

# Monitoring and Scaling Postgres at Datadog

Seth Rosenblum, Datadog

@SethRosenblum



Seth Rosenblum  
Data Reliability Engineering Lead

- Kafka
- Elasticsearch
- Cassandra
- Postgres!

@SethRosenblum





@datadoghq

SaaS-based monitoring

Trillions of data points per day

We're hiring!

<https://jobs.datadoghq.com>



Collecting data is cheap; not having it when you need it can be expensive



Collecting data is cheap; not  
having it when you need it can  
be expensive

...So instrument all the things!



# What metrics do we gather?

connections  
commits  
rollbacks  
disk\_read  
buffer\_hit  
rows\_returned  
rows\_fetched  
rows\_inserted  
rows\_updated  
rows\_deleted  
database\_size  
deadlocks  
temp\_bytes

temp\_files  
bgwriter.checkpoints\_timed  
bgwriter.checkpoints\_requested  
bgwriter.buffers\_checkpoint  
bgwriter.buffers\_clean  
bgwriter.maxwritten\_clean  
bgwriter.buffers\_backend  
bgwriter.buffers\_alloc  
bgwriter.buffers\_backend\_fsync  
bgwriter.write\_time  
bgwriter.sync\_time  
locks  
seq\_scans

seq\_rows\_read  
index\_scans  
index\_rows\_fetched  
rows\_hot\_updated  
live\_rows  
dead\_rows  
index\_rows\_read  
table\_size  
index\_size  
total\_size  
table.count  
max\_connections  
percent\_usage\_connections

replication\_delay  
replication\_delay\_bytes  
heap\_blocks\_read  
heap\_blocks\_hit  
index\_blocks\_read  
index\_blocks\_hit  
toast\_blocks\_read  
toast\_blocks\_hit  
toast\_index\_blocks\_read  
toast\_index\_blocks\_hit





# Scaling PostgreSQL at Datadog





Moar Resources!





Moar Instances!



Writes



Repl



Reads



Writes



Repl



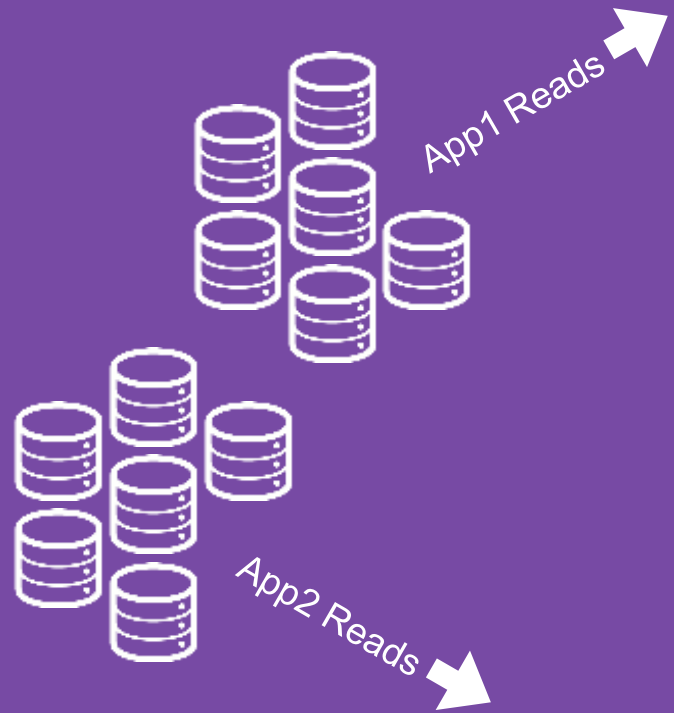
Reads



“Postgres performance-optimizes a lot better when it has a consistent workload”

Josh Berkus





How we do it

# Requirements

- ▶ Write master is writeable, read replicas are readable!



# Requirements

- ▶ Write master is writeable, read replicas are readable!
- ▶ Read replicas are up to date and don't lag



# Requirements

- ▶ Write master is writeable, read replicas are readable!
- ▶ Read replicas are up to date and don't lag
- ▶ Additional read replicas can be provisioned quickly



How we do it

# Solutions

- ▶ PostgreSQL!
  - ▶ <http://bit.ly/pg-repl-docs>
- ▶ WAL-E
  - ▶ <https://github.com/wal-e/wal-e>



How we do it

# Solutions

- ▶ PostgreSQL!

- ▶ <http://bit.ly/pg-repl-docs>

- ▶ ~~WAL-E~~ WAL-G

- ▶ <https://github.com/wal-g/wal-g>



# Requirements

- ▶ Write master is writeable, read replicas are readable!
- ▶ Read replicas are up to date and don't lag
- ▶ Additional read replicas can be provisioned quickly



# What are we alerting on?

- ▶ Write master is writeable, read replicas are readable!
  - ▶ Up/Down checks
  - ▶ Latency



# What are we alerting on?

- ▶ Read replicas are up to date and don't lag
  - ▶ Write master standby availability
  - ▶ Write master standby replication lag
  - ▶ Read replica lag



### [replica] replication delay (s)



# What are we alerting on?

- ▶ Additional read replicas can be provisioned quickly
  - ▶ Base backups are functioning properly





[i-0c3c018173bac31be] base-backup failed in 3524.65s

>>>> CMD <<<<

```
envdir /etc/wal-e/ /var/lib/postgresql/wal_e_venv/bin/wal-e backup-push
```

>>>> EXIT CODE <<<<

1

>>>> STDERR <<<<


```
botocore.vendored.requests.packages.urllib3.connectionpool INFO Starting new HTTP
connection (1): 169.254.169.254
botocore.vendored.requests.packages.urllib3.connectionpool INFO Starting new HTTP
connection (1): 169.254.169.254
botocore.vendored.requests.packages.urllib3.connectionpool INFO Starting new HTTPS
connection (1): s3.amazonaws.com
botocore.vendored.requests.packages.urllib3.connectionpool INFO Starting new HTTPS
connection (2): s3.amazonaws.com
botocore.vendored.requests.packages.urllib3.connectionpool INFO Resetting dropped
connection: s3.amazonaws.com
```



When do we care that a backup has failed?



# [master] Base backup age



22.14 hr	Avg: 11.85 hr	dd.postgres.base_backup_age	{role:ctx-pshard-master,tag:active,tag:failover,tag:master_id:ctx-pshard33}
17.93 hr	Avg: 12.13 hr	dd.postgres.base_backup_age	{role:ctx-pshard-master,tag:active,tag:master_id:ctx-pshard0}
19.4 hr	Avg: 11.88 hr	dd.postgres.base_backup_age	{role:ctx-pshard-master,tag:active,tag:master_id:ctx-pshard1}
20.14 hr	Avg: 11.88 hr	dd.postgres.base_backup_age	{role:ctx-pshard-master,tag:active,tag:master_id:ctx-pshard10}
21.84 hr	Avg: 12.07 hr	dd.postgres.base_backup_age	{role:ctx-pshard-master,tag:active,tag:master_id:ctx-pshard11}
21.93 hr	Avg: 12.06 hr	dd.postgres.base_backup_age	{role:ctx-pshard-master,tag:active,tag:master_id:ctx-pshard12}





## [Triggered] [postgres] ctx-pshard-master base backup is too old on master\_id:ctx-pshard14

#account:prod ...

The last wal-e base backup taken by this host is too old, it should not exceed 24 hours by more than a few minutes. It may be failing or just taking a long time. If we failed over this postgres master recently, it may have interrupted a base backup.

1. See if the process is still running on the host `ps awux | grep backup-push`. If it is it may just be taking unusually long.
2. Try running the command in `/etc/cron.d/wal_e_base_backup`. Run it in a screen as it may take a few hours, but it will let you know what's failing.

More information can be found here: <https://github.com/DataDog/devops/wiki/Postgres#backup>  
slack-bourbon-ops  
pagerduty-Ops-Bourbon





# Monitoring to Improve Performance



# Slow Queries

Name	Hits	Avg Latency ↓	Total time
☆ WITH sub_contexts SELECT key,...	90	2.17 s	196 s
☆ WITH sub_contexts SELECT key,...	58	1.49 s	86.7 s
☆ WITH sub_contexts SELECT key,...	49	1.09 s	53.2 s
☆ WITH sub_contexts SELECT key,...	194	759 ms	147 s
☆ WITH sub_contexts SELECT key,...	45	750 ms	33.7 s
☆ WITH sub_contexts SELECT key,...	22	740 ms	16.3 s
☆ WITH sub_contexts SELECT key,...	70	559 ms	39.1 s



# Slow Queries

Name	Hits	Avg Latency ↓	Total time
☆ WITH sub_contexts SELECT key,...	90	2.17 s	196 s
☆ WITH sub_contexts SELECT key,...	58	1.49 s	86.7 s
☆ WITH sub_contexts SELECT key,...	49	1.09 s	53.2 s
☆ WITH sub_contexts SELECT key,...	194	759 ms	147 s
☆ WITH sub_contexts SELECT key,...	45	750 ms	33.7 s
☆ WITH sub_contexts SELECT key,...	22	740 ms	16.3 s
☆ WITH sub_contexts SELECT key,...	70	559 ms	39.1 s



# Performance: RAM vs Disk

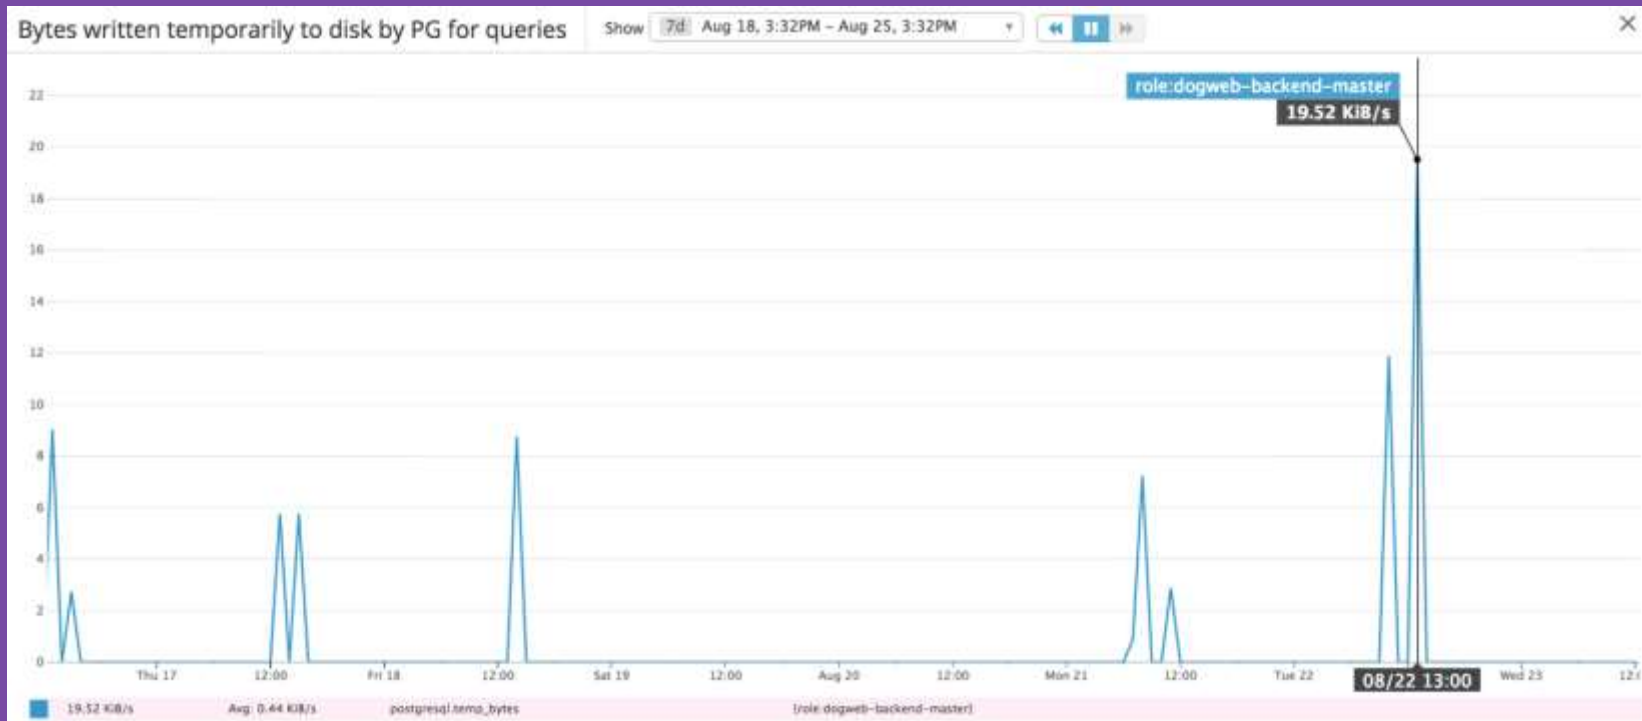


“Aside from `shared_buffers`, the most important memory-allocation parameter is `work_mem`... Raising this value can dramatically improve the performance of certain queries...”

Robert Haas



# Finding **\*\*Inefficient\*\*** Queries



# Finding **\*\*Inefficient\*\*** Queries



# EXPLAIN ANALYZE

<http://bit.ly/pg-explain>

- ▶ Explain displays the execution plan



# EXPLAIN ANALYZE

<http://bit.ly/pg-explain>

- ▶ Explain displays the execution plan
- ▶ Analyze runs it and gathers stats



# EXPLAIN ANALYZE

```
Merge Right Join (cost=25870.55..31017.51 rows=229367 width=92) (actual time=2884.501..5147.047 rows=354834 loops=1)
  Merge Cond: (a.uid = b.uid)
    -> Index Scan using foo on bar a (cost=0.00..537.29 rows=9246 width=27) (actual time=0.049..41.782 rows=9246 loops=1)
    -> Materialize (cost=25870.49..27204.80 rows=106745 width=81) (actual time=2884.413..3804.537 rows=354834 loops=1)
      -> Sort (cost=25870.49..26137.35 rows=106745 width=81) (actual time=2884.406..3099.732 rows=111878 loops=1)
        Sort Key: b.uid
        Sort Method: external merge  Disk: 8928kB
...
Total runtime: 5588.105 ms
(14 rows)
```

<http://bit.ly/pg-auto-explain>




“Aside from `shared_buffers`, the most important memory-allocation parameter is `work_mem`... Raising this value can dramatically improve the performance of certain queries, but it's important not to overdo it.”

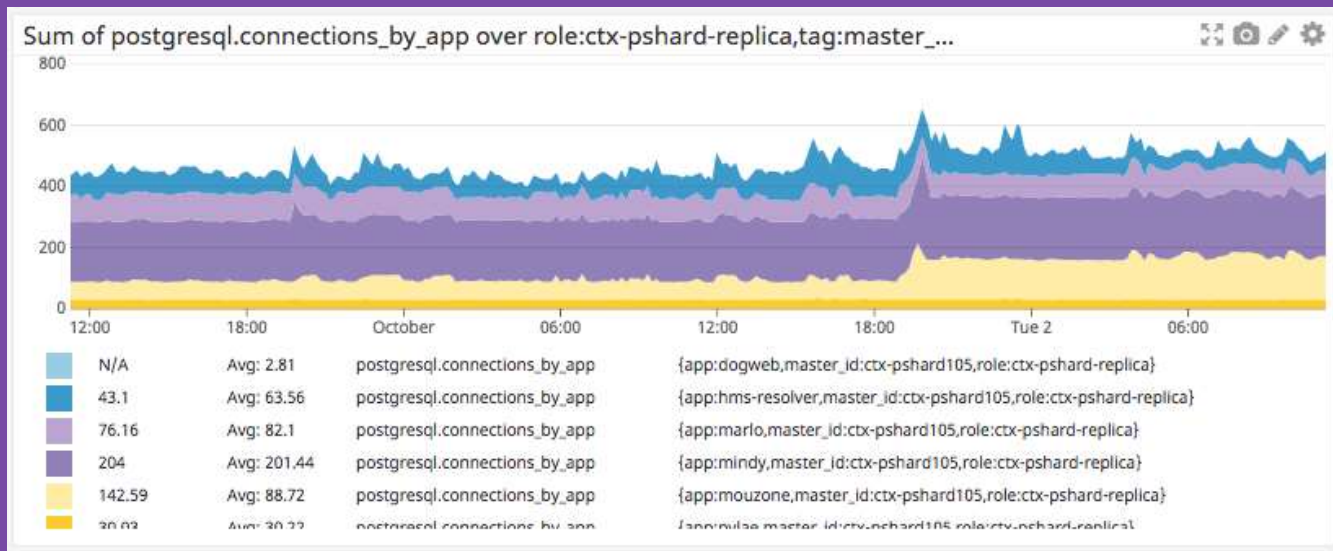
Robert Haas



# Track Connections



# Connections By Application



# Summary

1. Collect as many metrics as you can, before you need them
2. If the metrics that you have aren't providing the right value, build ones that do
3. Be aggressive in monitoring slow queries, catch them while they're easy to find



# Resources

- ▶ <http://dtdg.co/monitor-postgres>
- ▶ <https://dtdg.co/gcp-sql>
- ▶ <https://dtdg.co/postgresql-vacuums>



# Questions?

Seth Rosenblum

@SethRosenblum

seth@datadoghq.com

