESSIONS`
  - Must be same for each RAC instance, dynamic, and modifiable per-PDB
- `MAX_DATAPUMP_PARALLEL_PER_JOB` database parameter (**New**)
  - Default: 50, Range: 1 to 250 or Auto - 25% of `SESSIONS`
  - Can be different for each RAC instance, dynamic, and modifiable per-PDB

**Documentation:** Database Reference > [Initialization Parameters](#)

## Explicitly Enable Authenticated Roles



- Specify whether to use authenticated roles for export and import
- `ENABLE_SECURE_ROLES=YES | NO`
  - available for `expdp` and `impdp` clients, and for the Oracle Data Pump PL/SQL API
- Default: `NO` – does not enable authenticated protected roles
- Beginning with release 19c you must explicitly enable authenticated roles for an export or import job

**Documentation:** Database Readme

## Use Any Object Store Credentials



Import from Oracle Cloud, AWS or Azure object store into ADB

- `CREDENTIAL` Data Pump `impdp` client CLI parameter
  - No longer constrained to using the ADB default credential
  - Object store credentials added to database `DBMS_CLOUD.CREATE_CREDENTIAL()`
  - Data Pump validates the credential exists and the user has read access

## Use Any Object Store Credentials



- Example

- BEGIN

```
DBMS_CLOUD.CREATE_CREDENTIAL(  
  credential_name => 'MY_CRED_NAME',  
  username => 'adwc_user@oracle.com',  
  password => 'Auth token' );
```

- END; (or password for OCI/C)

- `impdp admin/password@ADWC1_high directory=data_pump_dir  
 credential=MY_cred_name ...`

Documentation: 19c & 18.3 Utilities Guide > Import > [Credential](#)

## Wildcards in Dumpfile Name




- Oracle Data Pump allows wildcards for dumpfile in object store
- Dumpfile Specification
  - A wildcard character can be specified in the file-name of a URL
  - simplifies importing multiple dump files, can't be used in bucket-name

### Example:

- ```
impdp admin/password@ATPC1_high
      directory=data_pump_dir credential=my_cred_name
      dumpfile= https://objectstorage.us-phoenix-
      1.oraclecloud.com/atpc/atpc_user/exp%u.dmp...
```

- **Documentation:** Utilities Guide > Database Readme






## Transportable Tablespaces Test Mode

- Test a TTS or Full Transportable `expdp` without setting the source `READ ONLY`
- `TTS_CLOSURE_CHECK: ON | OFF | FULL | TEST_MODE`
  - Preview an export for time to complete
  - Check for unforeseen closure issues
- Resulting dumpfile cannot be imported
  - If attempted, error will be issued


Documentation: Utilities Guide > Export> TRANSPORTABLE> [TTS\\_CLOSURE\\_CHECK](#)

## TTS Import Read-Only Tablespaces



- Allows Read-Only Tablespaces during Transportable Tablespaces import
- `TRANSPORTABLE=NEVER | ALWAYS | KEEP_READ_ONLY | NO_BITMAP_REBUILD`
  - Restores pre-12.2 ability to have tablespace files mounted on two databases at once
  - Note: Prevents fix-up of timezone data and rebuilding of bitmaps

## TTS Import Read-Only Tablespaces



Example:

```
impdp system DIRECTORY=dpump_dir DUMPFILE=dumpfile_name  
TRANSPORT_DATAFILES=datafile_name TRANSPORTABLE=KEEP_READ_ONLY
```


Documentation: Utilities Guide > Import > [TRANSPORTABLE](#)

## What's New in Data Pump 18c?

1. Data Pump Import `DATA_OPTIONS` parameter `CONTINUE_LOAD_ON_FORMAT_ERROR` option
2. Data Pump warns on import that encrypted fixed user database link passwords must be reset
3. Oracle Data Pump full or partial export and import operations can include a unified audit trail




## Continue Load on Format Error



- `CONTINUE_LOAD_ON_FORMAT_ERROR` is an import `DATA_OPTIONS` parameter
  - Condition: if a stream format error is encountered while loading table data
  - Action: Data Pump import skips forward to the start of next granule
- This parameter enables partial import of a corrupt dump file

Documentation: Utilities Guide > Import CLI > [Data\\_Options](#)




## TDE Encrypted Password Warning

- Oracle Data Pump import now (like export) generates an ORA-39395:  
Warning: object <database link name> requires password reset after import
- If fixed-user database passwords are encrypted, database link passwords are not exported  
- expdp stores a known invalid password
- Reset the database link password with the command:

```
ALTER DATABASE LINK database_link_name CONNECT TO schema_name  
IDENTIFIED BY password;
```

**Documentation:** Advanced Security Guide > [Using Oracle Data Pump With Encrypted Data Dictionary Data](#)



## Can Include a Unified Audit Trail

- A unified audit trail can be included for full or partial export and import
- Oracle Data Pump uses unified auditing to monitor and record specific user database actions and centralize all audit records in one place
- How to use a unified audit trail:
  - Create the policy: `SQL CREATE AUDIT POLICY` (or alter an existing policy)
  - Enable and disable the policy with SQL statements: `AUDIT` and `NOAUDIT`

**Documentation:** Utilities Guide > Overview > [Auditing Data Pump Jobs](#)



## What's New in Data Pump 12.2?

1. Parallel Export/Import of Metadata
2. Substitution Variables & Wildcards
3. REMAP\_DIRECTORY
4. Long Identifier support
5. TRUST\_EXISTING\_TABLE\_PARTITIONS
6. Validation & Verification options
7. ...and many more 12.2 Features



## Parallel Metadata Export: How it Worked Pre-12.2

- Start with ESTIMATE phase
  - Gatherable data objects
  - Other workers remain idle until data objects are gathered
- Metadata exported serially
- Data exported in parallel

## Parallel Metadata Export: New Feature in 12.2

- Start with Analysis step
  - Metadata objects passed immediately to workers as they are found
  - E.g. Worker 1 finds a set of TABLE definitions, they are handed off to worker 2
- ESTIMATE phase still happens, but metadata no longer held up by estimate

## Parallel Metadata Export: New Feature in 12.2

- Details:
  - Works for dumpfile jobs if source database is 12.2
  - Transportable and Full Transportable jobs are not (yet) parallel for metadata
  - ESTIMATE phase now uses STATISTICS only
  - Restart works as always

# Parallel Metadata Export: Logfile

## 12.1.0.2

```
18-SEP-16 10:53:16.733: Starting "SYSTEM"."MD_EXP_16_12102": system/***** parfile=md_exp_16_12102.par
18-SEP-16 10:53:17.600: Startup took 2 seconds
18-SEP-16 10:53:17.623: Estimate in progress using BLOCKS method...
```

```
18-SEP-16 10:54:37.945: Processing object type DATABASE_EXPORT/NORMAL_OPTIONS/VIEWS_AS_TABLES/TABLE_DATA
18-SEP-16 10:55:30.500: Estimated 10 TABLE_DATA objects in 0 seconds
18-SEP-16 10:55:30.502: Processing object type DATABASE_EXPORT/SCHEMA/TABLE/TABLE_DATA
18-SEP-16 10:55:56.008: Estimated 36026 TABLE_DATA objects in 79 seconds
18-SEP-16 10:55:56.380: Startup took 162 seconds
18-SEP-16 10:55:56.556: Startup took 162 seconds
18-SEP-16 10:55:56.757: Startup took 162 seconds
18-SEP-16 10:55:56.940: Startup took 162 seconds
```

```
18-SEP-16 10:56:01.566: Total estimation using BLOCKS method: 74.77 GB
18-SEP-16 10:56:02.015: Processing object type DATABASE_EXPORT/PRE_SYSTEM_IMPCALLOUT/MARKER
18-SEP-16 10:56:02.022: Completed 1 MARKER objects in 1 seconds
18-SEP-16 10:56:02.023: Processing object type DATABASE_EXPORT/PRE_INSTANCE_IMPCALLOUT/MARKER
18-SEP-16 10:56:03.534: Completed 1 MARKER objects in 0 seconds
18-SEP-16 10:56:03.535: Processing object type DATABASE_EXPORT/TABLESPACE
```

## 12.2.0.1

```
18-SEP-16 15:24:32.166: Starting "SYSTEM"."MD_EXP_16_12201": system/***** parfile=md_exp_16_12201.par
18-SEP-16 15:24:32.742: W-1 Startup took 2 seconds
18-SEP-16 15:24:35.601: W-3 Startup took 3 seconds
18-SEP-16 15:24:36.148: W-2 Startup took 3 seconds
18-SEP-16 15:24:36.205: W-4 Startup took 4 seconds
18-SEP-16 15:24:36.393: W-5 Startup took 4 seconds
18-SEP-16 15:24:36.490: W-6 Startup took 4 seconds
18-SEP-16 15:24:36.491: W-7 Startup took 4 seconds
18-SEP-16 15:24:36.650: W-8 Startup took 4 seconds
18-SEP-16 15:24:36.714: W-9 Startup took 4 seconds
18-SEP-16 15:24:36.715: W-10 Startup took 4 seconds
18-SEP-16 15:24:36.716: W-11 Startup took 4 seconds
18-SEP-16 15:24:37.153: W-12 Startup took 4 seconds
18-SEP-16 15:24:37.187: W-13 Startup took 4 seconds
18-SEP-16 15:24:37.220: W-14 Startup took 4 seconds
18-SEP-16 15:24:37.253: W-15 Startup took 4 seconds
18-SEP-16 15:24:37.286: W-16 Startup took 4 seconds
18-SEP-16 15:24:37.323: W-3 Processing object type DATABASE_EXPORT/PRE_SYSTEM_IMPCALLOUT/MARKER
18-SEP-16 15:24:37.324: W-3 Completed 1 MARKER objects in 0 seconds
18-SEP-16 15:24:37.350: W-2 Processing object type DATABASE_EXPORT/PRE_INSTANCE_IMPCALLOUT/MARKER
18-SEP-16 15:24:37.359: W-2 Completed 1 MARKER objects in 0 seconds
18-SEP-16 15:24:37.436: W-7 Processing object type DATABASE_EXPORT/PROFILE
18-SEP-16 15:24:37.509: W-8 Processing object type DATABASE_EXPORT/SYS USER/USER
18-SEP-16 15:24:37.580: W-4 Processing object type DATABASE_EXPORT/ROLE
18-SEP-16 15:24:37.584: W-7 Completed 3 PROFILE objects in 1 seconds
18-SEP-16 15:24:37.585: W-8 Completed 1 USER objects in 0 seconds
18-SEP-16 15:24:37.664: W-4 Completed 64 ROLE objects in 0 seconds
```

## Parallel Metadata Import

### Pre-12.2:

- One worker per partition/subpartition
- PQ used if partitions are large enough
- Package bodies loaded in parallel

### With patch for bug 22273229

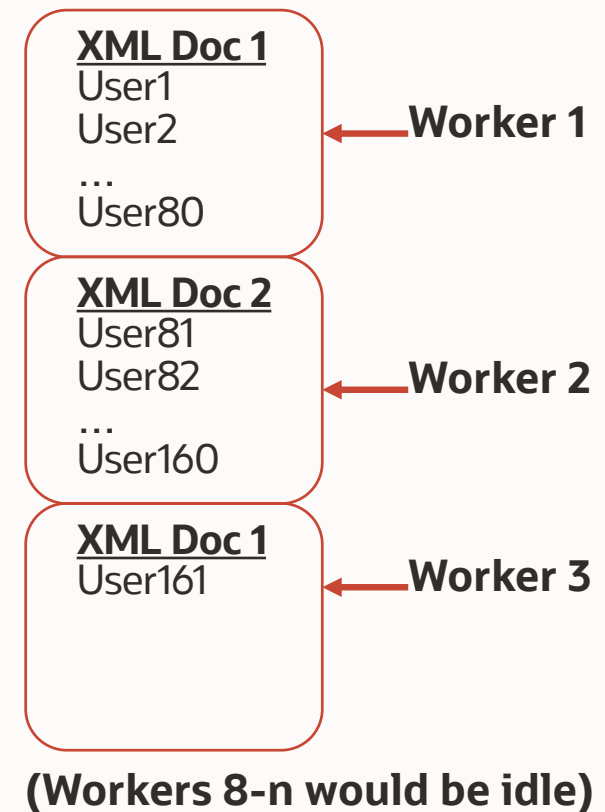
- Indexes built in parallel
- Constraints created in parallel
- Backport to 12.1.0.2, 11.2.0.4

### Starting with 12.2

- Added parallel import of **most** other metadata objects
- Some exceptions
  - Types (due to inheritance)
  - Schemas
  - Procedural actions

## Parallel Metadata Import: Internals

- Metadata is exported in XML documents
  - Each XML document in dumpfile contains  $n$  objects of a given type
- XML documents are allocated to workers 1 document at a time
- Example: 161 users to import
  - Users are exported with up to 80 users per XML document
  - What happens with `PARALLEL=8`?
- Notes:
  - Works for conventional (dumpfile) jobs
  - Not (until 20c) for TTS jobs or (not yet) network mode
  - Restart works same as always
  - Status command will show multiple workers on metadata





## Parallel Metadata Import: Logfile

- Comparison with `PARALLEL=8` for 27586 object grants and `METRICS=Y`

- 12.1.0.2

```
15-SEP-16 13:56:16.317: Processing object type DATABASE_EXPORT/SCHEMA/SEQUENCE/GRANT/OWNER_GRANT/OBJECT_GRANT
15-SEP-16 13:57:06.374:      Completed 27586 OBJECT_GRANT objects in 50 seconds
```

- 12.2.0.1

```
15-SEP-16 11:59:35.190: W-7 Processing object type DATABASE_EXPORT/SCHEMA/SEQUENCE/GRANT/OWNER_GRANT/OBJECT_GRANT
15-SEP-16 11:59:49.304: W-4      Completed 27586 OBJECT_GRANT objects in 10 seconds
15-SEP-16 11:59:49.304: W-4      Completed by worker 1 3426 OBJECT_GRANT objects in 10 seconds
15-SEP-16 11:59:49.304: W-4      Completed by worker 2 3440 OBJECT_GRANT objects in 10 seconds
15-SEP-16 11:59:49.304: W-4      Completed by worker 3 3440 OBJECT_GRANT objects in 10 seconds
15-SEP-16 11:59:49.304: W-4      Completed by worker 4 3440 OBJECT_GRANT objects in 9 seconds
15-SEP-16 11:59:49.304: W-4      Completed by worker 5 3440 OBJECT_GRANT objects in 10 seconds
15-SEP-16 11:59:49.304: W-4      Completed by worker 6 3520 OBJECT_GRANT objects in 10 seconds
15-SEP-16 11:59:49.304: W-4      Completed by worker 7 3440 OBJECT_GRANT objects in 10 seconds
15-SEP-16 11:59:49.304: W-4      Completed by worker 8 3440 OBJECT_GRANT objects in 10 seconds
```

## Performance: Parallel Metadata Import

- Examples from E-Business Suite test database

| Object Type             | Count  | 11.2.0.4<br>PARALLEL=32 | 12.1.0.2<br>PARALLEL=8 | 12.1.0.2<br>PARALLEL=32<br>With Patch | 12.2.0.1<br>PARALLEL=8 | 12.2.0.1<br>PARALLEL=32 | Comments                    |
|-------------------------|--------|-------------------------|------------------------|---------------------------------------|------------------------|-------------------------|-----------------------------|
| OBJECT_GRANT<br>(owner) | 27586  | 49                      | 50                     | 51                                    | 10                     | 22                      | Hard connect for each grant |
| SYNONYM                 | 43254  | 105                     | 109                    | 111                                   | 25                     | 44                      |                             |
| TYPE                    | 4364   | 108                     | 114                    | 119                                   | 111                    | 110                     | Handled by single worker    |
| PROCACT_SCHEMA          | 606    | 198                     | 216                    | 214                                   | 152                    | 175                     | Handled by single worker    |
| TABLE                   | 33164  | 923                     | 1160                   | 1298                                  | 368                    | 248                     |                             |
| OBJECT_GRANT<br>(table) | 358649 | 541                     | 543                    | 578                                   | 142                    | 157                     | Hard connect for each grant |
| INDEX                   | 53190  | 6721                    | 5770                   | 360                                   | 418                    | 272                     |                             |
| PACKAGE                 | 53217  | 424                     | 476                    | 474                                   | 114                    | 54                      |                             |
| VIEW                    | 34690  | 538                     | 583                    | 593                                   | 151                    | 184                     |                             |
| PACKAGE_BODY            | 52092  | 1363                    | 1974                   | 1186                                  | 1981                   | 959                     | Always parallel since 11.2  |

\*import time in seconds

## Substitution Variables for Dumpfile Name

- Substitution variables for dumpfile name:
  - Pre-12.2: %U generates a fixed-width 2-digit number
    - e.g. dumpfile=exp%U.dmp
  - New option for 12.2 expdp or impdp:
    - %l or %L: Incrementing number from 01 up to 2,147,483,646
  - New options in 12.2 expdp only:
    - %d or %D: Day of Month in DD format
    - %m or %M: Number of Month in MM format
    - %y or %Y: Year in YYYY format
    - %t or %T: Full date in YYYYMMDD format

```
$ expdp system/oracle directory=mydir \
  filesize=50K dumpfile=exp%T_%L.dmp full=y
...
...
...
. . exported "WMSYS"."WM$METADATA_MAP"
0 KB          0 rows
Master table "SYSTEM"."SYS_EXPORT_FULL_01" successfully
loaded/unloaded
*****
*****
Dump file set for SYSTEM.SYS_EXPORT_FULL_01 is:
/home/oracle/exp20160917_01.dmp
/home/oracle/exp20160917_02.dmp
/home/oracle/exp20160917_03.dmp
...
...
/home/oracle/exp20160917_67.dmp
/home/oracle/exp20160917_68.dmp
/home/oracle/exp20160917_69.dmp
/home/oracle/exp20160917_70.dmp
/home/oracle/exp20160917_71.dmp
/home/oracle/exp20160917_72.dmp
Job "SYSTEM"."SYS_EXPORT_FULL_01" successfully
completed at Sat Sep 17 23:47:31 2016 elapsed 0
00:03:00
```

## Wildcards for TRANSPORT\_DATAFILES

- Use a wildcard in 12.2 instead of listing every file

- Pre-12.2:

```
TRANSPORT_DATAFILES=users01.dbf
TRANSPORT_DATAFILES=users02.dbf
...
TRANSPORT_DATAFILES=data1.dbf
TRANSPORT_DATAFILES=data2.dbf
```

- 12.2: wildcards

- \* (asterisk) matches multiple characters
- ? (question mark) matches a single character

```
TRANSPORT_DATAFILES=users*.dbf
TRANSPORT_DATAFILES=data?.dbf
...
```

```
$ impdp system/oracle@pdb2 network_link=sourcedb \
version=12 full=y transportable=always metrics=y \
exclude=statistics \
directory=mydir \
logfile=pdb2.log \
transport_datafiles='/u02/oradata/CDB2/pdb2/user*.dbf'
```

## REMAP\_DIRECTORY

- Applies to DDL where directory specs are used
  - E.g. CREATE TABLESPACE
- Change directory spec **without** changing filenames
- Useful when moving between OS platforms
  - Example: importing dumpfile created on OpenVMS into database on Linux

```
REMAP_DIRECTORY="'DB1$: [HRDATA.PAYROLL] ':' /db1/hrdata/payroll/'"
```

## TRUST\_EXISTING\_TABLE\_PARTITIONS

- Pre-12.2
  - Importing into existing table was done serially
  - Data Pump couldn't be sure that partitioning in DB matched partitioning in dumpfile
- New 12.2 Parameter: `DATA_OPTIONS=TRUST_EXISTING_TABLE_PARTITIONS`
  - Big performance boost
  - If partitions don't match...error:

ORA-31693: Table data object "SH"."SALES\_BIG\_PT":"SALES\_2000" failed to load/unload and is being skipped due to error:

ORA-29913: error in executing ODCIEXTTABLEFETCH callout

ORA-14401: inserted partition key is outside specified partition

## Data Validation & Verification

### Extra Validation for Things That Should Never Happen

- DATA\_OPTIONS=VALIDATE\_TABLE\_DATA

- Import only
- Validates date & number formats of table data
- Default is no validation

```
ORA-02374: conversion error loading table "DPV"."TEST18"  
ORA-12899: value too large for column C1 (actual: 500,  
maximum: 498)  
ORA-02372: data for row: C8 : '
```

- DATA\_OPTIONS=VERIFY\_STREAM\_FORMAT

- Export only
- Default is no verification

```
Starting "SCOTT"."SYS_EXPORT_TABLE_01": scott/***** tables=t  
directory=dmpdir dumpfile=t.dmp reuse_dumpfiles=true  
Processing object type TABLE_EXPORT/TABLE/TABLE_DATA  
Processing object type TABLE_EXPORT/TABLE/STATISTICS/TABLE_STATISTICS  
Processing object type TABLE_EXPORT/TABLE/STATISTICS/MARKER  
Processing object type TABLE_EXPORT/TABLE/TABLE  
. . exported "SCOTT"."T" 5.570  
KB 1 rows  
ORA-31694: master table "SCOTT"."SYS_EXPORT_TABLE_01" failed to  
load/unload  
ORA-02354: error in exporting/importing data  
ORA-26009: stream verification error: [1], [0], [0], [0]
```

## Other 12.2 Features (1)

- Use direct path load in network mode!
  - Specify `ACCESS_METHOD=DIRECT_PATH` with `NETWORK_LINK=<dblink>`
  - Allows network import of `LONG` and `LONG RAW`
- Data Pump available in Instant Client
  - **Tools** package for Instant Client
  - Includes `SQL*Loader`, `expdp`, `impdp`, `exp`, `imp`
- Views that describe available transforms
  - `DBMS_METADATA_TRANSFORMS`
  - `DBMS_METADATA_TRANSFORM_PARAMS`
  - `DBMS_METADATA_PARSE_ITEMS`



## Other 12.2 Features (2)

- New interactive commands
  - TRACE parameter can be set for a running job
    - No need to stop/restart job for tracing to take effect
  - STOP\_WORKER command
    - Kill an individual worker you believe to be hung or stuck
  - Both documented in MOS notes
- Enhanced log files
  - When METRICS=Y
    - Show worker ID for each item processed
    - Show access method for each table
  - Include contents of parfile in logfile

```
18-SEP-16 15:24:30.950: ;;;  
*****  
18-SEP-16 15:24:30.951: ;;; Parfile values:  
18-SEP-16 15:24:30.953: ;;; parfile:  
job_name=md_exp_16_12201  
18-SEP-16 15:24:30.955: ;;; parfile: reuse_dumpfiles=Y  
18-SEP-16 15:24:30.957: ;;; parfile: logtime=all  
18-SEP-16 15:24:30.958: ;;; parfile: metrics=Y  
18-SEP-16 15:24:30.960: ;;; parfile: parallel=16  
18-SEP-16 15:24:30.962: ;;; parfile: full=Y  
18-SEP-16 15:24:30.963: ;;; parfile:  
logfile=md_exp_16_12201.log  
18-SEP-16 15:24:30.965: ;;; parfile:  
dumpfile=mdl6_12201_%U.dmp  
18-SEP-16 15:24:30.966: ;;; parfile: directory=EBSIMP  
18-SEP-16 15:24:30.968: ;;;  
*****
```

## Recap of 12.1 Features

- `VIEWS_AS_TABLES` parameter
  - Lets you export the contents of a view as a table
- `TRANSFORM` parameter options
  - `TRANSFORM=DISABLE_ARCHIVE_LOGGING:Y`
    - Will disable archive logging during import for tables and/or indexes
  - `TRANSFORM=LOB_STORAGE:SECUREFILE`
  - `TRANSFORM=STORAGE:N`
  - `TRANSFORM=TABLE_COMPRESSION:<compression_clause>`
- `LOGTIME=[ NONE | STATUS | LOGFILE | ALL ]` parameter
  - Will write timestamps on status and/or logfile messages

## Data Pump Best Practices - General

- For *full exports*:
  - Role `EXP_FULL_DATABASE` is required
- For *export consistency* use:
  - `FLASHBACK_SCN=<scn>`
  - `FLASHBACK_TIME=SYSTIMESTAMP`  
alternative:
    - `CONSISTENT=Y` [since Oracle 11.2 – Legacy Interface]
      - This will increase UNDO requirements for the duration of the export
- Always set parameters:
  - `EXCLUDE=STATISTICS`
  - `METRICS=YES`

## Data Pump Best Practices - Performance



- Use `PARALLEL=n`
  - Typically  $n = 2x$  <number of CPU cores>
- **New feature in 12.2:** Parallel Export/Import of Metadata
  - 11.2.0.4 and 12.1.0.2: apply patch for bug [22273229](#) to enable parallel import of constraints and indexes
- `EXCLUDE=STATISTICS` on export
- New feature in 12c: `TRANSFORM=DISABLE_ARCHIVE_LOGGING:Y`
  - 12.1.0.2: Apply patch for bug [20778442](#)

## Data Pump Best Practices - Network Mode

- Direct import via database link
  - Parameter: `NETWORK_LINK`
    - Run `impdp` on the target system - `no expdp` necessary
    - No dump file written, no disk I/O, no file transfer needed
  - **New Feature 12.2:** `ACCESS_METHOD=DIRECT_PATH` (including LONG data)
- Performance: Depends on network bandwidth and target's CPUs
- Restrictions of database links apply

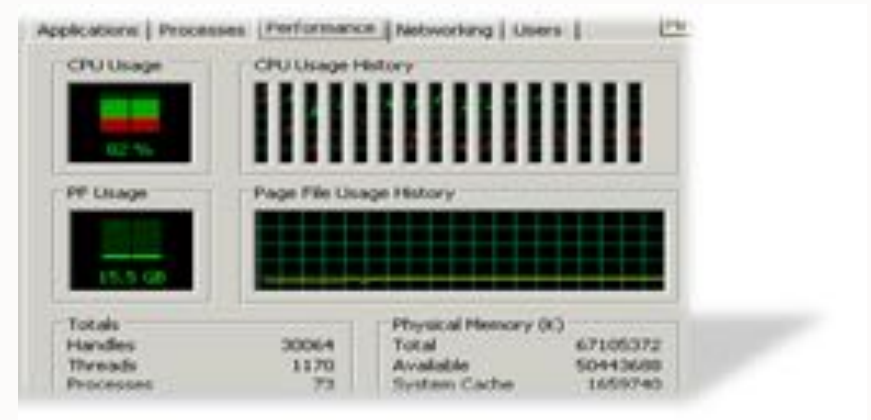


# Data Pump Best Practices - Network Mode

## Real World Case:

### Kaiser Permanente, Medicare (USA)

- `impdp` on `NETWORK_LINK` with 8 vs 16 CPU cores
  - 10Gbit connection leveraged up to 8 Gbit
  - 1 TB table copied in ~15 min ⇒ 4 TB/hour
- Network bandwidth and CPU bound



```

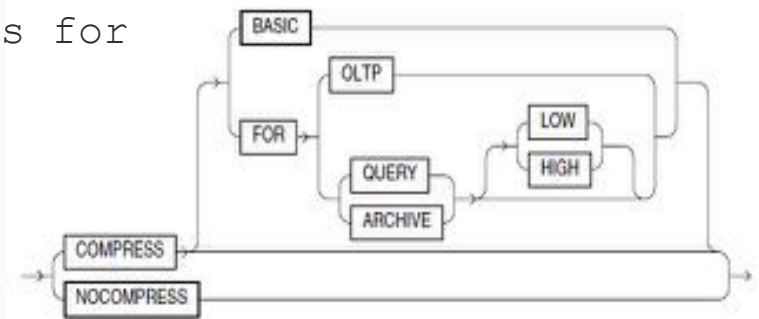
cciss/c6d0  F_NK" VARCHAR2(9 BYTE), "GL_UNIT" VARCHAR2(10 BYTE), "ACCOUNT" VARCHAR2(20 BYTE)
cciss/c6d1  ,"DEPT" VARCHAR2(10 BYTE), "DESCR1" VARCHAR2(255 BYTE), "GL_LOC" VARCHAR2(10 BY
cciss/c6d2  TE), "DESCR2" VARCHAR2(255 BYTE), "PERIOD" NUMBER, "YEAR_NR" NUMBER, "ITEM" VARC
cciss/c6d3  HAR2(12 BYTE), "LONG_DESCR" VARCHA
cciss/c7d0
cciss/c7d1  0.00      11.60     0.00      4.64      0.00      0.38
cciss/c7d2  0.00      47.60     0.00      18.07     0.00      0.38
cciss/c7d3  0.00      24.80     0.00      9.25      0.00      0.37
cciss/c8d0  0.00      29.20     0.00      10.41     0.00      0.36
cciss/c8d1  0.00      8.60      0.00      3.54      0.00      0.41
cciss/c8d2  0.00      21.60     0.00      9.62      0.00      0.45
cciss/c8d3  0.00      30.60     0.00      12.15     0.00      0.40
-----
SUM        0.00      705.80   0.00      278.96   0.00      0.40
AVG        0.00      22.06    0.00      8.72     0.00      0.40
    
```

```


23  74.5  2.2  0.0  23.4  0.000000
24  63.7  1.1  0.0  15.2  0.000000
Avg 51.6  1.1  1.5  45.8  0.000000
-----
Network I/O
I/F Name Recv=KB/s Trans=KB/s packin
lo        0.0    0.0    1.0
eth2      1.6    2.3   22.0
eth3      0.0    0.0    0.0
eth0      0.0    0.0    0.0
eth1 187626.9  739.2  23457.1
sit0      0.0    0.0    0.0
    
```

## Data Pump New in Oracle 12c

- TRANSFORM option to enable Advanced/HCC Compression
  - Example:
    - TRANSFORM=TABLE\_COMPRESSION:"compress for query high"
  - **But:** Granularity only on the entire import
    - Workarounds:
      - Precreate objects
        - Downside: Will slow down import!!!  
*or:*
      - Precreate the tablespace with COMPRESS option
        - create tablespace ARCHHIGH datafile 'archhigh.ora' size 100G default compress for archive high;
      - Then run Data Pump with TRANSFORM=TABLE\_COMPRESSION:N
        - This will drop all embedded compression attributes associated with the tables
        - Now tablespace compression option will be used for all newly created tables



## Exporting Views as Tables



```
expdp system/mgr  
views_as_tables=scott.view1 ...
```

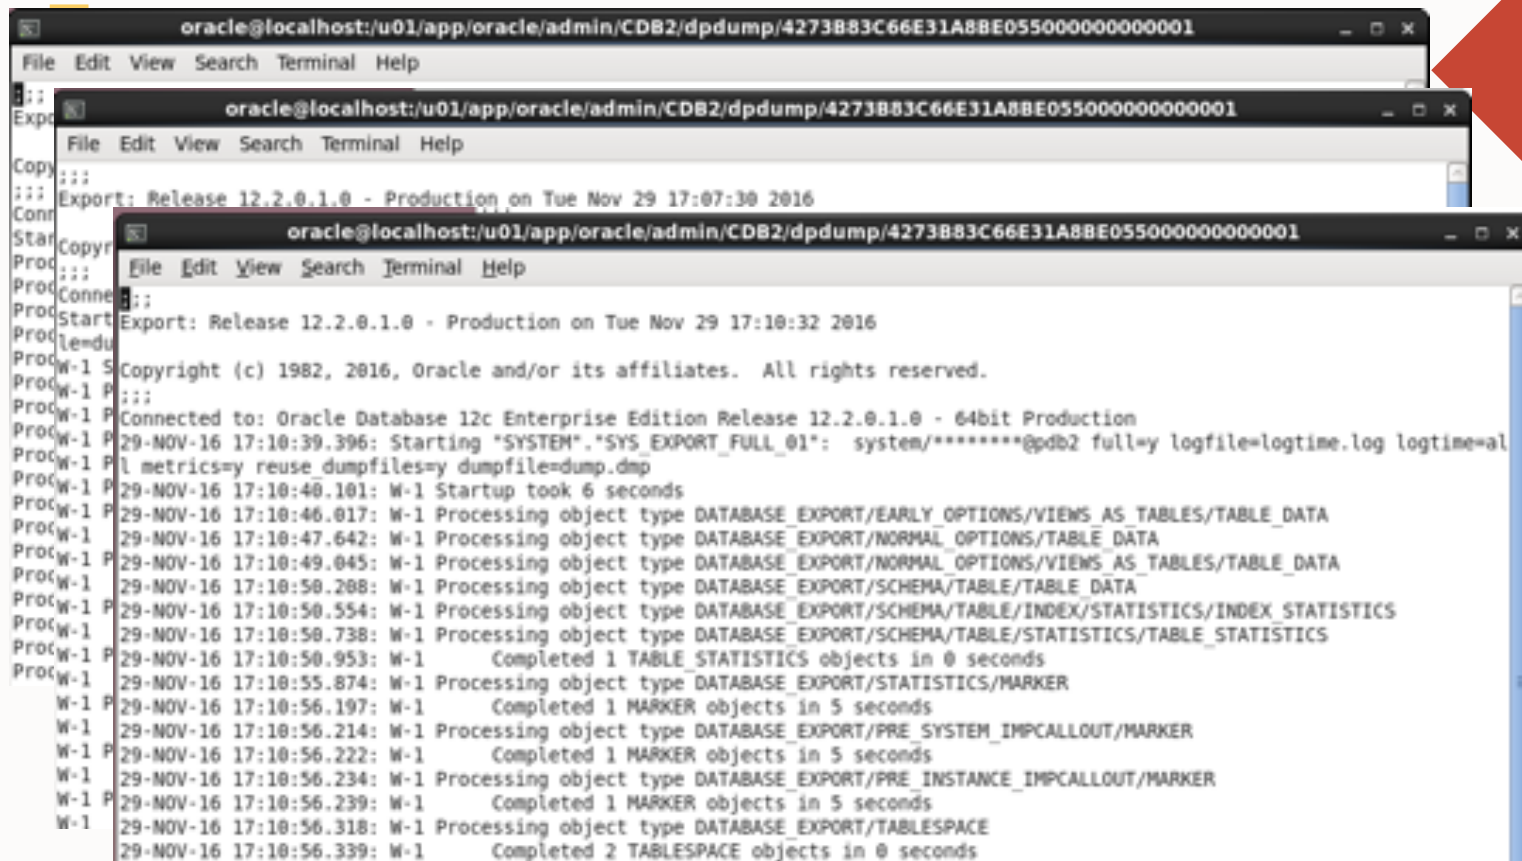
```
impdp system/mgr  
remap_table=view1:scott.table1 ...
```





## Improved Log Files:

METRICS=Y and LOGTIME=ALL



The screenshot displays three overlapping terminal windows showing Oracle Database export logs. The top window shows a basic log file. The middle window shows logs with the `METRICS=Y` option, and the bottom window shows logs with the `LOGTIME=ALL` option. Red arrows point from the text labels to the corresponding log entries in the terminal windows.

```
oracle@localhost:u01/app/oracle/admin/CDB2/dpdump/4273B83C66E31A8BE05500000000001
File Edit View Search Terminal Help
Export: Release 12.2.0.1.0 - Production on Tue Nov 29 17:07:30 2016

oracle@localhost:u01/app/oracle/admin/CDB2/dpdump/4273B83C66E31A8BE05500000000001
File Edit View Search Terminal Help
Export: Release 12.2.0.1.0 - Production on Tue Nov 29 17:10:32 2016
Copyright (c) 1982, 2016, Oracle and/or its affiliates. All rights reserved.
Connected to: Oracle Database 12c Enterprise Edition Release 12.2.0.1.0 - 64bit Production
29-NOV-16 17:10:39.396: Starting "SYSTEM"."SYS_EXPORT_FULL_01": system/*****@pdb2 full=y logfile=logtime.log logtime=all metrics=y reuse dumpfiles=y dumpfile=dump.dmp
29-NOV-16 17:10:40.101: W-1 Startup took 6 seconds
29-NOV-16 17:10:46.017: W-1 Processing object type DATABASE_EXPORT/EARLY_OPTIONS/VIEWS_AS_TABLES/TABLE_DATA
29-NOV-16 17:10:47.642: W-1 Processing object type DATABASE_EXPORT/NORMAL_OPTIONS/TABLE_DATA
29-NOV-16 17:10:49.045: W-1 Processing object type DATABASE_EXPORT/NORMAL_OPTIONS/VIEWS_AS_TABLES/TABLE_DATA
29-NOV-16 17:10:50.208: W-1 Processing object type DATABASE_EXPORT/SCHEMA/TABLE/TABLE_DATA
29-NOV-16 17:10:50.554: W-1 Processing object type DATABASE_EXPORT/SCHEMA/TABLE/INDEX/STATISTICS/INDEX_STATISTICS
29-NOV-16 17:10:50.738: W-1 Processing object type DATABASE_EXPORT/SCHEMA/TABLE/STATISTICS/TABLE_STATISTICS
29-NOV-16 17:10:50.953: W-1 Completed 1 TABLE_STATISTICS objects in 0 seconds
29-NOV-16 17:10:55.874: W-1 Processing object type DATABASE_EXPORT/STATISTICS/MARKER
29-NOV-16 17:10:56.197: W-1 Completed 1 MARKER objects in 5 seconds
29-NOV-16 17:10:56.214: W-1 Processing object type DATABASE_EXPORT/PRE_SYSTEM_IMPALLOCOUT/MARKER
29-NOV-16 17:10:56.222: W-1 Completed 1 MARKER objects in 5 seconds
29-NOV-16 17:10:56.234: W-1 Processing object type DATABASE_EXPORT/PRE_INSTANCE_IMPALLOCOUT/MARKER
29-NOV-16 17:10:56.239: W-1 Completed 1 MARKER objects in 5 seconds
29-NOV-16 17:10:56.318: W-1 Processing object type DATABASE_EXPORT/TABLESPACE
29-NOV-16 17:10:56.339: W-1 Completed 2 TABLESPACE objects in 0 seconds
```

Basic Logfile

METRICS=Y  
(Added 11.2)

LOGTIME=ALL  
(Added 12.1)

## Enhanced Compression Algorithm

- `COMPRESSION_ALGORITHM`
  - Defines the compression algorithm when compressing dump files
    - `BASIC` The same algorithm used in previous versions. Good compression, without severely impacting on performance
    - `LOW`: For use when reduced CPU utilization is a priority over compression ratio
    - `MEDIUM`: Recommended option. Similar characteristics to `BASIC`, but uses a different algorithm
    - `HIGH`: Maximum available compression, but more CPU intensive
  - Performance:
    - Compression ratio
    - CPU usage

```
$ expdp scott/tiger tables=emp directory=mydir  
dumpfile=emp.dmp logfile=expdp_emp.log compression=all  
compression_algorithm=medium
```

- Requires Advanced Compression Option license for export, not for import

## Enhanced Compression Algorithm

- Customer evaluation

- BASIC  
at 3.5 TB/hour

| Disk-Group-I/O   |       |         |                 |           |         |             |  |
|------------------|-------|---------|-----------------|-----------|---------|-------------|--|
| Name             | Disks | AvgBusy | Read Write-KB/s | TotalMB/s | xfers/s | BlockSizeKB |  |
| slot02           | 6     | 9.3%    | 123120.4 0.0    | 120.2     | 241.1   | 510.7       |  |
| slot03           | 6     | 6.7%    | 103354.8 0.0    | 100.9     | 202.2   | 511.1       |  |
| slot05           | 6     | 9.0%    | 130420.9 7.0    | 127.4     | 262.0   | 519.8       |  |
| slot06           | 6     | 10.5%   | 158841.9 175.3  | 155.3     | 329.3   | 511.0       |  |
| slot08           | 6     | 8.4%    | 130835.3 0.0    | 127.8     | 256.0   | 511.3       |  |
| slot09           | 6     | 10.1%   | 136525.9 0.0    | 133.3     | 267.0   | 510.6       |  |
| slot10           | 6     | 6.6%    | 140383.4 0.0    | 137.1     | 275.0   | 510.3       |  |
| slot11           | 6     | 6.8%    | 112600.0 2.0    | 110.0     | 220.7   | 510.3       |  |
| Groups= 8 TOTALS | 48    | 1.4%    | 1036082.5 184.3 | 1012.0    | 2053.3  |             |  |

- MEDIUM  
at 7.0 TB/hour

| Disk-Group-I/O   |       |         |                 |           |          |             |  |
|------------------|-------|---------|-----------------|-----------|----------|-------------|--|
| Name             | Disks | AvgBusy | Read Write-KB/s | TotalMB/s | xfers/s  | BlockSizeKB |  |
| slot02           | 6     | 14.5%   | 255770.4 0.0    | 249.8     | 500.9    | 510.7       |  |
| slot03           | 6     | 16.0%   | 273037.4 11.5   | 266.6     | 535.1    | 510.3       |  |
| slot05           | 6     | 15.4%   | 264851.1 17.5   | 258.7     | 519.0    | 510.3       |  |
| slot06           | 6     | 13.2%   | 222160.7 425.5  | 217.4     | 502.4    | 510.5       |  |
| slot08           | 6     | 15.0%   | 267156.6 1.5    | 260.9     | 523.3    | 510.6       |  |
| slot09           | 6     | 14.8%   | 263140.4 6.5    | 257.0     | 515.3    | 510.5       |  |
| slot10           | 6     | 14.6%   | 259603.7 2.5    | 253.5     | 508.5    | 510.4       |  |
| slot11           | 6     | 14.9%   | 258113.0 5.4    | 252.1     | 505.8    | 510.4       |  |
| Groups= 8 TOTALS | 48    | 2.5%    | 2063833.5 470.4 | 2015.9    | 4110.285 |             |  |

2x

Slides and more | <https://MikeDietrichDE.com>

The image shows a screenshot of a blog post on the website 'Mike Dietrich's Blog About Oracle Database Upgrades... Mostly'. The main heading is 'Upgrade your Database - NOW!'. Below the heading is a navigation menu with links for 'Blog', 'Slides', 'Hands-On Lab', 'Events', 'Papers / Docs', 'Videos', 'Scripts', 'Links', and 'Oracle Documentatio'. A red arrow points from the 'Slides' link to a QR code on the right side of the page. The QR code has the text 'Upgrade your Database - NOW!' and the URL 'https://mikedietrichde.com' overlaid on it. The main content of the blog post is titled 'AutoUpgrade and Data Guard, RAC, Restart and non-CDB to PDB'. It is dated 'Posted on July 19, 2019 by Mike Dietrich' and has a tag 'AutoUpgrade'. The text of the post reads: 'I need to bring this blog post forward about AutoUpgrade and Data Guard, RAC, Restart and non-CDB to PDB. Initially I planned to write this a bit later. But some of you had questions or were wondering why AutoUpgrade hasn't done certain tasks. Hence, I'd like to clarify what AutoUpgrade can do, what it can't and what you'll have to do at the moment. I refer to the AutoUpgrade tool as of July 2019. In later versions, one or the other restriction may be lifted. I will blog about it then as well.'

**Questions?**



**Thank you for your attention!**



# SAVE THE DATE

- OATUG Forum Live

December 7-11, 2020

Virtual Event

<https://forum.oatug.org>



- MOUS 2021



October 27, 2021

Schoolcraft College - VisTaTech Center,  
18600 Haggerty Rd, Livonia, MI



THANK YOU

[www.mous.us](http://www.mous.us)



# SURVEYS

- Session Surveys

Please complete the session survey for this session using the Zoom session survey link.

The survey link will be provided via email once the webinar is closed.



THANK YOU

[WWW.MOUS.US](http://WWW.MOUS.US)

