



What's new in Postgres 17

Tomas Vondra, Microsoft

vondratomas@microsoft.com / tomas@vondra.me

<https://vondra.me>

OpenAlt 2024, Brno

Development Schedule

- June 2023 - branch 16
- July 2023 - CF1
- September 2023 - CF2
- November 2023 - CF3
- January 2024 - CF4
- March 2024 - CF5
- September 2024 - RC1
- 2024-09-26 release 17.0

Current Status

- 2681 commits
- 3958 files changed, 419663 insertions(+), 217841 deletions(-)

<https://www.postgresql.org/docs/release/17.0/>

pgconf.dev 2024

<https://www.youtube.com/@pgconfdev>

New Features

- DBA and administration
- SQL and developer
- Backup and replication
- Performance

Breaking changes I.

- building
 - Windows MSVC builds removed (we now have meson)
 - AIX support removed (but people are working on it)
 - --disable-thread-safety (sufficiently good threading expected)
- removed
 - adminpack (not needed, pgadmin III is EOL)
 - db_user_namespace (pretty much never used)
 - snapshot too old (known correctness/performance issues, might be reintroduced)

Breaking changes II.

- `pg_stat_bgwriter`
 - Removed `checkpoints_timed` & `req` => `pg_stat_checkpoint`
 - Removed `write_time` & `sync_time` => `pg_stat_checkpoint`
 - Removed `buffers_checkpoint`, `backend` & `fsync` => `pg_stat_io`
- `search_path` during maintenance ops
 - expression indexes & materialized views
 - Secured by default!
 - Must be explicit!
 - `ANALYZE`, `CLUSTER`, `REINDEX`, `REFRESH MATERIALIZED VIEW`, `VACUUM`

DBA and administration

DBA and administration

- transaction_timeout
 - similar to statement_timeout, partial replacement for "snapshot too old"
 - not for PREPARED transactions
- event triggers
 - REINDEX
 - login (a bit of a footgun: bug => login impossible, long queries a problem too)
 - event_triggers=false (temporary disable, for debugging, fix for login trigger)
- wait events
 - pg_wait_events (descriptions of wait events, not aggregation/summary)
 - custom wait events (defined by extensions)
 - new wait events (checkpoint delays)

DBA and administration

```
test=# select * from pg_wait_events ;
```

type	name	description
Activity	ArchiverMain	Waiting in main loop of archiver process
Activity	AutovacuumMain	Waiting in main loop of autovacuum launcher process
Activity	BgwriterHibernate	Waiting in background writer process, hibernating
Activity	BgwriterMain	Waiting in main loop of background writer process
Activity	CheckpointerMain	Waiting in main loop of checkpointer process
Activity	LogicalApplyMain	Waiting in main loop of logical replication apply
Activity	LogicalLauncherMain	Waiting in main loop of logical replication launcher
...		
Timeout	SpinDelay	Waiting while acquiring a contended spinlock
Timeout	VacuumDelay	Waiting in a cost-based vacuum delay point
Timeout	VacuumTruncate	Waiting to acquire an exclusive lock to truncate ...
Timeout	WalSummarizerError	Waiting after a WAL summarizer error

```
(266 rows)
```

DBA and administration

- `pg_stat_bgwriter`
 - removed `checkpoints_timed` & `req`
 - removed `write_time` & `sync_time`
 - removed `buffers_checkpoint`, `backend` & `fsync`
- `pg_stat_checkpoint` (new thing)
- `pg_stat_statements`
 - add local block I/O timing (already had shared I/O)
 - add entry time, JIT, min/max statistics (can be reset separately)
 - normalize CALL / SAVEPOINT / PREPARE params (to the usual \$1, \$2, \$3, ...)
- `pg_stat_vacuum_progress`
 - shows index progress (`indexes_total`, `indexes_processed`)

DBA and administration

```
test=# select * from pg_stat_checkpoint ;
```

```
-[ RECORD 1 ]-----+-----
```

```
num_timed          | 6
```

```
num_requested      | 6
```

```
num_done           | 7
```

```
restartpoints_timed | 0
```

```
restartpoints_req   | 0
```

```
restartpoints_done  | 0
```

```
write_time         | 112467
```

```
sync_time          | 629
```

```
buffers_written    | 1734
```

```
slru_written        | 11
```

```
stats_reset        | 2024-11-01 17:56:38.163485+01
```

DBA and administration

- EXPLAIN (SERIALIZE)
 - Show time and memory to serialize data
- COPY
 - ON_ERROR 'ignore'

DBA and administration

```
CREATE TABLE t (a TEXT);
```

```
INSERT INTO t SELECT (  
    SELECT string_agg(md5(i::text), '') FROM generate_series(1,1000) S(i)  
) FROM generate_series(1,8192) x;
```

```
VACUUM ANALYZE t;
```

```
EXPLAIN ANALYZE SELECT * FROM t;
```

QUERY PLAN

```
-----  
Seq Scan on t (cost=0.00..125.08 rows=7208 width=32) (actual time=0.054..1.668 rows=8192 loops=1)  
Planning Time: 0.116 ms  
Execution Time: 2.229 ms  
(3 rows)
```

DBA and administration

```
\timing  
\o /dev/null  
SELECT * FROM t;  
Time: 784.096 ms
```

DBA and administration

```
EXPLAIN (ANALYZE, SERIALIZE) SELECT * FROM t;
```

QUERY PLAN

```
Seq Scan on t (cost=0.00..134.92 rows=8192 width=18) (actual time=0.136..2.905 rows=8192 loops=1)
```

```
Planning Time: 0.117 ms
```

```
Serialization: time=605.642 ms output=256048kB format=text
```

```
Execution Time: 609.632 ms
```

```
(4 rows)
```

DBA and administration

```
CREATE TABLE t (a INTEGER);
```

```
COPY t FROM '/home/user/test.csv' WITH (FORMAT csv);
```

```
ERROR:  invalid input syntax for type integer: "hello"
```

```
CONTEXT:  COPY t, line 5, column a: "hello"
```

```
COPY t FROM '/home/user/test.csv' WITH (FORMAT csv, ON_ERROR ignore);
```

```
NOTICE:  1 row was skipped due to data type incompatibility
```

```
COPY 4
```


DBA and administration

- grant maintenance tasks (permissions) to non-table-owners
 - VACUUM, ANALYZE
 - CLUSTER
 - REINDEX
 - REFRESH MATERIALIZED VIEW
 - LOCK TABLE
- MAINTAIN privilege (per table)
- pg_maintain
 - role granting MAINTAIN on all objects

DBA and administration

- builtin locale provider
 - Only for "C" and "C.UTF-8"
 - Faster! Stable! No external dependency.
 - <https://youtu.be/KTA6oau7tl8>
- Direct TLS handshake
 - sslnegotiation=direct
 - without negotiation, saves a roundtrip
 - friendlier to proxies
 - always with ALPN (https://en.wikipedia.org/wiki/Application-Layer_Protocol_Negotiation)

DBA and administration

- allow_alter_system
 - disable the ALTER SYSTEM command
 - NOT a security feature
 - mostly "protection" for Kubernetes (config managed by resource definition)
 - but (determined) superusers can still change configuration!

SQL and developer

SQL and developer

- PQchangePassword
 - New libpq function
 - Use to.... Change passwords!
 - Used to be psql-only (\password)

SQL and developer

- int to binary and octal
 - `to_bin(123)`, `to_oct(123)`
- infinite intervals
- `random(min, max)` range
 - used to be "double precision" in `[0.0, 1.0)`
 - now int, bigint, numeric
- MERGE
 - can modify updatable views
 - WHEN NOT MATCHED BY SOURCE (our extension of SQL standard, also BY TARGET)
 - MERGE RETURNING (`merge_action()` reports action performed)

SQL and developer

- JSONPATH
 - Many new operators
 - Convert between "data types"
 - e.g. .string() and .boolean()
- SQL/JSON functions
 - New functions from the standard
 - JSON_EXISTS()
 - JSON_QUERY()
 - JSON_VALUE()
- <https://youtu.be/-60ZceCdtSY>

SQL and developer

- JSON_TABLE
 - Convert JSON to relational
 - Like XMLTABLE
 - Single value to multiple columns
 - In one pass
- <https://youtu.be/-60ZceCdtSY>

Backup and replication

Backup and replication

- pg_dump
 - Get list of include/exclude from file
- Incremental pg_basebackup
 - back up only changed pages/blocks
 - uses wal summarizer
 - backup references manifest from full backup
 - to restore, use pg_combinebackup
 - <https://www.postgresql.eu/events/pgconfeu2024/schedule/session/5718-incremental-backup/>

Backup and replication

- Preserve subscriptions across upgrades
 - preserves full subscription state (publisher + subscriber)
 - pg_upgrade
 - upgrade without rebuilding subscribers
- Sync logical replication slots between physical replicas
 - failover enabled on each slot
 - pg_create_logical_replication_slot()
 - CREATE SUBSCRIPTION
 - enable sync_replication_slots on standby
 - configure standby_slot_names

Backup and replication

- pg_createsubscriber
 - cew commandline tool
 - convert physical to logical
 - much faster initial build!
 - before: create a new empty instance + copy all tables + setup slots + ...
 - now: pg_basebackup + pg_createsubscriber (faster, more efficient)
- <https://www.postgresql.eu/events/pgconfeu2024/schedule/session/5853-speeding-up-logical-replication-setup/>

Performance

Performance

- Many infrastructure improvements
 - no direct visibility
 - just runs faster
 - (almost every version)
- COPY performance
 - uuid_out
 - COPY TO when encoding matches

Performance

- VACUUM memory
 - VACUUM uses much less memory
 - internal datastructure changes
 - often an order of magnitude
 - fewer scans!
 - <https://youtu.be/-qPSN1YA19w>
- Redundant NOT NULL
- Parallel CREATE INDEX for BRIN

Performance

- SLRU caches
 - divide cache into banks
 - separate locking
 - configure each size independently (if needed, most grow with shared buffers)
 - xxxx_buffers
 - pg_stat_slru
 - <https://youtu.be/74xAqgS2thY>

Performance

- Vektored I/O
 - Numerous operations use it
 - Better performance for random
 - And foundation for aio
- Streaming I/O
 - Internal API for streamed I/O
 - Callback driven
 - Combines reads, Issues fadvise
 - More foundation for aio
- <https://youtu.be/7HHcD2shS4o>

There's always more

- Lots of smaller fixes
- Performance improvements
- etc, etc
- Can't mention them all!

Test, test, test!

Resources

- release notes
 - <https://www.postgresql.org/docs/release/17.0/>
- pgconf.dev 2024 (Vancouver, May)
 - <https://www.pgevents.ca/events/pgconfdev2024/schedule/>
 - <https://www.youtube.com/@pgconfdev>
- pgconf.eu 2024 (Athens, October)
 - <https://www.postgresql.eu/events/pgconfeu2024/schedule/>
 - <https://www.youtube.com/@pgeu>

Prague PostgreSQL Developer Day 2025 (P2D2)



<https://p2d2.cz/call-for-papers> / <https://p2d2.cz/call-for-sponsors>

