

MariaDB: The New M in LAMP?

Colin Charles, Monty Program Ab

colin@montyprogram.com

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

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

SCALE10x, Los Angeles, CA, USA

20 January 2012

whoami

- Chief Evangelist, MariaDB
- Formerly of MySQL AB/Sun Microsystems
- Past lives included FESCO (Fedora Project), OpenOffice.org

fb.com/MariaDB.dbms |
@mariadb

Monty Program

Aims

- Give you an update of what we've done with MariaDB
- Discuss the broader MySQL ecosystem

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

Monty Program

What is it?

- A branch of MySQL with more features?
- A branch of MySQL with better features?
- A fork?

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

Monty Program



MariaDB is...

- Community Developed
- Feature Enhanced
- Backward Compatible

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

Monty Program

Who's behind it?

- MySQL (database) owned by MySQL AB (company)
- Monty Program is just a sponsor of MariaDB
- maria-captains: 64% Monty Program, 36% community (+16% MP 3Q/2011)

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

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

Monty Program

Why MariaDB?

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

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

Monty Program

Aims of MariaDB

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

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

Monty Program



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

What's in MariaDB 5.1

- Released February 2010
 - Create external buildsystem using Buildbot + VMs for test +builds
 - Release filled with storage engines
 - *XtraDB, PBXT, FederatedX, (M)Aria*
 - *Croatian collations*
 - Numerous bug fixes
 - Test cases and coverage improvements
 - Removal of mutexes
 - Compiler warnings gone!
- fb.com/MariaDB.dbms | @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

- PrimeBase PBXT
 - Transactional, foreign keys, ACID, MVCC (read w/o locking)
 - Row-level locking during (SELECT FOR) UPDATE
- Rollback transactions, recovery after restart very fast (identifies garbage making *undo* unnecessary)
- Write once w/log-based storage; write data to DB, without first writing in transaction log
- <http://kb.askmonty.org/v/about-pbxt>

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_schemaprcesslist

```
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)
```

fb.com/MariaDB.dbms |
@mariadb

Monty Program

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)
 - 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>

fb.com/MariaDB.dbms |
@mariadb

Monty Program



What's in MariaDB 5.2

- Released November 2010
- MariaDB 5.2.10 (MariaDB 5.1.60 - changes up to MySQL 5.1.60)
- Pluggable authentication
 - authentication of users handled by plugins
 - mysql_native_password (20 byte) and mysql_old_password (8 byte) included by default
 - CREATE USER foo IDENTIFIED VIA mysql_native_password USING 'pass';
== CREATE user foo IDENTIFIED BY PASSWORD 'pass';
- Use socket_peercred - if already logged in via Unix credentials, MariaDB will trust it
- <http://kb.askmonty.org/v/development-pluggable-authentication>
- **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>

Monty Program



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)

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

Monty Program



```
MariaDB [test]> create table t1 (id int auto_increment primary key,  
-> name varchar(40),  
-> type enum ("shirt", "phone", "computer"),  
-> price decimal(10,2),  
-> dynstr mediumblob);
```

Query OK, 0 rows affected (0.05 sec)

```
MariaDB [test]> insert into t1 (name, type, price, dynstr) values  
-> ("Funny shirt", "shirt", 10.0, COLUMN_CREATE(1, "blue", 10, "XL")),  
-> ("nokia", "phone", 649, COLUMN_CREATE(1, "black", 2, "touchscreen")),  
-> ("htc Desire hd", "phone", 579, COLUMN_CREATE(1, "black", 3, "Android")),  
-> ("BM/Lenovo Thinkpad X60s", "computer", 419, COLUMN_CREATE(1, "black", 3, "Linux"));
```

Query OK, 4 rows affected (0.00 sec)

Records: 4 Duplicates: 0 Warnings: 0

```
MariaDB [test]> select id, name, type, price, length(dynstr) as len, column_list(dynstr) as list from t1;
```

id	name	type	price	len	list
1	Funny shirt	shirt	10.00	17	1,10
2	nokia	phone	649.00	27	1,2
3	htc Desire hd	phone	579.00	23	1,3
4	BM/Lenovo Thinkpad X60s	computer	419.00	21	1,3

4 rows in set (0.00 sec)

```
MariaDB [test]> SELECT name FROM t1 WHERE COLUMN_GET(dynstr, 1 as char(10)) = "black";
```

name
nokia
htc Desire hd
BM/Lenovo Thinkpad X60s

3 rows in set (0.00 sec)

```
MariaDB [test]> SELECT COLUMN_GET(dynstr, 1 as char(10)) as colour, count(*) FROM t1 group by COLUMN_GET(dynstr, 1 as char(10));
```

colour	count(*)
black	3
blue	1

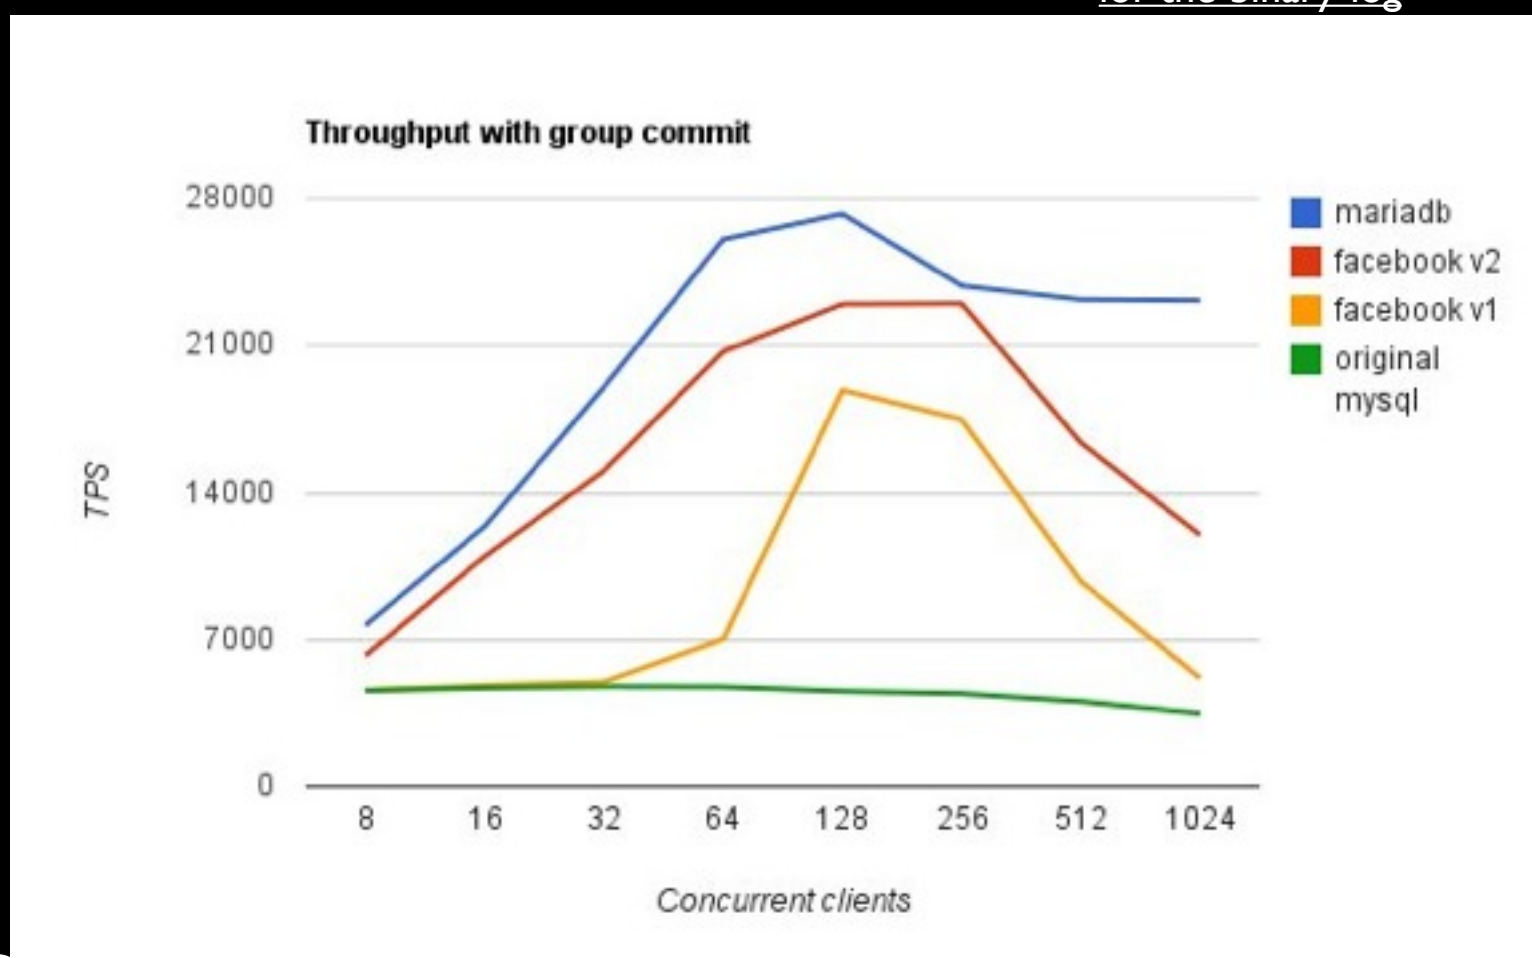
2 rows in set (0.00 sec)





MariaDB 5.3: group commit

- Group commit for 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+PBXT
- row-based replication for tables w/no primary key

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

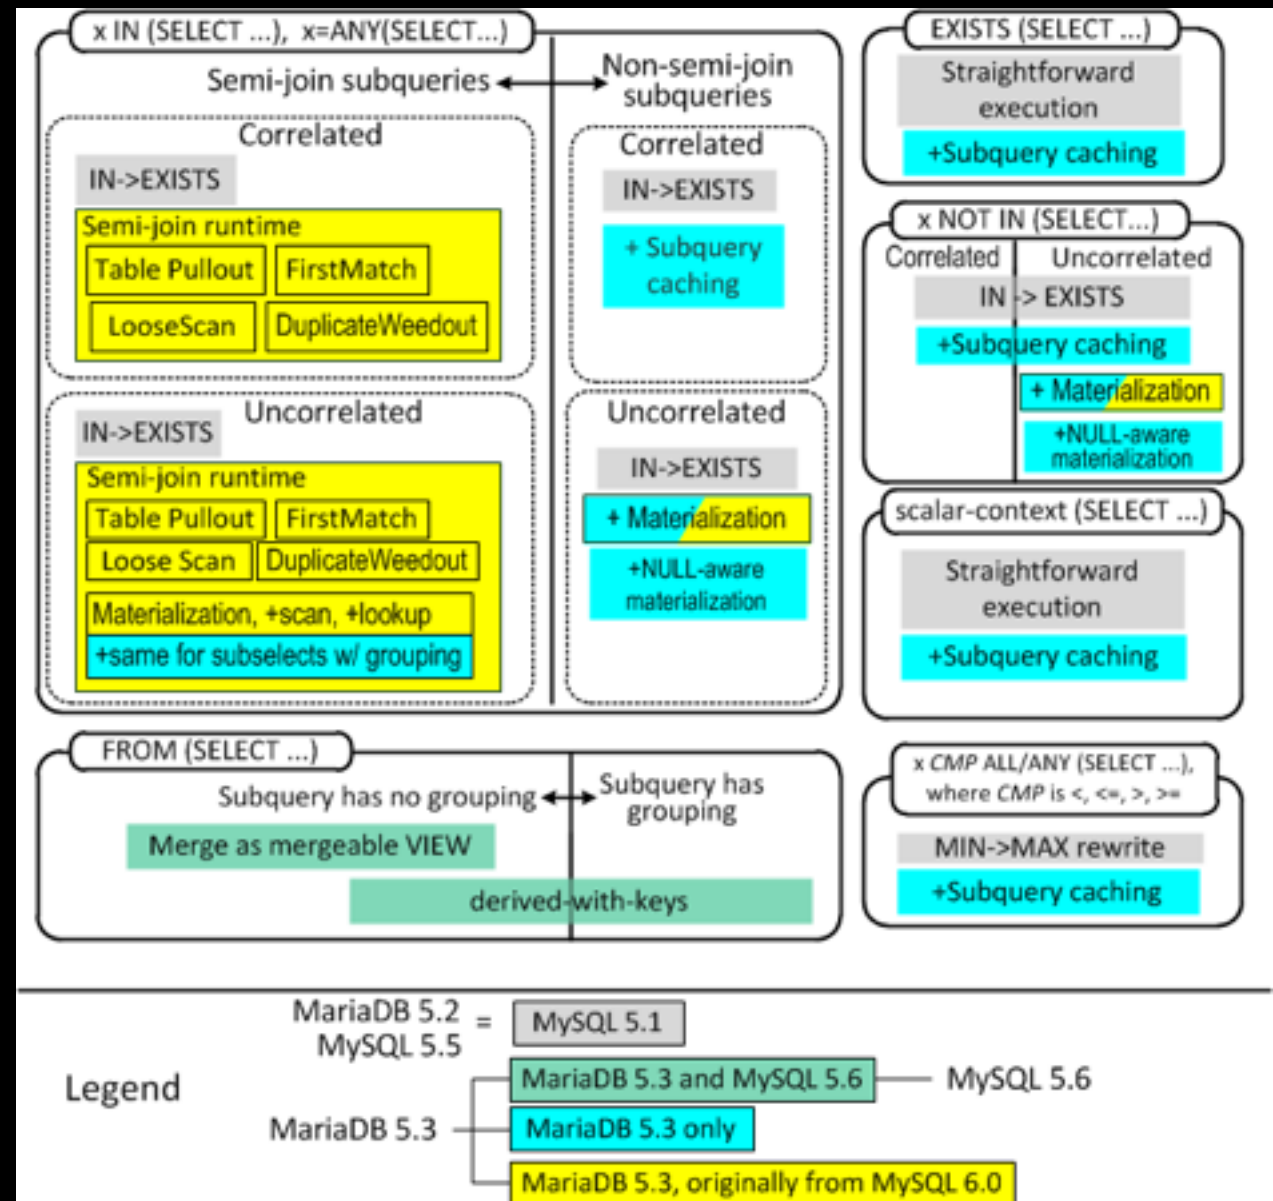
id	select_type	table	type	possible_keys	key	key_len	ref	rows	Extra
1	PRIMARY	<subquery2>	ALL	distinct_key	NULL	NULL	NULL	149637	Using where
1	PRIMARY	customer	ref	c_acctbal	c_acctbal	9	<subquery2>.max(c_acctbal)	1	
2	SUBQUERY	customer	index	NULL	i_c_nationkey	5	NULL	149637	

fb.com/MariaDB.dbms |
@mariadb

Monty Program

MariaDB 5.3: Subqueries

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




fb.com/MariaDB.dbms |
@mariadb

Monty Program

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/



fb.com/MariaDB.dbms |
@mariadb

Monty Program



MariaDB 5.3 & 5.5

- MariaDB 5.3 - RC now
 - MariaDB 5.3.3 - 21/12/2011
- MariaDB 5.5 - Beta any day now
 - <https://launchpad.net/maria/5.5>

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 the 'Support native IPv6' task page in the Monty Program Worklog. The page includes a header with navigation links, a search bar, and a 'Server: MariaDB' dropdown. The main content area is divided into several sections:

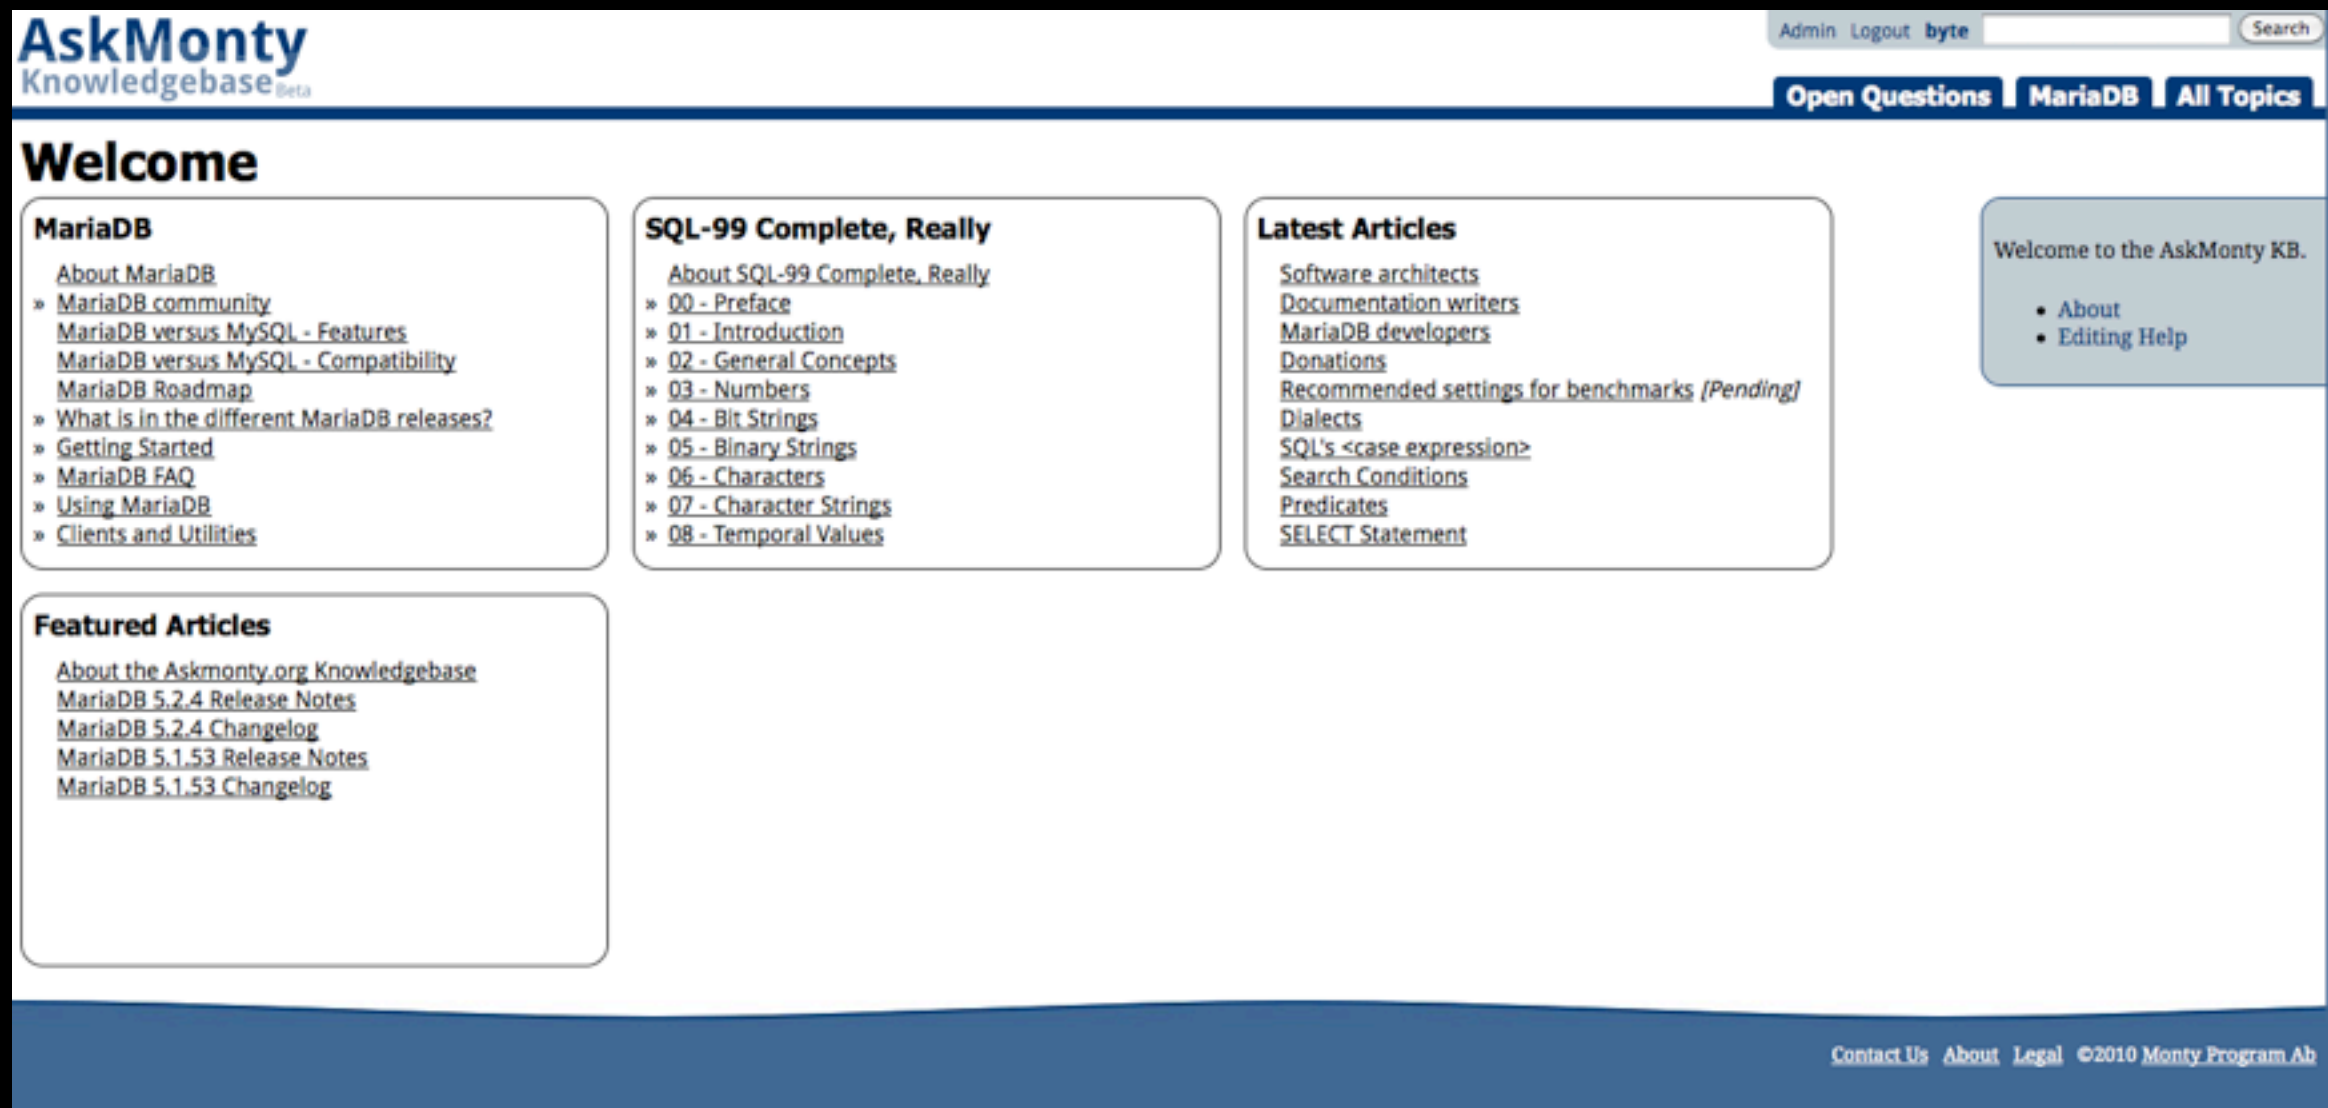
- Task Details:** Title 'Support native IPv6', Task ID '147', Queue 'Server: MariaDB', Version 'N/A', Status 'Un-Assigned', Priority 'N/A', and Copies to Sergej.
- Metadata:** Created by 'Dries', Supervisor 'N/A', Date '13 Sep 2010', Done status, and a list of roles (Lead Architect, Architecture Review, Implementer, Code Review, QA, Documentation) all assigned to 'NOBODY'.
- High-Level Description:** A paragraph explaining that IPv6 is gaining popularity but MySQL 5.5/5.6 do not support it, and that MariaDB 5.5 would be a clear advantage.
- Task Dependencies:** A section indicating 'Task 147 is waiting for' other tasks.
- High-Level Specification:** A section for detailed specifications.
- Low-Level Design:** A section for detailed design.
- File Attachments:** A table with columns for Name, Type, Size, By, and Date.
- User Comments:** A section for user feedback.
- Time Estimates:** A table with columns for Name, Hours Worked, Last Updated, Total, Hrs Worked, Progress, Current, and Original.
- Funding and Votes:** A section showing 'Votes: 1' and 'Funding: 0 offers, total 0 Euro'.
- Progress Reports:** A section with a report from 'Sergej - Mon, 13 Sep 2010, 11:55' stating 'Reviews changed: Sergej'.

fb.com/MariaDB.dbms |
@mariadb

Monty Program

Knowledgebase

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



The screenshot shows the AskMonty Knowledgebase website. At the top left is the AskMonty logo with 'Knowledgebase' in smaller text. To the right are links for 'Admin', 'Logout', and 'byte', along with a search box. Below this is a navigation bar with 'Open Questions', 'MariaDB', and 'All Topics'. The main content area is titled 'Welcome' and contains three columns of links. The first column is 'MariaDB' with links like 'About MariaDB', 'MariaDB community', 'MariaDB versus MySQL - Features', 'MariaDB versus MySQL - Compatibility', 'MariaDB Roadmap', 'What is in the different MariaDB releases?', 'Getting Started', 'MariaDB FAQ', 'Using MariaDB', and 'Clients and Utilities'. The second column is 'SQL-99 Complete, Really' with links for 'About SQL-99 Complete, Really' and a list of chapters from 00 to 08. The third column is 'Latest Articles' with links for 'Software architects', 'Documentation writers', 'MariaDB developers', 'Donations', 'Recommended settings for benchmarks [Pending]', 'Dialects', 'SQL's <case expression>', 'Search Conditions', 'Predicates', and 'SELECT Statement'. On the right side, there is a 'Welcome to the AskMonty KB.' message with links for 'About' and 'Editing Help'. At the bottom right, there are links for 'Contact Us', 'About', 'Legal', and '©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](https://twitter.com/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>

fb.com/MariaDB.dbms |
@mariadb

Monty Program


The future?

- GIS features making it into MariaDB 5.3
- Merging takes time
- Features (different implementations) make it into upstream
- Getting into distributions (Fedora/EPEL policies), Ubuntu/Debian, etc ?

fb.com/MariaDB.dbms |
[@mariadb](https://twitter.com/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 (here), FromDual, Open Query, SoftMethod GmbH, Percona

fb.com/MariaDB.dbms |
@mariadb

Monty Program



We're on social media

- facebook: fb.com/MariaDB.dbms
- twitter: [@mariadb](https://twitter.com/mariadb)
- google plus: search for mariadb

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

Monty Program

Q&A?

email: colin@montyprogram.com

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

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