



PARTNER

Webinar

What's new in PostgreSQL 17


All our microphones are muted

Ask your questions in Q&A, not in the Chat

Use Chat for discussion, networking or applause

1

Improvements in data processing in large environments



Performance improvements

- ▶ **Supports large-scale memory**

The SLRU(Simple Least Recently Used) cache is divided into multiple banks, which is expected to reduce the lock range and improve cache search speed.

- ▶ **Streaming I/O**

Introduce an abstraction allowing relation data to be accessed as a stream of buffers, with an implementation that is more efficient than the equivalent sequence of ReadBuffer() calls. Now used with ANALYZE command.

- ▶ **Adaptive Radix Trees and TIDStore**

20% faster Vacuum on systems with low maintenance_work_mem or autovacuum_work_mem

Performance improvements

- ▶ **Parallel build of BRIN index**
The multiple worker processes can now be used to build BRIN indexes.
- ▶ **Better Parallel DISTINCT Processing**
The Gather Merge plan is now available for parallel execution of SELECT statements with DISTINCT clauses.
- ▶ **GROUP BY optimization**
If the GROUP BY clause with multiple columns is not related to the sort order, the sorting process can now be reduced by swapping the columns.
- ▶ **Better IS [NOT] NULL Handling**
The optimizations have been implemented to reduce unnecessary evaluation of IS NULL and IS NOT NULL clauses.

Performance improvements

▶ UNION optimization

The Merge Append plan can now be used for lookups that contain a UNION clause in the subquery. Until now, UNION queries have often been suboptimal as the planner has only ever considered using an Append node and making the results unique by either using a Hash Aggregate, or by Sorting the entire Append result and running it through the Unique operator. Both of these methods always require reading all rows from the union subqueries.

▶ GiST Index optimization

The Incremental Sort plan is now available for GiST and SP-GiST indexes. This was previously possible only with btree indexes.

▶ Faster B-tree index scans

In Postgres 17 for queries that involve IN lists or other cases where multiple array values are being passed to Postgres. Can deliver three times faster queries.

2


Improvements in server administration



What's new in PostgreSQL 17

Incremental Backup


- ▶ The `pg_basebackup` is now able to do incremental backup. There is a new binary called `pg_combinebackup` that merges the base and incremental backups.
- ▶ To use incremental backup, the WAL summarize feature must be enabled.
- ▶ When this feature is enabled, a "walsummarizer" process is invoked in the instance and summary data of the WAL file is stored in the `${PGDATA}/pg_wal/summaries` directory.
- ▶ The WAL summary file is automatically removed when the parameter `wal_summary_keep_time` (default value '10d') is exceeded.



What's new in PostgreSQL 17

Configuration File


- ▶ The maximum length of tokens in the `pg_hba.conf` and `pg_ident.conf` files has been extended from 256 bytes to unlimited.
- ▶ In complex environments configured with LDAP, there were problems with exceeding the previous maximum length.



What's new in PostgreSQL 17

Logging

- ▶ **Recovery related logging**
The logs at the start and end of the recovery process have been added.
- ▶ **Password change**
A new function, `PQchangePassword()`, has been added to `libpq`. The password is not sent to the server in cleartext because it is "encrypted" on the client side. This is good because it ensures the cleartext password is never known by the server, and therefore won't end up in logs, `pg_stat` displays, etc.




New privilege and event trigger

- ▶ The login event trigger is now available that executes upon successful authentication. Useful for auditing.
- ▶ The MAINTAIN privilege has been added to perform maintenance operations such as VACUUM, ANALYZE, and REINDEX statements. The pg_maintain predefined role has been added to allow maintenance operations on objects for all users.
 - ▶ A role granted this privilege can execute VACUUM, ANALYZE, REINDEX, REFRESH MATERIALIZED VIEW, CLUSTER, and LOCK TABLE statements on the target relation.
 - ▶ GRANT MAINTAIN ON table TO user ;
 - ▶ The pg_maintain predefined role has been added. This role is able to grant MAINTAIN privilege on all relations.

What's new in PostgreSQL 17

New views

- ▶ The `pg_stat_checkpoint` view has been added. Stats moved from `pg_stat_bgwriter`.
- ▶ The `pg_wait_events` view provides the names and descriptions of registered wait events



Reindexdb

- ▶ The `--all` and other options can now be used at the same time.

```
reindexdb --all --schema=public
```

- ▶ The `--jobs` and `--index` options can now be used at the same time. Multiple indexes on different tables can be processed in parallel.

```
reindexdb --jobs=2 --index=idx_data1 --index=idx_data2
```

New GUCs


▶ `allow_alter_system`

This parameter determines whether the ALTER SYSTEM statement is allowed to be executed. The default value is 'on', which allows the ALTER SYSTEM statement to be executed. This parameter itself cannot be changed by the ALTER SYSTEM statement.

▶ `transaction_timeout`

Specify the maximum transaction execution time in milliseconds. If a timeout occurs, the session is forcibly disconnected. The default value is 0, which means no timeout will occur.

```
SET transaction_timeout = '10s' ;
```



Replication

▶ The `pg_createsubscriber`


- ▶ The new command `pg_createsubscriber` has been added. This command converts a streaming replication standby server to a logical replication standby server. The major advantage of using this command is to reduce the initial data copy overhead that occurs when setting up a logical replication environment.

▶ Streaming Replication Enhancements

- ▶ Data to be transferred to the standby instance is now retrieved from the WAL buffer if possible. If available, read directly from WAL buffers, avoiding the need to go through the filesystem. Only for physical replication for now, but can be expanded to other callers. In preparation for replicating unflushed WAL data.


Replication

- ▶ **Logical Replication Slot Synchronization**
 - ▶ The replication slot information used for logical replication can now be synchronized to the standby server. The "slotsync worker" process is launched.
- ▶ **Pg_upgrade**
 - ▶ The subscriber (SUBSCRIPTION) state can now be preserved. In previous versions, only metadata was preserved.
 - ▶ The logical replication slots are now recreated on the destination cluster during upgrade.



3


Incompatibility



What's new in PostgreSQL 17

MERGE statement


- ▶ The SELECT privilege on the target table is now required even when specifying the DO NOTHING clause in a MERGE statement. This specification will be backported to PostgreSQL 15 and later.



What's new in PostgreSQL 17

pg_walfile_name / pg_walfile_name_offset

- ▶ In previous versions, these functions returned the previous segment number when the LSN was on a segment boundary. It has been changed to always return the current segment number of the LSN. This fix is also reflected in previous versions.



Modified Parameters

The following parameters have been changed.

Parameter name	Changes
wal_sync_method	The configuration value fsync_writethrough is no longer available in Microsoft Windows environments.
log_connections	Information on trust connections is now output.
log_replication_commands	Replication slot addition information is now output.

Removed parameters

The following parameters have been removed

Parameter name	Reason
<code>db_user_namespace</code>	It was removed because it was determined that there were few users
<code>old_snapshot_threshold</code>	Removed due to accuracy and performance issues. It is a desirable feature, and improved implementations may be available in the future.
<code>trace_recovery_messages</code>	Removed because it can be replaced by <code>pg_waldump</code> command, etc.

4


SQL



What's new in PostgreSQL 17

MERGE statement


- ▶ The MERGE statements can now specify the RETURNING clause to return the data of updated tuples.
- ▶ A special function `merge_action()` has been added to determine the type (INSERT/UPDATE/DELETE) of the tuple returned by the RETURNING clause.



What's new in PostgreSQL 17

JSON Enhancements

- ▶ Among others the `JSON_TABLE` function has been added to convert JSON data to the relational view. This allows JSON data to be converted into a relational table format temporarily, facilitating more complex queries and data manipulation.



Partitioning


- ▶ Support for IDENTITY column
 - ▶ The INSERT statements are now supported on partitions of partitioned tables with IDENTITY columns.

```
CREATE TABLE test(column1 INT, column2 INT GENERATED ALWAYS AS IDENTITY)
PARTITION BY RANGE(column1) ;
CREATE TABLE test_p1 PARTITION OF test FOR VALUES FROM (0) TO (5000) ;
INSERT INTO test(column1) VALUES (1) ;
INSERT INTO test_p1(column1) VALUES (2) ;
ERROR: null value in column "column2" of relation "test_p1" violates not-null
constraint
```


What's new in PostgreSQL 17


Merge/Split partitions

- ▶ The ALTER TABLE statement for the RANGE / LIST partition table is now able to merge multiple partitions into a single partition or vice versa.
- ▶ This patch was reverted during beta period.



5

Misc



Built-in Locale Provider

- ▶ A 'built-in' locale provider has been added; it is independent of external libraries such as `libc` or `icu`. The `BUILTIN` keyword can now be specified in commands and SQL statements that specify a locale provider.
- ▶ Currently, the only locales offered by this locale provider are `C` and `C.UTF-8`. The built-in locale `C.UTF-8` provides advantages versus the `libc` locale, such as faster sorting, faster case conversion, and is platform-independent.


```
CREATE DATABASE zabbixdb LOCALE_PROVIDER BUILTIN LOCALE 'C.UTF-8' TEMPLATE template0 ;
initdb --locale-provider=builtin --builtin-locale=C.UTF8 -D data
```

What's new in PostgreSQL 17

\watch meta-command

- ▶ The iterative operation terminates when the specified number of tuples are no longer output.


```
\watch 10 min_rows=1
```



What's new in PostgreSQL 17

EXPLAIN statement

- ▶ The EXPLAIN statement with the MEMORY clause provides information on both the allocated memory and the memory that is actually in use.



6

Newly added parameters



Not a complete list of newly added parameters

Parameter name	Description (context)	Default value
allow_alter_system	Allow execution of ALTER SYSTEM statement (sighup)	on
event_triggers	Enable event triggers (superuser)	on
standby_slot_names	Replication slot name for standby database (sighup)	"
summarize_wal	Output WAL summary (sighup)	off
sync_replication_slots	Whether replication slots are synchronized (sighup)	off
transaction_timeout	Transaction execution timeout (user)	0

Not a complete list of newly added parameters






Parameter name	Description (context)	Default value
wal_summary_keep_time	WAL summary retention duration (sighup)	10d

What's new in PostgreSQL 17

Certified training and Support

Basic certified courses




Advanced certified courses

SQL Basics	Introduction to PostgreSQL	PostgreSQL Professional	PostgreSQL – administration and performance tuning	PostgreSQL – High Availability & Patroni
 <p>SQL Basics</p> <p>This training has been designed for people who want to get familiar with SQL. You will learn how SQL works and how to write proper SQL statements using practical examples that will be useful for your daily work.</p> <p>4 DAYS € 2,000 Excluding VAT</p> <p>Requirements: None Available online: Yes Certification: Yes</p> <p>Course dates: On request 13.–16. 1. 25 12.–15. 5. 25 3.–6. 11. 25</p> <p>REGISTER</p> <p>More info and dates</p>	 <p>Introduction to PostgreSQL</p> <p>This workshop has been designed for people who want to get familiar with SQL and PostgreSQL. You will learn how to use PostgreSQL and how to write proper SQL statements.</p> <p>4 DAYS € 2,000 Excluding VAT</p> <p>Requirements: None Available online: Yes Certification: Yes</p> <p>Course dates: On request 23.–26. 9. 24 27.–30. 1. 25 26.–29. 5. 25 24.–27. 11. 25</p> <p>REGISTER</p> <p>More info and dates</p>	 <p>PostgreSQL Professional</p> <p>This course provides a deep insight into advanced PostgreSQL topics like indexing, storage parameters, optimization, replication, monitoring and many more.</p> <p>3 DAYS € 1,500 Excluding VAT</p> <p>Requirements: None Available online: Yes Certification: Yes</p> <p>Course dates: On request 7.–9. 10. 24 10.–12. 2. 25 9.–11. 6. 25 8.–10. 12. 25</p> <p>REGISTER</p> <p>More info and dates</p>	 <p>PostgreSQL – administration and performance tuning</p> <p>This course is perfectly suitable for database administrators and sysadmins, dealing with topics related to administration and performance tuning.</p> <p>4 DAYS € 2,000 Excluding VAT</p> <p>Requirements: None Available online: Yes Certification: Yes</p> <p>Course dates: On request 21.–24. 10. 24 24.–27. 2. 25 23.–26. 6. 25</p> <p>REGISTER</p> <p>More info and dates</p>	 <p>PostgreSQL – High Availability & Patroni</p> <p>This course is intended for PostgreSQL users who are interested in fully automating high availability operations. The course first gives an overview of the general high availability landscape in PostgreSQL clusters and then focuses on installation, configuration and fully automated operation of very popular (open source) cluster manager "Patroni".</p> <p>3 DAYS € 1,500 Excluding VAT</p> <p>Requirements: advanced knowledge of PostgreSQL and OS Linux Available online: Yes Certification: Yes</p> <p>Course dates: On request 4.–6. 11. 24 10.–12. 3. 25 8.–10. 9. 25</p> <p>REGISTER</p> <p>More info and dates</p>

What's new in PostgreSQL 17

Certified training and Support

Expert certified courses

PostgreSQL in Kubernetes	Introduction to PostGIS	PostgreSQL - Migration tutorial
		
PostgreSQL in Kubernetes	Introduction to PostGIS	PostgreSQL - Migration tutorial
<p>This course provides an introduction to Kubernetes itself and to Kubernetes resources, which are needed in order to manage PostgreSQL.</p>	<p>This course provides an introduction to PostGIS and its most important features and capabilities. Along with gaining theoretical knowledge, participants will work with real-world datasets to deepen and reinforce their practical skills.</p>	<p>This training has been designed for people who want to switch to PostgreSQL.</p>
3 DAYS	3 DAYS	4 DAYS
€ 1,500 Excluding VAT	€ 1,500 Excluding VAT	€ 2,000 Excluding VAT
Requirements: None Available online: Yes Certification: Yes	Requirements: advanced knowledge of PostgreSQL Available online: Yes Certification: Yes	Requirements: advanced knowledge of PostgreSQL Available online: Yes Certification: Yes
Course dates: On request 18.–20. 11. 24 24.–26. 3. 25 22.–24. 9. 25	Course dates: On request 2.–4. 12. 24 7.–9. 4. 25 6.–8. 10. 25	Course dates: On request 16.–19. 12. 24 22.–25. 4. 25 20.–23. 10. 25
REGISTER	REGISTER	REGISTER
More info and dates	More info and dates	More info and dates

Certified training and Support



DBA – PostgreSQL

PostgreSQL is a powerful open-source object-relational database management system (ORDBMS). It is used for application development, data warehousing, analysis and other data-intensive tasks. Key features of PostgreSQL include a powerful engine, support for advanced data types and indexing methods, and support for stored procedures and triggers written in various programming languages, including PL/pgSQL, Tcl, and Python. Furthermore, PostgreSQL supports multiversion concurrency control (MVCC), allowing multiple users to access the same data simultaneously without conflicts, and offers robust data integrity and security support.

I AM INTERESTED IN THIS SERVICE →



Precise database management backed by experience



Reliable monitoring and notification system



Stability, availability and scalability



Questions?



What's new in PostgreSQL 17

Contact us:

Phone:



+420 800 244 442

Web:



<https://www.initmax.cz>

Email:



tomas.hermanek@initmax.cz

LinkedIn:



<https://www.linkedin.com/company/initmax>

Twitter:



<https://twitter.com/initmax>

Tomáš Heřmánek:



+420 732 447 184