

# A look at the Elephants Trunk

# PostgreSQL 13

Postgres Vision  
Virtual

Magnus Hagander  
*[magnus@hagander.net](mailto:magnus@hagander.net)*



# Magnus Hagander

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

# PostgreSQL 13

# PostgreSQL 13

Almost done!

# Development schedule

- July 2019 - branch 12
- July 2019 - CF1
- September 2019 - CF2
- November 201 - CF3
- January 2020 - CF4
- March 2020 - CF5
- June 2020 - Beta2

# New features

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

# DROP DATABASE

- Drop database
- And automatically disconnect users!

# DROP DATABASE

- Drop database
- And automatically disconnect users!

```
DROP DATABASE production
```



# DROP DATABASE

- Drop database
- And automatically disconnect users!

```
DROP DATABASE production
```

- Hangs...

# DROP DATABASE

- Drop database
- And automatically disconnect users!

```
DROP DATABASE production
```

- Hangs...

```
DROP DATABASE production WITH (FORCE)
```

# Breaking changes

- Not as bad as some..

# Wait events renamed

- Many wait events renamed
- More consistent now
- May cause issues with monitoring

# pg\_stat\_statements

- Timing columns renamed
  - total\_time -> total\_exec\_time
  - min\_time -> min\_exec\_time
  - max\_time -> max\_exec\_time
  - mean\_time -> mean\_exec\_time
  - stddev\_time -> stddev\_exec\_time

# New features

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

# ALTER STATISTICS

- Statistics target on extended statistics
- Separate from underlying columns

# ALTER STATISTICS

- Statistics target on extended statistics
- Separate from underlying columns

```
ALTER STATISTICS foo  
  SET STATISTICS 1000;
```



# ANALYZE

- Progress report for analyze
- How far along is the analyze
- How much has been generated

```
SELECT * FROM pg_stat_progress_analyze
```

# Parallel reindexdb

- Reindex multiple indexes in parallel
- (commandline only)
- Remember *maintenance\_work\_mem*

# Parallel VACUUM

- (partially so)
- Parallel processing of indexes
- Main vacuum is still serial

```
VACUUM PARALLEL mytable
```

# Autovacuum

## Trigger by INSERT

- Triggers VACUUM for insert only workloads
  - Not just for anti-wraparound
- More frequent, cheaper, runs
- `autovacuum_vacuum_insert_scale_factor`
- `autovacuum_vacuum_insert_threshold`

# Slow query sampling

- Log sample of slow queries
- Instead of all of them
- Prevent overload of system due to logging
- instead of *log\_min\_duration\_statement*

# Slow query sampling

```
log_min_duration_sample = 100ms  
log_statement_sample_rate = 0.1
```

- *log\_min\_duration\_statement* works independently

SSL

# OpenSSL

- Drop support for 1.0.0 and lower
- New minimum is 1.0.1



# SSL versions

- New default: TLSv1.2
  - Can still be lowered
- SSL version in libpq connection string
  - *sslminprotocolversion*
  - *sslmaxprotocolversion*

# SSL key passwords

- *sslpassword=abc123*
  - Provide password to unlock key
- PQsetSSLKeyPassHook
  - Call-back for password

# Channel binding

- Require channel binding
- *channel\_binding=require*
  - SCRAM authentication only

# postgres\_fdw

- Specify *sslkey* and *sslcert*
  - In of user mapping
- Specify *password\_required=false*
  - Superuser only
  - Allow use of e.g. peer or gss

# pg\_stat\_ssl

- Only contains rows for connections
  - Previously also background processes
- Same for pg\_stat\_gssapi

# pg\_stat\_slru

- Statistics about SLRU structures
- pg\_xact, pg\_subtrans, pg\_multixact etc

# pg\_stat\_activity

- New column *leader\_pid*
- Shows leader in parallel query
- NULL for non-parallel query

# Trusted extensions

- Extensions installed as non-superuser
- Also replaces PL templates
- Several built-in ones migrated
  - btree\_gist
  - intarray
  - pgcrypto
  - ...



# Extensions

- Remove support for unpackaged extensions

# New features

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

# GENERATED columns

- Ability to drop GENERATED

# GENERATED columns

- Ability to drop GENERATED

```
CREATE TABLE tab (  
  a int,  
  col int GENERATED ALWAYS AS (2) STORED  
)
```

# GENERATED columns

- Ability to drop GENERATED

```
CREATE TABLE tab (  
  a int,  
  col int GENERATED ALWAYS AS (2) STORED  
)
```

```
ALTER TABLE tab ALTER COLUMN col  
  DROP EXPRESSION
```

# UUID

- Random generation now in core
- Covers most use-cases
- No need for extension

# UUID

- Random generation now in core
- Covers most use-cases
- No need for extension

```
postgres=# SELECT gen_random_uuid();  
                gen_random_uuid
```

```
-----  
f34f4732-1ce9-46e1-a3e1-4da94d6e07db
```

# JSONPATH

- Now supports datetime



# JSONPATH

- Now supports datetime

```
postgres=# SELECT jsonb_path_query(  
    '{"t": "2020-03-05 14:15:10"}'::jsonb,  
    '$.t .datetime() < "2020-03-05 15:00:22" .datetime()'::jsonp  
    jsonb_path_query  
-----  
true
```

**FETCH FIRST WITH TIES**

# FETCH FIRST WITH TIES

```
postgres=# SELECT * FROM n ORDER BY i FETCH FIRST 2 ROWS ONLY;
 i
---
 1
 2
(2 rows)
```

```
postgres=# SELECT * FROM n ORDER BY i FETCH FIRST 2 ROWS WITH
 i
---
 1
 2
 2
(3 rows)
```

# New features

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

# Archive recovery

- Fails if PITR target not reached
- Server shuts down
- Usually misconfiguration or archive issues

# pg\_stat\_progress\_basebackup

- Shows process while basebackup is running
- Requires progress report enabled to show details
  - On by default with pg\_basebackup

# Backup manifests

- Included in base backups
  - By default
- List of files, checksums
- New tool:

`pg_verifybackup`

# ignore\_invalid\_pages

- Ignore errors on recovery
- **DO NOT USE**
  - Except when you use it



# Logical replication of partitioned tables

- Replicates individual partitions
  - Each partition replicated separately
  - Published individually or via root

```
publish_via_partition_root=true
```

- Receive into partitioned table

# logical\_decoding\_work\_mem

- Amount of RAM for logical decoding
- Before spill-to-disk
- *pg\_stat\_replication.spill\_count*
- *pg\_stat\_replication.spill\_bytes*
- Default: 64kB

# walreceiver replication slot

- Temporary slot used by default
- Unless explicit slot specified

# Limit slot disk usage

- `max_wal_slot_keep_size`
- Avoids running out of disk on primary
- Breaks replication/archiving if hit

# Online reconfiguration

- *primary\_conninfo* and *primary\_slot\_name*
- Change in file or with *ALTER SYSTEM*
- Only a *reload* needed

# New features

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

# Incremental sort

- Optimize multi-level sort
- When "lower level" in query is partially sorted
- Split in "buckets" and only sort individually

# Disk based hash aggregate

- Hash aggregates used in more cases
  - Even when `work_mem` is not enough
- Avoid OOMing on bad estimates
- GROUP BY, DISTINCT and grouping sets



# WAL usage tracking

```
postgres=# EXPLAIN (ANALYZE, WAL) INSERT INTO n VALUES (11);
              QUERY PLAN
-----
Insert on n  (cost=0.00..0.01 rows=1 width=4) (actual time=0.0094 ms)
  WAL: records=1 bytes=59
  -> Result (cost=0.00..0.01 rows=1 width=4) (actual time=0.0094 ms)
Planning Time: 0.026 ms
Execution Time: 0.094 ms
(5 rows)
```

# pg\_stat\_statements

- WAL usage tracking
- Planner statistics
- Renames execution statistics columns

# pg\_stat\_statements

```
postgres=# SELECT * FROM pg_stat_statements WHERE ...;
```

```
-[ RECORD 1 ]-----+-----  
query          | insert into n values ($1)  
plans          | 1  
total_plan_time | 0.033068  
min_plan_time  | 0.033068  
max_plan_time  | 0.033068  
mean_plan_time | 0.033068  
stddev_plan_time | 0  
.....  
wal_records    | 1  
wal_fpi        | 1  
wal_bytes      | 618
```

# Deduplication in btree

- Indexes with many duplicates
  - Full set of columns
- Reduced storage
- Improved performance

**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 the beta version!
  - apt packages available
  - rpm/yum packages available

# Thank you!

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

This material is licensed

