

The MySQL Ecosystem in 2023

Colin Charles, Consultant, Galera Cluster


colin.charles@galeracluster.com | byte@bytebot.net

<https://bytebot.net/blog> | @bytebot on Twitter

FOSDEM, Brussels, Belgium

5 February 2023

whoami



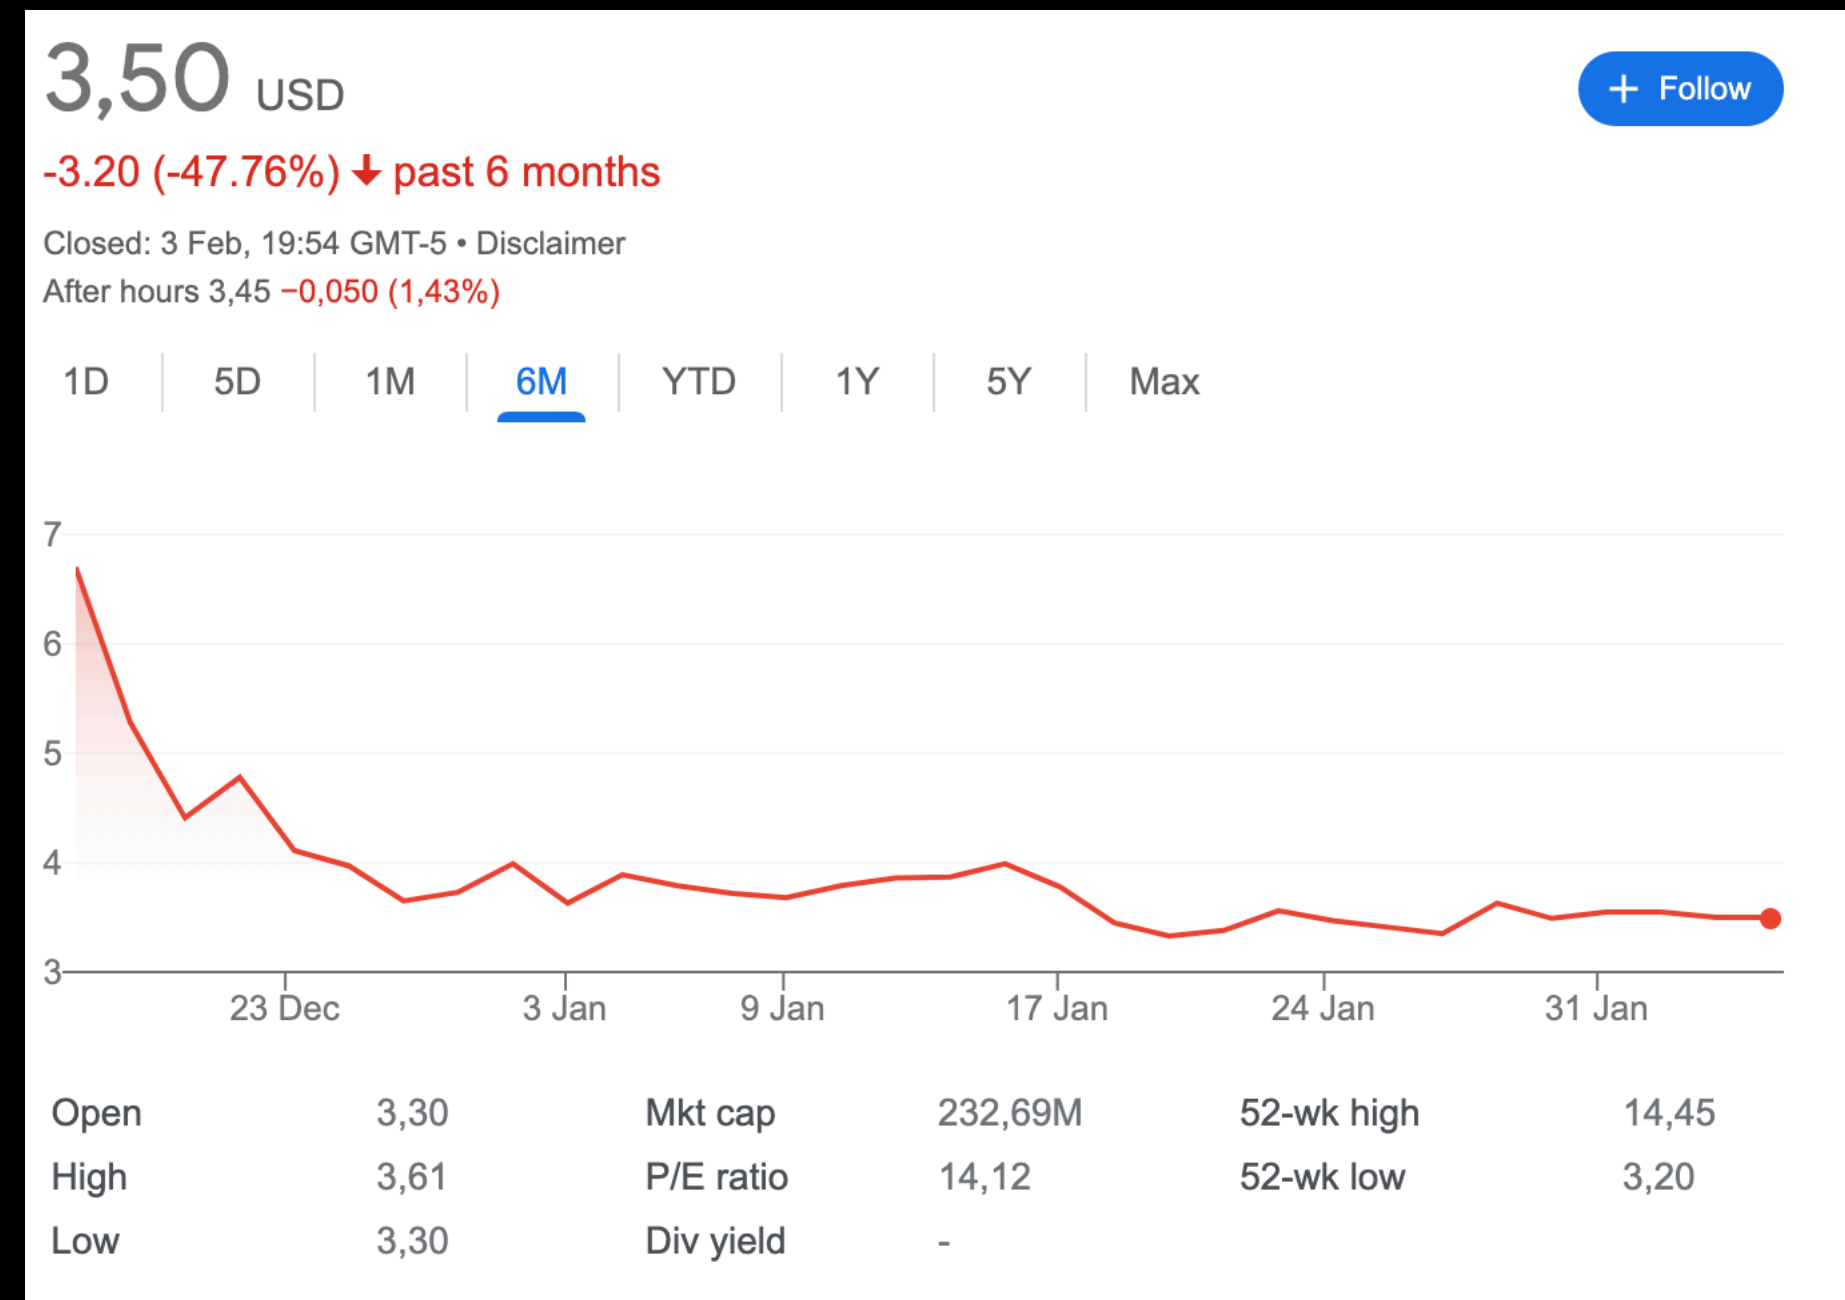
- Consultant at Codership, makers of Galera Cluster
- Active in the MySQL ecosystem: Founding team of MariaDB Server (2009-2016), early at MySQL AB (pre-Sun exit), Percona.
- Past lives include Fedora Project (FESCO), OpenOffice.org
- MySQL Community Contributor of the Year Award winner 2014

<ad> Codership </ad>

- Codership are the original makers and engineers of Galera Cluster, a multi-master, virtually synchronous replication solution for the MySQL ecosystem
- If you use Percona XtraDB Cluster (PXC) or MariaDB Galera Cluster, you directly benefit from the work done by the team at Codership
- Remember that beyond engineering, you may also purchase **24/7 support, training, consulting, Galera Cluster Enterprise Edition (EE)** and a whole lot more!
- Codership sponsored my travel to FOSDEM 2023

Fast forward, what changed in 3 years?

- MariaDB Corporation is now public as MariaDB plc, trading as MRDB; rapid release model + LTS releases
- Percona has re-branded
- Oracle has made 14 GA releases of MySQL 8
- Amazon RDS uses semi-sync
- Facebook has Raft-based replication



A mature ecosystem

- MySQL: 28 years — May 1995
- Percona Server: 15 years — November 2008
- MariaDB Server: 13 years — February 2010

410 systems in ranking, February 2023

Rank	Rank			DBMS	Database Model	Score		
	Feb 2023	Jan 2023	Feb 2022			Feb 2023	Jan 2023	Feb 2022
1.	1.	1.	1.	Oracle	Relational, Multi-model	1247.52	+2.35	-9.31
2.	2.	2.	2.	MySQL	Relational, Multi-model	1195.45	-16.51	-19.23
3.	3.	3.	3.	Microsoft SQL Server	Relational, Multi-model	929.09	+9.70	-19.96
4.	4.	4.	4.	PostgreSQL	Relational, Multi-model	616.50	+1.65	+7.12
5.	5.	5.	5.	MongoDB	Document, Multi-model	452.77	-2.42	-35.88
6.	6.	6.	6.	Redis	Key-value, Multi-model	173.83	-3.72	-1.96
7.	7.	7.	7.	IBM Db2	Relational, Multi-model	142.97	-0.60	-19.91
8.	8.	8.	8.	Elasticsearch	Search engine, Multi-model	138.60	-2.56	-23.70
9.	10.	10.	10.	SQLite	Relational	132.67	+1.17	+4.30
10.	9.	9.	9.	Microsoft Access	Relational	131.03	-2.33	-0.23
11.	12.	11.	11.	Cassandra	Wide column	116.22	-0.09	-7.76
12.	11.	15.	15.	Snowflake	Relational	115.65	-1.60	+32.47
13.	13.	12.	12.	MariaDB	Relational, Multi-model	96.81	-2.55	-10.30
14.	14.	13.	13.	Splunk	Search engine	87.08	-1.32	-3.73
15.	15.	17.	17.	Amazon DynamoDB	Multi-model	79.69	-1.87	-0.67
16.	16.	14.	14.	Microsoft Azure SQL Database	Relational, Multi-model	78.75	-1.62	-6.20
17.	17.	16.	16.	Hive	Relational	72.12	-2.22	-9.76
18.	18.	18.	18.	Teradata	Relational, Multi-model	63.03	-2.40	-5.54
19.	19.			Databricks	Multi-model	60.33	-0.49	
20.	20.	20.	20.	Neo4j	Graph	55.43	-0.41	-2.81

Branch vs. Fork



Release Matrix

MySQL 8 minor releases make all the difference

MariaDB	MySQL	Percona Server for MySQL
5.1: 1 Feb 2010	5.1: 14 Nov 2008	5.1.47-rel11.1: 25 Jun 2010
5.2: 10 Nov 2010		
5.3: 29 Feb 2012		
5.5: 11 Apr 2012	5.5: 3 Dec 2010	5.5.11-20.2: 18 Apr 2011
	5.6: 5 Feb 2013	5.6.14-62.0: 24 Oct 2013
10.0: 31 Mar 2014		
10.1: 17 Oct 2015		
	5.7: 21 Oct 2015	5.7.10-3: 23 Feb 2016
10.2: 23 May 2017		
	8.0: 19 Apr 2018	8.0.13-3: 21 Dec 2018
10.3: 25 May 2018		
10.4: 18 Jun 2019		
10.5: 24 Jun 2020		
10.6 LTS: 6 Jul 2021		
10.7: 9 Feb 2022		
10.8: 20 May 2022		
10.9: 22 Aug 2022		
10.10: 17 Nov 2022		
10.11:		
11.0:		

Open Source Community

- MariaDB Server: takes external contributors/committers, participates in Google Summer of Code to nurture new developers
- MySQL: takes external contributors, typically after signing a CLA; commits not welcome
- Percona Server: bug reports are welcome, commits aren't
- Contributor Agreements: Oracle Contributor Agreement (OCA), MariaDB Contributor Agreement (MCA), and the BSD New

MySQL 5.7

<http://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

MySQL 8.0

<https://mysqlserverteam.com/the-complete-list-of-new-features-in-mysql-8-0/>

- Recursive Common Table Expressions (CTEs)
- Window functions
- Instant ADD/RENAME COLUMN
- SET PERSIST
- Invisible indexes
- New REGEXP library + functions
- JSON support, schema validation
- caching_sha2_password
- Resource groups
- GIS support for Spatial Reference Systems
- UTF-8 (utf8mb4) as default character set (👍)
- Histograms
- EXPLAIN ANALYZE
- InnoDB improvements
- Transactional data dictionary
- Binary log compression
- Group replication
- SQL Roles
- MySQL Shell
- X Protocol

MariaDB Feature Highlights

- DML only Flashback - rollback instances/databases/tables to an older snapshot
- ed25519 password plugin
- simple_password_check, cracklib_password_check
- Accounts, passwords, and global privileges are stored in mysql.global_priv
- Progress reporting
- Table elimination (think anchor modelling — <https://mariadb.com/kb/en/what-is-table-elimination/>)
- Dynamic columns
- Column compression
- PROXY protocol support
- Instant DROP COLUMN
- LIMIT ROWS EXAMINED
- Usernames up to 80 characters, roles up to 128 characters
- MariaDB threadpool
- LIMIT ROWS EXAMINED
- Extended KILL syntax
- User statistics
- Oracle PL/SQL support
- Invisible columns
- System versioned tables, AS OF queries
- InnoDB AHI off by default

Storage Engines

- MyRocks: for write-intensive workloads
- SPIDER: for scalability and sharding*
- InnoDB: default for read/write operations (no longer Percona XtraDB since MariaDB 10.2); getting different from MySQL
- ColumnStore: analytical purposes
- Aria stores all system tables in 10.4+
- OQGRAPH: leaves algorithm
- PARTITION: updates to make SPIDER work better
- CONNECT: for ETL operations
- S3: store data in Amazon S3
- Cassandra: disabled in 10.5, removed in 10.6
- TokuDB: disabled in 10.5, removed in 10.6

Percona Server Feature Highlights

- Storage Engines: MyRocks, TokuDB (EOL 8.0.28-19)
- InnoDB full-text search improvements
- Extra diagnostic features (e.g. in INFORMATION_SCHEMA, user statistics)
- InnoDB improvements, e.g. parallel doublewrite buffer
- Column compression for VARCHAR/BLOB, JSON
- Compressed columns with dictionaries
- PAM authentication, audit logging
- Keyring in Hashicorp Vault
- A lot more, highlighted at https://www.percona.com/doc/percona-server/LATEST/feature_comparison.html but remember that comparison stops at 8.0.13 (and MySQL adds new features in latest releases... currently 8.0.32)

Governance

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

Mix & Match?

- Percona Server will work with MySQL since there is compatibility maintained
- MariaDB Server, your mileage will vary, but you can't have MySQL attached to it as a secondary
- Tools like MySQL Shell won't work on MariaDB Server (also, `mysql_ssl_rsa_setup`, etc.)
- ERROR 2059 (HY000): Authentication plugin 'caching_sha2_password' cannot be loaded: /usr/lib64/mysql/plugin/caching_sha2_password.so: cannot open shared object file: No such file or directory
- MariaDB MySQL 8 compat tracker: <https://jira.mariadb.org/browse/MDEV-28906>
- There can be strategies with separate deployments and proxies, but nowadays the path is clear — pick a server and stick with it

Master →	MariaDB-5.5	MariaDB-10.1	MariaDB-10.2	MariaDB-10.3	MariaDB-10.4	MySQL-5.6	MySQL-5.7	MySQL-8.0
Slave ↓								
MariaDB-5.5	✓	✗	✗	✗	✗	✗	✗	✗
MariaDB-10.1	✓	✓				✓		
MariaDB-10.2	✓	✓	✓			✓	✓	
MariaDB-10.3	✓	✓	✓	✓		✓	✓	
MariaDB-10.4	✓	✓	✓	✓	✓	✓	✓	
MySQL-5.6						*	*	*
MySQL-5.7						*	*	*
MySQL-8.0						*	*	*

- ✓: This combination is supported.
- ✗: This combination is **not** supported.
- *: MariaDB can't make any claims about MySQL-only combinations. Refer to the documentation for the specific MySQL version to determine supported combinations.

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

Clouds

Amazon RDS Multi-AZ deployments provide enhanced availability and durability, making them a natural fit for production database workloads. Deployment of Amazon RDS Multi-AZ with two readable standbys supports up to 2x faster transaction commit latencies than a Multi-AZ deployment with one standby instance. In this configuration, automated failovers typically take under 35 seconds. In addition, the two readable standbys can also serve read traffic without needing to attach additional read replicas.

- Amazon Web Services: RDS for MySQL (5.5, 5.6, 5.7, 8.0) and MariaDB Server (10.3-10.6), and of course Aurora MySQL
- Google Cloud SQL: MySQL (5.6, 5.7, 8.0 — default)
- Microsoft Azure: MySQL (5.7, 8.0) and MariaDB Server (10.2, 10.3)
- Alibaba Cloud: RDS MySQL (still supports 5.5->new upgrades), RDS MariaDB TX
- Oracle Cloud: MySQL Heatwave is exciting

High Availability Clustering Options

- MySQL 8
 - Group replication with MySQL InnoDB Cluster
 - InnoDB ClusterSet
- MariaDB Server
 - MariaDB Galera Cluster rolled into it
- Codership MySQL 8 with Galera Cluster
- Percona XtraDB Cluster (PXC) 8.0

Proxies

- MySQL integrates MySQL Router
- MariaDB Server recommends MariaDB MaxScale
- Percona recommends HAProxy, ProxySQL
- ProxySQL works with all the above

Ecosystem Tools

- Percona Toolkit
- Percona XtraBackup
- MariaBackup
- mydumper
- MHA (2018)
- vitess
- Orchestrator (2021)
- Signal18 Replication Manager
- dbdeployer

Commercial Ecosystem Tools

- Continuent Tungsten Clustering/Proxy/Replicator
- PlanetScale
- SeveralNines ClusterControl
- The continued need for Enterprise variants to move subscriptions
 - Oracle MySQL
 - MariaDB plc is a good example of pushing this around the open source MariaDB Server
 - Codership Galera Manager
- The cloud changes everything: Oracle MySQL Heatwave, MariaDB SkySQL

Distribution

- Linux distributions were golden, but even with them, you can see changes (e.g. Percona XtraDB Cluster (PXC) not being updated in Ubuntu LTS releases)
 - MariaDB Server was a choice during drop-in compatibility, but now is shipped alongside latest MySQL
- Docker
- Kubernetes
- OpenStack

Money Sloshing Around The Ecosystem

- Altinity (ClickHouse) raised \$4M (Accel lead)
- PlanetScale (vitess) raised \$105M in Series C (was: \$25M) (Kleiner Perkins, A16Z)
- PingCap (TiDB) raised \$341.6M (was: \$71.6M)
- MariaDB is public
- Cloud providers have huge revenue selling MySQL and MariaDB Server

Exciting stuff

- Go driver community contributions — <https://github.blog/2020-05-20-three-bugs-in-the-go-mysql-driver/>
- Facebook FlexiRaft — <https://www.cidrdb.org/cidr2023/papers/p83-yadav.pdf> and code: <https://github.com/facebook/mysql-5.6/search?q=raft>
 - yes, they have an 8.0 tree <https://github.com/facebook/mysql-8.0>
- A new, non-censored planet (thanks JF): <https://planet.oursqlcommunity.org/>

Providers galore

- There are service providers for the entire ecosystem
- People are going to ask for information about the software supply chain (even though the license is AS IS, there are expectations when companies exist)
- Trust the small player, they lack G&A overhead, they're unlikely to sell you something you don't need




At the [@EclipseFdn](#) we are seeing a troubling trend of information requests about our projects that ask things like:

- country of origin of the project
- country of origin of the contributors
- results from security audits, scans, and reviews
- SBOMs

These requests come from consumers of open source code that they've downloaded under open source licenses. For free. I am guessing that these are requests that they send to every supplier of code in their supply chain, regardless of license.

It really does feel like there is going to be relentless pressure to turn open source projects and foundations into "suppliers". I am pretty sure that the unintended consequences of that will not be positive.



I should add that of all of the info requests above, the one that bothers me the most is "country of origin of contributors". We don't track that. I don't know of any open source projects that track that. I consider it self-evident that starting to track that is a horrible idea

Several people have said variants of "just charge them for the information." But do not underestimate the effort required for a nonprofit to establish a new fee-for-service model. Worse, becoming a supplier changes the legal relationship



Also, having that information means investing the effort required to start systematically collecting the data, doing the audits & scans, and producing SBOMs for all our projects. Not a small undertaking or decisions to be taken lightly.

Conclusion

- Make yourself heard: Bugs databases are awesome
- 28 years, tools come and go
- MySQL and the ecosystem is not going anywhere :)
- Immerse yourself!
- Worrying though: new software packages, are they PostgreSQL or SQLite first, and MySQL later (or never)?



Thank You!

Colin Charles

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

<http://bytebot.net/blog> | @bytebot on twitter