

How to make ~~MySQL~~ Drizzle cool again?

Originally by:
Jonathan Levin

Adapted by:
Henrik Ingo

Drizzle Developer day
2011-04-15

Gearman

- parallelize queries / "multi-threaded php"
- "asynchronous triggers"
- MapReduce worker?

Hadoop

- RDBMS on local nodes:
<http://db.cs.yale.edu/hadoopdb/hadoopdb.html>
- As storage engine (ie Hive)

BashReduce

- SSH into N boxes and do grep|awk|magic

Column-based storage engines

- index-less, big data on community hw

Shiny new HW

- RAM, SSD

Out of the box good configuration (interactive)

- not to mention adaptive configuration?

Flexviews (Materialized views)

- read binlog to keep materialized view updated.

Replication

- parallel replication
 - Tungsten, Galera, MySQL56
- (semi)synchronous replication for HA
 - Galera, MySQL55
- synchronous replication for master-master
 - Galera
- Global trx id
 - Tungsten, Galera, Google patch

Sharding

- Transparent sharding
 - NDB, Spider, Swanhart, Wendel, Fremont?
- Online, automatic re-sharding/add node
 - NDB, others?

NoSQL

- HandlerSocket or Memcached?
- key-value store: `CREATE TABLE t (k INT, v BLOB);`
- schema-less / document store
 - secondary index into the BLOB:
 - virtual columns (MariaDB)
 - dynamic columns (MariaDB)
 - return "bags" instead of joined table
 - <http://openlife.cc/blogs/2010/october/mariadb-52-using-mariadb-column-store-and-virtual-columns-indexing>
 - Embed JSON parser
- storage engines: graphdb, columnar db, etc...
- parser plugins: graph query language, json...

How to make MySQL uncool again

- Online Backup
- Online ALTER TABLE, OPTIMIZE TABLE
- ANSI SQL (yes, sub queries)
- Stephane: Use Gearman as encrypting proxy for PCI compliance.

Sorry, but this is the stuff that old boring people actually need...

Observations

- Lot of this exists or is low hanging fruit -> packaging exercise
- Lot of this is shared MySQL vs Drizzle

Drizzle Suite

Drizzle

Vanilla MySQL
Percona, MariaDB

Drizzle Suite:
Memcached, Gearman, Tungsten, Shard-query, Maatkit...