

# Watch your elephants

## PostgreSQL Performance Analysis using collectd

Sebastian „tokkee“ Harl <sh@teamix.net>

teamix GmbH / collectd core team



PGConf.EU 2012  
October 24, 2012  
Prague

**WARNING:** I'm not a database (performance) expert!

This talk is an overview about a tool that may be used for the purpose of performance analysis.

## Solid IT-Infrastructure

Location: Nuremberg, Munich, Frankfurt

<http://teamix.net/>

Open-Source

N-IX

Riverbed

Monitoring

NetApp

VMWare

Network

Juniper

Trainings



## collectd Overview

Overview


Main Features

Feature Overview


PostgreSQL Processes

Querying statistics from PostgreSQL


- **collectd** collects performance data of systems
- some (simple) examples:
  - CPU utilization
  - memory utilization
  - network traffic
- **collectd** collects and stores the performance data
- stored data is usually used to generate graphs
- → performance analysis, capacity planing
- not to be confused with *monitoring!*
- Homepage: <http://collectd.org/>




- daemon
- free software (mostly GPL)
- portable (Linux, \*BSD, Solaris, ...)
- scalable (OpenWrt, ..., Cluster / Cloud)
- sophisticated network support
- efficient (default resolution: 10 seconds)
- flexible architecture
- modular (more than 100 plugins in Version 5.1)




- daemon
- free software (mostly GPL)
- portable (Linux, \*BSD, Solaris, ...)
- scalable (OpenWrt, ..., Cluster / Cloud)
- sophisticated network support
- **efficient** (default resolution: 10 seconds)
- flexible architecture
- modular (more than 100 plugins in Version 5.1)







- daemon
- free software (mostly GPL)
- portable (Linux, \*BSD, Solaris, ...)
- scalable (OpenWrt, ..., Cluster / Cloud)
- sophisticated network support
- efficient (default resolution: 10 seconds)
- flexible architecture
- **modular** (more than 100 plugins in version 5.1)



## available plugins (version 5.1)

amqp	apache	apcups	apple_sensors	ascent
battery	bind	contrack	contextswitch	cpu
cpufreq	csv	curl	curl_json	curl_xml
dbi	df	disk	dns	email
entropy	ethstat	exec	filecount	fscache
GenericJMX.java	gmond	hddtemp	interface	ipmi
iptables	ipvs	irq	java	libvirt
load	logfile	lpar	madwifi	match_empty_counter
match_hashed	match_regex	match_timediff	match_value	mbmon
md	memcached	memcached	memory	modbus
Monitorus.pm	multimeter	mysql	netapp	netlink
network	nfs	nginx	notify_desktop	notify_email
ntpd	numa	nut	olsrd	onewire
openvpn	OpenVZ.pm	oracle	perl	pinba
ping	postgresql	powerdns	processes	protocols
python	redis	routers	rrdcached	rrdtool
sensors	serial	snmp	swap	syslog
table	tail	tape	target_notification	target_replace
target_scale	target_set	target_v5upgrade	tcpconns	teamspeak2
ted	thermal	threshold	tokyotyrant	unixsock
uptime	users	uuid	varnish	vmem
vserver	wireless	write_graphite	write_http	write_mongodb
write_redis	xmms	zfs_arc		

- daemon collects data locally  $\Rightarrow$  runs on every client system (exceptions: SNMP, databases, etc.)
- one or more central servers
- clients push their data to the central servers
- first steps: `install`; `select plugins`; `start daemon`; `enjoy ;-)`

## C4


collection 4

- All instances
- All graphs
- Host "[redacted]"

## Graph "CPU utilization"

Instance "[redacted]"/0"

Instance: "[redacted] / cpu - 0 / cpu - all"

 Search

Hour ▾

JSON (gRaphaël)

RRDtool

Go


collection 4.0.0

- provide collected data through JSON
- different frontends possible
- efficiently handles large amounts of data
- flexible configuration of graphs

*graphite*

[Command Line Interface](#)  
Logged in as [admins](#) | [logout](#) | [edit profile](#)  
[Documentation](#)

[production](#)  
[pre-production](#)



collectd Overview

## Feature Overview

CPU, memory, network I/O

Networking Support

RRDtool Support

Generic Plugins (Overview)

PostgreSQL Processes

Querying statistics from PostgreSQL

- specialized read plugins
  - CPU, memory, network interfaces, ...
- IO plugins
  - network plugin
  - RRDtool, RRDcached
  - Graphite
  - MongoDB, Redis
  - AMQP
- generic plugins
  - SNMP
  - tail
  - PostgreSQL
- filter-chains

## configuration synopsis


```
LoadPlugin "cpu"  
LoadPlugin "memory"  
LoadPlugin "interface"
```





## configuration synopsis

```
LoadPlugin "cpu"  
LoadPlugin "memory"  
LoadPlugin "interface"
```

```
<Plugin interface>  
  Interface lo  
  Interface sit0  
  IgnoreSelected true  
</Plugin>
```







## modes of operation

- send data ("*client*")
- receive data ("*server*")
- forward data ("*proxy*")
- Unicast ("*point-to-point*")
- Multicast ("*point-to-group*")
- IPv4 and IPv6

## rule them all

Modes may be mixed arbitrarily.

## synopsis: client

```
LoadPlugin "network"
```

```
<Plugin "network">
```

```
  Server "collectd0.example.com"
```

```
  Server "collectd1.example.com"
```

```
  Server "ff18::efc0:4a42"
```

```
</Plugin>
```

## synopsis: server

```
LoadPlugin "network"  
  
<Plugin "network">  
  Listen "collectd0.example.com"  
  Listen "ff18::efc0:4a42"  
</Plugin>
```

### synopsis: proxy

```
LoadPlugin "network"
```

```
<Plugin "network">
```

```
  Listen "collectgw.extern.example.com"
```

```
  Server "collectd1.intern.example.com"
```

```
  Forward true
```

```
</Plugin>
```



- writes data to RRD files **efficiently** → caching
- functionality now also available in RRDtool as stand-alone RRD Caching Daemon (RRDCacheD)

### synopsis

```
LoadPlugin "rrdtool"
```

```
<Plugin "rrdtool">  
  DataDir "/var/lib/collectd/rrd"  
</Plugin>
```

### configuration synopsis

```
<Plugin "rrdtool">  
  DataDir "/var/lib/collectd/rrd"  
  
  CacheTimeout 3600 # 1 hour  
  CacheFlush 86400 # 1 day  
  
  WritesPerSecond 30  
</Plugin>
```

- FLUSH command allows for graphing of current values

- idea: generic approaches rather than specialized solutions
- → user configuration determines behavior
- ⇒ new equipment does not require a new version of **collectd**
- examples: SNMP, tail, curl, DBI, PostgreSQL

collectd Overview

Feature Overview

**PostgreSQL Processes**

Querying statistics from PostgreSQL

```
% ps ax | grep postgres
20177 ?  S    0:05 /usr/lib/postgresql/9.1/bin/postgres
        -D /var/lib/postgresql/9.1/main
        -c config_file=/etc/postgresql/9.1/main/postgresql.conf
20183 ?  Ss   0:09 postgres: writer process
20184 ?  Ss   0:05 postgres: wal writer process
20185 ?  Ss   0:04 postgres: autovacuum launcher process
20186 ?  Ss   0:13 postgres: stats collector process
20312 ?  Ss   2:04 postgres: collectd mail 127.0.0.1(33027) idle
```




- processes handling client connections:  
postgres: *user database host activity*

- The `processes` plugin collects various information about (groups of) processes
  - RSS and VM size
  - user and system time
  - number of page-faults
  - I/O estimates
- processes are selected either by process name or by regex of it's command line (`/proc/cmdline` on Linux)

## collectd.conf

```
<Plugin "processes">
  ProcessMatch pg_writer "postgres:.writer.process"
  ProcessMatch pg_wal_writer "postgres:.wal.writer.process"
  ProcessMatch pg_autovacuum "postgres:.*autovacuum"
  ProcessMatch pg_stats_collector \
    "postgres:.stats.collector.process"
  # database connections by 'user'
  ProcessMatch pg_user_mail "postgres:.user"
  # database connections to database 'mail'
  ProcessMatch pg_db_mail "postgres:.[A-Za-z0-9]+.mail"
</Plugin>
```

(versions before 5.0.1 did not support whitespace in regexes)





collectd Overview

Feature Overview

PostgreSQL Processes

Querying statistics from PostgreSQL

The PostgreSQL statistics collector

The collectd `postgresql` plugin

## postgresql.conf

```
# currently running command
track_activities = on
# access to tables and indices
track_counts = on
# user-defined functions
track_functions = none # none, pl, all
```

Storing statistics on, for example, flash storage:

```
stats_temp_directory = '/mnt/flash/pg_stat_tmp'
```

- server processes submit statistics before going idle
- collector generates report each `PGSTAT_STAT_INTERVAL` milliseconds
- during each transaction, a snapshot of the report will be used → see `pg_stat_clear_snapshot()`

- some predefined views
  - `pg_stat_bgwriter`
  - `pg_stat_database`
  - `pg_stat_all_indexes`
  - `pg_statio_all_tables`
  - many more (see table 27.1 in the documentation)
- also there are various functions to query single values

```
sh=# select datname, numbackends,  
sh-# xact_commit, xact_rollback  
sh-# from pg_stat_database;
```

datname	numbackends	xact_commit	xact_rollback
template1	0	0	0
template0	0	0	0
postgres	0	611	0
sh	1	661	12

(4 rows)

- generic plugin which collects arbitrary (numeric) values from a database
- by default, queries various values from the statistics collector
- configuration has two parts
  - SQL query and specifications how to interpret the values
  - database connection plus queries assigned to it

- each data-set uses a unique identifier
  - hostname
  - plugin name
  - plugin instance (optional)
  - type
  - type instance (optional)
- `hostname/plugin[-instance]/type[-instance]`
- the type specifies how **collectd** is supposed to handle the data-set (cf. RRDtool's data-source types)
- any type needs to be pre-defined (`types.db(5)`)
  
- example: `server1.ex.com/cpu-0/cpu-idle`

## collectd.conf

```
<Plugin postgresql>
  <Query disk_usage>
    Statement "SELECT pg_database_size($1) AS size;"
    Param database

    <Result>
      Type pg_db_size
      ValuesFrom "size"
    </Result>
  </Query>
</Plugin>
```




## collectd.conf

```
<Plugin postgresql>
  <Database mail>
    Host "db.ex.com"
    User "user"
    Password "secret"
    Query disk_usage
    Query disk_io
  </Database>
</Plugin>
```

Thanks for your attention!

## Any questions?



## Contact:

Sebastian „tokkee“ Harl  
teamix GmbH, Nürnberg  
<sh@teamix.net>

<collectd@verplant.org> — irc.freenode.net/#collectd — <http://identi.ca/collectd>

<http://github.com/collectd/collectd> — <http://github.com/tokkee/collectd>