



PostgreSQL monitoring with *pgwatch2*

Kaarel Moppel / PostgresConf US 2018

www.cybertec-postgresql.com

Why to monitor



- ▶ Failure / Downtime detection
- ▶ Slowness / Performance analysis
- ▶ Proactive predictions
- ▶ Maybe wasting money?

Different levels of Database monitoring



- ▶ Service availability level
- ▶ System level monitoring
- ▶ Database
- ▶ Application level would be nice also

- ▶ Log analysis
- ▶ Stats Collector
 - ▶ Dynamic views
 - ▶ `pg_stat_activity`, `pg_stat_(replication|wal_receiver)`,
`pg_locks`, `pg_stat_ssl`, `pg_stat_progress_vacuum`
 - ▶ Cumulative views
 - ▶ Most `pg_stat_*` views
 - ▶ Long uptimes cause “lag” for problem detection
 - ▶ NB! Not all `track_*` parameters enabled by default

PostgreSQL Monitoring Tools

No shortage of tools



<https://wiki.postgresql.org/wiki/Monitoring>

Approaches



- ▶ Ad hoc
- ▶ Continuous monitoring frameworks
 - ▶ Cloud / SaaS / APM / Generic
 - ▶ DIY / Open Source

Postgres specific Open Source continuous monitoring frameworks

Postgres specific



- ▶ pghero
- ▶ PoWa (server side, quite advanced - pg_qualstats, pg_stat_kcache)
- ▶ PgObserver (client side + ad hoc)
- ▶ pgwatch2 (client side)
- ▶ ...

pgwatch2

Main principles - why another tool?

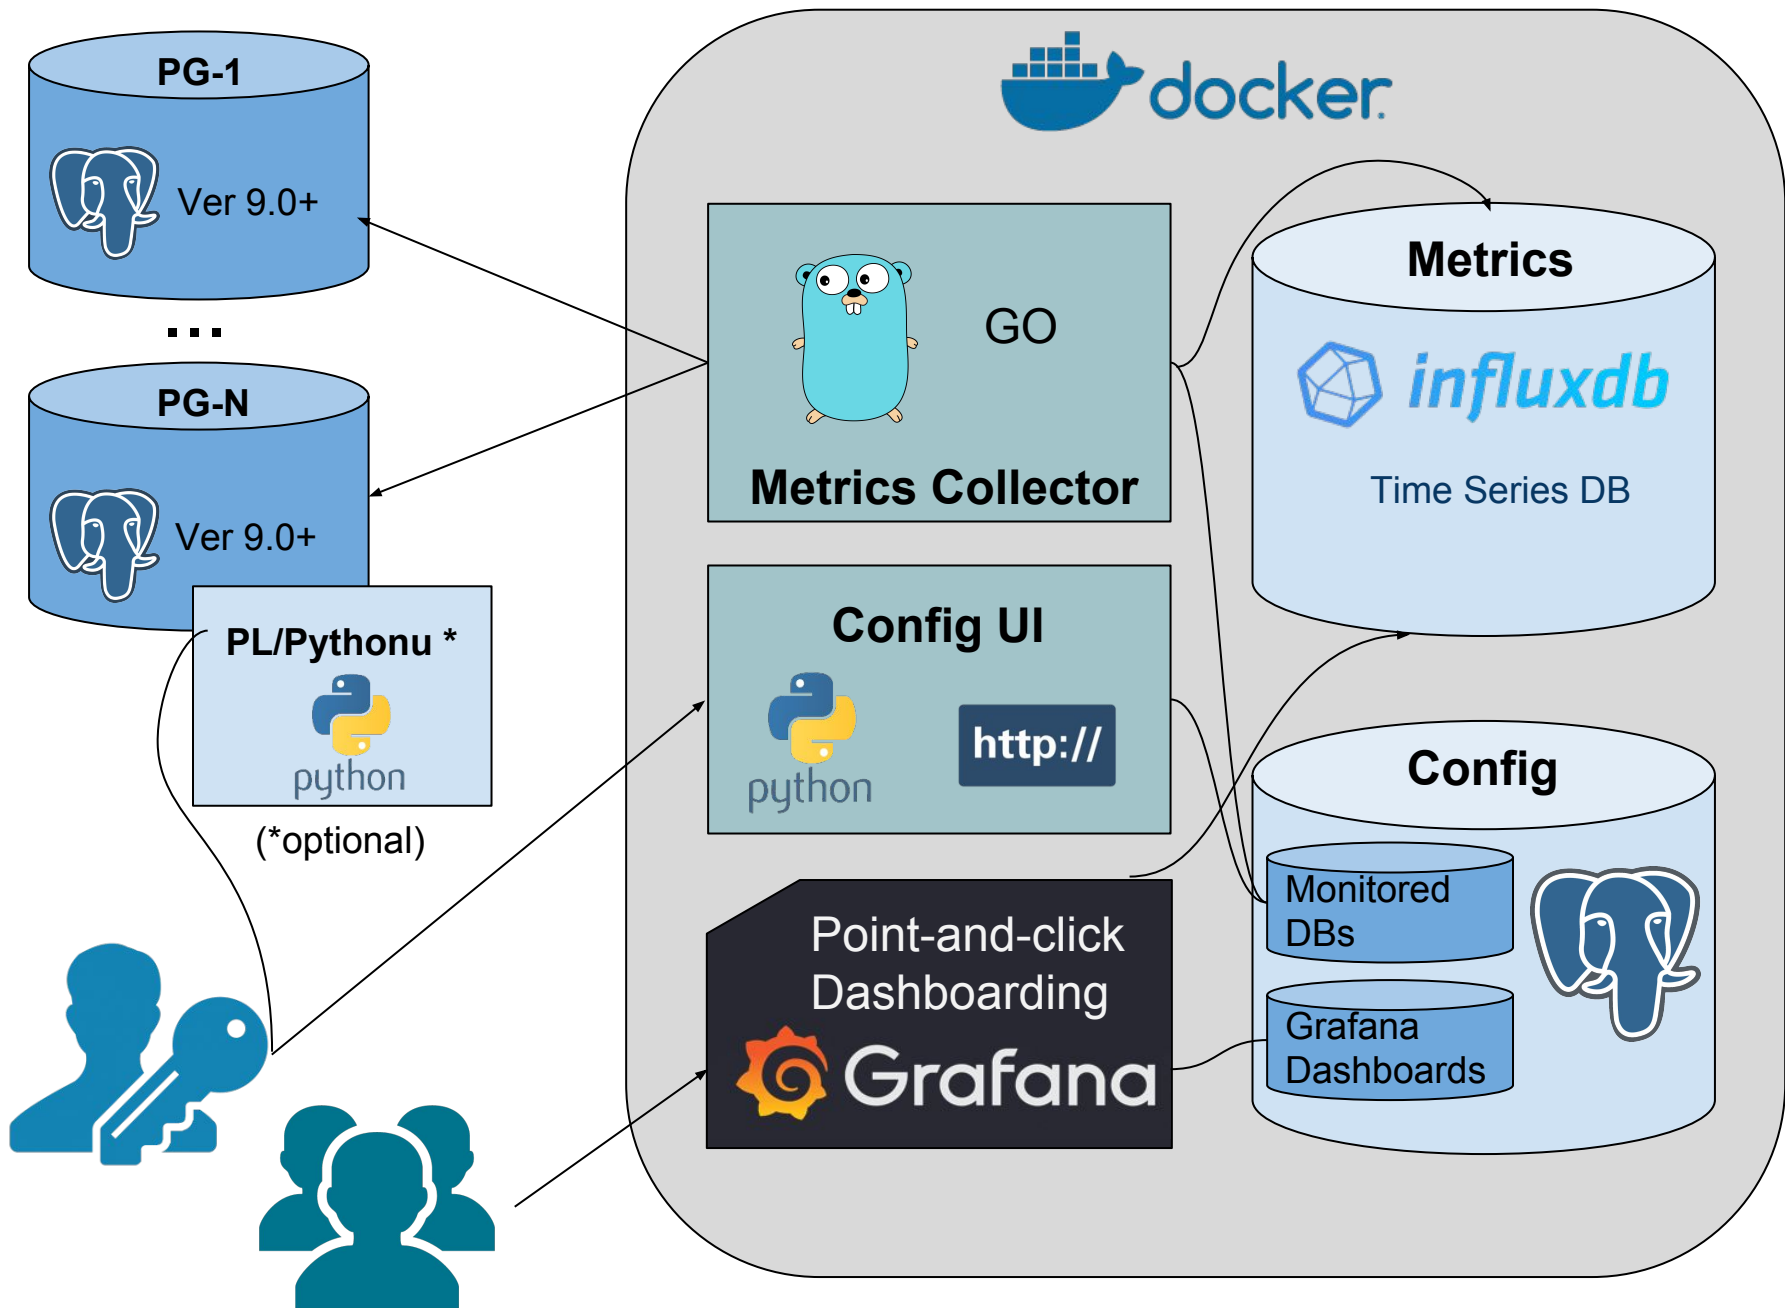


- ▶ 1-minute setup
 - ▶ Docker first
- ▶ User defined visuals / Dashboarding
- ▶ Non-invasive
 - ▶ No extensions for main functionality
- ▶ Easy extensibility
 - ▶ SQL defined metrics
- ▶ Do minimal work needed, use existing good software

Architecture components



- ▶ Web UI for administration
 - ▶ Python / Bootstrap
- ▶ Metrics gathering daemon
 - ▶ Go
- ▶ Config database
 - ▶ Postgres
- ▶ Metrics storage layer
 - ▶ InfluxDB (Graphite possible)
- ▶ Easy dashboarding with intuitive data discovery
 - ▶ Grafana



Features (1)



- ▶ “Ready to go”
 - ▶ Defaults cover almost all `pg_stat*` views
- ▶ Supports Postgres 9.0+ (older versions also possible)
- ▶ Security
- ▶ SSL between all components possible
- ▶ Admin / normal user separation
- ▶ Custom metrics via SQL, also for business domain!
- ▶ Reuse of existing Postgres, Grafana, InfluxDB installation possible
- ▶ Can be integrated with Kubernetes/OpenStack

Features (2)



- ▶ Possible to monitor all databases of a cluster automatically
- ▶ Change detection
 - ▶ Added/changed/deleted table/index/sproc/config events
- ▶ PgBouncer statistics support
- ▶ AWS CloudWatch support
- ▶ Alerting easily possible with Grafana
 - ▶ Kapacitor ("K" from InfluxData's TICK stack)
- ▶ Extensible data presentation - Grafana has plugins!











Getting started















1. `docker run -d -p 3000:3000 -p 8080:8080 \`
`--name pw2 cybertec/pgwatch2`
2. Wait some seconds and open browser at localhost:8080
3. Insert your DB connection strings and wait some minutes
4. Start viewing/editing dashboarding!



dbname test



dbname test





dbname test top 3

Top queries by total runtime

Total runtime	Query ID	Query
36.9 s	498864409	UPDATE pgbench_accounts SET abalance = abalance + ? WHERE aid = ?;
6.3 s	3513538518	UPDATE pgbench_tellers SET tbalance = tbalance + ? WHERE tid = ?;
4.7 s	2037868801	select (extract(? from now()) * ?)::int8 as epoch_ns, load_1min, load_5min, load_15min from public.get_load_average(); -- needs the plpythonu proc from "metric_fetching_helpers" folder

Top queries by avg. runtime

Avg. runtime	Query ID	Query
60.60 ms	811056313	select (extract(? from now()) * ?)::int8 as epoch_ns, approx_free_percent, approx_free_space as approx_free_space_b from public.get_table_bloat_approx() where approx_free_space > ?
13.20 ms	2037868801	select (extract(? from now()) * ?)::int8 as epoch_ns, load_1min, load_5min, load_15min from public.get_load_average(); -- needs the plpythonu proc from "metric_fetching_helpers" folder
12.52 ms	2320730284	Opens 'Single query details' dashboard for that queryid

Top queries by calls

Calls	Query ID	Query
37.4 K	2351822184	SELECT abalance FROM pgbench_accounts WHERE aid = ?;
37.4 K	2775512216	INSERT INTO pgbench_history (tid, bid, aid, delta, mtime) VALUES (?, ?, ?, ?, CURRENT_TIMESTAMP);
37.4 K	498864409	UPDATE pgbench_accounts SET abalance = abalance + ? WHERE aid = ?;

Top queries by IO

IO time	Query ID	Query
2.6 s	811056313	select (extract(? from now()) * ?)::int8 as epoch_ns, approx_free_percent, approx_free_space as approx_free_space_b from public.get_table_bloat_approx() where approx_free_space > ?
2.1 s	498864409	UPDATE pgbench_accounts SET abalance = abalance + ? WHERE aid = ?;
		select (extract(? from now()) * ?)::int8 as epoch_ns, quote_ident(table_schema) '?' quote_ident(table_name) as tag_table, md5((array_agg((c.*)::text order by ordinal_position)::text) from (SELECT current_database())::information_schema.sql_identifier AS table_catalog, nc.nspname::information_schema.sql_identifier AS table_schema, c.relname::information_schema.sql_identifier AS table_name, a.attname::information_schema.sql_identifier AS column_name, a.attnum::information_schema.cardinal_number AS ordinal_position, pg_get_expr(ad.adbin, ad.adrelid)::information_schema.character_data AS column_default, CASE WHEN a.attnotnull OR t.typtype = ?::"char" AND t.typtype = ?::"char" THEN ?::text ELSE ?::text END::information_schema.yes_or_no AS is_nullable, CASE WHEN t.typtype = ?::"char" THEN CASE WHEN bt.typlen < ?::oid AND bt.typlen = ?::integer THEN ?::text WHEN nbt.nspname = ?::name THEN format('type(%t %tbase%t %t::integer) ELSE ?::text END ELSE CASE WHEN t.typlen < ?::oid AND t.typlen =






Table changes

Time	event	table
2017-06-05 12:57:19	alter	pgwatch2.pgbench_accounts

Index changes

Time	event	index
2017-06-05 13:02:19	create	pgwatch2.pgbench_accounts_aid_idx

Sproc changes

Time	event	sproc
2017-06-05 12:57:19	alter	pgwatch2.test_func1







Config changes

No data to show




dbname pb

TPS 15	QPS 15	Avg. query runtime 1.0 ms	Pool wait time 0.17 μs	Incoming traffic 938.6 Bps	Outgoing traffic 3.2 kBps
-------------------------	-------------------------	--	--	---	--

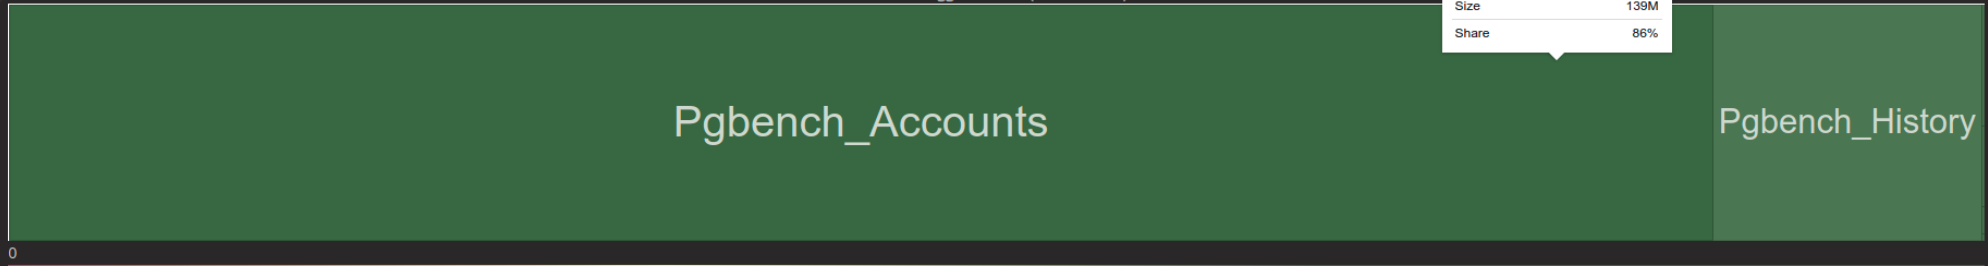


dbname kala_postgres


Biggest by total relation size



Biggest tables (w/o indexes)

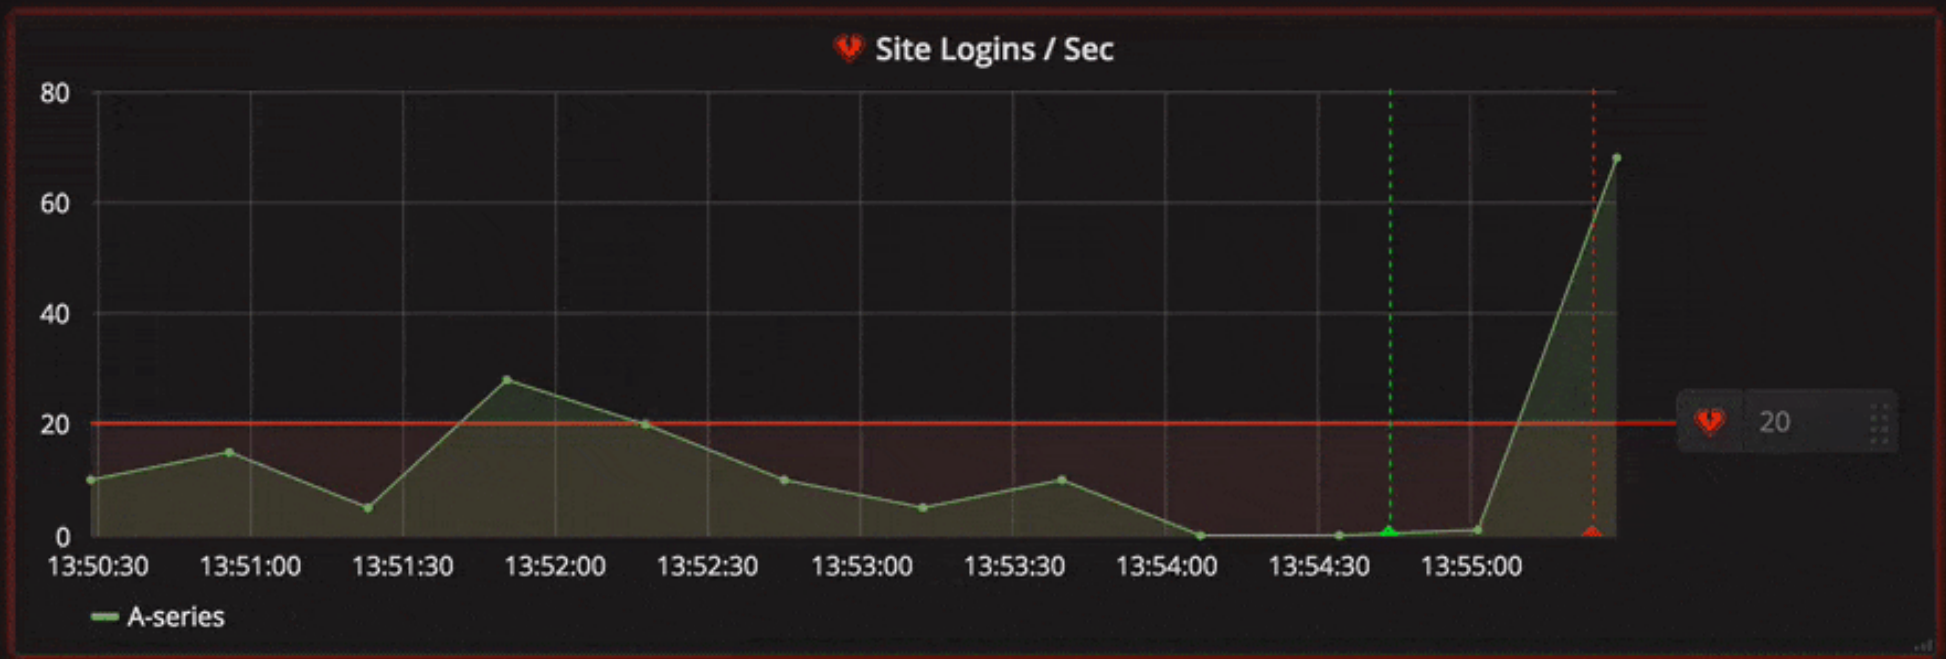


Biggest indexes



Alerting

- ▶ Easy setup, point and click
- ▶ Most important alerting services covered
 - ▶ Email
 - ▶ Slack
 - ▶ PagerDuty
 - ▶ Web hooks
 - ▶ Kafka
 - ▶ ...
- ▶ Graph panel only currently



Graph

General

Metrics

Axes

Legend

Display

Alert

Time range



Alert Config

Notifications (1)

State history

Delete

Alert Config

Name

Site Logins Too Low

Evaluate every

10s

Conditions

WHEN

avg ()

OF

query (A, 5m, now)

IS BELOW

20



+

If no data points or all values are null

SET STATE TO

No Data



Anomaly detection

Part of the InfluxData's TICK stack

- ▶ Harder to get going but very powerful!
- ▶ Features
 - ▶ Extensive math/string processing support
 - ▶ Statistical data mangling
 - ▶ UDF-s
 - ▶ Alert topics - pub/sub
 - ▶ Stream caching (e.g. last 10min moving average)
 - ▶ Stream redirection - store transformed data back into InfluxDB

Kapacitor sample - simple



```
stream
  |from()
    .measurement('cpu')
  |alert()
    .crit(lambda: "usage_idle" < 70)
    .log('/tmp/alerts.log')
    .email()
```

Kapacitor sample - complex (simplified)



```
|from()
    .measurement('cpu')
|groupBy('service', 'datacenter')
|window()
    .period(1m)
|percentile('load_1min', 95.0)
|eval(lambda: sigma("percentile"))
    .as('sigma')
|alert()
    .id('{{ .Name }}/{{ index .Tags "service" }}/{{ index .T
    .message('{{ .ID }} is {{ .Level }} cpu-95th:{{ index .F
    .crit(lambda: "sigma" > 3.0)
```

pgwatch2 - What's next?

Improvement areas



- ▶ More system level metrics
- ▶ Log parsing
- ▶ Per metric timeouts and activity times
- ▶ Fully automatic Docker updates
- ▶ ???

User input very much appreciated!

Contact us



github.com/cybertec-postgresql/pgwatch2

Austria / Switzerland / Estonia

Web: www.cybertec-postgresql.com

Github: github.com/cybertec-postgresql

Twitter: @PostgresSupport