thinkproject

Jonas Marasus
Senior Database Engineer

War Story: How Big Is Too Big For a Schema

8 May 2025

How Big is Too Big?

This is probably too big:

```
postgres=# SELECT count(*) FROM pg_class;
 count
-------
   10000000
(1 row)
```

What Actually Breaks?

thínkproject

- VACUUM and ANALYZE meet OOM-Killer
  - Postgres' relcache takes 4kb memory per relation visited by VACUUM or ANALYZE. (regardless of maintenance_work_mem)
  - Band aid: Use vacuumdb

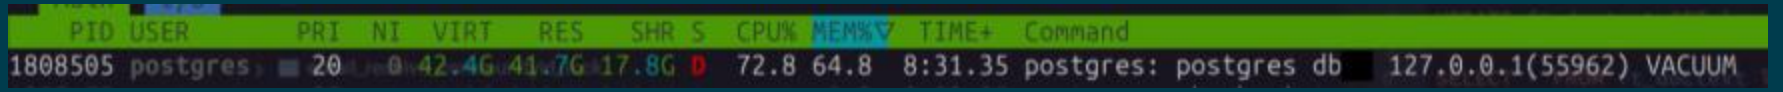

- pg_dump --schema-only takes hours.
  - Putting dumped objects in the right order requires identifying and fixing dependency cycles, this algorithm is non-linear.

- "out of shared memory" when accessing all tables in a transaction (e.g. pg_dump)
  - Band aid: increase max_locks_per_transaction.

How to Avoid This

thínkproject

- For dynamic schemas, consider using jsonb instead.
- When partitioning, avoid large numbers of small partitions.
- Anticipate some changes in usage patterns
- Avoid unused indexes even on empty tables

Thank you!