

What's new in MariaDB 5.5 and what's coming in 5.6

Colin Charles, Monty Program Ab
colin@montyprogram.com

<http://montyprogram.com> / <http://mariadb.org/>

<http://bytebot.net/blog> / @bytebot on Twitter

LinuxFest Northwest, Bellingham, WA

28 April 2012

whoami

- MariaDB dude at Monty Program Ab
- Formerly of MySQL AB/Sun Microsystems
- Past lives included FESCO (Fedora Project), OpenOffice.org

fb.com/MariaDB.dbms |
@mariadb

Monty Program

Aims

- Two years and four major releases - what is inside MariaDB?
- Discuss the broader MySQL ecosystem



MariaDB is currently...

- Community Developed
- Feature Enhanced
 - branch of MySQL with more
- Backward Compatible
 - i.e. not a fork

fb.com/MariaDB.dbms |
@mariadb

Monty Program

Who's behind it?

- MySQL (database) owned by MySQL AB (company)
- Monty Program is just a sponsor of MariaDB
- maria-captains: 52% Monty Program, 48% community

fb.com/MariaDB.dbms |
@mariadb

Monty Program

When did this begin?

- Sun buys MySQL AB in January 2008
- Oracle proposes acquisition of Sun Microsystems in April 2009
- Widenius decides that beyond the engine (February 2009), focus on MariaDB, a complete database solution

Why MariaDB?

- MySQL - an open source product or project?
- MariaDB is an open source project

Aims of MariaDB

- 100% compatible, drop-in replacement to MySQL
- Stable (bug-free) releases
- GPLv2



Compatibility with MySQL

- There is no NDB cluster storage engine
- Client libraries, client-server protocol, SQL dialect, replication master-slave all similar
- Data files are supported as long as its similar between versions
- Tools are similar (some additional tools for Aria, PBXT)
- XtraDB enabled by default; InnoDB included, but not enabled by default
- <http://kb.askmonty.org/v/mariadb-versus-mysql-compatibility>

fb.com/MariaDB.dbms |
@mariadb

Monty Program

Why history matters

Because current MariaDB 5.5 includes all changes in
MariaDB 5.1, 5.2, 5.3

What's in MariaDB 5.1

- Released February 2010
- Create external buildsystem using Buildbot + VMs for test +builds
- Release filled with storage engines
- *Croatian collations*
- Numerous bug fixes
- Test cases and coverage improvements
- Removal of mutexes
- Compiler warnings gone!

- *XtraDB, PBXT, FederatedX, (M)Aria*
fb.com/MariaDB.dbms |
[@mariadb](https://twitter.com/mariadb)

Monty Program

What's in MariaDB 5.1

- XtraDB
(ENGINE=InnoDB)
 - enhanced InnoDB designed to better scale on modern hardware; backward compatible, scale better on many cores, use memory more efficiently
- <http://www.percona.com/docs/wiki/percona-xtradb:start>

fb.com/MariaDB.dbms |
@mariadb

Monty Program

What's in MariaDB 5.1

- *Extended statistics for slow query log*

- based on *microslow* from Percona <http://kb.askmonty.org/v/slow-query-log-extended-statistics>

- *PROCESSLIST with microsecond precision*

- *TIME_MS in INFORMATION_SCHEMA.PROCESSLIST* http://kb.askmonty.org/v/time_ms-column-in-information_schemaprocesslist

```
MariaDB [(none)]> select id, time, time_ms, command, state from information_sche
ma.processlist, (select sleep(2)) t;
+-----+-----+-----+-----+-----+
| id | time | time_ms | command | state |
+-----+-----+-----+-----+-----+
| 2 | 2 | 2002.822 | Query | executing |
+-----+-----+-----+-----+-----+
1 row in set (2.01 sec)
```

What's in MariaDB 5.1

- Table elimination
 - Query highly normalised data, it is sometimes possible to resolve a query without accessing some tables the query refers to e.g. <http://kb.askmonty.org/v/what-is-table-elimination>
- Pool of threads (MySQL 6 backport - libevent based)
 - Use a limited set of threads to handle all queries vs one-thread-per-connection. Win if most queries are short running, and there are few table/row locks. <http://kb.askmonty.org/v/pool-of-threads>



What's in MariaDB 5.2

- Released November 2010
- Pluggable authentication
 - `mysql_native_password` (20 byte) and `mysql_old_password` (8 byte) included by default
 - **Now with a PAM authentication plugin!**
- <http://kb.askmonty.org/en/pam-authentication-plugin>
- *User Statistics (userstatsv2)*
 - Percona, Google, Weldon Whipple. Understand server activity better, identify database loads <http://kb.askmonty.org/v/user-statistics>
 - Client, user, index, table statistics

fb.com/MariaDB.dbms |
@mariadb

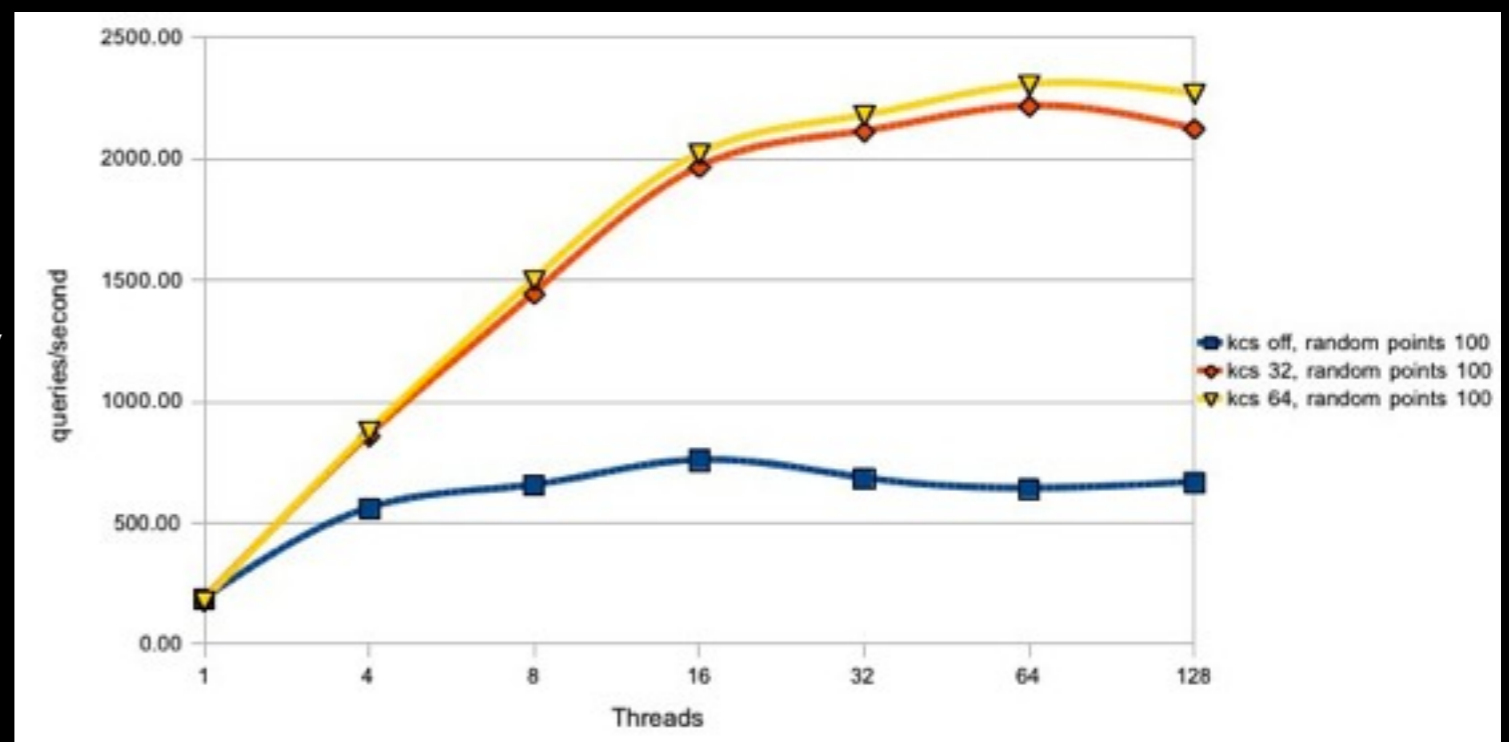
Monty Program

What's in MariaDB 5.2

- Segmented MyISAM keycache
- mitigates thread contention for key cache lock, with notable 250% performance improvements

[http://
kb.askmonty.org/
v/segmented-key-
cache](http://kb.askmonty.org/v/segmented-key-cache)

[fb.com/MariaDB.dbms](https://www.facebook.com/MariaDB.dbms) |
[@mariadb](https://twitter.com/mariadb)



Monty Program

MariaDB 5.2: SphinxSE

- `CREATE TABLE t1 (..) ENGINE=SPHINX CONNECTION="sphinx://localhost:9312/test";`
- `SELECT * from t1 WHERE query='test it;mode=any';`
- matching modes, limits, filters, ranges supported
fb.com/MariaDB.dbms | @mariadb
- monitor it - `SHOW ENGINE SPHINX STATUS;`
- can JOIN a SphinxSE search table and tables using other engines as well
- <https://kb.askmonty.org/v/about-sphinxse>

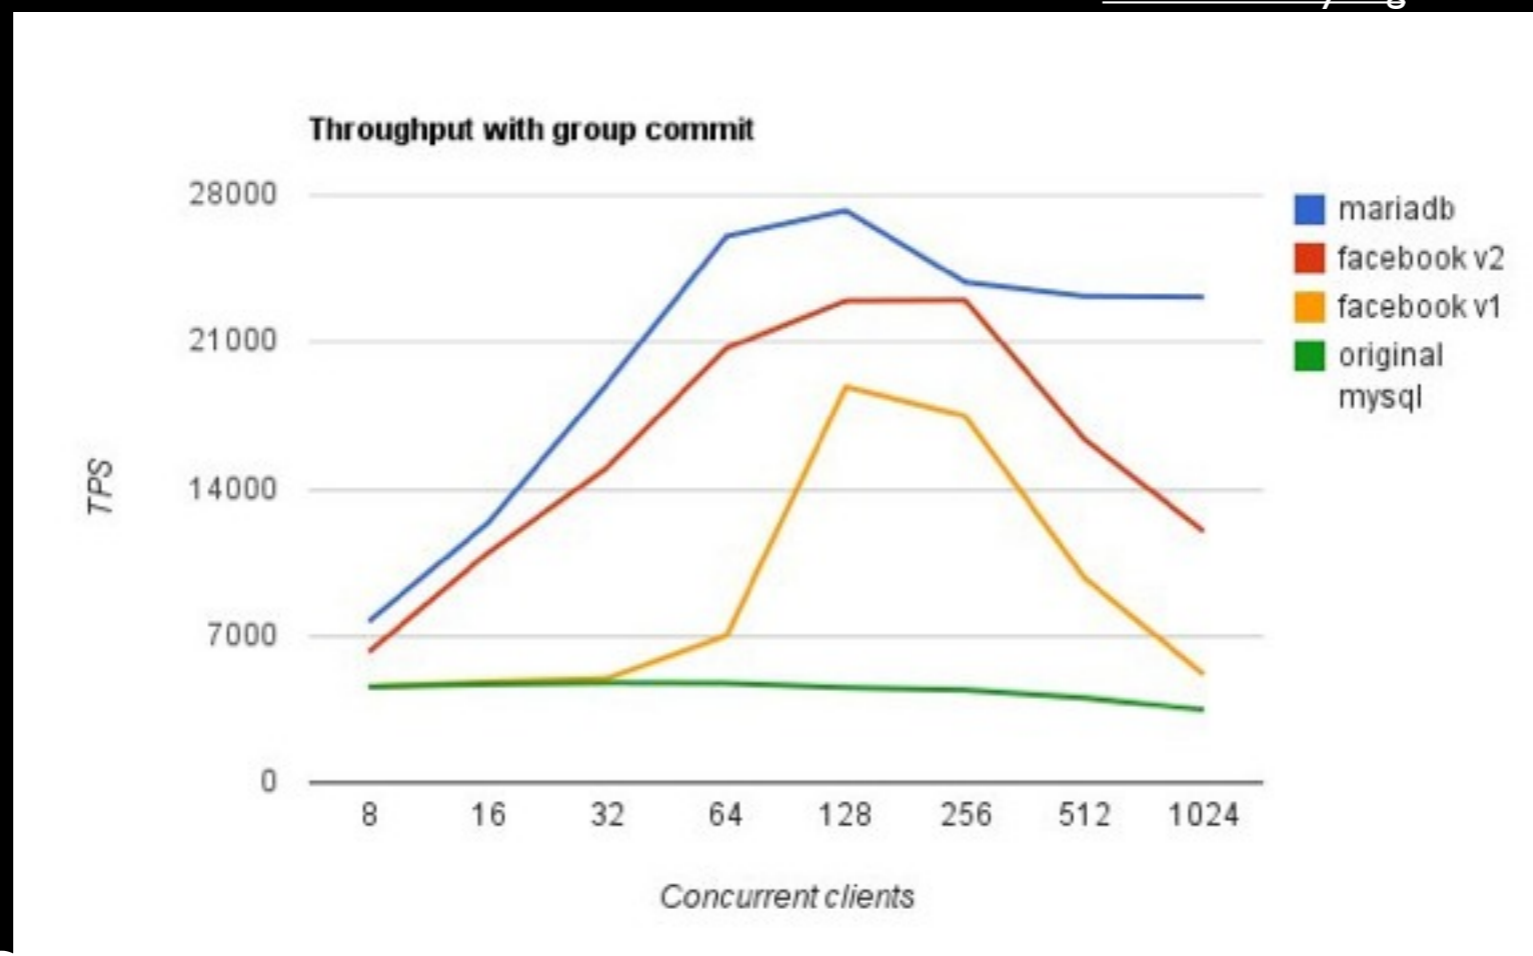
MariaDB 5.3: NoSQL

- HandlerSocket
 - direct access to InnoDB/XtraDB
 - no SQL, simple CRUD operations on tables
- Dynamic columns
 - different set of “virtual columns” for each row in your table
 - store different attributes for each item (see: e-commerce shops)



MariaDB 5.3: group commit

- Group commit in the binary log
- `sync_binlog=1,`
`innodb_flush_log_at_trx_commit=1`
- https://www.facebook.com/note.php?note_id=10150261692455933
- <http://kb.askmonty.org/en/group-commit-for-the-binary-log>



fb.com/mariadbadmins |
@mariadb

Monty Program



MariaDB 5.3: replication

- Checksums for binlog events (5.6 backport)
- START TRANSACTION WITH CONSISTENT SNAPSHOT works w/ binlog
- obtain binlog pos. w/o blocking other queries
- mysqldump --single-transaction --master-data => non-blocking backup! (new slave provisioning?) Works w/XtraDB
- row-based replication for tables w/no primary key

fb.com/MariaDB.dbms |
@mariadb

Monty Program

MariaDB 5.3

- Progress reporting for ALTER TABLE/LOAD DATA INFILE

```
MariaDB [test]> alter table my_mail engine=maria;  
Stage: 1 of 2 'copy to tmp table' 5.37% of stage done
```

- mytop included
 - Optimiser changes
 - @@optimizer_switch flags (default=off)
- fb.com/MariaDB.dbms |
@mariadb

- disk access: Index Condition Pushdown (ICP), Multi-Range Read (MRR)

- join: Classic Hash Join, Batched Key Access (ordered retrievals via MRR), and many many more optimizations

Monty Program



MariaDB 5.3:

Subqueries finally!

- DBT-3, 60M line item rows, 29GB XtraDB database
- optimizer_switch='semijoin=on' / materialization=on

Query: orders from customers with negative balance:

```
SELECT * FROM orders
WHERE o_custkey IN
  (SELECT c_custkey FROM customer
   WHERE c_acctbal < -500);
```

MariaDB 5.2 (any MySQL): 45 sec (slow)

id	select_type	table	type	key	ref	rows	Extra
1	PRIMARY	orders	index	i_o_custkey	NULL	1493631	Using where; Using index
2	SUBQUERY	customer	range	c_acctbal	NULL	10536	Using where; Using index

MariaDB 5.3: 0.43 sec (faster ~ 100x)

id	select_type	table	type	key	ref	rows	Extra
1	PRIMARY	customer	range	c_acctbal	NULL	10536	Using where; Using index
1	PRIMARY	orders	ref	i_o_custkey	dbt3sf1.customer.c_custkey	7	Using index

Query: find customers with top balance in their nations:

```
SELECT c_name, c_address
FROM customer
WHERE c_acctbal IN (SELECT max(c_acctbal)
                   FROM customer GROUP BY c_nationkey);
```

MariaDB 5.2 (any MySQL): > 1.5 hours (impossible)

id	select_type	table	type	possible_keys	key	key_len	ref	rows	Extra
1	PRIMARY	customer	ALL	NULL	NULL	NULL	NULL	149637	Using where
2	DEPENDENT SUBQUERY	customer	index	NULL	i_c_nationkey	5	NULL	3117	

MariaDB 5.3: 3.2 sec (faster ~ INF)

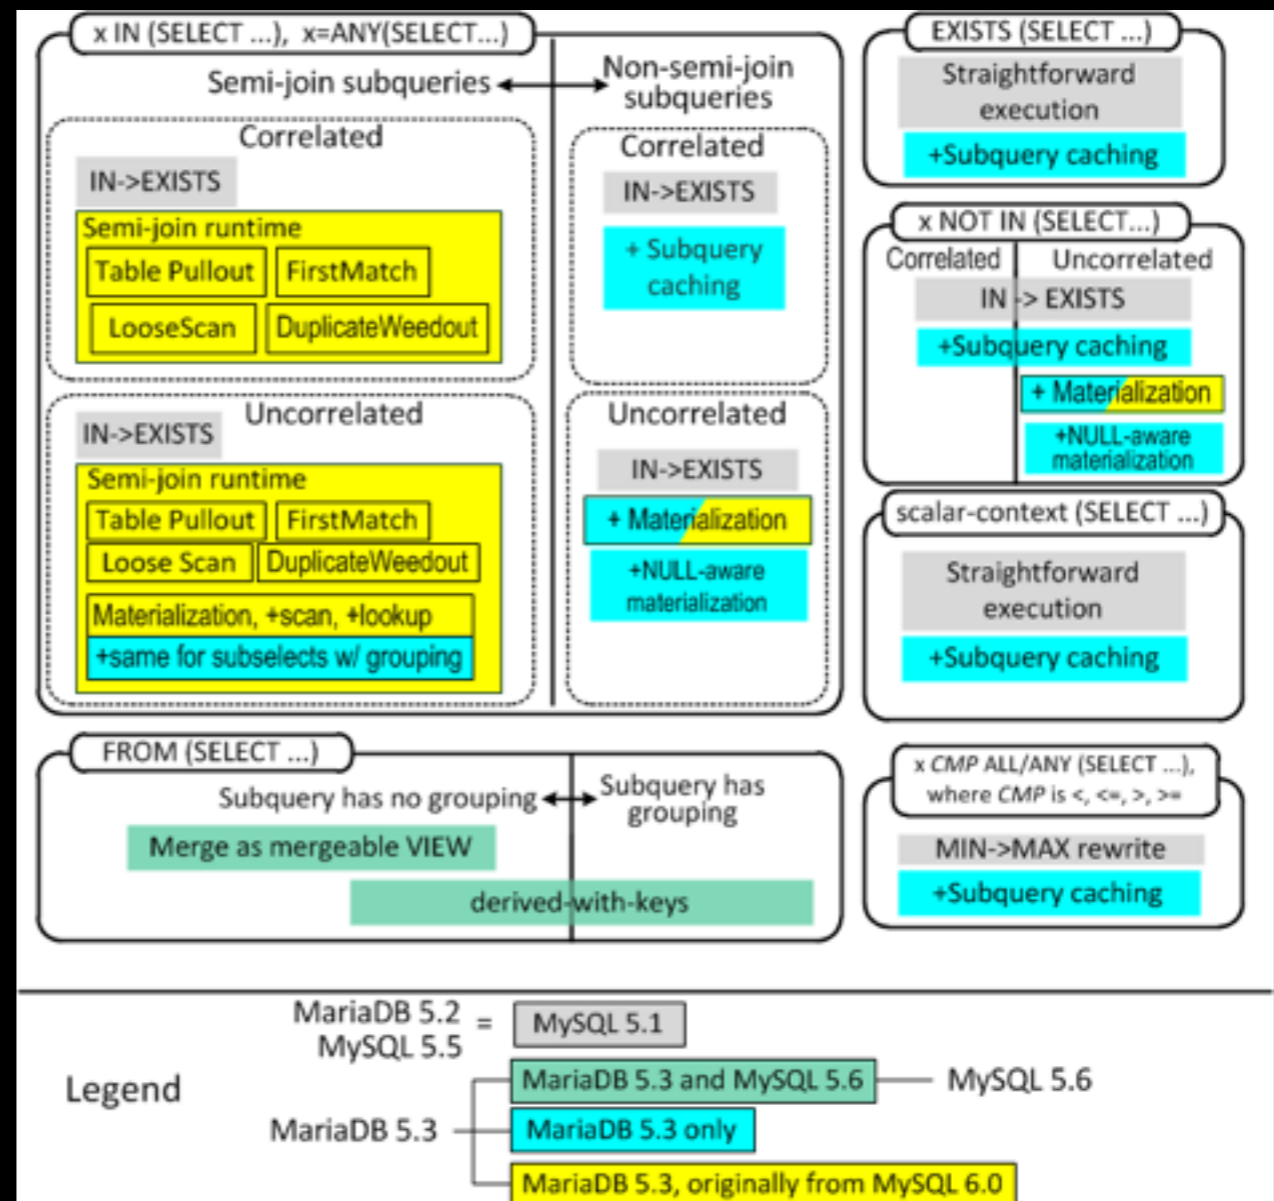
id	select_type	table	type	possible_keys	key	key_len	ref	rows	Extra
1	PRIMARY	<subquery2>	ALL	distinct_key	NULL	NULL	NULL	149637	Using where
1	PRIMARY	customer	ref	c_acctbal	c_acctbal	19	<subquery2>.max(c_acctbal)	1	
2	SUBQUERY	customer	index	NULL	i_c_nationkey	5	NULL	149637	

fb.com/MariaDB.dbms |
@mariadb

Monty Program


MariaDB 5.3: Subqueries

- Instant EXPLAIN
- <http://kb.askmonty.org/en/the-query-optimizer>



MariaDB 5.3: User Feedback Plugin

- New “phone home” plugin (works with Percona server, MySQL)
- disabled by default -- please consider enabling it (feedback=on)
- http://mariadb.org/feedback_plugin/



MariaDB 5.3: GIS

- MySQL has OpenGIS SFS (Simple feature access, SQL access method)
- Now, SQL with geometry types
- ST_ prefix
- Read more: <http://kb.askmonty.org/en/gis-features-in-533>



MariaDB 5.3: Aria storage engine

- Still a crash-safe MyISAM
- Has been included and improved upon since MariaDB 5.1
- Has group commit in MariaDB 5.2, to speed up multi-user inserts
- Is TRANSACTIONAL (crash safe, atomic), but doesn't do ROLLBACK yet (Aria 2.0)

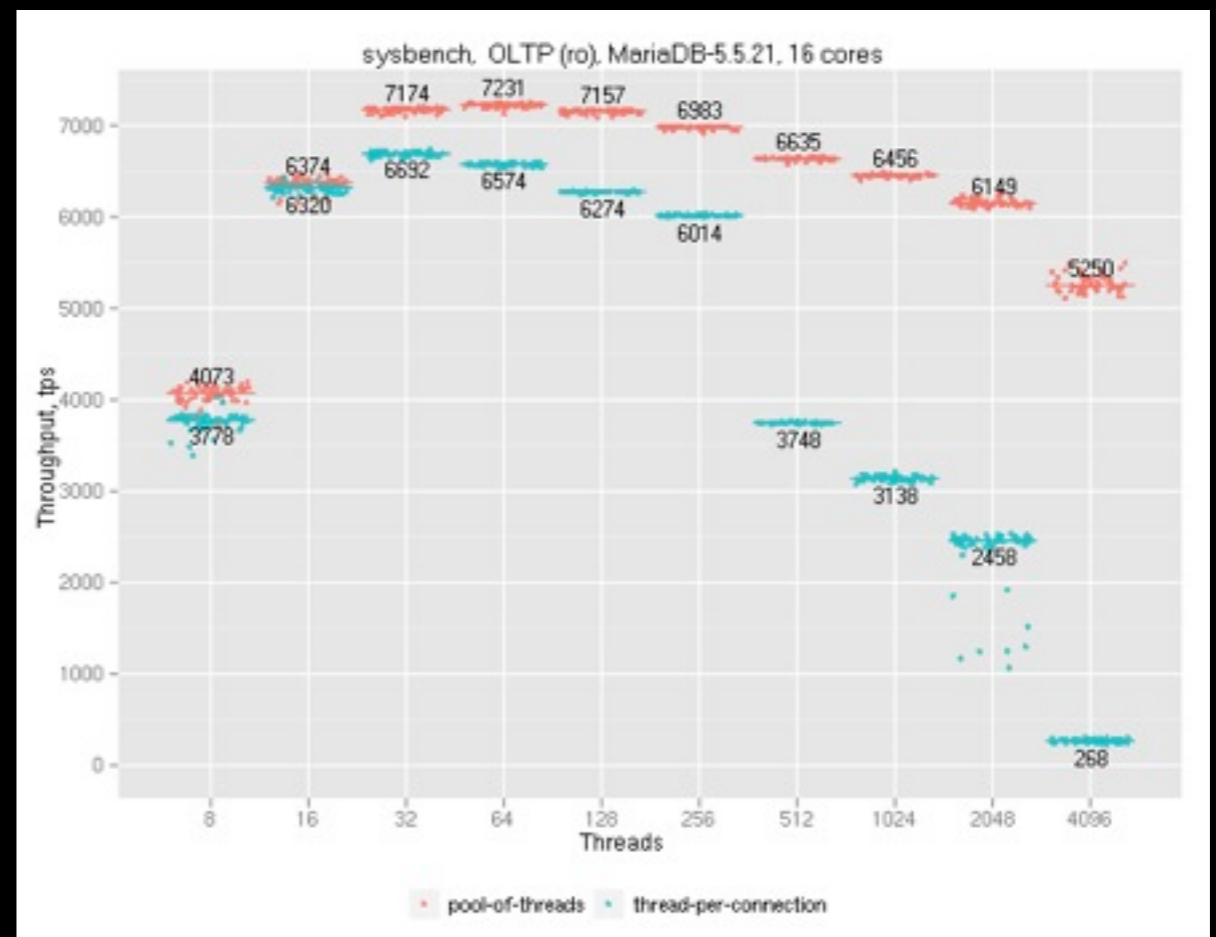
fb.com/MariaDB.dbms |
@mariadb

Monty Program



MariaDB 5.5: a new threadpool

- Modified from MariaDB 5.1, great for CPU bound loads, and short running queries
- No minimization of concurrent transactions!
- Windows (threadpool), Linux (epoll), Solaris (event ports), FreeBSD/OSX (kevent)
fb.com/MariaDB.dbms | @mariadb



More benchmarks: <http://kb.askmonty.org/en/threadpool-benchmarks>

Monty Program



MariaDB 5.5: better for DBAs

- Comes with a non-blocking client library
- start an operation, do other work in the same thread, operation processed in server, result travels back.
- multiple queries against a single server? (better utilise more CPUs?)
- queries against multiple servers? (eg. SHOW STATUS against all running servers)
- SELECT statements might examine too many rows and consume many resources. Terminate execution with LIMIT ROWS EXAMINED

fb.com/MariaDB.dbms |
@mariadb

Monty Program

MariaDB 5.5

- MariaDB 5.5 - **GA TODAY**
- <http://kb.askmonty.org/en/mariadb-5523-release-notes>



How do we support all these?

Version	Released	Supported till	Latest release	MySQL
5.1.42	Feb 2010	Feb 2015	5.1.62	5.1.62
5.2.3	Nov 2010	Nov 2015	5.2.12	5.1.62
5.3.5	Feb 2012	Feb 2017	5.3.6	5.1.62
5.5.23	Apr 2012	Apr 2017		5.5.23

fb.com/MariaDB.dbms |
@mariadb

Monty Program

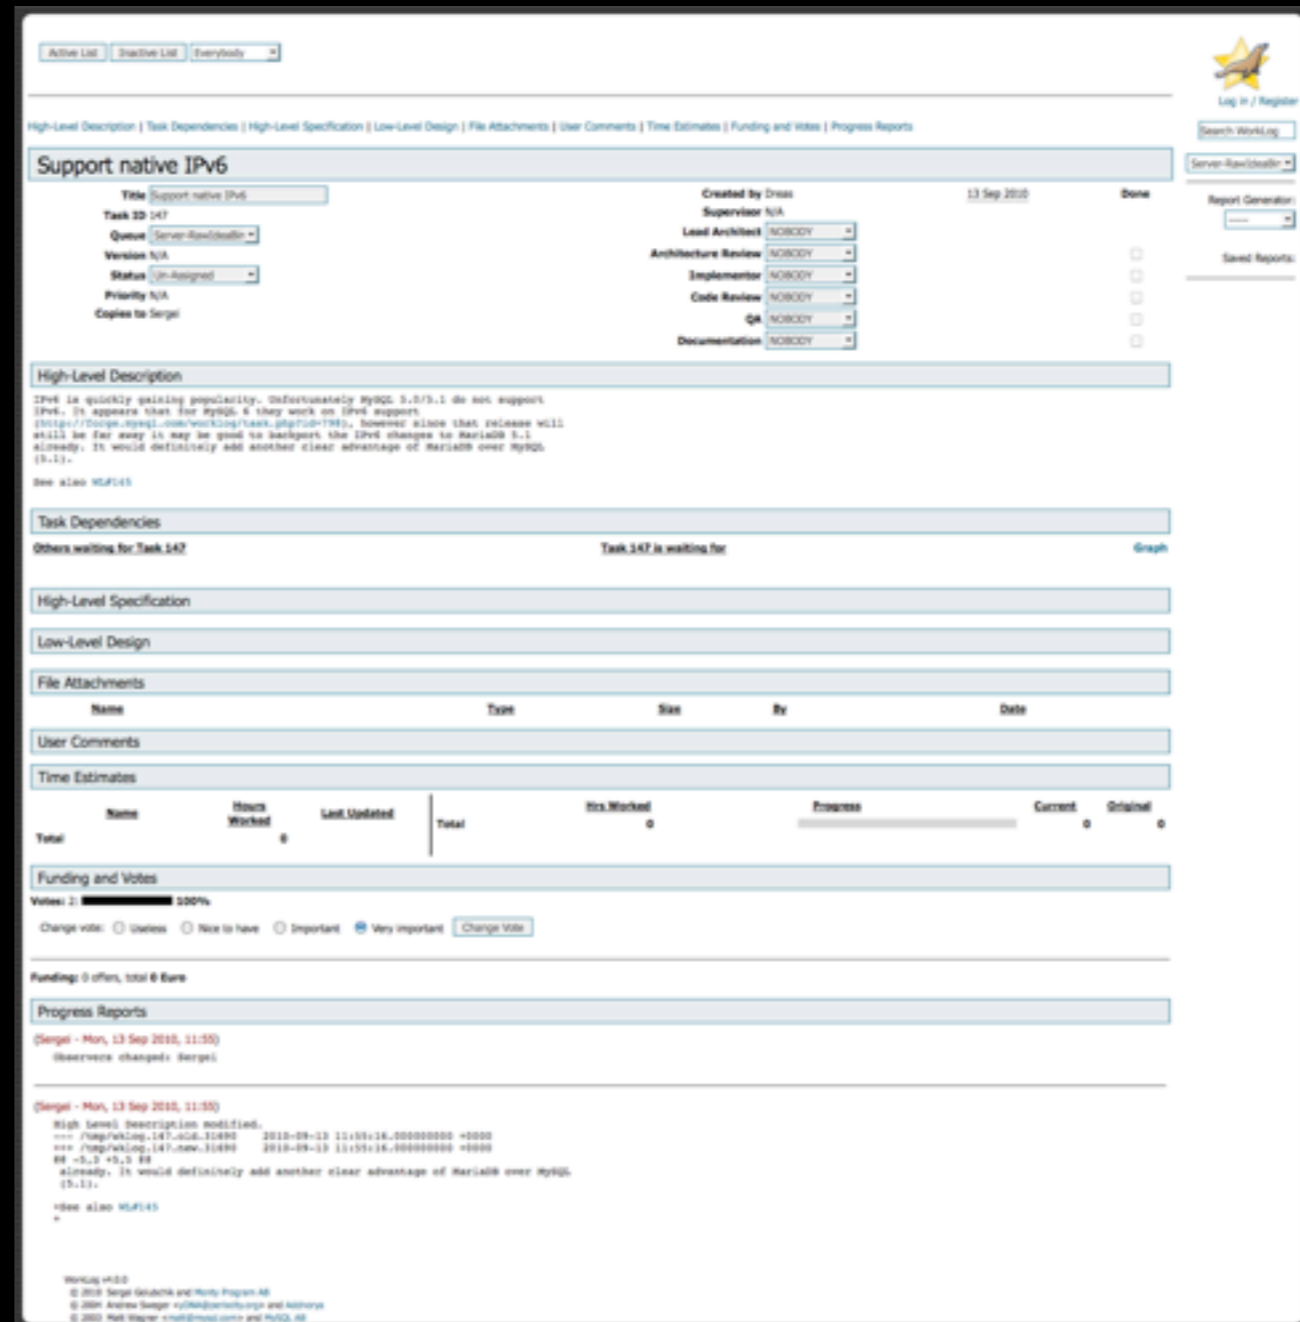
How open is MariaDB?

- Mailing lists: Launchpad
 - maria-developers@lists.launchpad.net
 - maria-discuss@lists.launchpad.net
- Code hosting:
Launchpad
fb.com/MariaDB.dbms |
[@mariadb](https://twitter.com/mariadb)
- <http://launchpad.net/maria>
- Bugs database:
 - <https://bugs.launchpad.net/maria/>
- #maria on FreeNode IRC

Monty Program

Worklog

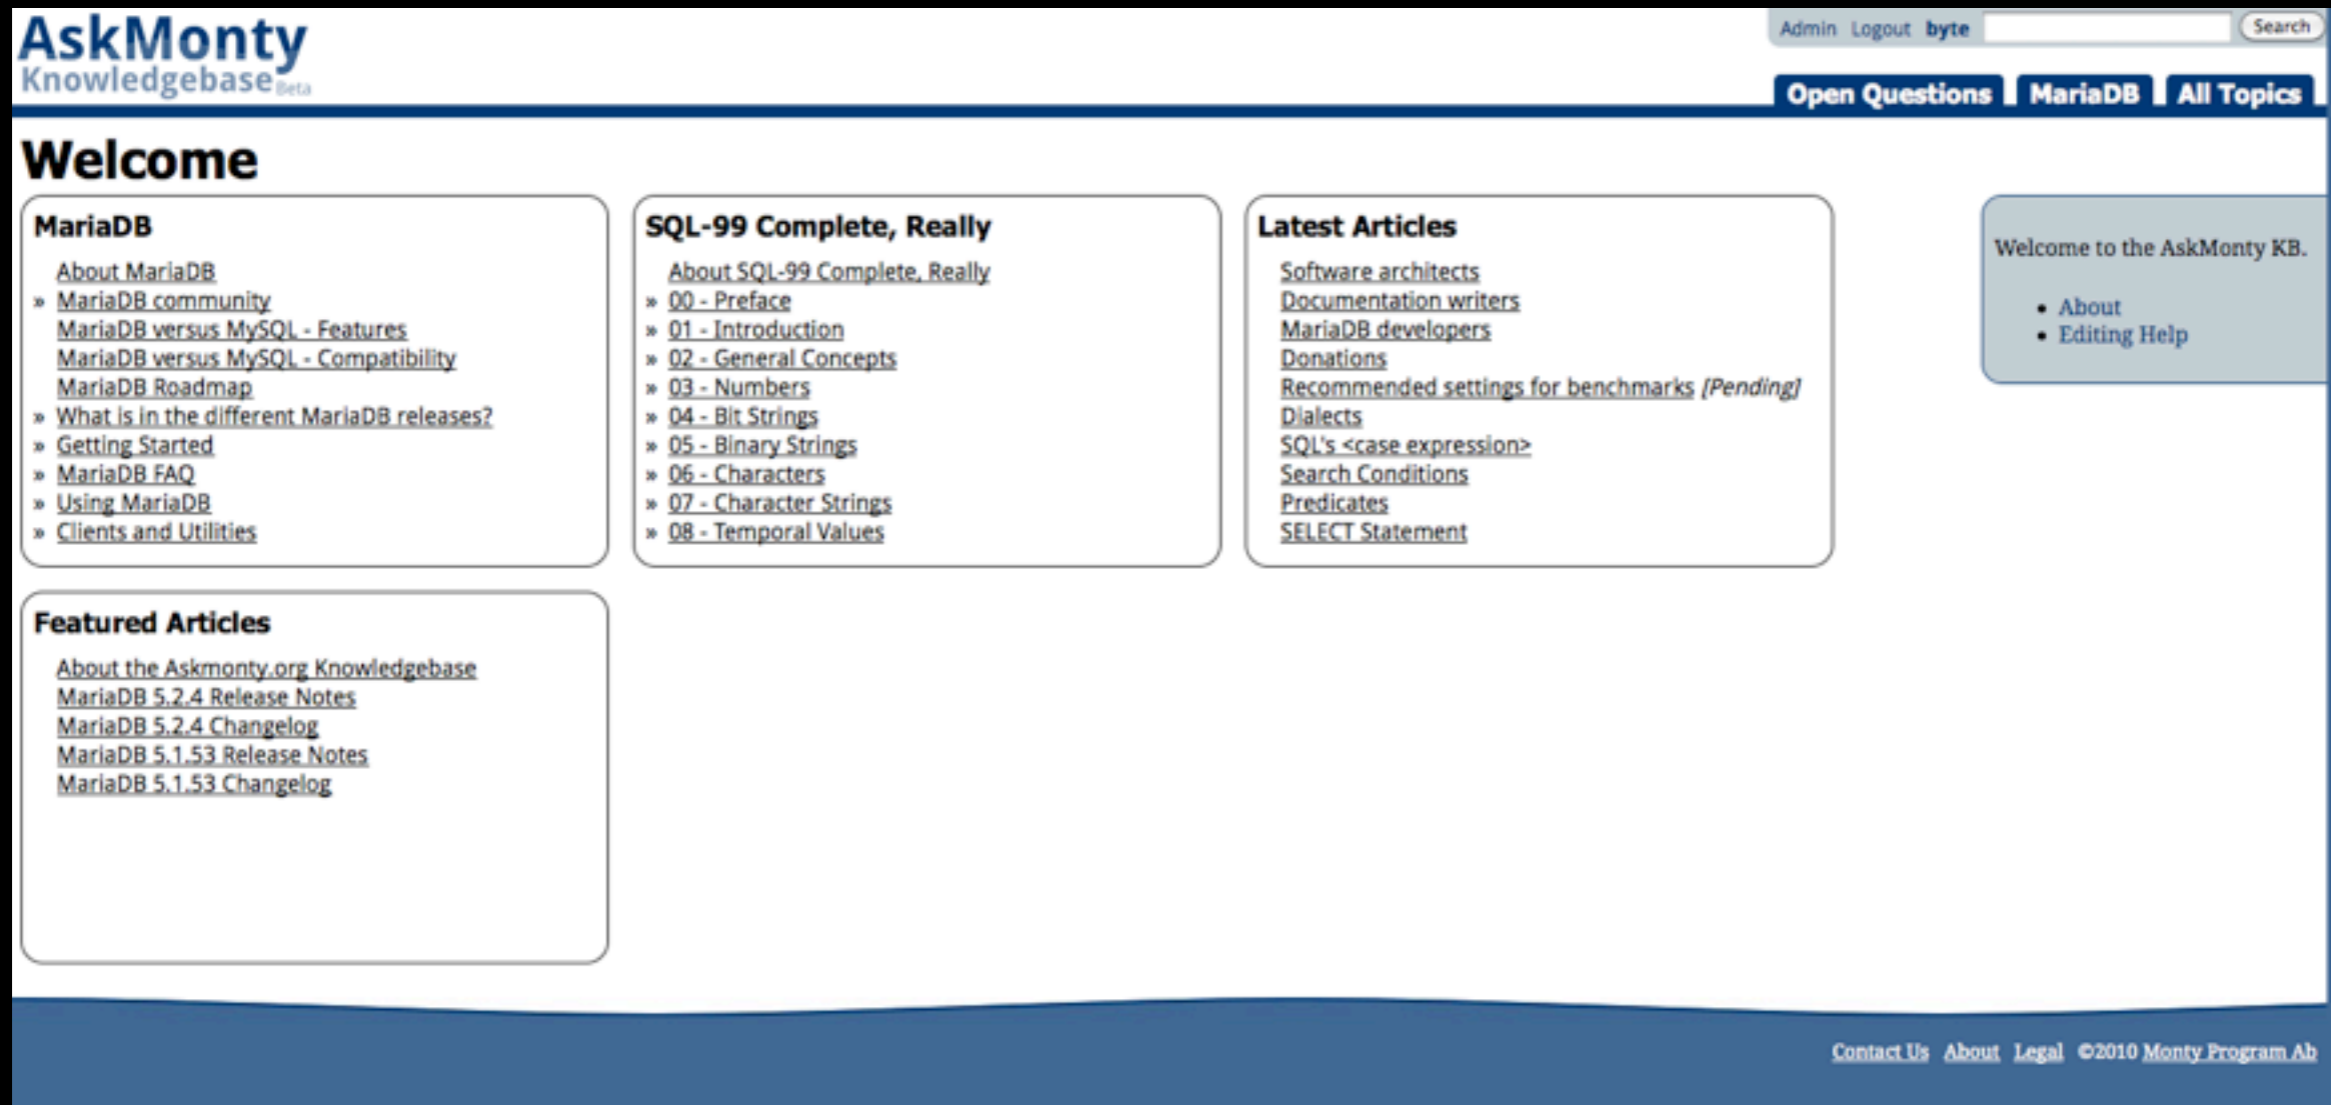
- Funding (bounties)
- Votes
- <http://askmonty.org/worklog/>



The screenshot shows a web-based worklog interface for a task titled "Support native IPv6". The interface includes a navigation bar with tabs for "Active List", "Inactive List", and "Everybody". The task details section shows the title "Support native IPv6", task ID "147", and various fields for "Queue", "Version", "Status", "Priority", and "Copies to Sergei". A table lists roles and their assigned users, such as "Lead Architect" (NOBODY), "Architecture Review" (NOBODY), "Implementor" (NOBODY), "Code Review" (NOBODY), "QA" (NOBODY), and "Documentation" (NOBODY). The "High-Level Description" section contains text about IPv6 support in MySQL 5.6/5.7.1 and a reference to WFP145. The "Task Dependencies" section shows "Others waiting for Task 147" and "Task 147 is waiting for". The "High-Level Specification" and "Low-Level Design" sections are currently empty. The "File Attachments" section has a table with columns for Name, Type, Size, By, and Date. The "User Comments" section is also empty. The "Time Estimates" section has a table with columns for Name, Hours Worked, Last Updated, Total, Hrs Worked, Progress, Current, and Original. The "Funding and Votes" section shows "Votes: 2" and "100%" with radio buttons for "Useless", "Nice to have", "Important", and "Very important". The "Funding" section shows "0 offers, total 0 Euro". The "Progress Reports" section shows a log of updates, including a comment from "Sergei" on "Mon, 13 Sep 2010, 11:55" stating "Serevets changed Sergei".

Knowledgebase

- <http://kb.askmonty.org/>



The screenshot shows the AskMonty Knowledgebase website. The header includes the AskMonty logo, navigation links for 'Admin', 'Logout', and 'byte', and a search bar. Below the header, there are tabs for 'Open Questions', 'MariaDB', and 'All Topics'. The main content area is titled 'Welcome' and contains several sections: 'MariaDB' with a list of links, 'SQL-99 Complete, Really' with a list of links, 'Latest Articles' with a list of links, and 'Featured Articles' with a list of links. A sidebar on the right contains a 'Welcome to the AskMonty KB.' message and links for 'About' and 'Editing Help'. The footer includes links for 'Contact Us', 'About', 'Legal', and a copyright notice for '©2010 Monty Program Ab'.

fb.com/MariaDB.dbms |
@mariadb

Monty Program



Deployments,



deployments, deployments

happy users: pap.fr, wabtec, Paybox Services, OLX, Jelastic, etc.

“MariaDB had these same bugs that we ran into with MySQL. However the big difference was that when we reported these bugs, they were quickly resolved within 48 hours!” -- Dreas van Donselaar, Chief Technology Officer, SpamExperts

B.V. after migrating over 300 servers from MySQL 5.0 to MariaDB 5.1.

“Migrating from MySQL 5.1 to MariaDB 5.2 was as simple as removing MySQL RPMs and installing the MariaDB packages, then running `mysql_upgrade`.” - Panayot Belchev, proprietor, Host Bulgaria on providing MariaDB to over 7,000 of their web hosting customers.

“We made the switch on Saturday – and we’re seeing benefits already – our daily optimization time is down from 24 minutes to just 4 minutes” -- Ali Watters, CEO, travelblog.org

fb.com/MariaDB.dbms |
@mariadb

Monty Program

Getting MariaDB

- <http://mariadb.org/> is the site + downloads
- Get it via OpenSUSE build service
- Its in Gentoo, FreeBSD, Homebrew
- <http://kb.askmonty.org/v/distributions-which-include-mariadb>
- Next up: Debian, Ubuntu, then Fedora/CentOS

fb.com/MariaDB.dbms |
@mariadb

Monty Program

The future?

- Merging takes time
- Features (different implementations) make it into upstream
- Probably moving to JIRA for bugs+worklog functionality


MariaDB next

- Like MariaDB 5.2 and 5.3, we need something before MySQL 5.6 is GA
- Cherry pick GA-ready features for next MariaDB
- MariaDB 10?
- MariaDB Galera Cluster

fb.com/MariaDB.dbms |
@mariadb

Monty Program

We have a book



fb.com/MariaDB

@mariadb

Community Program

Support

- Monty Program does NRE/engineering work for MySQL/MariaDB (not L1/L2 support)
- Support via capable **service providers**
 - <http://mariadb.org/service-providers/>
 - SkySQL (booth here!), Percona, FromDual, Open Query

fb.com/MariaDB.dbms |
@mariadb

Monty Program



Er.. what about other branches?

MySQL	Percona	MariaDB	MySQL	Percona	MariaDB
5.5.20	7.7M	61M	5.5.20	222299	1587843
5.5.22	16M	60M	5.5.22	438567	1540932

fb.com/MariaDB.dbms |
@mariadb

Monty Program



We're on social media

- facebook: [fb.com/MariaDB.dbms](https://www.facebook.com/MariaDB.dbms)
- twitter: [@mariadb](https://twitter.com/mariadb)
- google plus: <http://gplus.to/mariadb>

[fb.com/MariaDB.dbms](https://www.facebook.com/MariaDB.dbms) |
[@mariadb](https://twitter.com/mariadb)

Monty Program

Training: Administering a MySQL Database

Seattle, WA - June 25 - 29 2012

Q&A?

email: colin@montyprogram.com

<http://montyprogram.com/> | <http://mariadb.org/>

twitter: @bytebot / url: <http://bytebot.net/blog/>