



# EDB

## EDB Postgres™ Advanced Server

*Release 13*

### Release Notes

Nov 19, 2020

### Contents

<b>1</b>	<b>Installers and Documentation</b>	<b>2</b>
<b>2</b>	<b>Supported Platforms</b>	<b>3</b>
<b>3</b>	<b>Component Certification</b>	<b>3</b>
<b>4</b>	<b>EDB Postgres Advanced Server v13 Features</b>	<b>4</b>
4.1	Community PostgreSQL 13 Updates . . . . .	6
4.2	How to Report Problems . . . . .	6

---

With this release of EDB Postgres Advanced Server 13, EnterpriseDB continues to lead as the only world-wide company to deliver innovative and low cost open-source-derived database solutions with commercial quality, ease of use, compatibility, scalability, and performance for small or large-scale enterprises.

EDB Postgres Advanced Server 13 is built on open-source PostgreSQL 13, which introduces myriad enhancements that enable databases to scale up and scale out in more efficient ways. PostgreSQL 13 has significant performance improvements, which includes an improved indexing and lookup system, a level

up on query planning when using extended statistics, improved performance for queries that use aggregates or partitioned tables, and providing more ways to monitor activity within a PostgreSQL database. And along with highly requested features like parallelized vacuuming, incremental sorting, PostgreSQL 13 has a multitude of other new features and capabilities.

EDB Postgres Advanced Server 13 adds a number of new features that will delight developers and DBAs alike, including:

- Automatic LIST Partitioning
- PARTITION number or SUBPARTITION number clauses while creating a HASH partition table
- Forward declarations in the package body
- CREATE INDEX syntax contains column name and constant
- USING INDEX clause in CREATE TABLE and ALTER TABLE
- EDB Loader handles unique constraint violations
- STATS\_MODE aggregate function
- Support for utl\_http.end\_of\_body exception
- EDB Loader supports any connection parameters with the -c or CONNSTR option

## 1 Installers and Documentation

EDB Postgres Advanced Server v13 is packaged and delivered as interactive installers for Windows; visit the EnterpriseDB website:

<https://www.enterprisedb.com/advanced-downloads>

RPM Packages are available for Linux from:

<https://repos.enterprisedb.com/>

Debian/Ubuntu Packages are available for download from:

<https://repos.enterprisedb.com/>

To request the credentials required to access EDB repositories, visit:

<https://www.enterprisedb.com/repository-access-request>

Documentation is provided on the EDB website, visit:

<https://www.enterprisedb.com/edb-docs/>

## 2 Supported Platforms

EDB Postgres Advanced Server v13 installers support 64 bit Linux and Windows server platforms. The Advanced Server 13 RPM packages are supported on the following 64-bit Linux platforms:

- Red Hat Enterprise Linux (x86\_64) 7.x and 8.x
- CentOS (x86\_64) 7.x and 8.x
- OEL (x86\_64) 7.x and 8.x
- PPC64LE 8 running CentOS/RHEL 7.x

The EDB Postgres Advanced Server 13 native packages are supported on the following 64-bit Linux platforms:

- Debian 9.x and 10.x
- Ubuntu 18.04 and 20.04

Graphical installers are supported on the following 64-bit Windows platforms:

- Windows 2019
- Windows Server 2016

Additional information about supported platforms is available on the EDB website:

<https://www.enterprisedb.com/services-support/edb-supported-products-and-platforms>

## 3 Component Certification

The following components are included in the EDB Postgres Advanced Server v13 release:

- EPAS 13.1.4
- BART 2.6
- Cloneschema 1.14
- Connectors JDBC 42.2.12.3, ODBC 12.02.0000.02, .NET 4.1.5.1, OCL 13.1.4.1
- Edb-Modules 1.0
- EDBPlus 39.0.0
- EFM 4.0
- Hdfs\_fdw 2.0.7
- Mongo\_fdw 5.2.8
- MySQL\_fdw 2.5.5
- MTK 54.0.0
- Parallel Clone 1.8

- PEM 7.16
- PgAdmin 4.27
- PgAgent 4.2.0
- PgBouncer 1.14
- PgPool 4.1.2
- PostGIS 3.0.2
- Procedural Language Packs – PL/Perl 5.26, PL/Python 3.7, PL/TCL 8.6
- Slony 2.2.8

## 4 EDB Postgres Advanced Server v13 Features

The major highlights of this release are :

- Automatic LIST Partitioning is an extension to LIST partitioning that allows a database to automatically create a partition for any new distinct value of the LIST partitioning key. A new partition is created when data is inserted into the LIST partitioned table and the inserted value does not match any of the existing table partitions.
- Additional `smallint`, `int`, `bigint`, and `numeric` variants in the `MEDIAN` aggregate function.
- Added support for `CREATE INDEX` syntax that contains a column name and number, i.e. `(col_name, 1)`.
- Added support to alter the owner of the directory - `ALTER DIRECTORY <dir_name> OWNER TO <role_name>`.
- CSV and XML audit logs have been made consistent. It has been observed that there is a difference in terms of number of fields in the CSV and XML formatted audit logs. Now the audit log is consistent across both formats.
- Default behaviour for `dbms_output` made compatible with Redwood. By default, the `dbms_output` behaviour is `off` in Oracle; in EPAS it is always `on`. In this release, EPAS adds a GUC to control the default behaviour for `DBMS_OUTPUT` package.
- You can log the number of processed statements with `edb_log_every_bulk_value`. Currently, during bulk execution EPAS did not identify the number of rows processed when `edb_log_every_bulk_value` was `off`. This feature allows EPAS to log it both in the audit log file and server log file. This will help you analyze the logs for such cases.
- Previously, EDB Loader would abort the whole operation if any record insertion fails due to a unique constraint violation. This is fixed by using speculative insertion to insert rows. This behavior is enforced if `handle_conflicts` (a new parameter) is `true` and indexes are present.
- `SYSDATE` now behaves in a more compatible manner. `SYSDATE` output changes at every nesting level, so multiple copies of `SYSDATE` in the same SQL query will return the same value each time. `SYSDATE` changes on successive statements in the same procedure, or within a nested function call.

- `PARALLEL [n] | NOPARALLEL` option for `CREATE TABLE` and `INDEX`. Advanced Server now supports the `PARALLEL [n] | NOPARALLEL` clause in the `CREATE TABLE`, `ALTER TABLE`, `CREATE INDEX`, and `ALTER INDEX` commands to enable or disable parallelism on an index or a table.
- `PARTITION [n]` or `SUBPARTITION [n]` while creating a table. This feature allows you to automatically create [n] hash partitions at a subpartition level.
- EDB Loader supports any connection parameters. You can use the EDB Loader `-c` or `CONNSTR` options to specify any connection parameters supported by libpq. This includes SSL connection parameters.
- The `STATS_MODE` function takes a set of values as an argument and returns the value that occurs with the highest frequency. If multiple values appear with the same frequency, the `STATS_MODE` function arbitrarily chooses the first value and returns only that one value.
- Added support for `DBMS_SQL` function/procedures (`DEFINE_COLUMN_LONG`, `COLUMN_VALUE_LONG` and `LAST_ERROR_POSITION`).
- Added support for function `to_timestamp_tz()`.
- Added support for function/procedure specification inside package body.
- Added support for FM format in `to_number` function.
- Added support for AES192 and AES256 in the `DBMS_CCRYPTO` package.
- Added support for the spell mode in `to_char(timestampz, text)` function.
- Allow creating a compound trigger having `WHEN` clause with `NEW/OLD` variables and `STATEMENT` level triggering events. Enhanced Redwood compatible view.
- Log matched line of `pg_hba.conf` on successful client authentication. Enhanced `pg_catcheck` to test for and raise an error if a relation's `relfilenode` is missing from the `data` directory.
- Added support for `utl_http.end_of_body` exception. This feature declares the `end_of_body` exception into `utl_http` package and throws the same from `read_line`, `read_text`, and `read_raw` package procedures when no data is left in the response body.
- The `UNIQUE` and `PRIMARY KEY` constraint clauses now have a `CREATE INDEX` statement. This new syntax allows users to specify explicit index details like `fillfactor`, etc. Columns specified in the constraint and the columns specified in the index must be the same.

For information about Advanced Server features that are compatible with Oracle databases, see the following guides:

- *Database Compatibility for Oracle Developer's SQL Guide*
- *Database Compatibility for Oracle Developer's Reference Guide*
- *Database Compatibility for Oracle Developer's Built-in Package Guide*
- *Database Compatibility for Oracle Developer's Tools and Utilities Guide*
- *Database Compatibility Table Partitioning Guide*
- *Database Compatibility Stored Procedural Language Guide*

## 4.1 Community PostgreSQL 13 Updates

EDB Postgres Advanced Server 13 integrates all of the community PostgreSQL 13 features. To review a complete list of changes to the community PostgreSQL project and the contributors names, see the PostgreSQL 13 Release Notes at:

<https://www.postgresql.org/docs/13/release-13.html>

## 4.2 How to Report Problems

If you experience any problems installing the new software please contact Technical Support at:

- Email: [support@enterprisedb.com](mailto:support@enterprisedb.com)
- Phone:

US: +1-732-331-1320 or 1-800-235-5891

UK: +44-2033719820

Brazil: +55-2139581371

India: +91-20-66449612