



PostgreSQL 16: What is new?

By Charly Batista

Who am I?

I'm Charly Batista, currently the PostgreSQL Tech Lead at Percona and I'm passionate about databases and performance. I've been working with database and development for over 20 years, and participated in many projects, from the design of small sensors to count how many times a ball pass through a basketball hoop to complex database proxy applications.

Outside work I love sports, and as a good Brazilian I practice football (or soccer), volleyball, and surprisingly ping-pong. I'm also passionate about new cultures, their languages and traditions, and have lived and/or worked in few different countries around the globe.

Overview

PostgreSQL 16 contains many new features and enhancements

- Migration (be aware)
- Optimizer, General Queries and Performance
- Logical Replication
- Monitoring
- Privileges
- Server Configuration
- Other

Migration (be aware)

- Disallow NULLS NOT DISTINCT indexes for primary keys
- Change REINDEX DATABASE and reindexdb to not process indexes on system catalogs
- Rename server variable force_parallel_mode to debug_parallel_query
- Remove the ability to create views manually with ON SELECT rules
- Remove the server variables
 - vacuum_defer_cleanup_age
 - promote_trigger_file
 - lc_collate and lc_ctype

Optimizer, General Queries and Performance

- Allow parallelization of FULL and internal right OUTER hash joins
- Allow incremental sorts in more cases, including DISTINCT
- Add the ability for aggregates having ORDER BY or DISTINCT to use pre-sorted data
- During non-freeze operations, perform page freezing where appropriate

Logical Replication

- Optimizer, General Queries and Performance
- Allow logical replication initial table synchronization to copy rows in binary format
- Allow parallel application of logical replication
- Improve performance for logical replication apply without a primary key
- Allow logical replication subscribers to process only changes that have no origin
- Perform logical replication SELECT and DML actions as the table owner

Monitoring

- Add system view `pg_stat_io` view to track I/O statistics
- Record statistics on the occurrence of updated rows moving to new pages
- Add speculative lock information to the `pg_locks` system view
- Create subscription statistics entries at subscription creation time so `stats_reset` is accurate
- Correct the I/O accounting for temp relation writes shown in `pg_stat_database`
- Report stand-alone backends with a special backend type

Privileges

- Add predefined role `pg_create_subscription` with permission to create subscriptions
- Allow subscriptions to not require passwords
- Allow non-superusers to grant permissions using a granted-by user that is not the current user
- Add GRANT to control permission to use SET ROLE
- Prevent removal of superuser privileges for the bootstrap user
- Allow roles that create other roles to automatically inherit the new role's rights or the ability to SET ROLE to the new role

Server Configuration

- Improve performance of server variable management
- Allow autovacuum to more frequently honor changes to delay settings
- Allow the postmaster to terminate children with an abort signal
- Allow the server to reserve backend slots for roles with `pg_use_reserved_connections` membership
- Add `debug_io_direct` setting for developer usage

Other

- Many improvements in the JSON support
- Allow multiple libpq-specified hosts to be randomly selected
- Determine the default encoding from the locale when using ICU
- Add support for regular expression matching of user and database names in `pg_hba.conf`, and user names in `pg_ident.conf`
- Allow `VACUUM` and `vacuumdb` to only process TOAST tables
- Allow `postgres_fdw` to do aborts in parallel
- Make `ANALYZE` on foreign `postgres_fdw` tables more efficient

Conclusion

- It's very important to pay attention to the breaking changes in order to not break production when upgrading
- Always test, test, test, and test before going to production
- PostgreSQL 16 has many performance improvements in the Query Engine/Optimizer, including improvements in the indexes
- Long waiting changes in the logical replication to allow replicating from the replica
- Origin to the replicated data as start point for a possible cross-replication
- Improvements in the statistics collector and more data collected
- It's definitely worth testing PostgreSQL 16 and plan migration accordingly



Thank you all!

percona.com