MySQL Proxy

WebTech 2007
Варна, България

29-ти Юни 2007

Андрей Христов <andrey@mysql.com>

Лекторът

Андрей Христов

- работи за MySQL като разработчик
- в момента част от частта занимаваща с клиентските библиотеки
- автор на MySQL Event Scheduler, част от 5.1 (дипломна работа)
- бърка в кода на PHP от години
- разработил почти изцяло mysqlnd
- поддържа ext/mysql и ext/mysqli заедно с Georg Richter

- ... ще се опитам да отговоря на въпросите си - просто ги задайте! Ако сте срамежливи, ще ви очаквам след лекцията.
- ... отговарям на, някои, писма изпратени до мен на andrey@mysql.com

Agenda

- Use cases
- Implementation
- Usage
- Lua scripting
- Next Steps

Proxy (lat. "proximus", ...)

- a proxy is someone who executes something on behalf of someone else
- proxy-servers forward requests to servers and can:
  - transform
  - handle or
  - block them

History

- implemented in a weekend session in March 2007
- designed to handle thousands of parallel connections (c10k)
- uses a embedded scripting language for customizations

Load Balancing

- load balancing distributes the load across several slaves
- SQF is default:
  - Shortest Queue First
  - send new connections to the server with the least number of open connections
- Other balancers via scripts

Fail Over

- dead hosts are detected
- taking out of load balancing for 2min
- uses custom load balancers to decide how to handle a dead host:
  - hot + standby
  - normal load balancing

Scripting

- proxy embeds LUA
- allows analyzing and manipulating packets:
  - Inspection
  - Rewriting
  - Blocking
  - Injection

Query rewriting

- Macro Packages (ls, cd, who, ...)
- tagging queries with SQL_CACHE
- migrating table-names
- turn EXPLAIN UPDATE|DELETE into equivalent EXPLAIN SELECT

Query Injection

- SHOW SESSION STATUS around a Query
- EXPLAIN to track index usage
- logging all generated warnings
- Global Transaction IDs

Transparency

- goal is to be transparent to the application layer
- SHOW WARNINGS can be worked around with Query Injection
- SELECT USER() shows the connected user (the proxy, not the client) which can be corrected with result-set rewriting

Latency

- early tests via localhost
- same script run directly and through the proxy
- latency per mysql-packet: 0.4ms
- `ping` RTT on 1Gbit: 0.1ms

Connection Pooling

- reusing open connections between proxy and server
- reduces concurrency on the MySQL Server
- external connection pool for PHP
- routing of Queries

Statement routing

- split the query stream into reading and writing Queries
- `READ`s go to the slaves
- `WRITE`s and transactions to the master
- automatic scale-out
- sharding

Future Work

- Replication
  - Filtering
  - Routing
  - Manipulation
- Parallel Inserts
- Delayed Replication

Въпроси ?