Integrated Monitoring for PostgreSQL

Tim Retout

credativ limited

2nd April 2008

Outline

1. Introduction
2. Monitoring PostgreSQL
   - Statistics tables
   - Portability
3. Solutions
   - SNMP
   - Nagios
   - Munin
4. Conclusion

Who am I?

- Tim Retout <tim.retout@credativ.co.uk>
- Involved with free software
  - Debian
  - GNU
  - GNOME
- With credativ since September 2007
  - Systems administration
  - Development

Who am I?

- Tim Retout <tim.retout@credativ.co.uk>
- Involved with free software
  - Debian
  - GNU
  - GNOME
- With credativ since September 2007
  - Systems administration
  - Development

Who am I?

- Tim Retout <tim.retout@credativ.co.uk>
- Involved with free software
  - Debian
  - GNU
  - GNOME
- With credativ since September 2007
  - Systems administration
  - Development

What am I talking about?

How to integrate monitoring solutions with PostgreSQL

- Portability across different PostgreSQL versions
- Sometimes monitoring needs fine-tuning by hand

Outline

- Monitoring PostgreSQL
  - "What to monitor"
- Solutions and examples
  - "How to monitor"

What am I talking about?

How to integrate monitoring solutions with PostgreSQL

- Portability across different PostgreSQL versions
- Sometimes monitoring needs fine-tuning by hand

Outline

- Monitoring PostgreSQL
  - "What to monitor"
- Solutions and examples
  - "How to monitor"

What am I talking about?

How to integrate monitoring solutions with PostgreSQL

- Portability across different PostgreSQL versions
- Sometimes monitoring needs fine-tuning by hand

Outline

- Monitoring PostgreSQL
  - "What to monitor"
- Solutions and examples
  - "How to monitor"

Outline

1 Introduction

2 Monitoring PostgreSQL
- Statistics tables
- Portability

3 Solutions
- SNMP
- Nagios
- Munin

4 Conclusion

Types of monitoring

History of data

- Current status
  - Notifications
  - "Is the database accepting connections?"
- Status over time
  - Plotting graphs
  - "What was the highest number of concurrent connections this week? When was it?"

Types of monitoring

Source of data

- PostgreSQL statistics
- External sources
  - Disk usage via “du”
  - Connections via “netstat”

External sources of data will tend to be less portable. Some PostgreSQL statistics will not appear in earlier releases.

Connections

Check whether database is accepting connections (easy).
Count number of current connections:

```
SELECT COUNT(*) FROM pg_stat_activity;
```

Compare with value of ‘max_connections’.

Transactions

Number of commits and rollbacks for each database:

```
SELECT datname, xact_commit, xact_rollback
FROM pg_stat_database;
```

Queries

Summaries of number of inserts/updates/deletes:

```
SELECT SUM(n_tup_ins), SUM(n_tup_upd), SUM(n_tup_del)
FROM pg_stat_all_tables;
```

Summaries of query plans:

```
SELECT SUM(seq_scan), SUM(seq_tup_read), SUM(idx_scan),
    SUM(idx_tup_fetch) FROM pg_stat_all_tables;
```

Locks

Summaries of locks held:

```
SELECT mode, COUNT(mode) FROM pg_locks
    GROUP BY mode ORDER BY mode;
```

This may need post-processing to work out which are exclusive locks.

Disk I/O

Summary of disk I/O in terms of blocks read:

```
SELECT SUM(heap_blks_read) FROM pg_statio_user_tables;
SELECT SUM(idx_blks_read) FROM pg_statio_user_tables;
SELECT SUM(toast_blks_read) FROM pg_statio_user_tables;
SELECT SUM(tidx_blks_read) FROM pg_statio_user_tables;
```

Swap ‘read’ for ‘hit’ to get the number of blocks read from memory.

Caution!

This does not take the operating system’s disk cache into account!

Disk usage

Obtaining the disk usage of a database is more difficult to do in a portable way across PostgreSQL releases.
From version 8.1, there are SQL functions to do this:

`pg_database_size(name)`
`pg_tablespace_size(name)`

Previous releases had a tool 'oid2name' in contrib, but not all installations have this.
If all else fails, analysis of VACUUM information is possible.

Disk usage

Fall back gracefully:

```
SELECT CASE EXISTS(
    SELECT COUNT(*) FROM pg_proc
    WHERE proname='pg_database_size'
    )
WHEN true THEN ...
ELSE ...
END
```

See discussion at:
http://postgresql.org/docs/8.1/static/diskusage.html

Outline

1 Introduction

2 Monitoring PostgreSQL
- Statistics tables
- Portability

3 Solutions
- SNMP
- Nagios
- Munin

4 Conclusion

SNMP

- "Simple Network Management Protocol"
- Not simple
- Different implementations of agents possible

SNMP and PostgreSQL

- Need a MIB (Management Information Base)
- pgsnmpd - http://pgsnmpd.projects.postgresql.org/
- 1.0 released 20 August 2007
- Implements RFC 1697 (generic RDBMS MIB)
- Few PostgreSQL-specific statistics yet - plan is to create PGSQL-MIB
- Little visible activity currently

Nagios

- Popular monitoring software
- Sends notifications
- (Can draw graphs of statistics using a plugin)
- Very configurable, so can be complicated when learning
- Can speak SNMP

Nagios plugins

Included in Nagios

- "check_pgsql" command
- Tests whether PostgreSQL database is accepting connections

User-contributed

These do a lot more:

- http://bucardo.org/nagios/
- http://pgfoundry.org/projects/nagiosplugins/

Nagios plugins

Included in Nagios

- "check_pgsql" command
- Tests whether PostgreSQL database is accepting connections

User-contributed

These do a lot more:

- http://bucardo.org/nagios/
- http://pgfoundry.org/projects/nagiosplugins/

Munin

- Popular monitoring software
- Draws graphs
- Quite easy to set up
- Can send notifications, or can be integrated with Nagios

Munin and PostgreSQL

- PostgreSQL plugins in SVN
- Will not be in a stable release of Munin until they stabilize
- Should perhaps all be ported to POSIX shell, because not all installations have DBD::Pg.

Munin graphs

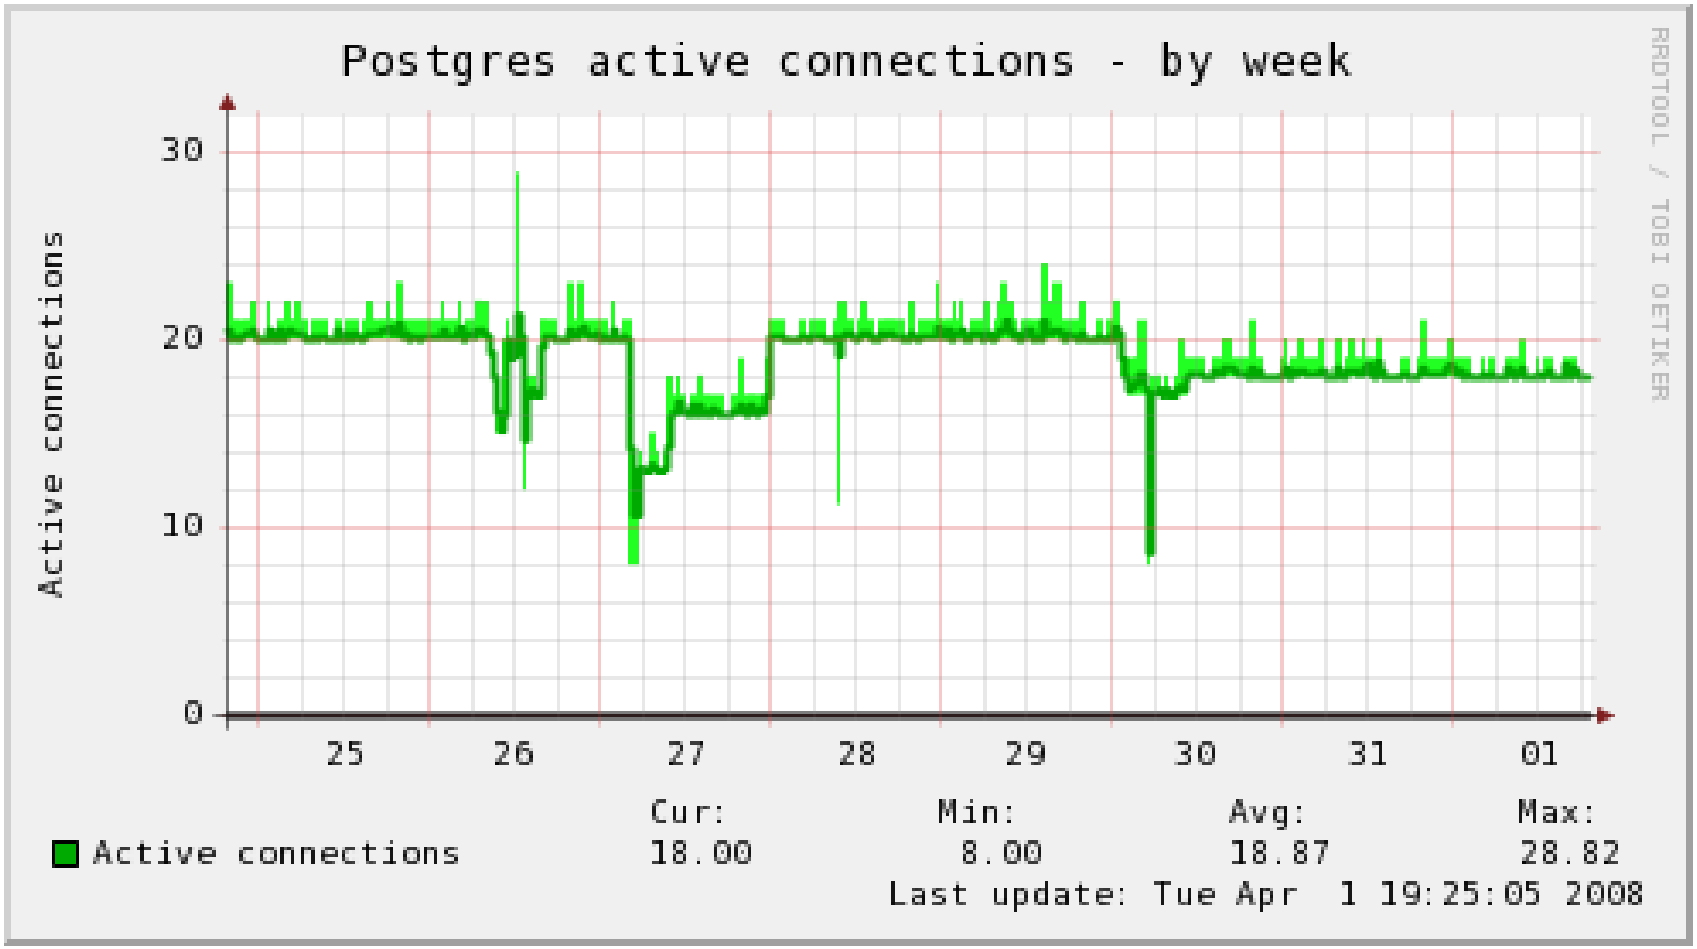

Munin graphs

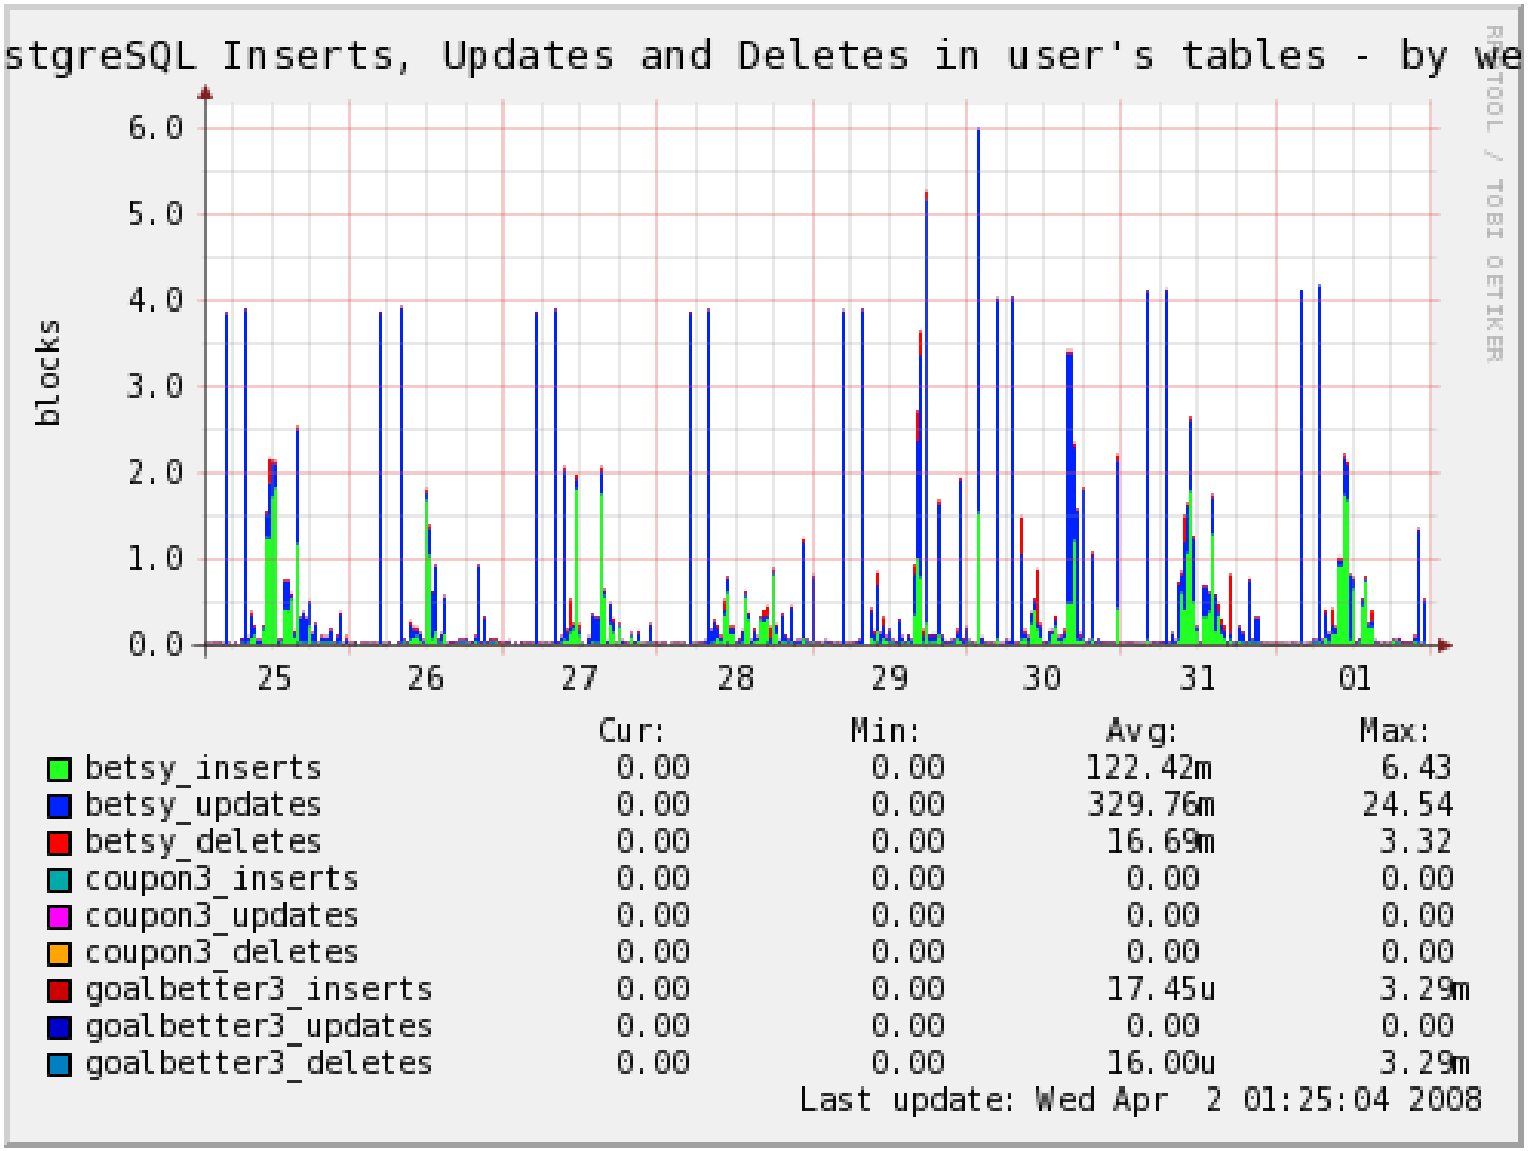

Munin graphs

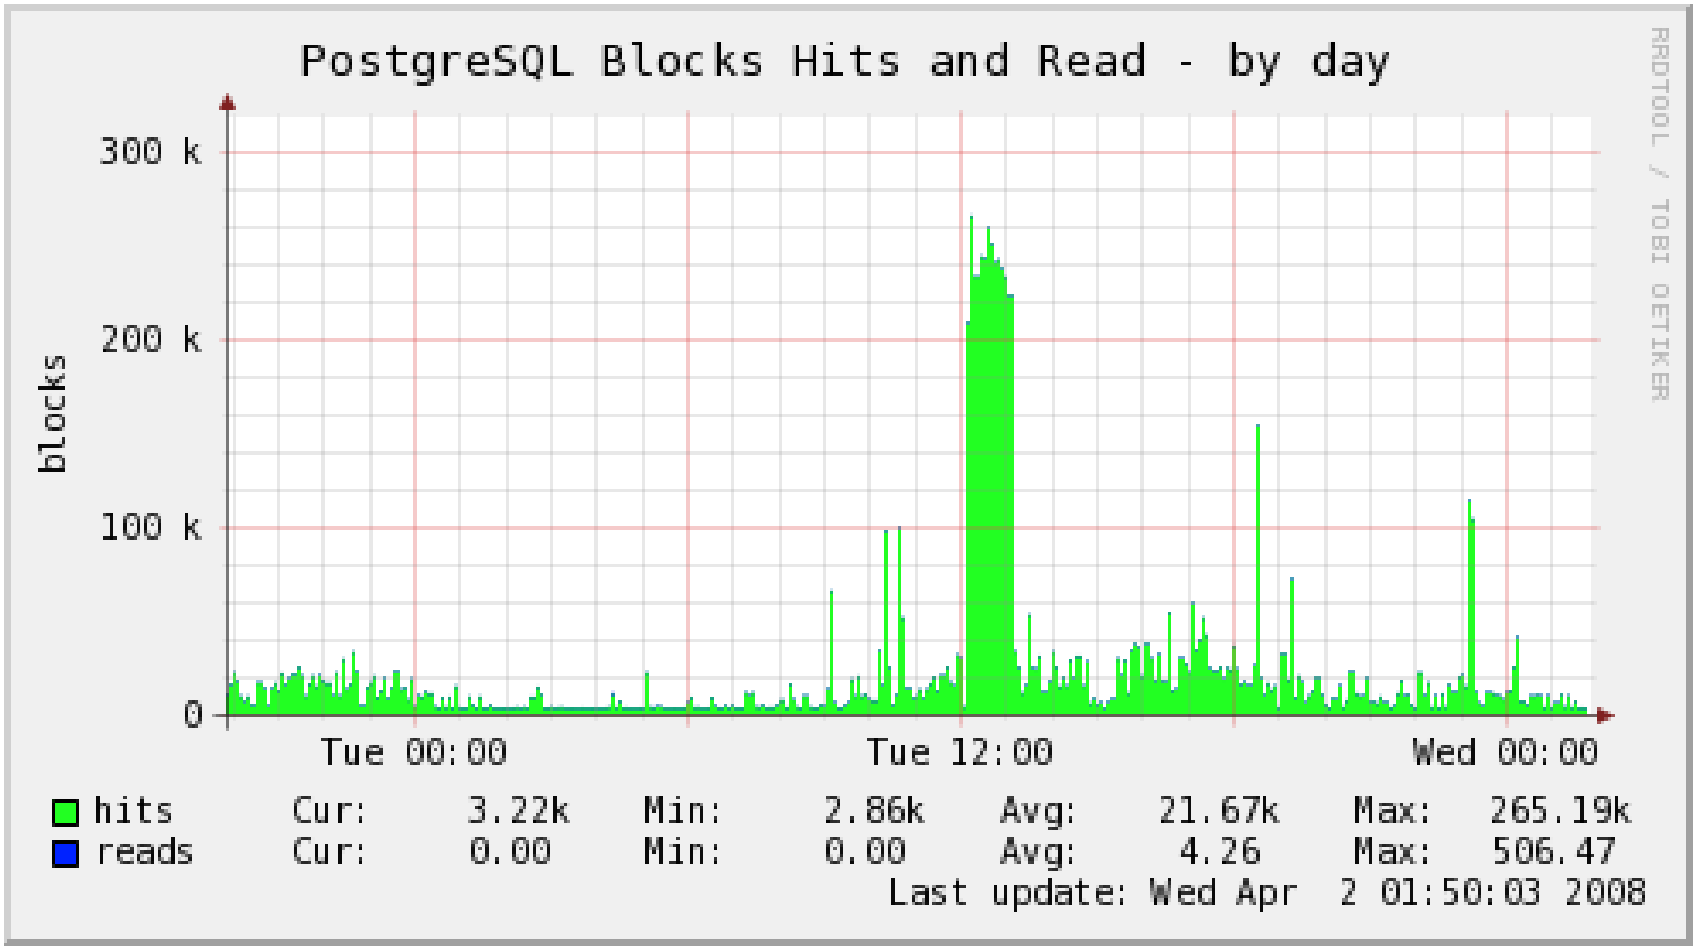

Outline

1. Introduction
2. Monitoring PostgreSQL
   - Statistics tables
   - Portability
3. Solutions
   - SNMP
   - Nagios
   - Munin
4. Conclusion

Conclusion

- Integrating monitoring with PostgreSQL can be a powerful tool for database administrators.
- One day, it could work out of the box.
- But knowing how to customize these solutions may be useful.

Thanks! Any questions?

Feel free to email me to discuss anything covered here:
<tim.retout@credativ.co.uk>

Conclusion

- Integrating monitoring with PostgreSQL can be a powerful tool for database administrators.
- One day, it could work out of the box.
- But knowing how to customize these solutions may be useful.

Thanks! Any questions?

Feel free to email me to discuss anything covered here:
<tim.retout@credativ.co.uk>