

What's New in Oracle Data Pump?

Oracle Database 12c Release 2

Roy F Swonger
Vice President, Database Upgrade & Utilities
December 6, 2016

Safe Harbor Statement

The following is intended to outline our general product direction. It is intended for information purposes only, and may not be incorporated into any contract. It is not a commitment to deliver any material, code, or functionality, and should not be relied upon in making purchasing decisions. The development, release, and timing of any features or functionality described for Oracle's products remains at the sole discretion of Oracle.

What's New in Data Pump?

Parallel Export/Import of Metadata

Substitution Variables & Wildcards

REMAP_DIRECTORY

Long Identifier support

TRUST_EXISTING_TABLE_PARTITIONS

Validation & Verification options

Other 12.2 Features

Recap of 12.1 Features



Parallel Metadata Export: How it USED TO Work

Pre-12.2

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

Parallel Metadata Export: How it Works Now

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
- Notes:
 - Works for dumpfile jobs, and for network jobs if destination database is 12.2
 - 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 10:55:56.949: 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

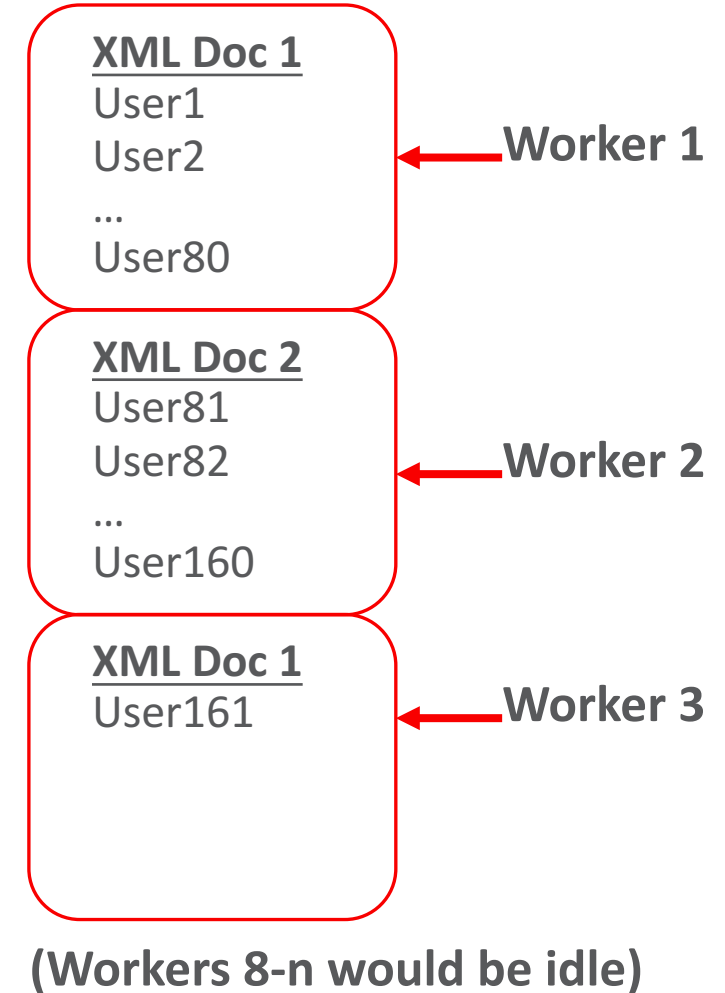
```
18-SEP-16 15:24:32.166: Starting "SYSTEM"."MD_EXP_16_12201": system/***** parfile=md_exp_16_12201.par
18-SEP-16 15:24:32.742: W-1 Startup took 2 seconds
18-SEP-16 15:24:35.601: W-3 Startup took 3 seconds
18-SEP-16 15:24:36.148: W-2 Startup took 3 seconds
18-SEP-16 15:24:36.205: W-4 Startup took 4 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 15:24:37.358: 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
18-SEP-16 15:24:37.665: W-6 Processing object type DATABASE_EXPORT/SCHEMA/SPR...
```

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
 - Available as 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 (yet) for transportable jobs or 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 PARALLEL=32	12.1.0.2 PARALLEL=8	12.1.0.2 PARALLEL=32 With Patch	12.1.0.2 PARALLEL=8	12.2.0.1 PARALLEL=32	Comments
OBJECT_GRANT (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
PROCTACT_SCHEMA	606	198	216	214	152	175	Handled by single worker
TABLE	33164	923	1160	1298	368	248	
OBJECT_GRANT (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 2147483646

- 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
```

```
...
```

- New Feature: 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/ ' "
```

Long Identifier Support

- Long (128-byte) identifiers are supported in 12.2 by
 - Data Pump expdp/impdp
 - SQL*Loader
 - ORACLE_LOADER access driver
 - ORACLE_DATAPUMP access driver
 - External Tables
- Importing database must support 128-byte identifiers
 - Be careful when exporting with VERSION=12.1 or earlier

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 and 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 will be 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=md16_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 News 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**

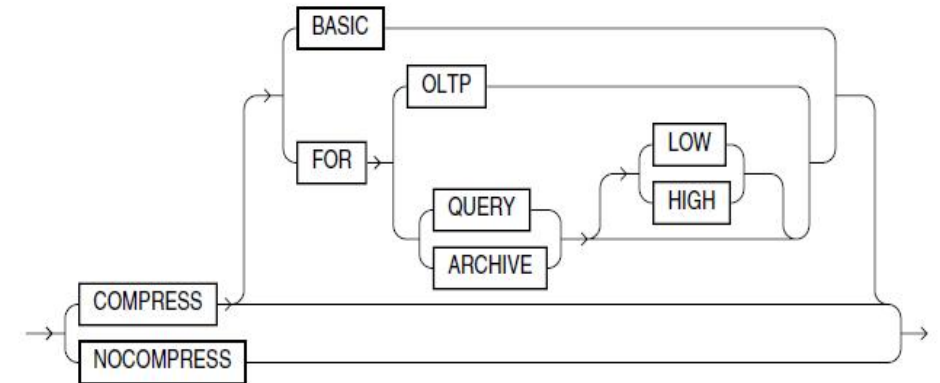
- Downside: Will slow down import!!!
or:

- **Precreate the tablespace with COMPRESS option**

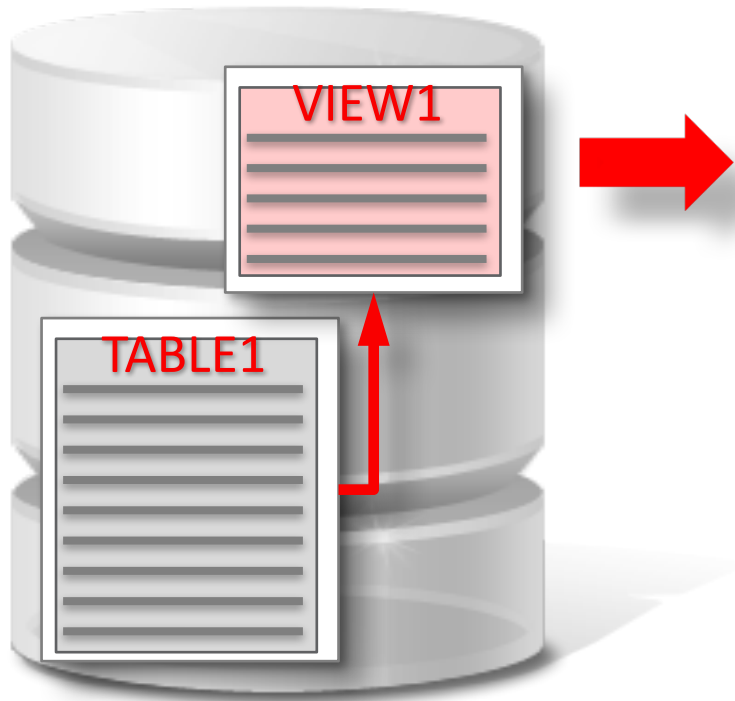
- create tablespace ARCHHIGH datafile 'archhigh.ora' size 100G default compress for archive high;

- 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

```
oracle@localhost:u01/app/oracle/admin/CDB2/dpdump/4273B83C66E31A8BE05500000000001
File Edit View Search Terminal Help
Expd
Copy
;;;
Export: Release 12.2.0.1.0 - Production on Tue Nov 29 17:07:30 2016
Conr
Star
Copyr
Proc
;;;
Export: Release 12.2.0.1.0 - Production on Tue Nov 29 17:10:32 2016
Proc
le=du
Proc
W-1 S Copyright (c) 1982, 2016, Oracle and/or its affiliates. All rights reserved.
Proc
W-1 P
Proc
W-1 P
Proc
W-1 P Connected to: Oracle Database 12c Enterprise Edition Release 12.2.0.1.0 - 64bit Production
Proc
W-1 P 29-NOV-16 17:10:39.396: Starting "SYSTEM"."SYS_EXPORT_FULL_01": system/*****@pdb2 full=y logfile=logtime.log logtime=al
Proc
W-1 P l metrics=y reuse_dumpfiles=y dumpfile=dump.dmp
Proc
W-1 P 29-NOV-16 17:10:40.101: W-1 Startup took 6 seconds
Proc
W-1 P 29-NOV-16 17:10:46.017: W-1 Processing object type DATABASE_EXPORT/EARLY_OPTIONS/VIEWS_AS_TABLES/TABLE_DATA
Proc
W-1 P 29-NOV-16 17:10:47.642: W-1 Processing object type DATABASE_EXPORT/NORMAL_OPTIONS/TABLE_DATA
Proc
W-1 P 29-NOV-16 17:10:49.045: W-1 Processing object type DATABASE_EXPORT/NORMAL_OPTIONS/VIEWS_AS_TABLES/TABLE_DATA
Proc
W-1 P 29-NOV-16 17:10:50.208: W-1 Processing object type DATABASE_EXPORT/SCHEMA/TABLE/TABLE_DATA
Proc
W-1 P 29-NOV-16 17:10:50.554: W-1 Processing object type DATABASE_EXPORT/SCHEMA/TABLE/INDEX/STATISTICS/INDEX_STATISTICS
Proc
W-1 P 29-NOV-16 17:10:50.738: W-1 Processing object type DATABASE_EXPORT/SCHEMA/TABLE/STATISTICS/TABLE_STATISTICS
Proc
W-1 P 29-NOV-16 17:10:50.953: W-1 Completed 1 TABLE_STATISTICS objects in 0 seconds
Proc
W-1 P 29-NOV-16 17:10:55.874: W-1 Processing object type DATABASE_EXPORT/STATISTICS/MARKER
Proc
W-1 P 29-NOV-16 17:10:56.197: W-1 Completed 1 MARKER objects in 5 seconds
Proc
W-1 P 29-NOV-16 17:10:56.214: W-1 Processing object type DATABASE_EXPORT/PRE_SYSTEM_IMPCALLOUT/MARKER
Proc
W-1 P 29-NOV-16 17:10:56.222: W-1 Completed 1 MARKER objects in 5 seconds
Proc
W-1 P 29-NOV-16 17:10:56.234: W-1 Processing object type DATABASE_EXPORT/PRE_INSTANCE_IMPCALLOUT/MARKER
Proc
W-1 P 29-NOV-16 17:10:56.239: W-1 Completed 1 MARKER objects in 5 seconds
Proc
W-1 P 29-NOV-16 17:10:56.318: W-1 Processing object type DATABASE_EXPORT/TABLESPACE
Proc
W-1 P 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

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	10.7	
slot03	6	6.7%	103354.8 0.0	100.9	202.2	11.1	
slot05	6	9.0%	130420.9 7.0	127.4	262.0	97.8	
slot06	6	10.5%	158841.9 175.3	155.3	329.3		
slot08	6	8.4%	130835.3 0.0	127.8	256.0	11.0	
slot09	6	10.1%	136525.9 0.0	133.3	267.0	11.3	
slot10	6	6.6%	140383.4 0.0	137.1	275.0	10.6	
slot11	6	6.8%	112600.0 2.0	110.0	220.7	10.3	
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	50.7	
slot03	6	16.0%	273037.4 11.5	266.6	535.1	50.3	
slot05	6	15.4%	264851.1 17.5	258.7	519.0	50.3	
slot06	6	13.2%	222160.7 425.5	217.4	502.4	43.1	
slot08	6	15.0%	267156.6 1.5	260.9	523.3	50.5	
slot09	6	14.8%	263140.4 6.5	257.0	515.3	50.6	
slot10	6	14.6%	259603.7 2.5	253.5	508.5	50.5	
slot11	6	14.9%	258113.0 5.4	252.1	505.8	50.4	
Groups= 8 TOTALS	48	2.5%	2063833.5 470.4	2015.9	4110.285		

2x

Download Slides from the Database Upgrade Blog

• <http://blogs.oracle.com/UPGRADE>

Upgrade your Database - NOW!
Ease your Oracle Database upgrades and migrations - Best Practices, Workshops, Projects - and something about the pleasures of traveling

Recent Posts

- Upgrade NOW! - OTN Interview at Collaborate16
- Are BPs, PSUs and Proactive BPs cumulative?
- Can I apply a BP on top of a PSU? Or vice versa?
- Upgrade to Oracle Database 12c: We don't insist :-)
- Incremental Statistics Collection in Oracle 12.1.0.2 - A True Story
- Incremental Statistics Collection in Oracle 12.1.0.2 - Upgrade Pitfalls
- MOS Note 1454618.1: Quick Reference to Database PSUs, CPUs, BPs and Patchsets
- New PREUPGRD.SQL is available - Upgrade 12c - Apr16
- Oracle Database BP April16 applied successfully
- Oracle April 2016 PSU and Proactive BPs are there

Are BPs, PSUs and Proactive BPs cumulative?
By Mike Dietrich-Oracle on May 04, 2016

Are Bundle Patches (BPs) and Patch Set Updates (PSUs) cumulative?

That is a question sounding trivial to many people but actually it does get asked quite often. And sometimes I forget to mention this during the workshops - and luckily usually somebody asks the question reminding me to explain it.

Yes, Bundle Patches and Patch Set Updates (and of course Proactive Bundle Patches and Critical/Security Patch Updates (CPUs/SPUs) are all cumulative.

You'll find this mentioned in the first paragraph of [MOS Note: 854428.1 - Patch Set Updates for Oracle Products](#)

My Oracle Support Note: 854428.1
January 15, 2013

Patch Set Updates (PSUs) are proactive **cumulative** patches containing recommended bug fixes t...

Interesting note on the side:
I would have expected this important piece of information in [MOS Note:1962125.1 - Oracle Database - Overview of Database Patch Delivery Methods](#) but I couldn't find it. So it's no wonder why people ask such a trivial question ... [irony!]

Two simple examples:

- You have the October 2015 PSU applied
- You'd like to apply the April 2016 PSU on top
 - Then you don't need the January 2016 PSU as it is included in the April 2016 PSU already
- You never applied a Proactive Bundle Patch
- You'd like to apply the April 2016 Proactive Bundle because a guy recommended it on an Oracle blog - and actually MOS notes mention it as well as highly recommended
 - You don't need to apply anything beforehand.
The April 2016 Proactive BP has all the fixes from all previous BPs included on top of Oracle Database 12.1.0.2

Further Information?

About

Mike Dietrich
Master Product Manager - Database Upgrade & Migrations - Oracle

Based in Germany. Interlink between customers/partners and the Upgrade Development. Running workshops between Arctic and Antarctica. Assisting customers in their reference projects onsite and remotely. Connect via:

[Twitter](#) [LinkedIn](#) [Xing](#)

Search

Enter search term:

Search only this blog

Slides Download Center

Comprehensive
Upgrade, Migrate & Consolidate to Oracle Database 12c
Refreshed 23-SEP-2015

Upgrade Best Practices - 12c
(latest update on 8-NOV-2014)

Upgrade Methods
(Refresh: 8-NOV-2014)

What's New with Upgrades to 12c?
Upload: 8-NOV-2014

Webcast for ISVs Apr-2015: Why Upgrade to Oracle 12c?
Upload: 21-APR-2015

Deep Dive

Parallel Multitenant Upgrades
c atct.pl *Internals*
Upload: 27-NOV-2014

Full Transportable Export/Import with RMAN Incrementals
Upload: 31-JUL-2015

Single Tenant for DBAs
Upload: 30-JUL-2015

Hands On Lab

Hands On Lab
Upgrade, Migrate, Consolidate to 12c
Uploaded 9-FEB-2015

Slides Download Center

Comprehensive
Upgrade, Migrate & Consolidate to Oracle Database 12c
Refreshed 23-SEP-2015

Upgrade Best Practices - 12c
(latest update on 8-NOV-2014)

Upgrade Methods
(Refresh: 8-NOV-2014)

What's New with Upgrades to 12c?
Upload: 8-NOV-2014

Webcast for ISVs Apr-2015: Why Upgrade to Oracle 12c?
Upload: 21-APR-2015

Deep Dive

Parallel Multitenant Upgrades
c atct.pl *Internals*
Upload: 27-NOV-2014

Full Transportable Export/Import with RMAN Incrementals
Upload: 31-JUL-2015

Single Tenant for DBAs
Upload: 30-JUL-2015

Hands On Lab

Hands On Lab
Upgrade, Migrate, Consolidate to 12c
Uploaded 9-FEB-2015



Integrated Cloud

Applications & Platform Services

ORACLE®