



# Using Prometheus and Grafana to build a Postgres Dashboard

Gregory Stark

October 25, 2018

# What is Monitoring?





- We want to alert on global properties such as
  - The fraction of the fleet currently operating well
  - The average response time across the fleet
  - The consistency of the data across the fleet
- We want to alert based on historical data
  - Average rates over time period
  - Compare current data with 24h ago or 7d ago
- We want to alert on comparisons between services
  - Ratio of rates of transactions in database to application requests
  - Are there any database servers for which S3 does not contain a recent backup

- Prometheus

*Database specifically designed for handling time series. It performs recorded queries regularly to synthesize new time series and to generate alerts.*

- Alertmanager

*Part of Prometheus project. Handles generating notifications for alerts.*

- node\_exporter

*Agent for system statistics. For more agents see:  
<https://prometheus.io/docs/instrumenting/exporters/>*

- postgres\_exporter

*Agent that exports statistics from pg\_stat\_\* views*

- mtail

*Useful to fill gaps where Postgres doesn't provide a statistics views to expose them. e.g. log\_min\_duration, log\_lock\_waits*

- Grafana

*WYSIWYG dashboard software.*

PostgreSQL

```
graph TD; PG[PostgreSQL] --> Tables[pg_stat_* tables];
```

pg\_stat\_activity

pg\_stat\_replication

pg\_stat\_wal\_receiver

pg\_stat\_subscription

pg\_stat\_ssl

pg\_stat\_progress\_vacuum

pg\_stat\_archiver

pg\_stat\_bgwriter

pg\_stat\_database

pg\_stat\_database\_conflicts

pg\_stat\_all\_tables


pg\_stat\_all\_indexes


pg\_stat\_user\_functions

pg\_statio\_all\_tables


pg\_statio\_all\_indexes

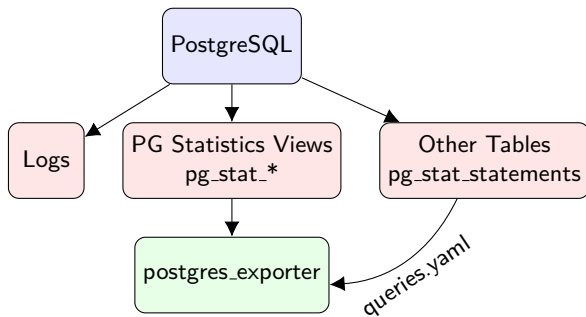
pg\_statio\_all\_sequences

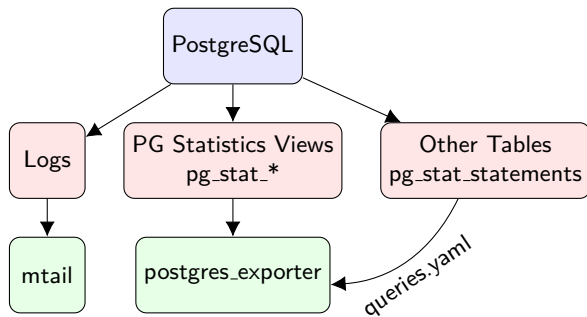


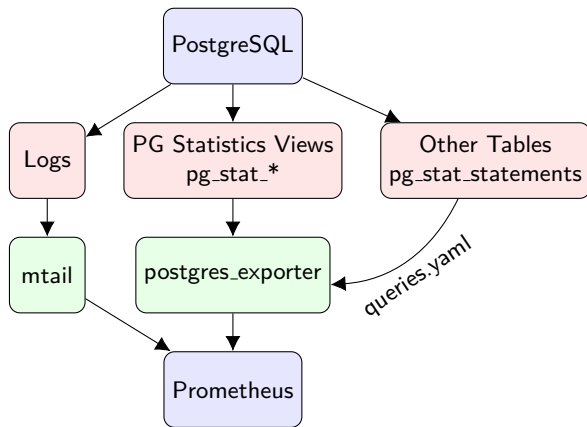


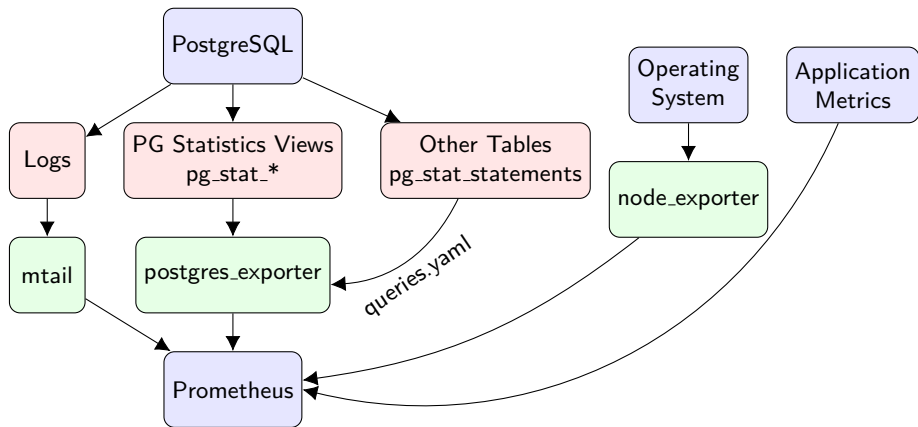


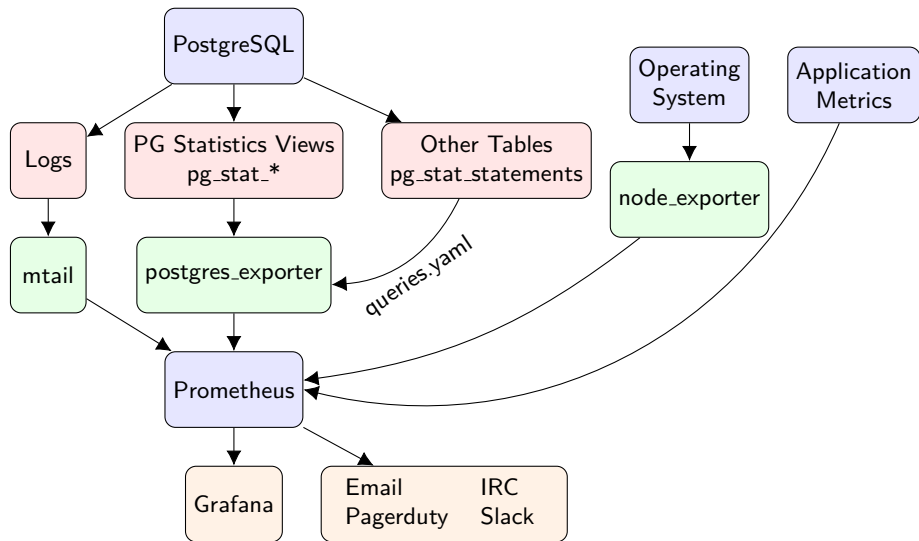












The USE method uses three key metrics for each component of a complex system:

- Utilization
- Saturation
- Errors

It was published in ACMQ as Thinking Methodically about Performance (2012):

<https://queue.acm.org/detail.cfm?id=2413037>

Further discussion:

<http://www.brendangregg.com/usemethod.html>

Presented at FISL13:

<http://dtrace.org/blogs/brendan/2012/09/21/fisl13-the-use-method/>

The RED model uses latency (duration) instead of utilization:

- Rate
- Errors
- Duration

From:

<https://www.weave.works/blog/the-red-method-key-metrics-for-microservices-architecture/>

See also:


<https://www.vividcortex.com/blog/monitoring-and-observability-with-use-and-red>



SRE Golden Signals are very similar:

- Latency
- Traffic
- Errors
- Saturation

Originally published in Site Reliability Book:



Edited by Betsy Beyer, Chris Jones,  
Jennifer Petoff & Niall Richard Murphy

Also see discussion at:

<https://medium.com/devopslinks/how-to-monitor-the-sre-golden-signals-1391cad7524>

☐ Enable query history

pg\_stat\_activity\_count

Load time: 290ms  
Resolution: 14s  
Total time series: 24

Execute

- insert metric at cursor - ▾

Graph

Console

Element	Value
pg_stat_activity_count{datname="postgres",environment="prd",instance="localhost:9187",job="postgres",state="active",tier="db",type="postgres"}	0
pg_stat_activity_count{datname="postgres",environment="prd",instance="localhost:9187",job="postgres",state="disabled",tier="db",type="postgres"}	0
pg_stat_activity_count{datname="postgres",environment="prd",instance="localhost:9187",job="postgres",state="fastpath function call",tier="db",type="postgres"}	0
pg_stat_activity_count{datname="postgres",environment="prd",instance="localhost:9187",job="postgres",state="idle",tier="db",type="postgres"}	0
pg_stat_activity_count{datname="postgres",environment="prd",instance="localhost:9187",job="postgres",state="idle in transaction",tier="db",type="postgres"}	0
pg_stat_activity_count{datname="postgres",environment="prd",instance="localhost:9187",job="postgres",state="idle in transaction (aborted)",tier="db",type="postgres"}	0
pg_stat_activity_count{datname="stark",environment="prd",instance="localhost:9187",job="postgres",state="active",tier="db",type="postgres"}	1

☐ Enable query history

```
count(pg_stat_activity_count) by (datname)
```

Load time: 291ms

Resolution: 14s

Total time series: 4

Execute

- insert metric at cursor - ▾

Graph

Console

Element	Value
{datname="postgres"}	6
{datname="stark"}	6
{datname="template0"}	6
{datname="template1"}	6

[Remove Graph](#)

Add Graph

☐ Enable query history

```
sum(pg_stat_activity_count{environment="prd", tier="db", type="postgres"}) by (instance)
```

Load time: 272ms

Resolution: 14s

Total time series: 1

[Execute](#)[- insert metric at cursor - ▾](#)[Graph](#)[Console](#)

Element	Value
{instance="localhost:9187"}	2

[Remove Graph](#)[Add Graph](#)

```
groups:
- name: postgresql.rules
  rules:
  - alert: PostgreSQL_CommitRateTooLow
    expr: |
      rate(pg_stat_database_xact_commit{datname="gitlabhq_production",
        environment="prd"}[1m]) < 1000
    for: 2m
    labels:
      severity: warn
      channel: database
    annotations:
      description: |
        Commits/s on {{$labels.instance}} database {{$labels.datname}}
        is {{$value | printf "%.0f" }} which is implausibly low.
        Perhaps the application is unable to connect
      runbook: troubleshooting/postgresql.md#availability
      title: 'Postgres seems to be processing very few transactions'
```

```
groups:
- name: postgresql.rules
  rules:
  - alert: PostgreSQL_RollbackRateTooHigh
    expr: |
      rate(pg_stat_database_xact_rollback{datname="gitlabhq_production"}[5m])
        / ON(instance, datname)
      rate(pg_stat_database_xact_commit{datname="gitlabhq_production"}[5m])
        > 0.02
    for: 5m
    labels:
      severity: warn
      channel: database
    annotations:
      description: |
        Ratio of transactions being aborted compared to committed is
        {{ $value | printf "%.2f" }} on {{ $labels.instance }}
      runbook: troubleshooting/postgresql.md#errors
      title: 'Postgres transaction rollback rate is high'
```

```
groups:
- name: postgresql.rules
  rules:
  - alert: PostgreSQL_ConnectionsTooHigh
    expr: |
      sum(pg_stat_activity_count) BY (environment, instance)
        > ON(instance)
        pg_settings_max_connections * 0.75
    for: 10m
    labels:
      severity: warn
      channel: database
    annotations:
      runbook: troubleshooting/postgresql.md#connections
      title: |
        Postgres has {{$value}} connections on {{$labels.instance}}
        which is close to the maximum
```

# Alerts - more Errors

```
# Count of specific types of errors -- notably statement timeouts
counter postgresql_logs_total by severity
counter postgresql_errors_total by type

/^[0-9.:_-]* [a-z0-9-]* postgresql: (?P<date>\d\d\d\d-\d\d-\d\d \d\d:\d\d:\d\d [A-Z]{3}) \[[0-9]*\]:
 \[[0-9]*-1\] (?P<severity>DEBUG[1-5]|INFO|NOTICE|WARNING|ERROR|LOG|FATAL|PANIC): / {
    postgresql_logs_total[$severity]++

    /ERROR: (?P<message>.*)$/ {
        /canceling statement due to statement timeout/ {
            postgresql_errors_total["statement_timeout"]++
        }
        /canceling autovacuum task/ {
            postgresql_errors_total["canceled_autovacuum"]++
        }
        /deadlock detected/ {
            postgresql_errors_total["deadlock_detected"]++
        }
        /duplicate key value violates unique constraint/ {
            postgresql_errors_total["duplicate_key"]++
        }
        otherwise {
            postgresql_errors_total["other"]++
        }
    }
}
```



```
groups:
- name: postgresql.rules
  rules:
  - alert: PostgreSQL_StatementTimeout_Errors
    expr: |
      rate(postgresql_errors_total{type="statement_timeout"}[1m]) > 0.5
    for: 5m
    labels:
      severity: warn
      channel: database
    annotations:
      description: |
        Database {{$labels.instance}} is logging
        {{ $value | printf "%.1f" }} statement timeouts per second
      runbook: troubleshooting/postgresql.md#errors
      title: 'Postgres transactions showing high rate of statement timeouts'
```

```
pg_replication:
  query: |
    SELECT EXTRACT(epoch FROM (
      now() - pg_last_xact_replay_timestamp()
    ))::int AS lag,
    CASE WHEN pg_is_in_recovery() THEN 1 ELSE 0 END AS is_replica
  metrics:
    - lag:
        usage: "GAUGE"
        description: "Replication lag behind primary in seconds"
    - is_replica:
        usage: "GAUGE"
        description: "Indicates if this host is a replica"
```

```
groups:
- name: postgresql.rules
  rules:
  - alert: PostgreSQL_ReplicationLagTooLarge
    expr: |
      (pg_replication_lag > 120)
      AND ON(instance)
      (pg_replication_is_replica == 1)
    annotations:
      description: |
        Replication lag on server {{$labels.instance}} is currently
        {{$value | humanizeDuration }}
      runbook: troubleshooting/postgres.md#replication-is-lagging-or-has-stopped
      title: 'Postgres Replication lag is over 2 minutes'
```

# Alerts - Exposing hidden problems

```
pg_replication_slots:
```

```
  query: |
```

```
    SELECT slot_name, slot_type,  
           case when active then 1.0 else 0.0 end AS active,  
           age(xmin) AS xmin_age,  
           age(catalog_xmin) AS catalog_xmin_age,
```

```
    FROM pg_replication_slots
```

```
  metrics:
```

```
    - slot_name:
```

```
      usage: "LABEL"
```

```
      description: "Slot Name"
```

```
    - slot_type:
```

```
      usage: "LABEL"
```

```
      description: "Slot Type"
```

```
    - active:
```

```
      usage: "GAUGE"
```

```
      description: "Boolean flag indicating whether this slot has a consumer streaming from it"
```

```
    - xmin_age:
```

```
      usage: "GAUGE"
```

```
      description: "Age of oldest transaction that cannot be vacuumed due to this replica"
```

```
    - catalog_xmin_age:
```

```
      usage: "GAUGE"
```

```
      description: "Age of oldest transaction that cannot be vacuumed from catalogs due to this replica (use
```

```
groups:
- name: postgresql.rules
  rules:
  - alert: PostgreSQL_UnusedReplicationSlot
    expr: 'pg_replication_slots_active == 0'
    for: 30m
    labels:
      severity: warn
      channel: database
    annotations:
      description: |
        Unused {{$labels.slot_type}} slot "{{$labels.slot_name}}"
        on {{$labels.instance}}
```

## Alerts - Exposing hidden problems

- alert: PostgreSQL\_SplitBrain  
expr: 'count(pg\_replication\_is\_replica == 0) BY (environment) != 1'  
annotations:  
  title: |  
    Split Brain: more than one postgres databases in environment  
    {{\$labels.environment}} in read-write (primary) mode
- alert: PostgreSQL\_SplitBrain\_Replicas  
expr: |  
  count(  
    count(pg\_stat\_wal\_receiver\_status >= 0) BY (environment, upstream\_host)  
  ) BY (environment) > 1  
annotations:  
  title: |  
    Split Brain: replicas in environment {{\$labels.environment}}  
    have different upstream databases configured


- alert: PostgreSQL\_FleetSizeChange  
expr: 'postgres:databases != postgres:databases OFFSET 2m'  
annotations:  
    description: 'There are now {{\$value}} databases in "{{\$labels.environment}}"  
    title: 'Number of PostgreSQL Databases in {{\$labels.environment}} has changed'
- alert: PostgreSQL\_RoleChange  
expr: 'pg\_replication\_is\_replica and changes(pg\_replication\_is\_replica[1m]) > 0'  
    title: 'Postgres Database replica promotion occurred in "{{\$labels.environment}}"
- alert: PostgreSQL\_ConfigurationChange  
expr: |  
    {\_\_name\_\_=~"pg\_settings\_.\*"} !=  
    ON(\_\_name\_\_, instance)  
    {\_\_name\_\_=~"pg\_settings\_.\*",\_\_name\_\_!="pg\_settings\_transaction\_read\_only"}  
    OFFSET 10m

# The GUI Dashboard



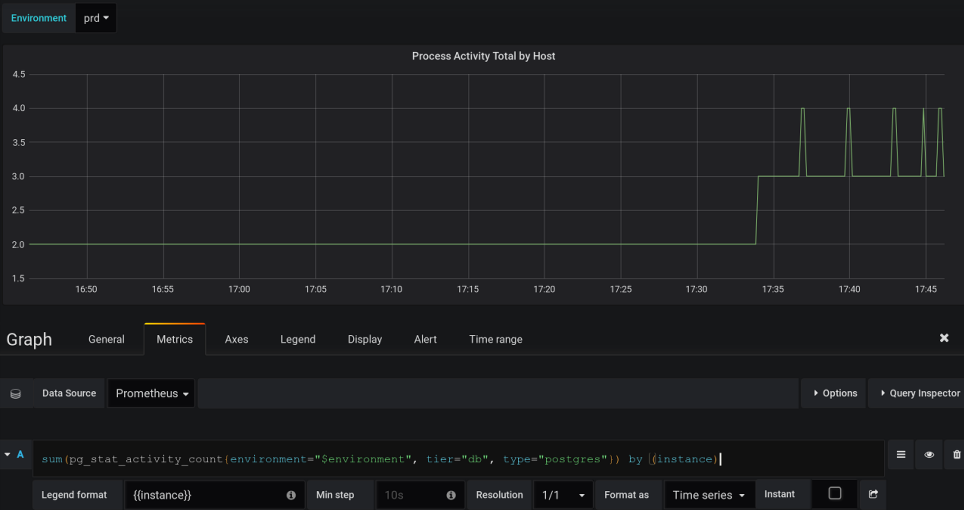


## Process Activity Total by Host



- postgres-01-db-gprd.c.gitlab-production.internal
- postgres-02-db-gprd.c.gitlab-production.internal
- postgres-03-db-gprd.c.gitlab-production.internal
- postgres-04-db-gprd.c.gitlab-production.internal
- postgres-05-db-gprd.c.gitlab-production.internal

# The GUI Dashboard



```
pg_stat_statements:
query: |
SELECT
    pg_get_userbyid(userid) as user,
    pg_database.datname,
    pg_stat_statements.queryid,
    pg_stat_statements.calls,
    pg_stat_statements.total_time as time_milliseconds,
    pg_stat_statements.rows,
    pg_stat_statements.shared_blks_hit,
    pg_stat_statements.shared_blks_read,
    pg_stat_statements.shared_blks_dirtied,
    pg_stat_statements.shared_blks_written,
    pg_stat_statements.local_blks_hit,
    pg_stat_statements.local_blks_read,
    pg_stat_statements.local_blks_dirtied,
    pg_stat_statements.local_blks_written,
    pg_stat_statements.temp_blks_read,
    pg_stat_statements.temp_blks_written,
    pg_stat_statements.blk_read_time,
    pg_stat_statements.blk_write_time
FROM pg_stat_statements
JOIN pg_database
    ON pg_database.oid = pg_stat_statements.dbid
```

```
metrics:
- user:
  usage: "LABEL"
  description: "The user who executed the statement"
- datname:
  usage: "LABEL"
  description: "The database in which the statement was executed"
- queryid:
  usage: "LABEL"
  description: "Internal hash code, computed from the statement's parse tree"
- calls:
  usage: "COUNTER"
  description: "Number of times executed"
- time_milliseconds:
  usage: "COUNTER"
  description: "Total time spent in the statement, in milliseconds"
- rows:
  usage: "COUNTER"
  description: "Total number of rows retrieved or affected by the statement"
- shared_blks_hit:
  usage: "COUNTER"
  description: "Total number of shared block cache hits by the statement"
- shared_blks_read:
  usage: "COUNTER"
  description: "Total number of shared blocks read by the statement"
- shared_blks_dirtied:
  usage: "COUNTER"
  description: "Total number of shared blocks dirtied by the statement"
- shared_blks_written:
  usage: "COUNTER"
  description: "Total number of shared blocks written by the statement"
- local_blks_hit:
  usage: "COUNTER"
  description: "Total number of local block cache hits by the statement"
```

This has some issues with Cardinality....

- 15 metrics
- for each of 5000 queryids (or more)
- for each database
- every 15s

This can quickly become performance issue for Prometheus.

Gitlab / PostgreSQL Top Queries -

Environment: prd Instance: All Database: stark User: stark Time Range: 10m

### Most Called Queries

datname	environment	Query ID	user	Calls/s
stark	prd	<a href="#">-8620724237793657752</a>	stark	575 ops
stark	prd	<a href="#">2397681704071010949</a>	stark	575 ops
stark	prd	<a href="#">-7810315603562552972</a>	stark	575 ops
stark	prd	<a href="#">432536356703194120</a>	stark	575 ops
stark	prd	<a href="#">-1391856223003275181</a>	stark	575 ops
stark	prd	<a href="#">-5636784524681433689</a>	stark	575 ops
stark	prd	<a href="#">6739068184171305142</a>	stark	575 ops
stark	prd	<a href="#">3888237132185418900</a>	stark	0 ops
stark	prd	<a href="#">6636882400379613332</a>	stark	0 ops

### Slowest Queries

Slowest Queries (Total Time Consumed per second)

datname	environment	Query ID	user	Time/s
stark	prd	<a href="#">-8620724237793657752</a>	stark	33.62 ms
stark	prd	<a href="#">-5636784524681433689</a>	stark	14.79 ms
stark	prd	<a href="#">6739068184171305142</a>	stark	12.59 ms
stark	prd	<a href="#">-1391856223003275181</a>	stark	6.41 ms
stark	prd	<a href="#">432536356703194120</a>	stark	5.78 ms
stark	prd	<a href="#">2397681704071010949</a>	stark	0.53 ms
stark	prd	<a href="#">1746008609887187818</a>	stark	0.46 ms
stark	prd	<a href="#">-5995353327424413110</a>	stark	0.36 ms

Slowest Queries (Average Runtime)

datname	environment	Query ID	user	Runtime
stark	prd	<a href="#">1746008609887187818</a>	stark	4.62 ms
stark	prd	<a href="#">-5995353327424413110</a>	stark	3.62 ms
stark	prd	<a href="#">8231563794068656526</a>	stark	2.83 ms
stark	prd	<a href="#">-4432434327813341498</a>	stark	2.72 ms
stark	prd	<a href="#">7213032506619474748</a>	stark	2.38 ms
stark	prd	<a href="#">786164931737574996</a>	stark	2.14 ms
stark	prd	<a href="#">-2701216582046352299</a>	stark	1.59 ms
stark	prd	<a href="#">-4492017811537561230</a>	stark	0.77 ms

Shared Blocks I/O

Gitlab / PostgreSQL Top Queries - Last 24 hours

datname	prn	prn	prn	prn
stark	prd	-4209465214022532407	stark	0 ops

Slowest Queries

Slowest Queries (Total Time Consumed per second)

datname	environment	Query ID	user	Time/s
stark	prd	-8620724237793657752	stark	33.62 ms
stark	prd	-5636784524681433689	stark	14.79 ms
stark	prd	6739068184171305142	stark	12.59 ms
stark	prd	-1391856223003275181	stark	6.41 ms
stark	prd	432536356703194120	stark	5.78 ms
stark	prd	2397681704071010949	stark	0.53 ms
stark	prd	1746008609887187818	stark	0.46 ms
stark	prd	-5995353327424413110	stark	0.36 ms

Slowest Queries (Average Runtime)

datname	environment	Query ID	user	Runtime
stark	prd	1746008609887187818	stark	4.62 ms
stark	prd	-5995353327424413110	stark	3.62 ms
stark	prd	8231563794068656526	stark	2.83 ms
stark	prd	-4432434327813341498	stark	2.72 ms
stark	prd	7213032506619474748	stark	2.38 ms
stark	prd	786164931737574996	stark	2.14 ms
stark	prd	-2701216582046352299	stark	1.59 ms
stark	prd	-4492017811537561230	stark	0.77 ms

Shared Blocks I/O

Highest shared\_bkls\_dirtied

datname	environment	Query ID	user	kb/s Dirtied
stark	prd	-8620724237793657752	stark	1.27 MB/s
stark	prd	432536356703194120	stark	29.30 kB/s
stark	prd	6739068184171305142	stark	1.88 kB/s
stark	prd	-5636784524681433689	stark	1.29 kB/s
stark	prd	-414016695874664502	stark	0 kB/s
stark	prd	-2701216582046352299	stark	0 kB/s

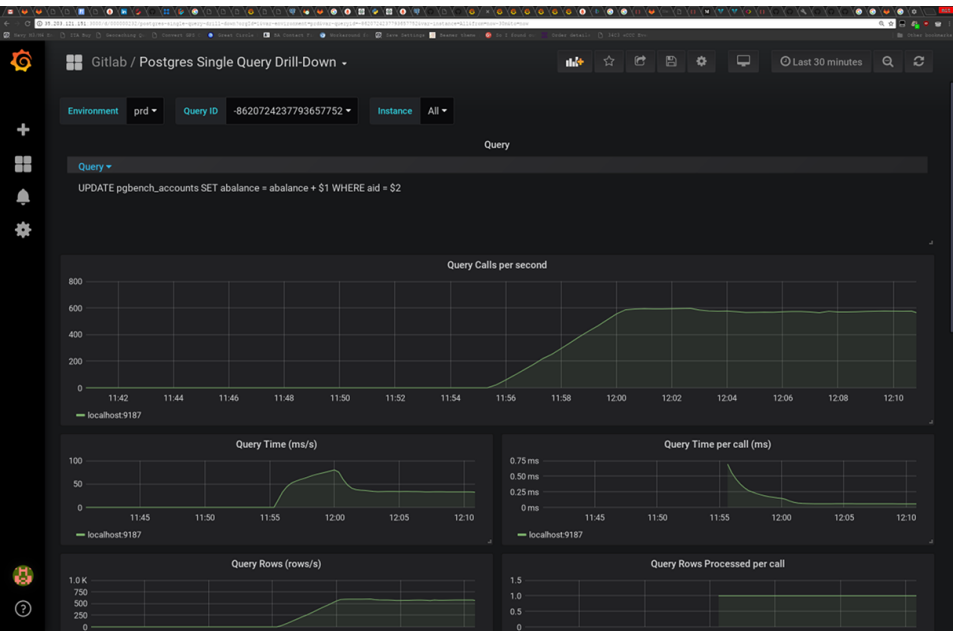
Highest shared\_bkls\_read

datname	environment	Query ID	user	kb/s Read
stark	prd	-8620724237793657752	stark	1.14 MB/s
stark	prd	432536356703194120	stark	29.30 kB/s
stark	prd	6739068184171305142	stark	1.15 kB/s
stark	prd	-5636784524681433689	stark	1.14 kB/s
stark	prd	-414016695874664502	stark	0 kB/s
stark	prd	-2701216582046352299	stark	0 kB/s

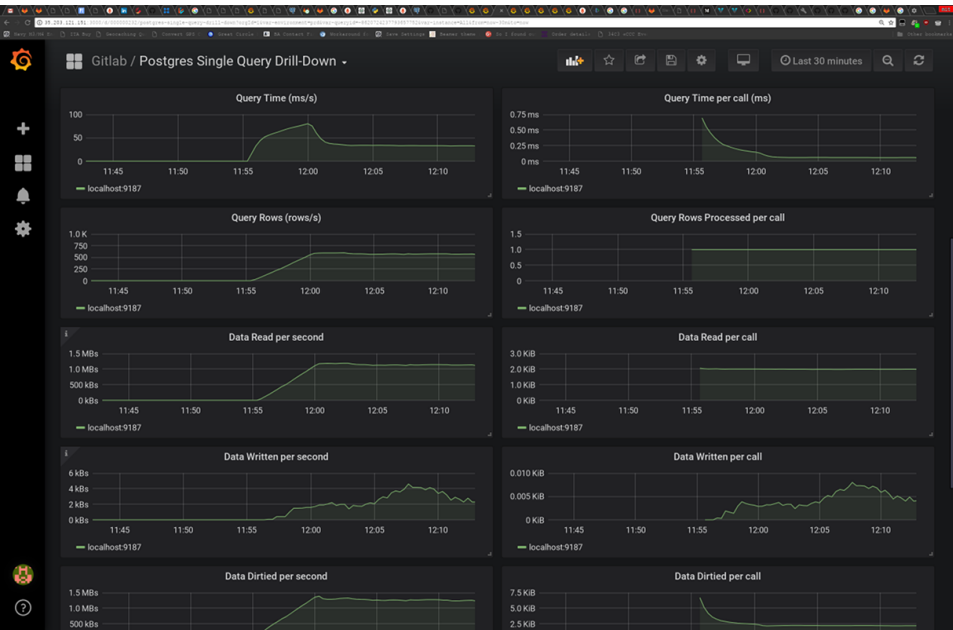
> Temp Data Read (2 panels)

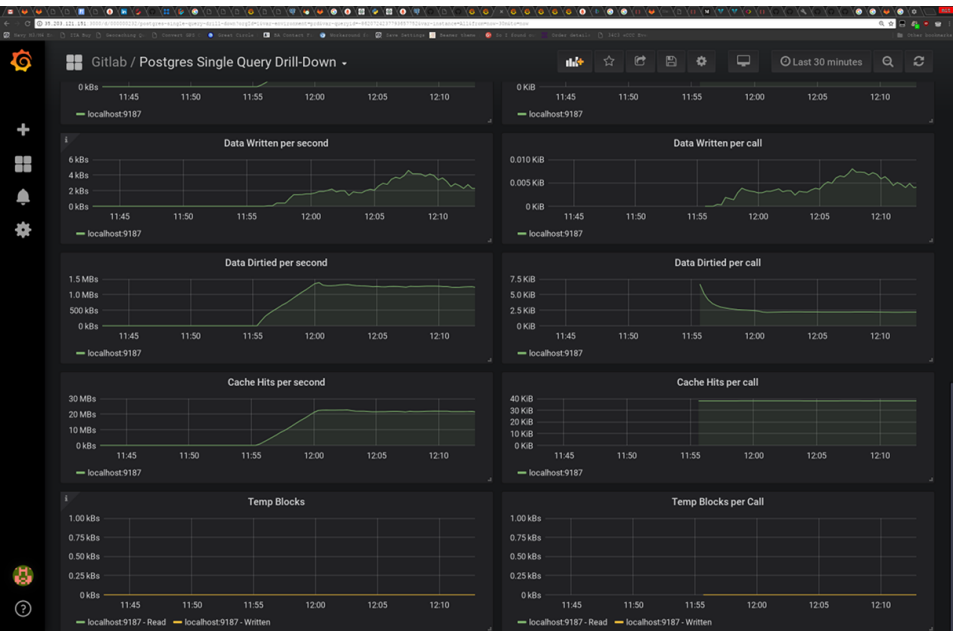
> Temp Data Written (2 panels)

> New Slow Queries (Very Slow) (1 panel)









- Missing stats in postgres\_exporter  
Queries.yaml requires a good understanding of Postgres \*and\* Prometheus to write. It also makes rules and dashboards non-portable which is a major downside. It's considered an "anti-pattern" in Prometheus exporter world.
- Missing data in pg\_stats\_\*: errors, lock timing, slow queries  
If you have more please tell me, I'll be working on adding these in the future.
- Saturation is basically impossible to measure in Postgres  
pg\_stat\_activity does not really represent saturation very well when applications keep persistent connections and use pooling of any form. If you filter on state=active it's useful but still very coarse and incomplete representation. There's a pgbouncer exporter as well and you can add instrumentation to your application to address this. But it would be good to identify standard ways of measuring Postgres saturation.

- postgres\_exporter should be eliminated entirely  
It would be much preferable to have Postgres speak common monitoring protocols directly. That would make the statistics more consistent, reliable, and easier to deploy.
- Distributed Tracing  
This is different from but complementary to monitoring and is a major gap that would help expose the connections between database metrics and application metrics.

- This presentation is online at:  
[https://\\_stark.gitlab.io/monitoring-postgres-pgconf.eu-2018/monitoring.pdf](https://_stark.gitlab.io/monitoring-postgres-pgconf.eu-2018/monitoring.pdf)
- Gitlab Project for presentation at:  
[https://gitlab.com/\\_stark/monitoring-postgres-pgconf.eu-2018](https://gitlab.com/_stark/monitoring-postgres-pgconf.eu-2018)
- Source code for presentation at:  
[https://gitlab.com/\\_stark/monitoring-postgres-pgconf.eu-2018.git](https://gitlab.com/_stark/monitoring-postgres-pgconf.eu-2018.git)
- Author contact address:  
[stark@mit.edu](mailto:stark@mit.edu)