

# MySQL Query Benchmarking

Jeremy Zawodny  
Yahoo! Inc.

O'Reilly Open Source Convention  
Portland, Oregon  
Friday, July 11th, 2003

[Jeremy@Zawodny.com](mailto:Jeremy@Zawodny.com)  
<http://jeremy.zawodny.com/mysql/>


# About Me



- Long-time MySQL user
  - Using for 5-6 years
- Active on MySQL mailing lists
- Yahoo's resident “MySQL Geek”
  - Part of the platform engineering group
  - MySQL support, planning, tuning, etc.
- Write about LAMP for Linux Magazine
- Random MySQL consulting
- E-Mail: [Jeremy@Zawodny.com](mailto:Jeremy@Zawodny.com)



# MySQL at Yahoo!

- 200-400 servers world-wide
- Mission-critical 
- FreeBSD and Linux on commodity hardware
- Replaces home-grown DB systems, sometimes Oracle
- Used both “live” and in batch processing
- Replication and load-balancing to scale

# Yahoo! Is Hiring




It's not just Amazon and Ticketmaster.  
We need good, open source-loving  
hackers to help build some of the world's  
most popular web destinations.

We use FreeBSD, C/C++ (gcc), Apache,  
Perl, PHP, MySQL, vi(m), Emacs, and  
even some Java here and there.

I'm collecting resumes...

# Outline

- Why
  - What
  - When
  - How
- Mysql-bench
  - super-smack
  - mybench
  - Book Plug



# Why Benchmark

- Know when to expect problems
- To answer important questions:
  - How much load can you handle?
  - What's likely to break first?
  - Where will it break? How?
  - Does the server scale gracefully?
  - How will the application behave when busy?
  - Does your monitoring system notice?
    - You do have one, right? :-)

# Why Benchmark

- Compare alternative implementations
  - Queries
  - Indexes, Table Types
  - Application Logic
- Test server hardware configurations
  - Should I add RAM or disks?
  - Is RAID-10 really faster than RAID-5?
  - Will dual CPUs really help?
    - How much?
- Because it's easy

# What to Benchmark

- Individual queries
- A normal “session”
- The whole application
  - Requires more work/tools
  - Web apps aren't too difficult
    - Many load/stress testing tools available
    - Apache Bench (ab) is a simple start



# When to Benchmark

- Benchmark early and often
- During design
  - Helps make decisions about table structure, indexes, etc.
  - Know early on if you have sufficient hardware
- Before launch
- Anytime you're wondering “which is faster, A or B?”

# Quick Question

Who here has already setup a system for benchmarking MySQL?



# How to Benchmark

- Change one parameter at a time
  - Have a plan
  - Know what you're testing
  - Record your findings and configuration
- Always repeat test runs (3 times min.)
- Separate client and server
- Be aware of external influences
- Use realistic configurations

# Mysql-bench

- A good overall benchmark
  - Useful for comparing servers and configurations in a general way
  - Doesn't test your specific application
    - Unless you add the code, of course
- Tests can be run separately or as a whole
- Look at the `sql-bench` subdirectory of the MySQL distribution
- Demo


# Super-smack

- Separate multi-threaded C++ tool
- Build a configuration file to specify test conditions
  - Takes time to configure
  - Can generate data if needed
- Simulate any amount of load necessary (within reason)
- Demo
- <http://advanced-mysql.com/super-smack> (coming soon)

# Mybench

- Perl-based benchmarking framework
- Very simple to use
- Quite flexible, since it's Perl
- You really need to know a bit about DBI to use it
- Demo
- <http://advanced-mysql.com/mybench/>  
(coming soon)

# Coming Soon



# More Info

- MySQL's BENCHMARK() function  
[http://www.mysql.com/doc/en/Miscellaneous\\_functions.html](http://www.mysql.com/doc/en/Miscellaneous_functions.html)
- Custom Benchmarks for MySQL  
[http://www.mysql.com/doc/en/Custom\\_Benchmarks.html](http://www.mysql.com/doc/en/Custom_Benchmarks.html)
- MySQL super-smack  
<http://advanced-mysql.com/super-smack/>
- MyBench  
<http://advanced-mysql.com/mybench/>



# Questions?

