Keep an eye on your PostgreSQL clusters

Open PostgreSQL Monitoring & PostgreSQL Workload Analyzer

Julien Rouhaud

Dalibo - www.dalibo.org

pgconf.ru 2015 - February

Monitoring ?

DALIBO
L'expertise PostgreSQL

- Service availability
  - Service, host, network...
  - Alerting
- Service performance
  - Graphing
  - Trending
  - Analysis

At what cost ?

DALIBO
L'expertise PostgreSQL

- Open source projects

What already exists

DALIBO
L'expertise PostgreSQL

- Different kind of tools
  - Command line tools / datasources
  - Generic solutions with probes
  - Dedicated solutions

Command line

And data sources

DALIBO
L'expertise PostgreSQL

- Nice features
- But more useful for emergency situations
- Or need some external tool for best usage

Command line

Some examples

DALIBO
L'expertise PostgreSQL

- command line
  - pg_view
  - pg_activity
  - pgstats
- Data source
  - all pg_stat* catalog
  - pg_stat_statements, pg_stat_plans
  - pg_proctab

Probe

DALIBO
L'expertise PostgreSQL

- Gather informations for another tool
- Can usually be used in standalone

Probe

Example

DALIBO
L'expertise PostgreSQL

- check_postgres
- pg_monz

Dedicated solution

DALIBO
L'expertise PostgreSQL

- Complete for its purpose
- More pertinent
- But not numerous

OPM and PoWA ?

DALIBO
L'expertise PostgreSQL

- OPM
  - aim to be a dedicated solution
  - relies on existing or new probes and components
- PoWA
  - Dedicated performance solution

Ecosystem

Generic solution based on probes

DALIBO
L'expertise PostgreSQL

Name	Native graphing	Alerting
Nagios	No	Yes
Zabbix	Yes	Yes
Munin	Yes	Yes
Cacti	Yes	No

General Overview

DALIBO
L'expertise PostgreSQL

- Pros
  - Robust and mature
  - Adaptable
  - Extendable
- Cons
  - UI not flexible
  - Data not available for querying
  - Except check_postgres, lack of really complete datasource

Nagios

Specific

DALIBO
L'expertise PostgreSQL

- Graphing (with third part tool like OPM V2)
- PostgreSQL compatibility with
  - check_postgres (Bucardo)
  - check_pg_activity (OPM)
- Hard to configure

Zabbix

Overview

DALIBO
L'expertise PostgreSQL

- PostgreSQL compatibility with
  - pg_monz (only PG 9.2+)
  - postbix (no news since 2013)

Munin

Overview

DALIBO
L'expertise PostgreSQL

- Native PostgreSQL compatibility, and with
  - pyMunin

Cacti

Overview

DALIBO
L'expertise PostgreSQL

- PostgreSQL compatibility with
  - Some workaround with check_postgres (MRTG format)

What about OPM

DALIBO
L'expertise PostgreSQL

- And for Open PostgreSQL Monitoring ?

OPM

Architecture

DALIBO
L'expertise PostgreSQL

- Nagios
  - Scheduler
  - Alerting
- Probes
  - PG specific : check_pgactivity
  - And monitoring-plugin
- Storage
  - PostgreSQL :)
  - 9.3 or more
- GUI
  - Dedicated GUI

OPM

Overview

DALIBO
L'expertise PostgreSQL

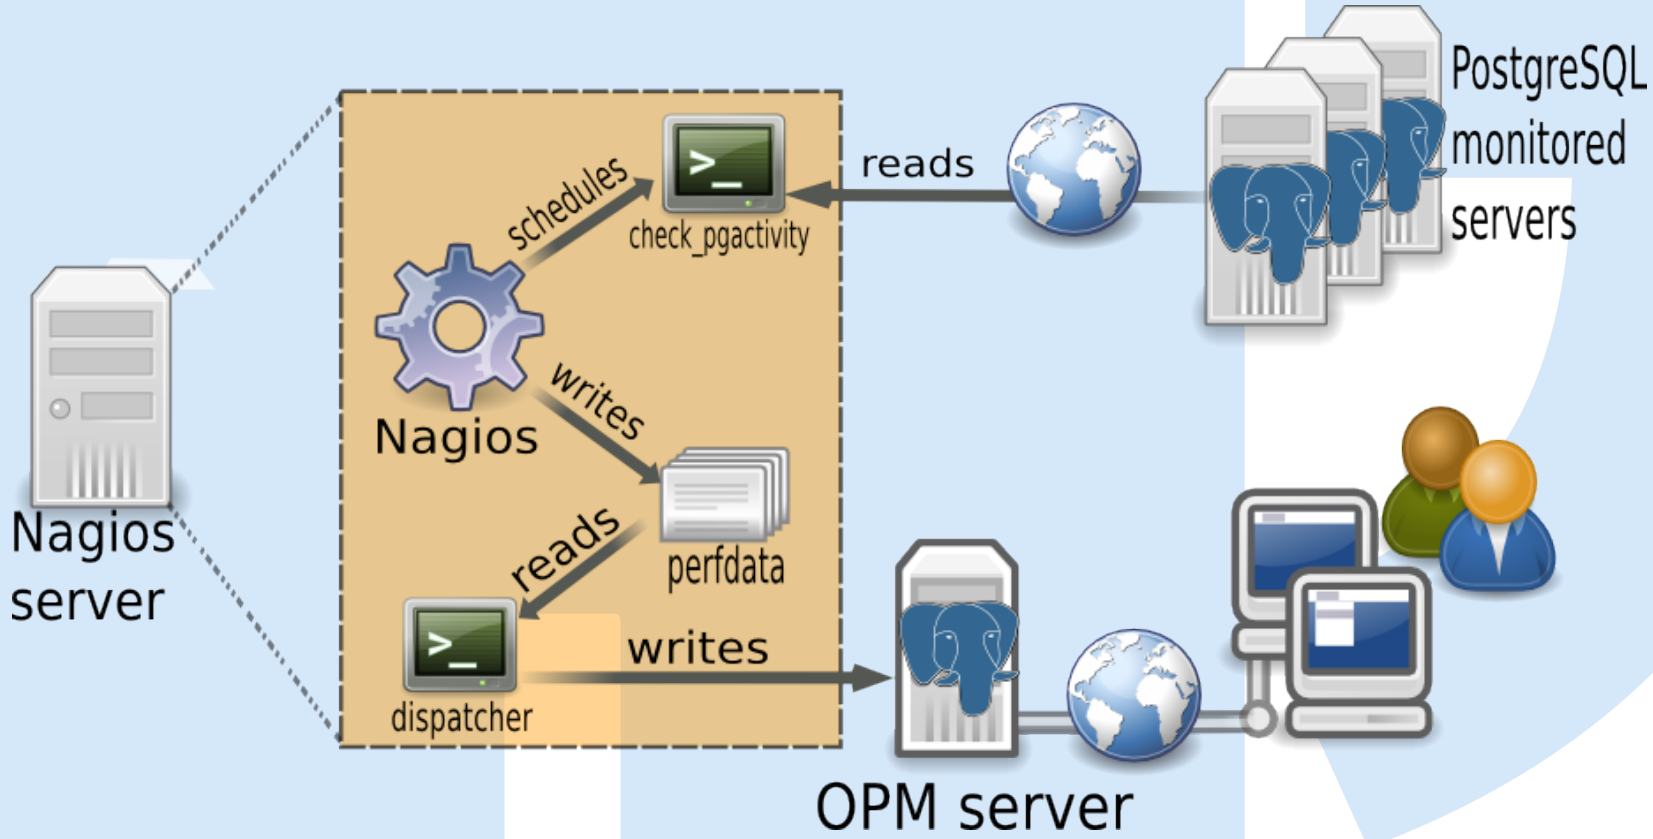

check_pgactivity

Design

DALIBO
L'expertise PostgreSQL

- Written from scratch, simpler code
- Provide lots of services
- Handle a small cache
- Better perfdata

check_pgactivity

New features

DALIBO
L'expertise PostgreSQL

- Handle multi-database connections transparently
- Can compute delta instead of raw values
- Handle units

check_pgactivity

New service examples

DALIBO
L'expertise PostgreSQL

- bgwriter statistics
- better bloat estimation
- Most important settings check

check_pgactivity

Enhanced service

DALIBO
L'expertise PostgreSQL

- better bloat estimation
- replication : pg_stat_replication or hot standby
- "Instant" hit ratio

User interface

DALIBO
L'expertise PostgreSQL

- Specific ACL
- More modern graphs
- Custom graph display

Ecosystem

Specific solution

DALIBO
L'expertise PostgreSQL

Name	Maintained	Aim	Alerting
pg_statsinfo	Yes	Generalist	Possible
pg_watch	Dead ?	Generalist	No
pgObserver	Yes	Performance	No
pgCluu	Yes	Generalist	No
PoWA	Yes	Performance	No

pg_statsinfo

Pros & cons

DALIBO
L'expertise PostgreSQL

- Pros
  - Useful metrics
- Cons
  - Reports on demand
  - Needs specific extension and reboot on each monitored server

pgObserver

Pros & cons

DALIBO
L'expertise PostgreSQL

- Pros
  - Useful metrics
  - And novel metrics (like stored proc)
  - Dynamic reports
- Cons
  - Focused on performance

pgCluu

Pros & cons

DALIBO
L'expertise PostgreSQL

- Pros
  - Useful metrics
  - General overview
  - Easy to use
- Cons
  - Can be storage greedy
  - Need to generate reports

PoWA 2.0

DALIBO
L'expertise PostgreSQL

- A complete solution
- Handle a lot of datasource
- Requires PG 9.4 or more
  - requires queryid exposure, since pg_stat_statements 1.2

What is PoWA

DALIBO
L'expertise PostgreSQL

- A background worker
- Dedicated snapshot, aggregation and purge functions
- Dedicated UI

PoWA

Extensions handled

DALIBO
L'expertise PostgreSQL

- Existing extensions
  - pg_stat_statements
  - pg_proctab (WIP)
- New one
  - pg_qualstats
  - pg_stat_kcache
  - pg_track_settings (WIP)

pg_qualstats

Presentation

DALIBO
L'expertise PostgreSQL

- Gather real-time statistics on where clauses per query
  - SeqScan / indexScan
  - Operator
  - Number of execution
  - Filter ratio
  - ...

pg_qualstats

What it allows

DALIBO
L'expertise PostgreSQL

- Find missing indexes
  - Which could be partial
  - Or on several columns if some other queries could benefit from it
  - ...

pg_qualstats

What it allows

DALIBO
L'expertise PostgreSQL

- EXPLAIN normalized queries with real values
  - Most frequents
  - Most/Least filtering
  - ...

pg_stat_kcache

Presentation

DALIBO
L'expertise PostgreSQL

- Gather real-time per query system statistics

pg_stat_kcache

What it allows

DALIBO
L'expertise PostgreSQL

- Compute a real hit ratio (shared_buffers, system cache and disk)
- Show CPU consumption
  - global, per database and/or per query and/or per user
  - over time

Demo

DALIBO
L'expertise PostgreSQL

For the future

DALIBO
L'expertise PostgreSQL

- Nagios optional from OPM
- Add ability to handle PoWA in OPM
- Better use of metrics
  - Trending / statistic analysis
  - Correlate informations
- Better index suggestion, global database analysis
- And much more...

Questions ?

DALIBO
L'expertise PostgreSQL

- julien.rouhaud@dalibo.com
- contact@opm.io - powa@dalibo.com
- github.com/OPMDG - github.com/dalibo/powa
- Thank you !