


# 2017 Year in Review for open source database software

Colin Charles, Chief Evangelist, Percona Inc  
[colin.charles@percona.com](mailto:colin.charles@percona.com) / [byte@bytebot.net](mailto:byte@bytebot.net)  
<http://bytebot.net/blog/> | @bytebot on Twitter

Percona Webminar  
13 December 2017

# Open Source in the Organisation


- Open source software is *everywhere*
- From a database standpoint, db-engines reports:



The above chart shows the historical trend of the popularity of open source and commercial database management systems.

# Depends on the database model

Popularity broken down by database model, December 2017



© 2017, DB-Engines.com

This chart demonstrates that the popularity of open source tools heavily depends on their database model.

# What do the analysts say?

- Gartner
  - Through 2020, relational technology will continue to be used for at least 70% of new applications and projects.
  - By 2019, the separation of storage and compute designed for cloud DBMS architectures becomes the dominant dbPaaS model and begins to appear on-premises, also.

# In 2016...

Open source database management systems (OSDBMSs) have matured into viable alternatives to proprietary, commercial solutions. It comes as no surprise that OSDBMSs are not-so-subtly pushing proprietary solutions to the periphery. According to Gartner's *The State of Open Source RDBMS, 2015*, the OSDBMS market is worth \$562 million and has grown 31% year over year since 2013. By contrast, the overall market for database management solutions, valued by IDC to be over \$50 billion by 2018, grew by only 5.4% year over year during that same period. In terms of market evolution, Gartner writes: "By 2018, more than 70% of new in-house applications will be developed on an OSDBMS, and 50% of existing relational DBMS instances will have been converted or will be in process." *Id.* In other words, the entire definition of the database management system market will shift from commercial, closed-source DBMSs to OSDBMSs – a paradigm shift with decades of future staying power.

# Thinking about open source software adoption

- Quality (code, community)
- Licensing
- Market uptake
- Interoperability
- The “hidden” cost of ownership

# Why is open source good for your organisation

- Shift developers to high-value work
- Given enough eyeballs, all bugs are shallow
- Lower TCO
- Modern software development practices
- Upstream improvements, reduce duplication of effort
- Attract talent!

# Viewing source code $\neq$ open source

- Read the Open Source Definition: <https://opensource.org/osd>
- The Four Essential Freedoms: <https://www.gnu.org/philosophy/free-sw.en.html>



# Picking software

- Who's using it?
- Does it meet needs?
- Do others talk about this?
- Who funds the project? (beware *fauxdations*)
- Community involvement? Users? Developers?
- Project communication methods

# What about the cloud?

- Amazon AWS RDS
  - Aurora
- Microsoft Azure
  - CosmosDB
- Google Cloud
  - Spanner
- Rackspace
- Alibaba Cloud
- Oracle Cloud

# Security & Distribution

- Follow the patch updates - MySQL has critical patch updates
- You may get “booted out” of a distribution if you don’t provide security updates, e.g. Elastic, MySQL
- Does distribution matter?
- Security - privacy of data - regulation driven

# Star companies open sourcing code in our space

- Github — Orchestrator, gh-ost
- Facebook — MongoRocks, MyRocks
- Alibaba Cloud — Flashback, and more in MariaDB Server
- Tencent Cloud — SPIDER enhancements

# Servers

- MySQL 8.0 release candidate
- MariaDB Server 10.2 GA - May 2017
- MySQL & Percona Server 5.7 continue getting updates, with new features like mysqlsh, X Protocol, etc.
- MariaDB Server 10.3 Alpha

# Tools

- Orchestrator
- Replication Manager
- ProxySQL
- Tungsten Replicator
- Percona Monitoring & Management (PMM)
- Vitess

# Analytics

- MariaDB Server 10.2 (Window functions, Common Table Expressions)
- MariaDB ColumnStore
- ClickHouse

# R.I.P.

- RethinkDB
  - 2009, raised \$12.2m, ended up at Stripe + CNCF
- Basho
  - 2008, raised \$61.2m, ended up at Bet365 + Riak all open



# Read This Week in Data

- <https://www.percona.com/blog/author/colin-charles/>

# Looking back at 2017

- MongoDB IPOs
- MariaDB Corporation raises a lot, now \$98.2m
- Still money going into the database world (e.g. BlazingDB)

# Looking at 2018

- Oracle compatibility (incl. PL/SQL) in MariaDB Server 10.3+
- [confluent.io](http://confluent.io)
- AntidoteDB - <http://syncfree.github.io/antidote/>
- more manageability in MySQL (auto-sharding, auto-scaling)
- MongoDB deprecates mmapv1
- More focus on efficiency (e.g. RocksDB storage efficiency vs InnoDB compression)

# Explosion in Time Series Databases

- Theme for Percona Live Santa Clara 2017
- Interesting business models, including InfluxDB
- Submit for Percona Live Santa Clara 2018! Till December 22 2017
  - <https://www.percona.com/live/18/>

# Closing thoughts

- It takes 7-10 years to make good database software (without a *pivot*)
- Beware marketing dollars
- Look at where the money is going, website redesigns, management changes, etc?
- There can be lock-in in open source if you don't choose carefully

# Thank You. Q&A?

colin.charles@percona.com / byte@bytebot.net  
@bytebot on Twitter | <http://www.bytebot.net/blog/>  
slides: [slideshare.net/bytebot](http://slideshare.net/bytebot)