

# Importance of PostgreSQL Vacuum Tuning To Optimize Database Performance

PRESENTED BY

**Suyog Pagare**

Lead DBA – PostgreSQL

Datavail



**datAvail**

# Suyog Pagare

Lead DBA -  
PostgreSQL,  
Datavail



- ✓ 15 years in Database Administration
- ✓ Specialties: Database Migrations, Performance tuning
- ✓ Microsoft Certified in Azure Solution Architect Expert
- ✓ AWS Certified Solutions Architect
- ✓ AWD DB specialty

Connect or follow me on LinkedIn:



<https://www.linkedin.com/in/suyog-pagare-5a226918/>

# Put the PostgreSQL Pieces Together

✓ Fill out your session evaluation form for a chance to win a LEGO set.

# Datavail at a Glance

Delivering a superior approach to leverage data through application of a tech-enabled global delivery model & deep specialization in databases, data management, and application services.



# 16<sup>+</sup>

**Years**

building and operating mission critical data and application systems



# \$25<sup>M</sup>

**Invested**

in IP that improves the service experience and drives efficiency



# 1,000<sup>+</sup>

**Employees**

staffed 24x7, resolving over 2,000,000 incidents per year



# Agenda

- ✔ Database Bloat
- ✔ What is Vacuum and Autovacuum
- ✔ Best Practices for Tuning Autovacuum
- ✔ Dead tuples Not Yet Removable
- ✔ Pg\_repack
- ✔ Other DB Parameter Tuning

# Introduction to PostgreSQL Maintenance

## Overview

- ✔ PostgreSQL is an advanced, open-source relational database system.
- ✔ Like any database, PostgreSQL requires regular maintenance to perform optimally.

## Key Maintenance Tasks

- ✔ Autovacuum
- ✔ Manual Vacuum
- ✔ Analyze
- ✔ Index Maintenance

# Understanding Database Bloat

- ✔ Definition: Database bloat refers to the unnecessary consumption of disk space in a database due to outdated or deleted data, and inefficiencies in storage.
- ✔ Causes of Bloat
  - Frequent updates and deletes: PostgreSQL uses a multi-version concurrency control (MVCC) system, which can leave behind dead tuples.
  - No Vacuum.
  - Inefficient use of indexes: Unused or poorly maintained indexes can also contribute to bloat.
- ✔ Symptoms of Database Bloat
  - Increase Disk Usage
  - Degraded Performance
  - Longer Backup and Restore Times
- ✔ How to Identify Bloat
  - Pgstattuple extension
  - Pg\_stat\_user\_tables: To monitor dead tuples and live tuple counts

# Query to Find Bloat

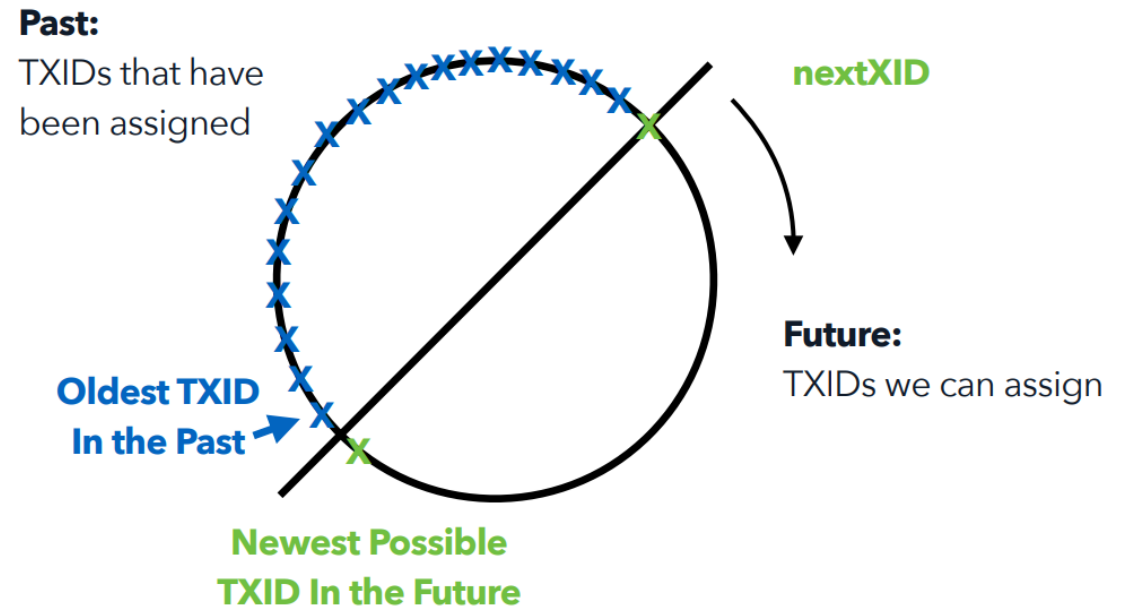
```
SELECT
  c.oid,
  n.nspname || '.' || c.relname AS ObjectName,
  pg_stat_get_live_tuples(c.oid) AS LiveTuples,
  pg_stat_get_dead_tuples(c.oid) AS DeadTuples,
  CASE
    WHEN pg_stat_get_live_tuples(c.oid) > 0
    THEN (pg_stat_get_dead_tuples(c.oid)::numeric / pg_stat_get_live_tuples(c.oid)::numeric(18,3)) * 100
    ELSE 0
  END AS DeadtupleRatio,
  pg_table_size(n.nspname || '.' || c.relname) AS tbl_sz,
  pg_size_pretty(pg_table_size(n.nspname || '.' || c.relname)) AS tbl_sz_prty,
  pg_size_pretty(pg_indexes_size(n.nspname || '.' || c.relname)) AS tbl_idx_sz_prty,
  (SELECT count(1) FROM pg_indexes i WHERE i.schemaname || '.' || i.tablename = n.nspname || '.' || c.relname) AS idx_cnt,
  CASE
    WHEN pg_table_size(n.nspname || '.' || c.relname) > 0
    THEN (pg_indexes_size(n.nspname || '.' || c.relname)::numeric / (pg_indexes_size(n.nspname || '.' || c.relname) + pg_table_size(n.nspname || '.' || c.relname))) * 100
    ELSE 0
  END AS idx_tbl_ratio,
  relhastriggers
FROM
  pg_class c
JOIN
  pg_catalog.pg_namespace n
ON
  c.relnamespace = n.oid
WHERE
  c.relkind = 'r'
  AND n.nspname NOT LIKE 'pg\_%'
ORDER BY
  pg_table_size(n.nspname || '.' || c.relname) DESC
LIMIT 30;
```

# Consequences of Ignoring Bloat:

- ✔ Performance Degradation
- ✔ High Storage Costs
- ✔ Operations Challenges
- ✔ Transaction Id Wraparound

# Transaction ID Wraparound

- ✓ What is Transaction ID (Xid)
- ✓ Wraparound Issue



# Preventing and Managing Bloat

- ✓ Vacuum....!!!
- ✓ It removes dead tuples
- ✓ Reclaim storage back to the OS
- ✓ Maintain visibility map

## ✓ Types of Vacuum

- Vacuum
- Vacuum Full

Syntax:

Vacuum verbose table\_name;

vacuum (index\_cleanup true) table\_name;

➤ Analyze

```
VACUUM [ ( option [, ...] ) ] [ table_and_columns [, ...] ]  
VACUUM [ FULL ] [ FREEZE ] [ VERBOSE ] [ ANALYZE ] [ table_and_columns [, ...] ]
```

where *option* can be one of:

```
FULL [ boolean ]  
FREEZE [ boolean ]  
VERBOSE [ boolean ]  
ANALYZE [ boolean ]  
DISABLE_PAGE_SKIPPING [ boolean ]  
SKIP_LOCKED [ boolean ]  
INDEX_CLEANUP { AUTO | ON | OFF }  
PROCESS_MAIN [ boolean ]  
PROCESS_TOAST [ boolean ]  
TRUNCATE [ boolean ]  
PARALLEL integer  
SKIP_DATABASE_STATS [ boolean ]  
ONLY_DATABASE_STATS [ boolean ]  
BUFFER_USAGE_LIMIT size
```

and *table\_and\_columns* is:

```
table_name [ ( column_name [, ...] ) ]
```

# Autovacuum

- ✓ Automate solution to execute Vacuum
- ✓ No manual intervention needed
- ✓ Triggers it based on the defined thresholds
- ✓ Does not reclaim the free space

# Understanding Autovacuum and Autoanalyze

## Benefits of Autovacuum

- Automates Maintenance
- Prevents Data Corruption
- Optimize Performance

## Formula:

**Autovacuum** = `autovacuum_vacuum_scale_factor` \* number of rows  
+ `autovacuum_vacuum_threshold`

**Autoanalyze** = `autovacuum_analyze_scale_factor` \* number of rows  
+ `autovacuum_analyze_threshold`

- Check status using table: `pg_stat_vacuum_progress`

# Tuning of Autovacuum Parameters

## ✓ Important Parameters:

- Autovacuum\_vacuum\_threshold/ Autovacuum\_scale\_factor
- Autovacuum\_analyze\_threshold/ Autoanalyze\_scale\_factor
- Autovacuum\_max\_workers
- Autovacuum\_naptime
- Autovacuum\_cost\_delay
- Max\_worker\_processes
- Max\_parallel\_workers
- Max\_parallel\_maintenance\_workers
- Default\_stats\_target

- ✓ Setup a cronjob daily to execute vacuum based of % of table bloat
- ✓ Use vacuumdb utility to run vacuum with parallel jobs
- ✓ Setup aggressive Autovacuum for specific tables
- ✓ Monitor vacuum process
- ✓ For faster operation, set high maintenance\_work\_mem and execute manual vacuum for faster execution

# Dead tuple Not Yet Removed

Sometimes even after manual vacuum, dead tuples are not removed.

- ✓ Long running transaction
- ✓ Abandoned replication slot
- ✓ DDL transaction is in process
- ✓ Orphan Prepared transactions

```
SELECT gid, prepared, owner, database, transaction AS xmin  
FROM pg_prepared_xacts  
ORDER BY age(transaction) DESC;
```

- ✓ Standby with hot\_standby\_feedback=on

```
SELECT application_name, client_addr, backend_xmin  
FROM pg_stat_replication  
ORDER BY age(backend_xmin) DESC;
```

- ✔ pg\_repack is a PostgreSQL extension used for reclaiming space and optimizing database performance without requiring a full table lock.
- ✔ Purpose: Efficiently removes table and index bloat, reducing storage footprint and improving query speed.
- ✔ **Online Operation:** Performs maintenance tasks without significant downtime, allowing for continuous database availability.
- ✔ **Bloat Reduction:** Reorganizes tables and indexes, eliminating wasted space caused by deleted or updated tuples.
- ✔ **Pre-checks:** Ensure sufficient disk space and backups before running.

# How to Run pg\_repack

- ✓ Install pg\_repack
- ✓ Add pg\_repack to shared\_preload\_libraries
- ✓ Create extension pg\_repack on database
- ✓ Run command to execute pg\_repack
  - `pg_repack -d mydb`  
You can execute pg\_repack for specific table as well
  - `pg_repack -t mytable -d mydb`

# Other DB Parameter Tuning

- ✓ <https://pgtune.leopard.in.ua/>
- ✓ <https://pgconfigurator.cybertec-postgresql.com/>
- ✓ <https://www.pgconfig.org>
- ✓ <https://www.dbtune.com/>

# Put the PostgreSQL Pieces Together

✓ Fill out your session evaluation form for a chance to win a LEGO set.



Q

&

A

# Thank You



**Suyog Pagare**

Lead DBA: PostgreSQL Services

[Suyog.Pagare@datavail.com](mailto:Suyog.Pagare@datavail.com)

+1 720-756-4881

**datAvail**