

MariaDB 10.0 stable (GA)



MariaDB 10.0

MariaDB & MySQL community day
3rd April, 2014

Michael “Monty” Widenius
CTO of MariaDB Foundation

monty@mariadb.org

<http://mariadb.org/>



Questions addressed in talk



This talk is aimed for people considering to move to MariaDB 10.0 from MySQL 5.5 or MariaDB 5.5.

The talk aims to address these points:

- 1) Why MariaDB
- 2) What is MariaDB
- 3) Features of the different MariaDB releases
- 4) Some benchmarks
- 5) The role and impact of the MariaDB foundation
- 6) Reasons to switch to MariaDB

Maria



Why Maria...

Why MariaDB was created



“Save the People, Save the Product”

- To keep the MySQL talent together
- To ensure that a free version of MySQL always exists
- To get one community developed and maintained branch
- Work with other MySQL forks/branches to share knowhow and code

MariaDB Foundation Overview



The Foundation is the new driver of the MariaDB project

Custodian of the code, Guardian of the community

Foundation can never to be controlled by a single entity or person

MariaDB Foundation Goals



That MariaDB be actively developed in the community and to:

- Increase adoption of MariaDB
- Ensure sustainable high-quality efforts to build, test and distribute MariaDB
- Ensure that community patches are reviewed and adopted
- Guarantee a community voice
- Keep MariaDB compatible with MySQL
- Maintain mariadb.org.

The MariaDB foundation now employes:

- Management team: CEO, CTO, lawyer
- 3 full time MariaDB developers (including the CTO)
- 1 documentation writer

MariaDB Foundation members



Many of the MariaDB 10.0 features has been developed together with the MariaDB foundation!

The foundation are very grateful to it's 2013 and 2014 members, who have made MariaDB 10.0 possible:

- Booking.com
- SkySQL (2 years!)
- Parallels (2 years!)
- Automattic
- Visma
- Zenimax

Additional sponsors can be found on mariadb.org



More members and sponsors are welcome!

If you care about the future of the MySQL ecosystem,
please contact us and
and ask how you can get involved!

Simon Phipps, CEO, simon@mariadb.org

Michael Widenius, CTO, monty@mariadb.org

MariaDB server is a branch of MySQL



- User level (data, API, replication..) compatible with MySQL
- Drop in replacement
- More plugins, more features, faster, better code quality.
- GPL-only server license.
- LGPL C and Java connectors.
- More open development
 - Source in public repository on launchpad
 - Active external contributors
 - All development plans public on mariadb.com, anyone can add to it or be part of the development of them

MariaDB server releases



- MariaDB 5.1 was released as stable in February 2010
- MariaDB 5.2 was released as stable in November 2010
- MariaDB 5.3 was released as stable in April 2012
- MariaDB 5.5 was released as stable in April 2012
- MariaDB 10.0 was released as alpha in November 2012
 - First beta (feature complete) released November 7, 2013
 - Now RC. Stable (GA) scheduled for end of March 2014.
- MariaDB-Galera (multi-master) was released as stable (GA) in February 2013 after a lot of testing.


Feedback plugin



- All recent MariaDB versions has the feedback plugin
- Enable by adding “plugin-load=feedback.so” and “enable-feedback” to the [mysqld] section in my.cnf.
- Feedback plugin will automatically send a report (basically SHOW STATUS)a few minutes after a startup and once a week
- This information is used to decide what features should be developed/expanded upon
- For more information see <https://mariadb.com/kb/en/feedback-plugin/>
- For statistics see http://mariadb.org/feedback_plugin/
- In February 2014 some 4200 people had it enabled

Feedback plugin

Some statistics



Major version of MariaDB in use.


Major new features in MariaDB 5.2



- **SphinxSE**: Text search within MariaDB
 - Built-in Sphinx client which allows MariaDB to talk to searchd, run search queries, and obtain search results.
- **Virtual columns**
 - Columns that are an expression and are calculated on retrieval.
- **Extended User Statistics**
 - Client, User, Index and Table statistics.
- **Segmented MyISAM key cache** (see separate slide)
- **Pluggable Authentication**
- Storage-engine-specific CREATE TABLE (Makes life easier for SE vendors)
- Fast 'copying to temp table' phase (speeds up GROUP BY and complex queries).
- **Group commit** & better recovery for the **Aria** engine.
 - Speeds up multi-user inserts.

Blue means developed by the community

MyISAM Segmented key cache



- Blue line is without segmented key cache.
- Solves one of the major read bottlenecks for MyISAM
- We see up to 250% performance gain depending on the amount of concurrent users.
- Fix applies to all MyISAM usage with many readers!

What's new in MariaDB 5.3



This is the biggest redesign of the MariaDB optimizer in 10 years and it will finally makes subqueries usable in MariaDB.

- Faster subqueries
 - Back porting and extending subquery optimization from MySQL 6.0
 - No materialization for many kinds of subqueries or VIEW's in the FROM clause. `SELECT * from (SELECT)`
 - Caching of subquery results

In applicable cases, you can get 10x – 100x speedups.

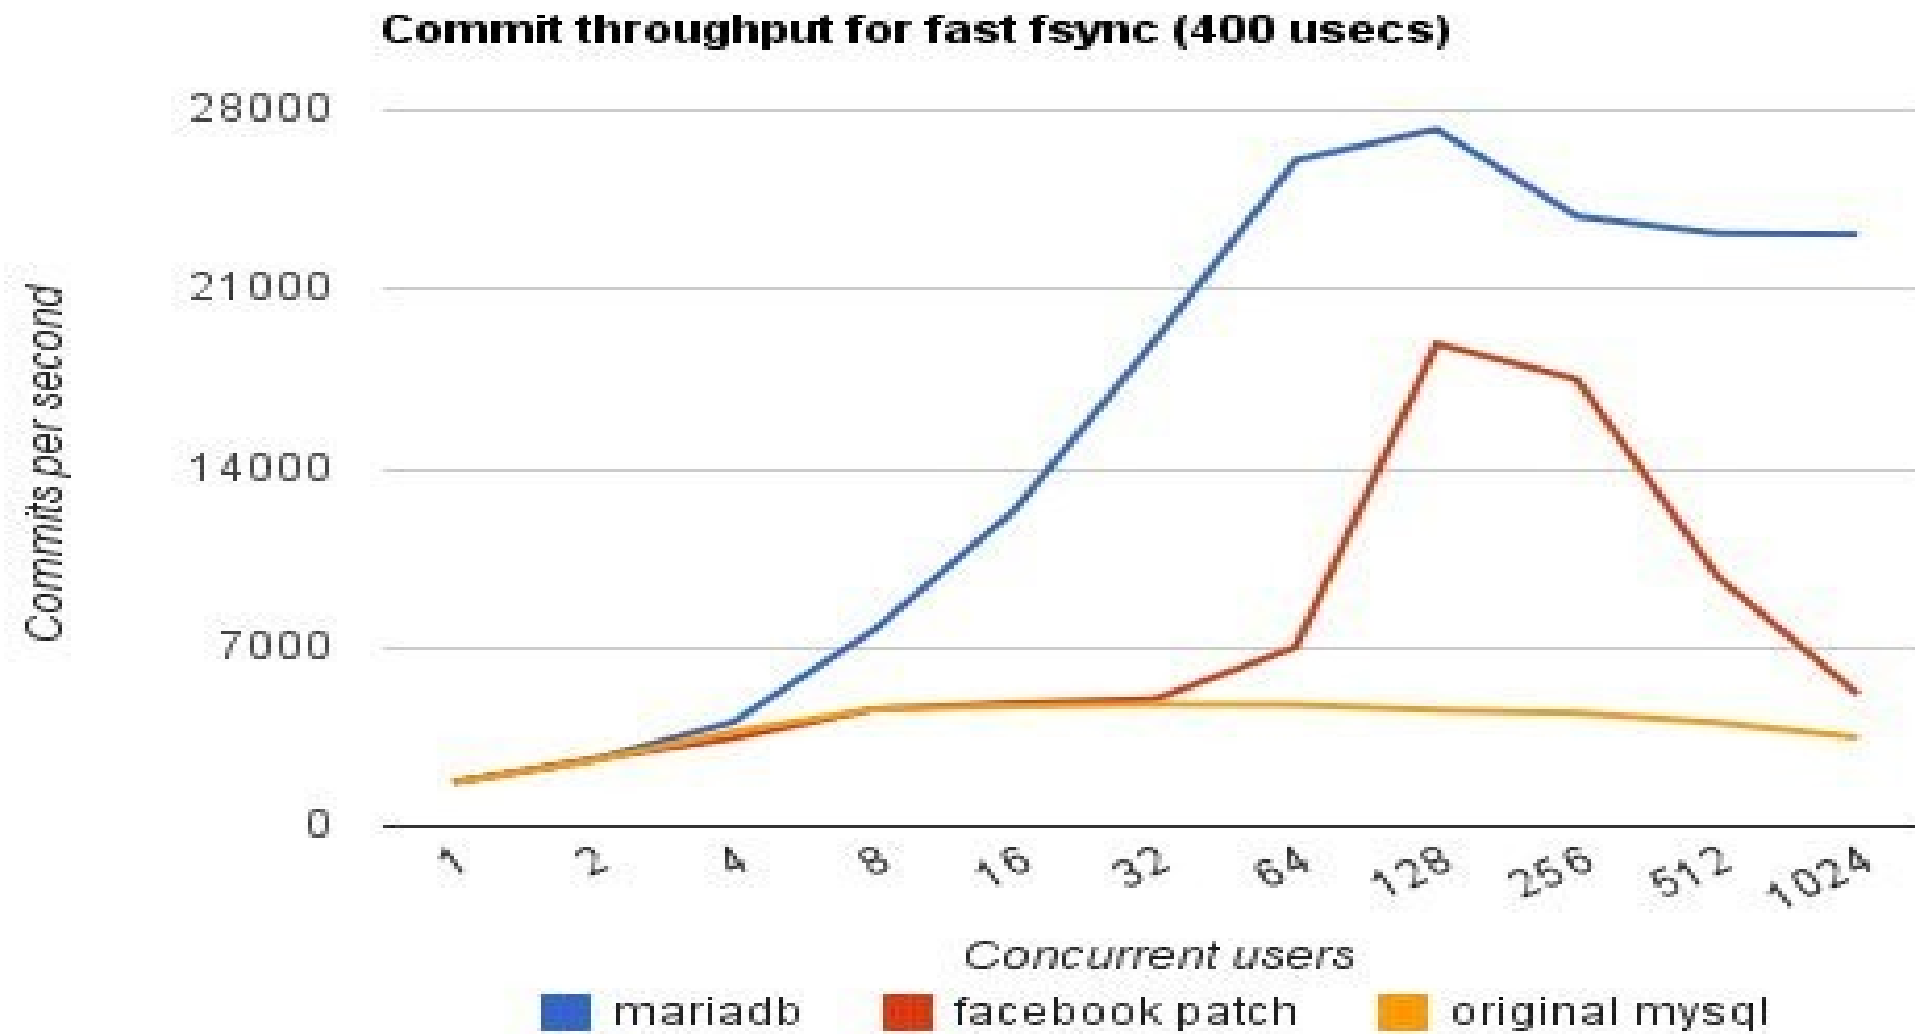
- Faster joins (of big tables) thanks to
 - Multi-Read-Range (MRR) access (better than in MySQL 5.6)
 - Batch key access (BKA)
 - Index condition pushdown
 - Classic Hash joins

What's new in MariaDB 5.3



- Full **microsecond** support. This includes `TIMESTAMP`, `TIME`, `DATETIME` types, `NOW()` and all `CAST` and `TIME` related functions, replication etc.
- **Group commit** between binary log and storage engines
 - `FASTER` and safer replication
- Progress report (with `PUSH` method) for `ALTER TABLE`, `LOAD DATA INFILE`, `REPAIR`, `OPTIMIZE & ANALYZE`.
- Precise GIS operations.
- Windows installer that includes Windows GUI-tool (HeidiSQL)
- For full list, see <http://kb.askmonty.org/en/what-is-mariadb-53>

Group commit, verified



Source: Marc Callaghan's facebook blog for a server with 400 microsecond fsync latency

Benchmark of new group commit Comparison with MySQL 5.6



- MySQL 5.6 also removes checkpoint-per-commit
- More complex thread scheduling, less grouping of commits
- Does not remove stall on LOCK_log, less grouping of commits

MariaDB 5.3 and NoSQL



The main reasons for using NoSQL are:

- Handling of unstructured data (not everything is table and fixed number of columns)
- Faster replication (usually with 'unconventional' shortcuts)
- The same way MySQL with its storage engine interface can handle both transactional and data warehousing, we are extending MariaDB to be a bridge between SQL and NoSQL.
- MariaDB 5.3 has now even better “NoSQL” support:
 - 50 % Faster HANDLER commands; HANDLER READ now also work with prepared statements.
 - HandlerSocket compiled in (Direct access to InnoDB)
 - Dynamic columns (each row can have different set of columns)

SQL doesn't solve all common problems



The (web) store problem:

All items needs: ID, Type, Price, Country, Manufacturer)

A T-Shirt has the following additional properties:

Size, color...

A computer has the following additional properties:

CPU, MHz, memory, Watt...

There is no easy way to store many different types into a relational database!

(It will not work by having one table/types as joins becomes impossible to manage).

Dynamic columns in MariaDB 5.3



- With dynamic columns all extra columns are stored in one or many packed blobs, maintained by the database.
- You can instantly add more columns, remove or query them for a row.
- You can access columns in the server or retrieve the full blob to the client and manipulate it there.
- You can use virtual columns to create indexes on some values.
 - True indexes for dynamic columns is planned for later.
- Implemented through functions to enable use by ODBC etc.
- First implementation uses an integer to access columns.
 - In MariaDB 10.0 columns can be accessed by name.

Dynamic columns in MariaDB 5.3



- Simple set of functions (available in server and client):
 - COLUMN_CREATE(column_nr, value, [column_nr,....])
 - COLUMN_ADD(blob, column_nr, value, [column_nr,...])
 - COLUMN_DELETE(blob, column_nr, column_nr...);
 - COLUMN_EXISTS(blob, column_nr);
 - COLUMN_LIST(blob, column_nr);
 - COLUMN_GET(blob, column_nr, type);

As a proof of concept we have in MariaDB 10.0 created storage engines for **Cassandra** and **LevelDB** where we use dynamic columns as a bridge.

We are working on doing the same with MongoDB.

New thread pool for 5.5

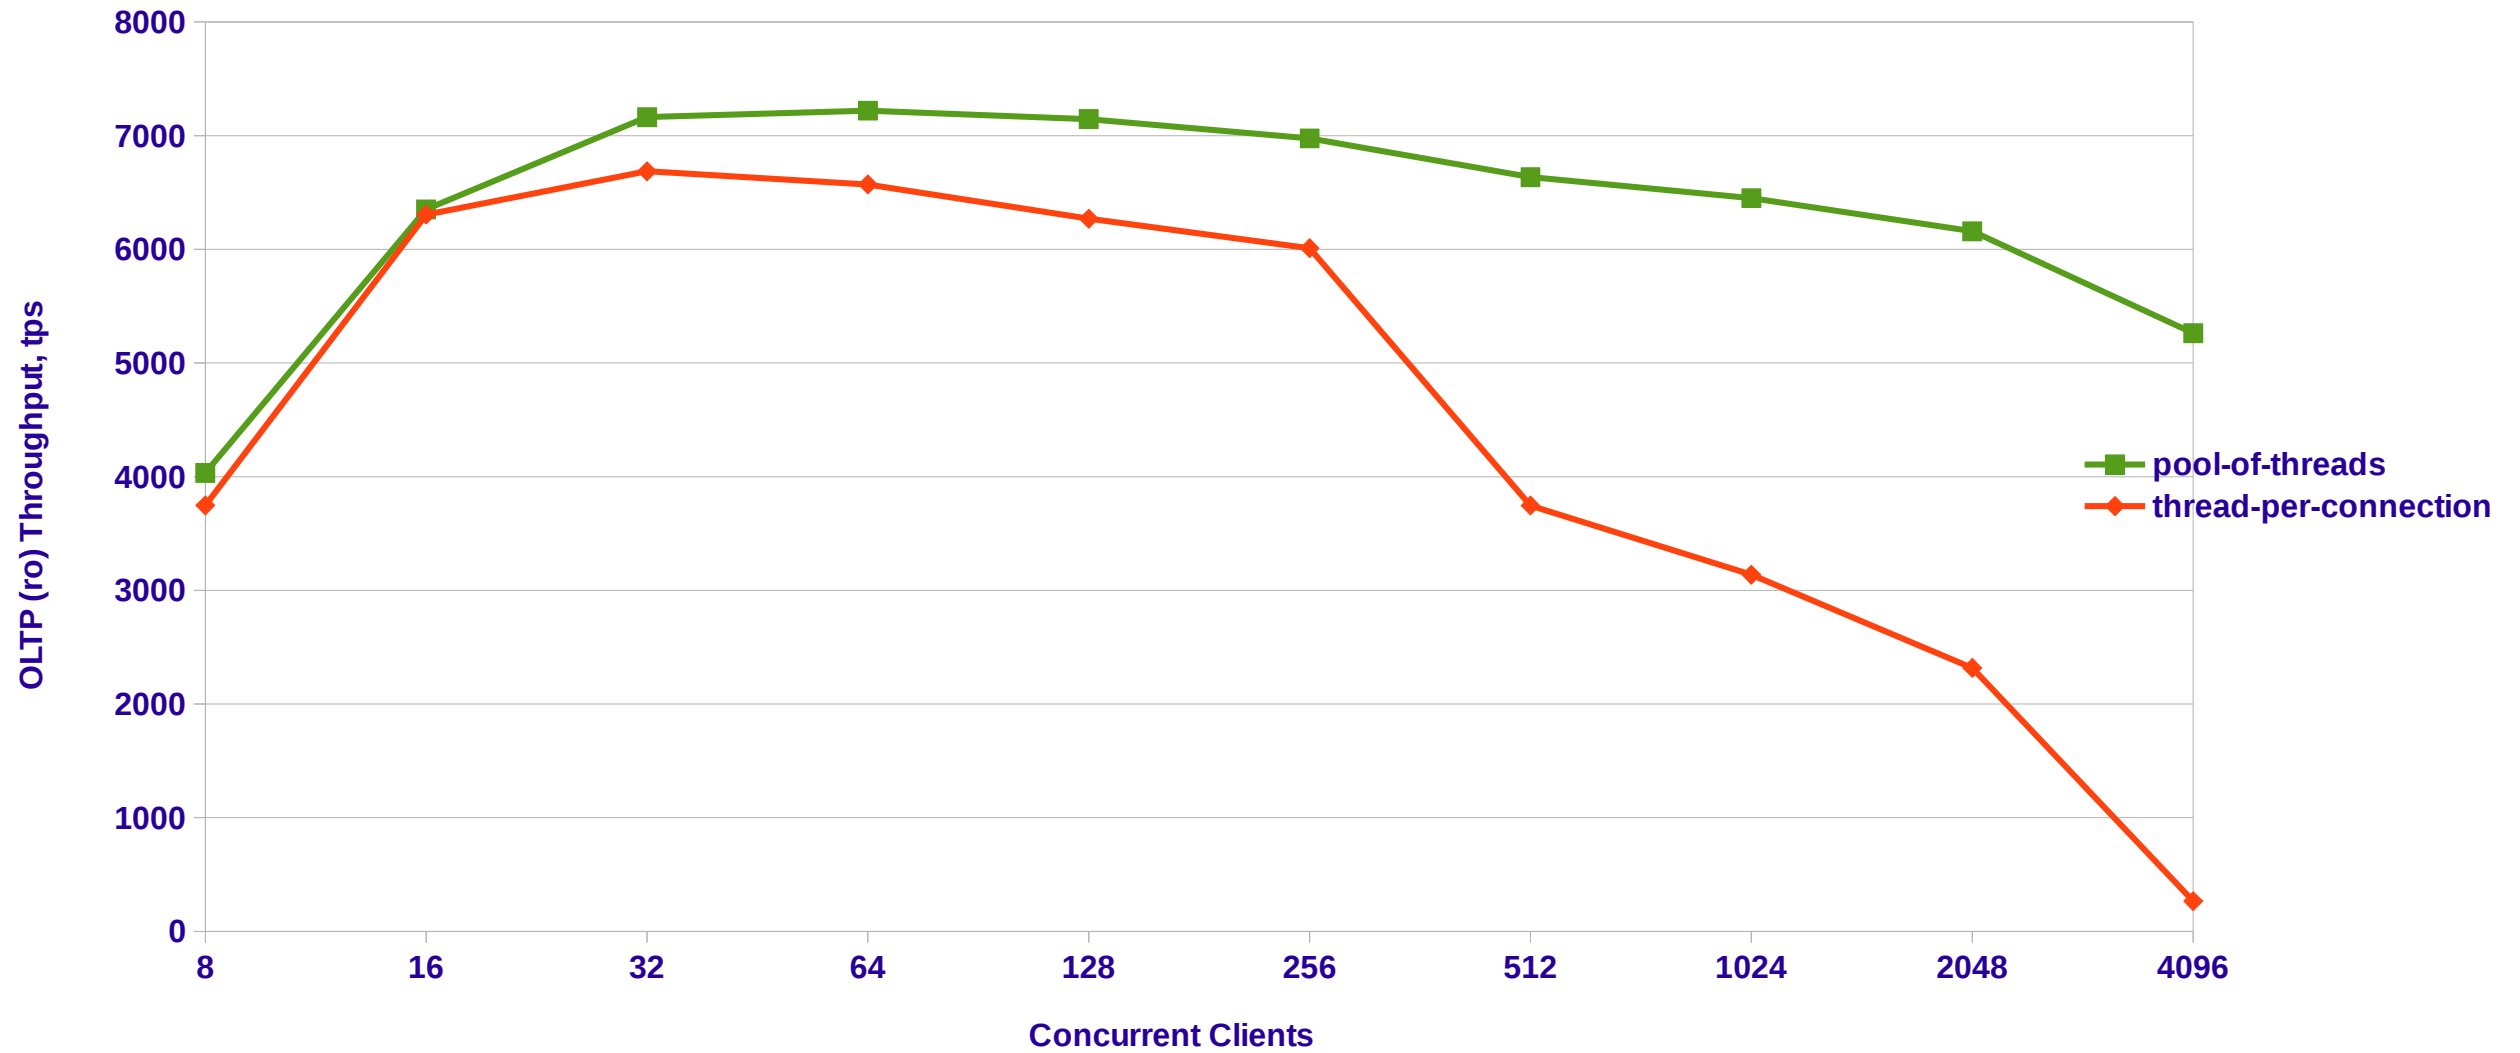


Thread pool exists in MariaDB 5.5 and the MySQL 5.5 Enterprise edition. It's **not** in the open source version of MySQL.

Thread pools solves a couple of problems:

- Allows you to limit the number of worker threads at your machines peak performance.
- More fair scheduler; Less query time distribution
- If too many queries, machine can run at 1% of peak performance

New thread pool for 5.5



Optimizations comparison



Feature	MariaDB 5.3	MySQL 5.5	MySQL 5.6
Index Condition Pushdown (ICP)	Yes		Yes
Disk-sweep Multi-range read (DS-MRR)	Yes		Yes
DS-MRR with Key-ordered retrieval	Yes		
Index_merge / Sort_intersection	Yes		
Cost-based choice of range vs. index_merge	Yes		
ORDER BY ... LIMIT <small_limit>	(In 10.0)		Yes
Use extended (hidden) primary keys for innodb/xtradb	5.5		
Batched key access (BKA)	Yes		Yes
Block hash join	Yes		
User-set memory limits on join buffers	Yes		
Apply early outer table ON conditions	Yes		
Null-rejecting conditions tested early for NULLs	Yes		

Optimizations comparison



Feature	MariaDB 5.3	MySQL 5.5	MySQL 5.6
Subquery: In-to-exists	Yes	Yes	Yes
Subquery: Semi-join	Yes		Yes
Subquery: Materialization	Yes		Yes
Subquery: NULL-aware Materialization	Yes		
Subquery: Cost choice of materialization vs. in-to-exists	Yes		
Subquery: Cache	Yes		
Subquery: Fast explain with subqueries	Yes		
Delayed materialization of derived tables / materialized views	Yes		Yes
Instant EXPLAIN for derived tables	Yes		Yes
Derived Table with Keys optimization	Yes		Yes
Fields of merge-able views and derived tables used	Yes		

Optimizations comparison



Feature	MariaDB 5.3/5.5	MySQL 5.5	MySQL 5.6
LIMIT ROWS EXAMINED rows_limit	5.5		
Systematic control of all optimizer strategies	Yes		Partial
Explain for DELETE, INSERT, REPLACE, and UPDATE			Yes
EXPLAIN in JSON format			Yes
More detailed and consistent EXPLAIN for subqueries	Yes		

Why MariaDB 10.0



- MariaDB 5.5 already have most (+ a lot more) of the optimizer features of MySQL 5.6
- MariaDB 5.5 is already a superset of MySQL 5.5. MySQL 5.6 has only a fraction of the MariaDB 5.5 new features.
- A full merge of MySQL 5.6 into MariaDB 5.6 is a one year project as a lot of the code has to be completely rewritten.
 - Features and usable code are removed, either intentionally or by mistake
 - New code is way to complex (you can do same thing much simpler)
 - Lots of new introduced bugs we have to get rid of.
 - It's clear that some of the new MySQL programmers doesn't understand the current code (see Kristian Nielsen's blog)
 - A lot of the new code is re-factoring we don't want to have.
- Better to do the merge in 2 steps into 10.0, and 10.1
- MariaDB 10.1 will have all important features of MySQL 5.6

Moving to MariaDB 10.0



MariaDB 10.0 beta should be binary compatible with MySQL 5.6 on:

- Data on disk
- Communication on the wire (ie, all MySQL connectors and applications following the MySQL specifications should work with MariaDB)
- Configure options (needs to be verified before gamma)
- SQL Syntax

Some configure options or optimizations may not be in MariaDB 10.0, but these should only affect performance. Some things will be faster, other things a bit slower in MariaDB than in MySQL.

We aim to have full compatibility for all important things that MySQL 5.6 supports in MariaDB 10.1.

MariaDB 10.0



MariaDB 10.0 is MariaDB 5.5 + some features from MySQL 5.6 + some new features

Features back ported from MySQL 5.6:

- All InnoDB changes (done)
- Performance schema changes (done)
- Read only transaction (significant InnoDB optimization) (done)
- Online ALTER TABLE (done)

Features from MySQL 5.6 that are reimplemented:

- Better error message (with system error string) (done)
- NOW() as default value for datetime (done)
- **Global transaction ID** for replication (done)
- **Parallel replication** (Base implementation exists)

New features in MariaDB 10.0



- SHOW EXPLAIN (see what other thread is doing)
- Multi source (one slave can have many masters)
- Faster ALTER TABLE with UNIQUE index
- DELETE ... RETURNING (See what you delete)
- Storage engine for **Cassandra**
- Storage engine **Connect**
- Full featured storage engine **TokuDB**
- Storage engine **Spider** (automatic sharding)
- Roles (group grant's)
- Much better GIS support (We can now run all Open Street Map queries)
- Engine independent statistics

For full list see <https://mariadb.com/kb/en/what-is-mariadb-100/>

Benefits of MariaDB 10 GTID (Global transaction id)



Global transaction ID allows you to easily promote a slave as a new master.

- Much simpler and more robust design than what is in MySQL 5.6
- GTID can be enabled or disabled independently and online for masters or slaves
 - In MySQL 5.6 one has to shut down all servers at once and enable GTID mode globally on both masters and slaves
- Slaves using GTID do not have to have binary logging enabled.
- Supports multiple replication domains (independent binlog streams)
 - Queries in different domains can be run in parallel on the slave.
- Supports multi-source replication

Benefits of MariaDB 10

Parallel replication



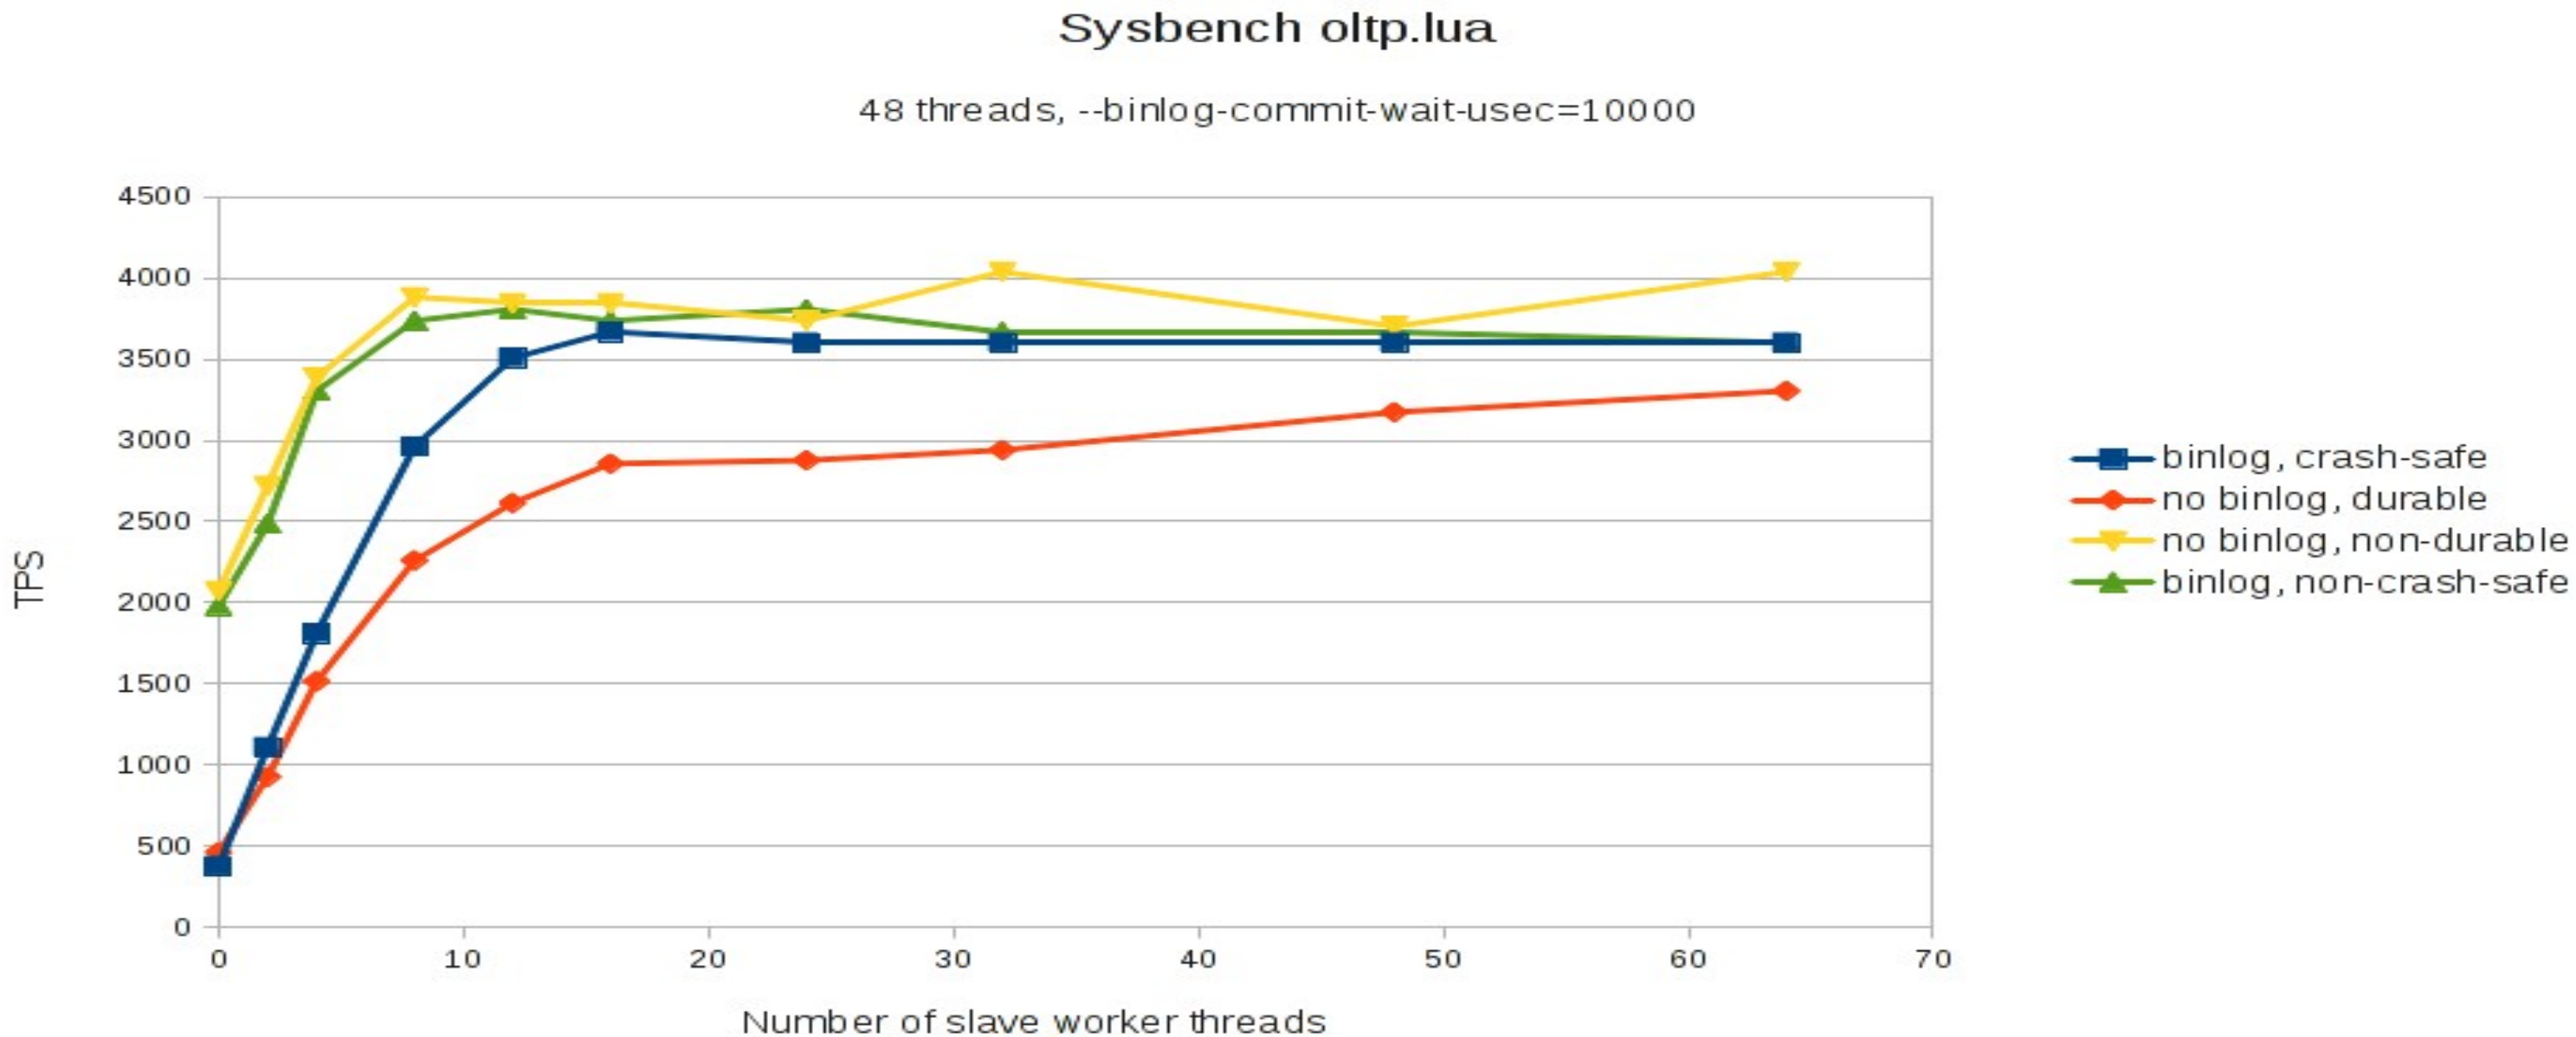
- Multi-source replication from different masters executed in parallel
- Queries from different domains are executed in parallel
- Queries that are run in parallel on the master are run in parallel on the slave (based on group commit).
 - Transactions modifying the same table can be updated in parallel on the slave!
- Supports both statement based and row based replication.

Conclusion:

- For the first time in MySQL/MariaDB history, the slave will be as fast as the master!

(By using dynamic domains, you could make the slave even faster than the master...)

Parallel slave benchmark 10x speedup with 12 threads



From: <http://kristiannielsen.livejournal.com/18435.html>


Crash safe: --log-slave-updates --sync-binlog=1 --innodb-flush-log-at-trx-commit=1

Parallel slave, no wait

Slave performance relative to master



slave performance relative to master



Benchmark: sysbench OLTP rw, 64 tables, data fits in memory

Master config: `trx_commit=1`, `sync_binlog=1`, SBR

Slave config: `trx_commit=1`

What are planned for MariaDB 10.x ?



This list of proposed features is still work in progress, and depends on the community and customers.

- GIS
 - OpenGIS compliance (largely done). MariaDB can already work with Open Streetmap data!
 - Deeper integration of GIS with optimizer
- More online operations
 - Analyze table
- Compatibility & usability
 - IPv6 type
 - Query logging and summary per query
 - Audit for specific users

What is planned for MariaDB 10.x ?



- Replication
 - Extend group commit to have on sync per group commit
 - ALTER TABLE's can be applied in parallel
- Statistics and monitoring
 - Better EXPLAIN
 - Persistent table statistics
 - Log all SQL errors
 - Progress indicator for SELECT

What is planned for MariaDB 10.x ?



- Optimizer
 - Implement UNION ALL without usage of a temporary table
 - Grace HASH join and Sort merge join (**Need sponsors**)
- Performance
 - Better multi CPU performance above 16 cores
 - More scalable query cache under higher concurrency
 - Faster VIEW (don't open & parse view for every query)
- Easy of use
 - VARCHAR and BLOB support for memory tables
 - Table functions

For full list, see <https://mariadb.com/kb/en/plans-for-10x>

New LGPL client libraries



- LGPL client libraries for C and Java
 - C is based on the LGPL library from MySQL 3.23
 - API compatible with latest MySQL client libraries.
 - Java is based on the drizzle driver.
- Works with MariaDB, Percona server, MySQL and drizzle
- Developed by Monty Program Ab and SkySQL.
- Announced and released 2012-10-29
- LGPL ODBC driver just released.

You can download these from <https://downloads.mariadb.org/>

MariaDB and TokuDB



TokuDB is now a native plugin for MariaDB.

This means that the official MariaDB binary is able to dynamically load the TokuDB storage engine directly (no patches needed for MariaDB) for all platforms that TokuDB support..

TokuDB is added to the MariaDB buildbot test suite to ensure that the combination is properly tested on all supported platforms.

Note that the TokuDB version for MariaDB 5.5 does not have all online features that the original TokuDB version from www.tokutek.com as we didn't want to do big changes in 5.5. The online features is in 10.0.

About TokuDB

- TokuDB uses Fractal Tree® indexing to improve insert and query speed, compression, replication performance, and online schema flexibility.
- TokuDB is created by Tokutek Inc. See www.tokutek.com for details.

Connect storage engine



MariaDB 10.0.2 includes the Connect storage engine by Olivier Bertrand.

You can read, write and update files in a lot of different storage formats:

- Various fixed and dynamic text formats
- .DBF (dBASE format)
- .CSV
- .INI
- .XML
- ODBC ; Table extracted from an application accessible with ODBC
-

With **MariaDB** and **CONNECT** you can for example join data in **INNODB**, **PostgreSQL**, **ORACLE** and **Cassandra** in the same query!

Free MariaDB/MySQL/SQL99 documentation at mariadb.com/kb



The screenshot shows the MariaDB Knowledge Base website. At the top, there is a navigation bar with links for 'KB Home', 'Downloads', 'Open Questions', 'MariaDB', and 'All Topics'. Below this, the main content area is divided into several sections:

- MariaDB**: A list of links including 'MariaDB Documentation', 'About MariaDB', 'About the MariaDB Foundation', 'MariaDB community', 'MariaDB versus MySQL - Features', 'MariaDB versus MySQL - Compatibility', 'MariaDB Case Studies', 'What is in the different MariaDB releases?', 'Recommended MariaDB / MySQL Books', 'Development', and 'More...'
- SQL-99 Complete, Really**: A list of links for SQL-99 topics, including '00 - Preface', '01 - Introduction', '02 - General Concepts', '03 - Numbers', '04 - Bit Strings', '05 - Binary Strings', '06 - Characters', '07 - Character Strings', '08 - Temporal Values', '09 - Boolean Values', and 'More...'
- Latest Articles**: A list of article titles such as 'sql99-banner', 'ST_BUFFER', 'ST_SYMDIFFERENCE', 'ST_INTERSECTION', 'ASWKT', 'ST_UNION', 'ST_OVERLAPS', 'ST_TOUCHES', 'ST_EQUALS', and 'ST_INTERSECTS'.
- Featured Content**: A list of links including 'MariaDB community', 'News and Information', 'About the AskMonty.org Knowledgebase', 'Contributing to the AskMonty Knowledgebase', 'Licensing FAQ', 'MariaDB Case Studies', 'Distributions Which Include MariaDB', and 'Recommended MariaDB / MySQL Books'.

On the right side of the page, there is a sidebar with the following sections:

- Welcome to the MariaDB Knowledgebase.**: A list of links for 'About', 'Editing Help', 'RSS Feeds', 'Articles [en]', 'Questions [en]', and 'Comments [en]'.
- Current Releases**: Two buttons for 'Download 10.0.4 Alpha' (released on 2013-08-16) and 'Download 5.5.32 Stable' (released on 2013-07-18).
- What's happening now?**: A section with statistics for the past week and past month, including 'Added/Changed Articles', 'Freenode #maria' activity, 'Launchpad Activity', and 'List Emails'.
- ...since launch:** A section showing 'Total Articles: 2950'.
- Social Media**: Buttons for 'Tweet', '+1', and 'Like'.

At the bottom of the page, there is a footer with links for 'Recent Changes', 'Contact Us', 'About', 'Legal', and copyright information: '©2010 - 2013 Monty Program Ab'.

The knowledgebase allows you to:

- Find answers to your problems
- Ask questions and get answers from others
- Add your own documentation or help with translations

There are a lot of others involved



- Most features in MariaDB 5.2 were contributed by the community!
- Many of the advanced features in MariaDB 5.3 are sponsored features
- In the mariadb.com knowledge base (free MariaDB and MySQL documentation) we have now 3000 (mostly English) articles

Statistics from the past month (March 2014):

- Added/Changed Articles: 262
- On Freenode #maria, 550 people wrote 10173 lines
- Launchpad Activity:
 - 25 active branches
 - 395 commits
- Hundreds of thousands of downloads of MariaDB. Probably >> 1M users
- Companies converting hundreds of machines to MariaDB in a few days
 - See <http://mariadb.com/kb/en/mariadb-case-studies>
- Google has employed people to work on MariaDB.

MariaDB popularity is increasing



- In December 2012
 - Wikipedia announced they are moving to MariaDB.
- In January-March 2013
 - DB at Mozilla blogged they have moved to MariaDB
 - Fedora voted 7-0 to make MariaDB the default MySQL database on Fedora.
 - OpenSuse 12.3 now includes MariaDB as default.
 - Slackware, Chakra Linux and Arch Linux have announced that MariaDB is default.
- In April 2013
 - Google is basing their new SQL offerings on MariaDB
 - FusionIO is showing benchmarks with MariaDB.
- June 2013
 - RedHat will include MariaDB in RedHat Enterprise.
- November 2013
 - MariaDB is in the Debian ftp queue to be included into Debian
- December 2013
 - RedHat Enterprise Linux 7 Beta has MariaDB as default

Reasons to switch to MariaDB today



- MariaDB has 30 man years of more development than MySQL (and the gap will continue growing).
- MariaDB is maintained by the people that originally created MySQL and has the best knowledge of the MySQL code.
- MariaDB is binary compatible (data and API) with MySQL, so its trivial to replace MySQL with MariaDB (minutes).
- Reasons to switch to MariaDB
 - Open source development: Anyone can be part of the development at all stages. Developer meetings are public!
 - Less risk as MariaDB will not remove features like MySQL is doing (thread pool, storage engines, older OS)
 - More features, including critical ones like microsecond, better statistics, dynamic columns and TokuDB, Cassandra and the Connect storage engines.
 - Faster queries thanks to XtraDB (InnoDB plugin fork from Percona), TokuDB, a much better optimizer, better replication and better code.

MariaDB 10.0 was made possible thanks to



- Oracle
 - Updated InnoDB, Updated Performance schema, online alter table
- SkySQL
 - Global Transaction Id, Parallel replication, Engine independent statistics, DELETE RETURNING, Cassandra engine, GIS support, QA
- MariaDB Foundation
 - Merge of MySQL 5.6, reviews of community patches, builds, speed optimizations, new character sets
- The MariaDB Community
 - CONNECT ENGINE by Olivier Bertrand
 - Spider by Kentoku Shiba
 - Multi-source replication, by Lixun Peng @ taobao
- Google, through Google summer of code
 - Roles by Vicențiu Ciorbaru
 - Move to new regexp library (PRCE) by Sudheera Palihakkara

Questions ?



For questions later, use the public MariaDB email list at maria-discuss@lists.launchpad.net or [#maria](#) on [Freenode](#).

For questions regarding the MariaDB foundation:
monty@mariadb.org