

A look at the Elephants Trunk

PostgreSQL 14

PGConf.NYC 2021
New York City, NY, USA

Magnus Hagander
magnus@hagander.net



Magnus Hagander

- Redpill Linpro
 - Principal database consultant
- PostgreSQL
 - Core Team member
 - Committer
 - PostgreSQL Europe

PostgreSQL 14

PostgreSQL 14

We're done!

Development schedule

- June 2020 - branch 13
- July 2020 - CF1
- September 2020 - CF2
- November 2020 - CF3
- January 2021 - CF4
- March 2021 - CF5
- September 2021 - Release

New features

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

Breaking changes

SCRAM

- scram-sha-256 now the default!
- Old clients may break!
- Change passwords!

SSL compression

- Removed
- Insecure
- (probably didn't work)

Version 2 wire protocol

- Removed
- (Last used in 7.3)

New features

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

Predefined roles

- New name!
 - Used to be "default roles"

Predefined roles

- pg_read_all_data
- pg_write_all_data

Predefined roles

- `pg_database_owner`

SSL

- SNI support
 - Set by default by libpq
- Use DN to log in
 - *clientname=dn* in `pg_hba.conf`

Idle connections

- *idle_session_timeout*

Hung connections

- *client_connection_check_interval*
- Only some platforms
 - Linux

Statistics

Statistics

pg_stat_statements

- pg_stat_statements_info
 - Reset info
- pg_stat_statements.toplevel
 - Differentiate between toplevel and nested

Statistics

pg_stat_wal

- Record count and size
- FPIs
- Buffers
- writes, syncs and times

Statistics

pg_stat_progress_copy

- yes..

Statistics

pg_backend_memory_contexts

- What's in my memory!
- Also pg_log_backend_memory_contexts()
 - What's in somebody else's memory!

Statistics

pg_stat_database

- Session statistics
 - Count
 - Time
 - Terminations

Statistics

pg_locks

- waitstart

New features

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

Generic subscribing

- jsonb, hstore etc

Generic subscripting

- jsonb, hstore etc

```
WITH t AS (  
  SELECT '{"foo": "bar"}'::jsonb AS j  
)  
SELECT j['foo'] FROM t;
```

multiranges

- More than one range!
- GiST indexable

multiranges

- More than one range!
- GiST indexable

```
SELECT int4multirange(  
    int4range(1, 4),  
    int4range(7, 10)  
);
```

Stored procedures

- OUT parameters
- SQL standard function bodies
 - SQL language only
 - *BEGIN ATOMIC*
 - Parsed at definition time
 - Tracks dependencies

date_bin()

- "Group" timestamp into a "bin"
- A bit like date_trunc
 - But much more granular!

date_bin()

- "Group" timestamp into a "bin"
- A bit like date_trunc
 - But much more granular!

```
postgres=# SELECT date_bin('15 minutes', now(), '2021-01-01');
2021-09-08 22:30:00+02
postgres=# SELECT date_bin('17 minutes 30 seconds',
postgres=#                               now(), '2021-01-01');
2021-09-08 22:27:30+02
```


New features

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

restore_command

- Can now change with reload
 - No need for restart

Incompatible params

- Incompatible change on primary
 - E.g. max_connections
- Now pauses recovery
 - Before shutting down
 - (still shuts down)

Logical replication

- Stream long transactions
- Off by default
- Enabled on individual subscription
- Streamed into files on subscriber

New features

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

Scalability improvements

- Faster MVCC snapshots
- On systems with many CPUs
- Particularly with many sessions
- (including idle ones)

Updates on partitioned tbl

- UPDATE and DELETE
- Gains exec time partition pruning
- Other speedups as well

DETACH

- Partition DETACH non-blocking
- `DETACH PARTITION x CONCURRENTLY`

BRIN indexes

- Multiple min/max pairs
- Separate opclass

BRIN indexes

- Multiple min/max pairs
- Separate opclass

```
CREATE INDEX ON t  
  USING brin (a int4_minmax_multi_ops(values_per_range=16));
```

Window functions

- Can now do incremental sorts
- Faster for semi-sorted data

Parallelism

- RETURN QUERY
- REFRESH MATERIALIZED VIEW
- Foreign tables
 - Scan multiple in parallel
 - Depends on FDW

Foreign tables

- INSERTs in bulk mode
 - Control with *batch_size*
 - On server or table
- TRUNCATE

LZ4 compression

- For TOAST
- Faster, sometimes better

LZ4 compression

- For TOAST
- Faster, sometimes better
- *default_toast_compression = 'pglz'*

LZ4 compression

- For TOAST
- Faster, sometimes better
- *default_toast_compression = 'pglz'*

```
ALTER TABLE t ALTER COLUMN x SET COMPRESSION lz4;
```


LZ4 compression

- For TOAST
- Faster, sometimes better
- *default_toast_compression = 'pglz'*

```
ALTER TABLE t ALTER COLUMN x SET COMPRESSION lz4;
```

```
CREATE TABLE t (... , x text COMPRESSION lz4);
```

libpq pipeline mode

- Send multiple queries
 - Before reading results
 - Before even sending the query
- Decreased latency
- Still sequential execution!

VACUUM

VACUUM

- Updated default values
 - vacuum_cost_page_miss
- TOAST vacuum can be skipped

VACUUM

- Skip index vacuum in more cases
 - Few removable index entries
 - 2% instead of just zero
- Recycle btree pages immediately

VACUUM

- Wrap-around failsafe
- Disables cost-based delay
- Stops regular vacuuming
- vacuum_failsafe_age,
vacuum_multixact_failsafe_age
- Defaults: 1.6B

There's always more

There's always more

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

Please help!

- Download and test!
 - apt packages available
 - rpm/yum packages available
 - installers available
 - (etc)

Thank you!

Magnus Hagander
magnus@hagander.net
@magnushagander
<https://www.hagander.net/talks/>

This material is licensed

