Leveraging pgBadger for effective PostgreSQL troubleshooting

Alicja Kucharczyk

PGConf.EU 2023 | 15.12.23

About me

- Principal Program Manager, Azure Database for Postgres – Flexible Server, Microsoft
- Postgres DBA, consultant, etc.
- 🐘 Warsaw PostgreSQL Users Group founder and co-organiser.
- Postgres song producer: Nothing Compares To VACUUM/The Ballad of Bloat, Explain Analyze (Feliz Navidad cover)

What this talk isn’t about

- Deep dive into particular Postgres features, i.e. how autovacuum works

What this talk is about

- Quickly solve the most common problems
- Become a hero in your organization

Troubleshooting Checklist

Snapshots of the system views (pg_stat_bgwriter, pg_stat_activity, pg_stat_user_tables, pg_stat_databases etc.)

OS level metrics: CPU, memory, IOPS usage.

General instance data: version, SKU, db and objects sizes, parameter settings.

Logs, especially: temporary files, locks, errors, vacuum activities, queries and query durations.

pgBadger

What is pgBadger?

- Small, standalone Perl script (works almost everywhere)
- Takes postgres **text logs** as an input and generates **html** report
- The most comprehensive tool for the task (in opposite to pg_stat_statements ;))

Installation

- Ubuntu: sudo apt-get install -y pgbadger
- macOS: brew install pgbadger
- Windows: use WSL - it will be easier

Configuration (postgresql.conf)

logging_collector = on # the only static one (or syslog)

log_line_prefix = '%t [%p]: [%l-1] db=%d,user=%u,app=%a,client=%h '

log_connections = on

log_disconnections = on

log_lock_waits = on

log_temp_files = 0

log_autovacuum_min_duration = 0

log_min_duration_statement = 60000 # pretty safe

log_line_prefix

```
log_line_prefix = '%m [%p] %q%u@%d '            # special values:
                                        #   %a = application name
                                        #   %u = user name
                                        #   %d = database name
                                        #   %r = remote host and port
                                        #   %h = remote host
                                        #   %b = backend type
                                        #   %p = process ID
                                        #   %P = process ID of parallel group leader
                                        #   %t = timestamp without milliseconds
                                        #   %m = timestamp with milliseconds
                                        #   %n = timestamp with milliseconds (as a Unix epoch)
                                        #   %Q = query ID (0 if none or not computed)
                                        #   %i = command tag
                                        #   %e = SQL state
                                        #   %c = session ID
                                        #   %l = session line number
                                        #   %s = session start timestamp
                                        #   %v = virtual transaction ID
                                        #   %x = transaction ID (0 if none)
                                        #   %q = stop here in non-session
                                        #        processes
                                        #   %% = '%'
                                        # e.g. '<%u%%%d> '
```

DEMO

ANIMALFOTO
STUDIO

Low-hanging fruits or How to quickly impress your peers?

- Sample report https://pgbadger.darold.net/examples/sample.html
- Temp Files - Queries generating the most temporary files (N)
- Vacuums – Vacuums per Tables
- Events - Most Frequent Errors/Events

Queries generating the most temporary files (N)

Rank	Count	Total size	Min size	Max size	Avg size	Query
1	70	859.11 MiB	12.27 MiB	12.28 MiB	12.27 MiB	SELECT "view_stadium_reports".* FROM "view_stadium_reports" WHERE (to_char (date, '') LIKE upper ('')) ORDER BY view_stadium_reports.checkin_date ASC, view_stadium_reports.id ASC LIMIT 0 offset 0; Examples
2	37	454.00 MiB	12.27 MiB	12.29 MiB	12.27 MiB	SELECT "view_stadium_reports".* FROM "view_stadium_reports" WHERE (to_char (date, '') LIKE upper ('')) ORDER BY view_stadium_reports.job_title_name ASC, view_stadium_reports.id ASC LIMIT 0 offset 0; Examples
3	27	331.23 MiB	12.27 MiB	12.28 MiB	12.27 MiB	SELECT count (*) FROM "view_stadium_reports" WHERE (to_char (date, '') LIKE upper ('')); Examples
4	27	343.77 MiB	12.27 MiB	24.59 MiB	12.73 MiB	SELECT count (*) FROM "view_stadium_reports" WHERE (upper (CAST ((view_stadium_reports.hotel) AS text)) LIKE upper ('')); Examples

Temp Files - Queries generating the most temporary files (N)

- Increase work_mem if temp files size is small
- If very high (GB)– look into the query itself – does it contain unnecessary clauses like ORDER BY? Amount of data is huge?
- Maybe you can set work_mem per query/function?
- Maybe you need Citus?

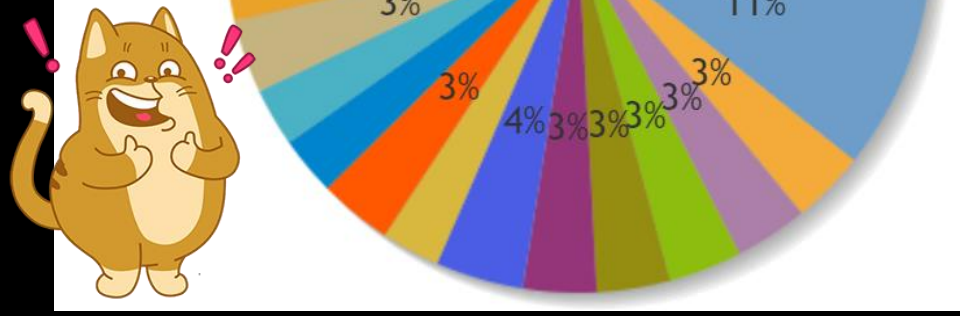

Vacuums – Vacuums per Tables

Vacuums – Vacuums per Tables

- Autovacuum is triggered when the number of dead tuples exceed: autovacuum_vacuum_threshold + autovacuum_vacuum_scale_factor * reltuples

- on a 20-GB table, this scale factor translates to 4 GB of dead tuples. On a 1-TB table, it’s 200 GB of dead tuples.

Vacuums – Vacuums per Tables

For both, very small and large (it works on my machine):

```
ALTER TABLE t SET (autovacuum_vacuum_threshold = 10000);
ALTER TABLE t SET (autovacuum_vacuum_scale_factor = 0);
```

Events - Most Frequent Errors/Events

⚠ Most Frequent Errors/Events

KEY VALUES

3,930
Max number of times the same event was reported

3,965
Total events found

Rank	Times reported	Error
1	3,930 Details	`ERROR: invalid input syntax for type date: ...` Examples
2	15 Details	`LOG: process ... still waiting for ExclusiveLock on extension of relation ... of database ... after ... ms` Examples
3	13 Details	`ERROR: syntax error at or near "..."` Examples
4	2 Details	`ERROR: column "..." does not exist` Examples
5	2 Details	`WARNING: pgstat wait timeout` Examples
6	2 Details	`ERROR: invalid input syntax for integer: ...` Examples
7	1 Details	`ERROR: relation "..." does not exist` Examples

Events - Most Frequent Errors/Events

- Might be dev errors, like **ERROR: invalid input syntax for type date: at character 32 or ERROR: syntax error at or near "GROUP" at character 17**
- Or DB problems

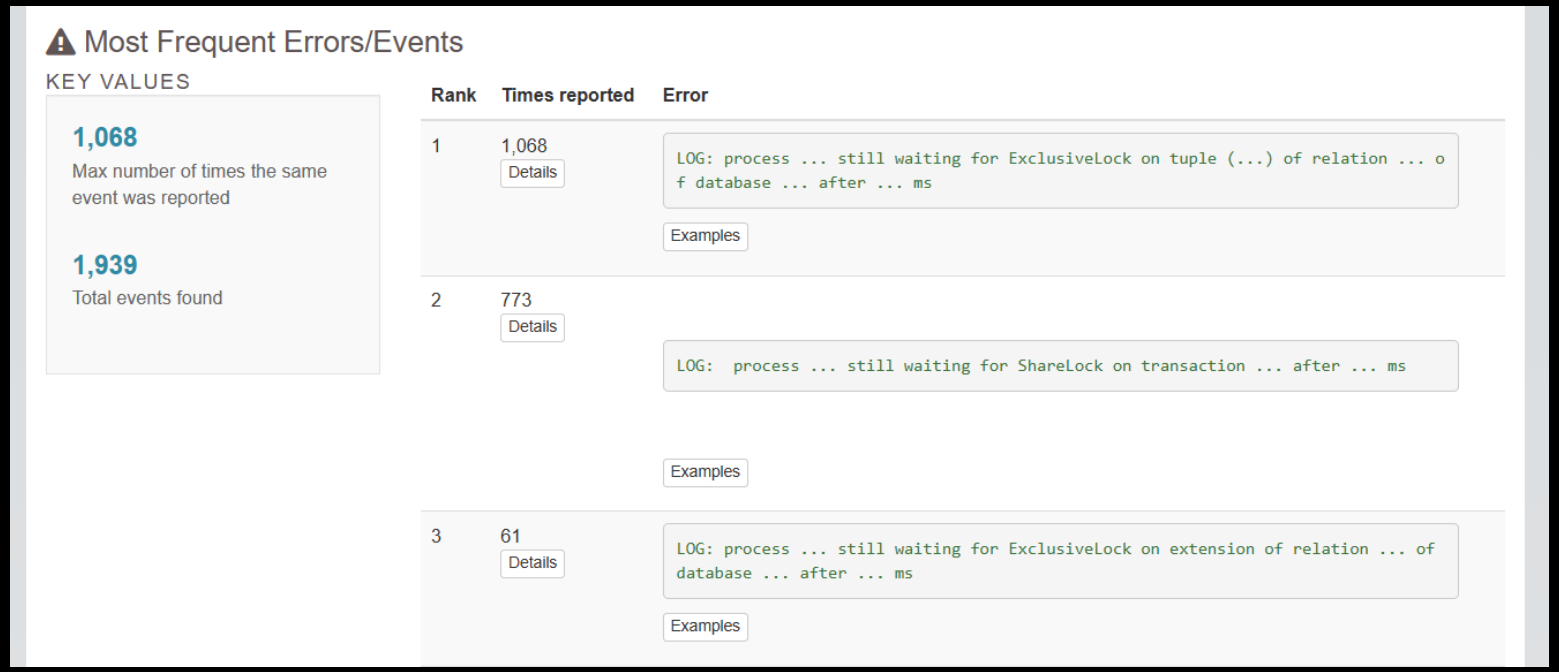

Events - Most Frequent Errors/Events

Locks

Locks by types

KEY VALUES

ExclusiveLock
Main Lock Type

1,875 locks
Total

Chart | Table

Type	Object	Count	Total Duration	Average Duration (s)
ExclusiveLock		1,104	2h1m24s	6s598ms
	extension	56	8m28s	9s79ms
	tuple	1,048	1h52m55s	6s465ms
ShareLock		771	52m26s	4s80ms
	transaction	771	52m26s	4s80ms
	Total	1,875	2h53m50s	5s562ms

Locks

- Bad News: pgBadger will show what was locked but not what caused the lock

Time consuming queries (N)

Rank	Total duration	Times executed	Min duration	Max duration	Avg duration	Query
1	4h9m39s	11,135 Details	300ms	23s855ms	1s345ms	SELECT count(*) FROM pg_catalog.pg_class AS c WHERE c.oid = pgpool_regclass (?) AND c.relpersistence = ?; Examples User(s) involved
2	3h42m17s	132 Details	30s641ms	4m28s	1m41s	COPY nppsmoketests (last_update, testname, current_status, suiteid, regressioncl, testcl, os, arch, build_type, branch, gpu, subtest, osversion) FROM STDIN WITH csv header escape ?; Examples User(s) involved
3	1h40m37s	6,895 Details	300ms	10s997ms	875ms	UPDATE gpu SET dd_date = ?, last_update = (SELECT date_trunc(?, now())) WHERE SYSTEM = ?; Examples User(s) involved
4	53m8s	3,406 Details	300ms	11s30ms	936ms	UPDATE systems SET daemon_update = now(), activated = ? WHERE vm = ?; Examples User(s) involved
5	43m5s	2,986 Details	300ms	10s743ms	865ms	UPDATE gpu SET dd_date = ?, last_update = (date_trunc(?, now())) WHERE SYSTEM = ? Examples User(s) involved
6	39m22s	2,816 Details	318ms	15s37ms	839ms	INSERT INTO proccheck (hostname, process, subtestid) VALUES (?, ?, ?); DELETE FROM proccheck WHERE age(now(), datestamp) > interval ?; Examples User(s) involved

Top - Time consuming queries (N)

Gotchas

- Log_line_prefix
- Changes! Log_min_duration
- Which parameters where enabled at given time?
- Size

Thank you!

- Questions?
- https://twitter.com/StiepanTrofimo
- alkuchar@microsoft.com
- More about my dogs: https://postgresql.pet/post/gustaw_stefan_parrot/