



Migration Toolkit

Version 55

1	Migration Toolkit	4
2	Release notes	5
2.1	Migration Toolkit 55.11.0 release notes	6
2.2	Migration Toolkit 55.10.0 release notes	7
2.3	Migration Toolkit 55.9.0 release notes	8
2.4	Migration Toolkit 55.8.0 release notes	9
2.5	Migration Toolkit 55.7.0 release notes	10
2.6	Migration Toolkit 55.6.1 release notes	11
2.7	Migration Toolkit 55.6.0 release notes	12
2.8	Migration Toolkit 55.5.0 release notes	13
2.9	Migration Toolkit 55.4.0 release notes	14
2.10	Migration Toolkit 55.3.0 release notes	15
2.11	Migration Toolkit 55.2.3 release notes	16
2.12	Migration Toolkit 55.2.2 release notes	17
2.13	Migration Toolkit 55.2.1 release notes	18
2.14	Migration Toolkit 55.2.0 release notes	19
2.15	Migration Toolkit 55.1.0 release notes	20
2.16	Migration Toolkit 55.0.0 release notes	21
3	Supported platforms and databases	22
4	Overview	24
5	Installing Migration Toolkit	27
5.1	Installing Migration Toolkit on Linux x86 (amd64)	29
5.1.1	Installing Migration Toolkit on RHEL 9 or OL 9 x86_64	30
5.1.2	Installing Migration Toolkit on RHEL 8 or OL 8 x86_64	31
5.1.3	Installing Migration Toolkit on AlmaLinux 9 or Rocky Linux 9 x86_64	32
5.1.4	Installing Migration Toolkit on AlmaLinux 8 or Rocky Linux 8 x86_64	33
5.1.5	Installing Migration Toolkit on SLES 15 x86_64	34
5.1.6	Installing Migration Toolkit on Ubuntu 24.04 x86_64	35
5.1.7	Installing Migration Toolkit on Ubuntu 22.04 x86_64	36
5.1.8	Installing Migration Toolkit on Ubuntu 20.04 x86_64	37
5.1.9	Installing Migration Toolkit on Debian 12 x86_64	38
5.1.10	Installing Migration Toolkit on Debian 11 x86_64	39
5.2	Installing Migration Toolkit on Linux AArch64 (ARM64)	40
5.2.1	Installing Migration Toolkit on RHEL 9 or OL 9 arm64	41
5.2.2	Installing Migration Toolkit on Debian 12 arm64	42
5.3	Installing Migration Toolkit on Linux IBM Power (ppc64le)	43
5.3.1	Installing Migration Toolkit on RHEL 9 ppc64le	44
5.3.2	Installing Migration Toolkit on RHEL 8 ppc64le	45
5.3.3	Installing Migration Toolkit on SLES 15 ppc64le	46
5.4	Installing on Windows	47
5.5	Installing a JDBC driver	49
6	Upgrading	51
7	Specifying connection properties	52
8	Invoking Migration Toolkit	61
8.1	Migration Toolkit command options	66
8.2	Specifying Migration Toolkit command options using a file	93
8.2.1	Creating a file with the command options	94
8.2.2	Executing a migration with a command options file	97

9	Migration errors	99
10	Error codes	105
11	FAQ	115

1 Migration Toolkit

Migration Toolkit is a powerful command-line tool that offers granular control of the migration process. Migration Toolkit helps with migrating database objects and data to an EDB Postgres Advanced Server or PostgreSQL database from:

- Oracle
- MySQL
- Microsoft SQL Server

You can also use Migration Toolkit to migrate database objects and data from Sybase Adaptive Server Enterprise to EDB Postgres Advanced Server or between EDB Postgres Advanced Server and PostgreSQL.

You can install Migration Toolkit with the EDB Postgres Advanced Server installer or by using Stack Builder. Stack Builder is distributed with both the EDB Postgres Advanced Server and the PostgreSQL one-click installer available from the [EnterpriseDB website](#).

Note

The term Postgres in this documentation refers to either an installation of EDB Postgres Advanced Server or PostgreSQL. The term Stack Builder refers to either StackBuilder Plus (distributed with EDB Postgres Advanced Server) or Stack Builder (distributed with the PostgreSQL one-click installer from EnterpriseDB).

2 Release notes

The Migration Toolkit documentation describes the latest version of Migration Toolkit 55 including minor releases and patches. The release notes in this section provide information on what was new in each release. For new functionality introduced in a minor or patch release, there are also indicators within the content about what release introduced the feature.

Version	Release Date
55.11.0	20 Feb 2025
55.10.0	21 Nov 2024
55.9.0	21 Aug 2024
55.8.0	16 May 2024
55.7.0	13 Dec 2023
55.6.1	06 Sep 2023
55.6.0	25 May 2023
55.5.0	14 Feb 2023
55.4.0	29 Nov 2022
55.3.0	06 Oct 2022
55.2.3	16 Jun 2022
55.2.2	10 Mar 2022
55.2.1	13 Jan 2022
55.2.0	2 Dec 2021
55.1.0	20 Sep 2021
55.0.0	19 Mar 2021

2.1 Migration Toolkit 55.11.0 release notes

Released: 20 Feb 2025

New features, enhancements, bug fixes, and other changes in Migration Toolkit 55.11.0 include:

Type	Description	Ticket
Enhancement	Added support for migration of user-defined data types from a source Microsoft SQL Server database to a target PostgreSQL/EDB Postgres Advanced Server database.	#41617
Enhancement	Added certification for installation and use on RHEL 9 ARM64 architectures.	
Bug fix	Fixed the issue where an offline data migration generated invalid COPY format contents when table-level parallel loading option was enabled.	#41558
Bug fix	Fixed an issue where an Oracle view based on a database link was skipped during migration.	#42798
Bug fix	Fixed an issue that caused a migration of a MySQL table containing a <code>TIMESTAMP</code> column with an <code>ON UPDATE</code> clause to fail.	
Bug fix	Fixed an issue where migrating a partitioned table in a PostgreSQL source database with a partition key column named 'partition' to a target EDB Postgres Advanced Server database failed.	#43595
Bug fix	Fixed an issue that caused a migration of a Microsoft SQL Server table containing a <code>SYSNAME</code> type column to fail.	#45034
Bug fix	Fixed an issue that caused a migration of an Oracle table containing a <code>VARRAY</code> type column to fail.	
Bug fix	Improved error handling to log error information when a Null Pointer Exception (NPE) occurs during a migration failure.	

2.2 Migration Toolkit 55.10.0 release notes

Released: 21 Nov 2024

New features, enhancements, bug fixes, and other changes in Migration Toolkit 55.10.0 include:

Type	Description	Ticket
Enhancement	Migration Toolkit now supports PostgreSQL version 17 and EDB Postgres Advanced Server version 17 both as a source and target database.	
Enhancement	Improved the handling and consistency in migrating virtual columns across different source and target database combinations. You can now choose whether to migrate virtual columns as virtual columns or map them to regular columns in the target. See mapVirtualToRegular for more information.	#38897
Bug fix	Fixed an issue that caused synonyms from all source schemas to be migrated even when only one schema was selected for migration.	#41178
Bug fix	Fixed an issue where the MySQL column type <code>BIGINT AUTO_INCREMENT</code> isn't mapped to the <code>BIGSERIAL</code> column type in the target PostgreSQL/EDB Postgres Advanced Server table.	
Bug fix	Fixed the issue where migration fails in parallel mode for a table with a primary key based on the float data type.	
Bug fix	Fixed an issue where the indexes are skipped from migration when multiple tables are selected for a source MySQL database.	

2.3 Migration Toolkit 55.9.0 release notes

Released: 21 Aug 2024

New features, enhancements, bug fixes, and other changes in Migration Toolkit 55.9.0 include:

Type	Description	Ticket
Enhancement	Migration Toolkit now allows you to specify one or more options via an options file. This makes it easier to manage a long list of object names.	
Enhancement	The validation capability of command-line options now reports more relevant error messages with a reference to the invalid option name/value and its usage.	#38264
Bug fix	Fixed an issue where a column value that falls in the daylight saving time slot for an MS SQL Server timestamp column was being changed when migrating to EDB Postgres Advanced Server.	#103780 #35739
Bug fix	Fixed an issue where the name of an Oracle database link was being truncated when migrated to EDB Postgres Advanced Server.	#35924
Bug fix	Fixed an issue that caused a failure when migrating a source EDB Postgres Advanced Server table column containing comments to a SQL Server destination database.	#37044
Bug fix	Fixed a corner case issue that caused an Oracle to EDB Postgres Advanced Server migration to fail with a “Connection closed” error when the source schema contained a fairly large number of objects.	#35593
Bug fix	Fixed an issue where an Oracle table with multi-byte character name failed to migrate in dblinkOra mode.	#101817 #35497
Bug fix	Fixed an issue that caused a migration from EDB Postgres Advanced Server to Oracle to fail for a TIMESTAMP type with explicit precision specification.	
Bug fix	Fixed a corner case issue where NULL constraints failed to migrate when migrating from a SQL Server source database to an EDB Postgres Advanced Server destination database.	

2.4 Migration Toolkit 55.8.0 release notes

Released: 16 May 2024

New features, enhancements, bug fixes, and other changes in Migration Toolkit 55.8.0 include:

Type	Description	Addresses
Enhancement	Added the ability to automatically attempt to reconnect to the target database and resume or restart the data migration of the table that was being migrated when the connection was lost.	
Enhancement	Migration Toolkit now allows the migration of stored <code>procedure</code> objects from a source PostgreSQL to a target PostgreSQL database.	
Enhancement	The exclude object capability (<code>-exclude</code> with a set of options) now supports excluding the subsets of <code>Sequences</code> , <code>Procedures</code> , <code>Functions</code> , <code>Packages</code> , <code>Synonyms</code> , and <code>Queues</code> from the migration.	#99271 / 31314
Enhancement	The connection sessions initiated by Migration Toolkit are now identifiable with an application name. This facilitates the debugging of issues relevant to a source or target database.	
Enhancement	Migration Toolkit now logs the names of missing tables as part of the error message when one or more tables are missing from the source database.	#98134 / 30362
Bug fix	Fixed an issue that caused migration failures for Oracle tables containing a numeric column with an empty default clause.	#100987 / 32874
Bug fix	Fixed an issue whereby user-defined NOT NULL check constraints were skipped from migration.	#102630 / 35543
Bug fix	Fixed an issue that caused migrations of MySQL tables with a JSON type column to fail.	#102690 / 35550
Bug fix	Fixed an issue where data migration for a partition table fails with duplicate errors when table-level parallel loading is enabled (<code>dataLoaderCount > 1</code>).	
Bug fix	Fixed the handling of the case-sensitive identifier when the entire MySQL table name is in upper case.	

2.5 Migration Toolkit 55.7.0 release notes

Released: 13 Dec 2023

New features, enhancements, bug fixes, and other changes in Migration Toolkit 55.7.0 include:

Type	Description
Enhancement	The Migration Toolkit is now certified to work with all the BigAnimal supported cluster types including the distributed high availability clusters.
Enhancement	Improved the performance of table-level parallel loading of many tables by updating the operation to fetch all the constraints in a single network round trip. [Support ticket #97379]
Bug fix	Fixed an issue where user-created triggers on the target PostgreSQL or EDB Postgres Advanced Server databases were getting disabled during the data migration phase. [Support ticket #97637]
Bug fix	Fixed an issue that caused reevaluation of the same table with an indefinite loop for table-level parallel loading mode. [Support ticket #96054]
Bug fix	Fixed an issue where data copy was failing while migrating from a source Oracle database configured with the Latin9 character set. [Support ticket #99475, #99506]
Bug fix	Fixed an issue where the copy API was failing to copy the data and a subsequent attempt to copy the data with the bulk insert mode was failing with a duplicate error in the target database.[Support ticket #99506, #97784]
Bug fix	Fixed an issue where a data copy was failing while migrating a source SQL server table with a case-sensitive column name in the EDB Postgres Advanced server target database.
Bug fix	Fixed an issue where a data copy was failing while migrating a source MySQL table due to a case-sensitive table name.
Bug fix	Fixed an issue for partitioned tables with a <code>SUBPARTITION TEMPLATE</code> clause that caused a partitioned table sub-partitioned by range to fail to migrate.
Bug fix	Fixed an issue where the <code>capture_admin</code> database system group was getting included when the <code>-allgroups</code> option was specified for migration from the source EDB Postgres Advanced Server version 16 database.

2.6 Migration Toolkit 55.6.1 release notes

Released: 06 Sep 2023

New features, enhancements, bug fixes, and other changes in Migration Toolkit 55.6.1 include:

Type	Description
Enhancement	Updated the table-level parallel loading capability with to reduce the overhead on the source database by using range-based criterion for loading each individual table data chunk instead of a fetch-offset approach. This optimization is applicable when the table primary key/unique constraint is based on a non-composite numeric type attribute. [Support ticket #93360]
Bug fix	Fixed the issue where a data copy fails when migrating a source Oracle table that has a full-width character in its name. [Support ticket #93552]
Bug fix	Fixed an issue that caused a data copy to fail when a source MySQL table name contains a dash character. [Support ticket #95794]
Bug fix	Fixed an issue that resulted in table creation failure for a source MySQL table that has a <code>DEFAULT</code> clause defined with one or more single-quote characters. [Support ticket #95794]
Bug fix	Fixed an issue that caused a data copy to fail when a source MySQL table uses a MySQL reserved word as its name. [Support ticket #94822]
Bug fix	Fixed the issue encountered during Microsoft SQL Server to EDB Postgres Advanced Server, Oracle to EDB Postgres Advanced Server, or EDB Postgres Advanced Server to Oracle <code>-dataOnly</code> migrations where foreign key constraints are not restored after the data copy completes if the <code>-tables</code> option is specified and the parent table isn't included in the list of specified tables.

2.7 Migration Toolkit 55.6.0 release notes

Released: 25 May 2023

New features, enhancements, bug fixes, and other changes in Migration Toolkit 55.6 include:

Type	Description
Enhancement	Migration Toolkit now provides two new command options to let you exclude a subset of objects during migration. The <code>-excludeTables</code> option skips a subset of tables when all tables are selected for migration, and the <code>-excludeViews</code> option skips a subset of views when all views are selected for migration.
Bug fix	Fixed the issue where a data migration fails when Migration Toolkit attempts to drop a system generated index. Any system generated indexes are now skipped from removal. [Support ticket: #90282]
Bug fix	Fixed the issue where a data migration fails with “Invalid LOB locator specified” error when migrating an Oracle table with an <code>XMLTYPE</code> column. [Support ticket: #92284]
Bug fix	Fixed the issue where a <code>TRUNCATE</code> is always applied on the target database when performing a data only migration from SQL Server if the source table contains a LOB type column. This behavior is now corrected so that <code>TRUNCATE</code> will only be applied when the user has specified the <code>-truncLoad</code> option.
Bug fix	Updated the Migration Toolkit packaging for Debian so that the JDBC drivers are no longer bundled with Migration Toolkit to make it consistent with the packaging for the other OS platforms supported by Migration Toolkit.

2.8 Migration Toolkit 55.5.0 release notes

Released: 14 Feb 2023

New features, enhancements, bug fixes, and other changes in Migration Toolkit 55.5 include:

Type	Description
Enhancement	Migration Toolkit now supports EDB Postgres Advanced Server version 15 as a source and/or target for different migration permutations.
Enhancement	Migration Toolkit has been enhanced to support the migration of data using the default <code>COPY</code> when the column ordering in the source and target databases are different. [Support ticket #88336]
Bug fix	Fixed an issue where a foreign key constraint based on an unique index constraint in a source PostgreSQL/EDB Postgres Advanced Server database is skipped from migration to the target database. [Support ticket #89491]
Bug fix	Fixed an issue that caused migration to halt for one or more tables in parallel data loading mode.

2.9 Migration Toolkit 55.4.0 release notes

Released: 29 Nov 2022

New features, enhancements, bug fixes, and other changes in Migration Toolkit 55.4 include:

Type	Description
Enhancement	Migration Toolkit is now certified to work with PostgreSQL version 15. PostgreSQL 15 is now supported both as a source and/or target for different migration permutations.
Bug fix	Fixed an issue that caused table constraint failure when multiple MySQL source tables share the same constraint name. [Support ticket #86413]
Bug fix	Fixed an issue that resulted in the data migration failure for a MySQL table row with a text column value ending with a backslash. [Support ticket #86412]
Other	Removed support of PostgreSQL 10 and EDB Postgres Advanced Server (EPAS) 10 for use with Migration Toolkit. These versions of PostgreSQL and EPAS have reached end of life and are no longer supported under the normal support agreements.
Other	Removed support of Oracle 10g for use with Migration Toolkit. This version had previously been tested and certified for use with Migration Toolkit. Although Oracle 10g may continue to work with Migration Toolkit, EDB is no longer performing ongoing verification of it.

2.10 Migration Toolkit 55.3.0 release notes

Released: 06 Oct 2022

New features, enhancements, bug fixes, and other changes in Migration Toolkit 55.3 include:

Type	Description
Enhancement	Migration Toolkit is now certified to work with the Microsoft provided SQL Server JDBC driver version 10.2. Although the open source jTDS driver can still be used for connections to Microsoft SQL Server with Migration Toolkit, it is recommended that the Microsoft provided SQL Server JDBC driver be used instead as it provides support for the more recent versions of SQL Server.
Enhancement	Migration Toolkit is now certified to work on the Debian 11 platform.
Security Fix	Migration Toolkit has been certified with the latest EDB JDBC driver version 42.5.0.1 and PostgreSQL community JDBC driver version 42.5.0 that contain the security fix for CVE-2022-31197 .
Security Fix	Migration Toolkit has been certified to work with the MySQL JDBC driver version 8.0.30 that contains the security fix for CVE-2022-21363 . References to the MySQL JDBC driver in the Migration Toolkit document have been correspondingly updated.
Bug Fix	Fixed an issue where an Oracle table migration to EPAS fails when it contains a virtual column. [Support Ticket # 82718]
Other	Removed support of Microsoft SQL Server 2008 and SQL Server Server 2012 for use with Migration Toolkit since these versions are no longer supported by Microsoft. These versions had previously been tested and certified for use with Migration Toolkit. Although they may continue to work with Migration Toolkit, EDB is no longer performing ongoing verification of them.

2.11 Migration Toolkit 55.2.3 release notes

Released: 16 Jun 2022

New features, enhancements, bug fixes, and other changes in Migration Toolkit 55.2.3 include:

Type	Description
Bug Fix	Fixed an issue where the Oracle table migration to EDB Postgres Advanced Server fails when a column definition includes a sequence with NEXTVAL call as the default clause.
Bug Fix	Fixed an issue that caused a migration failure for a table with a column named after the POSITION keyword.
Bug Fix	Corrected the Migration Toolkit Windows installer so that it removes the older version of commons-lang library as part of the upgrade process.
Bug Fix	Corrected some of the information/error messages to report schema-qualified table names.

2.12 Migration Toolkit 55.2.2 release notes

Released: 10 Mar 2022

New features, enhancements, bug fixes, and other changes in Migration Toolkit 55.2.2 include:

Type	Description
Enhancement	Provides support for SQL Server 2016, 2017, and 2019 and Sybase (SAP ASE) 16.0.
Security Fix	The installers used for Windows and macOS have been updated to include the latest version of the EDB JDBC Connector which addresses CVE-2022-21724 . For Linux, the JDBC driver is not packaged with Migration Toolkit and you will need to update the driver separately using the instructions for Linux documented in this EDB knowledge base article .
Bug Fix	Migration Toolkit now properly handles an exception if ENABLE TRIGGER fails while loading table data. If ENABLE TRIGGER fails for a particular table it is not skipped for the remaining tables.
Bug Fix	Migration Toolkit now escapes backslashes coming in TEXT fields while migrating from a MySQL database.

2.13 Migration Toolkit 55.2.1 release notes

Released: 13 Jan 2022

New features, enhancements, bug fixes, and other changes in Migration Toolkit 55.2.1 include:

Type	Description
Security Fix	The Migration Toolkit logging is now based on log4j 2.17.1. Previously, it had been based on log4j v1. This change removes the vulnerabilities reported against log4j v1 and ensures that the RCE vulnerabilities discovered and reported against older versions of log4j v2 are bypassed.
Bug Fix	Custom column type mapping is fixed from Oracle RAW type to PostgreSQL/Postgres Plus Advanced Server UUID type.
Bug Fix	<code>-loaderCount</code> parameter is disallowed for Sybase, MySQL, and SQL Server databases.
Bug Fix	Fixed duplicate index name issue while migrating from MySQL to PostgreSQL.
Bug Fix	Replace executeQuery with executeUpdate for the query executed on MySQL that is not returning a resultset.

2.14 Migration Toolkit 55.2.0 release notes

Released: 02 Dec 2021

New features, enhancements, bug fixes, and other changes in Migration Toolkit 55.2.0 include:

Type	Description
Enhancement	Migration Toolkit is certified to support PostgreSQL 14.
Enhancement	Added support for EPAS 14, which includes some of the new Oracle compatibility features.
Enhancement	Partition tables with the LIST BY PARTITION clause followed by the AUTOMATIC keyword can now be migrated to target EPAS 13 and 14 databases.
Enhancement	Migration Toolkit now successfully migrates partition tables with the SUBPARTITION TEMPLATE clause to target EPAS 14 databases.
Bug Fix	Fixed the ORA-01795 issue to allow successful migration of more than a thousand users.
Bug Fix	Fixed the ORA-01795 issue to allow successful migration of more than a thousand partitions.

2.15 Migration Toolkit 55.1.0 release notes

Released: 20 Sep 2021

New features, enhancements, bug fixes, and other changes in Migration Toolkit 55.1.0 include:

Type	Description
Enhancement	Migration Toolkit is enhanced to support migration from MySQL 5.7 and 8.0.
Bug Fix	For MySQL to PostgreSQL migration, the boolean data type was not migrating in LO mode. This issue is fixed.
Bug Fix	Fixed an issue to allow migration of table names starting with a number from MySQL to PostgreSQL.
Bug Fix	Fixed an issue where the TIMESTAMP value changes within the same time zone when migrating from MYSQL TO PostgreSQL.
Bug Fix	While migrating from MySQL to PostgreSQL, enum types were not being created on the target database with fastcopy mode. This issue is fixed.
Bug Fix	Fixed failures related to the migration of BIT(n) data type from MySQL to PostgreSQL when $n > 1$.
Bug Fix	Fixed an issue wherein while migrating from Oracle to PostgreSQL, a number format exception was reported (72863).
Bug Fix	Migration of the PostgreSQL table fails when the associated OID value exceeds the INTEGER range. This issue is fixed (72694).
Bug Fix	Fixed an issue related to the migration of JSON data type from PostgreSQL to Oracle.
Bug Fix	For Oracle 19c to EDB Postgres Advanced Server, certain system users/schemas were incorrectly picked for migration. This issue is fixed.
Bug Fix	Fixed an issue to allow migration of case-sensitive schema/table names in parallel loading mode from Oracle to EDB Postgres Advanced Server.
Bug Fix	Certain corner cases that were related to the use of the Custom Column Type mappings option, are fixed.

2.16 Migration Toolkit 55.0.0 release notes

Released: 19 Mar 2021

New features, enhancements, bug fixes, and other changes in Migration Toolkit 55.0.0 include:

Type	Description
Enhancement	Support for parallel data loading at the table level. This allows you to use multiple threads to load data in parallel from the source database table and apply in parallel on the target database table. This significantly improves the data load time when dealing with relatively large size tables. The lab benchmark results have revealed a 25-60% reduction in the data load duration for a large table when using multiple parallel threads. This feature is supported for source PostgreSQL, EDB Postgres Advanced Server, and Oracle databases. See Migration Options for Parallel Data Loading .
Enhancement	As part of the table-level parallel data loading enhancement, the <code>tableLoaderLimit</code> and <code>parallelLoadRowLimit</code> options have been added to the command-line options list.
Enhancement	Support of schema and data migration from Oracle versions 18c and 19c.
Enhancement	Total migration time for the the schema and data migration session is reported via the <code>Total Elapsed Migration Time</code> attribute.
Bug Fix	Summary report does not accurately report status when data migration fails for table(s).
Bug Fix	Summary report does not remove redundant invalid objects.
Bug Fix	Data size for a table reported in exponential format.
Bug Fix	Avoid failures when a duplicate table name is specified via '-tables' list.
Bug Fix	Reflect small table sizes that were previously displaying as 0.0 MB.

3 Supported platforms and databases

Supported database versions

You can use the following database product versions with Migration Toolkit:

- PostgreSQL versions 13, 14, 15, 16, and 17
- EDB Postgres Advanced Server versions 13, 14, 15, 16, and 17
- Oracle 11g
- Oracle 12c
- Oracle 18c
- Oracle 19c
- SQL Server 2014
- SQL Server 2016
- SQL Server 2017
- SQL Server 2019
- MySQL 5.5
- MySQL 5.7
- MySQL 8.0
- (Sybase) SAP Adaptive Server Enterprise 15.7
- (Sybase) SAP Adaptive Server Enterprise 16.0

Note

All the PostgreSQL and EDB Postgres Advanced Server versions available as EDB Postgres AI Cloud Service single-node, primary/standby high-availability, and distributed high-availability cluster types are also supported for use with Migration Toolkit. See the Cloud Service [documentation](#) for more information about Cloud Service's [supported cluster types](#). See the [database version policy documentation](#) for the versions of PostgreSQL and EDB Postgres Advanced Server available in Cloud Service.

The superuser role isn't available in EDB Postgres AI Cloud Service EDB Postgres Advanced Server distributed high-availability cluster databases. As a result, the following Migration Toolkit capabilities that require superuser access aren't supported when targeting Cloud Service distributed high-availability clusters:

- The `-copyViaDBLinkOra` option that enables data migration using the `db_link_ora` extension
- The Oracle-compatible SQL command `CREATE DATABASE LINK`

Contact your EnterpriseDB Account Manager or sales@enterprisedb.com if you require support for other database products.

Note

The Migration Toolkit isn't tested with and doesn't officially support use with Oracle Real Application Clusters (RAC) and Exadata. However, Migration Toolkit might work when it's connected to a single persistent node. For more information, contact your EDB representative.

Supported operating system versions

Migration Toolkit supports installations on Linux and Windows platforms. See [Product Compatibility](#) for details.

Note

The `ojdbc7.jar` can cause large data migrations to be incomplete due to a limit on records over `Integer.MAX_VALUE (2147483647)` while fetching from the `ResultSet`. To avoid this issue, we recommend upgrading to `ojdbc8.jar`.

Supported JDBC drivers

To perform a migration, you must [choose and install JDBC drivers](#) that support your source and target databases. These can be [Postgres JDBC drivers](#), or [non-Postgres JDBC drivers](#).

Supported Postgres JDBC drivers

For Postgres JDBC drivers, Migration Toolkit supports the following PostgreSQL and EDB Advanced Server drivers:

	PostgreSQL JDBC driver	EDB JDBC driver
MTK 55.10.0	JDBC 42.7.3 and later	JDBC 42.7.3.2 and later
MTK 55.9.0	JDBC 42.5.4 and later	JDBC 42.5.4.2 and later
MTK 55.8.0	JDBC 42.5.4 and later	JDBC 42.5.4.2 and later
MTK 55.7.0	JDBC 42.5.4 and later	JDBC 42.5.4.2 and later
MTK 55.6.0	JDBC 42.5.4 and later	JDBC 42.5.4.2 and later
MTK 55.5.0	JDBC 42.5.1 and later	JDBC 42.5.1.2 and later
MTK 55.4.0	JDBC 42.5.1 and later	JDBC 42.5.1.2 and later
MTK 55.3.0	JDBC 42.5.0 and later	JDBC 42.5.0.2 and later
MTK 55.2.0	JDBC 42.5.0 and later	JDBC 42.5.0.2 and later
MTK 55.1.0	JDBC 42.2.12 and later	JDBC 42.2.12.2 and later
MTK 55.0.0	JDBC 42.2.12 and later	JDBC 42.2.12.2 and later

See [Choosing a driver](#) for download information.

Supported non-Postgres database versions

If your source or target database is a third-party, non-Postgres database, you need a third-party JDBC driver.

To ensure compatibility of your Migration Toolkit version with the third-party JDBC driver version, use the latest available JDBC driver version that's compatible with the database you want to migrate from or to.

Check JDBC driver and database compatibility with the third-party provider, for example, in the corresponding Oracle, Microsoft SQL Server, MySQL, or SAP Adaptive Enterprise documentation.

See [Choosing a driver](#) for download information.

4 Overview

Migration Toolkit is a powerful command-line tool that offers granular control of the migration process. Using Migration Toolkit is a two-step process:

1. Edit the `toolkit.properties` file to specify the source and target database.
2. Invoke Migration Toolkit at the command line, specifying migration options.

Migration Toolkit helps with migration of database objects and data to an EDB Postgres Advanced Server or PostgreSQL database from:

- Oracle
- MySQL
- SQL Server

Migration Toolkit also allows you to migrate database objects and data to an EDB Postgres Advanced Server database from Sybase. You can also use Migration Toolkit to migrate between EDB Postgres Advanced Server and PostgreSQL. Migration Toolkit includes a number of options, allowing you granular control of the migration process:

- Use the `-safeMode` option to commit each row as it is migrated.
- Use the `-fastCopy` option to bypass WAL logging to optimize migration.
- Use the `-batchSize` option to control the batch size of bulk inserts.
- Use the `-cpBatchSize` option to specify the batch size used with the COPY command.
- Use the `-lobBatchSize` option to specify the batch size used for large object data types.
- Use the `-filterProp` option to migrate only those rows that meet a condition you specify.
- Use the `-customColTypeMapping` option to change the data type of selected columns.
- Use the `-dropSchema` option to drop the existing schema and create a new schema prior to migration.
- On EDB Postgres Advanced Server, use the `-allDBLinks` option to migrate all Oracle database links.
- On EDB Postgres Advanced Server, use the `-copyViaDBLinkOra` option to enable the `dblink_ora` module.
- Use the `-connRetryCount <connection_attempts>` option to specify the number of retry attempts to perform if the target database connection is lost during data migration.
- Use the `-connRetryInterval <seconds>` option to specify the seconds to wait before each target database reconnection attempt during a data migration.
- Use the `-abortConnOnFailure <true/false>` option to specify whether to abort the migration when all target database reconnection attempts fail during a data migration. The default is `true`, which aborts the session if the connection fails after the the specified `-connRetryCount` threshold.
- Use the `-pgIdleTxSessionTimeout <seconds>` to override the value of the PostgreSQL or EDB Postgres Advanced Server `idle_in_transaction_session_timeout` configuration option in the MTK connection session.

Object migration support

Migration Toolkit migrates object definitions (DDL), table data, or both. The following table contains a platform-specific list of the types of database objects that Migration Toolkit can migrate:

Object	Oracle	Sybase	SQL Server	MySQL
Schemas	X	X	X	X
Tables	X	X	X	X
List-partitioned table	X			
Range-partitioned table	X			
Hash partitioned table	X			
Constraints	X	X	X	X

Object	Oracle	Sybase	SQL Server	MySQL
Indexes	X	X	X	X
Triggers	X			
Table data	X	X	X	X
Views	X		X	
Materialized views	X			
Packages	X			
Procedures	X			
Functions	X			
Sequences	X			
Users/Roles	X			
Profiles	X			
Object types	X			
Object type methods	X			
Database links	X			
Queues	X			

For detailed information about the commands that offer granular control of the objects imported, see [Schema object selection options](#).

Online migration versus offline migration

Migration Toolkit can migrate immediately and directly into a Postgres database (*online migration*). You can also choose to generate scripts to use later to re-create object definitions in a Postgres database (*offline migration*).

By default, Migration Toolkit creates objects directly into a Postgres database. Alternatively, include the `-offlineMigration` option to generate SQL scripts you can use later to reproduce the migrated objects or data in a new database. You can alter migrated objects by customizing the migration scripts generated by Migration Toolkit before you execute them. With the `-offlineMigration` option, you can schedule the actual migration at a time that best suits your system load.

For more information about the `-offlineMigration` option, see [Offline migration options](#).

Limitations

EDB Postgres Advanced Server offers complete support for some Oracle features and partial support for others. Migration Toolkit can't migrate any object that uses an unsupported feature.

In some cases, Migration Toolkit can migrate objects that use features that offer partial compatibility. In other cases, EDB Postgres Advanced Server supports suitable workarounds.

Full-text search is an example of functionality that isn't fully compatible with Oracle. The EDB Postgres Advanced Server database has included support for full-text search for quite some time, but the implementation is different from Oracle's. Migration Toolkit can't migrate objects that use this feature.

EDB Postgres Advanced Server doesn't yet support other features. Features in this category include Automated Storage Management, table compression, and external tables. You can often implement a successful workaround:

- You can replace Automated Storage Management with system-specific volume management software.

- You can implement table compression by storing data in a tablespace that resides on a compressed filesystem.
- External tables don't exist in EDB Postgres Advanced Server, but you can load flat text files into staging tables in the database. We recommend using the EDB*Loader utility to load the data into an EDB Postgres Advanced Server database quickly.
- When migrating multiple profiles from an Oracle database into EDB Postgres Advanced Server, you must manually assign the profile to a user or users when the migration completes.

Unsupported Postgres features

Migration Toolkit doesn't support migration of the following Postgres features:

- `OPERATOR CLASS`
- `OPERATOR FAMILY`

For information about `OPERATOR CLASS` and `OPERATOR FAMILY`, see the [PostgreSQL core documentation](#).

5 Installing Migration Toolkit

Select a link to access the applicable installation instructions:

Linux [x86-64 \(amd64\)](#)

Red Hat Enterprise Linux (RHEL) and derivatives

- [RHEL 9](#), [RHEL 8](#)
- [Oracle Linux \(OL\) 9](#), [Oracle Linux \(OL\) 8](#)
- [Rocky Linux 9](#), [Rocky Linux 8](#)
- [AlmaLinux 9](#), [AlmaLinux 8](#)

SUSE Linux Enterprise (SLES)

- [SLES 15](#)

Debian and derivatives

- [Ubuntu 24.04](#), [Ubuntu 22.04](#), [Ubuntu 20.04](#)
- [Debian 12](#), [Debian 11](#)

Linux [IBM Power \(ppc64le\)](#)

Red Hat Enterprise Linux (RHEL) and derivatives

- [RHEL 9](#), [RHEL 8](#)

SUSE Linux Enterprise (SLES)

- [SLES 15](#)

Linux [AArch64 \(ARM64\)](#)

Red Hat Enterprise Linux (RHEL) and derivatives

- [RHEL 9](#)
- [Oracle Linux \(OL\) 9](#)

Debian and derivatives

- [Debian 12](#)

Windows

- [Windows Server 2022](#)
- [Windows 11](#)

5.1 Installing Migration Toolkit on Linux x86 (amd64)

Operating system-specific install instructions are described in the corresponding documentation:

Red Hat Enterprise Linux (RHEL) and derivatives

- [RHEL 9](#)
- [RHEL 8](#)
- [Oracle Linux \(OL\) 9](#)
- [Oracle Linux \(OL\) 8](#)
- [Rocky Linux 9](#)
- [Rocky Linux 8](#)
- [AlmaLinux 9](#)
- [AlmaLinux 8](#)

SUSE Linux Enterprise (SLES)

- [SLES 15](#)

Debian and derivatives

- [Ubuntu 24.04](#)
- [Ubuntu 22.04](#)
- [Ubuntu 20.04](#)
- [Debian 12](#)
- [Debian 11](#)

5.1.1 Installing Migration Toolkit on RHEL 9 or OL 9 x86_64

Prerequisites

Before you begin the installation process:

- Set up the EDB repository.

Setting up the repository is a one-time task. If you have already set up your repository, you don't need to perform this step.

To determine if your repository exists, enter this command:

```
dnf repolist | grep enterprisedb
```

If no output is generated, the repository isn't installed.

To set up the EDB repository:

1. Go to [EDB repositories](#).
2. Select the button that provides access to the EDB repository.
3. Select the platform and software that you want to download.
4. Follow the instructions for setting up the EDB repository.

Install the package

```
sudo dnf -y install edb-migrationtoolkit
```

Initial configuration

Before invoking Migration Toolkit, you must download and install JDBC drivers for connecting to the source and target databases. See [Installing a JDBC driver](#) for details.

5.1.2 Installing Migration Toolkit on RHEL 8 or OL 8 x86_64

Prerequisites

Before you begin the installation process:

- Set up the EDB repository.

Setting up the repository is a one-time task. If you have already set up your repository, you don't need to perform this step.

To determine if your repository exists, enter this command:

```
dnf repolist | grep enterprisedb
```

If no output is generated, the repository isn't installed.

To set up the EDB repository:

1. Go to [EDB repositories](#).
2. Select the button that provides access to the EDB repository.
3. Select the platform and software that you want to download.
4. Follow the instructions for setting up the EDB repository.

Install the package

```
sudo dnf -y install edb-migrationtoolkit
```

Initial configuration

Before invoking Migration Toolkit, you must download and install JDBC drivers for connecting to the source and target databases. See [Installing a JDBC driver](#) for details.

5.1.3 Installing Migration Toolkit on AlmaLinux 9 or Rocky Linux 9 x86_64

Prerequisites

Before you begin the installation process:

- Set up the EDB repository.

Setting up the repository is a one-time task. If you have already set up your repository, you don't need to perform this step.

To determine if your repository exists, enter this command:

```
dnf repolist | grep enterprisedb
```

If no output is generated, the repository isn't installed.

To set up the EDB repository:

1. Go to [EDB repositories](#).
2. Select the button that provides access to the EDB repository.
3. Select the platform and software that you want to download.
4. Follow the instructions for setting up the EDB repository.

Install the package

```
sudo dnf -y install edb-migrationtoolkit
```

Initial configuration

Before invoking Migration Toolkit, you must download and install JDBC drivers for connecting to the source and target databases. See [Installing a JDBC driver](#) for details.

5.1.4 Installing Migration Toolkit on AlmaLinux 8 or Rocky Linux 8 x86_64

Prerequisites

Before you begin the installation process:

- Set up the EDB repository.

Setting up the repository is a one-time task. If you have already set up your repository, you don't need to perform this step.

To determine if your repository exists, enter this command:

```
dnf repolist | grep enterprisedb
```

If no output is generated, the repository isn't installed.

To set up the EDB repository:

1. Go to [EDB repositories](#).
2. Select the button that provides access to the EDB repository.
3. Select the platform and software that you want to download.
4. Follow the instructions for setting up the EDB repository.

Install the package

```
sudo dnf -y install edb-migrationtoolkit
```

Initial configuration

Before invoking Migration Toolkit, you must download and install JDBC drivers for connecting to the source and target databases. See [Installing a JDBC driver](#) for details.

5.1.5 Installing Migration Toolkit on SLES 15 x86_64

Prerequisites

Before you begin the installation process:

- Set up the EDB repository.

Setting up the repository is a one-time task. If you have already set up your repository, you don't need to perform this step.

To determine if your repository exists, enter this command:

```
zypper lr -E | grep enterprisedb
```

If no output is generated, the repository isn't installed.

To set up the EDB repository:

1. Go to [EDB repositories](#).
2. Select the button that provides access to the EDB repository.
3. Select the platform and software that you want to download.
4. Follow the instructions for setting up the EDB repository.

- Activate the required SUSE module:

```
sudo SUSEConnect -p PackageHub/15.4/x86_64
```

- Refresh the metadata:

```
sudo zypper refresh
```

Install the package

```
sudo zypper -n install edb-migrationtoolkit
```

Initial configuration

Before invoking Migration Toolkit, you must download and install JDBC drivers for connecting to the source and target databases. See [Installing a JDBC driver](#) for details.

5.1.6 Installing Migration Toolkit on Ubuntu 24.04 x86_64

Prerequisites

Before you begin the installation process:

- Set up the EDB repository.

Setting up the repository is a one-time task. If you have already set up your repository, you don't need to perform this step.

To determine if your repository exists, enter this command:

```
apt-cache search enterprisedb
```

If no output is generated, the repository isn't installed.

To set up the EDB repository:

1. Go to [EDB repositories](#).
2. Select the button that provides access to the EDB repository.
3. Select the platform and software that you want to download.
4. Follow the instructions for setting up the EDB repository.

Install the package

```
sudo apt-get -y install edb-migrationtoolkit
```

Initial configuration

Before invoking Migration Toolkit, you must download and install JDBC drivers for connecting to the source and target databases. See [Installing a JDBC driver](#) for details.

5.1.7 Installing Migration Toolkit on Ubuntu 22.04 x86_64

Prerequisites

Before you begin the installation process:

- Set up the EDB repository.

Setting up the repository is a one-time task. If you have already set up your repository, you don't need to perform this step.

To determine if your repository exists, enter this command:

```
apt-cache search enterprisedb
```

If no output is generated, the repository isn't installed.

To set up the EDB repository:

1. Go to [EDB repositories](#).
2. Select the button that provides access to the EDB repository.
3. Select the platform and software that you want to download.
4. Follow the instructions for setting up the EDB repository.

Install the package

```
sudo apt-get -y install edb-migrationtoolkit
```

Initial configuration

Before invoking Migration Toolkit, you must download and install JDBC drivers for connecting to the source and target databases. See [Installing a JDBC driver](#) for details.

5.1.8 Installing Migration Toolkit on Ubuntu 20.04 x86_64

Prerequisites

Before you begin the installation process:

- Set up the EDB repository.

Setting up the repository is a one-time task. If you have already set up your repository, you don't need to perform this step.

To determine if your repository exists, enter this command:

```
apt-cache search enterprisedb
```

If no output is generated, the repository isn't installed.

To set up the EDB repository:

1. Go to [EDB repositories](#).
2. Select the button that provides access to the EDB repository.
3. Select the platform and software that you want to download.
4. Follow the instructions for setting up the EDB repository.

Install the package

```
sudo apt-get -y install edb-migrationtoolkit
```

Initial configuration

Before invoking Migration Toolkit, you must download and install JDBC drivers for connecting to the source and target databases. See [Installing a JDBC driver](#) for details.

5.1.9 Installing Migration Toolkit on Debian 12 x86_64

Prerequisites

Before you begin the installation process:

- Set up the EDB repository.

Setting up the repository is a one-time task. If you have already set up your repository, you don't need to perform this step.

To determine if your repository exists, enter this command:

```
apt-cache search enterprisedb
```

If no output is generated, the repository isn't installed.

To set up the EDB repository:

1. Go to [EDB repositories](#).
2. Select the button that provides access to the EDB repository.
3. Select the platform and software that you want to download.
4. Follow the instructions for setting up the EDB repository.

Install the package

```
sudo apt-get -y install edb-migrationtoolkit
```

Initial configuration

Before invoking Migration Toolkit, you must download and install JDBC drivers for connecting to the source and target databases. See [Installing a JDBC driver](#) for details.

5.1.10 Installing Migration Toolkit on Debian 11 x86_64

Prerequisites

Before you begin the installation process:

- Set up the EDB repository.

Setting up the repository is a one-time task. If you have already set up your repository, you don't need to perform this step.

To determine if your repository exists, enter this command:

```
apt-cache search enterprisedb
```

If no output is generated, the repository isn't installed.

To set up the EDB repository:

1. Go to [EDB repositories](#).
2. Select the button that provides access to the EDB repository.
3. Select the platform and software that you want to download.
4. Follow the instructions for setting up the EDB repository.

Install the package

```
sudo apt-get -y install edb-migrationtoolkit
```

Initial configuration

Before invoking Migration Toolkit, you must download and install JDBC drivers for connecting to the source and target databases. See [Installing a JDBC driver](#) for details.

5.2 Installing Migration Toolkit on Linux AArch64 (ARM64)

Operating system-specific install instructions are described in the corresponding documentation:

Red Hat Enterprise Linux (RHEL) and derivatives

- [RHEL 9](#)
- [Oracle Linux \(OL\) 9](#)

Debian and derivatives

- [Debian 12](#)

5.2.1 Installing Migration Toolkit on RHEL 9 or OL 9 arm64

Prerequisites

Before you begin the installation process:

- Set up the EDB repository.

Setting up the repository is a one-time task. If you have already set up your repository, you don't need to perform this step.

To determine if your repository exists, enter this command:

```
dnf repolist | grep enterprisedb
```

If no output is generated, the repository isn't installed.

To set up the EDB repository:

1. Go to [EDB repositories](#).
2. Select the button that provides access to the EDB repository.
3. Select the platform and software that you want to download.
4. Follow the instructions for setting up the EDB repository.

Install the package

```
sudo dnf -y install edb-migrationtoolkit
```

Initial configuration

Before invoking Migration Toolkit, you must download and install JDBC drivers for connecting to the source and target databases. See [Installing a JDBC driver](#) for details.

5.2.2 Installing Migration Toolkit on Debian 12 arm64

Prerequisites

Before you begin the installation process:

- Set up the EDB repository.

Setting up the repository is a one-time task. If you have already set up your repository, you don't need to perform this step.

To determine if your repository exists, enter this command:

```
apt-cache search enterprisedb
```

If no output is generated, the repository isn't installed.

To set up the EDB repository:

1. Go to [EDB repositories](#).
2. Select the button that provides access to the EDB repository.
3. Select the platform and software that you want to download.
4. Follow the instructions for setting up the EDB repository.

Install the package

```
sudo apt-get -y install edb-migrationtoolkit
```

Initial configuration

Before invoking Migration Toolkit, you must download and install JDBC drivers for connecting to the source and target databases. See [Installing a JDBC driver](#) for details.

5.3 Installing Migration Toolkit on Linux IBM Power (ppc64le)

Operating system-specific install instructions are described in the corresponding documentation:

Red Hat Enterprise Linux (RHEL)

- [RHEL 9](#)
- [RHEL 8](#)

SUSE Linux Enterprise (SLES)

- [SLES 15](#)

5.3.1 Installing Migration Toolkit on RHEL 9 ppc64le

Prerequisites

Before you begin the installation process:

- Set up the EDB repository.

Setting up the repository is a one-time task. If you have already set up your repository, you don't need to perform this step.

To determine if your repository exists, enter this command:

```
dnf repolist | grep enterprisedb
```

If no output is generated, the repository isn't installed.

To set up the EDB repository:

1. Go to [EDB repositories](#).
2. Select the button that provides access to the EDB repository.
3. Select the platform and software that you want to download.
4. Follow the instructions for setting up the EDB repository.

Install the package

```
sudo dnf -y install edb-migrationtoolkit
```

Initial configuration

Before invoking Migration Toolkit, you must download and install JDBC drivers for connecting to the source and target databases. See [Installing a JDBC driver](#) for details.

5.3.2 Installing Migration Toolkit on RHEL 8 ppc64le

Prerequisites

Before you begin the installation process:

- Set up the EDB repository.

Setting up the repository is a one-time task. If you have already set up your repository, you don't need to perform this step.

To determine if your repository exists, enter this command:

```
dnf repolist | grep enterprisedb
```

If no output is generated, the repository isn't installed.

To set up the EDB repository:

1. Go to [EDB repositories](#).
2. Select the button that provides access to the EDB repository.
3. Select the platform and software that you want to download.
4. Follow the instructions for setting up the EDB repository.

Install the package

```
sudo dnf -y install edb-migrationtoolkit
```

Initial configuration

Before invoking Migration Toolkit, you must download and install JDBC drivers for connecting to the source and target databases. See [Installing a JDBC driver](#) for details.

5.3.3 Installing Migration Toolkit on SLES 15 ppc64le

Prerequisites

Before you begin the installation process:

- Set up the EDB repository.

Setting up the repository is a one-time task. If you have already set up your repository, you don't need to perform this step.

To determine if your repository exists, enter this command:

```
zypper lr -E | grep enterprisedb
```

If no output is generated, the repository isn't installed.

To set up the EDB repository:

1. Go to [EDB repositories](#).
2. Select the button that provides access to the EDB repository.
3. Select the platform and software that you want to download.
4. Follow the instructions for setting up the EDB repository.

- Activate the required SUSE module:

```
sudo SUSEConnect -p PackageHub/15.4/ppc64le
```

- Refresh the metadata:

```
sudo zypper refresh
```

Install the package

```
sudo zypper -n install edb-migrationtoolkit
```

Initial configuration

Before invoking Migration Toolkit, you must download and install JDBC drivers for connecting to the source and target databases. See [Installing a JDBC driver](#) for details.

5.4 Installing on Windows

EDB provides a graphical interactive installer for Windows. You can access it two ways:

- Download the graphical installer from the [Downloads page](#) and invoke the installer directly. See [Installing directly](#).
- Use Stack Builder (with PostgreSQL) or StackBuilder Plus (with EDB Postgres Advanced Server) to download the EDB installer package and invoke the graphical installer. See [Using Stack Builder or StackBuilder Plus](#).

Prerequisites

You must have a Java JVM (version 1.8.0 or later of Red Hat OpenJDK or Oracle JDK) in place before Stack Builder can perform a Migration Toolkit installation.

The Java executable must be in your search path (`%PATH%` on Windows). On Windows, use the following command to set the search path, substituting the name of the directory that holds the Java executable for `javadir`:

```
SET PATH=javadir;%PATH%
```

Note

The installation process creates a configuration file at this location: `C:\Program Files\edb\mtk\etc\sysconfig\edbmtk-<version>.config`. This configuration file specifies the location of the Java executable. If you have multiple versions of Java installed on your system, you can modify the contents of `edbmtk-<version>.config` to ensure that Migration Toolkit is using the correct version.

Installing directly

After downloading the graphical installer, to start the installation wizard, assume sufficient privileges (superuser or administrator) and double-click the installer icon. If prompted, provide a password.

In some versions of Windows, to invoke the installer with Administrator privileges, you need to right-click on the installer icon and select **Run as Administrator** from the context menu.

Proceed to the [Using the graphical installer](#) section in this topic.

Using Stack Builder or StackBuilder Plus

If you are using PostgreSQL, you can invoke the graphical installer with Stack Builder. See [Using Stack Builder](#).

1. In Stack Builder, follow the prompts until you get to the module selection page.

On the Welcome page, use the drop-down listbox to select the target server installation from the list of available servers. If your network requires you to use a proxy server to access the internet, select `Proxy servers` and specify a server; if you do not need to use a proxy server, click `Next` to proceed.

2. Expand the **Registration-required and trial products** node.

3. Expand the **EnterpriseDB Tools** node and select **Migration Toolkit**.

4. Proceed to [Using the graphical installer](#).

If you are using EDB Postgres Advanced Server, you can invoke the graphical installer with StackBuilder Plus. See [Using StackBuilder Plus](#).

1. In StackBuilder Plus, follow the prompts until you get to the module selection page.

On the Welcome page, use the drop-down listbox to select the target server installation from the list of available servers. If your network requires you to use a proxy server to access the internet, select **Proxy servers** and specify a server; if you do not need to use a proxy server, click **Next** to proceed.

2. Expand the **Add-ons, tools, and utilities** node and select **Enterprise DB Migration Toolkit**.

3. Proceed to [Using the graphical installer](#).

Using the graphical installer

1. Select the installation language and select **OK**.

2. On the Setup Migration Toolkit page, select **Next**.

3. Read the license agreement. If you accept the agreement, select the **accept the agreement** option and select **Next**.

4. Browse to a directory where you want Migration Toolkit to be installed, or allow the installer to install Migration Toolkit in the default location. Select **Next**.

5. On the Ready to Install page, select **Next**.

An information box shows the installation progress of the selected components. This may take a few minutes.

6. When the installation has completed, select **Finish**.

After installing Migration Toolkit, you must install the appropriate JDBC drivers before performing a migration. See [Installing a JDBC driver](#) for more information.

5.5 Installing a JDBC driver

Migration Toolkit requires Java version 1.8.0 or later.

Choosing a driver

The JDBC drivers you use depend on the source and target databases for the migration. For example, to migrate from an Oracle Database instance to EDB Postgres Advanced Server, you need the Oracle Database JDBC driver and the EDB JDBC driver.

These are reference sources where JDBC drivers are freely available for download. However, there are other sources where you can obtain the same drivers.

Source or target database	JDBC driver
PostgreSQL	PostgreSQL JDBC driver
EDB Postgres Advanced Server	EDB JDBC connector
Oracle Database	Oracle Database JDBC driver
Microsoft SQL Server	Microsoft JDBC driver for SQL Server
MySQL	Connector/J JDBC driver
SAP Adaptive Server Enterprise (ASE)	Sybase ASE (jTDS) JDBC driver (See note)

Note

The open-source jTDS driver also supports older versions of Microsoft SQL Server (version 2012 and earlier) and may be used with Migration Toolkit. For the newer versions of SQL Server supported by Migration Toolkit, we recommend the Microsoft-provided JDBC driver for SQL Server.

Driver version compatibility

After selecting the appropriate drivers for your migration, verify that their versions are compatible with the Migration Toolkit you're using. We recommend using the latest available version of the Migration Toolkit to ensure broader compatibility with various database versions and JDBC drivers.

- If you're migrating from Oracle, MySQL, Microsoft SQL Server, or Sybase, use the freely available, source-specific JDBC driver:
 - [Oracle JDBC](#)
 - [MySQL JDBC](#)
 - [Microsoft SQL Server JDBC](#)
 - [Sybase ASE \(jTDS\) JDBC](#)

Adding the driver to `lib`

After downloading the driver, move the driver file into the `<mtk_install_dir>/lib` directory.

For the EDB JDBC driver

If the EDB JDBC driver is installed prior to installing Migration Toolkit you don't need to place a copy of it in the `<mtk_install_dir>/lib` directory. A symbolic link to the `edb-jdbc18.jar` file is created during installation of Migration Toolkit.

After installing Migration Toolkit, you can verify whether the symbolic link exists by executing the following command:

```
ls -l /usr/edb/migrationtoolkit/lib/edb-jdbc18.jar
```

6 Upgrading

You can upgrade an existing Linux installation by first upgrading the `edb.repo` file, which enables access to the current EnterpriseDB repository, and then upgrading Migration Toolkit.

To upgrade the `edb.repo` file:

```
# Update your repository configuration file
sudo <package-manager> -y upgrade edb-repo

# Upgrade the installed product
sudo <package-manager> -y edb-migrationtoolkit
```

Where `<package-manager>` is the package manager used with your operating system:

Package manager	Operating system
dnf	RHEL 8/9 and derivatives
zypper	SLES
apt-get	Debian and Ubuntu


7 Specifying connection properties

Migration Toolkit uses the configuration and connection information stored in the `toolkit.properties` file during the migration process to identify and connect to the source and target databases.

On Linux, the `toolkit.properties` file is located in `/usr/edb/migrationtoolkit/etc`.

On Windows, the file is located in `C:\Program Files\edb\mtk\etc`.

The figure shows a sample `toolkit.properties` file:

A screenshot of a text editor window titled '*toolkit.properties' with the path '/usr/edb/migrationtoolkit/etc' displayed below the title. The window has a menu bar with 'Open', 'Save', and a hamburger menu icon. The main text area contains the following properties:

```
1 SRC_DB_URL=jdbc:oracle:thin:@localhost:1521:mgmt
2 SRC_DB_USER=hr
3 SRC_DB_PASSWORD=hr
4
5 TARGET_DB_URL=jdbc:edb://localhost:5444/edb
6 TARGET_DB_USER=hrm
7 TARGET_DB_PASSWORD=hrm|
```

The status bar at the bottom shows 'Plain Text', 'Tab Width: 8', 'Ln 7, Col 23', and 'INS'.

Before executing Migration Toolkit commands, modify the `toolkit.properties` file with the editor of your choice. Update the file to include the following information:

- `SRC_DB_URL` specifies how Migration Toolkit connects to the source database. Details on forming the URL for each source database follows.
- `SRC_DB_USER` specifies a user name (with sufficient privileges) in the source database.
- `SRC_DB_PASSWORD` specifies the password of the source database user. Provide the password without any surrounding quotes.
- `TARGET_DB_URL` specifies the JDBC URL of the target database.
- `TARGET_DB_USER` specifies the name of a privileged target database user.
- `TARGET_DB_PASSWORD` specifies the password of the target database user. Provide the password without any surrounding quotes.

Note

Unless specified in the command line, Migration Toolkit expects the source database to be Oracle and the target database to be EDB Postgres Advanced Server. For any other source or target database, specify the `-sourcedbtype` or `-targetdbtype` options as described in [Migrating from a non-Oracle source database](#).

For specifying a target database on Cloud Service, see [Defining the URL and configuring the required parameters for Cloud Service](#).

URL and parameters for an Advanced Server

Migration Toolkit helps with migration from the following platforms to Advanced Server:

- Oracle
- PostgreSQL
- MySQL
- Sybase
- SQL Server

For a definitive list of the objects migrated from each database type, refer to [Functionality overview](#).

Migration Toolkit reads connection specifications for the source and the target database from the `toolkit.properties` file. Connection information for each must include:

- The URL of the database
- The name of a privileged user
- The password associated with the specified user

The URL conforms to JDBC standards and takes the form:

```
{TARGET_DB_URL | SRC_DB_URL}=jdbc:edb://<host>:<port><database_id>
```

An Advanced Server URL contains the following information:

- `jdbc` — The protocol is always `jdbc`.
- `edb` — If you're using Advanced Server, specify `edb` for the subprotocol value.
- `<host>` — The name or IP address of the host where the Postgres instance is running.
- `<port>` — The port number that the Advanced Server database listener is monitoring. The default port number is 5444.
- `<database_id>` — The name of the source or target database.

Ensure these values are set in the `toolkit.properties` file:

- `{TARGET_DB_USER | SRC_DB_USER}` — Specifies a user with privileges to create each type of object migrated. If migrating data into a table, the specified user might also require insert, truncate, and references privileges for each target table.
- `{TARGET_DB_PASSWORD | SRC_DB_PASSWORD}` — Set to the password of the privileged Advanced Server user. Provide the password without any surrounding quotes.

URL and parameters for PostgreSQL

Migration Toolkit helps with migration from the following platforms to PostgreSQL:

- Oracle
- Advanced Server
- MySQL
- SQL Server

For a definitive list of the objects migrated from each database type, refer to [Functionality overview](#).

Migration Toolkit reads connection specifications for the source and the target database from the `toolkit.properties` file. Connection information for each must include:

- The URL of the database
- The name of a privileged user
- The password associated with the specified user

A PostgreSQL URL conforms to JDBC standards and takes the form:

```
{SRC_DB_URL|TARGET_DB_URL}=jdbc:postgresql://<host>:<port>/<database_id>
```

The URL contains the following information:

- `jdbc` — The protocol is always `jdbc`.
- `postgresql` — If you're using PostgreSQL, specify `postgresql` for the subprotocol value.
- `<host>` — The name or IP address of the host where the Postgres instance is running.
- `<port>` — The port number that the Postgres database listener is monitoring. The default port number is 5432.
- `<database_id>` — The name of the source or target database.

Ensure these values are set in the `toolkit.properties` file:

- `{SRC_DB_USER|TARGET_DB_USER}` — Specify a user with privileges to create each type of object migrated. If migrating data into a table, the specified user might also need insert, truncate, and references privileges for each target table.
- `{SRC_DB_PASSWORD|TARGET_DB_PASSWORD}` — Set to the password of the privileged PostgreSQL user. Provide the password without any surrounding quotes.

URL and parameters for Cloud Service

Migration Toolkit helps with migration from the following platforms to Cloud Service:

- Oracle
- Advanced Server
- PostgreSQL
- MySQL
- SQL Server

Migration Toolkit reads connection specifications for the source and the target database from the `toolkit.properties` file. Connection information for each must include:

- The URL of the database
- The name of a privileged user
- The password associated with the specified user

When migrating to Cloud Service, `TARGET_DB_URL` takes the form of a JDBC URL. For example:

```
jdbc:<postgres_type>://<host_name>[:<port>]/<database_id>?sslmode=<sslmode>
```

Note

Many of the values you need for the target database URL are available from the Cloud Service portal. In Cloud Service, select your cluster and go to the **Connect** tab to find the values.

The URL contains the following information:

- `jdbc` — The protocol is always `jdbc`.
- `postgres_type` — The subprotocol is the Postgres type. Specify `edb` if you're using Advanced Server or `postgresql` if you're using PostgreSQL.
- `<host_name>` — The host name of your cluster. You can copy it from the **Host** field on the **Connect** tab in the Cloud Service portal.
- `<port>` — The port number that the database listener is monitoring. You can copy it from the **Port** field on the **Connect** tab in the Cloud Service portal.
- `<database_id>` — The name of the target database. Set this to the name of the database in your cluster that you want to use as your migration target database. The name of the default database for your cluster is shown in the **Dbname** field on the **Connect** tab in the Cloud Service portal. Often a separate database is created for use as the migration target.
- `sslmode` — Either "require" or "verify-full". See [Recommended settings for SSL mode](#). Listed at the end of the **Service URI** value on the **Connect** tab in the Cloud Service portal.

Ensure these values are set in the `toolkit.properties` file:

- `TARGET_DB_USER` — Specifies the name of a privileged database user. You can copy it from the **User** field on the **Connect** tab in the Cloud Service portal.
- `TARGET_DB_PASSWORD` — Contains the password of the specified user. Provide the password without any surrounding quotes.

URL and parameters for Oracle

Migration Toolkit helps with migration from an Oracle database to a PostgreSQL or Advanced Server database. When migrating from Oracle, you must specify connection specifications for the Oracle source database in the `toolkit.properties` file. The connection information must include:

- The URL of the Oracle database
- The name of a privileged user
- The password associated with the specified user

When migrating from an Oracle database, `SRC_DB_URL` must contain a JDBC URL, specified in one of two forms. The first form is:

```
jdbc:oracle:thin:@<host_name>:<port>:<database_id>
```

The second form is:

```
jdbc:oracle:thin:@//<host_name>:<port>{<database_id|service_name>}
```

An Oracle URL contains the following information:

- `jdbc` — The protocol is always `jdbc`.
- `oracle` — The subprotocol is always `oracle`.
- `thin` — The driver type. Specify a driver type of `thin`.
- `<host_name>` — The name or IP address of the host where the Oracle server is running.
- `<port>` — The port number that the Oracle database listener is monitoring.
- `<database_id>` — The database SID of the Oracle database.
- `<service_name>` — The name of the Oracle service.

Ensure these values are set in the `toolkit.properties` file:

- `SRC_DB_USER` — Specifies the name of a privileged Oracle user. The Oracle user needs read access to the source database objects you want to migrate. If you want to migrate users/roles and related profiles/privileges, grant the Oracle user SELECT privileges on the following Oracle catalog objects:
 - `DBA_ROLES`
 - `DBA_USERS`
 - `DBA_TAB_PRIVS`
 - `DBA_PROFILES`
 - `DBA_ROLE_PRIVS`
 - `ROLE_ROLE_PRIVS`
 - `DBA_SYS_PRIVS`
- `SRC_DB_PASSWORD` — Contains the password of the specified user. Provide the password without any surrounding quotes.

URL and parameters for an SQL Server

Migration Toolkit helps with migration from an SQL Server database to a PostgreSQL or Advanced Server database. Migration Toolkit supports migration of the following object definitions:

- Schemas
- Tables
- Table data
- Constraints
- Indexes

Migration Toolkit reads connection specifications for the source database from the `toolkit.properties` file. The connection information must include:

- The URL of the source database
- The name of a privileged user
- The password associated with the specified user

If you're connecting to a SQL Server database, `SRC_DB_URL` takes the form of a [Microsoft JDBC URL](#) (recommended) or a [JTDS URL](#).

Microsoft JDBC URL

By default, Microsoft JDBC Driver for SQL Server uses TLS encryption for all communication between the client and the SQL Server. Set `encrypt` to `false`, if you want to disable TLS encryption. For more information about connecting to a Microsoft SQL Server using JDBC type 4 driver, see [Building the connection URL](#).

```
jdbc:sqlserver://<server_name>:<port>[;databaseName=<database_id>][;encrypt=false]
```

A SQL server URL contains the following information:

- `jdbc` — The protocol is always `jdbc`.
- `sqlserver` — The server type is always `sqlserver`.
- `<server_name>` — The name or IP address of the host where the source server is running.
- `<port>` — The port number that the source database listener is monitoring.
- `<database_id>` — The name of the source database.

Ensure these values are set in the `toolkit.properties` file:

- `SRC_DB_USER` — Specifies the name of a privileged SQL Server user.
- `SRC_DB_PASSWORD` — Contains the password of the specified user. Provide the password without any surrounding quotes.

JTDS URL

```
jdbc:jtds:sqlserver://<server_name>:<port>/<database_id>
```

Tip

The JTDS driver is an open source JDBC 3.0 type 4 driver that supports older versions of Microsoft SQL Server. See <http://jtds.sourceforge.net/>. When connecting newer versions of Microsoft SQL Server with Migration Toolkit, the Microsoft JDBC Driver for SQL Server is recommended.

A SQL server URL contains the following information:

- `jdbc` — The protocol is always `jdbc`.
- `jtds` — The driver name is always `jtds`.
- `sqlserver` — The server type is always `sqlserver`.

- `<server_name>` — The name or IP address of the host where the source server is running.
- `<port>` — The port number that the source database listener is monitoring.
- `<database_id>` — The name of the source database.

Ensure these values are set in the `toolkit.properties` file:

- `SRC_DB_USER` — Specifies the name of a privileged SQL Server user.
- `SRC_DB_PASSWORD` — Contains the password of the specified user. Provide the password without any surrounding quotes.

URL and parameters for MySQL

Migration Toolkit helps with migration from a MySQL database to an Advanced Server or PostgreSQL database. When migrating from MySQL, you must specify connection specifications for the MySQL source database in the `toolkit.properties` file. The connection information must include:

- The URL of the source database
- The name of a privileged user
- The password associated with the specified user

When migrating from MySQL, `SRC_DB_URL` takes the form of a JDBC URL. For example:

```
jdbc:mysql://<host_name>[:<port>]/<database_id>
```

The URL contains the following information:

- `jdbc` — The protocol is always `jdbc`.
- `mysql` — The subprotocol is always `mysql`.
- `<host_name>` — The name or IP address of the host where the source server is running.
- `[<port>]` — The port number that the MySQL database listener is monitoring.
- `<database_id>` — The name of the source database.

Ensure these values are set in the `toolkit.properties` file:

- `SRC_DB_USER` — Specifies the name of a privileged MySQL user.
- `SRC_DB_PASSWORD` — Contains the password of the specified user. Provide the password without any surrounding quotes.

Note

- If datatype `tinyInt(1)` is used to store byte values other than 0 and 1 in the MySQL source database, make sure to append the optional parameter `?tinyInt1isBit=false` in the MySQL Connector/J JDBC Driver URL.
- Due to a bug in the MySQL Connector/J JDBC Driver 8.0.26, the migration of foreign key constraints fails in certain cases. EDB recommends that you don't use this driver for migrating data using Migration Toolkit. Instead, use MySQL Connector/J JDBC Driver 8.0.30 or later.

For detailed information about this bug, see the [MySQL bug report](#).

`TINYINT(1)` is mapped to `BIT(1)` in PostgreSQL/EDB Postgres Advanced Server (which might not be expected in some cases). But as the MySQL JDBC driver reports it as `BIT(1)`, Migration Toolkit maps it to `BIT(1)` in PostgreSQL/EDB Postgres Advanced Server.

Unlike `TINYINT(1)` that can hold the values -128 to 127, `BIT(1)` can hold only 0 and 1. So if you're storing any value other than 0 and 1 in this field, you might need to change the default behavior of the MySQL JDBC driver so that the driver reports it as `TINYINT`.

To change this behavior, set the following URL option in the connection string:

```
jdbc:mysql://localhost:3306/world?tinyInt1isBit=false
```

In this case, the JDBC driver reports `TINYINT(1)` as `TINYINT` and is mapped to `SMALLINT` in the target PostgreSQL/EDB Postgres Advanced Server.

URL and parameters for Sybase

Migration Toolkit helps with migration from a Sybase database to an Advanced Server database. When migrating from Sybase, you must specify connection specifications for the Sybase source database in the `toolkit.properties` file. The connection information must include:

- The URL of the source database
- The name of a privileged user
- The password associated with the specified user

When migrating from Sybase, `SRC_DB_URL` takes the form of a JTDS URL. For example:

```
jdbc:jtds:sybase://<host_name>[:<port>]/<database_id>
```

Tip

For an open source JDBC 3.0 type 4 driver for Sybase ASE and older versions of Microsoft SQL Server, see <http://jtds.sourceforge.net/>.

A Sybase URL contains the following information:

- `jdbc` — The protocol is always `jdbc`.
- `jtds` — The driver name is always `jtds`.
- `sybase` — The server type is always `sybase`.
- `<host_name>` — The name or IP address of the host where the source server is running.

- `<port>` – The port number that the Sybase database listener is monitoring.
- `<database_id>` – The name of the source database.

Ensure these values are set in the `toolkit.properties` file:

- `SRC_DB_USER` – Specifies the name of a privileged Sybase user.
- `SRC_DB_PASSWORD` – Contains the password of the specified user. Provide the password without any surrounding quotes.

8 Invoking Migration Toolkit

After installing Migration Toolkit and specifying connection properties for the source and target databases in the [toolkit.properties](#) file, Migration Toolkit is ready to perform migrations.

Note

Ensure the operating system user account running the Migration Toolkit is the owner of the `toolkit.properties` file with a minimum of read permission on the file. See [Invalid file permissions](#) for more information.

The Migration Toolkit executable is named `runMTK.sh` on Linux systems and `runMTK.bat` on Windows systems. On a Linux system, the executable is located in:

```
/usr/edb/migrationtoolkit/bin
```

On Windows, the executable is located in:

```
C:\Program Files\edb\mtk\bin
```

See [Migration Toolkit command options](#) for information on controlling details of the migration.

Importing character data with embedded binary zeros (NULL characters)

Migration Toolkit supports importing a column with a value of NULL. However, Migration Toolkit doesn't support importing NULL character values (embedded binary zeros 0x00) with the JDBC connection protocol. If you're importing data that includes the NULL character, use the `-replaceNullChar` option to replace the NULL character with a single, non-NULL, replacement character.

Note

- MTK implicitly replaces NULL characters with an empty string.
- The `-replaceNullChar` option doesn't work with the `-copyViaDBLinkOra` option.

Once the data is migrated, use a SQL statement to replace the character specified by `-replaceNullChar` with binary zeros.

Migrating schemas

When migrating schemas from the source database specified in the [toolkit.properties](#) file, you have to provide the `<schema_scope>` for the migration. You have these options:

- To migrate a single schema, specify the schema name.
- To migrate multiple schemas, provide a comma-delimited list of schema names.
- To migrate all the schemas, use the `-allSchemas` option.

These examples display how to execute the command depending on your operating system and schema scope.

	On Linux	On Windows
Single schema	<pre>\$./runMTK.sh <schema_name></pre>	<pre>> .\runMTK.bat <schema_name></pre>

	On Linux	On Windows
Multiple schemas	<code>\$./runMTK.sh <schema_name1>, <schema_name2>, <schema_name3></code>	<code>> .\runMTK.bat <schema_name1>, <schema_name2>, <schema_name3></code>
All schemas	<code>\$./runMTK.sh -allSchemas</code>	<code>> .\runMTK.bat -allSchemas</code>

Migrating schemas from Oracle

Unless specified in the command line, Migration Toolkit expects the source database to be Oracle and the target database to be EDB Postgres Advanced Server. To migrate schemas, navigate to the executable and invoke the following command.

On Linux:

```
$ ./runMTK.sh <schema_scope>
```

On Windows:

```
> .\runMTK.bat <schema_scope>
```

Where `<schema_scope>` is the [scope of schemas](#) to migrate which can be a single schema, multiple schemas, or all schemas from the specified database.

Single schema example

The following example migrates a schema (table definitions and table content) named `HR` from an Oracle database to an EDB Postgres Advanced Server database.

On Linux:

```
$ ./runMTK.sh HR
```

On Windows:

```
> .\runMTK.bat HR
```

Multiple schemas example

The following example migrates multiple schemas (named HR and ACCTG) from an Oracle database to a PostgreSQL database.

On Linux:

```
$ ./runMTK.sh -targetdbtype postgres HR,ACCTG
```

On Windows:

```
> .\runMTK.bat -targetdbtype postgres HR,ACCTG
```

All schemas example

The following example migrates all schemas from the Oracle database specified in the [toolkit.properties](#) file to an EDB Postgres Advanced Server database.

On Linux:

```
$ ./runMTK.sh -allSchemas
```

On Windows:

```
> .\runMTK.bat -allSchemas
```

Notes

- When the default database user of a migrated schema is automatically migrated, the custom profile of the default database user is also migrated if a custom profile exists. A custom profile is a user-created profile. For example, custom profiles exclude Oracle profiles `DEFAULT` and `MONITORING_PROFILE`.
- PostgreSQL default rows are limited to 8 KB in size. This means that each table row must fit into a single 8 KB block. Otherwise, an error occurs indicating, for example, that we can create a table with 1600 columns of INT and insert data for all the columns. However, we can't do the same with BIGINT columns because INT is stored as 4 bytes, but each BIGINT requires more space (8 bytes). For more information, see [PostgreSQL Limits](#) in the PostgreSQL documentation.

Migrating schemas from a non-Oracle source database

If you don't specify a source database type and a target database type, Postgres assumes the source database is Oracle and the target database is EDB Postgres Advanced Server.

To invoke Migration Toolkit, open a command window, navigate to the executable, and invoke the following command:

```
$ ./runMTK.sh -sourcedbtype <source_type> -targetdbtype <target_type> [options, ...] <schema_scope>;
```

`-sourcedbtype <source_type>`

`<source_type>` specifies the server type of the source database. `<source_type>` isn't case sensitive. By default, `<source_type>` is oracle. `source_type` can be one of the following values:

To migrate from:	Specify:
Oracle	oracle (the default value)
MySQL	mysql
SQL Server	sqlserver
Sybase	sybase
PostgreSQL	postgres or postgresql
EDB Postgres Advanced Server	enterprisedb

`-targetdbtype <target_type>`

`<target_type>` specifies the server type of the target database. `<target_type>` isn't case sensitive. By default, `<target_type>` is `enterprisedb`. `<target_type>` can be one of the following values:

To migrate to:	Specify:
EDB Postgres Advanced Server	enterprisedb
PostgreSQL	postgres or postgresql

<schema_scope>

Where <schema_scope> is the [scope of schemas](#) to migrate which can be a single schema, multiple schemas, or all schemas from the specified database.

Single schema example

The following example migrates a schema (table definitions and table content) named **HR** from a MySQL database on a Linux system to an EDB Postgres Advanced Server host. The command includes the `-sourcedbtype` and `-targetdbtype` options.

On Linux:

```
$ ./runMTK.sh -sourcedbtype mysql -targetdbtype enterprisedb HR
```

On Windows:

```
> .\runMTK.bat -sourcedbtype mysql -targetdbtype enterprisedb HR
```

Multiple schemas example

The following example migrates multiple schemas (named HR and ACCTG) from a MySQL database to a PostgreSQL database.

On Linux:

```
$ ./runMTK.sh -sourcedbtype mysql -targetdbtype postgres HR,ACCTG
```

On Windows:

```
> .\runMTK.bat -sourcedbtype mysql -targetdbtype postgres HR,ACCTG
```

All schemas example

The following example migrates all schemas from a PostgreSQL database specified in the [toolkit.properties file](#) to an EDB Postgres Advanced Server database.

Note

The `-allSchemas` parameter is supported only for Oracle, EDB Postgres Advanced Server, and PostgreSQL source databases.

On Linux:

```
$ ./runMTK.sh -sourcedbtype postgres -allSchemas
```

On Windows:


```
> .\runMTK.bat -sourcedbtype postgres -allSchemas
```

8.1 Migration Toolkit command options

To control details of the migration, append migration options when you run Migration Toolkit. For example, to migrate all schemas in a database, append the `-allSchemas` option to the command:

```
./runMTK.sh -allSchemas
```

Note

- The `-allSchemas` parameter is supported only for the Oracle, EDB Postgres Advanced Server, and PostgreSQL source database. It isn't supported for Sybase, MS SQL Server, and MySQL source databases.
- Migration Toolkit disables the user-created triggers while migrating the data. However, as PostgreSQL/EDB Postgres Advanced Server doesn't allow disabling inherited child partition triggers using the `DISABLE TRIGGER USER` clause in `ALTER TABLE` , the triggers activate during data migration on partition tables, which can cause unexpected results.

The command options that work with Migration Toolkit are grouped by their behavior, as shown in the table.

Feature	Relevant options
Offline migration options	-offlineMigration
Import options	-sourcedbtype, -targetdbtype, -schemaOnly, -dataOnly
Schema creation options	-dropSchema, -targetSchema
	-allTables, -tables, -excludeTables,
	-constraints, -ignoreCheckConstFilter,
	-skipCKConst, -skipFKConst,
	-skipColDefaultClause,
	-indexes, -triggers,
	-allViews, -views, -excludeViews,
	-allSequences, -sequences, -excludeSequences,
Schema object selection options	-allProcs, -procs, -excludeProcs,
	-allFuncs, -funcs, -excludeFuncs,
	-checkFunctionBodies,
	-allPackages, -packages, -excludePackages,
	-allDomains, -domains,
	-allQueues, -queues, -excludeQueues,
	-allRules,
	-allgroups, -groups

Feature	Relevant options
Migration options	-truncLoad, -enableConstBeforeDataLoad,
	-retryCount, -safeMode, -fastCopy,
	-analyze, vacuumAnalyze, -replaceNullChar,
	-copyDelimiter, -batchSize,
	-cpBatchSize, -lobBatchSize,
	-fetchSize, -filterProp,
	-customColTypeMapping, -customColTypeMappingFile,
	-mapVirtualToRegular
Connection retry options	-connRetryCount, -connRetryInterval, -abortOnConnFailure
Oracle-specific options	-allUsers, -users,
	-allProfiles, -profiles,
	-importPartitionAsTable,
	-objectTypes,
	-copyViaDBLinkOra, -allDBLinks,
	-allSynonyms, -allPublicSynonyms, -excludeSynonyms,
	-allPrivateSynonyms, -useOraCase,
	-skipUserSchemaCreation
Miscellaneous options	-help, -logDir, -logFileCount, -logFileSize, -logBadSQL -verbose, -version
Migration options for parallel data loading	-loaderCount, -parallelLoadRowLimit, -tableLoaderLimit
Specify options in file	-optionsFile

Note

If you're using several options for a migration or have to specify a long list of objects for an option, consider using the `-optionsFile` option to specify the values in a separate text file. See [Specifying options using a file](#) for more information.

Offline migration options

If you specify the `-offlineMigration` option in the command line, Migration Toolkit performs an offline migration. During an offline migration, Migration Toolkit reads the definition of each selected object and creates an SQL script that, when executed later, replicates each object in Postgres.

Note

The following examples invoke Migration Toolkit in Linux. To invoke Migration Toolkit in Windows, use the `runMTK.bat` command instead of the `runMTK.sh` command.

To perform an offline migration of both schema and data, specify the `-offlineMigration` keyword, followed by the [schema scope](#):

```
./runMTK.sh -offlineMigration <schema_scope>
```

Each database object definition is saved in a separate file with a name derived from each schema name and object type in your home folder. To specify an alternative file destination, include a directory name after the `-offlineMigration` option:

```
./runMTK.sh -offlineMigration <file_dest> <schema_scope>
```

To perform an offline migration of only schema objects (creating empty tables), specify the `-schemaOnly` keyword in addition to the `-offlineMigration` keyword when invoking Migration Toolkit:

```
./runMTK.sh -offlineMigration -schemaOnly <schema_scope>
```

To perform an offline migration of only data, omitting any schema object definitions, specify the `-dataOnly` keyword and the `-offlineMigration` keyword when invoking Migration Toolkit:

```
./runMTK.sh -offlineMigration -dataOnly <schema_scope>
```

By default, data is written in COPY format. To write the data in a plain SQL format, include the `-safeMode` keyword:

```
./runMTK.sh -offlineMigration -dataOnly -safeMode <schema_scope>
```

By default, when you perform an offline migration that contains table data, a separate file is created for each table. To create a single file that contains the data from multiple tables, specify the `-singleDataFile` keyword:

```
./runMTK.sh -offlineMigration -dataOnly -singleDataFile -safeMode <schema_scope>
```

Note

The `-singleDataFile` option is available only when migrating data in a plain SQL format. You must include the `-safeMode` keyword if you include the `-singleDataFile` option.

Executing offline migration scripts

You can use the `edb-psql` or `psql` command line to execute the scripts generated during an offline migration. The following example describes restoring a schema (named `hr`) into a new database (named `acctg`) stored in EDB Postgres Advanced Server.

1. Use the `createdb` command to create the `acctg` database, into which you'll restore the migrated database objects:

```
createdb -U enterprisedb acctg
```

2. Connect to the new database with `edb-psql`:

```
edb-psql -U enterprisedb acctg
```

3. Use the `\i` meta-command to invoke the migration script that creates the object definitions:

```
acctg=# \i ./mtk_hr_ddl.sql
```

4. If the `-offlineMigration` command included the `-singleDataFile` keyword, the `mtk_hr_data.sql` script will contain the commands required to re-create all of the objects in the new target database. Populate the database with the command:

```
acctg=# \i ./mtk_hr_data.sql
```

Import options

By default, Migration Toolkit assumes the source database is Oracle and the target database is EDB Postgres Advanced Server. Include the `-sourcedbtype` and `-targetdbtype` keywords to specify a nondefault source or target database.

By default, Migration Toolkit imports both the data and the object definition when migrating a schema. Alternatively, you can choose to import either the data or the object definitions.

```
-sourcedbtype <source_type>
```

The `-sourcedbtype` option specifies the source database type. For `source_type`, use one of the following values: `mysql`, `oracle`, `sqlserver`, `sybase`, `postgresql` or `enterprisedb`. `source_type` isn't case sensitive. By default, `source_type` is `oracle`.

```
-targetdbtype <target_type>
```

The `-targetdbtype` option specifies the target database type. For `target_type`, use one of the following values: `enterprisedb`, `postgres`, or `postgresql`. `target_type` isn't case sensitive. By default, `target_type` is `enterprisedb`.

```
-schemaOnly
```

This option imports the schema definition and creates all selected schema objects in the target database. You can't use this option with the `-dataOnly` option.

```
-dataOnly
```

This option copies only the data. When used with the `-tables` option, Migration Toolkit imports data only for the selected tables. You can't use this option with the `-schemaOnly` option.

Schema creation options

By default, Migration Toolkit imports the source schema objects or data into a schema of the same name. If the target schema doesn't exist, Migration Toolkit creates a schema. Alternatively, you can specify a custom schema name by using the `-targetSchema` option. You can choose to drop the existing schema and create a new schema using the following option:

```
-dropSchema [true|false]
```

With this option set to `true`, Migration Toolkit drops the existing schema and any objects in that schema and creates a new schema. By default, `-dropSchema` is `false`.

```
-targetSchema <schema_name>
```

Use the `-targetSchema` option to specify the name of the migrated schema. If you're migrating multiple schemas, specify a name for each schema in a comma-separated list with no intervening space characters. Without the `-targetSchema` option, the name of the new schema is the same as the name of the source schema.

You can't specify `information-schema`, `dbo`, `sys`, or `pg_catalog` as target schema names. These schema names are reserved for metadata storage in EDB Postgres Advanced Server.

Schema object selection options

Use the following options to select specific schema objects to migrate.

```
-allTables
```

Import all tables from the source schema.

```
-tables <table_list>
```

Import the selected tables from the source schema. `table_list` is a comma-separated list of table names with no intervening space characters.

For example: `-tables emp,dept,acctg`

When migrating multiple schemas or all schemas (with the `-allSchemas` option), the table name must be schema-qualified.

For example: `-allSchemas -tables comp_schema.emp,comp_schema.dept,finance_schema.acctg`

```
-excludeTables <table_list>
```

Exclude the selected tables from migration. The `table_list` is a comma-separated list of table names with no intervening space characters. This option applies when all tables from a schema are selected for migration using the `-allTables` option.

For example: `-allTables -excludeTables emp,jobhist`

When migrating multiple schemas or all schemas (with the `-allSchemas` option), the excluded table name must be schema-qualified.

For example: `-allSchemas -allTables -excludeTables comp_schema.emp,finance_schema.jobhist`

```
-constraints
```

Import the table constraints. This option is valid only when importing an entire schema or when you specify the `-allTables` or `-tables <table_list>` options.

-ignoreCheckConstFilter

By default, Migration Toolkit doesn't implement migration of check constraints and default clauses from a Sybase database. To migrate constraints and default clauses from a Sybase database, include the **-ignoreCheckConstFilter** parameter when specifying the **-constraints** parameter.

-skipCKConst

Omit the migration of check constraints. This option is useful when migrating check constraints that are based on built-in functions in the source database that aren't supported in the target database.

This option is valid only when importing an entire schema or when the **-allTables** or **-tables <table_list>** options are specified.

-skipFKConst

Omit migrating foreign-key constraints. This option is valid only when importing an entire schema or when the **-allTables** or **-tables <table_list>** options are specified.

-skipColDefaultClause

Omit migrating the column **DEFAULT** clause.

-indexes

Import the table indexes. This option is valid when importing an entire schema or when the **-allTables** or **-tables <table_list>** option is specified.

-triggers

Import the table triggers. This option is valid when importing an entire schema or when the **-allTables** or **-tables <table_list>** option is specified.

-allViews

Import the views from the source schema. This option migrates dynamic and materialized views from the source. Oracle and Postgres materialized views are supported.

-views <view_list>

Import the specified materialized or dynamic views from the source schema. Oracle and Postgres materialized views are supported. **view_list** is a comma-separated list of view names with no intervening space characters.

For example: **-views all_emp,mgmt_list,acct_list**

When migrating multiple schemas or all schemas (with the **-allSchemas** option), the view name must be schema-qualified.

For example: `-allSchemas -views comp_schema.all_emp,comp_schema.mgmt_list,finance_schema.acct_list`

`-excludeViews <view_list>`

Exclude the selected views from migration. The `view_list` is a comma-separated list of view names with no intervening space characters. This option applies when all views from a schema are selected for migration using the `-allViews` option.

For example: `-allViews -excludeViews all_emp,acct_list`

When migrating multiple schemas or all schemas (with the `-allSchemas` option), the excluded view name must be schema-qualified.

For example: `-allSchemas -allViews -excludeViews comp_schema.all_emp,finance_schema.acct_list`

`-allSequences`

Import all sequences from the source schema.

`-sequences <sequence_list>`

Import the selected sequences from the source schema. `<sequence_list>` is a comma-separated list of sequence names with no intervening space characters.

For example: `-sequences my_sequence,my_sequence_with_increment`

When migrating multiple schemas or all schemas (with the `-allSchemas` option), the sequence name must be schema-qualified.

For example: `-allSchemas -sequences comp_schema.my_sequence,finance_schema.my_sequence_with_increment`

`-excludeSequences <sequence_list>`

Exclude selected sequences from the migration. The `sequence_list` is a comma-separated list of sequence names with no intervening space characters. This option applies when all sequences from a schema are selected for migration using the `-allSequences` option.

For example: `-allSequences -excludeSequences my_sequence,my_sequence_with_increment`

When migrating multiple schemas or all schemas (with the `-allSchemas` option), the excluded sequence name must be schema-qualified.

For example: `-allSchemas -allSequences -excludeSequences
comp_schema.my_sequence,finance_schema.my_sequence_with_increment`

`-allProcs`

Import all stored procedures from the source schema.

-procs <procedures_list>

Import the selected stored procedures from the source schema. **procedures_list** is a comma-separated list of procedure names with no intervening space characters.

For example: **-procs show_notice,show_custom_notice**

When migrating multiple schemas or all schemas (with the **-allSchema** option), the procedure name must be schema-qualified.

For example: **-allSchema -procs comp_schema.show_notice,finance_schema.show_custom_notice**

-excludeProcs <procedures_list>

Exclude selected procedures from migration. The **procedures_list** is a comma-separated list of procedure names with no intervening space characters. This option applies when all procedures from a schema are selected for migration using the **-allProcs** option.

For example: **-allProcs -excludeProcs show_notice,show_custom_notice**

When migrating multiple schemas or all schemas (with the **-allSchemas** option), the excluded procedure name must be schema-qualified.

For example: **-allSchemas -allProcs -excludeProcs comp_schema.show_notice,finance_schema.show_custom_notice**

-allFuncs

Import all functions from the source schema.

-funcs <function_list>

Import the selected functions from the source schema. **function_list** is a comma-separated list of function names with no intervening space characters.

For example: **-funcs calculate_average_salary,add_two_numbers**

When migrating multiple schemas or all schemas (with the **-allSchemas** option), the function name must be schema-qualified.

For example: **-allSchemas -funcs comp_schema.calculate_average_salary,finance_schema.add_two_numbers**

-excludeFuncs <function_list>

Exclude selected functions from migration. The **function_list** is a comma-separated list of function names with no intervening space characters. This option applies when all functions from a schema are selected for migration using the **-allFuncs** option.

For example: **-allFuncs -excludeFuncs calculate_average_salary,add_two_numbers**

When migrating multiple schemas or all schemas (with the **-allSchemas** option), the excluded function name must be schema-qualified.

For example: **-allSchemas -allFuncs -excludeFuncs
comp_schema.calculate_average_salary,finance_schema.add_two_numbers**

```
-checkFunctionBodies [true/false]
```

When `false`, disables validation of the function body during function creation. Disabling this validation avoids errors if the function contains forward references. The default value is `true`.

```
-allPackages
```

Import all packages from the source schema.

```
-packages <package_list>
```

Import the selected packages from the source schema. `package_list` is a comma-separated list of package names with no intervening space characters.

For example: `-packages my_package1, mypackage2`

When migrating multiple schemas or all schemas (with the `-allSchemas` option), the package name must be schema-qualified.

For example: `-allSchemas -packages comp_schema.my_package1,finance_schema.mypackage2`

```
-excludePackages <package_list>
```

Exclude selected packages from migration. The `package_list` is a comma-separated list of package names with no intervening space characters. This option applies when all packages from a schema are selected for migration using the `-allPackages` option.

Example: `-allPackages -excludePackages my_package1, mypackage2`

When migrating multiple schemas or all schemas (with the `-allSchemas` option), the excluded package name must be schema-qualified.

For example: `-allSchemas -allPackages -excludePackages comp_schema.my_package1,finance_schema.mypackage2`

```
-domains <domain_list>
```

Import the selected domain, enumeration, and composite types from the source database. `<domain_list>` is a comma-separated list of domain names with no intervening space characters. This option is valid only when the source database is either PostgreSQL, EDB Postgres Advanced Server, or Microsoft SQL Server and the target database is PostgreSQL or EDB Postgres Advanced Server.

Note

Composite user-defined types or UDTs aren't supported in migrations from Microsoft SQL Server. Only simple UDTs are supported.

Composite types hold multiple fields, a table, or structured data, whereas simple types represent a single value.

-allDomains

Import all domain, enumeration, and composite types from the source database. This option is valid only when the source database is either PostgreSQL, EDB Postgres Advanced Server, or Microsoft SQL Server and the target database is PostgreSQL or EDB Postgres Advanced Server.

-allQueues

Import all queues from the source schema. These are queues created and managed by the DBMS_AQ and DBMS_AQADM built-in packages. When Oracle is the source database, you must also specify the **-objectTypes** option. When EDB Postgres Advanced Server is the source database, you must also specify the **-allDomains** and **-allTables** options. Oracle and EDB Postgres Advanced Server queues are supported.

-queues <queue_list>

Import the selected queues from the source schema. **queue_list** is a comma-separated list of queue names with no intervening space characters. These are queues created and managed by the DBMS_AQ and DBMS_AQADM built-in packages. When Oracle is the source database, you must also specify the **-objectTypes**. When EDB Postgres Advanced Server is the source database, you must also specify **-allDomains** and **-allTables**. Oracle and EDB Postgres Advanced Server queues are supported.

-excludeQueues <queue_list>

Exclude selected queues from migration. The **queue_list** is a comma-separated list of queue names with no intervening space characters. This option applies when all queues from a schema are selected for migration using the **-allQueues** option.

Example: **-allQueues -excludeQueues EMP_QUEUE1,EMP_QUEUE2**

-allRules

Import all rules from the source database. This option is valid only when both the source and target are stored on a Postgres host.

-allGroups

Import all groups from the source database.

-groups <group_list>

The selected groups from the source database. The **<group_list>** is a comma-separated list of group names e.g. **-groups acct_emp,mkt_emp**.

Migration options

Use the migration options to control the details of the migration process.

-truncLoad

Truncate the data from the table before importing new data. Use this option with the **-dataOnly** option.

-enableConstBeforeDataLoad

Include the **-enableConstBeforeDataLoad** option if a nonpartitioned source table is mapped to a partitioned table. This option enables all triggers on the target table, including any triggers that redirect data to individual partitions, before the data migration. **-enableConstBeforeDataLoad** is valid only if you also specify the **-truncLoad** parameter.

-retryCount [<value>]

If you're performing a multiple-schema migration, objects that fail to migrate during the first migration attempt due to cross-schema dependencies might successfully migrate during a later migration. Use the **-retryCount** option to specify the number of attempts for Migration Toolkit to make to migrate an object that failed during an initial migration attempt. Specify a value greater than 0. The default value is 2.

-safeMode

If you include the **-safeMode** option, Migration Toolkit commits each row as migrated. If the migration fails to transfer all records, rows inserted prior to the point of failure remain in the target database.

-fastCopy

Including the **-fastCopy** option specifies for Migration Toolkit to bypass WAL logging to perform the COPY operation in an optimized way. It is disabled by default. If you choose to use the **-fastCopy** option, you might not be able to recover the migrated data in the target database if the migration is interrupted.

-replaceNullChar <value>

The Migration Toolkit supports importing a column with a value of NULL. However, the Migration Toolkit doesn't support importing NULL character values (embedded binary zeros 0x00) with the JDBC connection protocol. If you're importing data that includes the NULL character, use the **-replaceNullChar** option to replace the NULL character with a single, non-NULL replacement character. Don't enclose the replacement character in quotes or apostrophes.

Once the data is migrated, use a SQL statement to replace the character specified by **-replaceNullChar** with binary zeros.

-analyze

Include the **-analyze** option to invoke the Postgres **ANALYZE** operation against a target database. The optimizer consults the statistics collected by the **ANALYZE** operation, using the information to construct efficient query plans.

-vacuumAnalyze

Include the **-vacuumAnalyze** option to invoke both the **VACUUM** and **ANALYZE** operations against a target database. The optimizer consults the statistics collected by the **ANALYZE** operation, using the information to construct efficient query plans. The **VACUUM** operation reclaims any storage space occupied by dead tuples in the target database.

-copyDelimiter

Specify a single character to use as a delimiter in the COPY command when loading table data. The default value is **'\t'** (tab).

-batchSize

Specify the batch size of bulk inserts. Valid values are 1 to 1000. The default batch size is 1000. Reduce the value of **-batchSize** if Out of Memory exceptions occur.

-cpBatchSize

Specify the batch size in MB to use in the COPY command. Any value greater than 0 is valid. The default batch size is 8MB.

-lobBatchSize

Specify the number of rows to load in a batch for LOB data types. The data migration for a table containing a large object type (LOB) column, such as **BYTEA**, **BLOB**, or **CLOB**, is performed one row at a time by default. This is to avoid an out-of-heap-space error in case an individual LOB column holds hundreds of megabytes of data. In case the LOB column average data size is at a lower end, you can customize the LOB batch size by specifying the number of rows in each batch with any value greater than 0.

-fetchSize

Use the **-fetchSize** option to specify the number of rows fetched in a result set. If the designated **-fetchSize** is too large, you might encounter Out of Memory exceptions. Include the **-fetchSize** option to avoid this pitfall when migrating large tables. The default fetch size is specific to the JDBC driver implementation and varies by database.

MySQL users note: By default, the MySQL JDBC driver fetches all of the rows in a table into the Migration Toolkit in a single network round trip. This behavior can easily exceed available memory for large tables. If you encounter an out-of-heap-space error, specify **-fetchSize 1** as a command line argument to force Migration Toolkit to load the table data one row at a time.

-filterProp <file_name>

<file_name> specifies the name of a file that contains constraints in key=value pairs. Each record read from the database is evaluated against the constraints. Those that satisfy the constraints are migrated.

The left side of the pair lists a table name:

- The table name must not be schema qualified.
- Oracle table names aren't case sensitive, but Oracle defaults to uppercase. Postgres, on the other hand, defaults to lowercase. When migrating Oracle data, ensure the case you're using matches the case in Oracle. Otherwise, Migration Toolkit can't match it and ignores the constraint.

The right side specifies a condition that must be true for each row migrated.

For example, this code migrates only those countries with a `COUNTRY_ID` value that isn't equal to `AR` :

```
COUNTRIES=COUNTRY_ID<>'AR'
```

This constraint applies to the COUNTRIES table.

You can also specify conditions for multiple tables. However, the condition for each table must be on a new line in the property file.

Example:

The following entries in the properties file migrate only the relevant data from EMPLOYEES and the DEPARTMENTS tables:

```
EMPLOYEES=(LAST_NAME IN ('Grant','Weiss') AND PHONE_NUMBER LIKE '650%')
```

```
DEPARTMENTS=(DEPARTMENT_ID BETWEEN 10 AND 30)
```

```
-customColTypeMapping <column_list>
```

Use custom type mapping to change the data type of migrated columns. The left side of each pair specifies the columns with a regular expression. The right side of each pair names the data type for that column to assume. You can include multiple pairs in a semi-colon-separated list for `<column_list>` . For example, to map any column whose name ends in `ID` to type `INTEGER` , use the following custom mapping entry:

```
.*ID=INTEGER
```

Custom mapping is applied to all table columns that match the criteria unless the column is table qualified.

The `'\\'` characters act as an escape string. Since `'.'` is a reserved character in regular expressions, on Linux use `'\\.'` to represent the `'.'` character. For example, to use custom mapping to select rows from the `EMP_ID` column in the `EMP` table, specify the following custom mapping entry:

```
EMP\\..EMP_ID=INTEGER
```

On Windows, use `'\.'` to represent the `'.'` character:

```
EMP\..EMP_ID=INTEGER
```

Alternatively, you can include multiple custom type mappings in a `<property_file>` :

```
-customColTypeMappingFile <property_file>
```

Specify each entry in the file on a separate line in a key=value pair. The left side of each pair selects the columns with a regular expression. The right side of each pair names the data type for that column to assume. When used in the `<property_file>` , the `'\\'` characters act as an escape string in any operating system.

```
-mapVirtualToRegular
```

Include the `-mapVirtualToRegular` option to migrate virtual, computed, or generated columns in the source database to regular columns in the target database.

Turning virtual columns to regular columns is helpful when the expression used to compute the values in the source database column isn't supported in the target database or can't be translated into an expression supported by the target database. With this option, you can still migrate the column's values to the target, but the values will no longer be computed dynamically, as they're stored in a regular column.

For schema migrations, including the option results in all virtual columns in the source database being migrated to regular columns in the target database.

For data migrations, including the option extracts the computed values from all computed columns in the source database and inserts those values into regular columns in the target database.

In the following example, a column `FullName` is computed from other two columns called `FirstName` and `LastName` in the source database. This column is migrated as a regular column called `FullName` in the target database.

```
./runMTK.sh -sourcedbtype sqlserver -targetdbtype enterprisedb -mapVirtualToRegular -allTables
<schema_scope>
```

In the target, this column is no longer dynamically generated but displays the `FullName` value the source database had at the time the migration was performed.

Connection retry options

Whenever there's a connection failure with the target database during a data migration, Migration Toolkit attempts to reconnect to the target database to ensure the migration completes without skipping any tables. When the connection is reestablished, Migration Toolkit restarts or resumes the data copy for the table that was being migrated when the connection was lost. It then performs the data copy for the remaining tables.

Migration Toolkit resumes the migration based on the mode that it was using to migrate the data:

- If Migration Toolkit was using Copy API to migrate the data, it continues copying data from the last failed row. Copy API is compatible with PostgreSQL and EDB Postgres Advanced Server as target databases.
- In other data migration modes, Migration Toolkit truncates the data on the target table and recopies the entire failed table.

Scope and limitations

Database scope: The connection retry capability allows Migration Toolkit to reconnect to the target database. Retry attempts for issues with the source database aren't supported.

Migration scope: This capability allows Migration Toolkit to retry migrating data. Retry attempts for issues with the schema migration aren't supported.

Modality scope: This reconnection capability is available with the data migration mode (`-dataOnly`). It's also available when you run a migration without specifying either the `-dataOnly` or `-schemaOnly` options.

You can specify several connection retry options:

```
-connRetryCount [<connection_attempts>]
```

Use the `-connRetryCount` option to specify the number of retry attempts to perform if the target database connection fails. The `[<connection_attempts>]` value must be a number between 0 and 50. The default is 3 retry attempts.

Since the retry applies to the migration of data, it isn't compatible with the `-schemaOnly` option.

Example:

```
./runMTK.sh -connRetryCount 2 -dataOnly -tables dept,emp,jobhist public
```

`-connRetryInterval [<seconds>]`

Use the `-connRetryInterval` option to specify the seconds to wait before each subsequent reconnection attempt if the target database connection fails. The `[<seconds>]` value must be a number between 0 and 900. The default is 30 seconds.

Since the retry applies to the migration of data, it isn't compatible with the `-schemaOnly` option.

Example:

```
./runMTK.sh -connRetryCount 2 -connRetryInterval 50 -dataOnly -tables dept,emp,jobhist public
```

`-abortConnOnFailure [true/false]`

Specify whether to abort the migration if all the reconnection attempts failed.

To skip migrating the failed table and proceed to the next table, set the `-abortConnOnFailure` to `false`. The default is `true`, which aborts the session if the connection fails after the specified `-connRetryCount` threshold.

Since the retry applies to the migration of data, it isn't compatible with the `-schemaOnly` option.

Example:

```
./runMTK.sh -connRetryCount 2 -connRetryInterval 50 -abortConnOnFailure false -dataOnly -tables dept,emp,jobhist public
```

`-pgIdleTxSessionTimeout [<seconds>]`

Specify the PostgreSQL or EDB Postgres Advanced Server `idle_in_transaction_session_timeout`, which defines the time after which the database terminates the session when a migration transaction is in idle state. `[<seconds>]` is a value greater than 0. The default is 180 seconds.

Example:

```
runMTK.sh -connRetryCount 2 -connRetryInterval 50 -pgIdleTxSessionTimeout 90 -dataOnly -tables dept,emp,jobhist public
```

Migration options for parallel data loading

When dealing with a relatively large number of tables or large tables, normal data migration takes a long time. As a solution to this problem, Migration Toolkit has a parallel data loading algorithm that allows migrating data in parallel at the schema level and at the table level.

Parallel data migration at the schema level uses multiple threads to migrate several tables at the same time. Parallel data loading at the table level uses multiple threads to migrate a single table from the source database to the target database more efficiently and in a shorter time.

Notes

- Parallel loading is useful only for large tables. For small tables, parallel threads might take more time. Thus we recommend using parallel data loading at the table level for relatively large tables.
- Parallel data loading doesn't apply to tables without a primary/unique key.
- You can't use the `fastCopy` parameter with parallel data loading.
- Parallel loading functionality applies only to the source Oracle, EDB Postgres Advanced Server, and PostgreSQL databases.

Migrations from Oracle databases

Because parallel data loading uses Oracle Flashback queries to perform data loading, ensure the user performing the migration has `FLASHBACK` and `READ / SELECT` privileges to the table objects to migrate. To allow queries on all tables, grant the `FLASHBACK ANY TABLE` privilege.

The following options control the way MTK migrates data in parallel.

`-loaderCount [<value>]`

Use the `-loaderCount` option to specify the number of parallel threads available in the pool to perform data load in parallel. This option is particularly useful if the host system that's running Migration Toolkit has high-end CPU and RAM resources. While `value` can be any nonzero, positive number, we recommend that value not exceed the number of CPU cores. For example, a dual-core CPU has an optimal value of `2`. The default is `1`.

The number of table-level threads can introduce overhead on the source database server depending on the table size and certain memory configurations, like `shared_buffers` and `work_mem` (memory configuration should be set to optimal values for the database server). It is recommended that for very large tables you try the lower `-loaderCount` option value, according to the total cores available on the machine, to calculate the optimal `-loaderCount` value.

Note

When multiple threads are used, depending on the actual number of threads and table data size, you might need to adjust the memory heap size for the Migration Toolkit.

`-parallelLoadRowLimit [<value>]`

Use the `-parallelLoadRowLimit` option to specify the minimum number of rows required in a table to perform data load in parallel at the table level. The `<value>` parameter must be greater than 0. The default is 100,000.

For example, if the value of `-parallelLoadRowLimit` is 10,000, only tables with the number of rows 10,000 and greater are loaded in parallel at the table level. The other tables migrate in sequential mode or at the schema level.

`-tableLoaderLimit [<value>]`

Use the `-tableLoaderLimit` option to specify the maximum number of threads to assign from the thread pool to load a single table data in parallel chunks at the table level. The `<value>` parameter must not be greater than the value of `-loaderCount`. The default value is equal to the value of the `-loaderCount` option. This option is implicitly changed based on the nondefault custom value of `-loaderCount`.

For example, if the `-loaderCount` is set to `4` and `-tableLoaderLimit` isn't specified, the `-tableLoaderLimit` is also set to `4`. Hence, if all threads aren't meant to be used for a single table load, reduce the `-tableLoaderLimit` to a lower value.

High-level overview of the parallel data-loading algorithm

The algorithm works in three steps:

1. A list of tables and table chunks is created from the ordered list of tables to migrate, sorted by the table name. The list is created based on the `-parallelLoadRowLimit` and `-tableLoaderLimit` option values.
2. A thread pool is created based on the `-loaderCount` option value.
3. The parallel data-loading algorithm creates tasks for each table/chunk from the list of tables and table chunks and executes them in the thread pool. As many tasks as the number of threads are executed in parallel. As soon as a task execution completes, its thread starts executing the next task from the list. Parallel data loading completes as soon as the last task from the list completes.

Example

In this scenario, migrate four tables—`tab1`, `tab2`, `tab3`, `tab4`—with parallel data loading:

```
./runMTK.sh -sourcedbtype enterprisedb -targetdbtype enterprisedb -loaderCount 4 -parallelLoadRowLimit 10000 -tableLoaderLimit 2 -tables tab1,tab2,tab3,tab4 public
```

Table `tab1` has 50,000 rows, table `tab2` has 9,500 rows, table `tab3` has 50,001 rows and table `tab4` has 9,999 rows. The list of tables for step 1 is `tab1`, `tab2`, `tab3`, and `tab4`.

Only tables `tab1` and `tab3` qualify for parallel loading on table level (`-parallelLoadRowLimit 10000`).

Since the value of the `-tableLoaderLimit` option is `2`, table `tab1` is loaded in two chunks, `tab1_chunk1` and `tab1_chunk2`. The number of rows in each chunk is determined based on the `totalRowCount/tableLoaderLimit`. In this case, both chunks have 25,000 rows ($50,000/2$).

Table `tab3` is also loaded in two chunks, `tab3_chunk1` and `tab3_chunk2`. As the total number of rows in table `tab3` is odd (`50,001`), the first chunk accommodates an additional row, so `tab3_chunk1` has 25,001 rows, and `tab3_chunk2` has 25,000 rows.

Here's the list of tables and table chunks created:

<code>tab1_chunk1</code>	(25,000 rows)
<code>tab1_chunk2</code>	(25,000 rows)
<code>tab2</code>	(9500 rows)
<code>tab3_chunk1</code>	(25,001 rows)
<code>tab3_chunk2</code>	(25,000 rows)
<code>tab4</code>	(9999 rows)

Thread pool with the four threads `T1`, `T2`, `T3`, `T4` is created.

Tasks migrating tables/chunks from the list created in step 1 are executed in the thread pool:

```
T1 → tab1_chunk1      (25,000 rows)
T2 → tab1_chunk2      (25,000 rows)
T3 → tab2              (9500 rows)
T4 → tab3_chunk1      (25,001 rows)
```

Thread **T3** finishes first and starts executing **tab3_chunk2**, the next chunk in the list:

```
T1 → tab1_chunk1      (25,000 rows)
T2 → tab1_chunk2      (25,000 rows)
T3 → tab3_chunk2      (25,000 rows)
T4 → tab3_chunk1      (25,001 rows)
```

Thread **T1** finishes second and starts executing **tab4**:

```
T1 → tab4              (9999 rows)
T2 → tab1_chunk2      (25,000 rows)
T3 → tab3_chunk2      (25,000 rows)
T4 → tab3_chunk1      (25,001 rows)
```

This way, four threads are used to migrate four tables in parallel at the mixed-schema and table level.

If you change command options to set **-tableLoaderLimit** to **1**, you can have pure schema-level parallelism:

```
./runMTK.sh -sourcedbtype enterprisedb -targetdbtype enterprisedb -loaderCount 4 -parallelLoadRowLimit
10000 -tableLoaderLimit 1 -tables tab1,tab2,tab3,tab4 public
```

Each table is processed simultaneously in one thread:

```
T1 → tab1      (50,000 rows)
T2 → tab2      (9500 rows)
T3 → tab3      (50,001 rows)
T4 → tab4      (9999 rows)
```

System resource recommendations and constraints

Choose the number of threads depending on the CPU and RAM resources available on the host system running Migration Toolkit. As there's a certain overhead on the source database in terms of CPU, IO, and disk use, exceeding parallel thread count beyond a certain threshold degrades Migration Toolkit performance.

Oracle-specific options

The following options apply only when the source database is Oracle.

-objectTypes

Import the user-defined object types from the schema list specified at the end of the **runMTK.sh** command.

-allUsers

Import all users and roles from the source database. The **-allUsers** option is supported only when migrating from an Oracle database to an EDB Postgres Advanced Server database.

-users <user_list>

Import the selected users or roles from the source Oracle database. **<user_list>** is a comma-separated list of user/role names with no intervening space characters (e.g., **-users MTK, SAMPLE, acctg**). The **-users** option is supported only when migrating from an Oracle database to an EDB Postgres Advanced Server database.

-allProfiles

Import all custom (that is, user-created) profiles from the source database. Other Oracle noncustom profiles such as **DEFAULT** and **MONITORING_PROFILE** aren't imported.

For the imported profiles, only the following password parameters associated with the profiles are imported:

FAILED_LOGIN_ATTEMPTS**PASSWORD_LIFE_TIME****PASSWORD_REUSE_TIME****PASSWORD_REUSE_MAX****PASSWORD_LOCK_TIME****PASSWORD_GRACE_TIME****PASSWORD_VERIFY_FUNCTION**

All other profile parameters, such as the Oracle resource parameters, aren't imported. The Oracle database user specified by **SRC_DB_USER** must have **SELECT** privilege on the Oracle data dictionary view **DBA_PROFILES**.

Note

The **-allProfiles** option is supported only when migrating from an Oracle database to an EDB Postgres Advanced Server database.

```
-profiles <profile_list>
```

Import the selected, custom (that is, user-created) profiles from the source Oracle database. `profile_list` is a comma-separated list of profile names with no intervening space characters (e.g., `-profiles ADMIN_PROFILE,USER_PROFILE`). Oracle noncustom profiles such as `DEFAULT` and `MONITORING_PROFILE` aren't imported.

As with the `-allProfiles` option, only the password parameters are imported. The Oracle database user specified by `SRC_DB_USER` must have `SELECT` privilege on the Oracle data dictionary view `DBA_PROFILES`.

Note

The `-profiles` option is supported only when migrating from an Oracle database to an EDB Postgres Advanced Server database.

```
-importPartitionAsTable <table_list>
```

Include the `-importPartitionAsTable` parameter to import the contents of a partitioned table that resides on an Oracle host into a single nonpartitioned table. `table_list` is a comma-separated list of table names with no intervening space characters (for example, `-importPartitionAsTable emp,dept,acctg`).

```
-copyViaDBLinkOra
```

The `dblink_ora` module provides EDB Postgres Advanced Server-to-Oracle connectivity at the SQL level. `dblink_ora` is bundled and installed as part of the EDB Postgres Advanced Server database installation. `dblink_ora` uses the `COPY API` method to transfer data between databases. This method is considerably faster than the `JDBC COPY` method.

This example uses the `dblink_ora COPY API` to migrate all tables from the `HR` schema:

```
./runMTK.sh -copyViaDBLinkOra -allTables HR
```

The target EDB Postgres Advanced Server database must have `dblink_ora` installed and configured. See [dblink_ora](#).

```
-allDBLinks [link_Name_1=password_1,link_Name_2=password_2,...]
```

Choose this option to migrate Oracle database links. The password information for each link connection in the source database is encrypted so, unless specified, the dummy password `edb` is substituted.

To migrate all database links using `edb` as the dummy password for the connected user:

```
./runMTK.sh -allDBLinks HR
```

You can alternatively specify the password for each of the database links through a comma-separated list of name=value pairs with no intervening space characters. Specify the link name on the left side of the pair and the password value on the right side.

To migrate all database links with the actual passwords specified on the command line:

```
./runMTK.sh -allDBLinks LINK_NAME1=abc,LINK_NAME2=xyz HR
```

Migration Toolkit migrates only the database link types that are currently supported by EnterpriseDB. These types include fixed user links of public and private type.

-allSynonyms

Include the **-allSynonyms** option to migrate all public and private synonyms from an Oracle database to an EDB Postgres Advanced Server database. If a synonym with the same name already exists in the target database, the existing synonym is replaced with the migrated version.

-allPublicSynonyms

Include the **-allPublicSynonyms** option to migrate all public synonyms from an Oracle database to an EDB Postgres Advanced Server database. If a synonym with the same name already exists in the target database, the existing synonym is replaced with the migrated version.

-excludeSynonyms <synonym_list>

Exclude selected synonyms from migration. The **synonym_list** is a comma-separated list of synonym names with no intervening space characters. This option applies when all synonyms from a schema are selected for migration using the **-allSynonyms** option.

Example: **-allSynonyms -excludeSynonyms SYNEMP1,SYNEMP2**

-allPrivateSynonyms

Include the **-allPrivateSynonyms** option to migrate all private synonyms from an Oracle database to an EDB Postgres Advanced Server database. If a synonym with the same name already exists in the target database, the existing synonym is replaced with the migrated version.

-useOraCase

Include the **-useOraCase** option to preserve the Oracle default, uppercase naming convention for all database objects when migrating from an Oracle database to an EDB Postgres Advanced Server database.

The uppercase naming convention is preserved for tables, views, sequences, procedures, functions, triggers, packages, and so on. For these database objects, the uppercase naming convention applies to:

- The names of the database objects
- The column names, key names, index names, constraint names, and so on, of the tables and views
- The **SELECT** column list for a view
- The parameter names that are part of the procedure or function header

Note

In the procedural code body of a procedure, function, trigger, or package, you might have to manually edit identifier references for the program to execute without an error. Such corrections are in regard to the proper case conversion of identifier references that might have occurred.

Note

When you specify the **-useOraCase** option, you might need to specify the **-skipUserSchemaCreation** option as well. For information, see the description of the **-skipUserSchemaCreation** option.

The default behavior of the Migration Toolkit without the **-useOraCase** option is that database object names are extracted from Oracle without enclosing quotation marks unless the database object was explicitly created in Oracle with enclosing quotation marks. The following is a portion of a table command generated by the Migration Toolkit with the **-offlineMigration** option:

```
CREATE TABLE DEPT
(
    DEPTNO NUMBER(2) NOT
NULL,
    DNAME VARCHAR2(14),
    LOC
VARCHAR2(13)
);
ALTER TABLE DEPT ADD CONSTRAINT DEPT_PK PRIMARY KEY (DEPTNO);
ALTER TABLE DEPT ADD CONSTRAINT DEPT_DNAME_UQ UNIQUE
(DNAME);
```

When you then migrate this table and create it in EDB Postgres Advanced Server, all unquoted object names are converted to lowercase letters, so the table appears in EDB Postgres Advanced Server as follows:

```
Table "edb.dept"
Column | Type          |
Modifiers
-----+-----+-----
deptno | numeric(2,0)   | not
null
dname  | character varying(14)
|
loc    | character varying(13)
|
Indexes:
    "dept_pk" PRIMARY KEY, btree
(deptno)
    "dept_dname_uq" UNIQUE CONSTRAINT, btree (dname)
```

If your EDB Postgres Advanced Server applications are referencing the migrated database objects using quoted uppercase identifiers, the applications fail since the database object names are now in lower case:

```
usepostcase=# SELECT * FROM "DEPT";
ERROR: relation "DEPT" does not
exist
LINE 1: SELECT * FROM "DEPT";
```

If your application uses quoted uppercase identifiers, perform the migration with the `-useOraCase` option. The DDL encloses all database object names in quotes:

```
CREATE TABLE "DEPT"
(
    "DEPTNO" NUMBER(2) NOT NULL,
    "DNAME" VARCHAR2(14),
    "LOC" VARCHAR2(13)
);
ALTER TABLE "DEPT" ADD CONSTRAINT "DEPT_PK" PRIMARY KEY
("DEPTNO");
ALTER TABLE "DEPT" ADD CONSTRAINT "DEPT_DNAME_UQ" UNIQUE
("DNAME");
```

When you then migrate this table and create it in EDB Postgres Advanced Server, all object names are maintained in uppercase letters, so the table appears in EDB Postgres Advanced Server as follows:

```
Table "EDB.DEPT"
Column | Type          | Modifiers
-----+-----+-----
DEPTNO | numeric(2,0)  | not null
DNAME  | character varying(14)
LOC    | character varying(13)

Indexes:
    "DEPT_PK" PRIMARY KEY, btree
    ("DEPTNO")
    "DEPT_DNAME_UQ" UNIQUE CONSTRAINT, btree
    ("DNAME")
```

Applications can then access the object using quoted uppercase names:

```
useoracase=# SELECT * FROM "DEPT";
DEPTNO | DNAME          | LOC
-----+-----+-----
10     | ACCOUNTING    | NEW YORK
20     | RESEARCH      | DALLAS
30     | SALES         | CHICAGO
40     | OPERATIONS    | BOSTON
(4 rows)
```

`-skipUserSchemaCreation`

When an Oracle user is migrated, a role (that is, a user name) is created in the target database server for the Oracle user if the role doesn't already exist. The role name is created in lowercase letters. When a new role is created, a schema with the same name is also created in lowercase letters.

Specifying the `-skipUserSchemaCreation` option prevents the automatic schema creation for a migrated Oracle user name. This option is particularly useful when the `-useOraCase` option is specified to prevent creating two schemas with the same name where only the case is different. Specifying the `-useOraCase` option results in creating a schema in the Oracle naming convention of uppercase letters for the source schema specified following the options list when Migration Toolkit is invoked.

Thus, if the `-useOraCase` option is specified without the `-skipUserSchemaCreation` option, the target database results in having two identically named schemas with one in lowercase letters and the other in uppercase letters. If the `-useOraCase` option is specified along with the `-skipUserSchemaCreation` option, the target database has only the schema in uppercase letters.

Miscellaneous options

Use these migration options to view Migration Toolkit help and version information. You can also use these options to control Migration Toolkit feedback and logging options.

`-help`

Display the application command-line usage information.

`-logDir <log_path>`

Include this option to specify where to write the log files. `<log_path>` is the location for saving log files. By default, on Linux log files are written to:

`$HOME/.enterprisedb/migration-toolkit/logs`

On Windows, the log files are saved to:

`%HOMEDRIVE%%HOMEPATH%\enterprisedb\migration-toolkit\logs``-logFileCount <file_count>`

Include this option to specify the number of files used in log file rotation. Specify a value of `0` to disable log file rotation and create a single log file. This file is truncated when it reaches the value specified using the `logFileSize` option. `<file_count>` must be greater than or equal to 0. The default is 20.

`-logFileSize <file_size>`

Include this option to specify the maximum file size limit in MB before rotating to a new log file. `file_size` must be greater than 0. The default is 50.

`-logBadSQL`

Include this option to save the schema definition (DDL script) of any failed objects to a file. The file is saved under the same path used for the error logs and is named in the format:

`mtk_bad_sql_<schema_name_timestamp>.sql`

Where `schema_name` is the name of the schema and `timestamp` is the timestamp of the Migration Toolkit run.

`-verbose [on|off]`

Display application log messages on standard output. By default, verbose is on.

`-version`

Display the Migration Toolkit version.

Example

This example performs a migration from Oracle to EDB Postgres Advanced Server.

The following is the content of the `toolkit.properties` file:

```
SRC_DB_URL=jdbc:oracle:thin:@192.168.2.6:1521:xe
SRC_DB_USER=edb
SRC_DB_PASSWORD=password

TARGET_DB_URL=jdbc:edb://localhost:5444/edb
TARGET_DB_USER=enterprisedb
TARGET_DB_PASSWORD=password
```

The following command invokes Migration Toolkit:

```
./runMTK.sh EDB
```

```
output
Running EnterpriseDB Migration Toolkit (Build 48.0.0) ...
Source database connectivity info...
conn =jdbc:oracle:thin:@192.168.2.6:1521:xe
user =edb
password=*****\*
Target database connectivity info...
conn =jdbc:edb://localhost:5444/edb
user =enterprisedb
password=*****\*
Connecting with source Oracle database server...
Connected to Oracle, version 'Oracle Database 10g Express Edition
Release 10.2.0.1.0 - Production'
Connecting with target EnterpriseDB database server...
Connected to EnterpriseDB, version '9.4.0.0'
Importing redwood schema EDB...
Creating Schema...edb
Creating Sequence: NEXT_EMPNO
Creating Tables...
Creating Table: BAD_TABLE
MTK-15013: Error Creating Table BAD_TABLE
DB-42704: ERROR: type "binary_double" does not exist at position 58
-- CREATE TABLE BAD_TABLE (
-- F1 NUMBER NOT NULL,
-- Line 3: F2 BINARY_DOUBLE
-- ^
Creating Table: DEPT
Creating Table: EMP
Creating Table: JOBHIST
Creating Table: "MixedCase"
Creating Table: "lowercase"
Created 5 tables.
Loading Table Data in 8 MB batches...
Loading Table: DEPT ...
[DEPT] Migrated 4 rows.
[DEPT] Table Data Load Summary: Total Time(s): 0.147 Total Rows: 4
Loading Table: EMP ...
[EMP] Migrated 14 rows.
[EMP] Table Data Load Summary: Total Time(s): 0.077 Total Rows: 14
Loading Table: JOBHIST ...
```

```

[JOBHIST] Migrated 17 rows.
[JOBHIST] Table Data Load Summary: Total Time(s): 0.042 Total Rows: 17
Total Size(MB): 9.765625E-4
Loading Table: "MixedCase" ...
["MixedCase"] Table Data Load Summary: Total Time(s): 0.098 Total Rows:0
Loading Table: "lowercase" ...
["lowercase"] Table Data Load Summary: Total Time(s): 0.066 Total Rows:0
Data Load Summary: Total Time (sec): 0.806 Total Rows: 35 Total
Size(MB): 0.001
Creating Constraint: DEPT_PK
Creating Constraint: DEPT_DNAME_UQ
Creating Constraint: EMP_PK
Creating Constraint: JOBHIST_PK
Creating Constraint: SYS_C008958
MTK-15001: Error Creating Constraint SYS_C008958
DB-42P01: com.edb.util.PSQLException: ERROR: relation "bad_table" does
not exist
Creating Constraint: EMP_REF_DEPT_FK
Creating Constraint: EMP_SAL_CHK
Creating Constraint: JOBHIST_REF_DEPT_FK
Creating Constraint: JOBHIST_REF_EMP_FK
Creating Constraint: JOBHIST_DATE_CHK
Creating Trigger: USER_AUDIT_TRIG
Creating Trigger: EMP_SAL_TRIG
MTK-13009:Warning! Skipping migration of trigger DROP_TRIGGER, currently
non-table triggers are not supported in target database.
Creating View: SALESEMP
Creating Function: EMP_COMP
Creating Package: EMP_ADMIN
MTK-16005:Package Body is Invalid, Skipping...
Schema EDB imported with errors.
MTK-12001: The user/role migration failed due to insufficient
privileges.
Grant the user SELECT privilege on the following Oracle catalogs:
DBA_ROLES
DBA_USERS
DBA_TAB_PRIVS
DBA_PROFILES
DBA_ROLE_PRIVS
ROLE_ROLE_PRIVS
DBA_SYS_PRIVS
One or more schema objects could not be imported during the migration
process. Please review the migration output for more details.
Migration logs have been saved to
/home/user/.enterprisedb/migration-toolkit/logs
***** Migration Summary *****
Sequences: 1 out of 1
Tables: 5 out of 6
Constraints: 9 out of 10
Triggers: 2 out of 3 (skipped 1)
Views: 1 out of 1
Functions: 1 out of 1
Packages: 1 out of 1
Total objects: 30
Successful count: 20
Failed count: 2
Skipped count: 1
Invalid count: 7
List of failed objects
=====

```

Tables

1. EDB.BAD_TABLE

Constraints

1. EDB.BAD_TABLE.SYS_C008958

List of invalid objects

=====

1. EDB.HIRE_CLERK (FUNCTION)

2. EDB.NEW_EMPNO (FUNCTION)

3. EDB.EMP_ADMIN (PACKAGE BODY)

4. EDB.EMP_QUERY (PROCEDURE)

5. EDB.EMP_QUERY_CALLER (PROCEDURE)

6. EDB.LIST_EMP (PROCEDURE)

7. EDB.SELECT_EMP (PROCEDURE)

Skipped and unsupported database objects are omitted. The migration information is summarized in the Migration Summary at the end of the run.

8.2 Specifying Migration Toolkit command options using a file

EDB supports using a text file to specify several [database migration options](#) instead of manually entering them at the command line. This is particularly helpful when:

- You have to specify a large number of command line options or a large number of values for complex migration scenarios. Entering all options or values at the command line becomes cumbersome, or you might reach command line length limits.
- You want to apply the same migration options and values to several different databases, or you might want to automate database migration. Running `runMTK.sh` with several options and values repeatedly at the command line becomes prone to errors.

The `-optionsFile <options_textfile>` argument allows you to reference a text file with options and values, making the migration process easier.

Execute a migration using an options file

To perform a migration with a file that specifies the command options:

1. [Create a text file with the command options.](#)
2. [Perform a migration by invoking the file with the command options.](#)

8.2.1 Creating a file with the command options

To specify a text file that contains command line options and values, create an `<options_textfile>` file with the options you want to execute. Then run the `runMTK.sh` command with the `-optionsFile <options_textfile>` argument.

Note

You can store the `<options_textfile>` file in the same folder as the Migration Toolkit utility or in a directory of your choice.

Creating the text file - guidelines

Create a text file in a location that's accessible by the user performing the migration. Then, add one option per line following these guidelines.

Skip the dash

Use the same option syntax as in the command line options without the dash (-). The options file recognizes each line as a new parameter and doesn't require a dash as a delimiter.

Option in CLI	Supported syntax for the option in file
<code>-tables hr</code>	tables hr

Use a line per option

Use a new line for each option. Use one line per key-value pair.

Option in CLI	Supported syntax for the option in file
<code>-tables hr -allViews</code> <code>-excludeViews all_emp,acct_list</code>	tables hr allViews excludeViews all_emp,acct_list

Specify values for an option in a single line or multiple lines

You can add a single line with an option and multiple values for that option.

You can add several lines for the same option, each time with a different value.

And you can combine both approaches, because Migration Toolkit combines all lines that specify the same option.

Option in CLI	Supported syntax for the option in file
<code>-tables TAB1,TAB2,TAB3</code>	tables TAB1 tables TAB2 tables TAB3 tables TAB1,TAB2,TAB3

Option in CLI	Supported syntax for the option in file
	tables TAB1 tables TAB2,TAB3

Note

Duplicating options is supported only for the options file.

When using `-tables` at the command line more than one time, only the last parameter is executed. For example, `./runMTK.sh -tables TAB2 -tables TAB3` includes only the `TAB3` table in the migration.

Use a space or equal sign to provide values

For options that require values (key-value pairs), separate the option from the value with a space or an equal sign (=). Both syntaxes are supported.

Option in CLI	Supported syntax for the option in file
<code>-views all_emp</code>	views=all_emp views all_emp

Don't specify an option in the file and at the command line

Don't include an option in the options file if you're specifying it as a flag at the command line.

Specifying an option both at the command line and in the text file causes the migration to fail.

Order of processing

Migration Toolkit reads command line options and option files in the order you provide them when running the command. Ensure you add the `schema scope` as the last parameter at the command line.

For example, if you run the following command, Migration Toolkit first recognizes the `-sourcedbtype oracle` option, and then reads the contents of `options_textfile` in order from top to bottom. The last parameter is the `<schema_scope>`.

```
runMTK.sh -sourcedbtype oracle -optionsFile options_textfile <schema_scope>
```

Using an options file means that you can employ different syntaxes to perform a migration where parameters are executed in the same way. The following alternatives perform the same migration.

Alternative 1

```
runMTK.sh -sourcedbtype oracle -optionsFile <options_textfile> <schema_scope>
```

Where the content of the `<options_textfile>` file is:

```
tables=TABLE1  
dataOnly
```

Alternative 2

```
runMTK.sh -sourcedbtype oracle -optionsFile <options_textfile> -dataOnly <schema_scope>
```

Where the content of the `<options_textfile>` file is:

```
tables=TAB1
```

Other considerations

- **Commenting out** — Use a pound sign (#) to comment out a line. Migration Toolkit doesn't execute lines that are commented out.
- **Limitations** — Don't use quote marks.
- **Space processing** — Migration Toolkit removes spaces between an option and its value. When you separate an option from its value with a space, only one space is treated as a delimiter.

For example, to run the `-views all_emp, mgmt_list, acct_list` command line option, you can add `views all_emp,mgmt_list,acct_list` or `views = all_emp,mgmt_list,acct_list` in the options file.

After you create the options file, [execute the migration with it](#).

8.2.2 Executing a migration with a command options file

After you create the options file, reference it when executing the migration command:

```
./runMTK.sh -optionsFile <options_textfile> <schema_scope>
```

Note

Provide the full path for `<options_textfile>` if you're running the Migration Toolkit command from a different folder from where you stored the file.

You can specify all options in the file or use a mix of the command line and the options file to specify the migration parameters.

Provide the scope for the schema migration

Specify the scope of the schemas to migrate:

- If you want to migrate all schemas, add the `-allSchemas` option at the command line or add `allSchemas` in your options file. For both methods, this option must be the last parameter.
- If you want to specify a subset of schemas, specify the schemas you want to migrate at the command line with no preceding option and as a comma-separated list. Schema specifications must be the last parameter at the command line.

See [Migrating schemas](#) for more information on `<schema_scope>`.

Here are some examples for specifying all options in the file.

Migrate a schema with specific tables

Content of the `options_textfile` text file:

```
tables comp_schema.emp,comp_schema.dept,finance_schema.acctg
```

Syntax of the migration command:

```
./runMTK.sh -optionsFile options_textfile schema1
```

Command line equivalent:

```
./runMTK.sh -tables comp_schema.emp,comp_schema.dept,finance_schema.acctg comp_schema,finance_schema
```

Use options file to exclude tables and include functions

You can use the options file to separate table names from configuration options. You can also use comments to clarify why parameters are used.

This example excludes tables and includes functions. It shows some comments to explain why tables are excluded.

Content of the `excludeInclude.options` text file:

```
# finances
excludeTables comp_schema.emp,finance_schema.jobhist

# temporary, #MIGRATION-123, ask Ben
excludeTables temp_schema.temp_table

excludeTables more_schema.more_tables

funcs finance_schema.add_two_numbers
```

Syntax of the migration command:

```
./runMTK.sh -allTables -optionsFile excludeInclude.options -safeMode -connRetryCount 7 schema1
```

Command line equivalent:

```
./runMTK.sh -allTables -excludeTables
comp_schema.emp,finance_schema.jobhist,temp_schema.temp_table,more_schema.more_tables -funcs
finance_schema.add_two_numbers -safeMode -connRetryCount 7 schema1
```

Offline migration

Run an offline migration with all schemas.

Content of the `options_textfile` text file:

```
offlineMigration file_dest
allTables
schemaOnly
```

Syntax of the migration command:

```
./runMTK.sh -optionsFile options_textfile -allSchemas
```

Command line equivalent:

```
./runMTK.sh -offlineMigration file_dest -allTables -schemaOnly -allSchemas
```

9 Migration errors

During the migration process, the migration summary displays the progress of the migration. If an error occurs, the summary displays information about the error. The migration summary is also written to a log file.

On Linux, the default location for the log files is:

```
$HOME/.enterprisedb/migration-toolkit/logs
```

On Windows, the log files are saved to:

```
%HOMEDRIVE%%HOMEPATH%\enterprisedb\migration-toolkit\logs
```

You can specify an alternative log file location with the `-logdir <log_path>` option in Migration Toolkit.

Connection errors

Migration Toolkit uses information from the `toolkit.properties` file to connect to the source and target databases. Most of the connection errors that occur when using Migration Toolkit are related to the information specified in the `toolkit.properties` file.

Invalid username/password

When I try to perform a migration from an Oracle database with Migration Toolkit, I get the following error:

```
MTK-11009: Error Connecting Database "Oracle"
DB-1017: java.sql.SQLException: ORA-01017: invalid username/password;
logon denied
The user name or password specified in the `toolkit.properties` file is
not valid to use to connect to the Oracle source database.
```

To resolve this error, edit the `toolkit.properties` file, specifying the name and password of a valid user with privileges to perform the migration in the `SRC_DB_USER` and `SRC_DB_PASSWORD` properties.

Invalid file permissions

When I try to perform a migration with Migration Toolkit, I get the following error:

```
MTK-11015: The connection credentials file ../etc/toolkit.properties is
not secure and accessible to group/others users. This file contains
plain passwords and should be restricted to Migration Toolkit owner user
only.
```

To resolve this error, ensure the operating system user account running the Migration Toolkit is the owner of the `toolkit.properties` file with a minimum of read permission on the file. In addition, there must be no permissions of any kind for group and other users.

The following is an example of the recommended file permissions, where user `enterprisedb` is running the Migration Toolkit.

```
-rw----- 1 enterprisedb enterprisedb 191 Aug 1 09:59 toolkit.properties
```

Connection rejected: FATAL: password

When I try to perform a migration with Migration Toolkit, I get the following error:

```
MTK-11009: Error Connecting Database "EDB Postgres EDB Postgres Advanced Server"
DB-28P01: com.edb.util.PSQLException: FATAL: password authentication
failed for user "name"
```

The user name or password specified in the `toolkit.properties` file is not valid to use to connect to the Postgres database.

To resolve this error, edit the `toolkit.properties` file, specifying the name and password of a valid user with privileges to perform the migration in the `TARGET_DB_USER` and `TARGET_DB_PASSWORD` properties.

Exception: ORA-28000: the account is locked

When I try to perform a migration from an Oracle database with Migration Toolkit, I get the following error message:

```
MTK-11009: Error Connecting Database "Oracle"
DB-28000: java.sql.SQLException: ORA-28000: the account is locked
The Oracle account associated with the user name specified in the
toolkit.properties file is locked.
```

To resolve this error, you can either unlock the user account on the Oracle server or edit the `toolkit.properties` file, specifying the name and password of a valid user with privileges to perform the migration in the `SRC_DB_USER` and `SRC_DB_PASSWORD` parameters.

Exception: oracle.jdbc.driver.OracleDriver

When I try to perform a migration with Migration Toolkit, the migration fails and I get the error message:

```
MTK-11009: Error Connecting Database "Oracle"
java.lang.ClassNotFoundException: oracle.jdbc.driver.OracleDriver
```

Before using Migration Toolkit, you must download and install the appropriate JDBC driver for the database that you are migrating from. See [Installing a JDBC driver](#) for complete instructions.

I/O exception: The Network Adapter could not establish the connection

When I try to perform a migration with Migration Toolkit, I get the following error:

```
MTK-11009: Error Connecting Database "Oracle"
DB-17002: java.sql.SQLException: Io exception: The
Network Adapter could not establish the connection
```

The JDBC URL for the source database specified in the `toolkit.properties` file contains invalid connection properties.

To resolve this error, edit the `toolkit.properties` file, specifying valid connection information for the source database in the `SRC_DB_URL` property. The `SRC_DB_URL` value is database specific. See [Building the toolkit-properties file](#) for detailed information for your source database type.

Exception: The URL specified for the “target” database is invalid

When I try to perform a migration with Migration Toolkit, I get the following error:

```
MTK-10045: The URL specified for the "target" database is invalid.
Check the connectivity credentials.
```

The JDBC URL for the target database (EDB Postgres Advanced Server) specified in the `toolkit.properties` file contains invalid connection properties.

To resolve this error, edit the `toolkit.properties` file, specifying valid connection information for the target database in the `TARGET_DB_URL` property. For information about forming a JDBC URL for EDB Postgres Advanced Server, see [Defining an EDB Postgres Advanced Server URL](#).

Migration errors

The following errors can occur after Migration Toolkit has successfully connected to the target and source database servers.

ERROR: Extra data after last expected column

When migrating a table online, I get the error message:

```
MTK-17001: Error Loading Data into Table: table_name
DB-22P04: com.edb.util.PSQLException: ERROR: extra data after last
expected column
Where: COPY table_name, line 5: "50|HR|LOS|ANGELES"
```

This error occurs when the data in a column in `table_name` includes the delimiter character. To correct this error, change the delimiter character to a character not found in the table contents.

Note

In this example, the pipe character (|) occurs in the text `LOS|ANGELES`, intended for insertion into the last column, and the Migration Toolkit is run using the `-copyDelimiter '|'` option, which results in the error.

Error loading data into table: `TABLE_NAME`

When performing a data-only migration, I get the following error:

```
MTK-17001: Error Loading Data into Table: TABLE_NAME
DB-42P01: com.edb.util.PSQLException: ERROR: relation
"schema.\ table_name" does not exist
*I also get the error:*
Trying to reload table: TABLE_NAME through bulk inserts with a batch
size of 100
MTK-17001: Error Loading Data into Table: TABLE_NAME
DB-42P01: com.edb.util.PSQLException: ERROR: relation
"schema.\ table_name" does not exist
```

You must create a table to receive the data in the target database before you can migrate the data. Verify that a table with a name of `TABLE_NAME` exists in the target database. Create the table if necessary, and retry the data migration.

Error creating constraint CONS_NAME_FK

When I perform a table migration that includes indexes and constraints, I get the following error message:

```
MTK-15001: Error Creating Constraint EMP_DEPT_FK
DB-42P01: com.edb.util.PSQLException: ERROR: relation "hr.departments"
does not exist
Creating Constraint: EMP_JOB_FK
MTK-15001: Error Creating Constraint EMP_JOB_FK
DB-42P01: com.edb.util.PSQLException: ERROR: relation "hr.jobs" does not exist
Creating Constraint: EMP_MANAGER_FK
Schema HR imported with errors.
One or more schema objects could not be imported during the migration
process. Please review the migration output for more details.
Migration logs have been saved to
/home/user/.enterisedb/migration-toolkit/logs
***** Migration Summary *****
Tables: 1 out of 1
Constraints: 4 out of 6
Total objects: 7
Successful count: 5
Failed count: 2
Invalid count: 0
List of failed objects
=====
Constraints
-----
1. HR.EMPLOYEES.EMP_DEPT_FK
2. HR.EMPLOYEES.EMP_JOB_FK
```

The table you're migrating includes a foreign key constraint on a table that doesn't exist in the target database. Migration Toolkit creates the table, omitting the foreign key constraint.

You can avoid generating the error message by including the `-skipFKConst` option in the Migration Toolkit command.

Error Loading Data into Table

I've already migrated the table definition; when I try to migrate the data into the table, I get an error:

```
MTK-17001: Error Loading Data into Table: DEPARTMENTS
DB-22P04: com.edb.util.PSQLException: ERROR: extra data after last
expected column
Where: COPY departments, line 1: "10 Administration 200 1700"
Trying to reload table: DEPARTMENTS through bulk inserts with a batch
size of 100
MTK-17000: Batch entry 0 INSERT INTO hr.DEPARTMENTS VALUES (10,
'Administration', 200, 1700); was aborted. Call getNextException to see
the cause.
DB-42601: java.sql.BatchUpdateException: Batch entry 0 INSERT INTO
hr.DEPARTMENTS VALUES (10, 'Administration', 200, 1700); was aborted.
Call getNextException to see the cause., Skipping Batch
MTK-17000:ERROR: INSERT has more expressions than target columns
Position: 48
[DEPARTMENTS] Table Data Load Summary: Total Time(s): 0.037 Total Rows:0
Data Load Summary: Total Time (sec): 0.168 Total Rows: 0 Total Size(MB):0.0
Schema HR imported with errors.
```

The table definition in the target database doesn't match the migrated data. If you altered the target or source table, confirm that the table definition and the data are compatible.

ERROR: value too long for type

I've already migrated the table definition. When I try to migrate the data into the table, I get the following error:

```
MTK-17001: Error Loading Data into Table: DEPARTMENTS
DB-22001: com.edb.util.PSQLException: ERROR: value too long for type
character(1)
Where: COPY departments, line 1, column location_id: "1700"
Trying to reload table: DEPARTMENTS through bulk inserts with a batch
size of 100
MTK-17000: Batch entry 0 INSERT INTO hr.DEPARTMENTS VALUES (10,
'Administration', 200, 1700); was aborted. Call getNextException to see
the cause.
DB-22001: java.sql.BatchUpdateException: Batch entry 0 INSERT INTO
hr.DEPARTMENTS VALUES (10, 'Administration', 200, 1700); was aborted.
Call getNextException to see the cause., Skipping Batch
MTK-17000:ERROR: value too long for type character(1)
```

A column in the target database is not large enough to receive the migrated data. This problem can occur if the table definition is altered after migration. The column name (in this example, `location_id`) is identified in the line that begins with **Where:**.

Where: `COPY departments, line 1, column location_id: "1700"`

To correct the problem, adjust the column size and retry the migration.

ERROR: Exception in thread:OutOfMemoryError

When migrating from a MySQL database, I encounter the following error:

```
Loading Table: big_range ...
Exception in thread "dataload_job_1" java.lang.OutOfMemoryError: Java
heap space
at com.mysql.jdbc.MySQLIO.nextRow(MySQLIO.java:1370)
at com.mysql.jdbc.MySQLIO.readSingleRowSet(MySQLIO.java:2369)
at com.mysql.jdbc.MySQLIO.getResultSet(MySQLIO.java:451)
at com.mysql.jdbc.MySQLIO.readResultsForQueryOrUpdate(MySQLIO.java:2076)
at com.mysql.jdbc.MySQLIO.readAllResults(MySQLIO.java:1451)
at com.mysql.jdbc.MySQLIO.sqlQueryDirect(MySQLIO.java:1787)
at com.mysql.jdbc.Connection.execSQL(Connection.java:3277)
at com.mysql.jdbc.Connection.execSQL(Connection.java:3206)
at com.mysql.jdbc.Statement.executeQuery(Statement.java:1232)
at com.edb.dbhandler.mysql.Data.getTableData(Data.java:90)
at com.edb.DataLoader.loadDataInFastMode(DataLoader.java:594)
at com.edb.DataLoader.run(DataLoader.java:186)
at java.lang.Thread.run(Thread.java:722)
```

By default, the MySQL JDBC driver fetches all of the rows in a table into Migration Toolkit in a single network round trip. This behavior can easily exceed available memory for large tables.

To correct the problem, specify `-fetchSize 1` as a command-line argument when you retry the migration.

Error: Read timeout during data migration

When I perform a migration to a target EDB Postgres Advanced Server or PostgreSQL server, Migration Toolkit reports a timeout error:

```
DB-08006: com.edb.util.PSQLException: An I/O error occurred while sending to the backend.
java.net.SocketTimeoutException: Read timed out
```

The target database server is either EDB Postgres Advanced Server or PostgreSQL and you are performing a data migration. The source tables' size ranges in the hundreds of gigabytes or the network between the database server and Migration Toolkit is slow, which causes the migration to take multiple hours.

The error message is reported after a successful data copy at the stage of enabling constraints.

This occurs because the server's default configuration closes idle TCP connections opened by Migration Toolkit to create constraints.

To work around this issue, modify the following parameters in the `postgresql.conf` file of the target database server:

```
tcp_keepalives_idle = 60          # TCP_KEEPIDL, in seconds;
tcp_keepalives_interval = 10     # TCP_KEEPINTVL, in seconds;
tcp_keepalives_count = 10       # TCP_KEEPCNT;
```

Then, reload the configuration to make the changes effective:

```
select pg_reload_conf();
```

This way, tcp_keepalive messages won't let the connection become idle and Migration Toolkit won't report a broken connection.

Incorrect timestamps and Daylight Saving Time

When migrating from SQL Server to PostgreSQL using the MSSQL JDBC driver, an error in the migration of timestamps may occur that is not reported. Specifically, if the source database contains rows with timestamp values that are within a range of time when Daylight Savings Time was in effect, those rows will be migrated to the target database with the wrong timestamps.

To resolve this issue, update the `runMTK.sh` file, and provide the option `-Duser.timezone=GMT`. This option will then be in effect when you run the toolkit.

For example, suppose this is the original line:

```
runJREApplication $JAVA_HEAP_SIZE -Dprop=$base/etc/toolkit.properties -cp $base/bin/edb-
migrationtoolkit.jar:$base/lib/* com.edb.MigrationToolkit "$@"
```

Update this line by inserting `-Duser.timezone=GMT` before the `-cp` option:

```
runJREApplication $JAVA_HEAP_SIZE -Dprop=$base/etc/toolkit.properties -Duser.timezone=GMT -cp
$base/bin/edb-migrationtoolkit.jar:$base/lib/* com.edb.MigrationToolkit "$@"
```


10 Error codes

When Migration Toolkit encounters a problem, it issues a notification in the form of an error code and a message.

Each error code begins with the prefix MTK- followed by five digits. The first two digits denote the error class, which is a general classification of the error. The last three digits represent the specific error condition.

The table lists the error classes.

Error Class	Description
00	Successful completion
01	Information
02	Warning
03	General drror
0x	Reserved
10	Invalid user input
11	Configuration error
12	Insufficient privileges
13	Unsupported feature
14	Missing object
15	Schema migration
16	Procedural language migration
17	Data loading

If there's an error reported back by a specific database server, this error message is prefixed with **DB-**. For example, if table creation fails due to an existing table in a Postgres database server, the error code **42P07** is returned by the database server. The specific error in the Migration Toolkit log appears as **DB-42P07**.

Error code summary

In the following tables, the column **Error code** lists the Migration Toolkit error codes. The **Message and resolution** column contains the message displayed with the error code. The message explains the cause of the error and how to resolve it.

In the **Message and resolution** column, **\$NAME** is a placeholder for information that is substituted at runtime with the appropriate value.

Class 02 - Warning

This class represents the warning messages.

Error code	Message and resolution
MTK-02000	All the warnings that relate to this class and don't have a specific error code binding use this error code.
MTK-02001	Run 'runMTK -help' to see the usage details.
MTK-02002	Warning! The offline migration path <i>\$OFFLINE_PATH</i> does not exist, the scripts will be created under the user home folder.

Class 10 - Invalid user input

This class represents invalid user input values.

Error Code	Message and Resolution
MTK-10000	All the errors that relate to this class and don't have a specific error code binding use this error code.
MTK-10001	You cannot select information_schema, dbo, sys, or pg_catalog as target schemas. These are used to store metadata information in <i>\$DATABASE</i> .
MTK-10002	The '-schemaOnly' and '-dataOnly' options cannot be specified at the same time.
MTK-10003	The "\t" is required as a copyDelimiter, when using escapeTabDelimiter option.
MTK-10004	The -truncLoad option can only be given with the -dataOnly option.
MTK-10005	The '-dataOnly' option is applicable only for -allTables/-tables option. Schema DDL cannot be copied when this option is in place.
MTK-10006	The -constraints, -indexes and -triggers options are applicable only in the context of -allTables/-tables option.
MTK-10007	The '-vacuumAnalyze' and '-analyze' options cannot be specified at the same time.
MTK-10008	The -skipFKConst option can only be given with -constraints option.
MTK-10009	The -skipCKConst option can only be given with -constraints option.
MTK-10010	The -fastCopy option cannot be specified with -schemaOnly option.
MTK-10011	The -skipColDefaultClause cannot be specified with -dataOnly option.
MTK-10012	The '-customColTypeMapping' and '-customColTypeMappingFile' options cannot be specified at the same time.
MTK-10013	Provided default date time must be in following format 'yyyy-MM-dd_HH:mm:ss'. Time portion is optional, to specify time, the underscore symbol '_' is necessary.
MTK-10014	You cannot select MySQL and Sybase as the target database for migration.
MTK-10015	Copy delimiter should be a single character.
MTK-10016	Batch size should be between 1 - 50000.
MTK-10017	Invalid number for batch size, use a number from 1-50000.
MTK-10018	Copy Batch Size should be in the range from 1 to 1024.
MTK-10019	Invalid number for Copy Batch Size, use a number > 0.
MTK-10020	Fetch Size should be greater than 0.
MTK-10021	Invalid number for Fetch Size, use a number > 0.
MTK-10022	One or more of the command-line arguments is invalid.

Error Code	Message and Resolution
MTK-10023	The -targetDBVersion value should be specified using <i>major.minor</i> pattern.
MTK-10024	The -targetDBVersion can only be used with -offlineMigration mode.
MTK-10025	Options (-constraints -indexes -triggers -tables -views -sequences -procs -funcs -packages -synonyms) cannot be used with multiple schemas option.
MTK-10026	The -allSchemas option should be specified as the last option in the command-line options list.
MTK-10027	You have specified invalid command-line arguments. Run 'runMTK -help' to see the usage details.
MTK-10028	The -replaceNullChar cannot be specified with -schemaOnly option.
MTK-10029	The -nullReplacementChar can only be specified with -replaceNullChar option.
MTK-10030	The -ignoreCheckConstFilter can only be given with -constraints option.
MTK-10031	The -enableConstBeforeDataLoad can only be given with -truncLoad option.
MTK-10032	The retryCount value should be greater than 0.
MTK-10033	Invalid number for retryCount, use a number > 0.
MTK-10034	The loaderCount value should be greater than 0.
MTK-10035	Invalid number for loaderCount, use a number > 0.
MTK-10036	Cannot use singleDataFile option for offline data migration in COPY format.
MTK-10037	The offline data migration in COPY format is supported only when PostgreSQL or EnterpriseDB (EDB Postgres Advanced Server) is the target database.
MTK-10038	The <i>\$SCHEMA</i> cannot be used as schema name in <i>\$DATABASE</i> . Choose a different schema name via -targetSchema option.
MTK-10039	Log file size should be greater than 0.
MTK-10040	Log file count should be greater than or equal to 0.
MTK-10041	Offline migration can only be used with schema only option.
MTK-10042	The copyViaDBLinkOra option can only be used for copying data from Oracle to EnterpriseDB.
MTK-10043	The -recreateConst option can only be given with the -dataOnly option.
MTK-10044	The <i>\$DATABASE</i> database type is not supported by Migration Toolkit. Specify a valid database type string (i.e., EnterpriseDB, Postgres, Oracle, SQLServer, Sybase, or MySQL).
MTK-10045	The URL specified for the <i>\$DATABASE</i> database is invalid. Check the connectivity credentials.
MTK-10046	The -escapeKeywords value may be true or false.
MTK-10047	The -useOraCase option can only be used for migration from Oracle.

Error Code	Message and Resolution
MTK-10048	The Postgres exported snapshot id is invalid.
MTK-10049	The URL specified for the Oracle database is not supported by dblink_ora. Check the connectivity credentials and provide a valid URL.
MTK-10050	The schema <i>\$\$SCHEMA</i> not found on source database.
MTK-10051	The skipUserSchemaCreation option can only be used for user migration from Oracle to EnterpriseDB.
MTK-10052	The -truncLoad option cannot be used with offline data migration.
MTK-10053	The -parallelLoadRowLimit should be greater than 0.
MTK-10054	Invalid number for -parallelLoadRowLimit, use an integer in the range of 1 to \$MAX_LIMIT.
MTK-10055	The -parallelLoadRowLimit cannot be specified with -schemaOnly option.
MTK-10056	The -loaderCount cannot be specified with -schemaOnly option.
MTK-10057	The -tableLoaderLimit should be greater than 0.
MTK-10058	Invalid number for -tableLoaderLimit, use an integer in the range of 1 to \$MAX_LIMIT.
MTK-10059	The -tableLoaderLimit cannot be specified with -schemaOnly option.
MTK-10060	The parallel loading options (-loaderCount, -parallelLoadRowLimit, -tableLoaderLimit) are only supported for Oracle, EnterpriseDB and PostgreSQL source databases.
MTK-10061	The -fastCopy cannot be specified with table-level parallel loading context (-tableLoaderLimit > 1).
MTK-10062	The '-copyViaDBLinkOra' option cannot be used with offline data migration.
MTK-10063	The permutation of source database type <i>\$DATABASE</i> and target database type <i>\$DATABASE</i> is not supported.
MTK-10064	The '-tables' and '-excludeTables' options cannot be specified at the same time. Please specify one of the options.
MTK-10065	The '-views' and '-excludeViews' options cannot be specified at the same time. Please specify one of the options.
MTK-10066	The table name should be schema qualified when used in multiple schemas context.
MTK-10067	The '-sequences' and '-excludeSequences' options cannot be specified at the same time. Please specify one of the options.
MTK-10068	The '-procs' and '-excludeProcs' options cannot be specified at the same time. Please specify one of the options.
MTK-10069	The '-funcs' and '-excludeFuncs' options cannot be specified at the same time. Please specify one of the options.
MTK-10070	The '-packages' and '-excludePackages' options cannot be specified at the same time. Please specify one of the options.
MTK-10071	The '-synonyms' and '-excludeSynonyms' options cannot be specified at the same time. Please specify one of the options.
MTK-10072	The '-queues' and '-excludeQueues' options cannot be specified at the same time. Please specify one of the options.

Error Code	Message and Resolution
MTK-10073	The connRetryCount value should be between 0 and 50.
MTK-10074	Invalid value for connRetryCount, use a number between 0 and 50.
MTK-10075	The connRetryInterval value should be between 0 and 900.
MTK-10076	Invalid value for 'connRetryInterval', use a number between 0 and 900.
MTK-10077	The '-pgIdleTxSessionTimeOut' value should be greater than 0.
MTK-10078	The argument '-\$OPTION' cannot start with '--'.
MTK-10079	The last argument should be the schema name without any option context.
MTK-10080	The value <i>\$VALUE</i> is specified without any option.
MTK-10081	The command-line argument '-\$OPTION' is not supported.
MTK-10082	The value of the option '-\$OPTION' can be either 'true' or 'false'.
MTK-10083	The value of the option '-\$OPTION' can be either 'on' or 'off'.
MTK-10084	The '-\$OPTION' option cannot be used without a value.
MTK-10085	The source schema name should not start with '-' or '--'.

Class 11 - Configuration issues

This class represents invalid configuration settings in the `toolkit.properties` file or in any other configuration file used by the Migration Toolkit.

Error code	Message and resolution
MTK-11000	All the errors that relate to this class and don't have a specific error code binding use this error code.
MTK-11001	The properties file containing table filter clause cannot be loaded.
MTK-11002	The custom type mapping file couldn't be loaded. Reason: <i>\$REASON</i> .
MTK-11003	The custom type mapping file contains an invalid mapping entry.
MTK-11004	The custom type mapping file doesn't contain any mapping entry.
MTK-11005	Connectivity information for source database is not available.
MTK-11006	Connectivity information for target database is not available.
MTK-11007	Unable to create the logs directory. <i>\$PATH</i> .

Error code	Message and resolution
MTK-11008	The properties file containing table columns filter list cannot be loaded.
MTK-11009	Error Connecting Database <i>\$DATABASE</i> .
MTK-11010	Error loading <i>\$DBLINKORA_FILE</i> file.
MTK-11011	Error while loading DBLink Ora module. <i>\$DBLINKORA_MODULE</i> . Verify that dblink_ora is installed/configured on target EnterpriseDB server. Please see the Database Compatibility for Oracle Developer's Guide for more information about installing and configuring the dblink_ora module.
MTK-11012	Error while loading given DBLink_Ora module. <i>\$DBLINKORA_MODULE</i> . Verify that dblink_ora is installed/configured on target EnterpriseDB server. Please see the Database Compatibility for Oracle Developer's Guide for more information about installing and configuring the dblink_ora module.
MTK-11013	Could not load DBLinkOra Module.
MTK-11014	Error connecting to DBLinkOra.
MTK-11015	The connection credentials file <i>\$TOOLKIT_PROP_FILE</i> is not secure and accessible to group/others users. This file contains plain passwords and should be restricted to Migration Toolkit owner user only.
MTK-11016	Data loader task failed to initialize. Reason: \$REASON

Class 12 - Insufficient privileges

This class represents insufficient privilege errors for loading database user/role information.

Error code	Message and resolution
MTK-12000	All the errors that relate to this class and don't have a specific error code binding use this error code.
MTK-12001	The user/role migration failed due to insufficient privileges. Grant the user SELECT privilege on the following Oracle catalogs: DBA_ROLES, DBA_USERS, DBA_TAB_PRIVS, DBA_PROFILES, DBA_ROLE_PRIVS, ROLE_ROLE_PRIVS, DBA_SYS_PRIVS.
MTK-12002	The migration of privileges failed due to insufficient privileges. Grant the user SELECT privilege on the following Oracle catalog: dba_tab_privs.

Class 13 - Unsupported features

This class represents errors related to the migration of unsupported objects and clauses. Either the target database has not implemented the given object, or the Migration Toolkit doesn't handle its migration.

Error code	Message and resolution
MTK-13000	All the errors that relate to this class and don't have a specific error code binding use this error code.
MTK-13001	Skipping index as Oracle does not allow multiple indexes against same column list.
MTK-13002	For <i>\$DATABASE</i> views migration is not supported.
MTK-13003	For <i>\$DATABASE</i> roles migration is not supported.

Error code	Message and resolution
MTK-13004	You cannot migrate triggers, sequences, procedures, functions, packages and synonyms for MySQL, SQL Server, and Sybase databases.
MTK-13005	You cannot migrate sequences and packages for SQL Server database.
MTK-13006	Given trigger is not migrated, the trigger is owned by a different user schema and cross-schema triggers are not supported by EnterpriseDB.
MTK-13007	The given trigger is not migrated, the trigger has WHEN clause which is not supported by EnterpriseDB.
MTK-13008	Skipping Database Link <i>\$DATABASE_LINK</i> . EnterpriseDB currently does not support this type of Database Link.
MTK-13009	Warning! Skipping migration of trigger <i>\$TRIGGER</i> , currently non-table triggers are not supported in target database.
MTK-13010	You cannot migrate procedures, packages, synonyms and database links to PostgreSQL database.
MTK-13011	Domain objects are not supported in target database.
MTK-13012	Rules are not supported in <i>\$DATABASE</i> database.
MTK-13013	The database type is not supported.
MTK-13015	<i>\$TYPE</i> is Not Supported by COPY.
MTK-13016	The migration to PostgreSQL is supported only when Oracle or PostgreSQL is the source database.
MTK-13017	Groups are not supported in <i>\$DATABASE</i> database.
MTK-13018	Profile migration is not supported in <i>\$SRC_DB</i> to <i>\$TARGET_DB</i> permutation.
MTK-13019	Cannot migrate unknown Profile <i>\$PROFILE</i> .
MTK-13020	Profiles cannot be migrated to database version <i>\$VERSION</i> .
MTK-13021	Password Profile verify function <i>\$MY_VERIFICATION_FUNCTION</i> must be explicitly migrated.
MTK-13022	Advanced Queues migration is not supported from <i>\$SRC_DB</i> database to <i>\$TARGET_DB</i> database.
MTK-13023	Advanced Queues cannot be migrated to database version <i>\$VERSION</i> .
MTK-13024	The INTERVAL partition is not supported in <i>\$DATABASE</i> , the table will be migrated without INTERVAL definition.
MTK-13025	Warning! User profile migration is not supported in target database version <i>\$VERSION</i> . The profile " <i>\$PROFILE</i> " for user " <i>\$USER</i> " will be skipped.
MTK-13026	Object Type migration is not supported from <i>\$SRC_DB</i> to <i>\$TARGET_DB</i> .
MTK-13027	Migration is not supported from <i>\$DATABASE</i> database for PDB Admin schemas/users with PDB_DBA role.
MTK-13028	PostgreSQL 10 and earlier versions do not support stored procedures.

Class 14 - Missing objects

This class represents failures to find metadata information in the source database.

Error code	Message and resolution
MTK-14000	All the errors that relate to this class and don't have a specific error code binding use this error code.
MTK-14001	One or more tables are missing from the source <i>\$DATABASE</i> database.
MTK-14002	One or more tables couldn't be found in the source <i>\$DATABASE</i> database. With <i>-tables</i> mode, the table name should be in uppercase unless it is case-sensitive.
MTK-14003	One or more users couldn't be found in the source <i>\$DATABASE</i> database. With <i>-users</i> mode, the user name should be in uppercase unless it is case-sensitive.

Class 15 - Schema migration

The class represents migration issues related to nonprocedural database objects such as tables, constraints, indexes, synonyms, views, users, and roles.

Error code	Message and resolution
MTK-15000	All the errors that relate to this class and don't have a specific error code binding use this error code.
MTK-15001	Error creating constraint <i>\$CONSTRAINT</i> .
MTK-15002	Error creating role <i>\$ROLE</i> .
MTK-15003	Error granting given privilege <i>\$PRIVILEGE</i> on <i>\$OBJECT</i> to <i>\$USER</i> .
MTK-15004	Error granting <i>\$ROLE</i> to <i>\$USER</i> .
MTK-15005	Error granting privilege <i>\$REASON</i> .
MTK-15006	Error creating user <i>\$USER</i> .
MTK-15007	Error creating index <i>\$INDEX</i> .
MTK-15008	Error creating view <i>\$VIEW</i> .
MTK-15009	Error creating materialized view <i>\$MVIEW</i> .
MTK-15010	Error creating public synonym <i>\$SYNONYM</i> .
MTK-15011	Error creating private synonym <i>\$SYNONYM</i> .
MTK-15012	Error creating sequence <i>\$SEQUENCE</i> .
MTK-15013	Error creating table <i>\$TABLE</i> .
MTK-15014	Error creating database link <i>\$DATABASE_LINK</i> .

Error code	Message and resolution
MTK-15015	Error creating given object type <i>\$OBJECT_TYPE</i> .
MTK-15016	The table <i>\$TABLE</i> could not be created in <i>\$DATABASE</i> database.
MTK-15017	The linked schema <i>\$LINKED_SCHEMA</i> doesn't exist in the target database. Create the schema and then retry.
MTK-15018	Error creating domain <i>\$DOMAIN</i> .
MTK-15019	Error creating custom data type <i>\$DATA_TYPE</i> .
MTK-15020	Error creating membership for group <i>\$GROUP</i> .
MTK-15021	Table name <i>\$TABLE</i> does not have a schema qualifier. With multiple schema migration context, each table should be schema qualified.
MTK-15022	Schema qualifier <i>\$SCHEMA</i> does not match the schema list.
MTK-15023	The table metadata information is not available.
MTK-15024	Tables list is not initialized yet.
MTK-15025	Error while getting database metadata information.
MTK-15026	Error creating group <i>\$GROUP</i> .
MTK-15027	Error creating Profile <i>\$PROFILE</i> .

Class 16 - Procedural language migration

The class represents migration issues related to database objects that are based on procedural languages such as procedures, functions, packages, anonymous blocks, triggers, and rules.

Error code	Message and resolution
MTK-16000	All the errors that relate to this class and don't have a specific error code binding use this error code.
MTK-16001	Error Creating Trigger <i>\$TRIGGER</i> .
MTK-16002	Error Creating Rule <i>\$RULE</i> .
MTK-16003	Error Creating Package Spec <i>\$PACKAGE</i> .
MTK-16004	Error Creating Package Body <i>\$PACKAGE</i> .
MTK-16005	Package Body is invalid, skipping... Note: This error message also appears when a package specification is successfully migrated, but there is no corresponding package body in the source database. In this case, the package specification should function properly in the target database despite the appearance of the error message.
MTK-16006	Error Creating Procedure <i>\$PROCEDURE</i> .
MTK-16007	Error Creating Function <i>\$FUNCTION</i> .

Error code	Message and resolution
MTK-16008	Error Creating Anonymous Block <i>\$BLOCK</i> .
MTK-16009	Error creating Queue <i>\$QUEUE</i> .

Class 17 - Data loading

This class represents errors that can occur while copying data from the source to the target database.

Error code	Message and resolution
MTK-17000	All the errors that relate to this class and don't have a specific error code binding use this error code.
MTK-17001	Error Loading Data into Table: <i>\$TABLE</i>
MTK-17002	Encoding Conversion encountered some characters that will be loaded using Bulk Inserts instead.
MTK-17003	Error in copy tables <i>\$REASON</i> .
MTK-17004	Invalid DataType.
MTK-17005	This Table Contains CLOB data, Marked for Bulk Insert Loading.
MTK-17006	One of the table loader threads failed while copying data for table <i>\$TABLE</i> .

11 FAQ

Does Migration Toolkit support the migration of packages?

Migration Toolkit supports the migration of packages from an Oracle database into EDB Postgres Advanced Server. See [Functionality overview](#) for information about the migration support offered by EDB Postgres Advanced Server.

Is there a way to transfer only the data from the source database?

Yes. Data transfer is supported as part of an online or offline migration.

Does Migration Toolkit support migration of tables that contain data of the CLOB data type?

Migration Toolkit does support migration of tables containing data of the CLOB type.

Does EDB Postgres Advanced Server support the enum data type?

EDB Postgres Advanced Server doesn't currently support the enum data type but will support it in future releases. Until then, you can use a check constraint to restrict the data added to an EDB Postgres Advanced Server database. A check constraint defines a list of valid values that a column can take.

The following code sample includes a simple example of a check constraint that restricts the value of a column to one of three dept types; `sales`, `admin` or `technical`.

```
CREATE TABLE emp
(
  emp_id INT NOT NULL PRIMARY
KEY,
  dept VARCHAR(255) NOT NULL,

  CHECK (dept IN ('sales', 'admin',
'technical'))
);
```

If you test the check constraint by entering a valid dept type, the INSERT statement works without error:

```
test=# INSERT INTO emp VALUES (7324,
'sales');

INSERT 0 1
```

If you try to insert a value not included in the constraint (`support`), EDB Postgres Advanced Server returns an error:

```
test=# INSERT INTO emp VALUES (7325,
'support');

ERROR: new row for relation "emp" violates check constraint
"emp_dept_check"
```

Does EDB Postgres Advanced Server support materialized views?

Postgres doesn't support materialized views compatible with Oracle databases. To set up a materialized view/summary table in Postgres you must manually create the triggers that maintain the summary table. Automatic query rewrite isn't currently supported. You must make the application aware of the summary table's existence.

When I try to migrate from a MySQL database that includes a TIME data type, I get the following error: Error Loading Data into Table: Bad format for Time. Does Postgres support MySQL ``TIME`` data types?

Postgres doesn't have a problem storing `TIME` data types as long as the value of the hour component isn't greater than 24.

Unlike Postgres, the MySQL `TIME` data type allows you to store a value that represents either a `TIME` or an `INTERVAL` value. A value stored in a MySQL `TIME` column that represents an `INTERVAL` value can potentially be out of the accepted range of a valid Postgres `TIMESTAMP` value. If, during the migration process, Postgres encounters a value stored in a `TIME` data column that it perceives as out of range, it returns an error.

What use cases does the connection retry capability support?

Database scope: The connection retry capability allows Migration Toolkit to reconnect to the target database. Retry attempts for issues with the source database are currently not supported.

Migration scope: This capability allows Migration Toolkit to retry migrating data. Retry attempts for issues with the schema migration are currently not supported.

Modality scope: This reconnection capability is available with the data migration mode (`-dataOnly`). It's also available for the data migration step when you run a migration without specifying either the `-dataOnly` and `-schemaOnly` options.

What do I do if there's an error during a data migration that results in the data for one or more of my tables not being fully migrated?

If Migration Toolkit can't reconnect successfully and one or more of the tables weren't fully migrated, restart the entire data migration (with the `-dataOnly` option) or configure Migration Toolkit to migrate only the tables that weren't fully migrated in the previous run. You can use `connRetryCount`, `connRetryInterval`, and `abortConnOnFailure` to alter the retry configuration options.

If the source database was accepting write transactions while the previous Migration Toolkit data migration was in process, perform the full (all tables) data migration again to ensure that data migrated to the destination database is in a consistent state.

If you can confirm that no write transactions were performed on the source database while the previous migration was being performed, then it's safe to migrate only those tables that weren't fully migrated.

What do I do if there's an error during a schema migration that results in not all of my schema objects being fully migrated?

Unfortunately, reconnection for schema migration errors isn't supported at this time. If Migration Toolkit can't migrate all schemas, restart the entire schema migration again (with the `-schemaOnly` option), or configure Migration Toolkit to migrate only the schemas that weren't fully migrated in the previous run.