


The MySQL Ecosystem - understanding it, not running away from it!

Colin Charles, Chief Evangelist, Percona Inc.
colin.charles@percona.com / byte@bytebot.net
<http://bytebot.net/blog/> | @bytebot on Twitter
FOSDEM, Brussels, Belgium
4 February 2018



whoami



- Chief Evangelist, Percona Inc
 - Focusing on the MySQL ecosystem (MySQL, Percona Server, MariaDB Server), as well as the MongoDB ecosystem (Percona Server for MongoDB) + **100% open source** tools from Percona like Percona Monitoring & Management, Percona xtrabackup, Percona Toolkit, etc.
- Founding team of MariaDB Server (2009-2016), previously at Monty Program Ab, merged with SkySQL Ab, now MariaDB Corporation
- Formerly MySQL AB (exit: Sun Microsystems)
- Past lives include The Fedora Project (FESCO), OpenOffice.org
- MySQL Community Contributor of the Year Award winner 2014

A Mature Ecosystem

- MySQL: nearly 23 years old — May 1995
- Percona Server: 9+ years old — November 2008
- MariaDB: 8 years old — February 2010
- Drizzle: 2008 - 2012 (R.I.P.)
- WebScaleSQL: March 2014 - December 2016

Learnings from WebScaleSQL

The collective goal of the companies was to enable the scale-oriented members of the MySQL community to work more closely together in order to prioritize the aspects that are most important to us. We learned over the course of this experiment, however, that sometimes each company need may not be similar enough to make sustained collaboration possible. Everyone was at different stages of MySQL lifecycle timelines; it wasn't productive to try to work on the same one. That's why, as you may have noticed, the WebScaleSQL collaborators are no longer contributing to the WebScaleSQL branch.

Many of the WebScaleSQL companies continue to make MySQL better through other projects: Alibaba with many changes for InnoDB performance, Facebook with MyRocks, Google with encrypted InnoDB. We'll continue a meet-up of WebscaleSQL participants at Percona Live once a year. It is always great to exchange ideas.



#mysqldevroom #FOSDE





Open source community

- MariaDB: takes external contributors/committers, participates in Google Summer of Code
- MySQL: 5.7 takes Generated Columns (virtual columns in MariaDB 5.2) from Andrey Zhakov
 - contributions welcome, commits not
 - see: MySQL Community Contributor Award Program
- Percona: bug reports welcome, commits not
- Contributor agreements: Oracle Contributor Agreement (OCA), MariaDB Contributor Agreement (MCA) and the BSD New

Is MySQL dying?

- *“The reports of my death have been greatly exaggerated”* — Mark Twain
- MySQL ecosystem development is at its **most vibrant now** than it has ever been
- Oracle has been a **great** steward of pushing MySQL development forward

MySQL 5.7 -

www.thecompletelistoffeatures.com

- Multi-source replication
- Dynamic replication filters
- Lossless semisync
- SHOW EXPLAIN for connection_id
- GIS functionality
- Statement timeouts
- Change master without stopping SQL thread
- Online GTID implementation
- GTID no longer requires log-slave-updates to be enabled
- Virtual columns (generated columns)
- Online buffer pool resize
- Username size increase
- LOCK/UNLOCK accounts
- JSON + MySQL Shell + X DevAPI
- Encryption at rest

MariaDB 10.1

- START TRANSACTION WITH CONSISTENT SNAPSHOT
- Integrated Galera Cluster
- Table/tablespace encryption
- Optimistic parallel replication
- Enhanced semi-sync replication
- InnoDB defragmentation
- ANALYZE <statement>
- Threadpool
- cracklib_password_check
- SQL error logging plugin
- Extended REGEXP (PCRE)
- Roles

MariaDB 10.2

- InnoDB as default InnoDB
- MyRocks alpha
- Window functions
- Recursive Common Table Expressions (CTEs)
- AWS Key Management plugin
- CHECK CONSTRAINT
- SHOW CREATE USER
- Multiple triggers for the same event
- EXECUTE IMMEDIATE (Oracle styled)
- DML only flashback

Percona Server 5.6/5.7

- Percona XtraDB (many improvements, parallel doublewrite buffer, etc.)
- variable: numa_interleave
- restrict # of binlog files - max_binlog_files
- Lock-free SHOW SLAVE STATUS NOLOCK
- Percona Toolkit UDFs
- Expanded fast index creation - expand_fast_index_creation
- Utility user - system access to admin tasks, limited access to user schema
- Slow query log enhancements
- Log all client commands to syslog
- Improved Memory storage engine
- Per-query variable statement
- PROXY protocol support
- Backup locks
- TokuBackup

A base blog post resource

- High level, *answer* to a whitepaper
- <https://www.percona.com/blog/2017/11/02/mysql-vs-mariadb-reality-check/>

The screenshot shows the Percona website header with navigation links: Services, Software, Solutions, Community, Resources, and About Percona. Below the header is an orange banner with the text "Percona Database Performance Blog". The main article title is "MySQL vs. MariaDB: Reality Check" by Chris Chiles, dated November 2, 2017. The article is categorized under "MySQL", "MySQL for DBAs", "MySQL", and "MySQL". The introduction states that the goal is to evaluate MySQL, MariaDB, and Percona Server for MySQL side-by-side to assist in decision-making. A table compares MySQL, Percona Server for MySQL, and MariaDB Server across four categories: Protocols, Community - Source Code, Community - Development, and a fourth category. The table shows that Percona Server for MySQL and MariaDB Server are both Open Source and support MySQL protocol over port 33060. The fourth category shows that Percona Server for MySQL and MariaDB Server are both Open Source and support contributions via the Oracle Contributor Program or the MariaDB Contributor Agreement.

	MySQL	Percona Server for MySQL	MariaDB Server
Protocols	MySQL protocol over port 3306, X Protocol over port 33060	MySQL protocol over port 3306, X Protocol over port 33060	MySQL protocol, MariaDB Server extensions
Community - Source Code	Open Source	Open Source	Open Source
Community - Development	Open Source, contributions via signing the Oracle Contributor Program or the MariaDB Contributor Agreement	Open Source	Open Source, contributions via the new BSD license or signing the MariaDB Contributor Agreement

#mysqldevroom #FOSDEM

Define: compatibility (OED)

- A state in which two things are able to exist or occur together without problems or conflict.

1.2 *Computing* The ability of one computer, piece of software, etc. to work with another.

'software compatibility is another important factor to consider'

Why this matters




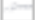
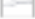
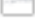
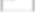


- MariaDB Server is the “default” MySQL in pretty much every Linux distribution except Ubuntu (shipping MySQL 5.7)
- Take a page from the cloud operators & their offerings
- Amazon Web Services (AWS) RDS MySQL & MariaDB, Microsoft Azure MySQL & MariaDB, Rackspace Cloud MySQL & MariaDB (and Percona Server)

Commitments

- Verbal commitments: **“MySQL 5.6, should be comparable to MariaDB Server 10.1. And for 10.2 it should be compatible with MySQL 5.7”** — Michael “Monty” Widenius, CTO of MariaDB Corporation and MariaDB Foundation, 7 October 2016, MariaDB Developer’s Meeting, Amsterdam
- <http://mariadb.org/about/>
- “It is an enhanced, drop-in replacement for MySQL.”

Governance

- MariaDB
 - MariaDB Corporation
 - MariaDB Foundation
 - <https://mariadb.org/about/governance/>
- MySQL
 - Oracle (by way of Sun, MySQL AB)
- Percona
 - Percona Inc

Number of Funding Rounds	9	Total Funding Amount	\$98.2M	
Announced Date	Transaction Name	Number of Investors	Money Raised	Lead Investors
Sep 29, 2017	 Series C - MariaDB	6	\$27M	Alibaba
May 8, 2017	 Venture Round - Mar...	1	€25M	European Investment Bank...
Apr 5, 2016	 Series B - MariaDB	1	\$3M	—
Jan 21, 2016	 Series B - MariaDB	2	\$9M	—
Feb 26, 2015	 Series B - MariaDB	1	€3M	—
Oct 23, 2013	 Series B - MariaDB	5	\$20M	Intel Capital
May 30, 2012	 Series A - MariaDB	1	\$2.5M	—
Apr 4, 2012	 Series A - MariaDB	3	\$4M	OnCorps, Inc.
Sep 1, 2010	 Venture Round - Mar...	1	\$2M	—

Releases

MariaDB	MySQL
5.1: 1 Feb 2010	5.1: 14 Nov 2008
5.2: 10 Nov 2010	
5.3: 29 Feb 2012	
5.5: 11 Apr 2012	5.5: 3 Dec 2010
	5.6: 5 Feb 2013
10.0: 31 Mar 2014	
10.1: 17 Oct 2015	
	5.7: 21 Oct 2015
10.2: 23 May 2017	

Replication

Replication Compatibility

Slave↓ Master→	MariaDB- 5.5	MariaDB- 10.0	MariaDB- 10.1	MariaDB- 10.2	MySQL- 5.6	MySQL- 5.7	MySQL-8.0
MariaDB-5.5	Ok	No	No	No	No	No	No
MariaDB-10.0	Ok	Ok			Ok		
MariaDB-10.1	Ok	Ok	Ok		Ok		
MariaDB-10.2	Ok	Ok	Ok	Ok	Ok	Ok	
MySQL-5.6					X	X	X
MySQL-5.7					X	X	X
MySQL-8.0					X	X	X

Note X: Refer to MySQL documentation

Note: When replication from MySQL in GTID mode, MariaDB will remove the MySQL GTID events and replace them with MariaDB GTID events.

Replication

- Default binlog format is now MIXED in MariaDB (ROW in MySQL)
- Default `replicate_annotate_row_events` is ON
- Time delayed replication arrives in MariaDB Server 10.2 (in MySQL 5.6)
- DML only Flashback - rollback instances/databases/tables to an older snapshot (via Alibaba!) - unique to MariaDB
- Synchronous replication
 - Galera Cluster: MariaDB Server has it built-in, Percona XtraDB Cluster 5.7 comes with ProxySQL + an admin tool
 - Group replication: works only on MySQL, Percona Server

X Protocol

- MariaDB Server has **no** support for the MySQL X Protocol
- This means you cannot use `mysqlsh` to access MariaDB Server either

Encryption

- MySQL 5.7 and MariaDB Server 10.1+ implement encryption differently (one is fully tablespace encryption, the other is based on the Google patch for tablespace encryption in addition to having table encryption)
- One does not encrypt logs, the other does
- MySQL requires `innodb_file_per_table`
- MySQL implementation works fully with Percona XtraDB Cluster

Security

- MySQL: sha256_password
- MariaDB: ed25519 password plugin
- validate_password is on by default in MySQL 5.7 (not in MariaDB)
- SSL: 5.7 has pre-generated keys, enable it in the client; disabled by default in MariaDB Server
- MySQL links against YaSSL, Percona Server against OpenSSL (changing in 8.0 for MySQL too)

mysql.user table changes

- MariaDB Server and MySQL differ here (not just by addition of roles)
 - `mysql.user.password` is just `mysql.user.authentication_string`
- Password expiry is coming? <https://jira.mariadb.org/browse/MDEV-7597>
- Password last changed? Lifetime?
- ACCOUNT LOCK/UNLOCK
- `VALIDATE_PASSWORD_STRENGTH()` SQL function doesn't work in MariaDB Server

Stuff that might matter to you

- Optimizer hints - <https://jira.mariadb.org/browse/MDEV-9078>
- RENAME INDEX - <https://jira.mariadb.org/browse/MDEV-7318>
- Query rewriting? - <https://jira.mariadb.org/browse/MDEV-5561>
- MySQL SUPER READONLY - <https://jira.mariadb.org/browse/MDEV-9458>
- Optimiser trace: <https://jira.mariadb.org/browse/MDEV-6111>
- PERFORMANCE_SCHEMA - 5.7 improvements, MariaDB still ships 5.6 P_S

Tools

- including new tools like `mysql_ssl_rsa_setup` ? `mysqlpump`?
- tools that require MySQL GTID don't work with MariaDB Server (e.g. `mysqlfailover`, `mysqlrpladmin`, MHA for GTID based failover, MySQL Router, etc.)

Ecosystem

- Percona Toolkit
- Percona XtraBackup
- ProxySQL
- MariaDB MaxScale
- mydumper
- MHA
- Tungsten Replicator
- vitess
- Tumblr JetPants
- MySQL Utilities
- MySQL Router
- PRM (w/Pacemaker)
- SeveralNines ClusterControl
- Orchestrator
- MySQL Sandbox
- MariaDB ColumnStore
- Numerous GUI tools: MySQL Workbench, phpMyAdmin, SQLYog, etc

Today, when to use MariaDB Server?

- MyRocks storage engine
- TokuDB storage engine
- MyISAM user? Segmented key caches will help
- CONNECT storage engine
- Threadpool
- PAM authentication
- GSSAPI authentication (Kerberos, Active Directory)
- Window functions
- PCRE Regular Expressions
- Optimistic parallel replication
- ANALYZE <statement>
- cracklib_password_check
- SQL Roles

Today, when to use Percona Server?

- When you need scalability, performance, flexibility, reliability, management, and diagnostic improvements
- MyRocks
- XtraDB
- Improved MEMORY engine
- TokuDB
- Enforcing storage engine, utility user
- Encryption via the Keyring Vault plugin
- Column compression
- Contributions go upstream to Oracle quite often, so less to maintain, and easier to develop new features

For everyone else, today?

MySQL 5.7 is truly what you should be using

The very near future

- MySQL 8 brings roles, CTEs, window functions, transactional data dictionary, histograms, SDIs (instead of FRMs), etc.
- <http://mysqlserverteam.com/the-mysql-8-0-0-milestone-release-is-available/>
- Will MariaDB Server 10.3 aim to be compatible with MySQL 8? Or more compatible with Oracle?
- <https://jira.mariadb.org/browse/MDEV-10137>, <https://jira.mariadb.org/browse/MDEV-11070>
- <https://mariadb.com/kb/en/library/changes-improvements-in-mariadb-103/>
- System versioned tables is an extremely neat feature, semisync via Alibaba merged in, invisible columns, sequences, sql_mode=ORACLE understands a subset of PL/SQL, column compression (storage engine independent), SPIDER merged, PROXY protocol support, etc.

What should you use?

- Think about the innovation **today**
- Think about the features you need today (with tolerance for the roadmap & what comes tomorrow)
- Beware vendor lock-in
- Ensure you are well supported
- Know the difference between a branch (Percona Server for MySQL) and a fork (MariaDB Server)



#mysqldevroom #FOSDEM

Thank you!

Colin Charles

colin.charles@percona.com / byte@bytebot.net

<http://bytebot.net/blog> | [@bytebot](#) on twitter

slides: slideshare.net/bytebot