

MySQL and the Open Source Sphere

2011

Shlomi Noach

openark.org

August 2011

MySQL and the Open Source Sphere


- MySQL plays a major role as a free RDBMS
 - It is very common in the open source world
 - It is labeled as “The most popular open source database”
- How does MySQL participate in the open source sphere?
 - In terms of cost-free?
 - In terms of source code?
 - In terms of forks and patches?
 - In terms of satellite projects?
 - In terms of community?

The 20 seconds history brief

Development begins, MySQL AB founded, 1994

```
#ifndef MySQL
#define MySQL
```

MySQL Releases 3.23 – 5.0.X 1995 – 2008
Oracle acquires InnoDB Oy, providers of InnoDB, 2005



Sun Microsystems acquires MySQL AB, 2008
MySQL 5.1 released



Oracle Corporation acquires Sun Microsystems 2010,
after long debate with EU regulator

ORACLE®

MySQL control

- The MySQL database and trademark are owned by Oracle Corporation
- Oracle:
 - Develops MySQL (including InnoDB)
 - Lays out milestones and release dates
 - Releases MySQL
- And so MySQL's development is fully under Oracle's control
- Some like to refer to MySQL as “*an open source product*”, as opposed to “*an open source project*”.

MySQL licensing

- MySQL is released under a dual license:
 - GPLv2
 - OEM Commercial license
- The GPL plays particularly well into the hands of web providers (**Facebook, Google, Flickr, Yahoo!** Etc.)
- Many other installations comply with the GPL
- Those who cannot or do not wish to comply with GPL may acquire a paid commercial license.
- *I am not an attorney.* Read more on:
 - <http://www.mysql.com/about/legal/licensing/index.htm>
 - <http://www.mysql.com/about/legal/licensing/oem/>

MySQL Source Code

- Since MySQL is available in GPL, source code is also freely available. This includes:
 - MySQL server
 - Storage engines: InnoDB, NDB, MyISAM, etc.
 - Client libraries: libmysqlclient, Connector/J, etc.
 - MySQL Workbench
- Enterprise tools, such as the MySQL Enterprise Monitor, are not open sourced.

MySQL forks

- There are several major forks to MySQL:
 - **Percona Server**, by Percona
 - **MariaDB**, by Monty Program
 - **Google Patches for MySQL**, by Google
 - **Facebook Patches for MySQL**, by Facebook
 - **Drizzle**, by various contributors
- The forks all attempt to improve the MySQL baseline.
- However they differ in orientation and usage.

Percona Server

- The *Percona Server* is a fork by **Percona**, a MySQL expert consulting company which plays a major role in this sphere.
- Percona Server offers *XtraDB*, a fork of the InnoDB storage engine, as well as numerous performance and stability improvements.
- Percona offers payed support for Percona Server.
- However any change is released as GPL

MariaDB

- *MariaDB* is developed by **Monty Program AB**, a company founded by Michael (Monty) Widenius, the original author and owner of MySQL.
- There is much development on new & rich feature set, including many patches by the community
- MariaDB maintains compatibility with MySQL, and has payed support.
- MariaDB is released under the GPL license

Google Patches for MySQL

- Developed by the MySQL team at **Google**.
- Developed for the purpose of making MySQL more useful to Google.
- Known for scalability patches, InnoDB patches, replication patches.
- Patches released under the BSD license.

Facebook Patches for MySQL

- Developed by the MySQL team at **Facebook**.
- Developed for the purpose of making MySQL more useful to Facebook.
- Similarly to Google, Facebook concentrate on patches for high scalability and usability.
- Patches released under the BSD license.

Confused?

- Much like in the UNIX/BSD/Linux history, patches created by one provider are quickly adopted by others.
- Sometimes such patches are upstreamed back to the MySQL baseline.
 - Upstream is not too common.
 - In some cases MySQL developers rewrite or re-code patches found in forks (e.g. InnoDB recovery time bugfix)
 - Some important solutions, existing in forks for years, are yet to be implemented in MySQL (e.g. global transaction IDs)

Open Source tools for MySQL

- There is a rich and vivid 3rd party tools development for MySQL in all possible fields:
 - Backup/restore
 - Monitoring
 - Analysis/profiling
 - Maintenance
 - Usability
 - Replication
 - Caching
 - More...
- See:
 - [10 essential MySQL tools for admins](#) By Daniel Nichter
 - [Making MySQL Administration a Breeze - A Look Into a MySQL DBA's Toolchest](#), by Lenz Grimmer

Notable open source projects: backup

- *Xtrabackup*: developed by **Percona**, forked off the InnoDB-Backup tool, this is a supported hot-backup tool for InnoDB. Features:
 - Backup+recovery
 - Streaming
 - I/O control
 - Single table backup
 - Incremental backups

Notable open source projects: backup/recovery

- *mylvmbackup*: A Perl script for making hot backups using *LVM*
 - Written by **Lenz Grimmer**.
- *MyDumper/MyLoader*: new kids on the block, allowing for parallel export/import of a read-only database, making for faster logical backups.
 - Written by **Domas Mituzas, Andrew Hutchings, Mark Leith**.
- InnoDB Recovery Tools: a set of tools offering manual disaster recovery of lost InnoDB data (have you accidentally dropped your database recently?)
 - Written by the **Percona** team

Notable open source projects: analysis/maintenance

- *Maatkit*: a very large set of tools aimed to provide “simple, predictable ways to do things you cannot otherwise do” with MySQL.
 - Includes some extremely useful tools, such as *mk-query-digest*, *mk-table-sync* and many more.
 - A popular, tested and active project
 - Originally written by **Baron Schwartz**.
 - In part maintained by **Percona**.

Notable open source projects: analysis/maintenance

- *Aspersa*: a set of tools, some of which are building-block tools, mostly aimed at analyzing and diagnosing MySQL performance, though also features some Linux-oriented diagnostics
 - Written by **Baron Schwartz**.
 - Soon to merge with *Maatkit* into “Percona Toolkit”.
- *openark-kit*: a set of tools aimed to fill in missing functionality within MySQL, or ease out daily DBA maintenance work.
 - Written by **Shlomi Noach**

Notable open source projects: analysis

- *mysqltuner*: A Perl script which examines a running server and provides tuning recommendations.
 - 2.0 version written by **Sheeri K. Cabral**
 - Earlier versions written by **Major Hayden**

Notable open source projects: monitoring

- *Better-cacti-templates*: MySQL-specific templates for the Cacti monitoring system. Offers deep insight into MySQL and InnoDB in particular.
 - Written by **Baron Schwartz**
- *mycheckpoint*: a lightweight, SQL-oriented monitoring solution for MySQL, offering custom queries, alerts, textual and graphic reports.
 - Written by **Shlomi Noach**

Notable open source projects: replication

- *Tungsten replicator*: a data replication engine for MySQL, featuring:
 - Multi master replication
 - Parallel replication
 - Global transaction Ids
 - More...
 - Developed by **Continuent, Inc.**

Notable open source projects: caching/NoSQL

- *HandlerSocket*: A NoSQL interface reading/writing directly from/to InnoDB tables, allowing for full featured transactional behavior and integrity along with faster, simpler access.
 - Written by **Yoshinori Matsunobu, DeNA**
 - MySQL **5.6** offers *memcached daemon plugin*, using similar concepts

Notable open source projects: test/install

- *mysqlsandbox*: a one-step setup tool for MySQL installations, including replication setup of several topologies, resulting in an isolated set of MySQL servers. Extremely useful in test environments.
 - Written by **Giuseppe Maxia**.

The MySQL Community

- MySQL enjoys a large and active community.
- There are many technical blogs covering MySQL aspects.
 - Aggregated in <http://planet.mysql.com/>
- Discussion groups, IRCs, forums are all available.

The MySQL Community

- There are various ways of meeting the community:
 - The annual *MySQL O'REILLY conference*
 - *OpenSQL camp*: the open databases meet-up
 - *Percona Live*: organized by Percona
 - Various Oracle conferences
 - Local user group meetings
 - There are many more conferences and meetings; keep an eye!

Conclusion

- This is not the “*MySQL is the best thing EVER*” kind of talk. It has its shortcomings.
- MySQL is not a completely open source project.
- However, the community has its say, and is a dominant part in the MySQL world.
- Oracle shows great development activity with MySQL.
- There is *a lot* going on, and it is getting harder to keep track.

Conclusion: you

- How can you help out?
 - As with all open source projects, the first and foremost way to support a product/tool is to install and use it.
 - But your feedback is essential! File bug reports, offer insights, request features, tell your mother.
 - If you like to write, people would like to read.
 - If you like to code, your help will be gladly accepted by many.