



# Finding the cursed query: slow log reports with pt-query-digest

Marcos Albe & Fernando Ipar  
MySQL Meetup Montevideo – Feb. 2012

# Agenda

---

- Who are we?
- What is pt-query-digest
- Setup
- Basic usage
- Some nice tricks
- Conclusions
- Questions?

# Who are we

---

- Fernando: senior consultant, founder of Percona-UY team and coffee connoisseur.
- Marcos: support engineer, ticket cruncher and beer connoisseur.

# What is pt-query-digest

- Part of the Percona Toolkit.
- Primary purpose is to give answers to “what queries should we improve or re-architecture”.
- Analyze query execution logs and generate a query report, filter, replay, or transform queries.

# What is pt-query-digest (contd.)

---

- Supports MySQL, PostgreSQL, memcached, and more.
- Consumes slow.log, tcpdump, binlogs and more.

# Setup

- Requirements:
  - Perl, DBI, DBD::mysql, Term::ReadKey
- Download tarball, RPM, or DEB packages:
  - `wget percona.com/get/percona-toolkit.rpm`
  - Not in Percona repositories (yet)

# Setup (contd.)

---

- OR Download the tool and set the right mode:
  - `wget percona.com/get/pt-query-digest`
  - `chmod +x pt-query-digest`

# Basic usage – collect data

- Enable slow log:
  - SET GLOBAL long\_query\_time=0; **REVERT WHEN DONE!!**
  - SET GLOBAL min\_examined\_row\_limit=0;
  - SET GLOBAL log\_output='FILE';
  - SET GLOBAL slow\_query\_log=1;
  - SET GLOBAL log\_queries\_not\_using\_indexes=1;



# Basic usage – collect data (cont.)

- Or capture tcpdump:
  - `tcpdump -i any -s 65535 -x -n -q -tttt \ 'port 3306 and tcp[1] & 7 == 2 and tcp[3] & 7 == 2'`

# Basic usage – report on data

---

- Get a report based on all the slow log
  - `pt-query-digest --report the-slow.log`
- Get a report limited to last 30 minutes of log
  - `pt-query-digest --report --since 1800s`

# Basic usage – report on data (cont.)

- Get a report limited between dates X and Y
  - pt-query-digest --report \ --since '2012-02-03 21:48:59' \ --until '2012-02-06 02:33:50'

# Some nice tricks

- Find all the SELECT's between date X and Y
  - Pt-query-digest \  
--filter '\$event->{fingerprint} =~ m/^select/i' \  
--since='2012-02-01 03:30:00' \  
--until='2012-02-10 03:30:00'
- Report on all queries from a given user
  - --filter '(\$event->{user} || "") =~ m/markus/'

# Some nice tricks

- Report queries doing full table scan or full join
  - `pt-query-digest --filter \`  
`'(($event->{Full_scan} || "") eq "Yes") || \`  
`((($event->{Full_join} || "") eq "Yes"))'`

---

# Sample reports

# Conclusions

- Free & easy to install and use (basically).
- Find the cursed queries...and fix them: Query tuning usually brings 80% of the possible gains in performance optimization!!
- Endless other possibilities, even custom attribute pairs to group/sort by (you can add these attribs as comments to your queries, and those will go to the slow log too ;-))

# Questions

---

Any questions? :-)



# Thank You to Our Sponsors

---

- GuruHub - INfraestructure
- Percona Inc. - Cofee break
- ...want to be here? Talk with Martin, Fernando or Marcos

# Proximo meetup!!!

- Los meetups en 2012 seran cada segundo Jueves del mes.
- El proximo es el 8 de Marzo
- Ya tenemos un request para una charla extendida (1 hora) que promete mucho
- Y muchas charlas chiquitas que esperemos que vayan llegando, asi que manden la suya!!



**marcos.albe@percona.com**  
**fernando.ipar@percona.com**

**We're Hiring! [www.percona.com/about-us/careers/](http://www.percona.com/about-us/careers/)**



PERCONA  
LIVE

[www.percona.com/live](http://www.percona.com/live)