

Questions from Cross Platform Migration – Transportable Tablespaces to the Extreme

Virtual Classroom 18

February, 20242024, Version [1.0]
Copyright © 2024, Oracle and/or its affiliates
Public

Disclaimer

This document in any form, software or printed matter, contains proprietary information that is the exclusive property of Oracle. Your access to and use of this confidential material is subject to the terms and conditions of your Oracle software license and service agreement, which has been executed and with which you agree to comply. This document and information contained herein may not be disclosed, copied, reproduced or distributed to anyone outside Oracle without prior written consent of Oracle. This document is not part of your license agreement nor can it be incorporated into any contractual agreement with Oracle or its subsidiaries or affiliates.

This document is for informational purposes only and is intended solely to assist you in planning for the implementation and upgrade of the product features described. 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, timing, and pricing of any features or functionality described in this document remains at the sole discretion of Oracle. Due to the nature of the product architecture, it may not be possible to safely include all features described in this document without risking significant destabilization of the code.

M5 procedure

What are the requirements for the source and target database release?

Both the source and target databases must be on 19.18.0 or later. Also, you must have the [Data Pump Bundle Patch](#) in both places.

What if my source database is Oracle Database 18c or lower?

Then, you must use the former process, XTTS v4 Perl script. We will not enable the use of M5 in lower releases or patch levels.

Can I migrate from a non-CDB and directly into a PDB?

Yes. Transportable tablespace has this feature, and it is supported by RMAN and Data Pump as well.

Does the M5 procedure support migrations to other clouds?

Yes. Any platform that allows you to use transportable tablespaces will work.

Does the M5 script support Microsoft Windows?

The M5 script is *just* a bash script. The core functionality is part of RMAN and Data Pump and works on Windows as well. If you have a bash-compatible terminal on Windows, the script should work (Oracle has not tested this). Otherwise, you can use the RMAN and Data Pump commands through a scripting language of your choice.

Do you recommend M5 over regular Data Pump export/import for smaller databases?

If you can perform the migration with a regular Data Pump export/import within the designated window, you should do that. That would most often be the simplest approach. If Data Pump export/import is too slow, you can use M5 instead.

Can I use M5 for same-endian/same-platform migration?

Yes. RMAN automatically determines whether a conversion of the data files is needed.

Can I use M5 for Red Hat Enterprise Linux migration to Exadata?

Yes. Although a regular Data Guard setup appears to be easier.

Will the procedure support the use of DBMS_FILE_TRANSFER?

No

Do you plan on porting the M5 scripts to Python, Perl, or other languages?

No. The core functionality used by the script exists in RMAN and Data Pump, and you can call this functionality through any language you prefer.

Is the M5 script part of the Oracle home or downloadable as a patch?

No. The M5 script is a separate package that you can download from My Oracle Support.

XTTS v4 Perl script has issues when Fast Recovery Area is enabled for a PDB. Does that apply to M5 as well?

No. Our testing shows that M5 procedure is not affected by this issue.

Can we group multiple self-contained tablespaces and run multiple backups at same time using M5?

You could use such an approach for the XTTS v4 Perl script to get better parallelism. However, in M5 we use a different type of parallelism, and this approach is not recommended, and it will interfere with the Full Transportable Export and Import.

Currently we have been using XTTS V4.3 script. Then what are the major difference between V4 and M5?

We are using newer and more modern technology which enables:

- Migrating encrypted databases
- Better parallelism

ORACLE

- Transparent procedure
- Better support for Data Guard
- Better support for multitenant architectures

Can you estimate the downtime window based on the database size?

No. There are many factors at play. In the webinar, we discuss how to estimate the RMAN part. But to get reliable numbers on RMAN and Data Pump jobs you must perform a test.

What are the requirements for the NFS share?

We recommend configuring according to [Mount Options for Oracle files for RAC databases and Clusterware when used with NFS on NAS devices \(Doc ID 359515.1\)](#).

Any requirements to the COMPATIBLE parameter? Do I need to set that to 19.18.0?

No. However, we generally recommend using the default COMPATIBLE parameter of a release, and that is what we tested as well.

Can I use M5 procedure for migrating packaged applications?

Yes. As long as the application vendor does not specify certain requirements.

M5 creates guaranteed restore points before the import. Do I need to enable FLASHBACK DATABASE?

No. M5 creates the restore point regardless of the state of FLASHBACK DATABASE.

Can I perform the final level 1 backup and the full transportable export simultaneously?

Yes. M5 will do that in a later release.

RMAN

Can I use my own backup as the starting point (level 0) for my migration?

No. You must perform the backup using the M5 driver script. Otherwise, essential information like SCN is not available to the script.

Can I continue my regular backups on the source database while I prepare for the migration with M5 backups?

Yes. You should still perform your regular backups. The M5 backups get recorded in the control file and recovery catalog (if used), so in case of a restore the database might decide to restore or recover from the M5 backups.

What is the default section size used by the M5 script?

The sample properties file uses 64G section size. You should use your knowledge of the database to set an optimal section size. Depending on the size of the data files RMAN might decide to use a different section size after all.

Does the M5 script supported compressed backup sets?

Version 1 does not allow you to use compressed backup sets. We have done tests already and expect to include it in a coming version. Compressed backup sets require a license for Advanced Compression Option.

Is the use of multisection backups mandatory?

Yes. The script always adds the section size clause.

How does M5 use parallel in RMAN?

In contrast to XTTS v4 Perl script, M5 uses parallel in RMAN by allocating multiple channels for the operation. RMAN handles the parallel activities internally.

Can I see the RMAN commands used by M5?

Yes. You can find this in the M5 script. Alternatively, when you perform an actual migration the RMAN backup and restore scripts are stored on disk and may be viewed.

Can we use a tape library as staging area instead of disk or ZDLRA?

In theory it should work, however, we have not tested that combination. Normally, we don't recommend this a tape library is unlikely to deliver the required performance for a migration.

Can I add tablespaces while I run the migration backups?

Yes, but you must add the new tablespace to the tablespace file in the M5 script folder. If you forget to add the tablespace to the list of files, you will get an error during the Data Pump transportable import.

Can I add a data file while I run the migration backups?

Yes. RMAN handles this automatically.

Can I use FROM TIMESTAMP instead of FROM SCN in the M5 backups?

No. The script uses FROM SCN. The M5 script handles this transparently for you, so you don't even see the difference.

If we have any issue with an L1 recovery, can we use flashback database to go back to the previous restore point?

In theory that should work. However, normally RMAN recovers from such situations without needing a flashback.

Does M5 script use SCN to track progress like XTTS v4 Perl script?

Yes.

Can I perform backups from a source standby database?

Yes.

Can I perform backups from a source standby database using block change tracking?

Yes. But be aware that using block change tracking on a standby database requires a license for Active Data Guard.

Can I perform backups from a cascading standby database?

Yes. That should be possible, however, we have not tested that combination.

What if I need to switch over to another database while I take L1 backups?

You can perform the M5 backups on any database in your Data Guard configuration.

We've seen issues with block change tracking. Anything to be worried about when doing the migration?

No. Our tests and our reference projects did not experience any issues with block change tracking. Generally, we advise you to patch your databases to the most recent Release Updates before starting the migration.

Is there a maximum number of level 1 backups that I can perform?

No.

If I migrate into a CDB with other PDBs, do I need to start a full backup of the entire CDB, including all PDBs, after the migration?

No. But we recommend starting a full backup of all the data files you are plugging in.

If we decide to postpone the migration after taking the final level 1 backup, can I continue with new level 1 backups?

Yes. There is no difference between level 1 and final level 1 backup, except the latter is taken with tablespaces marked as read-only.

Data Pump

Can I use Full Transportable Export/Import for my Oracle Database 11.2.0.4?

No. You can perform an FTEX export from 11.2.0.4, however, the source must be 12.1 or higher. You can still use transportable tablespaces, but you must use the traditional method.

What do I do if Full Transportable Export/Import is too slow for my downtime window?

In Oracle Database 21c and later you can run the transportable jobs in parallel which dramatically speed up the process. You can also ensure there is no superfluous data, especially metadata, in the source database. Also, archive as much data out of the source database as possible.

Can I perform a parallel FTEX import if the export was not parallel?

Yes. The export parallel degree is completely independent from the parallel import.

ORACLE

Is the Data Pump Bundle Patch a one-off patch for Oracle Database?

Yes. It is a merge patch available for each Release Update on Oracle Database 19c.

Does Oracle Database 19c support parallel transportable jobs?

No. This requires Oracle Database 21c.

Should parallelism always be a power of 2?

No.

Can I set `JOB_QUEUE_PROCESSES=0` when performing the Data Pump export and when tablespaces are read-only?

You could, but we don't see the benefit of this. You could on the other hand disable the database maintenance window (auto-tasks, like statistics gathering, SQL tuning task) if it is about to start.

Can I start the import driver script (`impdp.sh`) directly in *run* mode without using *test* mode first?

Yes. But we don't recommend this. You should always use test mode first to validate the Data Pump parameter file.

Does Data Pump transport information related to Oracle Automatic Data Optimization?

Partially. The ILM policies defined on the object is added in the target database, however, the heat map information is not transferred.

Does Data Pump transport AWR data?

No. You need to manually move this data if you want it in the target database. See [How to Export and Import the AWR Repository From One Database to Another \(Doc ID 785730.1\)](#).

Is there a dry-run option for Data Pump transportable jobs?

No. But the webinar covers a few ways to test the Data Pump part.

Will Full Transportable Export/Import replace traditional tablespaces exports and imports?

No. But we continue to invest time in making Full Transportable Export/Import faster and better.

Can I install Data Pump Bundle Patch on Exadata Database Service?

Yes.

Encryption

How do we deal with encrypted tablespaces?

It is fully supported to migrate a database with encrypted tablespace. Please see the MOS note for additional steps to export and import the TDE keys.

Does M5 support encrypting the tablespace during the migration?

No, the migration moves the tablespace "as-is". An unencrypted tablespace is still unencrypted when it reaches the target database. But after the migration, you can perform an online encryption of the tablespace.

Character Set

Can I change the character set to UTF-8 during the migration?

This would only be possible if the source character set was a proper binary subset of AL32UTF8. That means that all the characters in the source character set would have to have the exact same representation in AL32UTF8.

This is the case in some situations with US7ASCII. There may be others, check the documentation for the allowed cases.

E-Business Suite

Is the M5 procedure certified by EBS?

No, not at time of writing. However, we expect the EBS team to certify it shortly as it brings many benefits to EBS migration as well. You should always follow the up-to-date recommendation from the EBS team. Until M5 is certified you can use XTTS v4 Perl script.

Statistics

If we exclude statistics from the Data Pump job, when should we re-gather statistics?

Immediately after the Data Pump import, you can start to gather dictionary statistics and afterward proceed with statistics on user objects. While you are gathering statistics you can perform other activities on the target database.

Is there a permanent fix for the issue with expression statistics?

The fix is part of 19.18.0 but it must be selectively enabled using `DBMS_OPTIM_BUNDLE`.

Miscellaneous

Do you recommend disabling Unified Auditing for the duration of the migration?

No. But we do recommend that you archive as much audit data as possible, or moves the audit records outside the maintenance window.



Connect with us

Call **+1.800.ORACLE1** or visit **oracle.com**. Outside North America, find your local office at: **oracle.com/contact**.

 blogs.oracle.com

 facebook.com/oracle

 twitter.com/oracle

Copyright © 2024, Oracle and/or its affiliates. This document is provided for information purposes only, and the contents hereof are subject to change without notice. This document is not warranted to be error-free, nor subject to any other warranties or conditions, whether expressed orally or implied in law, including implied warranties and conditions of merchantability or fitness for a particular purpose. We specifically disclaim any liability with respect to this document, and no contractual obligations are formed either directly or indirectly by this document. This document may not be reproduced or transmitted in any form or by any means, electronic or mechanical, for any purpose, without our prior written permission.

Oracle, Java, MySQL, and NetSuite are registered trademarks of Oracle and/or its affiliates. Other names may be trademarks of their respective owners.