


PostgreSQL, PostgreSQL monitoring and monitoring postgresql.org



Stefan Kaltenbrunner
stefan@kaltenbrunner.cc
<http://www.kaltenbrunner.cc/blog>

Nagios conference
2008

Nuremberg, Germany
September 2008

About myself

- Stefan Kaltenbrunner
- Senior Systems Architect at conova communications
- Major contributor to PostgreSQL project
- Infrastructure and Webteam
- Patches and Portability

Agenda

- PostgreSQL
- PostgreSQL Internals
- Monitoring PostgreSQL
 - What, Why and How
 - Tools
- Monitoring [postgresql.org](https://www.postgresql.org)
- Q & A

Agenda


- PostgreSQL
- PostgreSQL Internals
- Monitoring PostgreSQL
 - What, Why and How
 - Tools
- Monitoring [postgresql.org](https://www.postgresql.org)
- Q & A

PostgreSQL

- BSD licenced ORDBMS
- Opensource Project
- 20 year old, 13 years of OSS
- Fast, reliable and strong community
- Adapted and used as the base for commercial and proprietary solutions

PostgreSQL

- ~900.000 lines of code
- hundreds of subprojects
- ~200 active developers



<http://ohloh.net/projects/postgres/>

Agenda

- PostgreSQL
- **PostgreSQL Internals**
- Monitoring PostgreSQL
 - What, Why and How
 - Tools
- Monitoring [postgresql.org](https://www.postgresql.org)
- Q & A

PostgreSQL Internals

- Process based architecture
- One process per connection
- Number of dedicated processes
- Shared memory
- OS level caching
- Highly portable (Unix, Windows, ...)

PostgreSQL Internals

```
PID TTY STAT TIME COMMAND  
3321 ? S 0:01 /usr/local/pgsql/bin/postgres -D /var/data/pg83  
3324 ? Ss 0:10 \_ postgres: writer process  
3325 ? Ds 0:06 \_ postgres: wal writer process  
3326 ? Ss 0:02 \_ postgres: autovacuum launcher process  
3327 ? Ss 0:01 \_ postgres: stats collector process  
27998 ? Ds 0:14 \_ postgres: postgres test [local] SELECT  
28583 ? Ss 0:01 \_ postgres: postgres test [local] INSERT
```

PostgreSQL Internals

```
PID TTY STAT TIME COMMAND  
3321 ? S 0:01 /usr/local/pgsql/bin/postgres -D /var/data/pg83  
3324 ? Ss 0:10 \_ postgres: writer process  
3325 ? Ds 0:06 \_ postgres: wal writer process  
3326 ? Ss 0:02 \_ postgres: autovacuum launcher process  
3327 ? Ss 0:01 \_ postgres: stats collector process  
27998 ? Ds 0:14 \_ postgres: postgres test [local] SELECT  
28583 ? Ss 0:01 \_ postgres: postgres test [local] INSERT
```

PostgreSQL Internals

```
PID TTY STAT TIME COMMAND
3321 ? S 0:01 /usr/local/pgsql/bin/postgres -D /var/data/pg83
3324 ? Ss 0:10 \_ postgres: writer process
3325 ? Ds 0:06 \_ postgres: wal writer process
3326 ? Ss 0:02 \_ postgres: autovacuum launcher process
3327 ? Ss 0:01 \_ postgres: stats collector process
27998 ? Ds 0:14 \_ postgres: postgres test [local] SELECT
28583 ? Ss 0:01 \_ postgres: postgres test [local] INSERT
```

PostgreSQL Internals

```
PID TTY STAT TIME COMMAND
3321 ? S 0:01 /usr/local/pgsql/bin/postgres -D /var/data/pg83
3324 ? Ss 0:10 \_ postgres: writer process
3325 ? Ds 0:06 \_ postgres: wal writer process
3326 ? Ss 0:02 \_ postgres: autovacuum launcher process
3327 ? Ss 0:01 \_ postgres: stats collector process
27998 ? Ds 0:14 \_ postgres: postgres test [local] SELECT
28583 ? Ss 0:01 \_ postgres: postgres test [local] INSERT
```

PostgreSQL Internals

<i>PID</i>	<i>TTY</i>	<i>STAT</i>	<i>TIME</i>	<i>COMMAND</i>
3321	?	S	0:01	/usr/local/pgsql/bin/postgres -D /var/data/pg83
3324	?	Ss	0:10	_ postgres: writer process
3325	?	Ds	0:06	_ postgres: wal writer process
3326	?	Ss	0:02	_ postgres: autovacuum launcher process
3327	?	Ss	0:01	_ postgres: stats collector process
27998	?	Ds	0:14	_ postgres: postgres test [local] SELECT
28583	?	Ss	0:01	_ postgres: postgres test [local] INSERT

PostgreSQL Internals

<i>PID</i>	<i>TTY</i>	<i>STAT</i>	<i>TIME</i>	<i>COMMAND</i>
<i>3321</i>	<i>?</i>	<i>S</i>	<i>0:01</i>	<i>/usr/local/pgsql/bin/postgres -D /var/data/pg83</i>
<i>3324</i>	<i>?</i>	<i>Ss</i>	<i>0:10</i>	<i>_ postgres: writer process</i>
<i>3325</i>	<i>?</i>	<i>Ds</i>	<i>0:06</i>	<i>_ postgres: wal writer process</i>
<i>3326</i>	<i>?</i>	<i>Ss</i>	<i>0:02</i>	<i>_ postgres: autovacuum launcher process</i>
<i>3327</i>	<i>?</i>	<i>Ss</i>	<i>0:01</i>	<i>_ postgres: stats collector process</i>
<i>27998</i>	<i>?</i>	<i>Ds</i>	<i>0:14</i>	<i>_ postgres: postgres test [local] SELECT</i>
<i>28583</i>	<i>?</i>	<i>Ss</i>	<i>0:01</i>	<i>_ postgres: postgres test [local] INSERT</i>

Agenda

- PostgreSQL
- PostgreSQL Internals
- **Monitoring PostgreSQL**
 - What, Why and How
 - Tools
- Monitoring [postgresql.org](https://www.postgresql.org)
- Q & A

Monitoring PostgreSQL

- Piece of software – needs CPU, RAM and Disk
 - Current state
 - Historic state
 - Trends
- Application specific information
 - Slow queries
 - Errors
 - Internal state

Agenda

- PostgreSQL
- PostgreSQL Internals
- Monitoring PostgreSQL
 - **What, Why and How**
 - Tools
- Monitoring [postgresql.org](https://www.postgresql.org)
- Q & A

PostgreSQL Monitoring

- Direct and indirect IO, CPU and Memory usage
- Maintenance tasks
- Background activity
- Average and worst case latency
- Application specific uses

PostgreSQL Monitoring

- Locks
- Active Queries
- Query Distribution
- Cache efficiency
- Disk Usage

Agenda

- PostgreSQL
- PostgreSQL Internals
- Monitoring PostgreSQL
 - What, Why and How
 - **Tools**
- Monitoring [postgresql.org](https://www.postgresql.org)
- Q & A

Monitoring PostgreSQL

- Statistics collector
- System catalogs
- Logfiles (role, database and session)
- Dtrace
- Operating System tools

Monitoring PostgreSQL

- `check_postgres.pl`
- `pgFouine`
- `Munin`
- `pgsnmpd`
- `cat/grep/awk` and `sed`

Tools - check_postgres.pl

- Swiss army knife of status monitoring
- Tablesizes, Transaction status, replication, maintenance operations, locks, number of connections and query runtime
- MRTG mode
- Must have for any nagios based PostgreSQL monitoring

Tools - check_postgres.pl

```
mastermind@mastermind:~$ check_postgres.pl -w "10 MB" -db test --action database_size
```

```
POSTGRES_DATABASE_SIZE WARNING: DB "test" test: 43719908 (42 MB) postgres: 4324580 (4223 kB) template0: 4243460 (4144 kB) template1: 4243460 (4144 kB) | time=0.04 test=43719908 postgres=4324580 template0=4243460 template1=4243460
```

```
mastermind@mastermind:~$ check_postgres.pl -w "16" -db test --action wal_files
```

```
POSTGRES_WAL_FILES OK: 8 | time=0.04
```

```
mastermind@mastermind:~$ check_postgres.pl -w "1 minute" -db test --action query_runtime --queryname "pg_class"
```

```
POSTGRES_QUERY_RUNTIME OK: DB "test" query runtime: 0.00089 seconds | time=0.04 qtime=0.00089
```

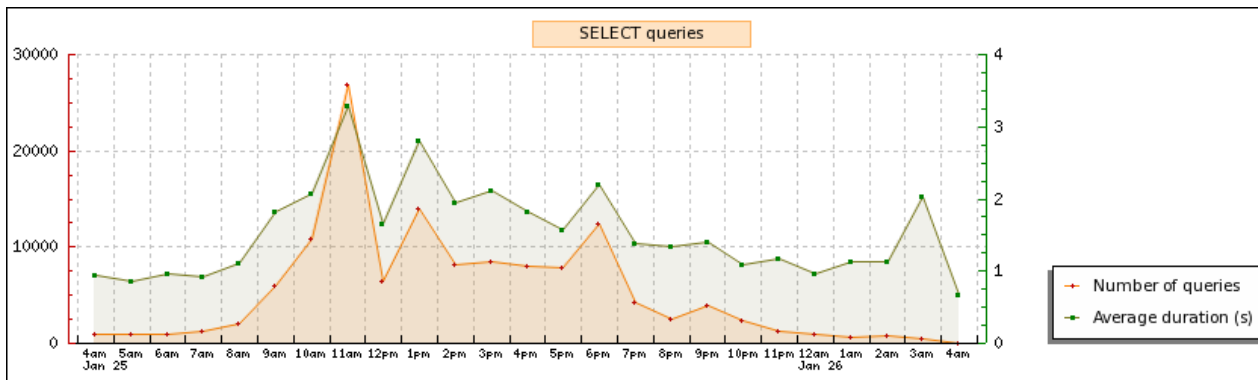
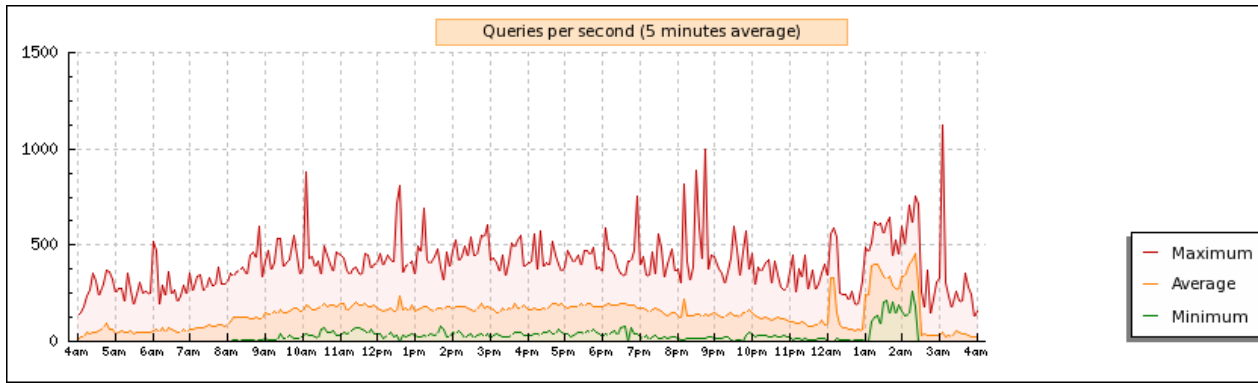
```
mastermind@mastermind:~$ check_postgres.pl -w "1 minute" -db test --action txn_time
```

```
POSTGRES_TXN_TIME WARNING: longest txn: 528s | time=0.04 maxtime:528
```


Tools - pgFouine

- PHP based logfile analyzer
- Creates nicely formatted reports
- replay/benchmark integration
- Requires query logging

Tools - pgFouine



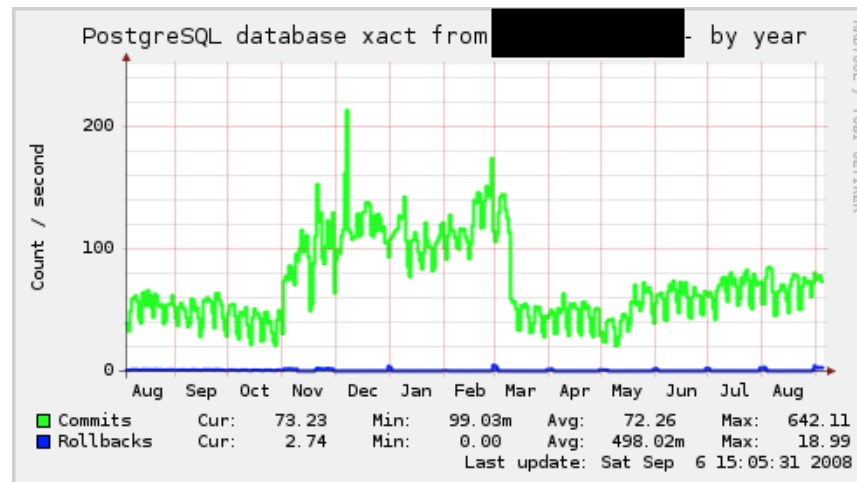
Tools - pgFouine

Queries that took up the most time (N) ^

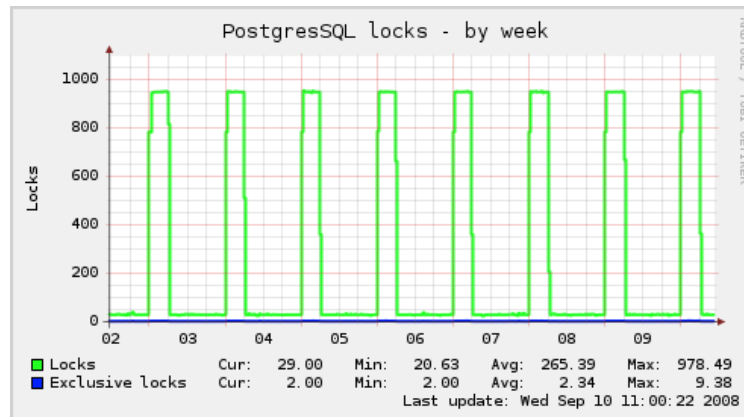
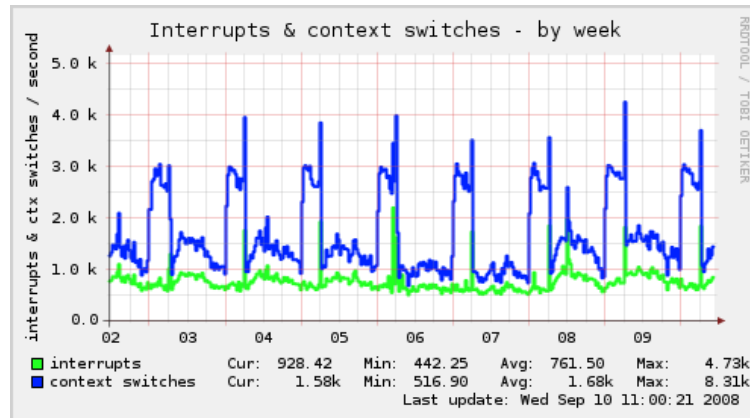
Rank	Total duration	Times executed	Av. duration (s)	Query
1	1933h26m41s	23,387	297.62	<p>UPDATE accounts SET filler=<i>lower</i>(') WHERE aid < 0;</p> <p>Hide examples</p> <p>4,777.68s UPDATE accounts SET filler=<i>lower</i>('teSt filIEr) WHERE aid < 1000;</p> <p>3,949.61s UPDATE accounts SET filler=<i>lower</i>('teSt filIEr) WHERE aid < 1000;</p> <p>3,763.44s UPDATE accounts SET filler=<i>lower</i>('teSt filIEr) WHERE aid < 1000;</p> <p>3,696.27s UPDATE accounts SET filler=<i>lower</i>('teSt filIEr) WHERE aid < 1000;</p> <p>3,687.32s UPDATE accounts SET filler=<i>lower</i>('teSt filIEr) WHERE aid < 1000;</p> <p>3,672.68s UPDATE accounts SET filler=<i>lower</i>('teSt filIEr) WHERE aid < 1000;</p> <p>3,539.89s UPDATE accounts SET filler=<i>lower</i>('teSt filIEr) WHERE aid < 1000;</p> <p>3,418.57s UPDATE accounts SET filler=<i>lower</i>('teSt filIEr) WHERE aid < 1000;</p> <p>3,371.81s UPDATE accounts SET filler=<i>lower</i>('teSt filIEr) WHERE aid < 1000;</p> <p>3,207.66s UPDATE accounts SET filler=<i>lower</i>('teSt filIEr) WHERE aid < 1000;</p>
2	17h14m20s	23,387	2.65	<p>UPDATE branches SET filler=<i>upper</i>(');</p> <p>Show examples</p>
3	17m13s	23,387	0.04	<p>SELECT history.* FROM accounts, history WHERE accounts.aid=0 AND accounts.aid=history.aid;</p> <p>Show examples</p>
4	15m4s	23,387	0.04	<p>SELECT accounts.* FROM accounts,history WHERE history.aid=0 AND accounts.aid=history.aid;</p> <p>Show examples</p>
5	13m56s	23,387	0.04	<p>SELECT history.* FROM accounts, branches, tellers, history WHERE accounts.aid=0 AND branches.bid=accounts.aid AND tellers.bid=accounts.aid AND history.aid=accounts.aid AND history.bid=branches.bid AND history.tid=tellers.tid;</p> <p>Show examples</p>
6	5m32s	46,774	0.01	<p>SELECT * FROM accounts WHERE aid=0;</p> <p>Show examples</p>
7	1m38s	23,387	0.00	<p>SELECT * FROM branches WHERE bid=0;</p> <p>Show examples</p>
8	1m6s	70,161	0.00	<p>SELECT * FROM tellers WHERE tid=0;</p> <p>Show examples</p>
9	0.1s	31	0.00	<p>SET search_path = public;</p>

Tools - Munin

- Simple monitoring solution
- Plugin based with agent on the host
- Postgresql plugins available
- Nagios integration for alerting



Tools - Munin



Tools - pgsnmpd

- SNMP based monitoring
- RFC 1697 (RDBMS-MIB)
- No PostgreSQL specific MIB yet
 - Mostly useful in SNMP only environments

Agenda

- PostgreSQL
- PostgreSQL Internals
- Monitoring PostgreSQL
 - What, Why and How
 - Tools
- **Monitoring postgresql.org**
- Q & A

Monitoring postgresql.org

- ~50 monitored hosts (about half fully community managed)
- 340 monitored services
- FreeBSD, Debian, Ubuntu, Slackware, CentOS, Redhat Enterprise Linux and VMware
- Austria, Canada, France, Israel, Panama, Sweden, USA

Monitoring postgresql.org

- Nagios based
- Smokeping latency tracking
- Munin graphing
- Trac
 - Configuration change management
 - Internal ticketing
 - Docs
- Coordination through IM and Mailinglists

Monitoring postgresql.org

- FreeBSD jails
- 2-8 jails/host
- Nagios + NRPE
- Munin/Smokeping
- Custom checks
 - Project specific things
 - Hardware checks
 - Security related checks

Thank you!



Questions?