



Java is a trademark of Sun Microsystems, Inc.



JavaOneSM

Drizzle: A New Database for The Cloud

Monty Taylor
Sun Microsystems, Inc.
Drizzle Core Team

- > Drizzle Vision
- > Drizzle Microkernel Architecture
- > Drizzle Protocol
- > Drizzle Storage Engines
- > Drizzle Logging
- > Drizzle Replication
- > Drizzle Authentication (3A's)
- > Drizzle Data Dictionary
- > Drizzle Protocol
- > Drizzle Roadman

MySQL™ Cab '05

2005 OSCON -> 1/(20+)

Goals

- > Pluggable/Infrastructure Aware
- > Community Developed
- > MultiCore/Concurrency
- > Focus on Web Applications/Enable Others
- > Modernize codebase for Manageability
- > C++, STL, re-use libraries, etc

Philosophies

Sun's Team Values

- > Have open and well-documented interfaces
- > Have transparent goals and processes, that are communicated publicly
- > Have fun and encourage collaboration
- > Remove barriers to contribution and participation for anyone
- > Enable contributors to build a business around Drizzle

Rethink Everything

Do not play
catchup.

The world is 64bit, and there is a lot of RAM.

Интернет
является UTF-8.

□ □ Web UTF —

8 □ □ □

The Web is

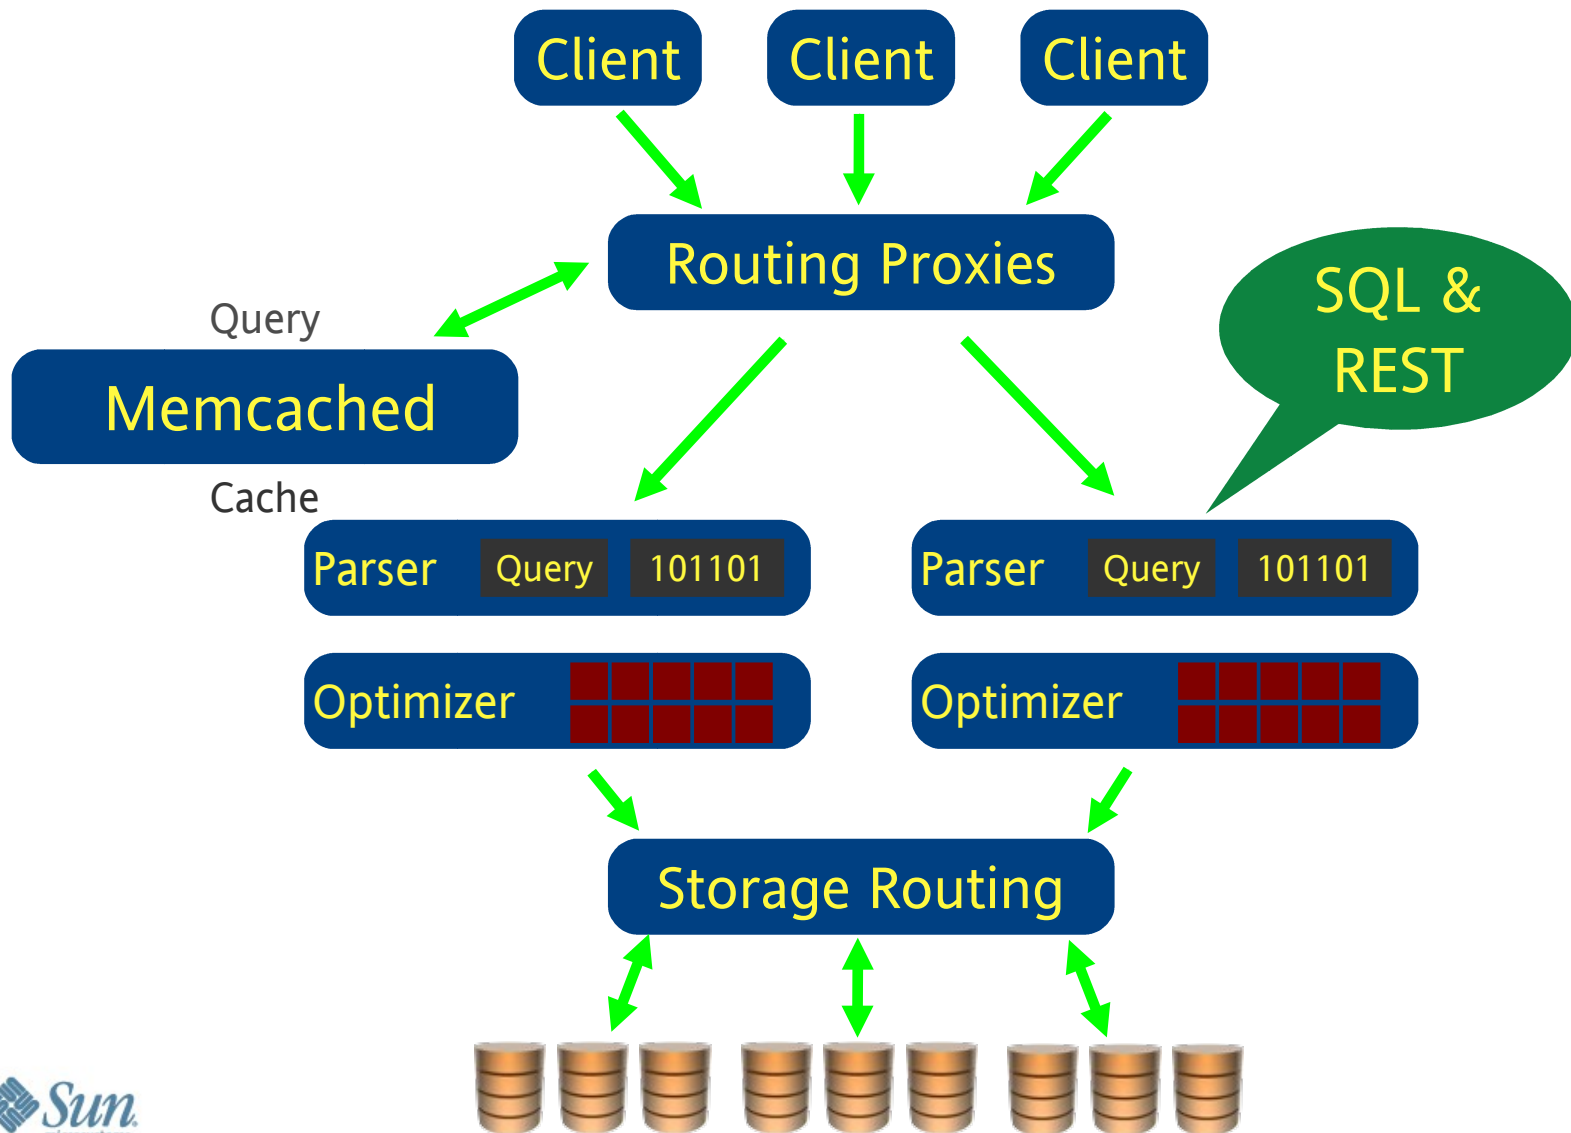
UTF-8.

We are part of an infrastructure

MicroKernel

- > Storage Engine
- > Replication
- > Logging
- > Authentication
- > Protocol
- > Parser
- > Error Messages
- > Scheduler

Architecture



How big is it?

6.0.5-alpha: 1,128,112

5.1.25-rc: 1,009,108

5.0.51b: 999,382

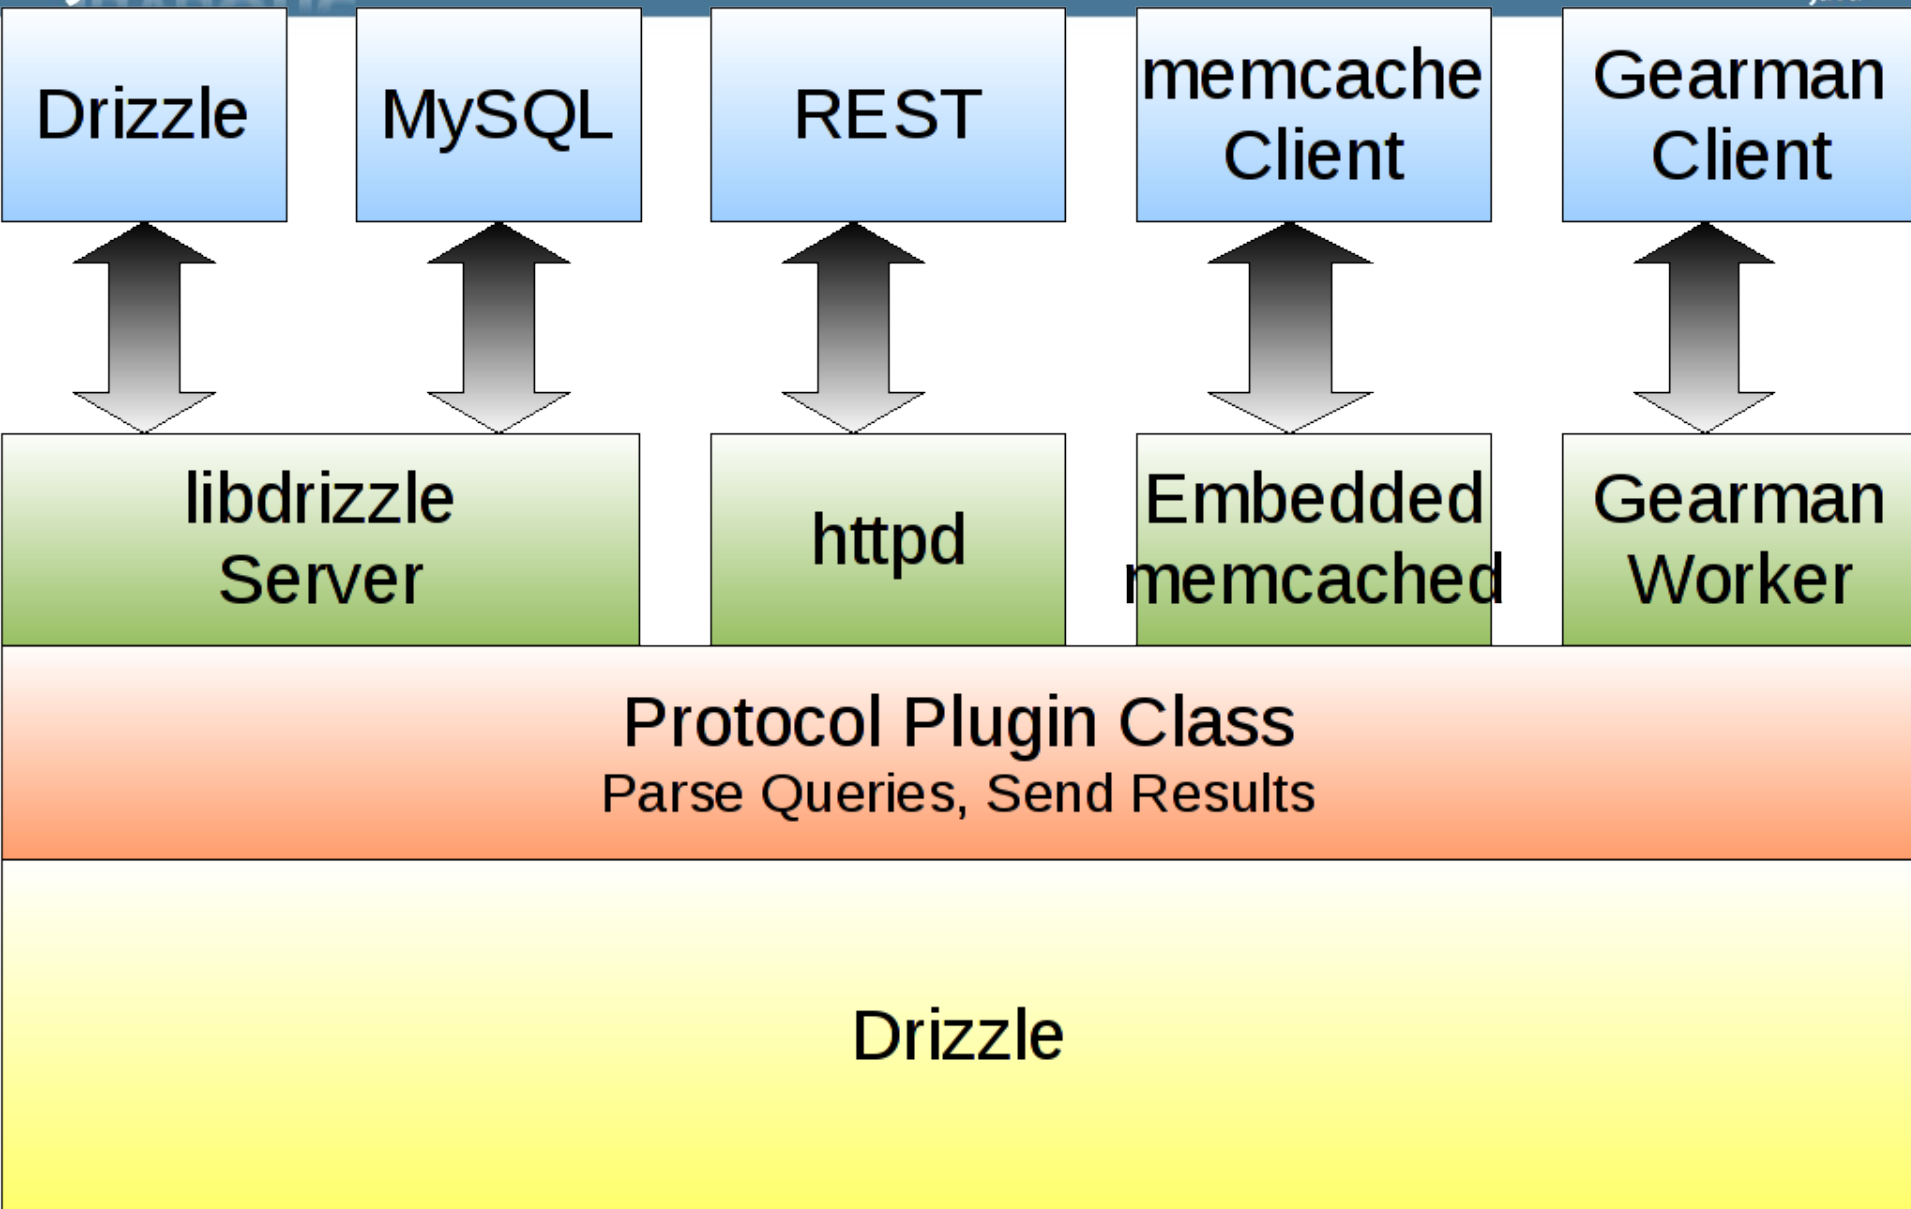
Drizzle: 315,560

generated using David A. Wheeler's 'SLOCCount'.

So where are we?

New Protocol

- > Asynchronous
- > Built In Sharding
- > UDP
- > Pluggable
 - Speak old/new/REST/other
- > Remove weak attack methods
- > Checksums



Storage Engines

- > Retain Multiple Engine Support
- > Default to ACID compliant Engine
 - InnoDB Plugin
- > Engines own their own MetaData

Logging

- > Pre/Post API
- > Gearman Logger
 - Query Analyzer
- > Syslog/etc (whatever you want to plugin)

Replication

- > New API interface to events
- > Google Protobuffer message format
 - Easy to extend/read from any language
 - Java
 - Python
 - C++
- > Multi-Database Synchronizer in Prototype

Security, 3A's

- > Focus on Authentication, Authorization, and Access Control
- > No cost for non-usage
- > Plugins:
 - PAM/ldap
 - HTTP AUTH

Extending SQL via Functions

- > Pluggable Interface via SWIG
 - Java via JNI
 - Python
 - Ruby
 - lua
- > Extensions allow for integration with Map/Reduce Systems
 - Gearman
 - Hadoop

Data Dictionary/Performance Interface

- > No Materialization
- > Single Execution path means easier to maintain
- > Extends capabilities of the current IS Plugin Interface
- > Leveraging work by MySQL Product Development

Open Source

Open Source Methodology

- > Internal/External Contributions are treated Equally
- > All Project Information is Public
- > Release Early/Release Often
 - 4 month Milestone cycles
 - 1~2 Week tarball Releases

Open Source Libraries

- > PCRE
- > Readline
- > Curl
- > Gnu TLS
- > Google tcmalloc
- > Intel TBB

Open Source Tools

> Launchpad

- Bugs
- BZR
- Blueprints
- Translations

> Automated/Distributed Regression Testing

- Buildbot
- Hudson

gettext translations

30+ Languages
(within 3 months)

100+ Contributors

(Over 500 people on the mailing list)

Innovation Happens Elsewhere

FederatedX

- > Patrick Galbraith (Principle Engineer/Lycos)
- > Antony Curtis (Google)
- > Transactional
- > Multi-System
- > Integrated
- > High-Performance

New JDBC Driver

- > Marcus Eriksson
- > <http://launchpad.net/drizzle-jdbc>

Transport/libdrizzle

- > Eric Day (Architect XO)
- > Asynchronous
- > Multi-protocol
- > Partitionable

Distributed Query Cache

- > Toru Maesaka (Research Engineer/Mixi.jp)
- > Non-locking
- > Cloud Distributed (Memcached)
- > Pluggable

Google Protocol Buffers

- > Internal Data Serialization
- > New Replication System

Roadmap: Cirrus

- > New Replication Events
- > Object Based Plugin Loader
- > Removal FRM
- > Libdrizzle
- > Adopt InnoDB Plugin
- > C++/STL

Roadmap: Aloha

- > New Table Discovery
- > Multi-Replication Applier
- > Dependency Checking for Plugins
- > New Information Schema/Backend
- > Refactor store_lock

More Information?

- > <http://launchpad.net/drizzle>
 - bzip branch lp:drizzle
- > <http://launchpad.net/drizzle-discuss>
 - (mailing list)
 - Freenode #drizzle



JavaOneSM

Thank You

Monty Taylor
mordred@sun.com

<http://launchpad.net/drizzle>

