



Sustainable Database Performance Profiling in PostgreSQL

25.06.2026

**Dirk Krautschick
Swiss PG Day**



#whoami



Dirk Krautschick
Senior Solution Architect




19 years

**DBA, Trainer, Consulting,
Sales Engineering**

PostgreSQL, Oracle,...

...and Kafka, Clickhouse, OpenSearch,...

Married, 2 Junior DBAs

**mountainbike, swimming,
movies, music,
hifi/homecinema  3 bit
computing**

Before we start...

Disclaimer

Different audience, different perspectives

My experience, my honest opinion

Let's stay open minded

Always open for discussions

Before we start...

What happened so far...?

There was a talk...

"Pro-Active Performance Analysis in PostgreSQL"

<https://www.youtube.com/watch?v=rgdA0FwVShI>




About

Performance problems overall

Different analysis approaches

Recommendations/usage of Extensions



Before we start...

What happened so far...?

Solid feedback, consensus in practical experiences

Many questions about "THAT LAST PART of the talk"

Sick of giving this talk after so many times ... :-)

Motivation to do a Spin-Off talk!

Before we start...

Performance Problems



**In PostgreSQL every relevant information is there...
...but only for NOW!**

Obvious Sources

Parameters, Sizing (at that time!)

Information_schema, system catalogues

Main Challenge: How to handle, keep and collect all that stuff!

Before we start...

What about monitoring...?

RECAP

For sure, monitoring is essential, but...

...it shows mostly

that something is slow, sometimes maybe...

what is exactly slow, but almost never...

why it is slow!

PostgreSQL insights necessary

Deep dive or investigation as a next step anyway

Before we start...

What about logging...?



PostgreSQL logging is awesome

Exhaustive possibilities

Straight and easy configuration

Be aware of storage and load

High maintenance

Evaluate Logging strategies

```
log_line_prefix = '%t [%p]:
user=%u,db=%d,app=%a,client=%h ,
...
log_parser_stats = off
log_planner_stats = off
log_executor_stats = off
log_statement_stats = on
...
log_checkpoints = on
log_connections = on
log_disconnections = on
log_lock_waits = on
log_temp_files = 0
log_autovacuum_min_duration = 0
log_error_verbosity = default
...
log_min_messages = debug5
log_min_error_statement = debug5
log_min_duration_statement = 0
log_min_duration_sample = 0
...
log_statement = 'all'
```

Before we start...

PG_STAT_STATEMENTS



Statement level statistics

Required by several monitoring tools

Statement based collection of e.g.

Executions

Execution times (min, max, average)

Rows

Blocks read/write

```
# \d pg_stat_statements
View "public.pg_stat_statements"
      Column          |          Type          | ...
-----+-----+-----
userid                |          oid           | ...
dbid                  |          oid           | ...
queryid               |         bigint        | ...
query                 |          text          | ...
total_plan_time       | double precision      | ...
...
calls                 |         bigint        | ...
total_exec_time       | double precision      | ...
min_exec_time         | double precision      | ...
max_exec_time         | double precision      | ...
mean_exec_time        | double precision      | ...
stddev_exec_time      | double precision      | ...
rows                  |         bigint        | ...
...
blk_read_time         | double precision      | ...
blk_write_time        | double precision      | ...
...
```

Before we start...



PG_STAT_STATEMENTS

```
# SELECT
  substring(query, 1, 50) as short_query,
  round(total_exec_time) as total_exec_time, calls,
  round(mean_exec_time) as mean_exec_time,
  round(100 * total_exec_time / (SELECT sum(total_exec_time) FROM stat_statements)) as percentage
FROM
  pg_stat_statements
ORDER BY
  percentage desc;
```

short_query	total_exec_time	calls	mean_exec_time	percentage
UPDATE pgbench_branches SET bbalance = bbalance +	7114	1500	5	63
UPDATE pgbench_tellers SET tbalance = tbalance + \$	2506	1500	2	22
copy pgbench_accounts from stdin	664	1	664	6
UPDATE pgbench_accounts SET abalance = abalance +	194	1500	0	2
alter table pgbench_accounts add primary key (aid)	193	1	193	2
vacuum analyze pgbench_accounts	138	1	138	1

...

Before we start...

PG_STAT_KCACHE



Statistics about

reads/writes on filesystem

Statistics about CPU usage

pg_stat_statements is required

```
postgres=# \d pg_stat_kcache_detail;
```

Column	Type	Collation	Nullable	Default
datname	name			
plan_user_time	double precision			
plan_system_time	double precision			
plan_minflts	numeric			
plan_majflts	numeric			
plan_nswaps	numeric			
plan_reads	numeric			
plan_reads_blks	numeric			
plan_writes	numeric			
plan_writes_blks	numeric			
...				
plan_nivcsws	numeric			
exec_user_time	double precision			
exec_system_time	double precision			
...				
exec_nsignals	numeric			
exec_nvcsws	numeric			
exec_nivcsws	numeric			

Before we start...

What is still missing...?

RECAP

Before we start...

Handling WAIT_EVENTS!!!

RECAP



Before we start...

PG_WAIT_SAMPLING



Wait events from `pg_stat_activity`

Sampled statistics of wait events

github.com/postgrespro/pg_wait_sampling

Views

```
pg_wait_sampling_current  
pg_wait_sampling_history  
Pg_wait_sampling_profile
```

Functions

```
pg_wait_sampling_get_current(pid)  
pg_wait_sampling_reset_profile()
```

Before we start...

PG_STAT_STATEMENTS




```
postgres=# select * from pg_wait_sampling_profile order by pid, count desc;
```

pid	event_type	event	queryid	count
1689	IO	DataFileWrite	2862011717192834034	4010499
1685	IO	DataFileRead	2862011717192834034	4010097
1686	IO	DataFileSync	-4888004026240188267	4007477
1684	Activity	BgWriterHibernate	2862011717192834034	3991477
1683	Activity	CheckpointerMain	1511417639870010300	3927957
1684	Activity	BgWriterMain	-4888004026240188267	88494
3720	Client	ClientRead	-4888004026240188267	2393
1685	IO	WALSync	6648255685428052402	65
3546	Client	ClientRead	-4888004026240188267	1
3546	IO	DataFileRead	2862011717192834034	1

...

The Idea

Getting sustainable?



Think outside the box

Let's pick a random example...

... let's say ... Oracle Database :-)

Collects almost everything per default (sometimes sampled)

Interpretation with

Querying Views (obviously!)

Statspack (basic, always available and "costless")

Think outside the box

Let's pick a random example...
... let's say ... Oracle Database :-)

Collects almost everything per default (sometimes sampled)

Interpretation with
Querying Views (obviously!)
Statspack (basic, always available and "costless")

Diagnostic and Tuning Pack (expensive Option)
Only for Enterprise Edition
Several Tools, like ASH, AWR,...



Think outside the box

Frequent snapshots of performance data in a repository

Defined time periods and retention

Creation of nice reports based on those snapshots

Time frame between two or more snapshots

Think outside the box

Frequent snapshots of performance data in a repository

Defined time periods and retention

Creation of nice reports based on those snapshots

Time frame between two or more snapshots



But hey, there was already...



Advanced logging analysis reporting

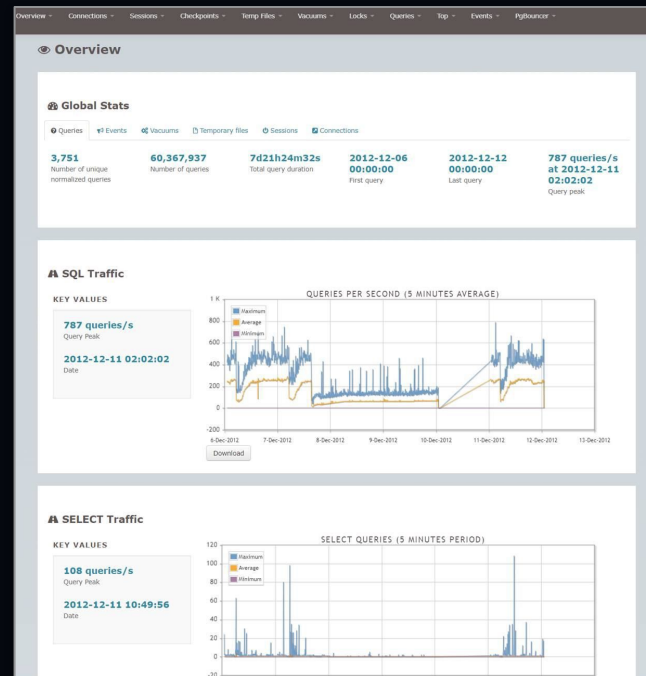
<https://pgbadger.darold.net/>

Incremental daily/cumulative weekly reports

The right direction, but still

Massive logging necessary

Log file handling



Getting sustainable

Ring-Buffer-like settings in extensions are volatile

Several different retention


Several Views, Tables...but also volatile

How to handle the collection of all that information?

The Idea

putting all in a repository/database

while handling the retention of all information



Development

Initiated by Andrei Zubkov

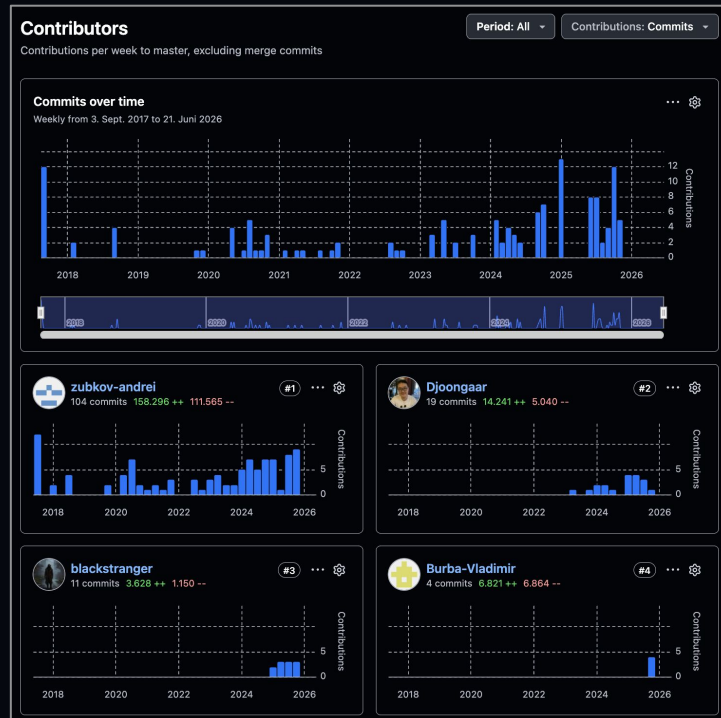
Individual Open Source license

https://github.com/zubkov-andrei/pg_profile

Starting Release v0.0.7 (Nov 2019)

Actual Release v4.11 (Nov 2025)

Pure PL/PGsql-based



Introduction

A good or the actual only(?) Example

**Sample collection of
System Catalogue, information_schema**

PG_STAT_STATEMENTS

PG_STAT_KCACHE

PG_WAIT_SAMPLING

Report Examples

Server Statistics

SQL query Statistics

Wait Event Statistics

Schema Object Statistics

User Function Statistics

Vacuum-related stats

Cluster settings

Report Examples

Contents Filter... Everywhere

Postgres profile report

Report details

Version	Server name	Interval (sample)	Interval (time)
		start	end
4.4	core16	10	11
		start	end
		2024-03-07 00:11:23+00	2024-03-07 00:21:57+00

Server statistics

Database statistics

Database	Transactions				Block statistics				Tuples			
	Commits	Rollbacks	Deadlocks		Hit(%)	Read	Hit	Ret	Fet	Ins	Upd	Del
postgres	531252	1	1	100	282	28129447	5258592	3875027	632778	1594587	105619	
Total	531252	1	1	100	282	28129447	5258592	3875027	632778	1594587	105619	

Cluster I/O statistics

Object	Backend	Context	Reads		Writes		Writebacks		Extends		Hits		
			Count	Bytes	Count	Bytes	Count	Bytes	Count	Bytes			
relation	auto-vacuum worker	normal							14	112 kB	16666		
relation	auto-vacuum worker	vacuum									30932		
relation	auto-vacuum worker	*							14	112 kB	47598		
relation	background worker	normal									893		
relation	background worker	*									893		
relation	checkpoint-proc	normal			3597	28 MB	3589	28 MB					
relation	checkpoint-proc	*			3597	28 MB	3589	28 MB					
relation	client backend	bulkwrite							1668	13 MB	1612		
relation	client backend	normal			202	2256 kB			3625	28 MB	18776552		
relation	client backend	vacuum									1959		
relation	client backend	*			202	2256 kB			5293	41 MB	18780223		
relation	Total				282	2256 kB	3597	28 MB	3589	28 MB	5387	41 MB	18828624
relation	* Total				282	2256 kB	3597	28 MB	3589	28 MB	5387	41 MB	18828624

Cluster SLRU statistics

Name	Zeroed	Hits	Reads	%Hits	Writes	Checked	Flushed	Truncated
MultiXactMember		42	100					
MultiXactOffset		42	100					
Sessions	259	4	100	105				
Xact	26	2678792	100					
Total	275	2678800	100	195				

Session statistics by database

Database	Timings				Sessions			
	Total	Active	Idle	Established	Abandoned	Faral	Killed	
postgres	1021	3160.90	352.0	10	1			
Total	1021	3160.90	352.0	10	1			

Cluster statistics

Metric	Value	Metric	Value
Scheduled checkpoints	2	WAL generated	270 MB
Requested checkpoints		WAL per second	437 kB
Checkpoint write time (s)	84.33	WAL records	4014621
Checkpoint sync time (s)	0.83	WAL FPI	2258
Checkpoint buffers written	3597	WAL buffers full	813
Background buffers written		WAL writes	532063
Backend buffers written	3949	Backend buffers written	839.22
Backend fsync count		WAL sync	531190
Bgwriter interrupts (too many buffers)		WAL syncs per second	837.84
Number of buffers allocated	5595	WAL write time (s)	
		WAL write duty	
		WAL sync time (s)	
		WAL sync duty	
		WAL segments archived	
		WAL segments archive failed	

Tablespace statistics

Tablespace	Path	Size	Growth
pg_default		75 MB	43 MB
pg_global		865 kB	

Wait sampling

Wait events types

Wait event type	Statements Wait(ed)	%Total	All Wait(ed)	%Total
Activity			2699.7	37.87
Client			1202.22	17.99
IO		0.04	377.7	2.41
Lock	2680.11	99.89	2680.11	37.65
LWLock	2.93	0.11	3.83	0.05
Timeout	0.01	0.01	284.24	3.99
Total	2683.09	100	7126.9	100

Top wait events (statements)

Top wait events detected in statements execution

Wait event type	Wait event	Wait(ed)	%Total
Lock	transactionid	2332.04	86.92
Lock	tuple	348.07	12.97
LWLock	BufferContent	1.53	0.06
LWLock	LockManager	1.4	0.05
IO	DataFileExtend	0.03	
IO	DataFileRead	0.01	
Timeout	SpinDelay	0.01	

WAL statistics

Metric	Value
WAL generated	270 MB
WAL per second	437 kB
WAL records	4014621
WAL FPI	2258
WAL buffers full	813
WAL writes	532063
Backend buffers written	839.22
WAL sync	531190
WAL syncs per second	837.84
WAL write time (s)	
WAL write duty	
WAL sync time (s)	
WAL sync duty	
WAL segments archived	
WAL segments archive failed	

SQL query statistics

Top SQL by execution time

Query ID	Database	User	Exec (s)	%Total	CPU time (s)	Rows	Mean	Min	Max	Execution times (ms)		
					Up	Sys				Mean	Min	Max
6c27814654fc54fd5	postgres	postgres	1502.85	53.1	26.03	10.78	531066	2.83	0.01	198.78		
17263a89808520b	postgres	postgres	1308.85	46.24	13.58	5.01	531066	2.46	0.01	331.14		
80a3065afcf839204	postgres	postgres	10.79	0.38	9.02	3.81	531066	0.02	0.01	31.32		
6a27827af51bc084	postgres	postgres	3.57	0.13	4.27	1.78	531066	0.01		24.12		
1cc31465487b6431	postgres	postgres	3.19	0.11	4.14	1.63	531066	0.01		20.22		
fff824867e71c826	postgres	postgres	0.42	0.01	0.35	0.01	1	419.36	419.36	419.36		
5c4b134935263c2f	postgres	postgres	0.19	0.01						0.14		
7d1e4c091bc57d8	postgres	postgres	0.18	0.01	0.16	0.01	1	176.73	176.73	176.73		
1ca7e89efcf7c423	postgres	postgres	0.18	0.01						0.46		
e8217574620ac4c2	postgres	postgres	0.13			100000	131	311	131			
1aaadf7524987c68	postgres	postgres	0.02					22.80	22.80	22.80		
7658fadac8d4b37	postgres	postgres	0.02					21.11	21.11	21.11		
f876c7178ed4f103	postgres	postgres	0.01	0.01			271	8.24	8.24	8.24		
580891d1f899c03	postgres	postgres	0.01	0.01			188	8.01	8.01	8.01		
e9063093427fae11	postgres	postgres	0.01	0.01			10	5.59	5.59	5.59		

Top SQL by executions

Query ID	Database	User	Executions	%Total	Rows	Mean(ms)	Min(ms)	Max(ms)	StdErr
1ca7e89efcf7c423	postgres	postgres	531063	14.29				0.46	
5c4b134935263c2f	postgres	postgres	531063	14.29				0.14	
6c27814654fc54fd5	postgres	postgres	531060	14.29	531060	2.83	0.01	198.78	3
17263a89808520b	postgres	postgres	531060	14.29	531060	2.46	0.01	331.14	4

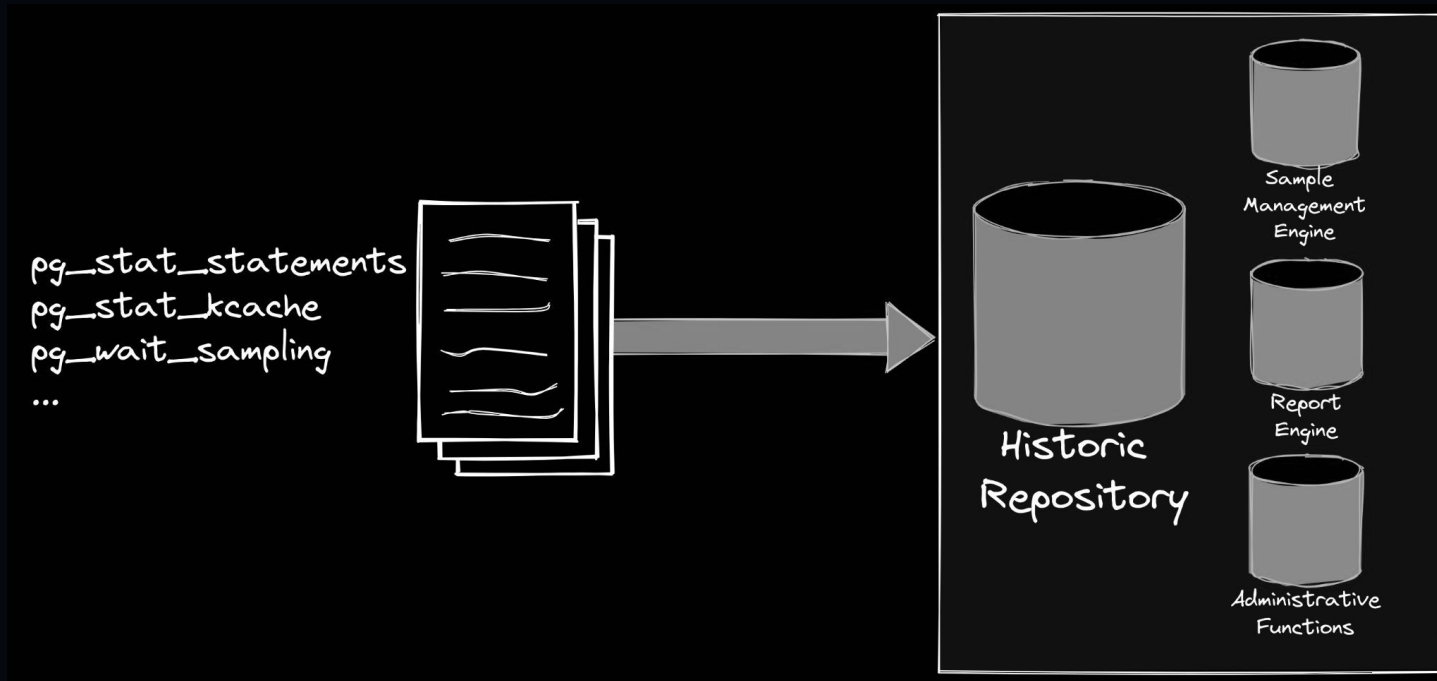
Top tables by estimated sequentially scanned volume

DB	Tablespace	Schema	Table	-SeqPages	SeqScan	IdxScan	IdxTf	Ins	Upd	Del	(p8)(B0T)
postgres	pg_default	public	pgbench_branches	2460	18	5640	525422	525422	1	531060	528651
postgres	pg_default	public	pgbench_counters	45	38	530673	530673	10	531060	531089	531089
postgres	pg_default	public	pgbench_tables	28	28	11408	1598083	376	258	376	1
postgres	pg_default	public	pgbench_accounts	27	18	2184238	1842338	188080	531060	531060	524844
postgres	pg_default	public	pgbench_sequences	12	12	11	2977	8314	44	109	542
postgres	pg_default	pg_catalog	pg_class	12	18	52	4861	4567	7	5	
postgres	pg_default	profile	sample_stat_tables_total	7160	8	179	-2	8	8		
postgres	pg_default	profile	pg_stat_statements	5276	8	647	734	734	6	6	
postgres	pg_default	profile	pg_stat_statements_2	4752	8	172	172	46	96	48	14
postgres	pg_default	profile	pg_stat_statements_1	4320	8	54	2	38	38		
postgres	pg_default	pg_catalog	pg_proc	1200	8	10	3190	3081	1	115	
postgres	pg_default	pg_catalog	pg_proc(TOAST)	1840	8	10	1686	1615	3		
postgres	pg_default	pg_catalog	pg_index	1808	8	21					
postgres	pg_default	pg_catalog	pg_statements(TOAST)	736	4	40	119	48			
postgres	pg_default	profile	sample_statements	720	8	15	610	593			
postgres	pg_default	pg_catalog	pg_statements(TOAST)	528	8	219	219	41			
postgres	pg_default	pg_catalog	pg_index(TOAST)	352	8	1	18	18			
postgres	pg_default	profile	pg_statements	336	6	6	16	16			

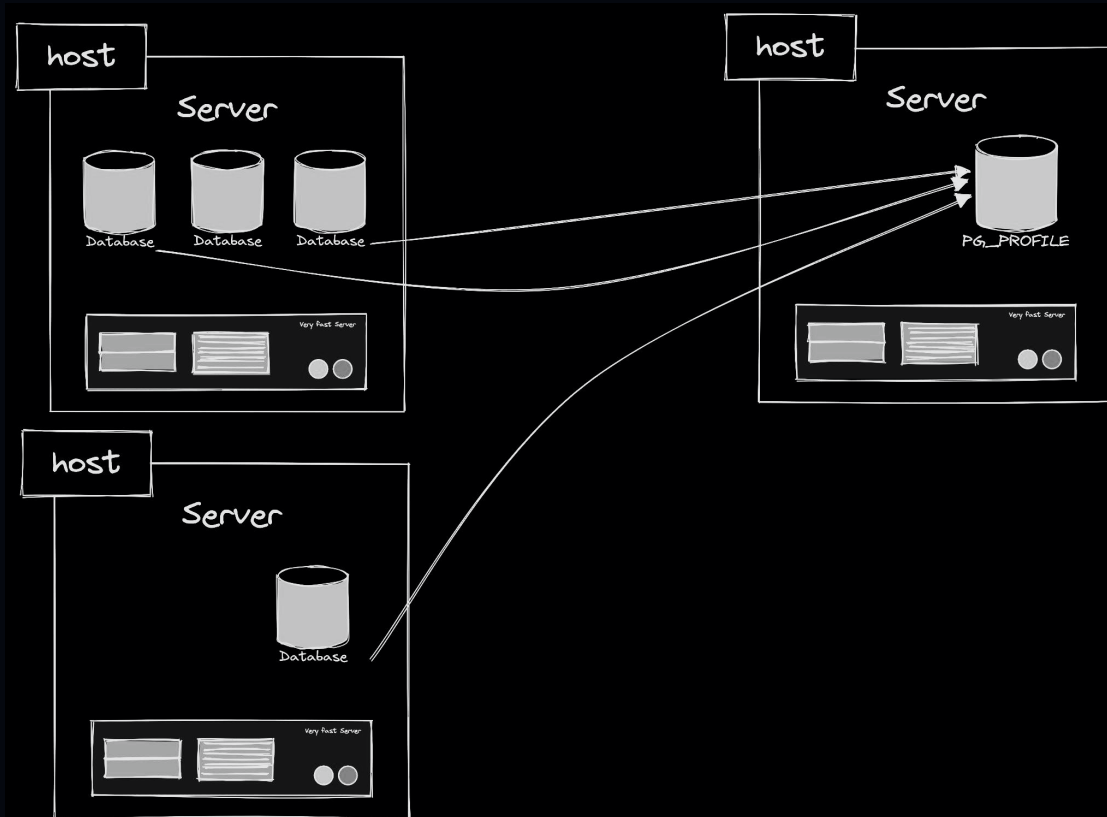
Top tables by blocks fetched

DB	Tablespace	Schema	Table	Heap	Idx	TOAST	TOAST-Idx
				Bkts	%Total	Bkts	%Total
postgres	pg_default	public	pgbench_branches	8077603	39.83	533675	2.65
postgres	pg_default	public	pgbench_counters	2178566	10.8	2142159	10.64
postgres	pg_default	public	pgbench_tables	1374600	6.77	530700	2.64
postgres	pg_default	profile	pg_stat_statements	1062277	5.2	29200	0.15
postgres	pg_default	public	pgbench_history	566530	2.76		
postgres	pg_default	profile	sample_stat_indexes	23764	1.2	5301	0.03
postgres	pg_default	pg_catalog	pg_class	5819	0.03	9655	0.05
postgres	pg_default	profile	pg_stat_indexes_2	6982	0.03	6640	0.03
postgres	pg_default	profile	tables_list	4379	0.02	6240	0.03
postgres	pg_default	pg_catalog	pg_statistic	3368	0.02	6304	0.03
postgres	pg_default	profile	sample_stat_tables	3094	0.02	4803	0.02
postgres	pg_default	profile	sample_stat_database	4222	0.02	2155	0.01
postgres	pg_default	pg_catalog	pg_index	3739	0.02	4221	0.02
postgres	pg_default	profile	sample_stat_indexes_1	1555	0.01	1545	0.01
postgres	pg_default	profile	pg_stat_statements_2	1588	0.01	623	
postgres	pg_default	profile	pg_stat_database_2	1481	0.01	740	
postgres	pg_default	pg_catalog	pg_proc	1175	0.01	1136	0.01
postgres	pg_default	pg_catalog	pg_statistics	736	0.01	14	33
postgres	pg_default	pg_catalog	pg_cast	173	0.00		
postgres	pg_default	pg_catalog	pg_statements	624	0.00		

Architecture



Architecture



Prerequisite and Setup

Extension `dblink` (part of contrib)

Repositories, e.g.

```
# sudo dnf install pg_profile_18
```

Direct from github

```
# curl -LJO  
https://github.com/zubkov-andrei/pg\_profile/releases/download/4.11/pg\_profile--4.11.tar.gz  
# sudo tar xzf pg_profile--4.11.tar.gz --directory $(pg_config --sharedir)/extension
```

Prerequisites

Create Schema for Repository (optional)

```
CREATE SCHEMA profile;
```

Activate necessary Extensions

```
CREATE EXTENSION pg_profile SCHEMA profile;  
CREATE EXTENSION dblink;
```

Prerequisites on source DBs

Preload Extensions (of your choice!)

Set few recommended Parameters

```
# vi $PGDATA/postgresql.conf  
  
...  
shared_preload_libraries = 'pg_stat_statements', ...  
  
...  
track_activities = on  
track_counts = on  
track_io_timing = on  
track_wal_io_timing = on  
track_functions = all
```

Configuration

Consider extension parameters

```
pg_profile.topn = 20  
pg_profile.max_sample_age = 7  
pg_profile.track_sample_timings = off  
pg_profile.max_query_length = 20000
```

As well for the related extensions, like e.g.

```
pg_stat_statements.max = 10000  
pg_stat_statements.track = 'all'
```

Adding Clusters/Servers for Collection

Add Server/Database

```
SELECT profile.create_server('cluster18', 'host=node0,dbname=postgres  
port=5432');
```

Other functions

```
profile.drop_server(server name)  
profile.enable_server(server name)  
profile.disable_server(server name)  
profile.show_servers()  
...
```

Collecting Data

Take a sample

```
select * from profile.take_sample();  
select * from profile.take_sample('core18');
```

Check existing samples

```
select * from profile.show_samples();  
select * from profile.show_samples('core18');
```

Baselines

Tagged Group of Samples

Independent Retention

e.g. for bulk operations, load testings,...

Example handling

```
select * from profile.show_baselines();  
select * from profile.create_baseline('cluster18', 'pgbench_run', 70, 71);
```

Collecting Data

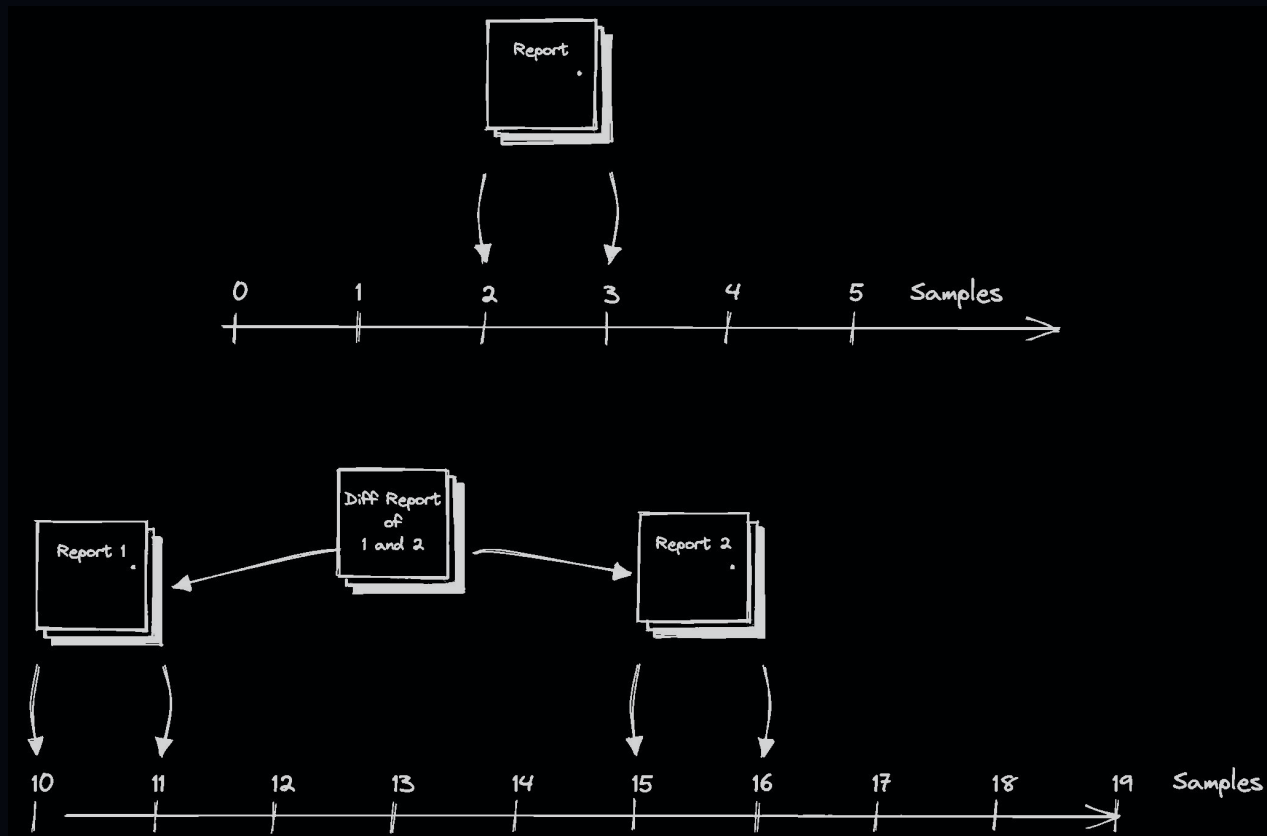
Best Practice Strategy

Frequented 30 Min Samples, starting point
Consider manual created Samples
Baselines

Putting into cron

```
*/30 * * * * psql -c 'SELECT profile.take_sample()' > /dev/null 2>&1
```

Creating Reports



Creating Reports

Standard Report

```
psql -Aqtc \  
"SELECT profile.get_report('core18',8, 9)" \  
-o report_8_9.html
```

Diff Report Report

```
psql -Aqtc \  
"SELECT profile.get_diffreport('core18', 8, 9, 11, 12)" \  
-o diff_report_8_9-11_12.html
```

A look into the repository schema

```

# \dt profile.*
List of relations
Schema | Name | Type | Owner
-----+-----+-----+-----
profile | baselines | table | postgres
profile | bl_samples | table | postgres
profile | funcs_list | table | postgres
profile | import_queries | table | postgres
profile | import_queries_version_order | table | postgres
profile | indexes_list | table | postgres
profile | last_stat_archiver | table | postgres
profile | last_stat_cluster | table | postgres
profile | last_stat_database | partitioned table | postgres
profile | last_stat_database_srv1 | table | postgres
profile | last_stat_database_srv2 | table | postgres
profile | last_stat_database_srv4 | table | postgres
profile | last_stat_indexes | partitioned table | postgres
profile | last_stat_indexes_srv1 | table | postgres
profile | last_stat_indexes_srv2 | table | postgres
profile | last_stat_indexes_srv4 | table | postgres
profile | last_stat_kcache | partitioned table | postgres
profile | last_stat_kcache_srv1 | table | postgres
profile | last_stat_kcache_srv2 | table | postgres
profile | last_stat_kcache_srv4 | table | postgres
profile | last_stat_statements | partitioned table | postgres
profile | last_stat_statements_srv1 | table | postgres
profile | last_stat_statements_srv2 | table | postgres
profile | last_stat_statements_srv4 | table | postgres
profile | last_stat_tables | partitioned table | postgres
profile | last_stat_tables_srv1 | table | postgres
profile | last_stat_tables_srv2 | table | postgres
profile | last_stat_tables_srv4 | table | postgres
profile | last_stat_tablespaces | partitioned table | postgres
profile | last_stat_tablespaces_srv1 | table | postgres
profile | last_stat_tablespaces_srv2 | table | postgres
profile | last_stat_tablespaces_srv4 | table | postgres
...
profile | last_stat_user_functions | partitioned table | postgres
profile | last_stat_user_functions_srv1 | table | postgres
profile | last_stat_user_functions_srv2 | table | postgres
profile | last_stat_user_functions_srv4 | table | postgres
profile | last_stat_wal | table | postgres
profile | report | table | postgres
profile | report_static | table | postgres
profile | report_struct | table | postgres
profile | roles_list | table | postgres
profile | sample_kcache | table | postgres
profile | sample_kcache_total | table | postgres
profile | sample_settings | table | postgres
profile | sample_stat_archiver | table | postgres
profile | sample_stat_cluster | table | postgres
profile | sample_stat_database | table | postgres
profile | sample_stat_indexes | table | postgres
profile | sample_stat_indexes_total | table | postgres
profile | sample_stat_tables | table | postgres
profile | sample_stat_tables_total | table | postgres
profile | sample_stat_tablespaces | table | postgres
profile | sample_stat_user_func_total | table | postgres
profile | sample_stat_user_functions | table | postgres
profile | sample_stat_wal | table | postgres
profile | sample_statements | table | postgres
profile | sample_statements_total | table | postgres
profile | sample_timings | table | postgres
profile | samples | table | postgres
profile | servers | table | postgres
profile | stmt_list | table | postgres
profile | tables_list | table | postgres
profile | tablespaces_list | table | postgres
profile | wait_sampling_total | table | postgres
(64 rows)

```

A look into the repository schema

Global Retention Policy

```
pg_profile.max_sample_age
```

Server Retention Policy

```
pg_profile.set_server_max_sample_age()
```

A look into the repository schema

Data Growth?

```
WITH schemas AS ( SELECT
    schemaname as name, sum(pg_relation_size(quote_ident(schemaname) || '.' ||
    quote_ident(tablename)))::bigint as size
FROM
    pg_tables GROUP BY schema name),db AS ( SELECT pg_database_size(current_database()) AS size
) SELECT schemas.name, pg_size_pretty(schemas.size) as absolute_size,
schemas.size::float / (SELECT size FROM db) * 100 as relative_size FROM schemas;
```

name	absolute_size	relative_size
public	41.00 MB	51.55
pg_catalog	0.50 MB	6.02
information_schema	0.09 MB	0.10
profile	18.00 MB	22.04

(4 rows)

Conclusion - My 2 Cents

Own Job handling would be nice, but cron is fine!

Availability in DBaaS offerings or contrib
...but usage always possible with self maintained repo db

Still room for even more information pieces in reports



Conclusion - What is missing?

Extensibility is a benefit, not a workaround!


Very tidy, pragmatic way to handle performance data

Sustainable repository approach

Not exactly the amount like the Oracle packs

But still a big point for considering Oracle folks on migrations

But perfect for nearly all common problems





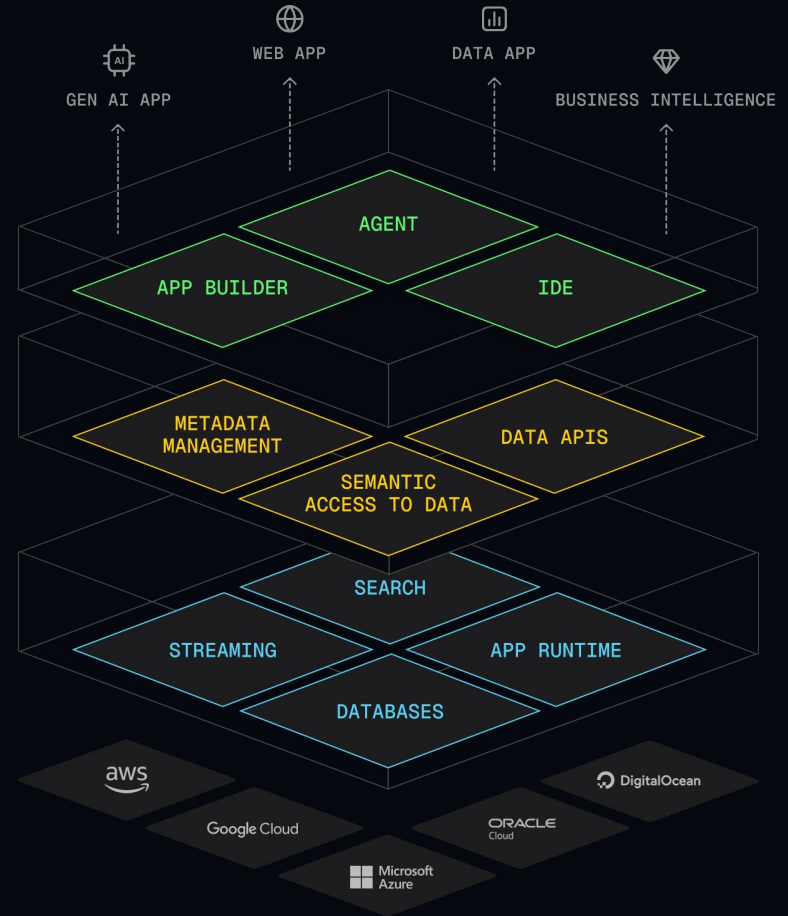
Thanks!

Questions?



YOUR AI-READY OPEN SOURCE DATA PLATFORM

Streaming | Database Optimization |
Analytics | Search | Data Warehousing |
In-Memory Caching



AIVEN PLATFORM

UNIFIED PLATFORM

STREAM	DATABASES	DEPLOY
<ul style="list-style-type: none">Apache Kafka®Apache Kafka® ConnectApache Kafka® MirrorMaker 2KarapaceKlaw	<ul style="list-style-type: none">PostgreSQL®MySQLValkeyDataHub®Apps®ClickHouse®OpenSearch®MetricsGrafana®	<ul style="list-style-type: none">awsGoogle CloudMicrosoft AzureORACLE CloudDigitalOceanUpCloudBYOC
THIRD-PARTY INTEGRATION		TOOLING
<ul style="list-style-type: none">DatadogPrometheusAmazon CloudWatchGCP MonitoringMongoDBAWS S3GCP BigQueryCouchbaseSnowflakeSplunkSumologicDebeziumGoogle Pub/SubGCS		<ul style="list-style-type: none">Terraformkubernetesavien CONSOLECLI API

TRUSTED BY OVER 1000 **CUSTOMERS** WORLDWIDE

AVAYA

idealo

Wolt

priceline®

LARE
DOU
TEU

Norauto

DECATHLON

SWIFT
SOLUTIONS

energy

MIRAKL

GOV.UK

DOJO

spare

Schibsted

TOYOTA

paf

CONRAD

adeo

CLAROTY

WÄRTSILÄ