MySQL server replication, optimization and monitoring

Igor Ajdišek
Lead System Operations Engineer
Central European Media Enterprises

Central European Media Enerprises

osvajalec

osvajalec.si

Brief introduction...

- Leading media company in CEE region w/ operations in 6 countries
- US funded; NASDAQ: CETV
- Owns 20+ TV stations
- Operates POPTV, Kanal A and TV Pika in Slovenia
- Combined audience reach = 50 Million people

Central European Media Enterprises

Well, what about Internet?

Central European Media Enerprises

24ur.com, poptv.si, **zadovoljna.si**, bibaleze.si, cekin.si, vizita.si, maxtv.si, moskisvet.com, dominvrt.si, **dnevnik.hr**, **blog.hr**, novatv.hr, zadovoljna.hr, gle.to, **markiza.sk**, **tvnoviny.sk**, **nova.cz**, tn.cz, **blog.cz**, **galerie.cz**, doma.cz, dobyvatel.cz, vybereme.cz, mtv.cz, **protv.ro**, **sport.ro**, acessatv.ro, **stirileprotv.ro**, mtv.ro, profm.ro, kombat.ro, conquiztador.ro, **btv.bg**, probg.bg, **ring.bg**, profm.bg, ...

Central European Media Enerprises

A few facts...

- Over 40 online products and services
- 2.5 Million users / month
- 450 Million page views / month
- Substantial YoY growth
- Dedicated infrastructure in 4 data centers (Ljubljana/Prague/Sofia/Bucharest)
- Built upon LAMP stack

Central European Media Enerprises

Average PV/month (24ur.com)

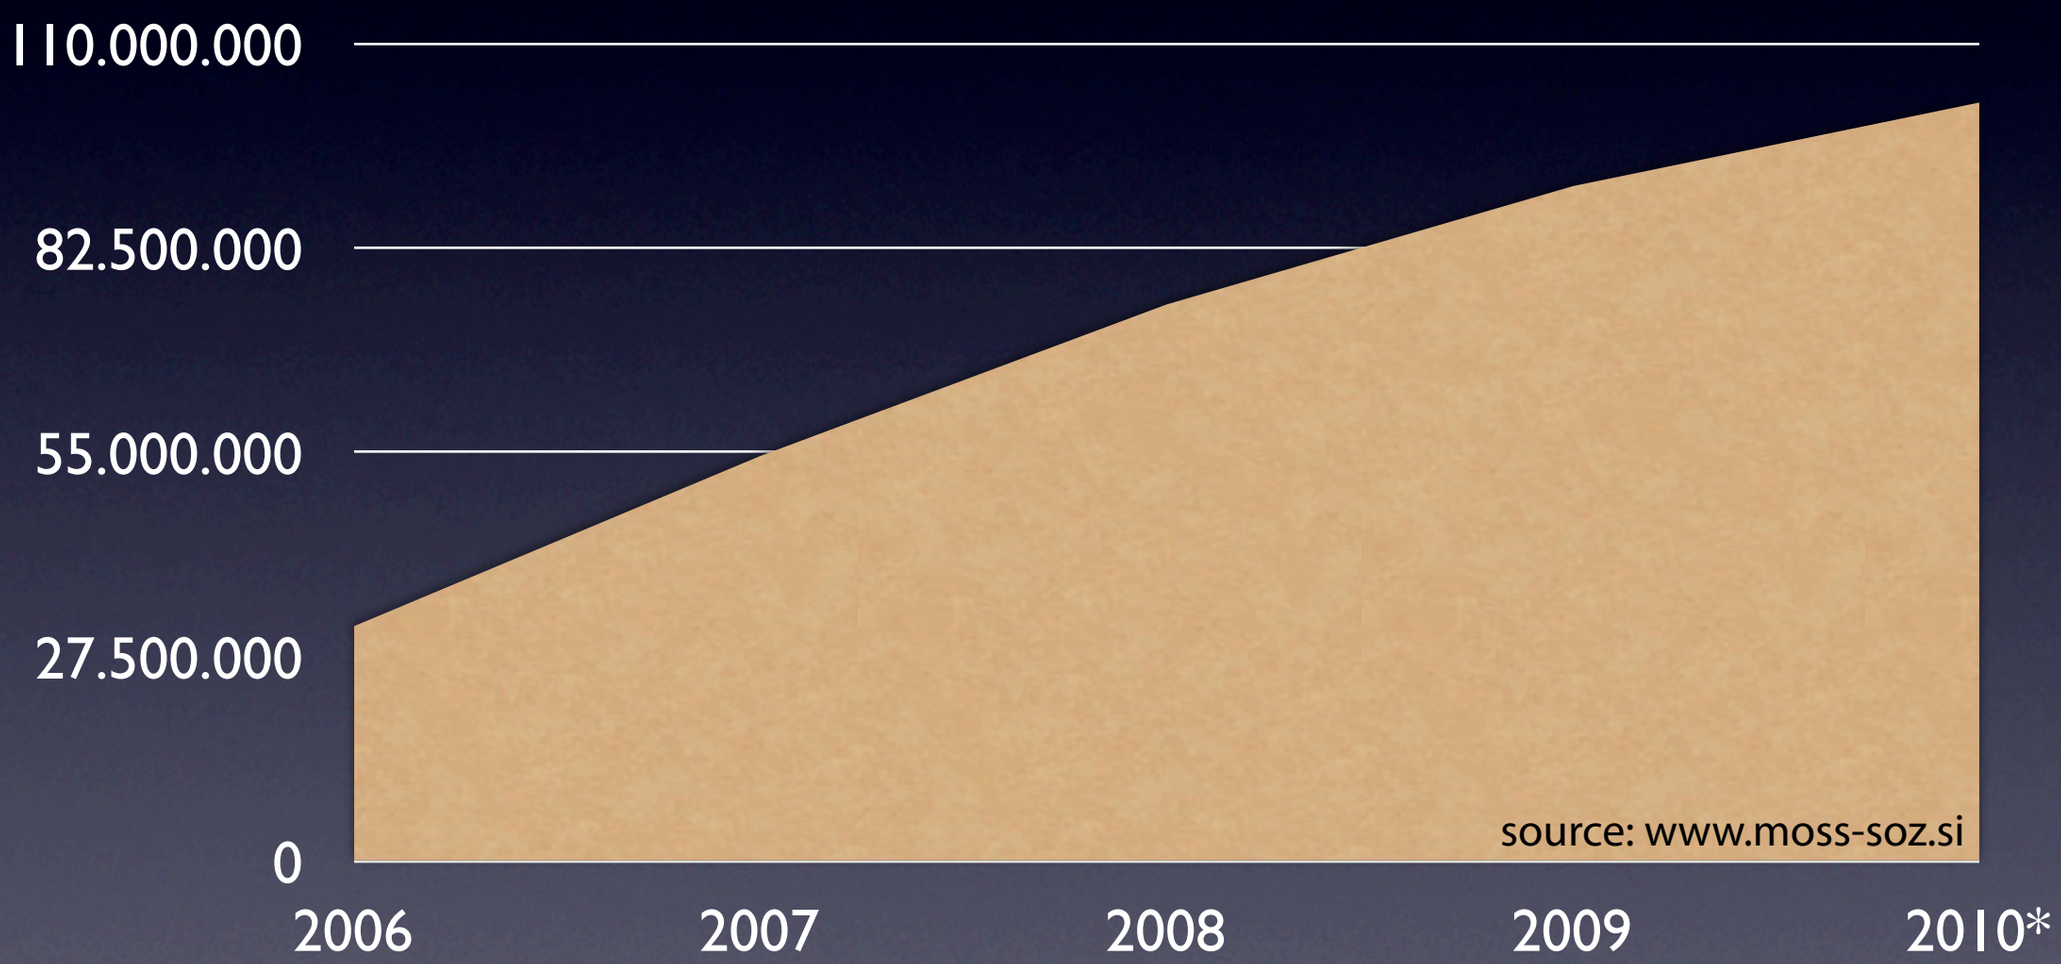

* data only from Q1/2010

Central European Media Enerprises

Adding 30 Million new PV/month every year

Central European Media Enerprises

How to keep up?

Central European Media Enerprises

Scale.

Central European Media Enerprises

Part 1: MySQL server replication

About MySQL replication...

- Foundation for building large, scalable web application on top of MySQL
- It’s simple. Has almost no overhead
- Great for web applications that have lots of reads and little writes

Central European Media Enerprises

How replication works?

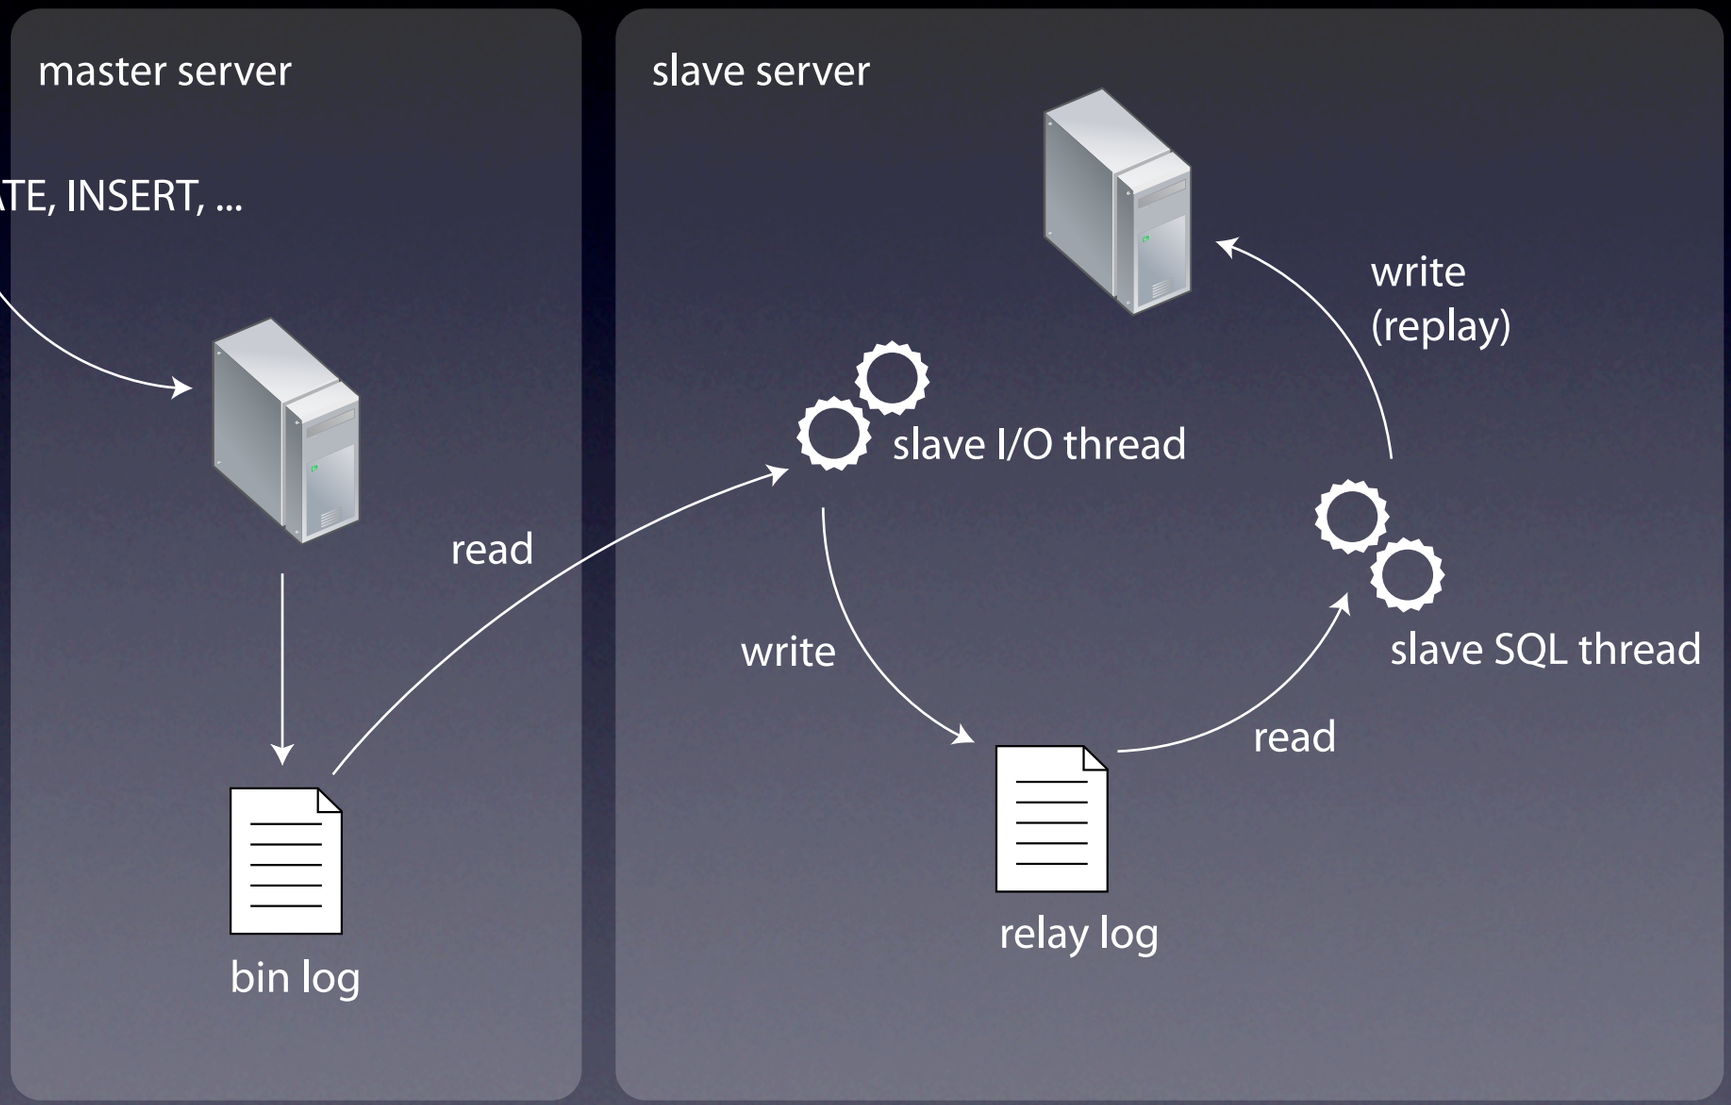

Central European Media Enerprises

Demo

(Setting up basic replication)

Central European Media Enerprises

Replication Topologies

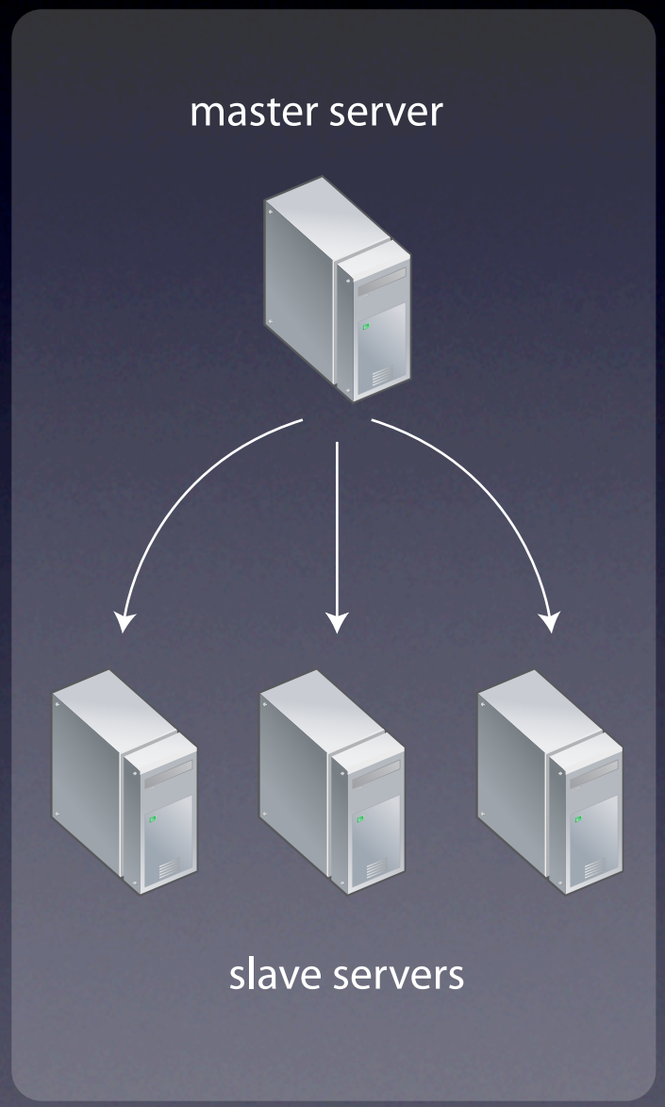

Master and Multiple Slaves

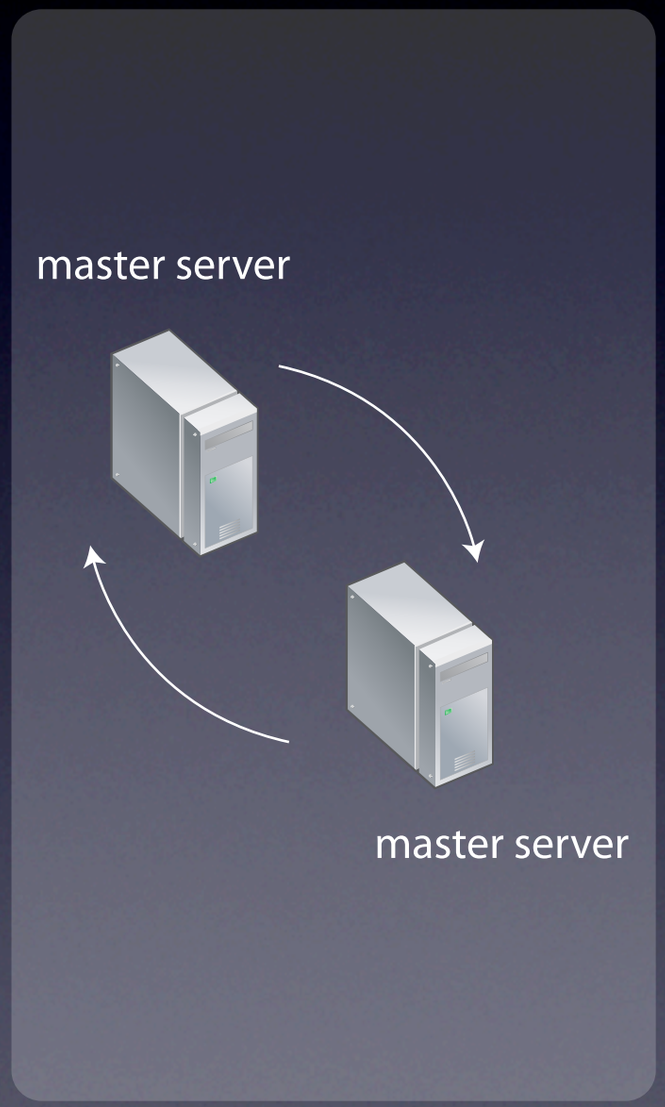

Master-Master

Central European Media Enerprises

Demo

(Going for a Master-Master setup)

Central European Media Enerprises

Distributing read queries across Slave servers

- Round-robin through DNS
- Linux Virtual Server (LVS)
  http://www.linuxvirtualserver.org
- MySQL Proxy
  http://dev.mysql.com/doc/refman/5.1/en/mysql-proxy.html
- Proprietary solutions
- Own custom solution

Central European Media Enerprises

Using PHP as you load-balancer

- Load-balancing is spread across all our webserver, making every webserver a load-balancer -> **no single point of failure!**
- `mysql.connect_timeout`
  Connect timeout in **seconds!!!**
  default = 60 sec
  minimum = 1 sec

Central European Media Enerprises

Is slave alive?

- Overcome `mysql.connect_timeout` issue with PHP Socket functions
  http://www.php.net/manual/en/intro.sockets.php

```
function is_slave_alive( $host ) {
   sock = @socket_create(AF_INET, SOCK_STREAM, SOL_TCP);
   socket_set_nonblock($sock);
   socket_connect($sock, $host, 3306);
   $error = socket_last_error();
   if ($error != SOCKET_EINPROGRESS && $error != SOCKET_EALREADY)
   {   socket_close( $sock );
       return FALSE;
   }
   socket_close( $sock );
}
```

Central European Media Enerprises

Part 2: Server-side optimization

Central European Media Enerprises

Inspect Server Status Variables

- The best way to tune MySQL server:
  `SHOW GLOBAL STATUS`
- There are tools that help you with that:
  - mysqltuner.pl
    http://blog.mysqltuner.com
  - tuning-primer.sh
    http://www.day32.com/MySQL/
- Predefined config files come with MySQL
  my-huge.cnf, my-large.cnf, my-small.cnf

Central European Media Enerprises

Demo

Central European Media Enerprises

"There is no way to optimize MySQL enough to compensate for poor application design."

Central European Media Enerprises

Part 3: Monitoring

Central European Media Enerprises

mtop

http://mtop.sourceforge.net

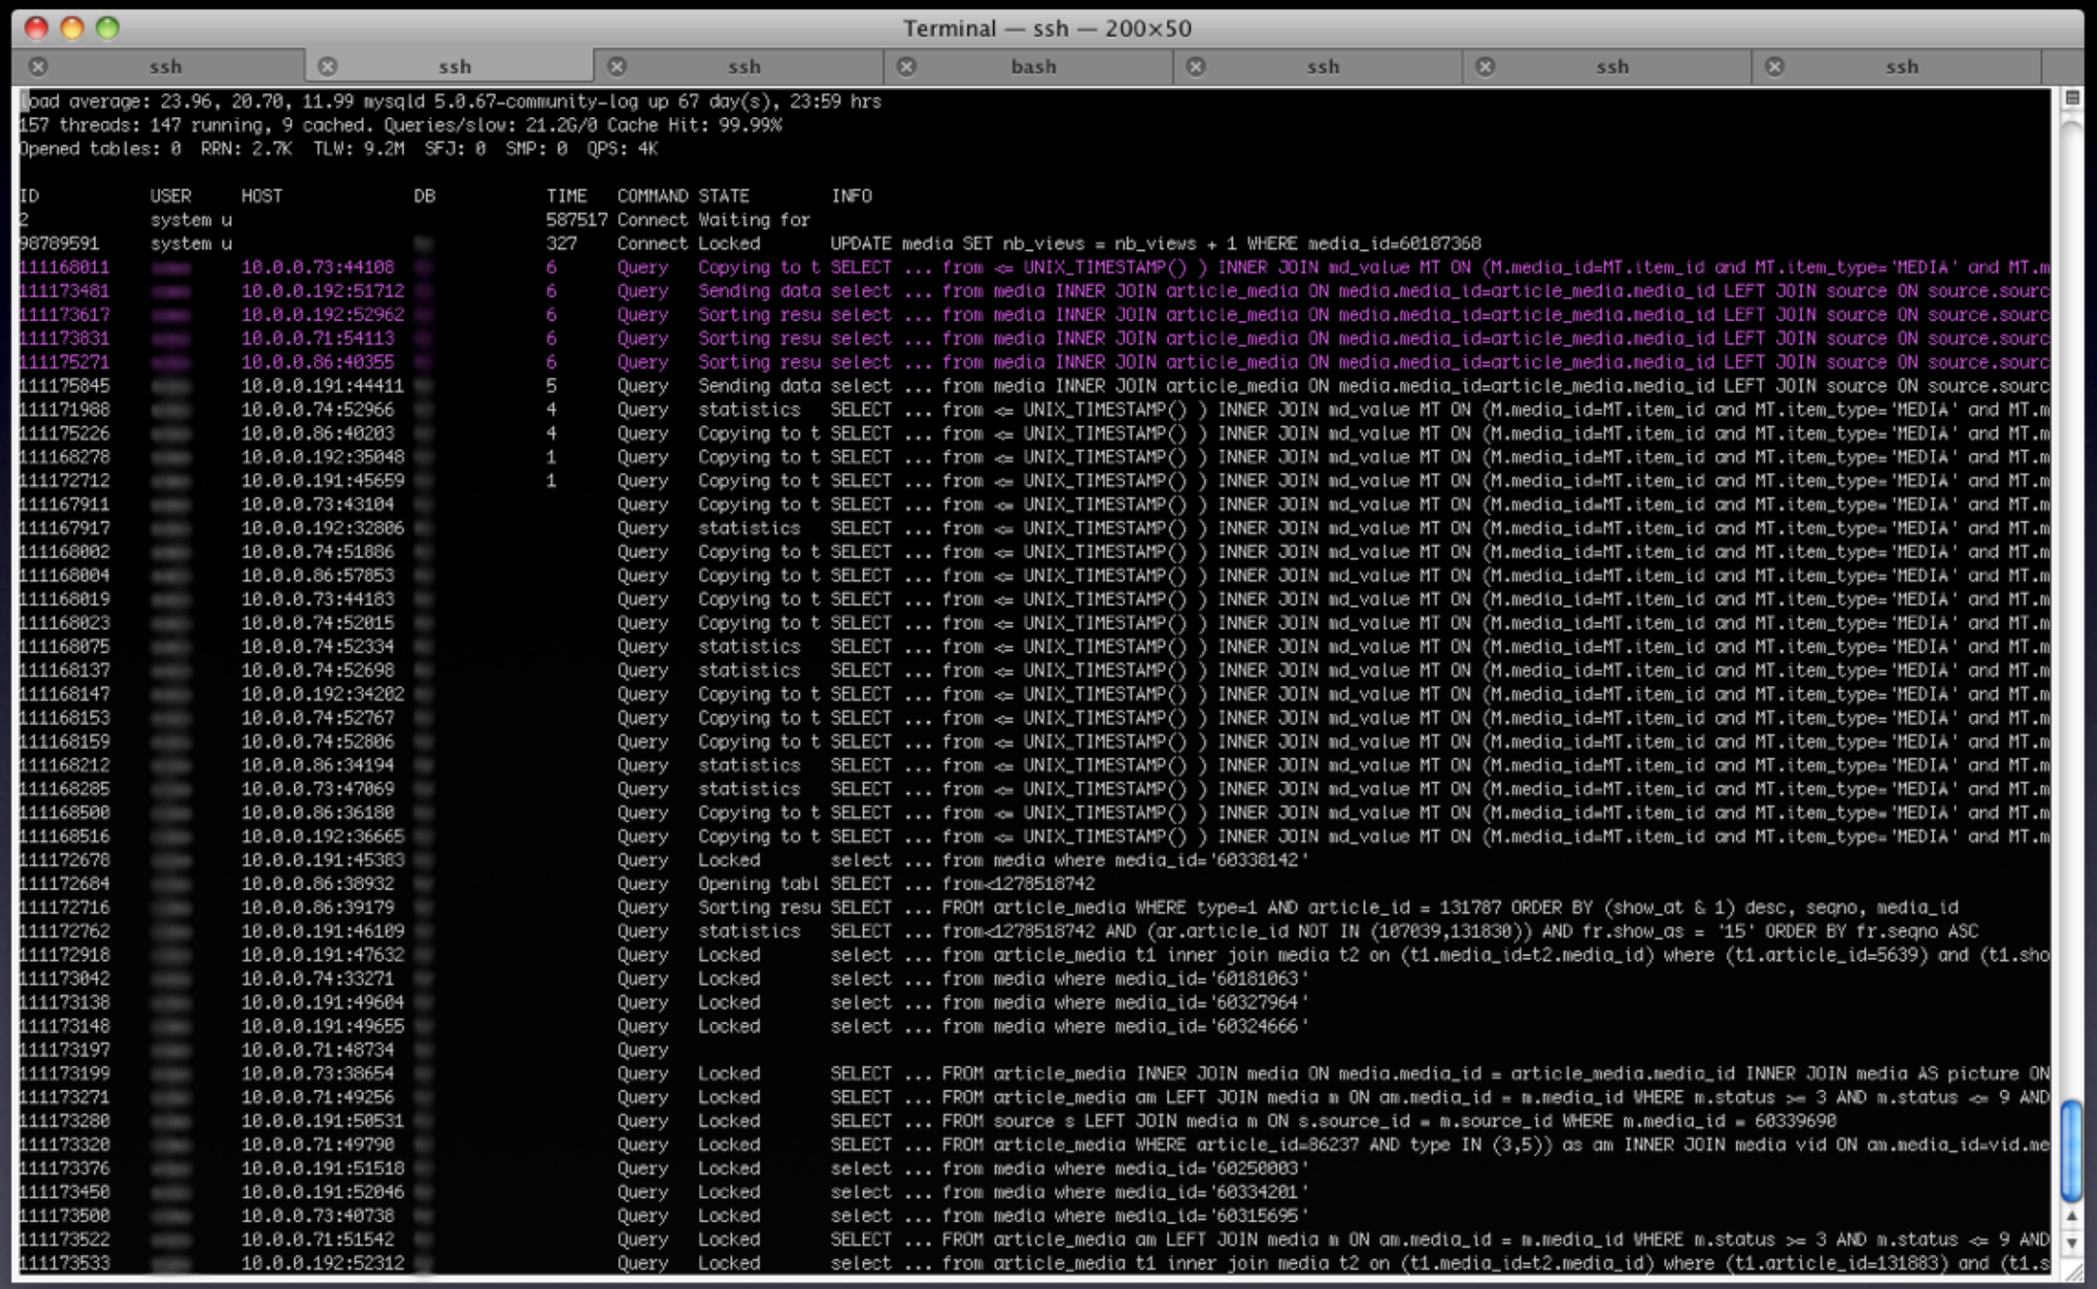

Central European Media Enterprises

Cacti

http://www.cacti.net

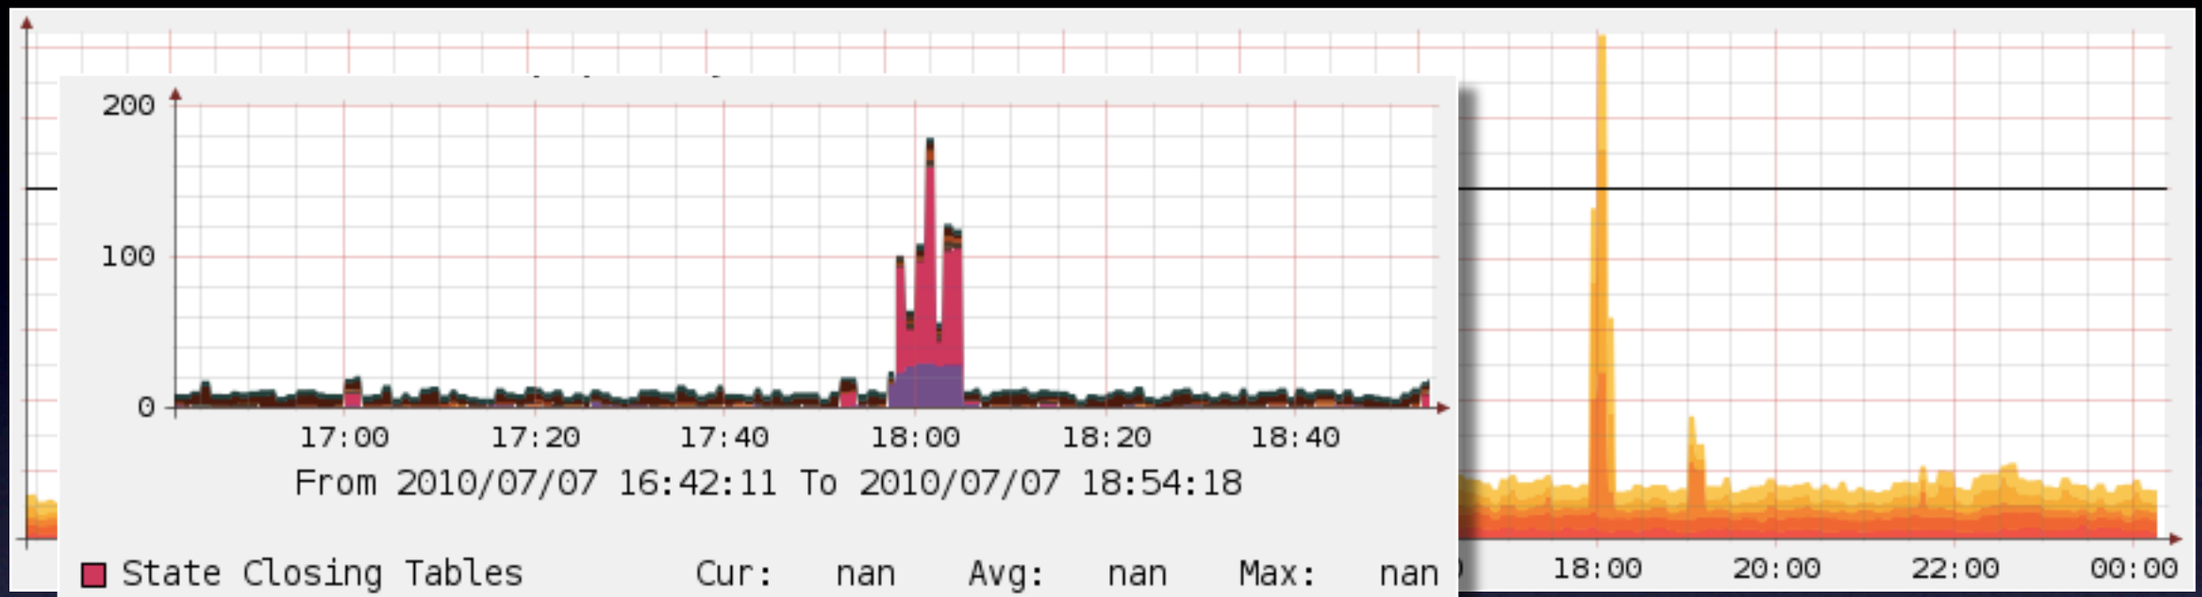

State Closing Tables	Cur:	nan	Avg:	nan	Max:	nan
State Copying To Tmp Table	Cur:	1.0	Avg:	6.2	Max:	29.0
State End	Cur:	nan	Avg:	nan	Max:	nan
State Freeing Items	Cur:	1.0	Avg:	1.0	Max:	1.0
State Init	Cur:	1.0	Avg:	1.0	Max:	1.0
State Locked	Cur:	8.0	Avg:	30.3	Max:	129.9
State Login	Cur:	nan	Avg:	nan	Max:	nan
State Preparing	Cur:	3.0	Avg:	1.7	Max:	3.0
State Reading From Net	Cur:	2.0	Avg:	1.1	Max:	2.0
State Sending Data	Cur:	1.0	Avg:	1.4	Max:	4.9
State Sorting Result	Cur:	2.0	Avg:	1.3	Max:	5.8
State Statistics	Cur:	1.0	Avg:	1.1	Max:	2.0
State Updating	Cur:	1.0	Avg:	1.0	Max:	1.0
State Writing To Net	Cur:	1.0	Avg:	1.0	Max:	1.0
State None	Cur:	6.8	Avg:	4.2	Max:	11.8
State Other	Cur:	3.0	Avg:	3.1	Max:	4.9

Central European Media Enerprises

dbmon.pl

http://faemalia.net/mysqlUtils

```
usage: ./dbmon.pl {--daemon|--terminal} [--help] [--rotate] [--killlong=N]
[--killvolatile] [--killconn] [--host=host[:port]] [--license]

Required argument:
  --daemon       run as a daemon, logging into
  --terminal     run in terminal, STDOUT is statistics, STDERR is queries

Optional arguments:
  --host=h[:p]   connect to a specific host and (optional) :port
  --killlong=N   kill >N-second-running things ('thing'==query by default,
                 but can be connection)
  --killvolatile also kill volatile (insert/update/delete) long-running
                 queries
  --killconn     kill long-connected connections, not long-running queries
```

Central European Media Enerprises

Other useful tools...

- innotop (like mtop)

  http://code.google.com/p/innotop

- maatkit (advanced tools for MySQL databases)

  http://code.google.com/p/maatkit

- and many others... :)

  http://www.mysqlperformanceblog.com/tools

Central European Media Enerprises

Further reading

- High Performance MySQL book (O'Reilly)
- MySQL Performance Blog
  http://mysqlperformanceblog.com
- The Twitter Engineering Blog
  http://engineering.twitter.com
- Digg Technology Blog
  http://about.digg.com/blog/technology

Central European Media Enerprises

Thank you.

twitter.com/igorajdisek
twitter.com/labs24ur